

Sensitivity Simulations and Half-wave Plate Polarization Modulators for Cosmic Microwave
Background Observatories

by

Charles Alexander Hill

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Physics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Adrian Lee, Chair

Professor Uros Seljak

Professor Aaron Parsons

Professor Akito Kusaka

Fall 2020

Sensitivity Simulations and Half-wave Plate Polarization Modulators for Cosmic Microwave
Background Observatories

Copyright 2020
by
Charles Alexander Hill

Abstract

Sensitivity Simulations and Half-wave Plate Polarization Modulators for Cosmic Microwave Background Observatories

by

Charles Alexander Hill

Doctor of Philosophy in Physics

University of California, Berkeley

Professor Adrian Lee, Chair

Cosmic microwave background (CMB) polarization anisotropies are at the forefront of modern cosmology, and a detection of B-mode polarization due to inflationary gravitational waves is among the most sought-after discoveries in astrophysics. Current constraints on these primordial B-modes set their amplitude at $\lesssim 10$ nK, posing extraordinary technological challenges to the design, characterization, and performance of today's CMB observatories. An effective CMB telescope requires both excellent sensitivity and tight control of systematic effects, and advancements in telescope instrumentation, calibration hardware, and analysis techniques are needed for a robust extraction of the inflationary signal.

In this dissertation, we present two primary research products that improve the scientific prospects of today's CMB experiments. First, we describe a Python simulation code to optimize the mapping speed of both existing and future instruments. This calculator conglomerates the best aspects of several existing codes to offer a generalized, feature-filled software both for modeling future experiments and for characterizing operating ones. We discuss its use within Simons Array (SA) and Simons Observatory (SO), as well as its broader utility to upcoming experiments, including CMB-S4. Second, we present new polarization modulators for SA, including an ambient-temperature, continuously rotating half-wave plate (HWP) for POLARBEAR-2a and a cryogenic HWP for POLARBEAR-2b. Continuously rotating HWPs are powerful tools to mitigate atmospheric $1/f$ noise and telescope-induced intensity-to-polarization leakage, and their effectiveness is increasingly important to a precise characterization of cosmic polarization. The HWPs described in this thesis introduce several optical, mechanical, and electrical hardware advancements, and we discuss how these HWPs are paving the way for similar modulators on POLARBEAR-2c and SO's small aperture telescopes.

To Lee Roberts,
who saw me as a physicist before I did...and then made it possible for me to become one.

Contents

List of Figures	viii
List of Tables	xiv
Preface	xv
Acknowledgments	xvi
Introduction	xxii
1 Scientific motivation	1
1.1 Big Bang cosmology	2
1.2 Cosmic microwave background	3
1.3 Inflation	6
1.4 Primordial perturbations	10
1.5 CMB anisotropies	12
1.6 CMB power spectrum	16
1.7 Foregrounds	18
1.8 State of the field	20
2 CMB instrument overview	21
2.1 Observing the CMB	21
2.1.1 Simons Array	23
2.1.2 Simons Observatory	23
2.2 Observation site	24
2.2.1 Sky intensity	24
2.2.2 Chile site	26
2.3 Optics	27
2.3.1 Telescope and optics tube	28
2.3.2 Anti-reflection coatings	30
2.3.3 Focal plane optics	33
2.3.4 Telescope-to-focal plane coupling	35
2.4 Detectors	38

2.4.1	TES operation	38
2.4.2	TES readout	41
2.5	IR-blocking design	42
2.6	Technical motivations	44
3	Instrument sensitivity	45
3.1	Noise spectrum	45
3.2	Photon statistics	46
3.3	Photon NEP	49
3.4	Optical power	51
3.4.1	Emissivity and absorptivity	51
3.4.2	Spillover fraction	53
3.4.3	Scattering fraction	53
3.4.4	Reflectivity	55
3.4.5	Transmissivity	55
3.4.6	Optical throughput and optical efficiency	55
3.4.7	Sky temperature and telescope temperature	56
3.5	Bolometer thermal carrier NEP	56
3.6	Readout NEP	58
3.7	Johnson NEP	59
3.8	NET	60
3.9	Photon noise correlation factor	62
3.10	N_ℓ and mapping speed	63
3.11	Discussion	65
4	Photon noise correlations	67
4.1	Theoretical foundations	68
4.1.1	Photon correlations in a circuit	68
4.1.2	Photon correlations in free space	70
4.1.3	Photon noise correlations	73
4.2	Model optical system	73
4.2.1	Telescope	74
4.2.2	Focal plane	76
4.3	Correlation calculation	77
4.3.1	Aperture radiation	78
4.3.2	Stop radiation	80
4.3.3	Correlation patterns	80
4.4	Impact on sensitivity	83
4.4.1	Mapping speed	83
4.4.2	Pixel size optimization	85
4.5	Implications for experiment design	90
4.6	Discussion	92

5	BoloCalc: a sensitivity calculator	93
5.1	Calculator design	94
5.1.1	Structure	94
5.1.2	Input and output parameters	96
5.1.3	User interface	98
5.2	Calculator features	101
5.2.1	Foregrounds	101
5.2.2	Site and atmosphere	103
5.2.3	Monte Carlo simulations	104
5.2.4	Parameter sweeps	105
5.2.5	Custom bandpasses	106
5.3	Calculator validation	107
5.4	Informing the design of SO	108
5.4.1	LATR architecture	108
5.4.2	Camera magnification	110
5.4.3	Pixel pitch	111
5.4.4	LAT primary spillover	112
5.4.5	SAT stop temperature	113
5.4.6	Detector saturation power	113
5.4.7	Scan strategy	114
5.5	Informing PB-2c Psats	115
5.5.1	Estimating P_{opt}	117
5.5.2	Maximizing $\text{MS} \times \eta_{\text{obs}}$	119
5.6	Other applications	121
6	Half-wave plate polarization modulation	123
6.1	Atmospheric 1/f noise	125
6.2	Polarized beam systematics	127
6.3	Sapphire achromatic HWP	128
6.3.1	Stokes polarization	129
6.3.2	Birefringent HWP	130
6.3.3	Mueller matrices	131
6.3.4	Achromatic HWP	133
6.4	AHWP performance	135
6.4.1	Modulation efficiency	137
6.4.2	Frequency-dependent phase	137
6.5	Continuous polarization modulation	142
6.5.1	HWP synchronous signals	144
6.6	HWPs for Simons Array	146
6.6.1	The POLARBEAR HWP	146
6.6.2	The SA HWP	147
6.7	SA HWP requirements	148

6.7.1	Modulation efficiency	149
6.7.2	Rotation velocity	150
6.7.3	Angle encoding noise	152
6.8	Discussion	153
7	A warm HWP for PB-2a	154
7.1	Requirements	155
7.1.1	Clear-aperture diameter	155
7.1.2	Emissivity and reflectivity	157
7.2	Optical design	158
7.2.1	Sapphire	159
7.2.2	Anti-reflection coating	161
7.3	Mechanical assembly	164
7.3.1	Mount	165
7.3.2	Rotation stage	166
7.3.3	Angle encoder	166
7.3.4	Mechanical validation	167
7.4	Optical evaluation	168
7.4.1	Sapphire index	168
7.4.2	Sapphire loss	170
7.4.3	Transmissivity	172
7.4.4	Polarization efficiency	173
7.5	Discussion	175
8	A cryogenic HWP for PB-2b	176
8.1	Precedent	178
8.2	Requirements	179
8.2.1	Spatial and optical requirements	179
8.2.2	Thermal requirements	181
8.2.3	Noise requirements	183
8.2.4	Operational requirements	185
8.3	Cryo-mechanical assembly	185
8.3.1	Location	185
8.3.2	Optical design	187
8.3.3	Bearing	187
8.3.4	Gripper	188
8.4	Motor	190
8.4.1	Motor cryogenic assembly	191
8.4.2	Motor driver	192
8.4.3	Motor efficiency	195
8.5	Angle encoder	195
8.6	Thermal design	196

8.6.1	Rotor thermometry	200
8.7	Integration	201
9	PB-2b CHWP evaluation	203
9.1	Prototype evaluation	203
9.2	PB-2b CHWP evaluation	205
9.3	Operational performance	205
9.3.1	Cooldown	206
9.3.2	Start-up	207
9.3.3	Continuous rotation	208
9.3.4	Thermal impact	209
9.3.5	Shutdown and recovery	211
9.4	Data quality	212
9.4.1	Angle encoder performance	212
9.4.2	Magnetic interference	214
9.4.3	Rotor temperature stability	215
9.5	Discussion	216
10	Sapphire anti-reflection coating	218
10.1	AR coating design	219
10.1.1	Tolerances	221
10.1.2	Objectives	221
10.2	Epoxy + epoxy AR	223
10.2.1	Fabrication	224
10.2.2	Performance	231
10.2.3	Assessment	232
10.3	Epoxy + Duroid AR	234
10.3.1	Fabrication	234
10.3.2	Performance	240
10.3.3	Assessment	243
10.4	Ongoing work	245
10.4.1	Epoxy + Duroid	245
10.4.2	Mullite + Duroid	245
10.4.3	Vespel SF-0940	246
11	Research impact	249
11.1	BoloCalc	249
11.2	HWPs	250
11.3	Conclusion	252
	Bibliography	253

A	Quantum and classical expectation values	274
A.1	Thermal photon density matrix	276
B	Classical intensity correlations	277
B.1	Polarization correlations	278
B.2	Stokes orthogonality	279

List of Figures

1.1	A schematic of the comoving grid.	3
1.2	The original Hubble diagram alongside a modern version.	4
1.3	The FIRAS CMB spectrum (1990).	6
1.4	Mode exit and re-entry in the inflationary paradigm.	7
1.5	The R^2 inflationary potential, which serves as an instructive example of slow-roll inflation.	9
1.6	Current constraints on the tensor-to-scalar ratio and the scalar index.	11
1.7	Full-sky maps of CMB temperature and polarization anisotropies from Planck (2018).	13
1.8	A schematic of CMB polarization due to Thomson scattering of quadrupolar anisotropies.	14
1.9	A schematic of E-modes and B-modes in the flat-sky approximation.	15
1.10	The CMB temperature power spectrum D_ℓ^{TT} from Planck (2018).	16
1.11	The power spectrum D_ℓ^{EE} of CMB E-mode anisotropies as measured by Planck (2018).	17
1.12	The power spectrum D_ℓ^{BB} of CMB B-mode anisotropies as measured by several experiments.	18
1.13	The brightness temperature and angular power spectra of galactic foregrounds.	19
2.1	The sky temperature at the Atacama observation site assuming 1.4 mm PWV and 50 deg elevation, plotted over SO's detector bands.	25
2.2	A photo of the Atacama observation site along with its annual PWV distribution and its accessible sky area.	27
2.3	The PB-2b telescope and receiver cryostat as a representative example of the SA optics.	29
2.4	CAD cross-sections of SO's LAT, LATR, and SAT.	31
2.5	A sampling of AR-coating technologies used within SA and SO.	32
2.6	SA's and SO's focal plane optics.	34
2.7	A schematic of detector-pixel optics.	35
2.8	Beam coupling efficiencies for SA's and SO's detector pixels.	37
2.9	A schematic and demonstration of TES bolometer operation.	39
2.10	A schematic of SA's dfMUX readout.	41

2.11	SA's 50 K alumina filter and its lowest-cutoff metal mesh filters.	43
3.1	The relationship between photon occupation number and electromagnetic mode temperature.	48
3.2	A schematic of the one-dimensional optical system assumed in the optical power calculation.	51
3.3	A schematic of the forward-time-sense and reverse-time-sense calculations. . .	52
3.4	Dependence of NEP_g and P_{sat} on bolometer transition temperature T_c for various bath temperatures T_b	57
3.5	Conversion factor from RJ temperature fluctuations to those of a blackbody source with temperature T_{phys}	61
3.6	Wave-noise fraction assuming a top-hat band with 20 GHz bandwidth and an optical throughput between the source and detector of 30%.	63
3.7	Forecasted N_ℓ noise curves for the SO LAT in temperature and the SO SAT in polarization.	65
4.1	A schematic of the quantum circuit formalism.	69
4.2	A simplified optical model used for all calculations in this chapter.	74
4.3	The assumed layout of detector pixels on the focal plane for the correlation calculation.	77
4.4	A schematic of the aperture, stop, and pixel radiation models used for the correlation simulation.	78
4.5	Simulated correlation at the focal plane due to radiation within the aperture and from the stop.	81
4.6	An example "classic" mapping speed calculation, which ignores the impact of HBT correlations.	85
4.7	The assumed source temperatures for the correlation calculation and their corresponding Bose noise fractions.	87
4.8	The HBT correlation across various focal plane layouts.	88
4.9	Mapping speed vs. pixel size in each frequency band for a variety of external sources.	89
4.10	The impact of stop temperature, telescope temperature, and internal detector noise on mapping speed vs. pixel size.	91
5.1	The generic layout of a BoloCalc project.	95
5.2	A BoloCalc optical power table output for an arbitrarily constructed example experiment.	99
5.3	The file structure of a BoloCalc experiment, which mimics the calculator's underlying class structure.	100
5.4	A screenshot of BCG's interface for editing the optics' input parameters. . . .	101
5.5	An example window for editing an input parameter using BCG.	102

5.6	A comparison of the Atacama and the South Pole atmospheric profiles used in BoloCalc.	103
5.7	An example MC-generated BoloCalc parameter sweep.	106
5.8	A sensitivity comparison between three candidate SO LATR configurations. . .	109
5.9	Relative mapping speed vs. camera F/# in each frequency band within SO's LAT and SAT.	110
5.10	Sensitivity vs. pixel pitch in both SO's LAT and SAT.	111
5.11	The impact of the SO LAT's primary spillover on optical power and mapping speed.	112
5.12	Relative optical power and mapping speed in each frequency band vs. the SO SAT's aperture stop temperature.	113
5.13	Relative mapping speed for each frequency band in SO's LAT and SAT as a function of saturation power.	114
5.14	Normalized mapping speed vs. PWV and elevation at the Atacama observation site in the 90 and 150 GHz bands of SO's LAT.	115
5.15	Probability distributions of sky temperature and telescope temperature in PB-2c's 220 and 270 GHz bands.	117
5.16	Probability distributions for optical power in PB-2c's 220 and 270 GHz bands.	118
5.17	Mapping speed vs. detector saturation power in PB-2c's 220 and 270 GHz bands for various detector optical efficiencies.	119
5.18	Mapping speed vs. detector saturation power in PB-2c's 220 and 270 GHz bands for various saturation-power cutoffs and detector optical efficiencies.	120
5.19	A BoloCalc-generated sensitivity plot for the initial US LiteBIRD proposal to NASA in 2016.	121
6.1	An idealized simulation to show how detector differencing rejects unpolarized atmospheric fluctuations.	124
6.2	A schematic of differential beam systematics.	128
6.3	The relationship between Jones vectors and Stokes vectors.	130
6.4	A depiction of how an HWP manipulates polarization orientation.	132
6.5	A depiction of the Pancharatnam achromatic HWP using the Poincaré sphere.	134
6.6	A cartoon of the sapphire achromatic HWP.	135
6.7	Detected intensity vs. HWP angle for a one-, three-, and five-stack HWP between 50 and 170 GHz.	136
6.8	Comparison of the band-averaged modulation efficiency and phase for one-, three-, and five-stack HWPs.	138
6.9	The estimated angle error induced by the AHWP when calibrated using Tau-A and a sparse wire grid.	140
6.10	A toy demonstration of amplitude modulation.	143
6.11	Differential thermal emission from a three-stack sapphire AHWP at 300 K and 100 K.	145
6.12	An overview of the POLARBEAR HWP's implementation.	147

6.13	The demonstrated performance of the POLARBEAR HWP.	148
6.14	The PB-2a/b band-averaged modulation efficiency vs. plate orientation for the 3- and 5-stack AHWPs.	149
6.15	The impact of thickness and relative-alignment tolerances on the modulation efficiency of PB-2a/b's AHWP.	150
6.16	A notional depiction of the HWP's modulation band referenced to the PSD of dummy-generated atmospheric data.	151
7.1	The PB-2a ray-trace simulation and the resulting WHWP clear-aperture diameter requirement.	156
7.2	The impact of WHWP emissivity and reflectivity on PB-2a's mapping speed. .	158
7.3	An exploded view of the PB-2a WHWP's optical stack.	159
7.4	Photographs of the GHTOT large-diameter sapphire.	160
7.5	AR coating candidates for the PB-2a WHWP and their measured indices and loss tangents.	162
7.6	The impact of a potential glue layer on the PB-2a WHWP's emissivity and reflectivity.	163
7.7	A cross-sectional view of the PB-2a WHWP's vacuum module.	164
7.8	Photos of the PB-2a WHWP's mechanical assembly.	165
7.9	The PB-2a WHWP's encoder mechanism and performance.	167
7.10	The PB-2a AHWP's modulation efficiency and phase compared to that of a single HWP.	169
7.11	PB-2a's sapphire crystal axis orientation measurement.	170
7.12	The PB-2a sapphire's loss tangent measured at 300 K.	171
7.13	A schematic of the FTS apparatus used to measure PB-2a's WHWP and a plot of the resulting transmissivity.	173
7.14	A measurement of the WHWP's polarization modulation efficiency in PB-2a's observation bands.	174
8.1	CAD of the CHWP in PB-2b.	177
8.2	The reverse-time-sense illumination pattern onto the CHWP aperture plane from detector pixels on the focal plane.	180
8.3	The impact of the CHWP's rotor temperature on its emissivity and on PB-2b's mapping speed.	182
8.4	The CHWP rotation mechanism without the sapphire stack.	186
8.5	A detailed view of a single gripper subassembly.	189
8.6	A CAD cross-section of the CHWP's motor and a photograph of the encoder read heads.	191
8.7	A signal diagram for a single phase of the CHWP's motor driver and for a single angle encoder output.	193
8.8	A photograph of the CHWP's driver PCB.	194

8.9	The CHWP's motor efficiency vs. the rotor's radial misalignment and rotation frequency.	196
8.10	The CHWP's thermal circuit during nominal operation.	198
8.11	The expected rotor temperature and CHWP-induced power on the field lens. .	200
8.12	A photograph of the CHWP being installed into the PB-2b optics tube.	201
9.1	The CHWP prototype developed at LBNL prior to constructing the PB-2b system.	204
9.2	The CHWP mounted in the PB-2b receiver at UCSD.	206
9.3	Temperature vs. time and radial displacement of each gripper finger vs. rotor and stator temperature during LBNL testing.	207
9.4	CHWP rotation frequency vs. time during start-up for various H-driver voltages in the LBNL cryostat.	208
9.5	Δf_{HWP} during one hour of continuous rotation in the LBNL cryostat.	209
9.6	Measured 50 K dissipation and mK heating during continuous rotation.	210
9.7	CHWP spin-down tests from various rotation frequencies using various braking voltages.	211
9.8	A measurement of encoder performance during 1 hour of testing in the LBNL setup.	213
9.9	A measurement of the magnetic field 1.5 m behind the CHWP assembly in the LBNL cryostat.	215
9.10	A power spectrum of the expected rotor temperature variations.	216
10.1	A schematic for a quarter-wave AR coating.	219
10.2	A comparison of one-, two-, and three-layer AR coatings for the PB-2b CHWP. .	220
10.3	The impact of index and thickness tolerances on the PB-2b AR coating.	222
10.4	A cartoon schematic of the epoxy + epoxy AR coating.	223
10.5	Measured performance of the epoxy + epoxy AR coating on a 50 mm-diameter, 3 mm-thick alumina sample at both 300 K and 140 K.	224
10.6	A microscope comparison of the sapphire surface after sandblasting vs. after sanding.	225
10.7	Photographs of the upgraded AP-134 application technique used for the PB-2b sapphire.	226
10.8	Photographs of the epoxy-AR application and cure.	227
10.9	Photographs of epoxy-AR machining and metrology.	228
10.10	Photographs of the epoxy-coated sapphire after fabrication and before thermal cycling.	230
10.11	Photographs of cryogenic delamination of epoxy + epoxy on sapphire.	232
10.12	Partial delamination of the epoxy + epoxy AR coating from alumina and the resulting impact on reflectance vs. temperature.	233
10.13	A cartoon schematic of the epoxy + Duroid AR coating.	234
10.14	A comparison of the epoxy + epoxy and epoxy + Duroid AR coatings.	235

10.15	Photographs of the EpoTek 301-2 application for bonding the Duroid 5880LZ top AR layer to the machined Stycast 2850FT bottom AR layer.	236
10.16	Microscope photos of the laser dicing lines, showing the impact of decreasing pulse duration and the number of passes.	238
10.17	A demonstration of the impact of annealing on epoxy stress.	240
10.18	An ambient measurement of epoxy + Duroid on a 150 mm-diameter alumina sample.	242
10.19	A cryogenic measurement of epoxy + Duroid on a 150 mm-diameter alumina sample.	243
10.20	The results of thermal cycling epoxy + Duroid on 500 mm-diameter sapphire to cryogenic temperatures.	244
10.21	Photograph and performance of the Vespel SF-0940 foam.	247
11.1	Photograph of the Simons Observatory small aperture telescope's CHWP. . . .	250
11.2	CAD of the polarization modulator unit for the LiteBIRD satellite.	251

List of Tables

1.1	The values of the cosmological parameters according to the Planck satellite. . .	5
8.1	The CHWP's numerical requirements and achieved values.	179
8.2	Parameter definitions and values for the CHWP thermal model.	199

Preface

Here is what I know - please know it also!

This dissertation has two primary objectives. The first is to leave a somewhat-detailed, well-organized account of my research endeavors at Berkeley. While the Graduate Division may disagree, this first objective is the lesser of the two, as nitty-gritty details about how I mix epoxies, probe printed circuit boards, and tighten screws are laid out in our collaborations' shared documentation. The second and more important objective is to teach the interested reader (you!) about the cosmic microwave background (CMB), CMB telescopes, instrument sensitivity, polarization modulation, sapphire half-wave plates, anti-reflection coatings, and other things I've latched onto during my PhD.

I vividly remember what it was like to be a new CMB student. The field is huge and busy, and a mountain of terminologies, concepts, and implicit understandings surround every telecon discussion. In addition, the task of finding help, asking questions, and retaining answers can feel utterly daunting. Attempting to assuage this dilemma, professors often refer their fresh graduate and undergraduate students to theses as an (efficient) pedagogy technique. While I believe that reading is at the center of learning and that studying papers is an important ritual, I've found most dissertations to be sporadic assemblages of disparate topics written in an esoteric tone. When I joined the Berkeley CMB group, this reality left me scrambling between ~ 10 theses, patching together the best parts of each to paint clear pictures. And even after these escapades, I often finished my study sessions with more unanswered questions than answered ones. What I'm saying is, *I know how you feel*, and I'm here to try and help.

I believe in a simple yet powerful antidote to generalized confusion among CMB researchers: context. Context is the most underrated aspect of technical mastery, constructive collaboration, and impactful innovation, as knowing *why* is just as important as knowing what or how. Therefore, I try to put a *substantial* emphasis on context in this dissertation, dedicating pages in each chapter to in-depth prefaces (like this one), auxiliary discussions, and illuminating examples. I've tried to imagine your questions and address them as if I were there with you, and I hope my textbook-like approach maximizes this document's usefulness to you. Thanks for reading, and happy learning!

Acknowledgments

I had *lots* of help.

First, I want to thank my advisor, Professor Adrian Lee. During my first year of grad school in 2013, Adrian replied to my initial email inquiry about the group—in which I naively asked about MAXIPOL instead of POLARBEAR—with enthusiasm and an open-minded attitude despite my having nearly zero experience either in astrophysics or with instrumentation. He invested in me over those first several months despite my rookie deficiencies, sending me to Chile, encouraging me to take ownership of the PB-2a warm half-wave plate (WHWP), and providing me with resources to do impactful science. During my time with Adrian, the group has grown in scope and influence, creating an *immensely* rich atmosphere to learn about and participate in a huge variety of cutting-edge experimentation. The research in this thesis exists within the framework of Simons Array, Simons Observatory, LiteBIRD, and the UC Berkeley experimental cosmology group, and it could not exist without the many opportunities made available to me under Adrian’s leadership.

Next, I want to thank Professor Akito Kusaka, who has been my teacher and comrade throughout my graduate career. Akito came to Lawrence Berkeley National Laboratory (LBNL) only a few months after I joined the Berkeley CMB group, and in the fall of 2014, he invited me to begin developing cryogenic half-wave plates (CHWPs) for Simons Array in his new group. When I first began spending time at LBNL, Akito’s lab (and office, for that matter) was nearly empty, but nonetheless, Akito brought in a steady flow of funding and undergraduate talent. This support enabled me to build a small CHWP prototype and launched my learning of mechanical, electrical, cryogenic, and design engineering. Under Akito’s active tutelage, I—along with a group of exceptionally talented undergrads, including Suhas Ganjam, Alex Madurowicz, and Mael Flament—built up the lab over time, and six years later, I can safely say that Akito has been and may forever be the most influential figure of my career as a scientist. Without his financial, intellectual, and professional support, this thesis would undoubtedly be less interesting and less impactful.

I want to thank Dr. Aritoki Suzuki, who really got me going with CMB instrumentation. Toki supervised and mentored me for my first couple of years in the group, and he taught me about *so* many things: how to run a wet dewar, how use an FTS, how to design and characterize antennas, TES bolometers, and anti-reflection coatings, how to program a sensitivity calculator, how to present a great slide deck, and much more. Most importantly, Toki demonstrated to me that with dedication, passion, and hard work, I can overcome any challenge.

Next, I want to thank my fellow UC Berkeley labmates, who have been amazing coworkers and teachers throughout my time in the campus lab. Thanks to Shawn Beckman for building the PB-2a WHWP with me and for working out the most difficult aspects of the mechanical assembly. Shawn has quite a talent for hardware design, fabrication, and testing, as well as for working with machinists and companies to meet tight specs. Without his partnership,

the WHWP would never have come to be. Thanks to John Groh for being a man of many (all?) trades for Simons Array and for being a close friend in the lab. John has provided a model for leadership on POLARBEAR that I have aspired to and learned from throughout my time working on PB-2a and PB-2b. Thanks to Neil Goeckner-Wald for answering (almost always seriously) my analysis- and science-related questions, for providing the data-driven advice I often needed for our hardware developments, and for also being a great friend and confidant both inside and outside of the lab. Thanks to Oliver Jeong, who has been a fellow optics guru and whose optics-tube related work was central to getting the PB-2b CHWP tested and deployed. Thanks to Ben Westbrook, who has been an ultra-friendly, ultra-helpful labmate and group leader over the years, and thanks especially to Ben for writing a GUI for BoloCalc. Ben's BoloCalc efforts have enabled a much broader user base for the software, and there is no way I could have written the GUI without his heroic effort. Thanks to Dick Plambeck, who has been a mentor and voice of solid reason throughout my graduate career. Most everything I know about coherent mm-wave detection, waveguides, feedhorns, mixers, mechanical oscillators, etc. comes from my time in Campbell Hall. In addition, Dick has taught me how to approach scientific problems with a balance of creativity and caution, a skill that I will continue to emulate throughout my career. And while I wish I could write a sentence to all my Berkeley labmates, let me instead simply say thank you to Aamir Ali, Jessica Avva, Tylor Adkins, Darcy Barron, Yuji Chinone, Lance Corbett, Kevin Crowley, Ari Cukierman, Nicole Farias, Sam Guns, Tijmen de Haan, Grant Hall, Nick Harrington, Kaori Hattori, Bill Holzapfel, Yasuto Hori, Nick Huang, Michael Ludlam, Aashrita Mangu, Mike Myers, Chris Raum, Paul Richards, Kaja Rotermund, Ian Shirley, Brian Steinbach, Leo Steinmetz, Giovanni Vizcarra, Nathan Whitehorn, Yuyang Zhou and many others I have worked with, if only briefly. It has been a tremendous pleasure being a part of the Berkeley group.

I want to thank the many amazing folks I've worked with at LBNL, in addition to Akito. First, let me thank Peter Ashton. When Peter showed up in 2017, he put my research into hyperdrive, as he was immensely helpful with CHWP construction and testing as well as with lab management tasks and undergraduate mentorship. Without Peter's help, it would have taken me *years longer* to commission the PB-2b modulator, and the task would have been a lot less fun too. Thanks to Paul Barton, who helped transform the PB-2b CHWP from a concept into a real-life object. Paul began providing much-needed engineering assistance to my hardware endeavors in 2017, and I am grateful to have absorbed valuable lessons about CAD design, hardware specifications, manufacturing tolerances, soldering, PCB design and testing, and many other things while working with him. I want to thank the hardworking, talented undergrads and interns I had the pleasure of working with at LBNL, including Richard Chen, Hawkins Clay, Alexander Droster, Andrew Fischer, Mael Flament, Chingam Fong, Samantha Gilbert, Arian Jadbabaie, Alex Madurowicz, Adam Rutkowski, and Danielle Sponseller. One of the best parts of my job has been coming to the lab and office to both teach and learn from the brightest young minds in the physics community. I want to *especially* thank Suhas Ganjam and Raymond Tat for their exceptional efforts on the CHWP system. I sometimes joke that my research is enabled by undergrads who are more talented than I

am, and in the case of Suhas and Raymond, this statement is maximally true. Thanks to Suhas for designing, prototyping, and teaching me about the CHWP motor system and for being a great fellow student in the early days of Akito’s lab, and thanks to Raymond for being a “jack of all trades” and for remaining amazingly consistent and productive for \sim two key years of CHWP assembly and testing. Thanks also to Mayuri Sathyanarayana Rao for being a fantastic lab partner and for showing me that exceptional competence and empathy can be leveraged into otherworldly teaching skills. Not many people explain things better than Mayuri does, and I often try to emulate her when leading a journal club or giving a presentation. And thanks also to Yuyang Zhou for taking on the PB-2c CHWP work and for pushing the project ambitiously, a job I never would have been able to do myself.

Next I want to thank my fellow PB-2b collaborators at UC San Diego (UCSD). While the PB-2b CHWP was built and tested at LBNL and was integrated into the optics tube on the UC Berkeley campus, much of its verification and readying for deployment took place at UCSD. Thanks to Bryce Bixler, who has been invaluable to much of the CHWP commissioning and will be in charge of the modulator after I am gone. I rest assured knowing that Bryce will carry this dissertation’s work forward to a science product. Thanks to Calvin Tsai for spending long days and nights with me in the UCSD lab, turning screws, sharing stories, and eating great tacos. Thanks to Praween Siritanasak for being the local expert on data-acquisition tasks and for being a big help when computer and network problems would creep up. Thanks to Tucker Elleflot for helping babysit and build up the CHWP when it first arrived in UCSD, and thanks to Jen Ito, Lindsay Lowry, and Logan Howe, whose leadership has enabled the entire PB-2b program. Thanks also to Professor Kam Arnold for leading PB-2b and for facilitating the construction and testing of the CHWP system at UCSD. Without the UCSD team, much of the work in this thesis would be either impossible or irrelevant!

I want to thank a few of the Japanese collaborators I’ve worked with more closely during my time on Simons Array. Thank you to Masaya Hasegawa for leading the PB-2a effort and for welcoming me with graciousness each time I’ve visited the KEK lab. Thank you to Yuki Inoue for spending several long days and nights to measure sapphire for the PB-2a WHWP. Thank you to Fred Matsuda (who was at UCSD but now is at Kavli IPMU) for providing a never-ending stream of optics simulations and for being one of the most fun collaborators to work with in Simons Array. Thank you to Satoru Takakura for being my go-to expert about polarization modulation-related matters. And thank you to Haruki Nishino for advising and helping with a wide variety of tasks related to HWP-associated data acquisition.

In addition to the KEK team, I want to thank Tomotake Matsumura and Yuki Sakurai, whose partnership has been truly invaluable throughout the CHWP’s development. Tomo is a pioneer of CHWPs for CMB polarimetry, and many of the modulator-related research products of this dissertation are, in a way, next-generation realizations of his pioneering work. In addition, Akito and I have consulted Tomo at various points throughout the PB-2b CHWP development, and his insights have been totally essential to the successes of our hardware designs. Yuki has been working with Tomo ever since I started working on waveplates, and during these years I’ve had so many fruitful collaborations and conversations

with him. Yuki is now building CHWPs for Simons Observatory, and I couldn't feel better about his shepherding the PB-2b design into bigger and better instruments for future CMB experiments.

Next, I'd like to thank several Simons Observatory (SO) collaborators that I've had the pleasure of working with over the years. First, let me thank the University of Michigan (now the UChicago) lab, with whom I've worked on sensitivity simulations and AR coating investigations. Thanks to Professor Jeff McMahon for always bringing an energetic, clever perspective that I could turn to for technical, interpersonal, or career advice. Thanks to Kevin Coughlin and Joey Golec for teaching me everything I know about metamaterial AR coatings and waveplates and for many fun conversations during my visits to Michigan. Thanks to Sara Simon for playing a central advising role to the presented sensitivity work and for teaching me much of what I know about feedhorns and beam-coupling efficiency. Thanks to Grace Chesmore and Shreya Sutariya for making a few of the AR coating measurements in this dissertation, and thanks to Carlos Sierra for being my most active and helpful BoloCalc user. Carlos is such a kind and thorough debugger! Other SO folks I'd like to thank: Sarah Marie Bruno at Princeton, with whom I got BoloCalc going for SO, Joe Seibert, who has also been helpful with BoloCalc among many other things, and Max Silva-Feaver, who has been a go-to readout expert as well as one of my closest friends in SO. There is a long list of others that I've enjoyed learning from during my time on SO, and while I wish I could list them all, I won't out of fear of forgetting folks!

I have had the pleasure of working with an exceptionally talented bunch of technical experts, without whom many of the hardware efforts in this thesis would not have been possible. Thank you to the UC Berkeley Physics machine shop team, including Warner Carlisle, Gordon Long, Tommy Gutierrez, and Abel Gonzalez, for grinding on several projects that other shops said couldn't be done. Thank you especially to Gordon Long, one of the most brilliant and important contributors to this dissertation's anti-reflection coating work. Consistently hitting $\sim 10\ \mu\text{m}$ tolerances on 20" parts composed of toxic, non-standard substrates is not a simple task, but no matter the challenge, Gordon always found a way to come through for me when I needed him to. Thank you to Chris Scholz, who worked tirelessly for months on the epoxy anti-reflection coatings for the POLARBEAR-2b lenses. I sincerely enjoyed my hours in the lab with Chris, and I learned so much from his attention to detail, his commitment to methodicalness, and his consistent temperament. Last but certainly not least, thank you to Cory Lee, who played an integral part in the CHWP's design and development, especially during the prototyping phase, and who machined a huge number of parts for me at LBNL. Much of what I know about how to translate CAD drawings into comprehensible machining instructions comes from Cory, and I also learned a lot from him about how to design practical, fabricable hardware assemblies.

Much of the work in this thesis would not have been possible without industry partnerships. I am grateful to dozens of companies that I've worked with over the years, but I would like to thank a few folks who have been exceptionally influential to the presented research. First, a hearty thank you to the team at GHTOT, including Carl Johnson, Dr. Yong Ji, Carl Zhang, and Valerie Wang, who provided the $\sim 510\ \text{mm}$ sapphire in this dissertation.

Back in 2014, when I first started working on the PB-2a warm half-wave plate, the lack of large-diameter sapphire at an affordable price (for science) was a crippling hindrance to the project. However, since meeting Dr. Ji at the SPIE Photonics West trade show in 2014, the prospect of half-wave plates for CMB science has totally transformed. I especially want to thank Carl Johnson, who has fielded my every question, whether over emails or beers, has accepted every one of my specs, has agreed to try many difficult manufacturing endeavors, and has, through it all, supported our science by offering low prices and short lead times. It is not an understatement to say that this thesis's work would not have been possible without GHTOT's partnership. Next, I want to thank ATZ GmbH in Germany, who fabricated the superconducting magnetic bearing (SMB) for the PB-2b cryogenic half-wave plate. Frank and Uta have been incredibly generous with their time and attention as I, a young scientist, learned how to interface my own crazy hardware ideas to the ATZ hardware. Frank went back and forth with me dozens of times, discussing the nitty-gritty technical and behavioral details of the SMB, and we are grateful to ATZ for providing such a high-end product fabricated with our needs at heart. Next, I want to thank Grey Brooks and Ryne Tacker at Laserod in Torrance. We rely on Laserod to strain relieve our anti-reflection coatings using their femtosecond, UV laser, and they have cut *dozens* of pieces for us over the years with the highest quality of work. A very special thanks to Grey, who was an active contributor to our epoxy R&D. Without Grey's rapid-turnaround machining, insights about the "side effects" of laser ablation, and willingness to answer my many phone calls and emails, the epoxy AR coatings for PB-2b and Simons Observatory would never have matured as much as they did.

While I would love to write a sentence for every company I've worked with, let me instead list several of my favorites: thanks to Dave Weeden, Russ Wooddell, and Joy Harper at Saint-Gobain Crystals for working on thin-sapphire grinding; thanks to Muhammad Suri at SMC for assisting me many times with the linear actuators in our CHWP's gripper; thanks to Teresa Ostroski and Socrates Gonis for helping us place large, non-standard orders of laminates and for helping us understand the material science of the composites; thanks to Brian Stull, Gayle Rakosky, and Justin Bond at Lord for teaching me how to optimize our adhesion promoter application processes; thanks to Brad Spraw and Diana Villanueva at Designetics for giving us a wide variety (including prototypes!) of applicator tips for our adhesion promoter processes; many thanks to Paul Tehan for providing critical consultations about the chemistry of epoxies; and thanks to Edwin Goyzueta, whose extra efforts and dedication enabled the development of the SF-0940 polyimide foam for CMB anti-reflections coatings, a technology that is gaining increased traction in the field and may become a game changer for CMB observation at low frequencies.

Without support from the amazing administrative staffs at UC Berkeley and LBNL, my research would've progressed slowly and painfully. I'd like especially to thank Carlos Bustamante, the UCB physics department's shipping and receiving guy (among other things), for always going above and beyond the call of duty, sending my shipments out quickly and getting critical packages to me promptly. Moving parts around the globe on tight timelines has been central to the success of my research, and I can't even imagine how lost I

would've been without Carlos's unwavering partnership. At LBNL, I want to thank Anytra Henderson, Steve Dardin, and Ken Wilson for answering my many questions and frequently helping with anything from administrative tasks to equipment needs. Behind every scientific publication lies a mountain of logistics and paperwork, and I am grateful to have had the support of amazing staffers during my research endeavors.

I want to offer a special thank you to my Boston University mentors and advisors, without whom my Berkeley experience would not have been possible. Thank you to Lee Roberts, to whom this dissertation is dedicated, for teaching me the beauty and fun of physics while I was still a trumpet major. Lee's guidance enabled my transition into the physics department and provided me with research opportunities along my way to grad school. Thank you to Eric Hazen for teaching me how to code (I've gotten lots of mileage out of that skill!) and for putting my pre-grad-school research experience into high gear. And thank you to Jim Rohlf for getting me involved with CMS, which was central to my graduate school applications. Thank you to Ken Lane and So-Young Pi, who, despite my deficiencies as a fast-tracked physics undergrad, believed in me and helped propel me into a successful physics career. Thank you to Terry Everson, who was fully supportive of my transition from fine arts into physics, despite his pouring his heart and soul (as he always does) into helping me become a professional trumpeter. Finally, thank you to my good friend Nafiu Awal, who first opened my eyes to how fun it is to be a physicist. My BU community laid the bedrock on which this dissertation's research is built, and I am forever grateful for their investments in me.

Next, I want to thank my friends within the Berkeley physics program who have made grad school an absolute blast. Thanks to Erik and Heather Urban for frequently spending late Friday nights to decompress after a long week in lab, and thanks to Chris and Chelsey Benson for many times listening to me vent for hours over pizza and beer. Thank you to Crystal Noel, who was so often exactly the roommate and confidant I needed after a long day, and thank you to Pourya Assem, Matt Gilbert, Claire Ziewitz, Justin Gerber, and Scott Ellis for also being supportive living partners and friends during the often-dark days of grad school. Thank you Alison Saunders for single-handedly carrying me (like a wrecking ball) through my first year at Berkeley, and thank you to many other of my cohort members for providing an intellectually rich and fun atmosphere during classes, preliminary exams, and beyond. Thank you also to Esther Cho, whose companionship and support have been a steadying force during my push to graduation.

Thank you to Word By Word, of which my mother is the founder and CEO, for proofing this dissertation. I'm grateful to have a second set of eyes to catch my misspellings and double words, but I'm even more grateful to have a mother who read my dissertation cover to cover. You're the best, Mom, and you're also now an expert in CMB instrumentation!

Finally, I want to thank my family, without whom pursuing a PhD would never have been possible. Thank you especially to my mom, dad, and sister Emily, who have been supportive of my workaholicism, reckless ambition, and often-tortured psyche throughout grad school. In addition to the innumerable sacrifices my family has made to get me a world-class education, they have been my confidants through both the low and high points of my past seven years.

Introduction

This dissertation is composed of three parts: background and motivation, instrument sensitivity and simulations, and polarization modulation.

In the first part, Chapters 1 and 2 provide scientific and technical context for the original research in the chapters that follow. Specifically, Chapter 1 overviews Big Bang cosmology, the theoretical paradigm of inflation, and how the cosmic microwave background (CMB) can be used to constrain the universe’s history and evolution. We focus in particular on a detection of primordial gravitational waves, which would both prove inflation and probe its physical mechanism. Chapter 2 reviews state-of-the-art CMB telescope technology and presents the trajectory of ongoing technological development. We broadly categorize CMB instrumentation into the observation site, telescope optics, detectors and readout, and cryostat thermal design, and we discuss how these subsystems enable the inflationary gravitational wave measurements presented in Chapter 1.

In the second part, Chapters 3-5 discuss how the sensitivity of CMB telescopes are calculated and how those calculations inform the design, fabrication, and evaluation of the optical, thermal, and detector subsystems. Chapter 3 details how photon, thermal, and amplifier noise impact detector performance, how to estimate the instantaneous sensitivity of detector arrays, and how those estimates influence science forecasts. While segments of this sensitivity calculation are well-established in the literature, we deliver a more comprehensive, top-to-bottom presentation to provide proper context for the following chapters. Chapter 4 presents the impact of Hanbury Brown-Twiss correlations on the sensitivity of densely packed CMB detector arrays. This calculation builds on the work of Zmuidzinas (2003) [241] and applies the concept of photon correlations to the problem of focal plane optimization. Chapter 5 presents a novel sensitivity calculator for CMB telescopes called BoloCalc. While CMB collaborations have traditionally written their own calculators, the need for a unified sensitivity simulator has arisen with the advent of large CMB conglomerations such as Simons Observatory and CMB-S4. BoloCalc is designed to satisfy this need, providing a flexible, thoroughly-featured, well-documented, open-source Python package for use across the CMB community.

In the third part, Chapters 6-10 discuss the mitigation of polarized systematic effects in the Simons Array (SA) telescopes using half-wave plate (HWP) polarization modulation. Chapter 6 overviews the theory, implementation, and optical requirements of continuously rotating sapphire HWPs, emphasizing the benefits and challenges of polarization modulation on SA. Chapter 7 presents an ambient-temperature HWP for SA’s first installment, POLARBEAR-2a (PB-2a), including its requirements, design, construction, and laboratory evaluation. PB-2a’s HWP is the largest deployed for CMB observation to date and introduces innovative techniques for sapphire characterization, multi-layer anti-reflection (AR) coatings, and HWP laboratory evaluation. Chapter 8 presents the design of a *cryogenic* HWP (CHWP) for SA’s second installment, POLARBEAR-2b (PB-2b), laying out its me-

chanical requirements and operational philosophies, and Chapter 9 presents the CHWP's laboratory evaluation. The PB-2b CHWP development advances the research area of large-diameter cryogenic rotation stages, and these advancements significantly improve the feasibility of polarization modulation for modern CMB experiments. Finally, Chapter 10 discusses cryogenic AR coating technologies for the PB-2b CHWP. The development of broadband, high-performance, cryogenically robust AR coatings for sapphire is among the most challenging tasks facing CMB instrumentation researchers today, and we discuss the design, fabrication, and evaluation of a few candidate AR technologies for SA. Chapters 6-10 together represent substantial advancements in hardware-based systematic error mitigation for CMB polarization measurements and better enable inflation science for various telescope systems.

In conclusion, Chapter 11 discusses the implications of the presented research and its impact throughout the broader CMB community.

Chapter 1

Scientific motivation

Our understanding of the universe continues to evolve, and the few-hundred-year rise of precision astronomy has propelled us towards an increasingly sophisticated theory of the cosmos. While there are many ways to probe the universe’s evolution—including mapping planet orbits, characterizing stars, inspecting galactic motions, and analyzing interstellar dust—the field of **cosmology**, or the study of the universe on its largest scales, has emerged as a cornerstone of modern physics.

The prevailing theory of contemporary cosmology is the so-called **Λ CDM model**, which states that the universe is composed of dark energy—quantified by the cosmological constant Λ —and cold dark matter (CDM), in addition to the familiar baryonic matter and radiation energy. A tenet of Λ CDM is the **Big Bang** theory, which asserts that the universe expanded from a hot, dense plasma into the statistically isotropic cosmic web of galaxies that we observe today. Big Bang cosmology has revolutionized our understanding of the universe and has been confirmed by multitudes of observations, including those of the **cosmic microwave background (CMB)**. The CMB is an isotropic microwave remnant of the Big Bang [155], whose **spatial anisotropies** [187] have been one of the most powerful tools to constrain the universe’s composition and history.

Despite Λ CDM’s success, it has several issues, including the so-called horizon, flatness, and magnetic monopole problems. An elegant solution to these issues is a theoretical paradigm called inflation [66], which predicts a period of super-luminal expansion $\sim 10^{-36}$ s after the Big Bang. If true, inflation would have imprinted a unique polarization signature onto the CMB [174, 235], motivating mm-wave observatories to deploy increasingly sensitive polarimeters. However, as telescope technologies and data analysis techniques improve, the upper limit on this inflationary signature is becoming buried beneath galactic foregrounds [210], necessitating not only an exquisite measurement of the CMB but also of interstellar dust [210] and synchrotron radiation [16]. Therefore, the CMB community is building increasingly sophisticated telescope arrays with larger detector counts and tighter control of systematic effects. In this chapter, we review Big Bang cosmology, introduce the cosmic microwave background, discuss inflation, and pose the experimental challenges that motivate this dissertation’s research.

1.1 Big Bang cosmology

General relativity relates spacetime’s curvature—parameterized by the Einstein tensor $G_{\mu\nu}$ —to its energy composition—parameterized by the energy-momentum tensor $T_{\mu\nu}$ —via the Einstein field equations

$$G_{\mu\nu} + \Lambda g_{\mu\nu} = 8\pi G T_{\mu\nu} , \quad (1.1)$$

where $g_{\mu\nu}$ is the metric tensor, Λ is the **cosmological constant**, G is the gravitational constant, and where we have set the speed of light to $c = 1$. Under the assumptions large-scale of homogeneity and isotropy, or that there are no “special” locations or directions in the universe,¹ the field equations are solved by the Friedmann-Robertson-Walker (FRW) metric, whose spacetime invariant is defined in hyperspherical coordinates as

$$ds^2 = -dt^2 + a^2(t)d\vec{x}^2 = -dt^2 + a^2(t) [dr^2 + S_k^2(r)d\Omega^2] . \quad (1.2)$$

Here, t is proper time, $a(t)$ is the **scale factor**, r is the hypersphere’s radial coordinate, Ω is solid angle with Jacobian $d\Omega^2 = d\theta^2 + \sin^2\theta d\phi^2$, and

$$S_k(r) \equiv \frac{\sin(\sqrt{k}r)}{\sqrt{k}} , \quad (1.3)$$

where k is the present-time spacetime curvature. The scale factor $a(t)$ quantifies the expansion of space as a function of time, and therefore it is convenient to define the **comoving distance**

$$\chi = \int_{t_1}^{t_2} \frac{dt}{a(t)} , \quad (1.4)$$

which for objects with zero peculiar velocity is constant in time. A depiction of the **comoving grid** and its relationship to proper distance are shown in Figure 1.1.

The solution to Einstein’s field equations in the FRW metric are

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

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

where the overdots denote time derivatives, ρ is energy density, and p is isotropic pressure. These **Friedman Equations** relate the universe’s expansion history to that of its constituents and are the cornerstone of Big Bang cosmology. The expansion factor at present time t_0 is defined to be $a(t_0) = 1$, and its evolution is parameterized by the **Hubble parameter**

$$H \equiv \frac{\dot{a}}{a} , \quad (1.7)$$

which is its logarithmic time derivative.

¹The assumption of homogeneity and isotropy is often called the **cosmological principle**.




Figure 1.1: A schematic of the comoving grid. In the standard cosmological paradigm, the scale factor $a(t)$ grows with time, enlarging the distance between otherwise stationary objects, which are represented by red circles. However, in comoving coordinates, the distance between these objects remains constant.

The **Hubble constant** H_0 , or the present-time value of the Hubble parameter, was first measured by Edwin Hubble in 1929 [88] using the redshift-distance relationship of extragalactic nebulae shown in Figure 1.2a.² Figure 1.2b shows a more recent constraint on H_0 using an assortment of modern measurements [8], revealing a value of $H_0 \approx 70$ km/s/MPc. Using the **equations of state** $w \equiv p/\rho = (1/3, 0, -1/3, -1)$ for (radiation, matter, curvature, dark energy), we can conveniently rewrite the Friedmann equations in terms of the Hubble constant as

$$H^2 = H_0^2 \left(\frac{\Omega_r}{a^4} + \frac{\Omega_m}{a^3} + \frac{\Omega_k}{a^2} + \Omega_\Lambda \right), \quad (1.8)$$

where $(\Omega_r, \Omega_m, \Omega_k, \Omega_\Lambda)$ are the present-time densities of (radiation, matter, curvature, dark energy) divided by the **critical density**

$$\rho_c \equiv \frac{3}{8\pi G} H_0^2. \quad (1.9)$$

The matter density Ω_m includes that of dark matter, and the radiation density Ω_r includes that of relativistic particles. The best current estimates of the universe's constituents—which come from the Planck satellite mission—are shown in Table 1.1.

1.2 Cosmic microwave background

One of the most prominent predictions of the Big Bang model is the existence of remnant background radiation. At early times, energy densities are high, and therefore the universe's

²While Hubble's measurement of the Hubble constant turned out to be inaccurate, it was the first firm evidence for an expanding universe and birthed the modern era of precision cosmological measurements.

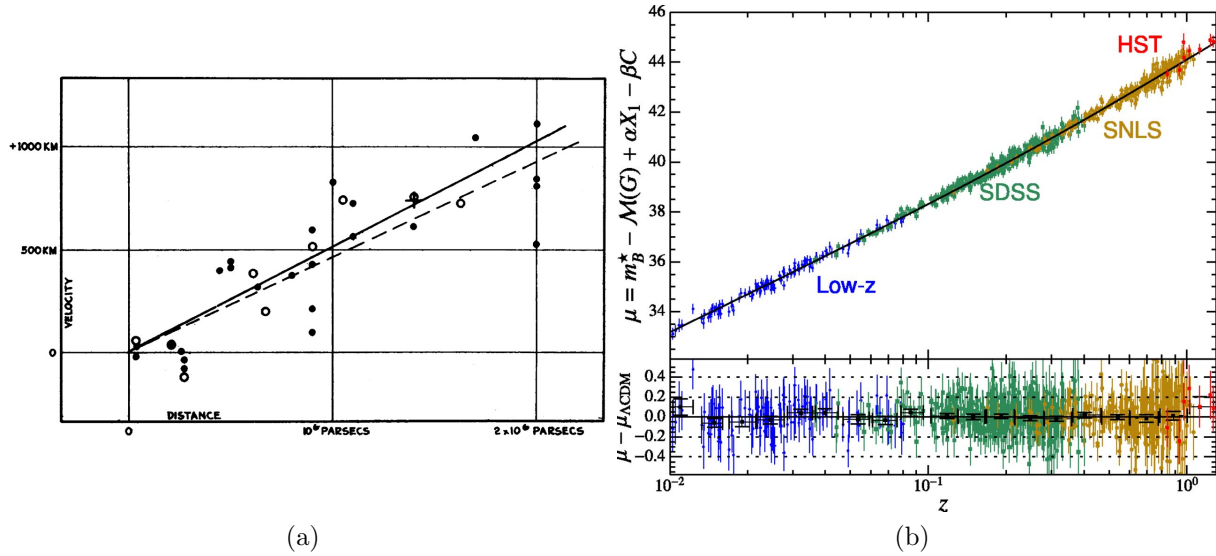


Figure 1.2: Two representations of the **redshift-velocity relationship**, which is quantified by the Hubble constant. Figure 1.2a shows Hubble’s original diagram, published in 1929 [88]. The filled circles and solid line show individual extragalactic nebulae and the resulting distance-velocity relationship, respectively. The open circles and the dotted line show these nebulae grouped and fit to. The positive slope demonstrates that the universe is expanding. Figure 1.2b shows a modern version of Hubble’s plot, with collections of data from various experiments—Low- z , the Sloan Digital Sky Survey (SDSS), the Supernova Legacy Survey, and the Hubble Space Telescope (HST)—which use supernovae to calibrate distances [8].

constituents (photons, electrons, positrons, protons, neutrons, and neutrinos) are in thermal equilibrium, forming what is often called the **primordial plasma**. As the universe expands, its energy dilutes (except for that of the cosmological constant) as in Equation 1.8, and light nuclei³ are **frozen out** of the primordial plasma according to each species’ binding energy. This process is called **Big Bang Nucleosynthesis**, and it determines much of what we know about cosmic particle abundances today.

As the universe cools to $T \lesssim 4000$ K, protons and electrons begin to combine and form neutral hydrogen, which in turn reduces the coupling between photons and baryonic matter. By the time the universe reaches $T \sim 3000$ K, enough neutral hydrogen has formed for the radiation and baryonic matter to be effectively decoupled, enabling photons to stream freely.⁴ This hydrogen-formation event is called **recombination**, and the remnant, free-streaming photons are called the **cosmic microwave background (CMB)**. After recombination, the CMB’s energy density dilutes as $\Omega_r/a^4(t)$, and its wavelength redshifts as

$$\lambda(z) = \frac{\lambda_0}{(1+z)}, \quad (1.10)$$

³Light nuclei typically include deuterium, helium, and lithium.

⁴The universe is often said to be “transparent” following hydrogen formation.

Parameter	Value
H_0	67.66 ± 0.42
$\Omega_b h^2$	0.02242 ± 0.00014
$\Omega_c h^2$	0.14240 ± 0.00087
Ω_Λ	0.6889 ± 0.0056

Table 1.1: The values of the cosmological parameters according to Planck [210]. The error bars represent 68% confidence limits, and Ω_k is consistent with zero, indicating a flat universe.

where λ_0 is the radiation’s wavelength today ($z = 0$), and where the **redshift** z is related to the scale factor as

$$a = \frac{1}{1+z}. \quad (1.11)$$

Therefore, given that the CMB decoupled at $T \sim 3000$ K and a redshift of $z \sim 1100$, we expect to observe a $T_{\text{CMB}} \sim 2.7$ K background radiation today.

The CMB was first discovered by Arno Penzias and Robert Wilson in 1964 [155] and was found to be uniform across the sky with a temperature 3.5 K higher than expected. Many years later in 1990, the CMB’s spectrum was measured by the Far Infrared Absolute Spectrometer (FIRAS) on the Cosmic Background Explorer (COBE) satellite [128]. The measurement, shown in Figure 1.3, was performed using a carefully calibrated Fourier-transform spectrometer (FTS) and is beautifully described by a blackbody spectrum

$$B_\nu(T) = \varepsilon \frac{2\nu^2}{c^2} \frac{h\nu}{e^{h\nu/(k_B T)} - 1}, \quad (1.12)$$

with emissivity $\varepsilon = 1$ and temperature $T = T_{\text{CMB}} = 2.73$ K. The current best estimate of the CMB’s temperature is

$$T_{\text{CMB}} = 2.7255 \pm 0.0006 \quad (1.13)$$

and is estimated from a combination of FIRAS and Wilkinson Microwave Anisotropy Probe (WMAP) satellite data [55].

In addition to FIRAS, the COBE satellite contained the Differential Microwave Radiometer (DMR), which made the first measurement of **CMB anisotropy** [186]. The CMB’s spatial intensity fluctuations were measured to be one part in ten thousand, demonstrating the background radiation’s extraordinary uniformity across the sky. Though small, these intensity fluctuations offer a “snapshot” of the universe $\sim 370,000$ years after the Big Bang, as the CMB’s hot and cold spots correspond to dark-matter over- and under-densities in the primordial plasma. As we will discuss in Section 1.5, CMB anisotropies trace the seeds of **structure formation**, making them among the most powerful observational phenomena in modern cosmology.




Figure 1.3: A preliminary spectrum of the CMB measured by the FIRAS instrument on the COBE satellite [128]. The box widths denote frequency bins for each data point, and the box heights denote a preliminarily stated error bar of 1%. The error bar of the final spectrum, published in 1994, was orders of magnitude smaller [56].

1.3 Inflation

Despite its success, Λ CDM fails to explain several outstanding mysteries. Some of these mysteries are at the interface of cosmology and particle physics, such as the particle nature of dark matter and dark energy. However, some standard cosmological shortcomings are “internal” inconsistencies, where the model’s assumptions contradict the observations that support it. The most prominent of these problems are:

- **Horizon problem:** the universe appears to have equilibrated across super-horizon length scales, or distances larger than could have ever been in causal contact.
- **Flatness problem:** The universe observed today is very flat, and because any deviation from flatness grows with time (see Equation 1.8), the universe must have begun with a flatness of one part in $\sim 10^{14}$, posing a fine-tuning issue.
- **Magnetic monopole problem:** magnetic monopoles are predicted by many grand unified theories, but we do not observe any.
- **Initial conditions:** while Λ CDM describes how density perturbations in the early universe evolved into the large-scale structure that we see today, it does not describe how these perturbations were seeded.




Figure 1.4: A qualitative depiction of the relationship between the Hubble radius—or equivalently the causal horizon—in units of comoving distance vs. the scale factor during both the inflationary and standard cosmological epochs. During inflation, the Hubble radius shrinks, moving length scales that were once in causal contact outside the causal horizon, “freezing” their evolution. During the standard cosmological evolution, the Hubble radius grows, and as these length scales reenter the horizon, they begin to evolve once again. The horizontal lines show that large-scale fluctuations exit/reenter the horizon first/last while small-scale fluctuations exit/reenter the horizon last/first. The kink in the Hubble radius at late times represents the transition from a radiation-dominated universe to a matter-dominated one.

There have been many proposed extensions to Λ CDM, but the most popular is **inflation** [66], which states that the universe went through a brief period of exponentially rapid expansion before the standard cosmological evolution began.

Figure 1.4 shows a schematic of how inflation impacts causal distances. During inflationary expansion, all energy other than that of the inflation field—often called the **inflaton**—is diluted away, enabling spacetime to expand exponentially.⁵ Because an exponential rate of expansion can surpass the speed of light, the **Hubble radius**

$$R_H \equiv 1/H, \quad (1.14)$$

or the length scale that could have at any time been in causal contact,⁶ *shrinks* in comoving coordinates.⁷

⁵This exponential expansion is not so different from that due to the dominance of dark energy at present time.

⁶Note that $1/H_0$ roughly corresponds to the size of the observable universe today.

⁷Stated another way, when two causally connected points move away from one another faster than the speed of light, they can no longer “see” each other.

The dynamics of the inflationary epoch are governed by the inflaton's evolution. If we assume that the inflaton is a scalar field ϕ , as is the case in many inflationary theories, the Einstein field equations become

$$\rho_\phi = \frac{1}{2}\dot{\phi}^2 + V[\phi] \quad (1.15)$$

$$H^2 = \frac{\rho_\phi}{24\pi G} \quad (1.16)$$

$$\ddot{\phi} + 3H\dot{\phi} + V'[\phi] = 0, \quad (1.17)$$

where $V[\phi]$ is the inflationary potential and ρ_ϕ is the inflaton's energy density. The final equation, which governs the inflaton's temporal evolution, is that of a harmonic oscillator damped by the Hubble parameter.

In order to “solve” Λ CDM's flatness and horizon problems, the inflaton must obey several slow-roll conditions,⁸ which are typically parameterized by the **slow-roll parameters**

$$\epsilon \equiv \frac{m_{\text{pl}}^2}{2} \left(\frac{V'[\phi]}{V[\phi]} \right)^2 \ll 1 ; \quad \eta \equiv m_{\text{pl}}^2 \left| \frac{V''[\phi]}{V[\phi]} \right| \ll 1, \quad (1.18)$$

where $m_{\text{pl}} \equiv \sqrt{\hbar c/G} \sim 10^{19}$ GeV is the Planck mass. These slow-roll parameters in turn determine the number of **e-folds**—or the logarithmic duration—over which inflation occurred

$$N_e = \frac{1}{m_{\text{pl}}} \left| \int_{\phi_1}^{\phi_0} \frac{d\phi}{\sqrt{2\epsilon(\phi)}} \right| \simeq 60 - \log \frac{10^{16} \text{ GeV}}{V^{1/4}}, \quad (1.19)$$

where (ϕ_0, ϕ_1) is the inflaton value at the (beginning, end) of inflation, and V is the **energy scale of inflation**. The number of e-folds needed to rectify the standard model is $N_e \simeq 60$, which constrains possible values for V .

As an instructive example, let us consider Alexei Starobinsky's R^2 inflationary potential $V[\phi]$ [190, 220]. This particular theory is classified as a **small-field** model⁹ and is derived from a proposed modification to the Einstein-Hilbert action

$$S = \frac{1}{2\kappa} \int R \sqrt{|g|} d^4x \rightarrow S_{R^2} = \frac{1}{2\kappa} \int \left(R + \frac{R^2}{6M^2} \right) \sqrt{|g|} d^4x, \quad (1.20)$$

where $\kappa \equiv 8\pi G$, R is the Ricci curvature, and M is the model parameter. Solving the R^2 action gives the potential

$$V_{R^2}[\phi] = V_0^4 \left[1 - \exp \left(-\sqrt{2/3} \phi / m_{\text{pl}} \right) \right]^2, \quad (1.21)$$

where V_0 is a constant. The R^2 potential is shown in Figure 1.5 along with a cartoon of how the inflaton evolves according to Equation 1.17. At $t = 0$, the inflaton is said to be in a

⁸The term “slow-roll” intentionally conjures up the familiar dynamics of a particle in a potential well.

⁹“Small-field” and “large-field” broadly categorize the energy scale of a given inflationary potential.




Figure 1.5: The R^2 inflationary potential, which is an instructive example of slow-roll inflation. The inflaton slowly “rolls” down the potential, enabling exponential expansion, and during reheating, it decays and its kinetic energy is converted into the particles of the standard model.

“false vacuum,”¹⁰ and as inflation begins, $V[\phi]$ slowly decreases such that the kinetic term $3H\dot{\phi} \approx 0$, enabling exponential expansion. However, as $\phi/m_{\text{pl}} \rightarrow 0$, H becomes very large and the inflaton’s energy decays into the constituents of the standard model. This phase transition is called **reheating** and marks the end of inflation and the beginning of standard cosmological evolution.

There are several compelling aspects of the inflation paradigm. First, the inflationary expansion of spacetime allows the entire observable universe to have thermalized prior to the Λ CDM evolution, solving the horizon problem. Second, inflationary expansion would have vastly suppressed any initial value for Ω_k , permitting the flat universe we observe today. Third, inflationary growth would have dramatically reduced the density of magnetic monopoles, presently making them very rare. And finally, Gaussian-random **quantum fluctuations** of the inflaton field would have rippled spacetime, seeding the density fluctuations that generate CMB anisotropies.

Despite inflation’s elegance and wide popularity, it is not the only proposed extension to the standard model. A prominent competitor is Ekpyrosis [219], which states that the universe is trapped in a never-ending cycle of expansion and collapse. While Ekpyrosis also solves the horizon, flatness, magnetic monopole, and initial conditions problems, there

¹⁰Generally speaking, an inflationary false vacuum can be a saddle point, a local minimum in $V[\phi]$, or a flat region, such as in the R^2 model.

is one signature that, if observed, would be smoking-gun-proof of the inflationary paradigm: primordial gravitational waves.

1.4 Primordial perturbations

Perturbations to the metric are typically classified into three categories: scalar, vector, and tensor [87]. Scalar perturbations arise from density fluctuations, which are sourced by gravity and therefore evolve with underlying dark matter. Vector perturbations arise from vortical motions of matter, which are not enhanced by gravity and therefore rapidly die out as the universe expands. Tensor perturbations are transverse-traceless perturbations to the metric $g_{\mu\nu}$ that propagate by stretching and compressing spacetime [87] and are often referred to as **gravitational waves**. While both inflation and Ekpyrosis produce Gaussian-random scalar fluctuations, only inflation could have produced a background of primordial gravitational waves.

It is both convenient and powerful to discuss perturbations δ in terms of their **power spectrum** $P_\delta(\vec{k})$ as

$$\langle \delta(\vec{k})\delta(\vec{k}') \rangle \sim \frac{\delta^3(\vec{k} - \vec{k}')}{k^3} P_\delta(\vec{k}), \quad (1.22)$$

where the \sim denotes the limit of no “mode coupling,” $k = |\vec{k}|$, and $\delta^3(\vec{x})$ is the three-dimensional Dirac delta function. Here, k is a wavenumber, and therefore **large k corresponds to short length scales** and vice versa. As ϕ evolves¹¹ during inflation and its energy $V[\phi]$ decreases, as shown in Figure 1.5, perturbations exit the horizon from low k to high k , as shown in Figure 1.4. Therefore, the **scalar spectrum of primordial perturbations** $P_\zeta(k)$ is “tilted” and can be written as

$$P_\zeta(k) \equiv A_s \left(\frac{k}{k_*} \right)^{n_s(k)-1} \sim \frac{H_*^2}{\epsilon_* m_{\text{pl}}^2}, \quad (1.23)$$

where k_* is a mode number often chosen as $k_* = 0.05 \text{ Mpc}^{-1}$, and where H_* and ϵ_* denote the Hubble parameter and first slow-roll parameter, respectively, when the mode k_* exited the horizon during inflation. The spectral tilt is parameterized by the **scalar index**

$$n_s(k) \equiv 1 + \frac{d \log P_\zeta(k)}{d \log k}. \quad (1.24)$$

In a similar manner, the **tensor spectrum** $P_t(k)$ can be written as

$$P_t(k) \equiv A_t \left(\frac{k}{k_*} \right)^{n_t(k)} \sim \left(\frac{H_*}{m_{\text{pl}}} \right)^2, \quad (1.25)$$

¹¹As a reminder, a given k -mode is able to evolve as long as its wavelength $\lambda = k/2\pi$ is shorter than the Hubble radius R_H such that $k \lesssim aH$.




Figure 1.6: Current constraints on the tensor-to-scalar ratio $r_{k_*=0.002}$ and the scalar index n_s to 95% confidence according to the Planck 2018 release [207], the BICEP/Keck 2015 release [202], and baryon acoustic oscillation data from the 6dF [17] survey and the Sloan Digital Sky Survey [164]. Contours of various inflationary potentials are also plotted for reference. The convex/concave line delineates the large/small-scale inflation regions.

where the tensor index is defined to be

$$n_t(k) \equiv \frac{d \log P_t(k)}{d \log k}. \quad (1.26)$$

Using these conventions, we expect $n_s < 1$ and $n_t < 0$, and we expect each to have a slightly negative slope with k .

The primordial tensor amplitude is typically quantified by the **tensor-to-scalar ratio**

$$r_{k_*} \equiv \frac{P_t(k_*)}{P_\zeta(k_*)} = \frac{A_t}{A_s} \sim 16 \frac{V_*}{A_s m_{\text{pl}}^4}, \quad (1.27)$$

where V_* is the inflationary energy scale when mode k_* exits the horizon. In turn, the **energy scale of inflation** can be written in terms of the tensor-to-scalar ratio as

$$V_*^{1/4} = 10^{16} \text{ GeV} \left(\frac{r_{k_*}}{0.01} \right)^{1/4}, \quad (1.28)$$

and therefore a detection of r would be a direct measurement of inflation and its energy scale.¹² Figure 1.6 shows the 95%-confidence contours on the tensor-to-scalar ratio and the

¹²On the contrary, if V_* is much lower than the grand unified energy scale 10^{19} GeV , then r may be undetectably small.

scalar index from the Planck satellite and the BICEP/Keck (BK) experiment. The leading constraint on inflationary gravitational waves comes from the latest (as of December 2020) BK data release [203] and is

$$r < 0.072 \text{ (95\% CL)}, \quad (1.29)$$

which Figure 1.6 shows is compatible with Starobinsky’s R^2 inflation.

1.5 CMB anisotropies

As discussed in Section 1.2, the CMB is a snapshot of the primordial plasma at the epoch of recombination, and its fluctuations trace perturbations in the underlying metric. The latest full-sky maps of CMB temperature and polarization anisotropies from Planck are shown in Figure 1.7, and the statistics of these fluctuations are powerful probes of both Λ CDM and inflation.

Because the CMB fills the sky, we can conveniently decompose its temperature fluctuations¹³ into **spherical harmonics** $Y_{\ell m}(\hat{n})$ as

$$\Delta T(\hat{n}) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m} Y_{\ell m}(\hat{n}), \quad (1.30)$$

where \hat{n} denotes a sky location, (ℓ, m) are the **angular mode** numbers, and $a_{\ell m}$ are the mode amplitudes. As predicted by inflation, CMB temperature anisotropies are isotropic¹⁴ and largely Gaussian,¹⁵ and therefore their statistics are well-described by their power spectrum C_ℓ as

$$\langle \Delta T(\hat{n}) \Delta T(\hat{n}') \rangle = \frac{1}{4\pi} \sum_{\ell=0}^{\infty} (2\ell + 1) C_\ell P_\ell(\hat{n} \cdot \hat{n}'), \quad (1.31)$$

where $P_\ell(\cos \theta)$ are Legendre polynomials and the mode variances are

$$\langle a_{\ell m} a_{\ell' m'}^* \rangle = \delta_{\ell m} \delta_{\ell' m'} C_\ell. \quad (1.32)$$

In this formalism, the power spectrum amplitude C_ℓ measures the variance of the m -values within mode ℓ and therefore quantifies fluctuation vs. **angular scale**. Because $\ell = 1$ corresponds to the dipole, the relation

$$\theta_{\text{sky}} \sim \frac{180^\circ}{\ell}, \quad (1.33)$$

relates mode number to angular scale on the sky.¹⁶

¹³A CMB temperature fluctuation modulates the blackbody spectrum in Figure 1.3, which appears as an *intensity* fluctuation when observing at a fixed frequency. Therefore, CMB scientists often use the terms “intensity” and “temperature” interchangeably.

¹⁴Except for the dipole, which arises due to our local motion with respect to the cosmic rest frame.

¹⁵As we discuss in Section 1.7, secondary anisotropies give rise to small levels of non-Gaussianity.

¹⁶For reference, the moon has an angular extent of $\approx 1^\circ$, corresponding to $\ell \sim 180$.

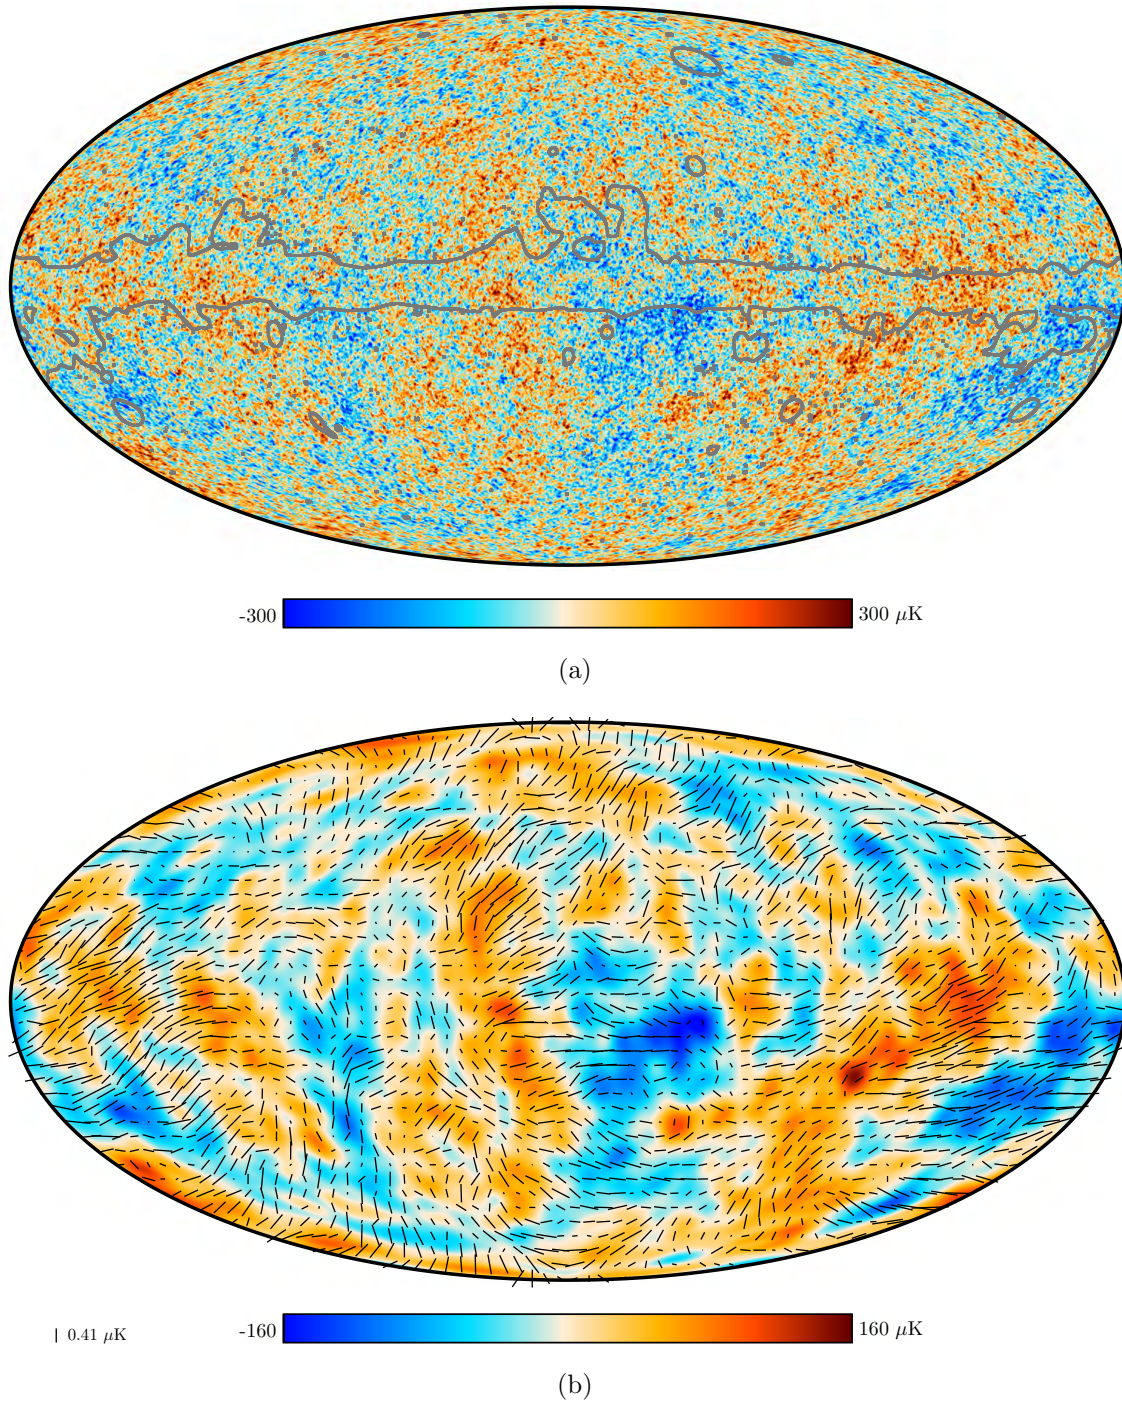


Figure 1.7: Full-sky maps of the CMB as presented in the Planck satellite’s 2018 data release [207]. Figure 1.7a shows temperature anisotropies, with foreground emission near the galactic plane removed. Figure 1.7b shows polarization anisotropies smoothed to an angular resolution of 5° for visual clarity [207].




Figure 1.8: A schematic of how temperature anisotropies give rise to linear CMB polarization. Hot/cold spots correspond to gravitational over/underdensities, and as photons climb/fall out of these potential wells, they lose/gain momentum $p = h\lambda$, which in turn causes their wavelength to shift down/up. Therefore, “hot” photons emerge from *underdensities*, while “cold” photons emerge from *overdensities*. At recombination, these photons Thomson scatter from ionized particles, and when the temperature anisotropy is quadrupolar $|\vec{E}_1| < |\vec{E}_2|$ about a scattering electron e^- , fractional linear polarization is generated.

There are only $2\ell + 1$ m ’s for a given ℓ , imposing a limit to how well one can measure C_ℓ . This limit is called **cosmic variance** and is defined as

$$\Delta C_\ell = \sqrt{\frac{2}{(2\ell + 1)f_{\text{sky}}}} C_\ell, \quad (1.34)$$

where f_{sky} is the measured **sky fraction**.¹⁷ Cosmic variance has two important implications for CMB anisotropy measurements. First, because $\Delta C_\ell \propto C_\ell$, one cannot estimate the CMB’s angular power spectrum with infinite precision, no matter how big the signal-to-noise. Second, because larger angular scales have fewer m ’s, their measurement is more likely to become **cosmic variance limited**, motivating larger sky coverage when measuring at low- ℓ .

In addition to temperature anisotropies, the CMB also contains polarization fluctuations. The CMB is polarized by Thomson scattering of **quadrupolar anisotropies** at the last scattering surface [87], as shown in Figure 1.8. The resulting polarization pattern on the sky is conveniently decomposed into so-called **E-modes** and **B-modes** [175, 235], which in the

¹⁷The sky fraction is $0 \leq f_{\text{sky}} \leq 1$ and quantifies the (effective) sky area both observed and utilized during power-spectrum estimation.

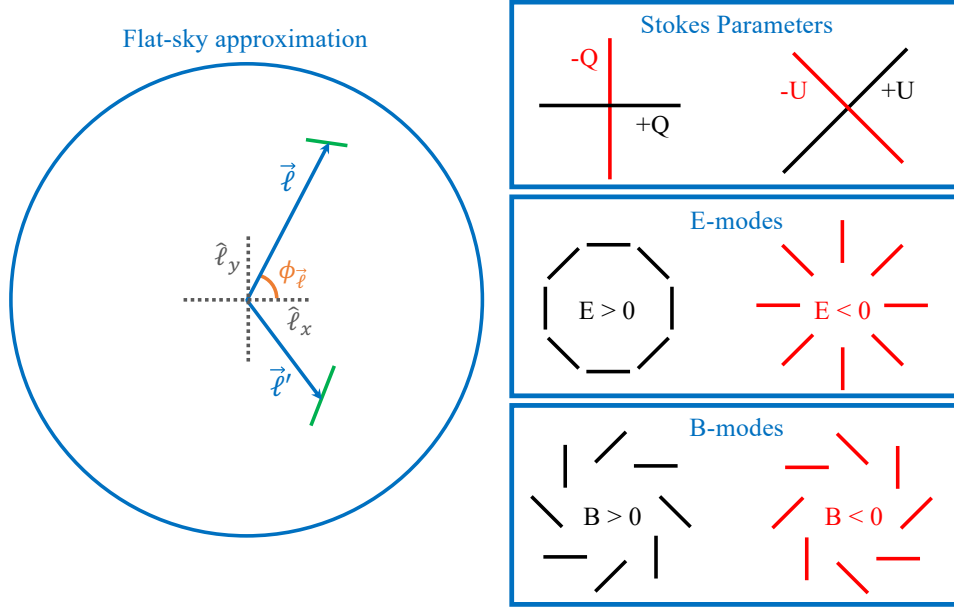


Figure 1.9: A schematic of the flat-sky approximation, which quantifies the spherical-harmonic modes using vector $\vec{\ell}$ instead of (ℓ, m) . In this scheme, the Stokes vectors Q/U are converted to E-modes and B-modes using Equation 1.36. As shown in the right panel, E-modes are parity-even, while B-modes are parity-odd.

flat-sky approximation¹⁸ are written as

$$E(\vec{\ell}) = Q(\vec{\ell}) \cos(2\phi_{\vec{\ell}}) + U(\vec{\ell}) \sin(2\phi_{\vec{\ell}}) \quad (1.35)$$

$$B(\vec{\ell}) = U(\vec{\ell}) \cos(2\phi_{\vec{\ell}}) - Q(\vec{\ell}) \sin(2\phi_{\vec{\ell}}), \quad (1.36)$$

where $Q = (T_x - T_y)/2$ and $U = (T_a - T_b)/2$ are the linear-polarization **Stokes parameters** shown in Figure 1.9, $\vec{\ell}$ is a two-dimensional Fourier vector, which in the flat-sky approximation replaces the (ℓ, m) indices, and $\phi_{\vec{\ell}}$ denotes the mode vector's orientation. There are two important properties of the E/B composition. First, they are spin-2 fields, meaning they flip sign under a $\pi/2$ rotation. Second, unlike with Q and U , E and B are coordinate-system independent but are instead defined by their symmetry. As shown in Figure 1.9, E-modes have zero curl and non-zero divergence, while B-modes have zero divergence and non-zero curl. In other words, B-modes have a handedness while E-modes do not.

The E-mode/B-mode decomposition is a powerful tool for studying CMB polarization. While primordial gravitational waves generate both E-modes and B-modes, density fluctuations only generate E-modes [175], making **primordial B-modes** a **null channel** through which primordial gravitational waves can be probed. In other words, a primordial B-mode detection is not limited by cosmic variance and would therefore be definitive proof of inflation. In tandem, E-modes offer their own wealth of information about the universe and

¹⁸The flat sky approximation is valid given a sufficiently small sky area and is often used for analysis of ground-based datasets.




Figure 1.10: The power spectrum of CMB temperature anisotropies from the Planck satellite [207]. The fit corresponds to the Λ CDM model, and the exquisite precision over a wide range of angular scales is largely cosmic-variance limited.

its evolution. Just as temperature anisotropies trace density fluctuations, E-modes trace the primordial plasma’s **velocity field**, and because the CMB is only $\sim 10\%$ polarized, the E-mode signal is substantially smaller than that of intensity, reducing cosmic variance and enabling tighter cosmological constraints.

1.6 CMB power spectrum

As discussed in Section 1.5, CMB temperature fluctuations are largely Gaussian and are therefore characterized by their power spectrum C_ℓ . These angular fluctuations are often plotted as

$$D_\ell^{XX} = \frac{\ell(\ell+1)C_\ell^{XX}}{2\pi}, \quad (1.37)$$

where $XX = (TT, EE, BB, TE, TB, EB)$ represent various auto- and cross-spectra between temperature, E-mode, and B-mode fluctuations.¹⁹

The temperature power spectrum, as measured by the Planck satellite [208], is shown in Figure 1.10. At $\ell \lesssim 100$, or on angular scales $\theta_{\text{sky}} \gtrsim 1^\circ$, the spectrum is nearly flat. These modes are outside the horizon at the time of last scattering and therefore have not evolved since the inflationary epoch.²⁰ Therefore, in principle, these low- ℓ modes could be used to measure r , as primordial gravitational waves generate temperature fluctuations. However, D_ℓ^{TT} is dominated by gravitational redshifting, also called the **Sachs-Wolfe Effect**, at low- ℓ , and its characterization is limited by cosmic variance.

¹⁹The addition of the $\ell(\ell+1)/2\pi$ factor is conventionally applied such that D_ℓ^{TT} is relatively flat.

²⁰Just as for k -modes, small/large ℓ -modes are the last/first to exit during inflation and are the first/last to reenter the horizon during standard evolution, as shown in Figure 1.4.




Figure 1.11: The power spectrum D_ℓ^{EE} of the CMB E-mode anisotropies as measured by Planck [207]. The oscillatory features are due to the BAO velocity field, and the fit uses the same cosmological parameters as that of Figure 1.10.

At $\ell \gtrsim 100$, modes reenter the horizon before recombination and evolve before the CMB decouples. **Baryon acoustic oscillations (BAO)** are acoustic waves that arise from gravitation-induced compression and photon-induced rarefaction in the primordial plasma. Peaks in the D_ℓ^{TT} spectrum between $100 \lesssim \ell \lesssim 1500$ correspond to scales of maximum rarefaction (odd peaks) and maximum compression (even peaks). When the CMB decouples, baryonic matter is “frozen” into place, and its oscillatory signature is imprinted onto the CMB anisotropies. As the universe evolves after recombination, the underlying dark matter densities accumulate baryons, forming the BAO-seeded large-scale structure that we observe today. At $\ell \gtrsim 1000$, BAO are suppressed by **Silk damping**, which arises when an oscillation’s length scale is shorter than the primordial plasma’s photon diffusion length.²¹ As is evident in Figure 1.10, the measurement of D_ℓ^{TT} is exquisite and is described to high precision by the standard cosmological parameters.

As measurements of D_ℓ^{TT} become increasingly cosmic-variance limited, CMB polarization measurements become increasingly important. The EE spectrum, shown in Figure 1.11, is also generated by BAO but is the result of photon velocities, resulting in peaks that are hundreds of times smaller than and out of phase with those in D_ℓ^{TT} . There is no polarization analog to the Sachs-Wolfe Effect, and therefore D_ℓ^{EE} from $10 \lesssim \ell \lesssim 100$ has a slope of $\simeq \ell(\ell+1)$, corresponding to a (nearly) scale-independent variance C_ℓ^{EE} . The excess anisotropy at $\ell < 10$ corresponds to modes that reentered the horizon after the epoch of **reionization**, when young stars reionize the interstellar medium at redshift at $z = 7.68 \pm 0.79$ [209], “re-illuminating” very-large-angular-scale fluctuations at very late times.

²¹The diffusion length can also be thought of as the photon’s mean free path.




Figure 1.12: The power spectrum D_ℓ^{BB} of CMB B-mode anisotropies as measured by several experiments. The solid line shows the gravitational lensing of E-modes into B-modes, the lower dotted line shows primordial B-modes with $r = 0.06$, and the upper dotted line shows the sum of the two. This figure is courtesy of Yuji Chinone.

The B-mode power spectrum D_ℓ^{BB} is shown in Figure 1.12 and has two components. First, **lensing B-modes**, which are represented by the solid line, arise due to **gravitational lensing** of E-modes into B-modes. Large-scale structure distorts the metric between the surface of last scattering and the Earth, deflecting CMB photons, mixing E-modes and B-modes, and generating BB power that peaks at $\ell \simeq 1000$. Second, primordial B-modes, which are represented by the lower dotted line, are hypothetically generated by inflationary gravitational waves and peak at $\ell \simeq 100$. Because tensor modes are not enhanced by gravitational structure, they are quickly washed out after reentering the horizon, suppressing D_ℓ^{BB} at $\ell \gtrsim 100$. While lensing B-modes are becoming well-characterized, primordial B-modes remain undetected and the value of r remains unknown.

1.7 Foregrounds

The B-mode spectrum D_ℓ^{BB} , while theoretically a null probe of inflationary gravitational waves, is in reality obfuscated by **foregrounds**. Therefore, an accurate characterization and removal of these foregrounds is becoming increasingly critical to a robust primordial B-mode measurement.

The first B-mode foreground is polarized emission from the galaxy. The interstellar medium (ISM) is filled with microscopic grains of **galactic dust** that absorb and emit

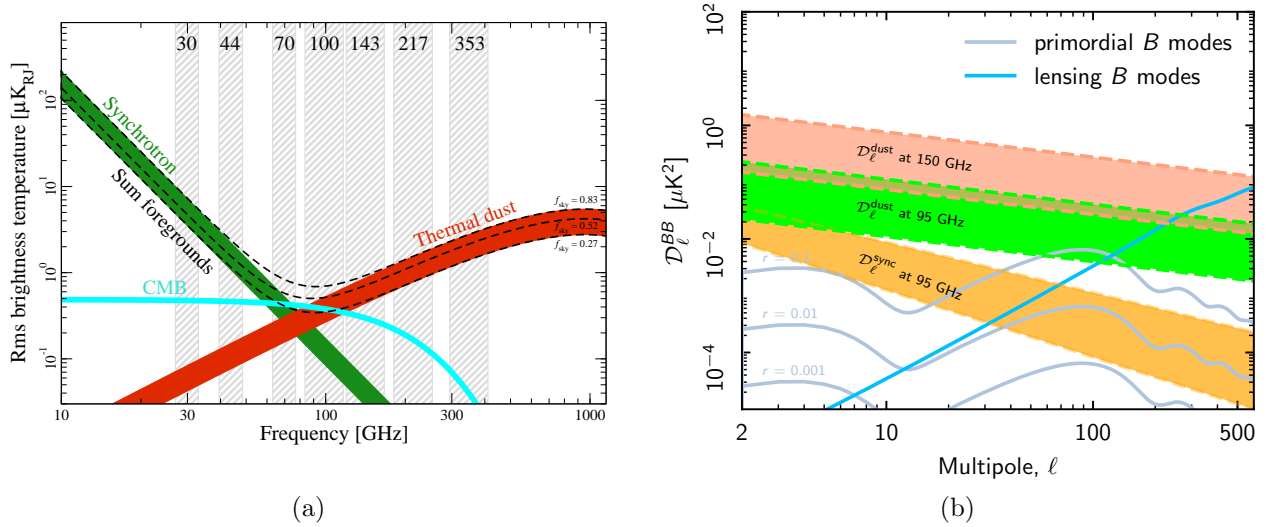


Figure 1.13: The brightness temperature and angular power spectra of galactic foregrounds compared to those of the CMB. Figure 1.13a shows the spatially averaged foreground brightness vs. frequency for various fractions of the sky [208]. Figure 1.13b shows the angular power spectrum of dust and synchrotron emission at two common CMB observation frequencies, compared to the lensing B-mode spectrum and a possible primordial B-mode spectra [210]. The shaded regions represent variations with sky fraction.

thermal radiation [218]. While this dust is concentrated around the galactic plane, it is bright enough to contaminate even the cleanest portions of the sky [210]. In addition, the ISM is filled with **synchrotron radiation** generated by relativistic charged particles whirling through the galaxy’s magnetic field. While the galactic magnetic field is not completely understood, it is known to have a coherent structure, and therefore synchrotron radiation is highly polarized [16]. Dust and synchrotron’s brightness vs. microwave frequency are shown in Figure 1.13a, while their angular power spectra are shown in Figure 1.13b [210]. Even near the **foreground minimum** at ~ 100 GHz, galactic foregrounds are large compared to the inflationary BB signal when $r \lesssim 0.1$ and therefore must be carefully removed in the search for primordial gravitational waves.

The second B-mode foreground is the gravitational lensing of E-modes into B-modes [236]. While gravitational lensing is a powerful tool to measure the gravitational potential along the line of sight [174] and reveal a precise evolutionary history of dark matter, dark energy, and cosmic neutrinos [86, 118], **lensing B-modes** obfuscate the primordial B-mode signal. As shown in Figure 1.12, when $r \lesssim 0.01$, lensing B-modes dominate at $\ell \gtrsim 50$ and must be subtracted via a process called **delensing** [104, 107, 178, 231]. An alternative strategy to delensing is to measure the reionization peak $\ell \lesssim 10$, but such large angular scales are difficult to access from the ground due to atmospheric low-frequency noise and limited sky coverage.²²

²²Nonetheless, the reionization peak is being targeted by some ground-based experiments, including the Cosmology Large Angular Scale Surveyor (CLASS) [74, 73] and GroundBIRD [146].

1.8 State of the field

The search for inflationary B-modes is among the most lucrative in modern-day physics. While CMB anisotropies are used to study a wide range of phenomena [204], an exquisite measurement of primordial B-mode polarization is the “holy grail” of CMB cosmology. For $r = 0.01$, the peak of D_ℓ^{BB} corresponds to fluctuations of ~ 100 nK, which is $\gtrsim 100$ *million* times fainter than emission from the atmosphere. Therefore, an accurate B-mode measurement is an enormous technical challenge that requires the most sophisticated telescopes, observation strategies, and analysis techniques.

The current generation of CMB observatories are often referred to as **Stage-3** experiments and include Simons Array [196], Advanced ACTPol [38, 32, 120], CLASS [234], SPT-3G [188], and BICEP Array [172]. These experiments employ $\mathcal{O}(1,000\text{--}10,000)$ detectors and have brought the E-mode spectrum into focus while also detecting lensing B-modes with high confidence. Even as these Stage-3 experiments continue to collect cutting-edge data, a new generation of **Stage-4** experiments, including Simons Observatory [59, 215] and CMB-S4 [205], aims to deploy $\mathcal{O}(100,000\text{--}1,000,000)$ detectors and make the ultimate ground-based measurement of CMB polarization. These increasingly sophisticated observatories require increasingly sophisticated technologies to meet their scientific objectives, motivating the research advancements presented in this dissertation.

Chapter 2

CMB instrument overview

High-precision measurements of CMB temperature and polarization anisotropies are among the most powerful tools to study the birth and evolution of the universe. Such measurements require telescopes that can map the CMB with high signal-to-noise and low systematic bias, and therefore cutting-edge instrumentation is at the heart of cutting-edge cosmology. In this chapter, we overview how the CMB is measured, and we focus on four areas of mm-wave instrumentation that are particularly relevant to this dissertation’s research: the atmosphere, optics, detectors, and infrared (IR) blocking. To provide concrete examples, we present these subsystems in the context of the Simons Array and Simons Observatory.

2.1 Observing the CMB

The fundamental goal of any CMB telescope is to collect as many CMB photons as possible and focus them onto a detector array that generates a crisp sky image. Examples of CMB telescope systems are shown in Figures 2.3 and 2.4, but before we detail their hardware components, we first provide a high-level overview their basic functionalities.¹

As CMB photons stream towards the telescope, they first traverse the earth’s atmosphere, whose water and oxygen absorb at millimeter wavelengths (Section 2.2). Not only does the atmosphere attenuate the CMB signal, it also emits **parasitic photons** that obfuscate the cosmic image (Section 3.3). CMB photons that survive the atmosphere then propagate through the telescope’s optical system (Section 2.3), which translates the sky image—whose *angular* dimension is defined by the telescope’s **field of view**—onto a **focal plane**—whose *linear* dimension is defined by the telescope’s **plate scale** or **magnification** (Section 2.10). The telescope component that forms the focal-plane image is called the **optics tube** (Section 2.3.1), which a dewar that cryogenically cools and precisely aligns the lenses and detectors. Cooling the optics is necessary for the same reason as minimizing atmospheric emission: lenses absorb and emit, attenuating the CMB signal and injecting parasitic photons into the system. Modern lenses for mm-wave astronomy typically have

¹Not all CMB experiments have the characteristics described in this chapter, but many do.

large refractive indexes to enable large fields of view, but as a result, these optics require **anti-reflection coatings** to maximize the telescope’s throughput (Section 2.3.2).

The coupling between the telescope and focal plane is governed by the optical system’s **f-number**, which quantifies both the telescope’s magnification and the optimal detector-pixel density (Section 2.3.4). The focal plane is populated with **coupling optics**—which are typically either feed horns or antennas—that sense the telescope’s image and feed it to the detectors (Section 2.3.3). Modern coupling elements sense both polarizations and often sense multiple “colors,” or frequency bands, enabling broadband characterization of CMB polarization (Figure 2.1). The light collected by the focal plane optics is detected using **transition-edge sensors**, which are extraordinarily sensitive, highly-linear bolometers that convert an input power signal into a current signal (Section 2.4). This current signal is then converted to a voltage using a **superconducting quantum interference device (SQUID)** amplifier, and the SQUID output is further amplified and converted to a digital signal that is finally saved to disk (Section 2.4.2).

In order to suppress parasitic photons due to thermal emission, cryogenic optical components operate between 50 K and 0.5 K, and in order to suppress thermal noise, detectors operate at ~ 100 mK. As previously stated, the optics tube is a cryostat for the lenses and bolometers, and in order to minimize heat on cryogenic components, it must *reject* IR radiation while *admitting* mm-wave radiation (Section 2.5). To accomplish this, CMB experiments often use optical materials whose dielectric loss rises sharply above ~ 300 GHz, and they employ arrays of absorptive, reflective, and scattering **IR blockers** to reject THz radiation.

Modern CMB instruments are **background limited**, meaning that their sensitivity is limited by photon noise rather than “internal” detector noise. In this background-limited paradigm, improving the sensitivity of the detectors themselves has a diminishing impact on the experiment’s overall sensitivity. Instead, the most effective way to improve the signal-to-noise is to decrease parasitic optical power and increase the number of detectors,² whose data can be coadded to average down background interference. This paradigm is in contrast to earlier CMB instrumentation efforts, which focused on making a small number of detectors as sensitive as possible. Of course, maximizing per-detector sensitivity is still critical to a successful experiment, but the CMB community’s transition to large detector arrays drives many of its most active research areas. Instrument sensitivity and detector-array optimization are covered in Chapters 3 and 4, respectively.

In addition to maximizing sensitivity, CMB instruments must minimize **systematic effects**, which can create non-celestial signals and biases in the detector data. There are two general methods to mitigate systematics: via calibration hardware and via analysis techniques. Hardware-based systematic suppression is, in theory, most effective, as instrument-based techniques can be combined with analysis techniques; however, complex hardware

²This assertion assumes a fixed project lifetime, as observing the CMB for longer is another way to improve sensitivity. However, instrumentation technologies advance rapidly, and it is often cheaper to build a new telescope with more detectors than to continue operating an antiquated one.

systems can create their *own* systematic effects and therefore carry the danger of doing more harm than good if not implemented carefully. Two particularly prominent systematic effects for ground-based CMB telescopes are **atmospheric 1/f noise** (Section 6.1), which degrades detector sensitivity on long time scales, and instrumental polarization (Section 6.2), which generates spurious signals that can degrade data quality. Many methods have been used to mitigate instrument-induced errors in CMB polarization measurements, but Chapters 6-9 focus specifically on **continuously rotating half-wave plates**, a hardware technique that both Simons Array and Simons Observatory use to control a variety of systematic effects.

2.1.1 Simons Array

The Simons Array (SA) is a CMB observatory located on Cerro Toco in the Atacama Desert of Chile at an elevation of 5,200 m. SA consists of three telescopes in total, each housing a **dichroic**³ POLARBEAR-2 (PB-2) receiver cryostat: **PB-2a** and **PB-2b** observe at 90 and 150 GHz, while **PB-2c** observes at 220 and 270 GHz [196]. Each telescope has a 2.5 m primary mirror in an off-axis Gregorian configuration [217] and generates 5.2, 3.5, 2.7, and 2.2 arcmin **full width at half maximum (FWHM)** beams⁴ at 90, 150, 220, and 270 GHz, respectively. Each receiver cryostat has a 0.5 m vacuum window aperture, three 4 K alumina reimaging lenses, a 4 K Lyot stop, and a 365 mm diameter, 0.3 K focal plane. Each focal plane contains seven **detector wafers**⁵ totaling 1,897 detector pixels, and each pixel contains four transition-edge sensor (TES) bolometers, which together sense two colors and two polarizations. PB-2a achieved first light in January 2019 [100], PB-2b [85] has been deployed to Chile, and PB-2c is under development. The PB-2 experiments are the successor of the POLARBEAR experiment [102], which deployed in 2012 and has produced multiple landmark B-mode measurements [213, 212, 211].

2.1.2 Simons Observatory

The Simons Observatory (SO) is an upcoming CMB observatory also located on Cerro Toco in the Atacama Desert of Chile. SO consists of four telescopes in total: one **large aperture telescope (LAT)** that can house up to 13 individual optics tubes (OTs) and three **small aperture telescopes (SATs)** [59]. The LAT has a ≈ 6 m primary aperture and an **angular resolution**⁶ of ≈ 1 arcmin at 150 GHz [154], while the SAT has a ≈ 0.5 m aperture with an angular resolution of ≈ 20 arcmin at 150 GHz [4]. Given this construction, the LAT will study structure formation, galaxy clusters, and gravitational lensing, while the SAT will search for primordial B-modes [215]. All SO OTs are dichroic and observe at one of three

³In the context of CMB instrumentation, “dichroic” refers to the simultaneous observation of two distinct colors or frequency bands.

⁴In this context, “beam” refers to the telescope’s far-field angular response function.

⁵A detector wafer is a 150 mm-diameter, 0.7 mm-thick silicon “wafer” on which the detectors are fabricated. The detector array is then composed of detector wafers.

⁶As suggested in Section 2.1.1, a telescope’s angular resolution is determined by its FWHM beam size.

frequency varieties: **low-frequency (LF)** at 30 and 40 GHz, **mid-frequency (MF)** at 90 and 150 GHz, and **ultra-high-frequency (UHF)** at 220 and 270 GHz. Each LAT OT is designed to house three detector wafers, while each SAT OT is designed to house seven. The LAT and SAT detector wafer designs are identical, and each LF wafer has 37 dichroic, dual-polarization detector pixels while each MF and UHF wafer has 432. The selection of OT varieties within the LATR is discussed in Section 5.4, but assuming two LF OTs, eight MF OTs, and three UHF OTs, the LAT can hold up to $\approx 58,000$ detectors, while the SATs, assuming one LF OT, two MF OTs, and one UHF OT,⁷ can hold up to $\approx 37,000$. Therefore, SO represents a substantial leap in detector count above those of Stage-3 experiments like SA, leading to unprecedented sensitivity but also unprecedented technological challenges.

2.2 Observation site

Observing the CMB from the ground is challenging. Complex atmospheric structures scatter, absorb, and emit mm-wave radiation, resulting in a temporally- and spatially-varying sky intensity that obfuscates the CMB signal [50]. To limit this obfuscation, CMB observatories operate at high elevation and in arid climates, where the air is thin and humidity is low. In this section, we discuss the atmosphere’s mm-wave properties and assess the observation site for SO and SA.

2.2.1 Sky intensity

The most important parameter when evaluating a CMB observation site is the sky’s mm-wave brightness. When discussing the sky’s intensity, radio astronomers often refer to its “temperature”

$$T_{\text{ATM}} = \varepsilon_{\text{ATM}} T_{\text{phys}} , \quad (2.1)$$

where T_{phys} is the sky’s physical temperature, ε_{ATM} is the atmosphere’s emissivity, and where T_{ATM} is an **effective temperature** or **Rayleigh-Jeans temperature**. Equation 2.1 is valid in the Rayleigh-Jeans (RJ) limit where $h\nu \ll k_{\text{B}}T_{\text{phys}}$, and given a diffraction-limited entendue⁸ of $A\Omega = \lambda^2$, the sky’s optical power P_{sky} and effective temperature are proportionally related as

$$P_{\text{sky}} = k_{\text{B}}T_{\text{ATM}}\Delta\nu , \quad (2.2)$$

where k_{B} is the Boltzmann constant and $\Delta\nu$ is the detection bandwidth. Atmospheric emission is a major source of parasitic mm-wave photons for ground experiments and is therefore a major source of photon noise. In addition, atmospheric transmissivity is $\eta_{\text{ATM}} \approx 1 - \varepsilon_{\text{ATM}}$, and therefore minimizing the atmosphere’s emission also maximizes its transparency to CMB photons.

⁷There are four SATs being developed, but only three can observe simultaneously, necessitating a staged deployment program.

⁸CMB optical systems are typically diffraction limited, as discussed further in Section 2.3.1.




Figure 2.1: The sky temperature at the Atacama observation site assuming 1.4 mm PWV and 50 deg elevation, plotted over SO’s detector bands. The bands are placed within **atmospheric windows** between the oxygen absorption/emission lines at ~ 60 , 120, 180, and 330 GHz. The precipitous increase in atmospheric brightness temperature at high frequencies is due to the increasing emissivity of water vapor.

There are two primary constituents in the atmosphere that absorb/emit mm-wave radiation: water and oxygen. The impact of water follows a power law below ~ 1 THz and becomes increasingly prominent at higher frequencies [223], while oxygen has absorption lines at ~ 60 , 120, 180, and 330 GHz that are due to spin-rotational quanta of the O_2 molecule [221, 223]. As a result, water has a broadband impact on atmospheric opacity, while oxygen has impacts which are limited to relatively narrow frequency ranges. In addition, the observed sky temperature depends on the line integral of its optical depth along the telescope’s line of sight, which to leading order scales as

$$T_{\text{ATM}} \propto \frac{1}{\sin \theta_{\text{el}}}, \quad (2.3)$$

where θ_{el} is the telescope’s **boresight elevation** above the horizon. Given this relation, observation strategies must consider atmospheric power vs. elevation when optimizing how to scan the sky (see Section 5.4.7).

In order to quantify the amount of water in the atmosphere, we monitor the **precipitable water vapor (PWV)**, which is the depth of liquid water equal to the water vapor in an imaginary column stretching from the ground to the upper atmosphere.⁹ Microwave emission due to water—and therefore the impact of PWV—increases with increasing frequency,

⁹In other words, if you placed a bucket on the ground and condensed all of the water vapor directly above it from $z = 0$ to $z \rightarrow \infty$, the PWV would be the depth of water in the bucket.

and at higher frequencies where water is more important than oxygen, the sky intensity approximately scales as [223]

$$T_{\text{ATM}} \propto \text{PWV}. \quad (2.4)$$

While Equations 2.3 and 2.4 are useful guides, they are in practice replaced by molecular simulations of the atmosphere’s absorption and scattering profile, one of which is shown in Figure 2.2.

2.2.2 Chile site

In order to minimize both CMB signal attenuation and atmospheric mm-wave emission, it is advantageous to observe through as dry and as little air as possible. While the most effective approach is to measure the CMB from beyond the Earth’s atmosphere [21, 150, 96, 208], launching and operating a satellite costs \sim billions of dollars and is risky, as the instrument cannot be serviced after launch.¹⁰ Another effective strategy is to launch a high-altitude balloon, which can float at $\sim 30,000$ m elevation above $> 99\%$ of the atmosphere [160, 99, 206, 65]. However, balloons have limited flight durations, and while they are much cheaper than satellites, balloons also cannot be serviced after launch and are therefore also quite risky.¹¹ Given these practical concerns, CMB observations from the ground are competitive if performed at the right location.

Both SA and SO observe on **Cerro Toco** in the Atacama Desert of Chile, and a photo of the site is shown in Figure 2.2a. Located at 5,200 m elevation in an exceptionally arid climate, Cerro Toco offers some of the best mm-wave and infrared observing conditions on Earth. Additionally, with its latitude of $\approx -23^\circ$, the site can view $> 70\%$ of the sky, offering opportunities both to reduce cosmic variance (see Equation 1.34) and to **cross correlate** with overlapping surveys. A map of SO’s observation areas is shown in Figure 2.2c, along with survey coverages of a few independent experiments. Also, because the sky rises, sets, and rotates in Chile, CMB patches can be scanned¹² in different directions—a technique called **cross linking**—which is a powerful tool to mitigate scan-dependent systematic errors.¹³ Figure 2.2b also shows a probability distribution of PWV at the Chile site during POLARBEAR’s second observation season [212], and the median value is ≈ 1 mm.¹⁴

¹⁰Unlike orbital observatories, such as the Hubble Space Telescope, which can be repaired, CMB satellites operate at the second Lagrange point (L2), which is ~ 1 million miles from Earth.

¹¹Even the balloon launch itself is risky, requiring both perfect weather and sophisticated flight hardware to avoid damaging the payload at take off.

¹²A typical mode of CMB observation is the **constant-elevation scan**, where the telescope modulates in azimuth while the sky patch rises/sets through the telescope’s field of view.

¹³Chile’s large sky coverage and sky rotation are in contrast to the South Pole, another widely-used CMB observation site, which has access to a smaller fraction of the sky and offers no sky rotation.

¹⁴As a comparison, the annual median PWV in Oakland, CA is ≈ 20 mm [176].

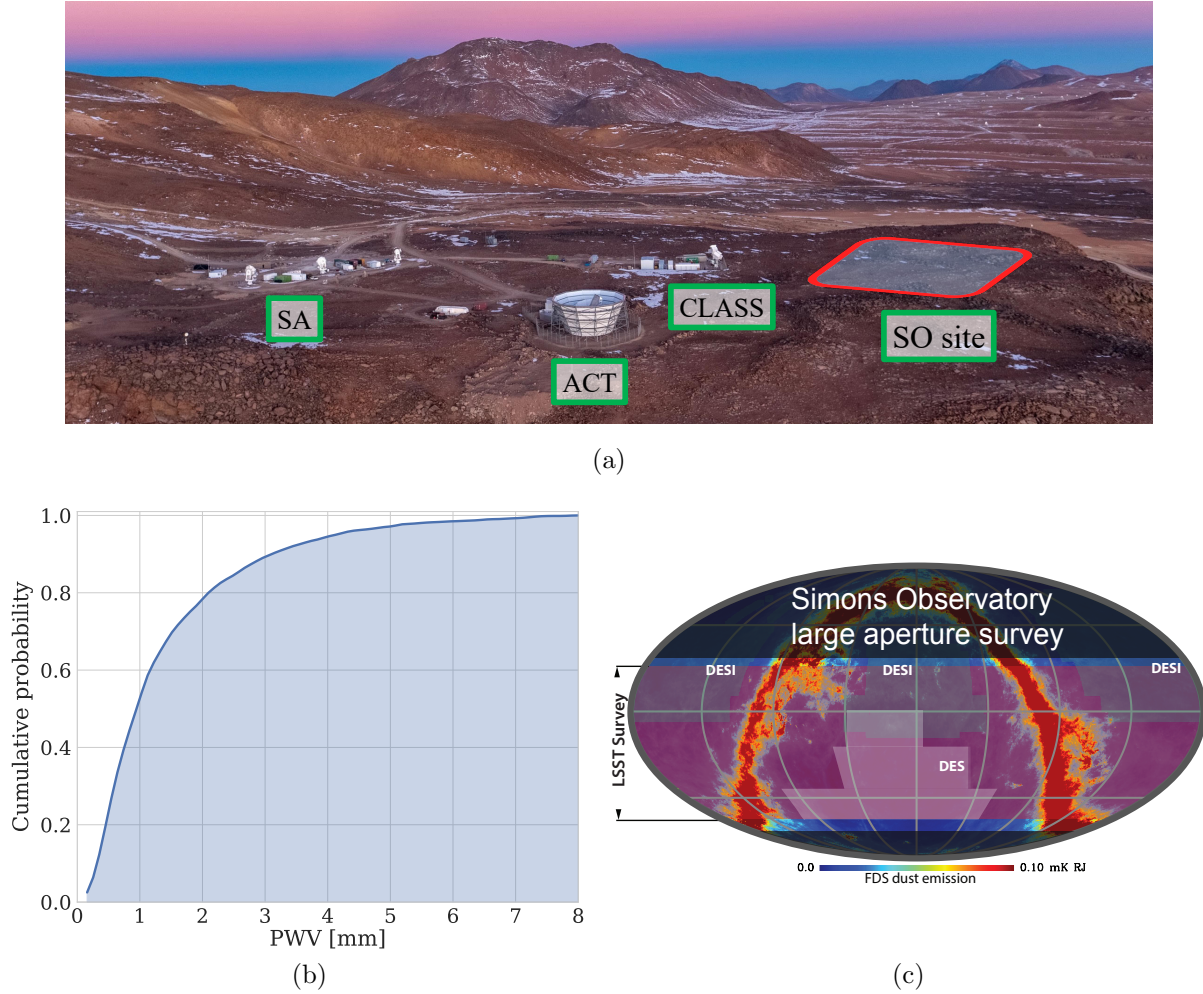


Figure 2.2: Figure 2.2a shows an image of the observation site for both SA and SO. Photo courtesy of Debra Kellner. Figure 2.2b shows the cumulative probability distribution for PWV at the Cerro Toco site [212]. The median value is ≈ 1 mm, and the value is < 2 mm nearly 80% of the time. Figure 2.2c shows the survey area in galactic coordinates of SO’s LAT [215]. Enabled by its moderate latitude, the Atacama site offers access to 70% of the sky, enabling overlap with other experiments, including the Dark Energy Spectroscopic Instrument (DESI) and the Dark Energy Survey (DES).

2.3 Optics

A CMB telescope’s fundamental purpose is to image the sky with low distortion and high throughput. Because 100 GHz optics are not widely available in the commercial sector,¹⁵ CMB telescopes are designed and built in-house using materials with favorable mm-wave

¹⁵Microwave optics *are* receiving increasing industry attention with the advent of 5G telecommunications, but these technologies are in the early stages of proprietary development and are often incompatible with the diameters and temperatures of CMB telescopes.

properties. In addition, in order to minimize parasitic photons and enable sub-Kelvin detectors (see Section 2.4), CMB optical systems are cooled to cryogenic temperatures,¹⁶ both improving their transparency and limiting their thermal emission. In this section, we overview SA’s and SO’s optical systems, focusing on aspects relevant to this dissertation’s research. We first discuss SA and then highlight SO’s differences.

2.3.1 Telescope and optics tube

The SA telescope is shown in Figure 2.3a and in the left panel of Figure 2.3b. It is composed of two monolithic mirrors that together form an off-axis Gregorian configuration [217]. The primary mirror is a parabolic reflector that focuses incoming parallel rays¹⁷ onto the **prime focus**. The sky’s image at the prime focus is diffraction limited over a small field of view and is therefore **reimaged** by an ellipsoidal secondary mirror onto the **Gregorian focus**. While improved with respect to that of the primary mirror, this focus is not **telecentric**¹⁸ and has only a moderate **diffraction-limited field of view (FOV)**. Therefore, an **optics tube** again reimages onto the detector array and employs a **Lyot stop** to define the primary-mirror illumination. The primary and secondary mirrors form a compact assembly shown in Figure 2.3, which is designed to be nimble and provide scan-strategy flexibility.¹⁹ In addition, the off-axis Gregorian design satisfies the **Mizuguchi-Dragone condition** [138, 47], which limits cross polarization over a wide FOV.

The SA optics tube (OT), which is the tubular section of the receiver cryostat shown in Figure 2.3b, contains three reimagining lenses: the field lens, aperture lens, and collimator lens. The field lens adjusts the speed²⁰ of the optics at the Gregorian focus, while the aperture and collimator lenses form an image of the primary mirror at the Lyot stop²¹ before forming a high-fidelity, telecentric, large-FOV **sky image** at the focal plane. The Lyot stop is cooled to 4 K and limits the primary mirror’s optically-active area. This function is critical to minimizing detected parasitic power, as the Lyot’s cryogenic aperture²² “hides” ambient-temperature telescope structures from the detectors’ view. All three lenses are made of alumina, which is polycrystalline aluminum oxide Al_2O_3 with a large refractive index at ~ 100 GHz ($n_{\text{alumina}} \approx 3.1$), and are also cooled to ≈ 4 K. To maximize the telescope’s **optical throughput**, each lens surface has a dual-layer **anti-reflection (AR) coating** that achieves $\approx 1\%$ reflectivity per optic. AR coatings are a central topic of this dissertation and are overviewed further in Section 2.3.2.

¹⁶Except for large mirrors, which typically exist in ambient conditions.

¹⁷In a physical-optics sense, parallel rays correspond to plane waves emitted from a source in the telescope’s far field, such as the CMB.

¹⁸A telecentric image is one that is **orthographic** or has constant magnification and image quality across a planar surface. Telecentricity is important for CMB telescopes, which use planar detector arrays.

¹⁹Tunable scan parameters include azimuthal velocity, turnaround acceleration, and elevation range.

²⁰“Fast” optics have shorter focal lengths, while “slow” optics have longer ones.

²¹The image at the primary mirror and Lyot stop are the Fourier transform of the sky and are sometimes called **conjugates** of the “real” image.

²²A 4 K surface emits $\approx 0.7\%$ the power of an equivalent 300 K surface at 150 GHz.




Figure 2.3: The PB-2b telescope and receiver cryostat as a representative example of the SA optics. Figure 2.3a is a photograph of an SA-style telescope. Baffling and shields limit stray light, and the secondary mirror and receiver cryostat are hidden within their respective enclosures. Figure 2.3b shows computer-aided design (CAD) cross-sections of the full SA system (left) and the receiver alone (right). The optics and detectors are discussed in this chapter, while the cryogenic half-wave plate (CHWP) is discussed in Chapters 6-9.

The SO telescopes are shown in Figure 2.4 and come in two varieties: a large-aperture telescope (LAT) and a small-aperture telescope (SAT) [59]. The LAT is shown in Figure 2.4a and is composed of two polythitic mirrors²³ which together form a **Crossed Dragone (CD)** configuration [154, 60]. While less compact than the off-axis Gregorian configuration that SA uses, the CD arrangement offers many appealing optical properties, including outstanding image quality over a large FOV.²⁴ The **LAT receiver (LATR)**, shown in Figure 2.4b, reimages onto an array of (up to) 13 discrete focal planes, each with its own cryogenic OT [240]. In order to maximize the LAT’s FOV and image quality, each off-center optics tube has an alumina wedge to correct wavefront errors away from the telescope’s chief ray [44]. Each OT has three single-crystal silicon ($n_{\text{Si}} \approx 3.4$) reimaging lenses at $1 \sim 4$ K that reimage through a Lyot stop onto three detector wafers.

The SO SAT is shown in Figure 2.4c and is designed to measure large-angular-scale CMB fluctuations, a goal that enables a smaller system free of ambient-temperature mirrors and reimaging optics.²⁵ Its **aperture stop**²⁶ is cooled to ≈ 2 K near the vacuum window, and the optics tube contains three silicon lenses at $1 \sim 2$ K that image the sky onto a ≈ 400 mm-diameter focal plane of seven detector wafers [4]. The SAT is much cheaper to build than the LAT, and therefore three SATs are planned to operate in parallel on separate platforms.

Each SA receiver and SO SAT employs a sapphire **continuously-rotating half-wave plate (HWP)** polarization modulator. The SA HWPs are located near the receiver’s vacuum window [82, 80, 79], which is near the telescope’s Gregorian focus, while the SO SAT HWPs are located directly in front of the aperture stop [170]. Polarization modulation is a powerful tool to reject the unpolarized atmosphere, suppressing atmospheric low-frequency noise and improving sensitivity on large angular scales. HWPs are a central research topic of this dissertation and are detailed in Chapters 6-9.

2.3.2 Anti-reflection coatings


SA and SO use alumina and silicon lenses, respectively, each of which has a large refractive index ($n_{\text{alumina}} \approx 3.1$, $n_{\text{Si}} \approx 3.4$) compared to that of vacuum $n_{\text{vacuum}} = 1$. Therefore, the vacuum-to-alumina/silicon interface has a reflectivity of $R = (n_{\text{lens}} - 1)^2 / (n_{\text{lens}} + 1)^2 \sim 30\%$ at normal incidence and even larger at oblique angles. Because each lens has two surfaces and there are three lenses per OT, the total **transmissivity** of barefaced refractors would be only $T = (1 - R)^6 \sim 15\%$. In addition, such large reflectivities encourage multiple reflections that can generate **ghosting**, or the formation of multiple images within the cryostat. Therefore, **anti-reflection (AR) coatings** play a critical role in the optical performance of the receiver.

²³Polythitic mirrors usually comprise discrete panels assembled with high precision.


²⁴For this reason, some CMB telescopes use a CD optical system without any reimaging optics, such as the Atacama B-mode Search (ABS) [51] and the LiteBIRD satellite [173].

²⁵The SAT instead simply contains “imaging” optics.


²⁶In the absence of reimaging optics, a cold aperture stop both defines the telescope’s angular resolution and limits its ambient optical load.



(a)



(b)



(c)

Figure 2.4: CAD cross-sections of the SO telescope systems. Figure 2.4a shows the LAT, whose primary and secondary mirrors focus onto the LATR. Figure 2.4b isolates the LATR, showing three of its thirteen identical optics tubes. Figure 2.4c shows the SAT, which features a cryogenic half-wave plate (CHWP) directly in front of its aperture stop.




Figure 2.5: A sampling of AR coating technologies used within SA and SO. Figure 2.5a shows the PB-2b field lens coated with a two-layer epoxy AR coating (see Section 10.2) and mounted in a cryostat before thermal testing. Figure 2.5b shows an example three-layer metamaterial AR coating used for SO’s MF silicon lenses, and the photo is courtesy of Joey Golec.

AR coatings are detailed in Chapter 10 and Section 7.2.2, but we overview SA’s and SO’s AR technologies briefly here.

SA’s alumina optics are AR coated using a few technologies that have matured during the course of the project. All coatings are dichroic and comprise two distinct layers, each with an optimized index and thickness. In PB-2a, the lenses’ curved sides are coated with two layers of Stycast **epoxy** [163], which is molded onto each surface and milled to its target thickness. A photo of an epoxy-coated field lens is shown in Figure 2.5a. The flat sides of the lenses are coated with a bottom layer of plasma-sprayed Mullite and a top layer of an expanded polyimide foam [93]. PB-2b and PB-2c aim to²⁷ **plasma spray** alumina powders loaded with various compositions of fillers that adjust each layer’s index [97]. The layers are applied directly onto the optic’s surface using an optimized deposition process, and the plasma-spray parameters are finely tuned to control each layer’s density and uniformity. All three of these technologies have advantages and disadvantages, but each attains $\sim 1\%$ reflectivity per surface, amounting to a all-lenses throughput of $\gtrsim 90\%$.

SO uses **metamaterial** AR coatings for their alumina and silicon optics [42]. This technique involves cutting **sub-wavelength structures** into the optic’s surface to tune each layer’s *effective* refractive index via its silicon/alumina-to-vacuum ratio. An example of an MF silicon metamaterial coating is shown in Figure 2.5b. This sub-wavelength-structure method has several advantages over the more “traditional” film-based approach. First, there are no concerns about cryogenic delamination because there is no **coefficient of thermal expansion (CTE) mismatch** between the AR coating and the substrate. Second, each layer’s refractive index is tuned by simply adjusting each layer’s sub-wavelength geometry.²⁸

²⁷Multiple AR techniques are still under evaluation for SA. See Section 10.4 for more details.

²⁸The flexibility of geometric tuning is in sharp contrast to the limited selection of dielectric films for mm-wave applications (see Section 7.2.2).

Third, it is relatively²⁹ straightforward to add more AR layers, which further reduces reflectivity [36]. Given these advantages, SO plans to cut three-layer AR coatings into its silicon optics and two-layer coatings into its alumina optics, achieving $< 1\%$ reflectivity per surface. While metamaterial coatings are well demonstrated on silicon optics, SO will be the first CMB experiment to field sub-wavelength structures on alumina.

2.3.3 Focal plane optics

As discussed in Section 2.3.1, a CMB telescope images the sky onto a focal plane, and the focal plane optics must feed the telescope’s image to detectors. The focal plane is composed of discrete **detector pixels**, each of which is sensitive to two polarizations and two colors. PB-2a, PB-2b, two SATs, and some LAT OTs observe at 90 and 150 GHz (mid-frequency, or “MF”), PB-2c, one SAT, and some other LAT OTs observe at 220 and 270 GHz (ultra-high-frequency, or “UHF”), and one SAT and the remaining LAT OTs observe at 30 and 40 GHz (low-frequency, or “LF”). Given these **dichroic** OTs and detector pixels, the focal plane optics must achieve high throughput across broad, well-defined bandwidths.

Each SA detector array consists of seven **detector wafers**, each of which contains 271 detector pixels [224], as shown in Figure 2.6a. The sensing element for SA’s detector pixels is a planar **sinuous antenna** (Figure 2.6a) coupled to a hemispherical silicon **lenslet** (Figure 2.6b) [197]. The reimaging lenses focus the sky onto the tip of the lenslet, which in turn “focuses” light onto the antenna feed.³⁰ In order to optimize focal-plane coupling, the lenslets are offset from the sinuous antenna by a silicon spacing wafer and are oversized to ensure efficient collimation. The sinuous is a **fractal antenna**, whose logarithmically repeating pattern makes it sensitive over a broad bandwidth, and its high- and low-frequency limits are determined by its smallest feature and overall size, respectively [197]. After the antenna collects the incident radiation, transmission lines carry the light through **on-chip bandpass filters**, which separate 90 and 150 GHz light while rejecting out-of-band radiation [224]. The band-admitted light is then dissipated onto a **thermistor** whose output is amplified, digitized, and stored.

The SO LF detector arrays use a similar lenslet-coupled-antenna architecture as SA, but its MF and UHF arrays instead employ **feedhorns** coupled to polarization-discriminating **orthomode transducers** [183, 182]. A feedhorn is effectively an impedance-matched waveguide, designed to transform free-space spherical waves into wave-guide modes, which can then be routed, manipulated, and sensed. Feedhorns come in many different varieties, the most common of which uses corrugated walls to enforce a symmetric angular response [41, 77]. While **corrugated feedhorns** are highly effective, a substantial fraction of their aperture

²⁹There are practical challenges to cutting layers with either very thin grooves or very thin pillars, and so the process of adding metamaterial layers is not trivial. Nonetheless, concerns about complexity and durability for a three-layer sub-wavelength AR coating are substantially less than those of a three-film coating [142].

³⁰The use of the term “focus” here is strictly inappropriate to describe the lenslet-antenna system, whose separation is of order one wavelength and whose behavior is only describable using near-field electrodynamics. Nonetheless, ray-trace analyses describe the lenslet’s geometric dependencies reasonably well [197].

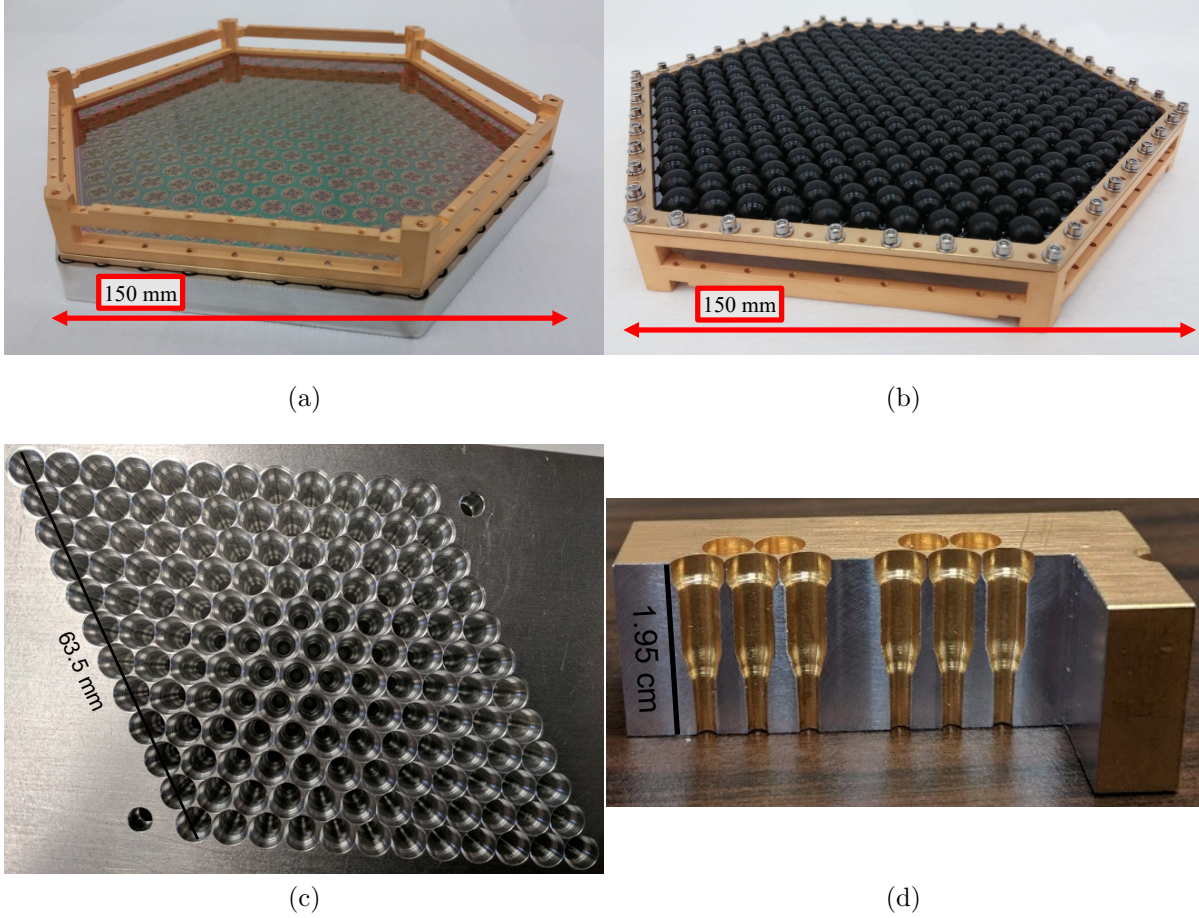


Figure 2.6: The SA and SO focal plane optics. Figure 2.6a shows a PB-2a wafer, which has 271 detector pixels comprising sinuous antennas and 90/150 GHz TES bolometers, and Figure 2.6b shows the lenslet array that focuses light onto the antennas. The lenslets are coated with an epoxy AR coating, which is that used by PB-2a and PB-2b, but PB-2c and SO are developing alternative lenslet coating technologies. Figures 2.6c and 2.6d show a prototype feedhorn array for 90/150 GHz pixels on SO. The spline profile offers a symmetric, efficient response while allowing for dense pixel packing. The sinuous and lenslet photos are courtesy of Aritoki Suzuki, and the feedhorn photos are courtesy of Sara Simon [182].

is optically inactive, hence “wasting” focal plane space. Therefore, SO uses smooth-walled feedhorns with spline-profiled shapes that maximize coupling to the telescope optics, minimize cross polarization, and symmetrize optical response [183]. Without corrugations, the array of feed horns can be very dense, with only hundreds of microns separating adjacent detector pixels. This configuration maximizes the density of detectors on the focal plane, which improves overall experiment sensitivity (see Section 4.4.1) [81].




Figure 2.7: A schematic of the detector pixel optics. The green detector pixel with diameter D_{pix} illuminates (in the reverse-time sense) the aperture stop of diameter D_{ap} and distance z_{ap} from the focal plane. The beam size $w(z)$ is inversely proportional to its beam waist w_0 , and the approximation $w(z) \approx z\lambda/\pi w_0$ applies in the **Rayleigh limit** where $z \gg w_0$.

2.3.4 Telescope-to-focal plane coupling

Given a high-level overview of SA's and SO's pixel architectures, we now formalize the coupling between the focal plane and the telescope, which is central to many of the sensitivity calculations in Chapters 3 and 5. Assuming a **diffraction-limited optical system**, the optics tube's **magnification** is quantified by its **f-number** (also called the **focal ratio**) at the focal plane

$$F = \frac{f}{D_{\text{ap}}} , \quad (2.5)$$

where D_{ap} is the diameter of the aperture stop and f is the focal length of the focal plane image.³¹ Provided a fixed aperture size, a larger (smaller) f-number corresponds to a smaller (larger) focal plane or equivalently a smaller (larger) magnification.

As long as the detector pixels are large compared to the wavelength, they will induce minimal diffractive ringing, and therefore their far-field **beam**—or their angular response—is well approximated by a Gaussian function (Figure 2.7)

$$B_{\text{pix}}(\theta, \lambda) = \exp \left(\frac{-\theta^2}{2\theta_{\text{pix}}^2(\lambda)} \right) , \quad (2.6)$$

³¹In the SA telescopes and SO LAT, which employ reimaging optics, radiation propagating through the Lyot stop is not necessarily an orthogonal set of plane waves, and therefore the situation is more complicated than Figure 2.7 suggests. Nonetheless, focal plane designers typically quote a Lyot system's **effective f-number**, which, to a good approximation, can be treated just like the simple f-number in Equation 2.5.

where θ is zenith angle and where

$$\theta_{\text{pix}}(\lambda) = \frac{\lambda}{\pi w_0} . \quad (2.7)$$

Here, the pixel's beam width θ_{pix} is inversely proportional to its **beam waist** w_0 , which can be related to the pixel's diameter D_{pix} via its **waist factor** w_f

$$w_0 = \frac{D_{\text{pix}}}{w_f} . \quad (2.8)$$

The most important result of this diffraction-limited formalism is that a larger (smaller) pixel has a tighter (wider) beam, impacting its coupling to the telescope.

There are two useful ways to conceptualize the relationship between the telescope and focal plane optics. The first way is via the **forward-time sense**, which is to imagine photons propagating from the sky to the detectors. In this paradigm, the telescope's image becomes smaller (larger) with smaller (larger) f-number/magnification, which in turn favors smaller (larger) pixels. The second and equally valid way is via the **reverse-time sense**, which is to imagine photons propagating from the detectors to the sky.³² In this paradigm, a smaller (larger) f-number corresponds to a larger (smaller) aperture, which in turn favors a wider (narrower) pixel beam and therefore a smaller (larger) pixel. While these two paradigms are equivalent, it is often most convenient to think in the reverse-time sense.³³

From the perspective of the detector pixel, the aperture stop's half-angle is

$$\theta_{\text{st}} = \arctan\left(\frac{1}{2F}\right) \approx \frac{1}{2F} , \quad (2.9)$$


where the approximation applies in the **paraxial limit**, or the small-angle regime. The coupling between the telescope and the detector pixel's beam is often called the **beam coupling efficiency**

$$\eta_{\text{apert}} = \frac{\int_0^{\theta_{\text{st}}} B_{\text{pix}}(\theta, \lambda) d\theta}{\int_0^{\pi/2} B_{\text{pix}}(\theta, \lambda) d\theta} = 1 - \exp\left[-\frac{\theta_{\text{st}}^2}{2\theta_{\text{pix}}^2}\right] \approx 1 - \exp\left[-\frac{1}{2}\left(\frac{\pi D_{\text{pix}}}{w_f F \lambda}\right)^2\right] , \quad (2.10)$$

where the second equivalence again relies on the paraxial approximation. Note that this beam-spill treatment only applies to a diffraction-limited optical system, which has a **single-moded etendue** of $\Lambda\Omega = \lambda^2$. In this diffraction-limited, single-moded paradigm, changing a pixel's diameter does not change the amount of light it can collect but only the size of its far-field illumination.

³²The equivalence of the forward-time and reverse-time senses is due to the time-reversal symmetry of electromagnetism.

³³The reverse-time-sense paradigm is standard in antenna engineering, as a receiver's response is most conveniently studied via its radiation pattern.



Pixel	Coupling	D_{pix} [mm]	$D_{\text{pix}}/F\lambda$	η_{apert}
SA 90/150 GHz	Lenslet + sinuous	6.8	90 GHz: 1.1 150 GHz: 1.8	0.54 0.85
SA 220/270 GHz	Lenslet + sinuous	6.8	220 GHz: 2.6 270 GHz: 3.2	0.98 0.99
SO LF	Lenslet + sinuous	18.3	30 GHz LAT: 0.8 30 GHz SAT: 1.0 40 GHz LAT: 1.1 40 GHz SAT: 1.5	0.25 0.37 0.48 0.65
SO MF	Feedhorn	5.3	90 GHz LAT: 0.8 90 GHz SAT: 1.0 150 GHz LAT: 1.2 150 GHz SAT: 1.5	0.25 0.37 0.51 0.71
SO UHF	Feedhorn	5.3	220 GHz LAT: 1.8 220 GHz SAT: 2.3 270 GHz LAT: 2.2 270 GHz SAT: 2.9	0.79 0.92 0.90 0.97

Figure 2.8: Beam coupling efficiencies for SA's and SO's detector pixels. The figure presents the pixel sizes in units of $F\lambda$ along with the corresponding η_{apert} according to Equation 2.10. The shaded region corresponds to $2.8 \leq w_f \leq 3.2$. The table presents the detector's labels along with their D_{pix} and η_{apert} values. Many of SA's and SO's pixel sizes are clustered near $1F\lambda$, which is a natural optimum detailed in Chapter 4.

A sampling of pixel diameters in units of $F\lambda$ from SA and SO along with their respective beam coupling efficiencies η_{apert} is shown in Figure 2.8. As evidenced by the table’s values, many CMB detector pixel diameters are near $1F\lambda$ and therefore substantially spill over the aperture stop. While undersized pixels degrade *per-detector* sensitivity, they allow more detectors to be packed onto the focal plane,³⁴ which improves *overall* sensitivity. When η_{apert} is low, minimizing the stop’s temperature is critical to limiting parasitic power on the detectors.

2.4 Detectors

As discussed in Section 2.3, telescope optics couple to the detectors through an antenna or feedhorn, and the to-be-detected microwave frequencies are defined by on-chip bandpass filters. This filtered radiation is then transduced using a **transition-edge sensor bolometer (TES)**. TESes are exceptionally powerful mm-wave power detectors and have revolutionized CMB telescopes over the past decade or so. In this section, we discuss TES operation and how their signals are amplified and digitized [197, 5].

2.4.1 TES operation

Figure 2.9 shows a schematic of TES operation. A superconducting film is **voltage biased**³⁵ such that the time-averaged electrical power on the bolometer is

$$P_{\text{bias}} = \frac{V_{\text{bias}}^2}{R_{\text{bolo}}}, \quad (2.11)$$

where V_{bias} is the root-mean square (RMS) **bias voltage** and R_{bolo} is the bolometer’s operating resistance. The total power on the bolometer is the sum of the bias power P_{bias} and sky power P_{opt}

$$P_{\text{tot}} = P_{\text{bias}} + P_{\text{opt}}. \quad (2.12)$$

As shown in Figure 2.9a, the TES operates at its **superconducting transition temperature** T_c , and its **saturation power** P_{sat} defines the power it dissipates to a thermal bath with temperature T_b via a thermal link with conductivity G

$$P_{\text{sat}} = G(T_c - T_b). \quad (2.13)$$

³⁴This statement is true in the commonplace limit of a fixed focal plane area.

³⁵While it is also possible to current bias the bolometer, doing so negates some of the TES’s most powerful properties.

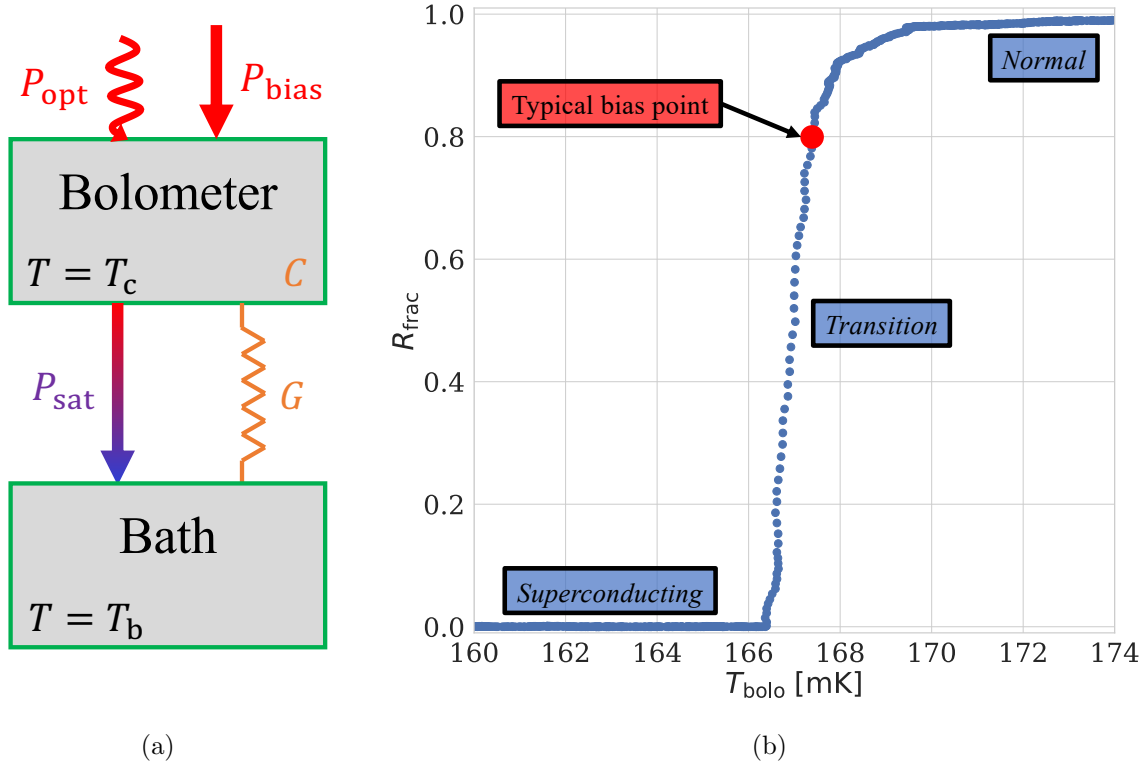


Figure 2.9: A description and demonstration of the TES bolometer. Figure 2.9a shows a schematic of bolometer operation, and Figure 2.9b shows the measured superconducting transition of a sample TES for SO. The bolometer is held at its superconducting transition temperature by an optical load P_{opt} , an electrical load P_{bias} , and a tuned thermal conduction G to a thermal bath with temperature T_b . During normal operation, $P_{\text{sat}} = P_{\text{opt}} + P_{\text{bias}}$. The bias point is chosen such that, in the presence of an optical power fluctuation dP_{opt} , dR/dP_{opt} is large. In the SO sample measurement, the bias point is chosen to be $R_{\text{frac}} = 0.8$, or 80% of the bolometer's normal resistance.

During normal operation, P_{bias} is tuned such that $P_{\text{tot}} = P_{\text{sat}}$.³⁶ If $P_{\text{tot}} > P_{\text{sat}}$, the bolometer's operating temperature $T_{\text{bolo}} > T_c$ and it saturates,³⁷ and if $P_{\text{tot}} < P_{\text{sat}}$, $T_{\text{bolo}} < T_c$ and the bolometer **latches**³⁸, rendering it inoperable.

To most easily understand the bolometer's behavior in the presence of changing sky signal, consider a single Fourier mode of the optical power with angular frequency ω

$$P_{\text{tot}} + \delta P_{\text{opt}} e^{i\omega t} + \frac{dP_{\text{bias}}}{dT_{\text{bolo}}} \delta T_{\text{bolo}} e^{i\omega t} = G(T_c - T_b) + (g + i\omega C) \delta T_{\text{bolo}} e^{i\omega t}. \quad (2.14)$$

³⁶Because P_{opt} depends on sky conditions such as weather and elevation, bolometers are often retuned during the course of an observation.

³⁷If the bolometer is *intentionally* saturated by setting $P_{\text{bias}} > P_{\text{tot}} - P_{\text{opt}}$, it is said to be **overbiased**.

³⁸In the latched state, the bolometer is fully superconducting and therefore cannot dissipate electrical power. When this happens, the recovery mode is to warm $T_b > T_c$, rebias the bolometer, and cool the bath back down.

Here, $g = \delta P / \delta T$ is the **dynamic thermal conductance** and C is the bolometer's heat capacity, as shown in Figure 2.9a. We can rewrite the electrical coefficient as

$$\frac{dP_{\text{bias}}}{dT_{\text{bolo}}} = -\frac{V_{\text{bias}}^2}{R_{\text{bolo}}^2} \frac{dR_{\text{bolo}}}{dT_{\text{bolo}}} = -\frac{P_{\text{bias}}\alpha}{T_c}, \quad (2.15)$$

where $\alpha = d \log(R_{\text{bolo}}) / d \log(T_{\text{bolo}})$ parameterizes the steepness of the TES's superconducting transition, and where we note that, on average, $T_{\text{bolo}} = T_c$. Isolating the time-varying parts of Equation 2.14 gives

$$\delta P_{\text{opt}} = \left[\frac{P_{\text{bias}}\alpha}{T_c} + g + i\omega C \right] \delta T_{\text{bolo}}. \quad (2.16)$$

This equation represents a transducer amplifier that converts fluctuations in optical power δP_{opt} to fluctuations in bolometer temperature δT_{bolo} with a **loop gain** of

$$\mathcal{L}(\omega) = -\frac{\delta P_{\text{bias}}}{\delta P_{\text{opt}}} = \frac{P_{\text{bias}}\alpha}{gT_c(1 + i\omega\tau_0)} = \frac{\mathcal{L}}{1 + i\omega\tau_0}, \quad (2.17)$$

where $\mathcal{L} = P_{\text{bias}}\alpha / (gT_c)$ is the **open loop gain**. The bolometer's **responsivity**, or its conversion from optical power to an electrical current signal, is

$$S_I = \frac{dI}{dP_{\text{opt}}} = \frac{-\tilde{S}_{\text{fact}}}{V_{\text{bias}}} \frac{\mathcal{L}}{\mathcal{L} + 1} \frac{1}{1 + i\omega\tau}, \quad (2.18)$$

where the time constant is $\tau = \tau_0 / (\mathcal{L} + 1) = (C/g) / (\mathcal{L} + 1)$ and where the responsivity factor depends on whether the bolometer is AC- or DC-biased

$$\tilde{S}_{\text{fact}} = \begin{cases} 1 & \text{if } V_{\text{bias}} \text{ is DC} \\ \sqrt{2} & \text{if } V_{\text{bias}} \text{ is AC RMS.} \end{cases} \quad (2.19)$$

In the limit of high loop gain, the responsivity becomes

$$S_I \approx -\frac{\tilde{S}_{\text{fact}}}{V_{\text{bias}}} \quad \text{if } \mathcal{L} \gg 1. \quad (2.20)$$

There are a few key characteristics of the TES. First, because it is voltage biased, $dP_{\text{bias}}/dT_{\text{bolo}}$ is negative, as shown in Equation 2.15, which stabilizes the bolometer via **electrothermal feedback**.³⁹ Second, this inherent stability enables a very large dR/dT , which enables a large detector responsivity. Third, because they are nearly superconducting, TESes operate with low impedance, allowing current signals to be amplified using SQUID amplifiers, which we discuss in the next session.

³⁹Stated explicitly, if optical power on the bolometer increases, so too does the bolometer's temperature and resistance. However, the bolometer's increased resistance *decreases* its electrical power as $1/R_{\text{bolo}}$, therefore keeping the *total* power on the bolometer constant.

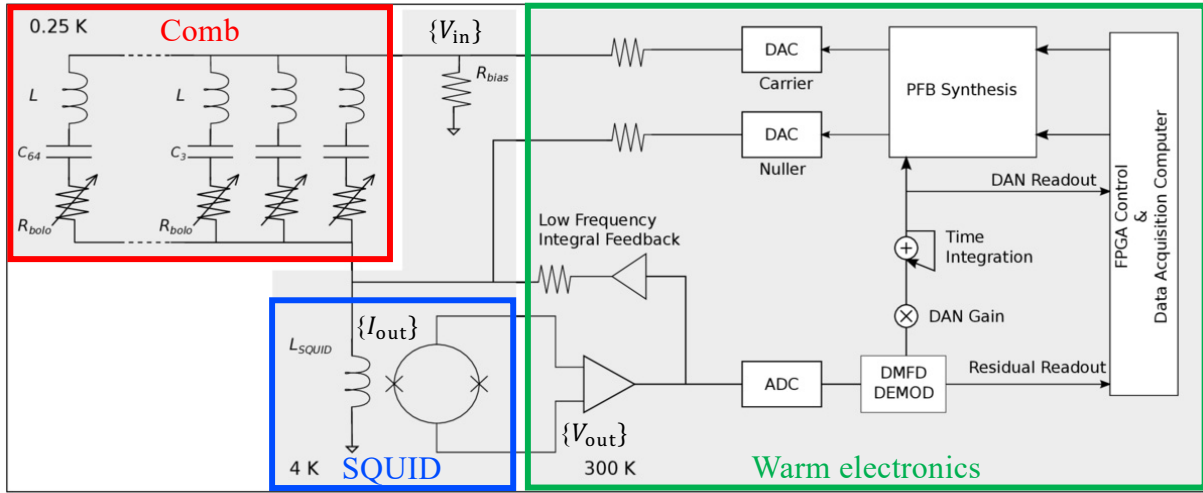


Figure 2.10: A schematic of the dfMUX readout for SA [14]. Thirty-eight bolometers with operating resistance R_{bolo} are in each series with an LC resonator, and each bolometer’s capacitor tunes its resonance frequency to form a “comb” of low-impedance peaks between 1-5 MHz. A series of matching input voltage waveforms $\{V_{\text{in}}\}$ bias the bolometers, and changes in R_{bolo} amplitude modulate the output current waveforms $\{I_{\text{out}}\}$. The output current waveforms are transimpedance amplified by a single SQUID into a series of output voltages $\{V_{\text{out}}\}$, which are then further amplified and digitized by ambient electronics.

2.4.2 TES readout

In order to digitize the detector signals, both SA and SO **multiplex** the readout of their TESes using **superconducting quantum interference devices (SQUIDs)** [193]. SQUIDs are transimpedance amplifiers with low input impedance and large gain, and they operate by sensing current changes across the bolometer (via Equation 2.18) using **Josephson junctions**. SQUIDs are able to sense μA currents across a large bandwidth (~ 10 GHz), and their low input impedance makes them well-suited for TESes, which operate at low resistances ($R_{\text{bolo}} \sim 1 \Omega$ for SA, $R_{\text{bolo}} \sim 10 \text{ m}\Omega$ for SO). Multiplexed readout enables many mK detectors to be sensed on each 4 K SQUID amplifier, decreasing the number of wires to the mK stage. This reduced wiring reduces the thermal load on the focal plane while also reducing interconnect complexity in the cryostat. SA uses a technique called **digital frequency multiplexing (dfMUX)** [45], while SO uses a technique called **microwave multiplexing (μMUX)** [46].

SA’s dfMUX scheme is shown in Figure 2.10 and reads out 38 bolometers per SQUID using AC bias voltages at MHz frequencies. Each bolometer is in series with a tuned **LC resonator**, creating a “comb” of 38 low-impedance, non-overlapping “peaks” between 1-5 MHz [165]. These frequency channels are then voltage biased by a matching comb of AC waveforms, and when sky power changes on a given channel, the AC current across the bolometer also changes. All 38 current waveforms are fed to one SQUID, are transimpedance amplified into voltages, and are demodulated by a system of ambient electronics to extract the detector signal.

In an effort to deploy even larger, denser detector arrays, SO’s μ MUX scheme reads out $\sim 1,000$ bolometers on each 4 K-to-mK connection [46]. Instead of using the same waveform to both bias the bolometer and sense its resistance changes, as is done in SA’s dfMUX scheme, SO’s TESes are DC-biased in parallel⁴⁰ and are probed via GHz waveforms on a separate line. All 1,000 **probe tones** travel to the mK stage via a single coaxial cable and are admitted to each detector via a comb of GHz resonators. The detector is *inductively* coupled to an **rf-SQUID**, which converts changes in TES current into changes in the resonator’s *inductance* and hence its resonance frequency. This modulation of the comb’s impedance **frequency modulates**⁴¹ the voltage output, which is amplified by a high-bandwidth, high-electron-mobility transistor (HEMT) amplifier at 4 K [171] and digitized by ambient-temperature electronics [103, 76].

While both dfMUX and μ MUX are powerful techniques to multiplex bolometer arrays, each scheme has many challenges, and we briefly discuss two that are relevant to this dissertation. The first challenge is optimizing each bolometer’s responsivity. According to Equation 2.18, responsivity is governed by the bolometer’s loop gain and depends on implementation-specific details, including series impedances due to wiring [75], the electromagnetic environment at the focal plane [222], and the uniformity and consistency of the detector and resonator fabrication [224]. Parasitic effects at high frequencies can be difficult to control and therefore can lead to substantial variations in how sky power is digitized and calibrated. The second challenge is **readout noise**, which depends on a plethora of implementation-specific details, including SQUID noise, the electrical grounding scheme, and detector bias parameters. Many of these factors are particularly prominent at higher frequencies and are therefore critical to the efficacy of dfMUX and μ MUX.⁴² Because SA and SO push to very low thermal and optical noise, readout noise characterization and modeling are essential to an accurate assessment of instrument sensitivity [13].

2.5 IR-blocking design

In addition to generating a high-fidelity image of the sky at the focal plane, SA’s and SO’s OTs must reject infrared (IR) radiation in order to cool the detectors. Motivated by the thermal discussions in Chapter 8, we first focus on SA’s IR blocking system before briefly discussing SO’s. As shown in Figure 2.3b, the SA receiver has a 0.5 m-diameter window, and given a ≈ 273 K average ambient temperature at the site, the total load entering the cryostat is ~ 100 W. At the other end of the receiver, the sub-K refrigerator that cools the detectors only has ~ 5 μ W of cooling power. Therefore, the optics tube must reject power at one part $\sim 10^8$ while also being highly transparent to CMB photons at ~ 100 GHz.

⁴⁰The uniformity of the TES resistances is critical when DC-biasing many bolometers on one bias line [48].

⁴¹A feedback technique called **tone tracking** is used to keep the probe tones and resonator comb well-matched [76].

⁴²Such high-frequency concerns are much less important in DC readout schemes like **time-domain multiplexing (TDM)**.




Figure 2.11: Example IR filters from SA. Figure 2.11a shows the un-AR-coated 50 K alumina filter. Twenty-three 6N aluminum straps conduct absorbed IR radiation to the fixture plate, and the straps are flexible in order to absorb differential thermal contraction between the alumina filter and the aluminum mount. Figure 2.11b shows the reflective metal mesh filters that are directly above PB-2a’s detector wafers, which have a cutoff frequency of $\approx 5.7 \text{ cm}^{-1}$ or $\approx 170 \text{ GHz}$.

The receiver cryostat consists of three stages shown in the right panel of Figure 2.3b: the 50 K stage, 4 K stage, and mK stage [84]. The 50 and 4 K stages are cooled by two Cryomech PT-415⁴³ pulse-tube refrigerators (PTRs), one connected near the field lens and another connected near the focal plane [85]. Each PTR itself has two stages, one of which sinks \sim tens of watts from the 50 K stage and another that sinks $\sim 1 \text{ W}$ from the 4 K stage. Given this configuration, the cryostat is built like a “Russian doll” of concentric shells from the outer 300 K shell (the vacuum shell) to the inner 4 K shell. Each stage is isolated from the next by hollow G10 tubes, which have a low thermal conductivity, and the cryogenic stages are covered with **multi-layer insulation (MLI)**⁴⁴ to limit radiative transfer between them. The mK stage is cooled by a **Simon Chase refrigerator** [18], which leverages the low boiling point of the ^3He isotope to cool the focal plane to $\approx 0.3 \text{ K}$. The cooling power of the mK fridge is only $\sim 5 \mu\text{W}$ at base temperature, and the fridge must be **cycled** every day or so to recondense helium after it has boiled off.⁴⁵

The SA optics include a combination of scattering, absorptive, reflective **IR filters** that limit radiative power on the focal plane. The scattering filters are the first line of defense and include a 300 K, $\sim 300 \text{ mm}$ -thick vacuum window of closed-cell polyethylene foam⁴⁶ and 8 layers of **radio-transparent multi-layer insulation (RT-MLI)** at $\sim 170 \text{ K}$ [31]. The

⁴³Cryomech PT-415: <https://www.cryomech.com/products/pt415/>

⁴⁴RUAG MLI: <https://www.ruag.com/>

⁴⁵The helium-3 fridge system is closed such that the gas can be recycled indefinitely.

⁴⁶HD-30 Zotefoam: <https://www.zotefoams.com/>

absorptive filters are the second line of defense and include a 2 mm-thick alumina flat at 50 K [93], shown in Figure 2.11a, as well as the alumina lenses at 4 K. These alumina optics absorb $\gtrsim 1$ THz and are mounted to maximize thermal conduction to the PTR. In addition, alumina itself has a high thermal conductivity [232], which minimizes temperature gradients between the optic’s center and its edge, where it is thermally sunk. The effectiveness of the absorptive filters depends on their conductivity to the cryogenic stage, and therefore much attention is paid to how the alumina optics are heatsunk. The reflective filters are developed at Cardiff University and comprise capacitive metamaterial layers that act as low-pass reflectors [2]. These **mesh filters** consist of layered metallic sub-wavelength structures embedded in a plastic matrix to create a large (small) impedance mismatch for incident free-space radiation above (below) its cutoff frequency. Mesh filters are ubiquitous throughout GHz and THz astronomy, as they efficiently reject IR radiation while maintaining $\gtrsim 95\%$ mm-wave transmissivity. The SA optics tubes employ three sets of mesh filters whose cutoff frequencies progressively decline towards the detectors [94]. The final set of these, which reside directly above the detector wafers, are shown in Figure 2.11b.

While SO’s IR-blocking design shares many of SA’s principles, there are a few differences worth highlighting. First, SO uses silicon lenses with metamaterial AR coatings. Silicon’s IR absorptivity is generally lower than that of alumina [122, 114] and trends strongly with its resistivity, which decreases with temperature. Given these constraints, SO relies more heavily on additional filters to reject IR radiation [4]. Second, SO’s focal planes are cooled by **dilution refrigerators (DRs)**, which operate continuously and provide $\sim 100 \mu\text{W}$ of cooling power at ≈ 0.1 K. In addition, the DR has an intermediate stage with ~ 10 mW of cooling power at 1 K, which is used to cool the LAT’s (SAT’s) Lyot (aperture) stop and lenses. As a result, the parasitic optical load due to thermal emission from SO’s OT components is expected to be lower than in SA.

2.6 Technical motivations

As detailed in Chapter 1, detecting inflationary B-modes requires a CMB observatory with exquisite sensitivity. This sensitivity requirement has driven CMB telescopes to become larger and more technologically sophisticated. The process of designing these systems, setting their requirements, defining their interfaces, and evaluating their performance requires an equally sophisticated set of simulation tools. Chapters 3 and 4 detail how optical, thermal, detector, and readout noise are calculated, and Chapter 5 describes how SA and SO use sensitivity estimates to inform instrument design and fabrication. In addition, as telescopes become increasingly powerful and complex, tight control of systematic effects is becoming increasingly important. In particular, with the advent of CMB polarimetry, it is critical that hardware systems mitigate instrumental polarization and atmospheric contamination. We discuss these two systematic effects in Chapter 6 and we discuss mitigation techniques in Chapters 6-10.

Chapter 3

Instrument sensitivity

This chapter overlaps substantially with Hill, Bruno, Simon, et al. (2018) [81], which we encourage the reader to cite instead of or in addition to this dissertation.

Telescopes for CMB polarization observation are complex systems whose success relies on high-quality optical, thermal, and detector performance. Instrument evaluation is typically broken into two categories: **sensitivity** and **systematics**. Sensitivity is a measure of the instrument’s signal-to-noise and is essentially governed by two factors: the detectors’ noise level and the telescope’s optical throughput. Sensitivity analyses often assume a white noise spectrum¹ and a fixed sky signal. Systematics, on the other hand, are substantially more complex and involve system-dependent effects such as optical sidelobes,² intensity-to-polarization leakage, cross polarization, scan-synchronous signals, low-frequency noise, detector gain drift, focal plane temperature changes, readout pathologies, and more. While both sensitivity and systematics are critical components of instrument evaluation, sensitivity is usually easier to quantify and understand. Therefore, accurate signal-to-noise calculations are at the center of CMB telescope forecasting and verification. Chapter 5 presents a sensitivity calculator for Simons Array (SA), Simons Observatory (SO), and the broader CMB community, but as a precursor to that discussion, this chapter presents the sensitivity calculation itself, highlighting a few key advancements.

3.1 Noise spectrum

An instrument’s sensitivity is often quantified in terms of its **signal-to-noise ratio (SNR)**. In experiments that use photodetectors, such as mm-wave telescopes, SNR is often expressed as a **noise-equivalent power (NEP)**, which is the signal power that gives a $\text{SNR} = 1$ given

¹White noise assumes that all noise fluctuations are Gaussian, uncorrelated, and of equal amplitude on all time scales.

²The “main” lobe is the predominant component of the telescope’s angular response function, while “side” lobes are artifacts of diffraction effects in the optical system.

1 Hz of **output bandwidth**³ [161]. In this sense, NEP is spectral quantity, and the amount of noise power in a given sample depends on both the noise spectrum and the integration bandwidth⁴

$$\langle \sigma_P^2 \rangle = \int_{f_{lo}}^{f_{hi}} |\text{NEP}(f)|^2 df, \quad (3.1)$$

where f is the **audio frequency**.⁵

When quoting NEP, the noise is often approximated as “white,” or independent of frequency. In this case, the integral in Equation 3.1 becomes

$$\langle \sigma_P^2 \rangle = |\text{NEP}|^2 \Delta f, \quad (3.2)$$

where $\Delta f = f_{hi} - f_{lo}$ is the output bandwidth, which is related to the detector’s sampling rate Δt_{samp} as

$$\Delta f = \frac{1}{2\Delta t_{\text{samp}}}. \quad (3.3)$$

NEP in CMB experiments is generally decomposed into **photon noise**, **bolometer thermal noise**, **readout/amplifier noise**, **TES Johnson noise**, and extraneous noise sources, each of which we detail in the following sections. These noise sources are typically assumed to be white and uncorrelated, allowing them to be simulated independently and added in quadrature.⁶ After estimating or measuring a detector’s NEP, it is often referenced to a sky signal as a **noise-equivalent temperature (NET)**. NET is a ubiquitous figure of merit for CMB telescopes, and an accurate NET estimate relies on an accurate understanding of many instrument characteristics, including the ambient optical environment, telescope + receiver optical throughput, bolometer and amplifier noise, spillover both inside and outside of the cryostat, and beam coupling efficiency. We discuss these components of the sensitivity calculation in the following sections and assess their relative SNR impacts.

3.2 Photon statistics

It is often convenient to describe the propagation of radiation through an optical system in terms of **modes**. A single mode of the *classical* electric field is given by

$$E_p(\vec{r}, t) = E_0 \sin(\vec{k} \cdot \vec{r} - \omega t), \quad (3.4)$$

³Output bandwidth is the bandwidth at which the data is sampled. According to the Nyquist Sampling Theorem, a sinusoid with frequency f_{sig} is resolved when sampled at least twice per cycle, or at $f_{\text{samp}} \geq 2f_{\text{sig}}$. Therefore, for the definition of NEP, an output bandwidth of 1 Hz corresponds to 0.5 s of integration time.

⁴Integration bandwidth and output bandwidth are functionally equivalent. The former term is often used in the CMB community, as bolometers naturally integrate over a frequency range defined by their time constant (see Section 2.4.1).

⁵The *audio* frequency f describes data sampling and is distinct from the *microwave* frequency ν , which describes electromagnetic radiation.

⁶In other words, the noise functions for each source are assumed to be orthogonal.

where \vec{r} is position, t is time, $\vec{k} = (2\pi\nu/c)\hat{k}$ is the wave vector, $\omega = ck$ is the light's angular frequency, and the index p denotes polarization. In other words, an electromagnetic mode is described by a single vibrational component of the electric field with microwave frequency ν .

The *quantum mechanical* analog to electromagnetic modes are **coherent states**, which is a basis of full coherence and comprises eigenvectors of the photon creation operator $a_{\vec{k}p}^\dagger$

$$a_{\vec{k}p}^\dagger \left| \alpha_{\vec{k}p} \right\rangle = \alpha_{\vec{k}p} \left| \alpha_{\vec{k}p} \right\rangle . \quad (3.5)$$

The quantum electric field operates on the coherent basis as

$$\hat{E}_p^{(+)}(\vec{r}, t) \left| \alpha_{\vec{k}p} \right\rangle = E_p(\vec{r}, t) \left| \alpha_{\vec{k}p} \right\rangle , \quad (3.6)$$

where the observable quantity $E_p(\vec{r}, t)$ is the classical mode defined in Equation 3.4. Using Equations 3.5 and 3.6, we can relate the electric field operator for a single electromagnetic mode to the photon operator for a coherent state as

$$\hat{E}_p^{(+)}(\vec{r}, t) \propto \sqrt{\frac{\hbar\omega}{2\varepsilon_0}} a_{\vec{k}p}^\dagger \sin(\vec{k} \cdot \vec{r} - \omega t) , \quad (3.7)$$

where here \hbar is the reduced Planck constant and ε_0 is the permittivity of free space. Equations 3.7 and 3.4 demonstrate that by understanding the *classical* fluctuations of the field $E_p(\vec{r}, t)$, we can also understand the *quantum* fluctuations of the coherent state $\left| \alpha_{\vec{k}p} \right\rangle$. This relationship is powerful, as it enables us to use electromagnetic waves to quantify photon noise.

In order to study the statistics of thermal light for CMB applications, we note that the probability of finding a photon in a state n is governed by the Boltzmann probability distribution [123]

$$P_n = \frac{\exp(-E_n/k_B T)}{\sum_{n=0}^{\infty} \exp(-E_n/k_B T)} , \quad (3.8)$$

where each state's energy is that of a harmonic oscillator

$$E_n = \left(n + \frac{1}{2} \right) \hbar\omega . \quad (3.9)$$

Noting that the zero-point energy cancels out when Equation 3.9 is substituted into Equation 3.8, it is convenient to define the factor

$$U = \exp(-\hbar\omega/k_B T) , \quad (3.10)$$

such that the Boltzmann distribution can be rewritten as

$$P_n = \frac{U^n}{\sum_{n=0}^{\infty} U^n} = (1 - U) U^n . \quad (3.11)$$

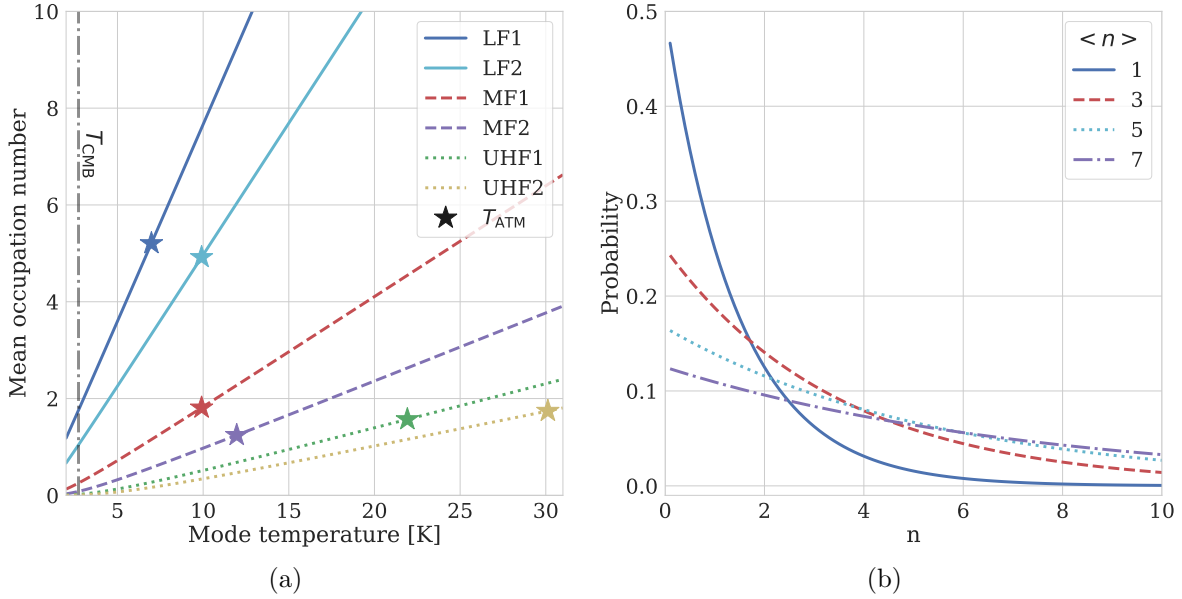


Figure 3.1: The relationship between photon occupation and electromagnetic mode temperature. Figure 3.1a shows the mean occupation number vs. temperature integrated over the SO bands in Figure 2.1. The atmospheric temperature T_{ATM} for each band is marked with a star, and the CMB temperature T_{CMB} is marked with a vertical line. Figure 3.1b shows the occupation probability distribution—or equivalently, the population density—for a few mean occupation numbers $\langle n \rangle$. The probability broadens with increasing $\langle n \rangle$, giving rise to larger mode fluctuations, while at smaller $\langle n \rangle$ the distribution looks increasingly Poissonian.

Here, we have invoked the geometric-series relation $\sum_{n=0}^{\infty} U^n = (1 - U)^{-1}$. Given this simplified form for the probability distribution, we can calculate the mean photon occupation number as

$$\langle n \rangle = \sum_{n=0}^{\infty} n P_n = (1 - U) \sum_{n=0}^{\infty} n U^n = \frac{U}{1 - U}, \quad (3.12)$$

where we note that $\sum_n n U^n = U \partial / \partial U \sum_n U^n$. Substituting Equation 3.10 into Equation 3.12 gives the mean photon number for thermal light with temperature T at frequency ω

$$\langle n \rangle = \frac{1}{\exp(\hbar\omega/k_B T) - 1}, \quad (3.13)$$

which is the familiar **Bose-Einstein distribution**.

Armed with an expression for the *mean* photon occupation number, we now find an expression for root-mean-square (RMS) photon count fluctuations [123]

$$(\Delta n)^2 = \sum_n (n - \langle n \rangle)^2 P_n = \langle n^2 \rangle - \langle n \rangle^2. \quad (3.14)$$

To simplify this expression, we rewrite P_n by plugging $U = \langle n \rangle / (1 + \langle n \rangle)$ into Equation 3.11 to obtain

$$P_n = \frac{\langle n \rangle^n}{(1 + \langle n \rangle)^{1+n}}. \quad (3.15)$$

At this point, it is useful to write P_n in terms of its factorial moments

$$\left\langle \frac{n!}{(n-r)!} \right\rangle = \sum_n n(n-1)(n-2)\dots(n-r+1)P_n = r! \langle n \rangle^r, \quad (3.16)$$

where the last equality comes from plugging Equation 3.15 in for P_n [123]. We can then use Equation 3.16 with $r = 2$ to find

$$\langle n(n-1) \rangle = 2 \langle n \rangle^2, \quad (3.17)$$

which we finally plug into Equation 3.14 to obtain

$$(\Delta n)^2 = \langle n \rangle + \langle n \rangle^2. \quad (3.18)$$

When the mean photon occupation number is $\langle n \rangle \ll 1$, which is true when $\hbar\omega \gg k_B T$, Equation 3.18 reduces to $(\Delta n)^2 \approx \langle n \rangle$, which is true of a **Poisson distribution**. Therefore, in the high-frequency and/or low-temperature limit, photon noise is predominantly **shot noise** and its fluctuations are *uncorrelated*. On the other hand, when the occupation is $\langle n \rangle \gg 1$, which is true when $\hbar\omega \ll k_B T$, Equation 3.18 reduces to $(\Delta n)^2 \approx \langle n \rangle^2$, which is true of an exponential distribution. Therefore, in the low-frequency and/or high-temperature limit, photon noise is predominantly **Bose noise** (also called **wave noise**) and fluctuations are, in general, *correlated*. The regime where $\hbar\omega \ll k_B T$ is called the **Rayleigh-Jeans (RJ) limit** and is invoked widely in radio astronomy, especially $\lesssim 100$ GHz.

3.3 Photon NEP

According to Equations 3.9 and 3.18, the variance of the energy in an electromagnetic mode is determined solely by the mode's mean photon occupation number. In this section, we relate these photon statistics to fluctuations in optical power and quantify photon noise using NEP [113, 161, 15, 241].

Assuming that all N_{elem} optics are isothermal blackbodies,⁷ the detected power spectral density $p_i(\nu)$ for an optical element i is determined by its physical temperature T_i , the aggregate transmissivity of all optics between it and the detector $[\eta_{i+1}(\nu), \dots, \eta_{N_{\text{elem}}}(\nu)]$, its

⁷Isothermality is not generally true for optics within CMB telescopes, which are illuminated at the center and heatsunk at the edges, but spatial temperature variations across the optic are typically $< 10\%$ and are often averaged over the focal plane. Therefore the assumption of isothermal optical elements is appropriate for sensitivity estimations.

emissivity $\epsilon_i(\nu)$, its spillover fraction $\beta_i(\nu)$, the effective temperature of its spilled radiation $T_{\beta;i}$, its scattering fraction $\delta_i(\nu)$, and the effective temperature of its scattered radiation $T_{\delta;i}$

$$p_i(\nu) = \prod_{j=i+1}^{N_{\text{elem}}} \eta_j(\nu) [\epsilon_i(\nu)S(T_i, \nu) + \beta_i(\nu)S(T_{\beta;i}, \nu) + \delta_i(\nu)S(T_{\delta;i}, \nu)] . \quad (3.19)$$

The power spectral density $S(T, \nu)$ of the emitted, spilled, and scattered radiation into a mode with frequency ν is given by the **Planck spectral density**

$$S(T, \nu) = A\Omega \frac{h\nu^3}{c^2} n(T, \nu) = h\nu n(T, \nu) . \quad (3.20)$$

Here, $n(T, \nu)$ is the mean Bose occupation number defined in Equation 3.13, and the entendue $A\Omega$ for a diffraction-limited, single-moded detector, is given by the square of the detected wavelength [177, 161]

$$A\Omega = \left(\frac{c}{\nu}\right)^2 = \lambda^2 . \quad (3.21)$$

The RMS of power fluctuations in a single thermal mode is

$$(\Delta P_{\text{mode}})^2 = (h\nu)^2 (\langle n \rangle + \langle n \rangle^2) , \quad (3.22)$$

and we can convert this relation into an NEP using Equations 3.19 and 3.20 as

$$\text{NEP}_{\text{ph}} = \sqrt{2 \int_0^\infty \left[h\nu B(\nu) \sum_{i=1}^{N_{\text{elem}}} p_i(\nu) + \left(B(\nu) \sum_{i=1}^{N_{\text{elem}}} p_i(\nu) \right)^2 \right] d\nu} , \quad (3.23)$$

where $B(\nu)$ is the detector's transmissivity vs. frequency, also called the **detector band**. This relation is often approximated as

$$\text{NEP}_{\text{ph}} \approx \sqrt{2 \left[h\nu_c P_{\text{opt}} + \frac{P_{\text{opt}}^2}{\Delta\nu} \right]} , \quad (3.24)$$

where ν_c and $\Delta\nu$ are the central frequency and bandwidth of the detector band $B(\nu)$, where the total detected optical power is

$$P_{\text{opt}} = \int_0^\infty \left[\sum_{i=1}^{N_{\text{elem}}} p_i(\nu) \right] B(\nu) d\nu , \quad (3.25)$$

and where the approximation remains valid in the RJ limit where $h\nu \ll k_B T$. Note that the first term in Equations 3.23 and 3.24 represents shot noise $\text{NEP}_{\text{shot}}^2 \propto \langle n \rangle$, while the second term represents wave noise $\text{NEP}_{\text{wave}}^2 \propto \langle n \rangle^2$.




Figure 3.2: A schematic of the one-dimensional optical system assumed in the optical power calculation. CMB photons are attenuated by the atmosphere and telescope optics, while atmospheric emission is attenuated by just the telescope optics. The dewar itself also emits, and all attenuated radiation impinges on the detector pixel, where it is finally sensed.

3.4 Optical power

An accurate calculation of photon noise relies on an accurate understanding of each optical element's power spectral density $p(\nu)$, which is defined in Equation 3.19. Assume that the optical stack, which comprises the telescope's mirrors, the receiver's lenses and filters, and the focal plane's optics, can be represented as a one-dimensional array of blackbody absorbers/emitters/attenuators in **thermal equilibrium**, as shown in Figure 3.2. While isothermal equilibrium is not true in reality, it is a reasonable approximation for a white-noise, single-mode NET estimate. There are four primary components to the detector's optical load that we detail in the following subsections: emissivity, spillover, scattering, and transmissivity.

3.4.1 Emissivity and absorptivity

Emissivity quantifies the propensity of an optical element to emit thermal radiation. The emissivity of a dielectric is determined by its **loss tangent**

$$\tan \delta = \frac{2\pi\nu\epsilon'' + \sigma}{2\pi\nu\epsilon'} \approx \frac{\epsilon''}{\epsilon'}, \quad (3.26)$$

where the conductivity is σ , the dielectric constant is

$$\epsilon = \epsilon' - i\epsilon'', \quad (3.27)$$

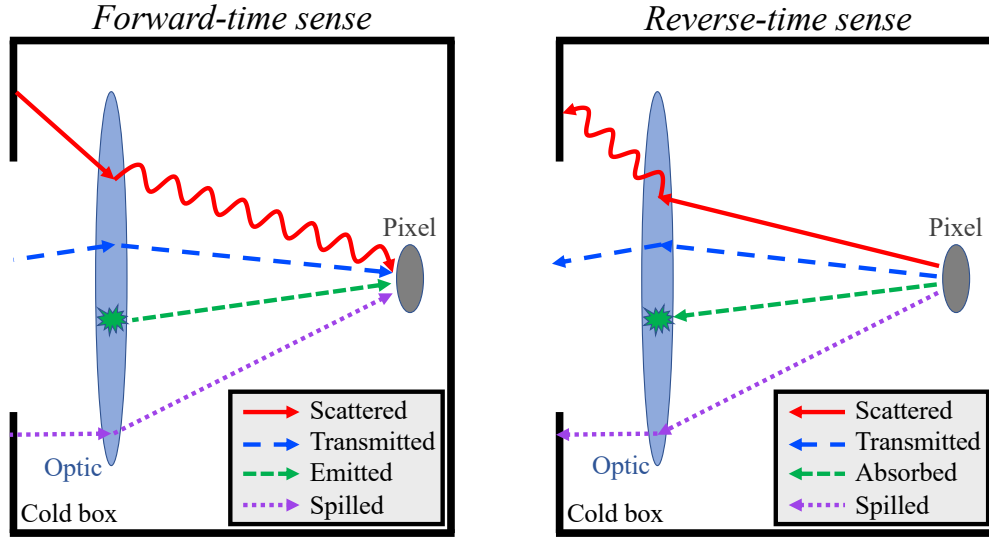


Figure 3.3: A schematic of forward-time-sense and reverse-time-sense calculations as they apply to transmissivity, absorptivity/emissivity, scattering fraction, and spillover fraction. When estimating sensitivities, it is often useful to think in the reverse-time sense, or in the paradigm of the detector pixel radiating towards the sky.

and the approximation in Equation 3.26 applies when $\sigma \ll 2\pi\nu\epsilon''$. The dielectric's emissivity can then be written as

$$\epsilon_{\text{diel}}(\nu) = 1 - \exp\left(-\frac{2\pi\nu nt}{c} \tan \delta\right) \approx \frac{2\pi\nu nt}{c} \tan \delta, \quad (3.28)$$

where t is the optical thickness, c is the speed of light in a vacuum, n is the dielectric's refractive index, and the approximation applies when $\tan \delta \ll 1$.

For conductors, such as mirrors, emissivity is generated by Ohmic losses and is estimated using the **Hagens-Ruben relation** [68, 180]

$$\epsilon_{\text{cond}}(\nu) = 4\sqrt{\frac{\pi\nu\epsilon_0}{\sigma}}, \quad (3.29)$$

where σ is the conductivity and ϵ_0 is the permittivity of free space. Note that at oblique incident angles, the conductor's skin depth is different for the S and P polarizations, which creates a slight degree of polarized emission. This polarization leakage is an important systematic effect, but it is irrelevant when estimating optical NEP, which is concerned with spatially- and polarization-averaged loading estimates.

Under the assumption of thermal equilibrium, each optic's emission results from the thermodynamic motions of its composing molecules, which are Gaussian random and wide-

band.⁸ In this paradigm, emissivity $\epsilon(\nu)$ and **absorptivity** $\alpha(\nu)$ are equivalent

$$\alpha(\nu) = \frac{P_{\text{abs}}(\nu)}{P_{\text{in}}(\nu)} = \epsilon(\nu), \quad (3.30)$$

where $P_{\text{in}}(\nu)$ and $P_{\text{abs}}(\nu)$ are the input and absorbed power, respectively. Therefore, not only does dielectric emissivity lead to parasitic loading, but it also leads to signal attenuation.

3.4.2 Spillover fraction

Spillover fraction quantifies the amount of incident power that spills over an optical element’s optically viable area⁹

$$\beta(\nu) = \frac{P_{\text{spill}}(\nu)}{P_{\text{in}}(\nu)}, \quad (3.31)$$

where $P_{\text{in}}(\nu)$ and $P_{\text{spill}}(\nu)$ are the input and spilled power, respectively. Spillover is most easily imagined in the reverse-time sense (see Figure 3.3). If some power from a detector pixel falls outside an element’s optically viable area, this radiation will not propagate to the sky but will instead terminate elsewhere. Because this spillover allows the detector to “see” non-optical surfaces in the telescope system, photons emitted (in the forward-time sense) from such surfaces are detected, hence dissipating non-sky optical power on the bolometer.

The spillover’s intensity is governed by its **effective temperature** T_{eff} , or the effective temperature on which the spilled reverse-time-sense photons terminate

$$T_{\text{eff}} = \epsilon_{\text{surface}} T_{\text{surface}}. \quad (3.32)$$

Here, T_{surface} is the surface’s physical temperature, $\epsilon_{\text{surface}}$ is its emissivity, and for spillover in Equation 3.19, $T_{\beta} = T_{\text{eff}}$. Equation 3.32 relies on the RJ approximation, which is usually accurate enough for spillover estimates during sensitivity forecasting. Some common examples include spill over the aperture (which terminates on the stop), spill over a lens (which terminates somewhere in the cryostat), and spill over the primary mirror (which terminates on the sky or ground).

3.4.3 Scattering fraction

Scattering fraction quantifies the scattering in a given optical element

$$\delta(\nu) = \frac{P_{\text{scatt}}(\nu)}{P_{\text{in}}(\nu)}, \quad (3.33)$$

⁸These broadband, thermomolecular phenomena are much more important than any atom-light interactions, which can produce narrow-band features in the emission spectrum.

⁹Often times, the optically viable area of a lens, thermal filter, or a wave plate is the optic’s clear-aperture diameter.

where $P_{\text{in}}(\nu)$ and $P_{\text{scatt}}(\nu)$ are the input and scattered power, respectively. In a similar sense to spillover, scattering is most easily imagined in the reverse-time sense. If some power from a detector pixel is scattered by an optical element, this radiation will not necessarily propagate to the sky but may instead terminate on some surface that will, in the forward-time sense, emit optical power on the bolometer. The scattering's intensity is similarly governed by the effective temperature in Equation 3.32, and for scattering in Equation 3.19, $T_{\delta} = T_{\text{eff}}$.

There are many mechanisms for scattering, but the two most common in mm-wave experiments are Mie scattering and Ruze scattering. **Mie scattering** arises due to any irregularity in an otherwise homogeneous, well-ordered medium. Common examples in CMB telescopes include voids in dielectric substrates, air bubbles or gaps between anti-reflection coating layers, and fillers in composite materials. If the scattering center is much smaller than the wavelength, which is often a good approximation at radio frequencies, Mie scattering is well-described by the **Rayleigh scattering cross section** [121]

$$\sigma_{\text{Ray}} = \frac{2\pi^5}{3} \frac{d^6}{\lambda^4} \left(\frac{n_2^2 - n_1^2}{n_2^2 + 2n_1^2} \right)^2, \quad (3.34)$$

where d is the particle/void/deformity's diameter, n_2 is its refractive index, and n_1 is the refractive index of the medium. The aggregate impact of Rayleigh scatterers is then given by the **Beer-Lambert law**

$$\delta_{\text{Ray}}(\nu) \approx 1 - \exp(-\sigma_{\text{Ray}} N z), \quad (3.35)$$

where N is the number density of scatterers, z is the optical path length in the scattering medium, and the approximation assumes that the medium is uniform. Scattering is an especially important parameter in cryogenic systems because it does not improve with decreasing temperature (unlike absorption loss).

Scattering from reflectors typically arises due to surface roughness and is quantified by **Ruze scattering**

$$\delta_{\text{Ruze}}(\nu) = 1 - \exp \left[\left(\frac{4\pi\sigma_{\text{surf}}\nu}{c} \right)^2 \right], \quad (3.36)$$

where σ_{surf} is the reflector's RMS roughness. At millimeter wavelengths, Ruze scattering tends to be negligible for terrestrial experiments,¹⁰ but its impact on sensitivity can become more prominent in low-load environments, such as in outer space, where polished mirrors may be beneficial.

¹⁰In fact, many terrestrial experiments *intentionally* roughen their mirrors as a safety against accidentally pointing the telescope at the sun. In this scenario, the roughened mirror will diffuse the sun's short wavelengths, preventing heat-induced damage to the experiment.

3.4.4 Reflectivity

Reflectivity is the fraction of incident power that is reflected by any dielectric optic

$$r(\nu) = \frac{P_{\text{refl}}(\nu)}{P_{\text{in}}(\nu)} . \quad (3.37)$$

Reflections arise at interfaces between media with mismatched refractive indices, and anti-reflection (AR) coatings are designed to limit these reflections. An in-depth discussion of AR coatings and their realized reflectivities is presented in Chapter 10.

3.4.5 Transmissivity

Transmissivity is the ratio of transmitted optical power to incident optical power

$$\eta(\nu) = \frac{P_{\text{trans}}(\nu)}{P_{\text{in}}(\nu)} \quad (3.38)$$

and is the product of absorptivity, spillover fraction, scattering fraction, and reflectivity

$$\eta(\nu) = [1 - \alpha(\nu)] [1 - \beta(\nu)] [1 - \delta(\nu)] [1 - r(\nu)] . \quad (3.39)$$

Transmissivity is effectively synonymous with transparency, and maximizing transmission is a core principle of high-throughput telescopes.

3.4.6 Optical throughput and optical efficiency

Optical throughput is the total transmission through the telescope system and is defined as

$$\eta_{\text{through}} = \prod_{i=0}^{N_{\text{inst}}} \eta_i \quad (3.40)$$

where N_{inst} includes all optical elements, including the detector.¹¹ Throughput can be thought of as the number of sky photons detected divided by the number of in-band photons within the telescope’s beam.

As noted in Section 2.3.3, the aperture stop (or the Lyot stop in telescopes with reimaging optics) defines the telescope’s angular resolution by truncating the beams from detector pixels. Unlike other transmissivity degradations, that of the aperture is *intentional*, and therefore CMB instrumentalists often quote an instrument’s **optical efficiency** as

$$\eta_{\text{eff}} = \frac{\eta_{\text{through}}}{\eta_{\text{apert}}} , \quad (3.41)$$

where η_{apert} is the aperture efficiency (also called the beam-coupling efficiency) defined in Equation 2.10. The theoretical maximum of η_{eff} , in the limit of perfectly transparent optics, is 100%, while that of η_{through} is $(1 - \eta_{\text{apert}})$.

¹¹Note that throughput excludes any attenuation due to the atmosphere.

3.4.7 Sky temperature and telescope temperature

In the broadest sense, there are two “sources” of optical loading on the bolometer: the sky and the telescope

$$P_{\text{opt}} = P_{\text{sky}} + P_{\text{tel}} . \quad (3.42)$$

Here, sky power is given by

$$P_{\text{sky}} = \int_0^\infty \left[\sum_i^{N_{\text{sky}}} p_i(\nu) \right] B(\nu) d\nu , \quad (3.43)$$

where $B(\nu)$ is the detector’s band and where the summation runs over all N_{sky} sky sources, including the CMB, galactic dust emission, synchrotron emission, and atmospheric emission. In a similar manner, instrument-induced power is given by

$$P_{\text{tel}} = \int_0^\infty \left[\sum_i^{N_{\text{inst}}} p_i(\nu) \right] B(\nu) d\nu , \quad (3.44)$$

where the summation runs over all N_{inst} instrument sources, such as the telescope mirrors, vacuum window, lenses, thermal filters, aperture stop, and focal plane coupling optics.

Given Equation 3.42, CMB instrumentalists often describe the sky and instrument loads in terms of their effective RJ temperatures

$$P_{\text{opt}} = \eta_{\text{through}} k_B \Delta\nu (T_{\text{sky}} + T_{\text{tel}}) , \quad (3.45)$$

where $\Delta\nu$ is the telescope’s effective bandwidth¹² and η_{through} is defined in Equation 3.40. Equation 3.45 describes the total power from the sky (telescope) as being due to a perfect blackbody with temperature T_{sky} (T_{tel}) placed in front of a 0 K telescope with throughput η_{through} . This scheme allows parasitic power from telescope optics to be quickly compared against sky power, and a typical goal is to make $T_{\text{tel}} \ll T_{\text{sky}}$.

3.5 Bolometer thermal carrier NEP

Thermal carrier noise arises due to fluctuations in heat flow between the bolometer’s absorbing element and the bath to which it is weakly connected and is given by the equation [129]

$$\text{NEP}_g = \sqrt{4k_B F_{\text{link}} T_c^2 g} , \quad (3.46)$$

where T_c is the bolometer’s operating temperature—which for a TES is equivalent to its **superconducting transition temperature**— g is the **dynamic thermal conductance**

¹²The telescope’s effective bandwidth is typically defined as the bandwidth between the detector band’s $B(\nu)$ -3 dB points.




Figure 3.4: Dependence of NEP_g and P_{sat} on the bolometer's transition temperature T_c for various bath temperatures T_b . Figure 3.4a assumes a constant P_{sat} , while Figure 3.4b assumes $Ak_0/(n+1) = 240 \text{ pW/mm K}^4$ and a bolometer “leg” length of 1 mm. Reducing the bath temperature substantially improves thermal carrier noise while having little impact on saturation power. For a given T_b , the optimum transition temperature is $T_c \approx 1.7T_b$.

from the bolometer's absorbing element to the bath (see Equation 2.14), and F_{link} is a numerical factor that depends on the link's **thermal conduction index** n [5]

$$F_{\text{link}} = \frac{\int_{T_b}^{T_c} \left[\frac{Tk(T)}{T_c k(T_c)} \right]^2 dT}{\int_{T_b}^{T_c} \left[\frac{k(T)}{k(T_c)} \right] dT} = \frac{n+1}{2n+3} \frac{1 - (T_b/T_c)^{2n+3}}{1 - (T_b/T_c)^{n+1}}. \quad (3.47)$$

Here, T_b is the bath temperature, and the conductivity between the bolometer and the bath is assumed to be $k(T) = k_0 T^n$.

The power flowing from the bolometer to the thermal bath is called the **saturation power**¹³

$$P_{\text{sat}} = \int_{T_b}^{T_c} \frac{A}{l} k(T) dT = \frac{A}{l} \frac{k_0}{(n+1)} (T_c^{n+1} - T_b^{n+1}), \quad (3.48)$$

¹³It is worth noting that different TES bolometer engineers sometimes use different nomenclature. For example, some TES experts will refer to the saturation power P_{sat} as the power at which the TES saturates or becomes normal. However, in the Simons Array collaboration, saturation power is defined as the power dissipated by the TES during operation, and the power at which the bolometer saturates is instead called P_{turn} . Therefore, $P_{\text{sat}} \approx 0.9P_{\text{turn}}$.

where A and l are the thermal link's cross-sectional area and length, respectively. We can then calculate the dynamic thermal conductance as

$$g = \frac{\partial P_{\text{sat}}}{\partial T_c} = \frac{A}{l} k_0 T_c^n. \quad (3.49)$$

While A , l , and k_0 are familiar quantities to those who fabricate bolometers, saturation power is what matters for bolometer operation, and therefore we rewrite g as

$$g = P_{\text{sat}}(n+1) \frac{T_c^n}{T_c^{n+1} - T_b^{n+1}}. \quad (3.50)$$

Plugging Equation 3.50 into Equation 3.46 gives a more phenomenological form for thermal carrier noise

$$\text{NEP}_g = \sqrt{4k_B F_{\text{link}} P_{\text{sat}}} \sqrt{\frac{(n+1) T_c^{n+2}}{T_c^{n+1} - T_b^{n+1}}}. \quad (3.51)$$

While Equation 3.51 is determined by a small set of parameters, in reality NEP_g is also influenced by the specifics of the bolometer's geometry, composition, and fabrication. For example, transition-edge sensors (TESes) have well-known pathological noise sources, such as flux flow noise and non-equilibrium Johnson noise, that often increase the measured NEP_g beyond the theoretical formalism presented here [58]. These effects are empirically absorbed into F_{link} during bolometer characterization both in the lab and in the field (e.g., [32]).

3.6 Readout NEP

As presented in Section 2.4, many modern CMB detectors are low-impedance, voltage-biased bolometers read out using superconducting quantum interference device (SQUID) transimpedance amplifiers [193]. SQUIDs are current sensors, and a power fluctuation on a voltage-biased bolometer gives rise to a current fluctuation at the amplifier's input. Therefore, **readout noise** is typically characterized as a **noise-equivalent current (NEI)**, which has units of A/ $\sqrt{\text{Hz}}$. In order to refer NEI_{read} to an NEP, we must consider the bolometer's responsivity $S_I = dI/dP$ (see Equation 2.18). We conveniently write the responsivity as

$$S_I = -S_{\text{fact}} \frac{1}{V_{\text{bias}}}, \quad (3.52)$$

where V_{bias} is the bias voltage across the bolometer defined in Equation 2.11. In the presence of electrothermal feedback, the bolometer's responsivity can typically be written in terms of its DC loop gain \mathcal{L} (see Equation 2.17), time constant τ , and the optical-power modulation mode frequency ω (see Equation 2.14) as

$$S_{\text{fact}} = -\tilde{S}_{\text{fact}} \frac{\mathcal{L}}{\mathcal{L} + 1} \frac{1}{1 + i\omega\tau}, \quad (3.53)$$

where \tilde{S}_{fact} is 1 if V_{bias} is DC (e.g., time-division multiplexing or microwave multiplexing) or $\sqrt{2}$ if it is AC (e.g., frequency-domain multiplexing), as shown in Equation 2.19.

Characterization of the bolometer’s loop gain \mathcal{L} and time constant τ are hugely important to determining detector responsivity and hence to correctly relating amplifier output to sky power. As a result, detector responsivity is calibrated frequently in the field using both celestial sources and calibration hardware. Given an estimate of S_{fact} , **readout NEP** can be calculated as

$$\text{NEP}_{\text{read}} = \frac{\text{NEI}_{\text{read}}}{S_{\text{I}}} = \frac{\sqrt{R_{\text{bolo}} P_{\text{bias}}}}{S_{\text{fact}}} \text{NEI}_{\text{read}}, \quad (3.54)$$

where the RMS voltage bias is written in terms of the bolometer’s operating resistance R_{bolo} as $V_{\text{bias}} = \sqrt{R_{\text{bolo}} P_{\text{bias}}}$.

Readout noise comprises a multitude of noise sources and electrical effects, and therefore readout NEI is *not* synonymous with **SQUID NEI**. SQUID noise is typically quite small with $\text{NEI}_{\text{SQUID}} \sim 5 \text{ pA}/\sqrt{\text{Hz}}$, while noise in dfMUX systems can range anywhere from $10 \sim 30 \text{ pA}/\sqrt{\text{Hz}}$ [13] and noise in μMUX systems can be even higher [46]. An effective technique to maximize responsivity is to minimize R_{bolo} . SO bolometers are DC biased and have a resistance of $\mathcal{O}(10) \text{ m}\Omega$ [192], while SA bolometers are AC biased—necessitating a larger resistance to trivialize the impact of parasitic series impedances¹⁴—and have a resistance of $\mathcal{O}(1) \Omega$ [224]. Minimizing readout noise is a tenet of modern CMB experiments, which face challenges associated with large multiplexing factors at high frequencies.

3.7 Johnson NEP

Johnson noise arises due to thermal fluctuations in the bolometer that cause fluctuations in its resistance. **Johnson noise equivalent current** $\text{NEI}_{\text{Johnson}}$ is given by

$$\text{NEI}_{\text{Johnson}} = \frac{1}{\mathcal{L}} \sqrt{\frac{4k_{\text{B}} T_{\text{c}}}{R_{\text{bolo}}}}, \quad (3.55)$$

where again, T_{c} the bolometer’s operating temperature and R_{bolo} is its (average) operating resistance. This current noise can be converted to an NEP using the bolometer responsivity S_{I} defined in Equation 3.53 as

$$\text{NEP}_{\text{Johnson}} = \frac{\text{NEI}_{\text{Johnson}}}{S_{\text{I}}} = \tilde{S}_{\text{fact}} \frac{\mathcal{L} + 1}{\mathcal{L}^2} (1 + i\omega\tau) \sqrt{4k_{\text{B}} T_{\text{c}} P_{\text{bias}}}, \quad (3.56)$$

where P_{bias} is defined in Equation 2.11.

We highlight two important notes regarding Equation 3.56. First, in the limit of a large loop gain $\mathcal{L} \gg 1$, which is enabled by the TES bolometer’s large dR/dT , $\text{NEP}_{\text{Johnson}} \rightarrow 0$. Second, Johnson noise is suppressed by a factor of $1/\mathcal{L}$ with respect to the readout noise

¹⁴Typical parasitic impedances for AC-biased bolometers include series inductance and capacitance in the cables and connectors between the SQUID on 4 K stage and the detectors on mK stage.

NEP_{read} in Equation 3.54. For these two reasons, it is often customary to ignore Johnson noise when estimating the NEP of modern TESes. Even in the worse-possible scenario of $\mathcal{L} = 1$, the estimated $\text{NEI}_{\text{Johnson}}$ for an AC-biased (DC-biased) bolometer is ~ 5 (20) pA/ $\sqrt{\text{Hz}}$, which is far below the NEI_{read} values discussed in Section 3.6. Furthermore, $\text{NEP}_{\text{Johnson}}$ becomes negligible when assuming a healthier, more typical loop gain of $\mathcal{L} \sim 10$.

3.8 NET

Assuming that all noise sources are white and uncorrelated, a detector's total NEP is

$$\text{NEP}_{\text{det}} = \sqrt{\text{NEP}_{\text{ph}}^2 + \text{NEP}_{\text{g}}^2 + \text{NEP}_{\text{read}}^2}, \quad (3.57)$$

where NEP_{ph} is the photon NEP, NEP_{g} is the thermal carrier NEP, and NEP_{read} is the readout NEP. Because a bolometer measures fluctuations in sky temperature, NEP_{det} is conventionally converted into a **noise-equivalent temperature (NET)**

$$\text{NET}_{\text{det}} = \frac{\text{NEP}_{\text{det}}}{\sqrt{2} (dP/dT_{\text{sky}})}, \quad (3.58)$$

where the $\sqrt{2}$ arises due to a unit conversion from output bandwidth in $1/\sqrt{\text{Hz}}$ to integration time in \sqrt{s} , as described in Section 3.1.

CMB experiments often reference sky power either to the CMB's physical temperature $T_{\text{phys}} = T_{\text{CMB}} = 2.725$ K or to some other source's RJ temperature T_{RJ} . The conversion factor in W/K from fluctuations in physical temperature to fluctuations in power is

$$\frac{dP}{dT_{\text{phys}}} = \xi \int_0^\infty \left[\frac{1}{k_{\text{B}}} \left(\frac{h\nu}{T_{\text{phys}} (\exp[h\nu/k_{\text{B}}T_{\text{phys}}] - 1)} \right)^2 \exp[h\nu/k_{\text{B}}T_{\text{phys}}] \right] B(\nu) d\nu, \quad (3.59)$$

where $\xi \leq 1$ is an overall signal degradation factor,¹⁵ and $B(\nu)$ is the detector's band (examples in Figure 2.1). The conversion factor in W/K_{RJ} is

$$\frac{dP}{dT_{\text{RJ}}} = \xi \int_0^\infty k_{\text{B}} B(\nu) d\nu, \quad (3.60)$$

where $B(\nu)$ is the detector bandpass.

When reconstructing the sky during analysis, data from each detector are coadded¹⁶ to improve signal-to-noise in the map domain. To quantify this SNR increase in the time

¹⁵For example, the overall signal degradation factor might be associated with imperfect image formation at the focal plane.

¹⁶Coaddition is a process of synthesizing maps from individual detectors and adding them to a common grid to produce an output image with improved signal-to-noise.




Figure 3.5: The conversion factor from RJ temperature fluctuations to those of a blackbody with temperature T_{phys} in the SO detector bands (see Figure 2.1). The most commonly used factor is $dT_{\text{CMB}}/dT_{\text{RJ}}$, which converts NET_{RJ} to NET_{CMB} , but the conversion factor for a 10 K atmosphere is also plotted for reference. The RJ approximation works well at high temperatures and low frequencies and is therefore often used to quantify noise in radio telescopes and the intensity of bright sources.

domain, we define **array NET** as the **inverse-variance-weighted average**¹⁷ of individual-detector NETs

$$\text{NET}_{\text{arr}} = \sqrt{\frac{1}{\sum_i^{N_{\text{det}}} 1/\text{NET}_{\text{det},i}^2}} \Gamma, \quad (3.61)$$

where the sum is over all **yielded** detectors in the array and where Γ quantifies the degree to which white noise is correlated between detector pixels on the focal plane (see Section 3.9). If we assume that all operational detectors have the same NET_{det} , an assumption that is often invoked when forecasting *median* sensitivity estimates, the inverse-variance-weighted average becomes the simple average

$$\text{NET}_{\text{arr}} = \frac{\text{NET}_{\text{det}}}{\sqrt{Y N_{\text{det}}}} \Gamma, \quad (3.62)$$

where N_{det} is the *total* number of detectors and Y is the **detector yield**. Note that nonoperational detectors effectively have infinite NET_{det} .

¹⁷In order to maximize signal-to-noise in the coadded map, each detector's data is weighted by its inverse variance, hence giving more weight to high-SNR detectors and minimizing the variance in the final map.

3.9 Photon noise correlation factor

It is possible to **oversample** the focal plane by deploying more detector pixels than there are independent **spatial modes** in the telescope optics [149]. When the pixel density is higher than the mode density, photon wave noise correlates between detector outputs [241]. This correlation slows the rate at which noise is averaged down during detector coaddition and is quantified by Γ in Equation 3.62. A detailed derivation of **photon noise correlations** is discussed in Chapter 4, but we briefly review the calculation here for the completeness of the sensitivity presentation.

The number of independent spatial modes at the focal plane is determined by the telescope's beam size and field of view, and it is functionally synonymous with the number of non-overlapping telescope beams. For example, if the telescope's field of view is 5.0 deg and its beam size is 0.1 deg, the number of independent modes is $\approx (5.0/0.1)^2$. This approximation is rough and depends on the details of the beam's profile and sidelobes, but it is nonetheless intuitive and functional. The telescope's magnification determines the spatial modes' linear dimension at the focal plane, and to a good approximation, the correlation length is $\approx 1.2 F\lambda$. Therefore, if neighboring detectors pixels are $< 1.2 F\lambda$ apart, their noise will correlate.

Intensity correlations are determined by the **Hanbury Brown-Twiss (HBT) coefficient** [24, 23]

$$\gamma_{ij} = \frac{\langle |e_i|^2 |e_j|^2 \rangle - \langle |e_i|^2 \rangle \langle |e_j|^2 \rangle}{\text{RMS}(|e_i|^2) \text{RMS}(|e_j|^2)}, \quad (3.63)$$

where angle brackets denote the statistical average. The field amplitude e_i is an integral of the incoming electric field for detector i with beam $b_i(x, y)$ and optical path length to the source plane $\ell_i(x, y)$

$$e_i = \iint dx dy e^{2\pi i \ell_i(x, y)} b_i(x, y) E(x, y). \quad (3.64)$$

Because the stop is not 0 K, correlations can arise due to radiation outside the aperture as well as within it, and these details are discussed in Chapter 4. In this section, we focus only on radiation within the aperture, which is usually the dominant noise contributor for ground-based telescopes.

The **cumulative correlation coefficient** γ is a summation of the correlation coefficients between all N_{pix} detector pixels on the focal plane

$$\gamma = \frac{1}{N_{\text{pix}} - 1} \sum_i \sum_{j \neq i} \gamma_{ij}. \quad (3.65)$$

These correlations then propagate to NET_{arr} by impeding the degree to which wave noise is suppressed during inverse-variance averaging

$$\text{NET}_{\text{arr}} = \sqrt{\frac{\text{NET}_{\text{shot}}^2 + (1 + \gamma)\text{NET}_{\text{wave}}^2 + \text{NET}_{\text{g}}^2 + \text{NET}_{\text{read}}^2}{Y N_{\text{det}}}}, \quad (3.66)$$




Figure 3.6: Wave-noise fraction, which is defined as $\text{NEP}_{\text{wave}} / \sqrt{\text{NEP}_{\text{wave}}^2 + \text{NEP}_{\text{shot}}^2}$, assuming a 20 GHz-wide top-hat band and an optical throughput between the source and detector of 30%.

where NET_{shot} and NET_{wave} are related to NEP_{shot} and NEP_{wave} (see Section 3.3) via Equations 3.59 or 3.60. We can now write the array **NET correlation suppression factor** Γ in Equation 3.62 as

$$\Gamma = \sqrt{1 + \frac{\gamma \text{NET}_{\text{wave}}^2}{\text{NET}_{\text{shot}}^2 + \text{NET}_{\text{wave}}^2 + \text{NET}_{\text{g}}^2 + \text{NET}_{\text{read}}^2}}. \quad (3.67)$$

As is evident in Equation 3.67, the impact of correlations on NET_{arr} depends not only on the HBT coefficient but on the relative contribution of wave noise to the total noise, necessitating an accurate estimate of P_{opt} for an accurate estimate of Γ .

3.10 N_ℓ and mapping speed

As presented in Section 3.8, NET has units of $\mu\text{K}\sqrt{\text{s}}$ and measures noise in the detector time stream with respect to fluctuations in sky temperature. It is often useful to convert this time-domain noise into a **map depth**, or noise in a sky-map domain, which is defined as

$$\sigma_{\text{map}} = \text{NET}_{\text{arr}} \sqrt{\frac{4\pi (10800/\pi)^2 f_{\text{sky}}}{\eta_{\text{obs}} t_{\text{obs}}}}, \quad (3.68)$$

in units of K-arcmin. Here, we convert NET from a temporal spectral density to a spatial spectral density by dividing by the **integration time**—or the product of the observatory’s lifetime t_{obs} and its **observation efficiency** η_{obs} —and by multiplying by the **sky area**—or the total number of arcminutes on the sky times the observed **sky fraction** f_{sky} . In addition to improving NETs, map depth is improved both by integrating for longer and by observing smaller sky patches. However, as shown in Equation 1.34, a smaller sky fraction leads to larger cosmic variance, and therefore trade-offs between map depth and sky coverage are important considerations when optimizing observation strategies.

The figure of merit when measuring CMB power spectra (see Figures 1.10, 1.11, and 1.12) is \mathbf{N}_ℓ , which measures instrument noise in K^2 as a function of angular multipole number ℓ . The power spectrum error is

$$\Delta C_\ell = \sqrt{\frac{2}{(2\ell + 1) f_{\text{sky}}}} (C_\ell + N_\ell) , \quad (3.69)$$

where the first term is due to cosmic variance (see Section 1.6) and the second term is due to instrumental noise. If N_ℓ is due to white noise alone, it can be written using the **Knox Formula** [106]

$$N_\ell^{\text{white}} = 4\sigma_{\text{map}}^2 e^{\ell^2 \sigma_{\text{beam}}^2} , \quad (3.70)$$

where the exponential term is called the **beam transfer function** and σ_{beam} is the experiment’s angular resolution. The sensitivity degradation at angular scales smaller than the telescope’s beam is an important driver of the primary aperture size. For example, a large-aperture telescope with $\sigma_{\text{beam}} = 0.1$ deg beam will make a much more sensitive measurement at $\ell = 2,000$ than a small-aperture telescope with $\sigma_{\text{beam}} = 1$ deg, even if the small-aperture telescope has better NETs.

It is important to note that Equations 3.68 and 3.70 are simplistic approximations and that the conversion from NET to map depth is, in general, a complicated product of instrument properties, observation strategies, and analysis techniques. For example, practical observations cover the sky non-uniformly, and each sky pixel’s “weight” is related to its “hit count,” or the number of times it is observed. In addition, a sky patch’s location¹⁸ and orientation¹⁹ when scanned further complicates map-domain processing.

Despite the assumptions throughout this chapter, real-world noise is *not* white but instead increases at low frequencies. This phenomenon is called **1/f noise** and is typically modeled in ℓ -space as [70, 215]

$$N_\ell = N_\ell^{\text{white}} + N^{\text{red}} \left(\frac{\ell}{\ell_{\text{knee}}} \right)^{\alpha_{\text{knee}}} , \quad (3.71)$$

¹⁸Because the atmosphere is brighter closer to the horizon (see Equation 2.3), detector sensitivity is typically better at higher elevations. In addition, the ground signal depends, in general, on elevation, further complexifying how the detector data are filtered.

¹⁹When observing from anywhere but the poles, sky patches “rotate” throughout the day. As an explicit example, the constellation Centaurus spins about his own body as he arcs across the sky.




Figure 3.7: Forecasted N_ℓ **noise curves** for the SO LAT in temperature (Figure 3.7a) and the SO SAT in polarization (Figure 3.7b). The assumed NETs in these forecasts are calculated using BoloCalc, which is presented in Chapter 5, and the ℓ dependencies were quantified using Equations 3.70 and 3.71.

where the $1/\ell$ knee²⁰ is ℓ_{knee} , the $1/\ell$ spectral index is $\alpha_{\text{knee}} < 0$, and the low-frequency-noise amplitude is N^{red} , all of which are usually deduced empirically.

The figure of merit for any CMB power spectrum measurement is the noise in power spectrum space (Equation 3.69), and therefore CMB instrumentalists often quantify instrument sensitivity in terms of **mapping speed**

$$\text{MS} = \frac{1}{\text{NET}_{\text{arr}}^2} = \frac{N_{\text{det}} Y}{\text{NET}_{\text{det}}^2 \Gamma^2} \propto \frac{1}{N_\ell^{\text{white}}}. \quad (3.72)$$

Mapping speed quantifies the number of detector-hours needed to reach a specified CMB map depth and is proportional to detector yield observation efficiency. As a result, mapping speed is a powerful merit figure during experiment design and optimization. Quantification and mitigation of N^{red} and ℓ_{knee} are also important to CMB power spectrum measurements and are discussed in Chapter 6.

3.11 Discussion

There are two critical aspects to CMB instrument's performance: sensitivity and systematic effects. While both aspects are important to forecasting, sensitivity is often easier to quantify and understand, making it a powerful tool to inform telescope design, validation, and in-field characterization. In this chapter, we have overviewed how to calculate detected optical power P_{opt} , noise-equivalent temperature NET, and mapping speed MS, laying out many of

²⁰The $1/\ell$ knee is conceptually identical to the $1/f$ knee described in Section 6.1.

the details and assumptions that flow into the calculation. A primary goal of this dissertation is to apply the presented sensitivity formalism into a general, public sensitivity calculator, which we present in Chapter 5. In addition, as we mentioned in Section 3.9, another goal of this dissertation’s research is to investigate the impact of photon noise correlations on NET_{arr} , which is detailed in Chapter 4. Finally, we use mapping speed to set specifications and tolerances for various hardware components in Chapters 6-10. Therefore, the presented sensitivity discussion lays a foundation for research products throughout the remainder of this thesis.

Chapter 4

Photon noise correlations

This chapter overlaps substantially with Hill and Kusaka (exp. 2021) [78], which we encourage the reader to cite instead of or in addition to this dissertation.

Modern CMB telescopes are often limited by fluctuations in the background radiation, which are described in Sections 3.2 and 3.3. This statement is especially true for ground-based experiments, where emission from the atmosphere and telescope are substantial. At high frequencies (e.g., optical wavelengths), the mode occupation number (Equation 3.13) $\langle n \rangle \ll 1$ and photon fluctuations are dominated by **uncorrelated shot noise** such that $\Delta n \approx \sqrt{\langle n \rangle}$. At low frequencies (e.g., radio wavelengths), $\langle n \rangle \gg 1$ and photon fluctuations are dominated by **correlated wave noise** such that $\Delta n \approx \langle n \rangle$. Millimeter wavelengths lie in the cross-over regime where $\langle n \rangle \sim 1$, as shown in Figure 3.6, making an accurate calculation of the array-averaged sensitivity nontrivial.

In addition, modern CMB telescopes are typically **field-of-view-limited** and therefore aim to increase detector count by increasing pixel density.¹ In the high-pixel-density paradigm, it is possible to **oversample the focal plane** such that neighboring detectors measure the same **spatial mode**. As we show in this chapter, mode sharing introduces **photon noise correlations** when the pixel spacing is less than the telescope’s Nyquist distance $D_{\text{pix}} < 1.2F\lambda$, and these correlations can have substantial impacts on the array-averaged noise-equivalent temperature (NET).

The theory of intensity correlations from incoherent sources has been studied extensively [54, 61, 62, 127, 29, 28], and the phenomenon was experimentally demonstrated by Hanbury Brown and Twiss via angular-diameter measurements of distant radio sources (1952) [23, 25, 24]. A circuit-based formalism for thermal photon correlations in quantum detectors has been developed by Zmuidzinas (2003) [241], and their impact on millimeter-wave detector arrays has been noted by Padin (2010) [149]. In this chapter, we apply the theory of intensity correlations to the calculation of photon noise averaging in Section 3.8, and we use this application to estimate the sensitivity of a model CMB telescope. We then

¹One can also deploy more detectors by building more telescopes, but increasing pixel density is usually more cost-efficient.

discuss the implications of the presented sensitivity optimization for the design of mm-wave detector arrays.

4.1 Theoretical foundations

In this section, we review the theory of photon-count statistics and reformulate the quantum optical calculation using classical fields. We first adopt the treatment of thermal photon correlations derived by Zmuidzinas (2003) [241], which uses the machinery of transmission lines and scattering matrices to calculate the propagation of quantum modes in a **linear optical system**.² We then write these circuit-based mode operators in a basis of free-space field operators, the statistics of which can be found using the mixed-state thermal density matrix, and we invoke the optical equivalence theorem [194] to equate the statistics of the quantum and classical fields. Finally, we use these classical fields to calculate the the Hanbury Brown-Twiss coefficient, which we then show can be used to calculate photon correlations.

A more rigorous, compact theoretical treatment can be found in Hill and Kusaka (exp. 2021) [78], but in this chapter, we somewhat simplify the calculation to provide an alternative, more intuition-focused supplement to the published formalism.

4.1.1 Photon correlations in a circuit

Following the prescription laid out by Zmuidzinas (2003) [241], consider a linear, lossy network of $k = 1, \dots, N$ input ports³ detected at an output port i .⁴ Input modes enter the network along N semi-infinite transmission lines via the photon creation operator $a_k^\dagger(\nu)$ and are mapped onto outputs via a **scattering matrix** S_{ik} . Loss in the system is modeled by an orthogonal scattering matrix S'_{ik} , which governs the noise injected between input mode k and output mode i . Given this structure, the creation operator $b_i^\dagger(\nu)$ at output port i and mode frequency ν is

$$b_i^\dagger(\nu) = \sum_k S_{ik}(\nu) a_k^\dagger(\nu) + \sum_{k'} S'_{ik'}(\nu) a_{k'}^\dagger(\nu). \quad (4.1)$$

When the loss during mode propagation is uncorrelated, $S'_{ik} = c_i \delta_{ik}$, where c_i is the noise amplitude. As shown in Figure 4.1, there is no fundamental distinction between the in-

²The assumption of linearity mandates that optical elements do not shift or mix frequencies and that electric fields are describable using the superposition principle. These assumptions hold true for the lenses, waveplates, and polarizers found in CMB telescopes.

³Throughout this chapter, values and variables will be *italicized*, while indexes will not. For example, k refers to a wavenumber, while k is an integer.

⁴While we will relate the input ports to the optical input and the output ports to detectors, there is no fundamental difference between the inputs and outputs, and every port has both incoming and outgoing photons.




Figure 4.1: A schematic of the quantum circuit formalism. The creation operator for incoming modes is a^\dagger , while that of the outgoing modes is b^\dagger . The scattering matrix S maps the input modes onto the output modes, while the noise matrix S' calculates noise and loss within the system. S and S' are assumed to be orthogonal.

put/output ports and the lossy ports, and therefore for simplicity, we hereafter fold the lossy scattering matrix S'_{ik} into S_{ik} and treat them through a single unified scattering matrix.

The two-photon expectation value at detector outputs i and j is given by

$$\langle b_i^\dagger(\nu) b_j(\nu') \rangle = \sum_k \sum_m S_{ik}^*(\nu) S_{jm}(\nu') \langle a_k^\dagger(\nu) a_m(\nu') \rangle. \quad (4.2)$$

Here, the expectation values $\langle \dots \rangle$ are taken over quantum-statistical mixed states, which are governed by the density matrix, and represent the **quantum coherence** of the photon modes. See Appendix A.1 for further discussion regarding the thermal photon density matrix.

When the mixed states are in thermal equilibrium, which is a good approximation for the photon sources in the calculations that follow, $\langle a_k^\dagger(\nu) a_m(\nu') \rangle = n(T_k, \nu) \delta_{km} \delta(\nu - \nu')$, where T_k is the temperature of port k , and $n(T, \nu) \equiv (e^{h\nu/k_B T} - 1)^{-1}$ is the mean occupation number of a blackbody at temperature T for frequency ν [241]. We can write the two-photon output expectation value as

$$\langle b_i^\dagger(\nu) b_j(\nu') \rangle \equiv B_{ij}(\nu) \delta(\nu - \nu'), \quad (4.3)$$

where $B_{ij}(\nu)$ is the **quantum mutual intensity** and satisfies

$$B_{ij}(\nu) = \sum_k S_{ik}^*(\nu) S_{jk}(\nu) n(T_k, \nu). \quad (4.4)$$

Thermal detectors for CMB detection integrate photon count over time τ (see Equation 2.17) and sense mean intensity

$$\langle d_i \rangle = \frac{1}{\tau} \int_0^\tau dt \langle b_i^\dagger(t) b_i(t) \rangle \simeq \int_{\nu_1}^{\nu_2} d\nu h\nu B_{ii}(\nu), \quad (4.5)$$

where we define the time-dependent operators as

$$b_i(t) \equiv \int_{\nu_1}^{\nu_2} d\nu \exp[2\pi i \nu t] b_i(\nu) \sqrt{h\nu}, \quad (4.6)$$

$$b_i^\dagger(t) \equiv \int_{\nu_1}^{\nu_2} d\nu \exp[-2\pi i \nu t] b_i^\dagger(\nu) \sqrt{h\nu}. \quad (4.7)$$

Here, the integration limits are set by the detection bandwidth $\Delta\nu = \nu_2 - \nu_1$, and the factors of $\sqrt{h\nu}$ arise due to power detection as opposed to photon counting. The second equality in Equation 4.5 is a good approximation when $\tau \gg 1/\Delta\nu$, which is often true in CMB detectors for which $\tau \sim \mathcal{O}(10^{-2}-10^{-3} \text{ s})$ and $1/\Delta\nu \sim \mathcal{O}(10^{-10} \text{ s})$.

The covariance for quantum thermal detectors $\langle \Delta d_i \Delta d_j \rangle = \langle d_i d_j \rangle - \langle d_i \rangle \langle d_j \rangle$ can be written as

$$\langle \Delta d_i \Delta d_j \rangle \simeq \frac{1}{\tau} \int_{\nu_1}^{\nu_2} d\nu (h\nu)^2 [B_{ij}(\nu)\delta_{ij} + |B_{ij}(\nu)|^2] , \quad (4.8)$$

as shown in Zmuidzinas (2003) [241]. The first term in the integrand of Equation 4.8 represents uncorrelated shot noise, while the second term represents wave noise, which can correlate between output ports.⁵ Finally, the **intensity correlation** is defined as

$$\gamma_{ij}^2 = \frac{\langle \Delta d_i \Delta d_j \rangle \tau}{\langle d_i \rangle \langle d_j \rangle} \simeq \frac{|B_{ij}(\bar{\nu})|^2}{B_{ii}(\bar{\nu})B_{jj}(\bar{\nu})} \quad (4.9)$$

where $\bar{\nu} \simeq (\nu_1 + \nu_2)/2$ is the mean frequency, and the second equality is a good approximation when the fractional bandwidth is small enough that variations of the integrated function in Equation 4.8 are negligible.

4.1.2 Photon correlations in free space

The circuit model presented in Sec. 4.1.1 requires knowledge of $S(\nu)$ and $S'(\nu)$, and while scattering matrices fully describe a linear optics system, we seek to formulate Equation 4.8 within the framework of **classical fields**, which apply more directly to CMB telescope design. Therefore, we relate the photon creation operators a_i^\dagger to the free-space electric field operator $\hat{\mathbf{E}}^{(+)}(\vec{r}, t)$, and we show that the statistics of the quantum field are described by those of the classical field.

According to Glauber (1963) [61], we can define the electric field operator's positive and negative frequency components in the basis of photon creation operators $a_{\vec{k}p}^\dagger$ for the spatial mode defined by wave vector \vec{k} and polarization p as

$$\hat{\mathbf{E}}^{(+)}(\vec{r}, t) = [\hat{\mathbf{E}}^{(-)}(\vec{r}, t)]^\dagger = \sum_{\vec{k}} \sum_{p=1,2} i \sqrt{\frac{h\nu_k}{2\epsilon_0 v}} \hat{\epsilon}_{\vec{k}p} a_{\vec{k}p}^\dagger e^{i(\vec{k}\vec{r} - 2\pi\nu_k t)} , \quad (4.10)$$

where ϵ_0 and v are the vacuum permittivity and the volume of the space of interest, $a_{\vec{k}p}^\dagger$ is the creation operator for wave number \vec{k} and polarization state p, $\nu_k \equiv (c/2\pi)|\vec{k}|$ is the frequency, and $\hat{\epsilon}_{\vec{k}p}$ is the polarization vector. In this chapter, we focus on single-moded, single-polarization detectors, which most modern CMB imagers employ. In contrast to the vector field $\hat{\mathbf{E}}^{(+)}$, which has two degrees of freedom for propagating waves, a single-moded

⁵The second term in Equation 4.8 $|B_{ij}(\nu)|^2$ is often referred to as the **bunching term**, as it quantifies the degree to which photon arrival times are correlated or are “bunched.”

detector is sensitive to a single degree of freedom of the electric field, which can be written as

$$\hat{E}_i^{(+)}(t) = [\hat{E}_i^{(-)}(t)]^\dagger = \sum_{\vec{k}} \sum_{p=1,2} i \sqrt{\frac{h\nu_k}{2\epsilon_0 v}} g(\nu_k) G_p(\theta_k, \phi_k) a_{\vec{k}p}^\dagger e^{i(\vec{k}\vec{r}_i - 2\pi\nu_k t)}, \quad (4.11)$$

where the index i labels a detector with position \vec{r}_i and where $g(\nu_k)$ and $G_p(\theta_k, \phi_k)$ are the detector's frequency and polarization-dependent angular responses, respectively. The angles θ_k and ϕ_k are defined such that $\theta_k \equiv \arccos(k_z/|\vec{k}|)$ and $\phi_k \equiv \arctan(k_y/k_x)$. As discussed later in this section and in Appendix A, according to the optical equivalence theorem, the response functions $g(\nu_k)$ and $G_p(\theta_k, \phi_k)$ defined in the context of the quantum-operator formalism are equal to those of the classical electromagnetic field.

We can rewrite Equation 4.11 by decomposing the summation over \vec{k} into angular and frequency components⁶ and by taking a simple bandpass function of $g(\nu_k)$ that is unity for $\nu_1 \leq \nu \leq \nu_2$ and zero otherwise

$$\hat{E}_i^{(+)}(t) = [\hat{E}_i^{(-)}(t)]^\dagger = \int_{\nu_1}^{\nu_2} i \sqrt{\frac{h\nu_k}{2\epsilon_0 v}} d\nu_k \iint d\Omega_k \sum_{p=1,2} G_p(\theta_k, \phi_k) a_{\vec{k}p}^\dagger e^{i(\vec{k}\vec{r}_i - 2\pi\nu_k t)}, \quad (4.12)$$

with

$$\iint d\Omega_k \equiv \int d(\cos \theta_k) \int d\phi_k. \quad (4.13)$$

In other words, the mapping from Equations 4.1 and 4.7 to Equation 4.12 is

$$a_k^\dagger(\nu) \rightarrow a_{\vec{k}p}^\dagger, \quad S_{ik} \rightarrow \mathcal{C} G_p(\theta_k, \phi_k) e^{i\vec{k}\vec{r}_i}, \quad \sum_k \rightarrow \iint d\Omega_k \sum_{p=1,2}, \quad b_i^\dagger(t) \rightarrow \hat{E}_i^{(+)}(t), \quad (4.14)$$

where \mathcal{C} is an overall normalization constant. Equation 4.12 is consistent with a correspondence between Equation 4.5 and the fact that power detectors measure $\langle \hat{E}_i^{(+)}(t) \hat{E}_i^{(-)}(t) \rangle$.

The quantum mutual intensity for this system can be calculated as

$$B_{ij}(\nu) = |\mathcal{C}|^2 \iint d\Omega_k \sum_{p=1,2} |G_p(\theta_k, \phi_k)|^2 e^{i\vec{k}(\vec{r}_j - \vec{r}_i)} n(T_{\vec{k}p}, \nu), \quad (4.15)$$

where $T_{\vec{k}p}$ is the temperature of a mode with wave vector \vec{k} , polarization p , and creation operator $a_{\vec{k}p}^\dagger$. In general, source modes do not always correspond to plane waves, as some

⁶We omit a factor $\propto |\vec{k}|^2$ that shows up in this decomposition and let it be absorbed in the overall normalization. It sometimes cancels with the frequency dependence of $G_p(\theta, \phi)$, whose integral, for a single-moded detector, often scales with the inverse of $|\vec{k}|^2 \propto \nu_k^2$. A more specifically engineered detector with a wide-band response may have a less trivial scaling. These details do not affect the first-order approximation of the intensity correlation.

sources may not be in the telescope's far field.⁷ Equation 4.12 can be rewritten for these cases by using a general scattering matrix $S_{i,kp}$

$$\hat{E}_i^{(+)}(t) = [\hat{E}_i^{(-)}(t)]^\dagger = \mathcal{C}' \int_{\nu_1}^{\nu_2} \sqrt{h\nu} d\nu \sum_k \sum_{p=1,2} S_{i,kp} e^{-2\pi i \nu t} a_{kp}^\dagger, \quad (4.16)$$

where \mathcal{C}' is an overall normalization constant, k runs over the spatial modes with frequency ν , and p indexes polarization.

For the purposes of this dissertation, we simplify the calculation of the generalized quantum mutual intensity in Equation 4.16 by adopting the same approximation used for Equation 4.9. If we assume that variations as a function of ν within the detection band $g(\nu)$ are small, then

$$\langle \hat{E}_i^{(+)}(t) \hat{E}_j^{(-)}(t) \rangle \simeq |\mathcal{C}'|^2 \Delta\nu h\bar{\nu} \sum_{kp} S_{i,kp}^*(\bar{\nu}) S_{j,kp}(\bar{\nu}) n(T_{kp}, \bar{\nu}) \quad (4.17)$$

$$= |\mathcal{C}'|^2 \Delta\nu h\bar{\nu} B_{ij}(\bar{\nu}), \quad (4.18)$$

where $\Delta\nu \equiv \nu_2 - \nu_1$ and T_{kp} is the temperature of the mode specified by k and polarization p .

The **optical equivalence theorem**, which is described in Appendix A, allows us to calculate the **quantum expectation value** in Equations 4.17 and 4.18 using the classical field as

$$\langle \hat{E}_i^{(+)}(t) \hat{E}_j^{(-)}(t) \rangle = \langle E_i(t) E_j^*(t) \rangle_c, \quad (4.19)$$

where $\langle \cdots \rangle_c$ denotes the **classical expectation value** over the classical complex electric field $E_i(t)$, and where the electric-field amplitude of source mode kp follows a complex Gaussian random distribution with variance $n(T_{kp}, \bar{\nu})$. The combination of Equations 4.9, 4.18, and 4.19 allows us to evaluate the intensity correlation between *quantum* detectors using the first-order coherence of the *classical* field, with the scattering matrix $S_{i,kp}$ obtained through classical wave propagation. Equations 4.17 and 4.19 are a representation of the **van Cittert-Zernike theorem (VCZT)** [35, 238], where $\sum_{kp} \cdots n(T_{kp}, \bar{\nu})$ corresponds to the integral over the intensity profile of far-field sources.

The **classical intensity covariance** can be calculated as (see Appendix B)

$$V_{ij} \equiv \langle \Delta I_i(t) \cdot \Delta I_j(t) \rangle_c = \frac{|\mathcal{C}'|^2}{\tau} \int_{\nu_1}^{\nu_2} d\nu (h\nu)^2 |B_{ij}(\nu)|^2 \simeq \frac{|\mathcal{C}'|^2}{\tau} \Delta\nu (h\bar{\nu})^2 |B_{ij}(\bar{\nu})|^2, \quad (4.20)$$

with $I_i(t) \equiv |E_i(t)|^2$ and $\Delta I_i(t) \equiv I_i(t) - \langle I_i(t) \rangle_c$. In contrast to Equation 4.8, this classical treatment lacks the shot noise term, which is a purely quantum mechanical phenomenon.⁸

⁷As an example, thermal emission from the telescope's own optical elements is not in the optical system's far field and therefore needs to be treated more generally.

⁸"Shot noise" is due to the uncorrelated arrival times of photons and therefore is due to light's packet-like nature, while "wave noise" is due to light's wave-like nature.

Here, the equivalence between photon bunching $|B_{ij}(\bar{\nu})|^2$ and intensity correlations V_{ij} is referred to as the **Hanbury Brown-Twiss (HBT) effect** and has been well researched over the past several decades [23, 25, 24, 54, 61, 62, 127, 140, 29, 28]. Using Equation 4.20, the *quantum* intensity correlation γ_{ij}^2 can also be calculated as a correlation coefficient of *classical* intensity fluctuations

$$\gamma_{ij}^2 \simeq \frac{\langle \Delta I_i \cdot \Delta I_j \rangle_c \tau}{\langle I_i \rangle_c \langle I_j \rangle_c} = \frac{\langle \Delta I_i \cdot \Delta I_j \rangle_c}{\sqrt{\langle \Delta I_i^2 \rangle_c \langle \Delta I_j^2 \rangle_c}} = \frac{V_{ij}}{\sqrt{V_{ii} V_{jj}}}, \quad (4.21)$$

where the second equality holds only for the classical expectation values. As described in Appendix B, the equivalence between the quantum and classical intensity correlations shown in Equation 4.21 arises from the **Gaussian distribution** of *classical* amplitudes due to the Bose-Einstein distribution of *quantum* amplitudes. Thus, what we call “classical” here does have a quantum aspect to it. For the remainder of this chapter, we refer to γ_{ij} as the **HBT correlation coefficient**.

4.1.3 Photon noise correlations

Using the approximation in Equation 4.9, we can decompose the photon-count covariance in Equation 4.8 into shot- and wave-noise terms using the HBT coefficient $\gamma_{ij}(\bar{\nu})$

$$\begin{aligned} \sigma_{ij}^2 &\equiv \langle \Delta d_i \Delta d_j \rangle \\ &= \frac{1}{\tau} \int_{\nu_1}^{\nu_2} d\nu (h\nu)^2 B_{ij}(\nu) \delta_{ij} + \frac{1}{\tau} \int_{\nu_1}^{\nu_2} d\nu (h\nu)^2 |\gamma_{ij}(\bar{\nu})|^2 B_{ii}(\nu) B_{jj}(\nu) \\ &= \sigma_{ij,\text{shot}}^2 + \sigma_{ij,\text{wave}}^2. \end{aligned} \quad (4.22)$$

Equation 4.22 can be rewritten in terms of the mode occupation number $n(T, \nu)$, where again T is the mode temperature, and the covariance between output ports $\{i, j\}$ can be written as

$$\begin{aligned} \sigma_{ij,\text{shot}}^2 &= \frac{1}{\tau} \int_{\nu_1}^{\nu_2} d\nu (h\nu)^2 \eta_i(\nu) n(T_i, \nu) \delta_{ij} \\ \sigma_{ij,\text{wave}}^2 &= \frac{1}{\tau} \int_{\nu_1}^{\nu_2} d\nu (h\nu)^2 |\gamma_{ij}(\bar{\nu})|^2 \eta_i(\nu) \eta_j(\nu) n(T_i, \nu) n(T_j, \nu), \end{aligned} \quad (4.23)$$

where $\eta_i(\nu)$ is the quantum efficiency at output i . As shown in Equation 4.23, the problem of calculating wave noise correlations is reduced to finding the HBT coefficient $|\gamma_{ij}(\bar{\nu})|^2$, the mode temperature T , and the detector’s **quantum efficiency** $\eta(\nu)$.

4.2 Model optical system

Using the photon noise formulation in Equation 4.23, we now quantify the impact of HBT correlations on the sensitivity of CMB instruments. The topic of intensity correlations from




Figure 4.2: A simplified optical model used for all calculations and simulations in this chapter. The model assumes an objective lens, aperture stop, and focal plane filled with an array of sensing antennas coupled to planar detectors (see Section 2.3.3). In the reverse-time sense (see Figure 3.3), each pixel emits a collection of rays defined by spherical coordinate $(\theta_{\text{pix}}, \phi_{\text{pix}})$, which the objective lens uniquely maps onto an aperture-plane coordinate \vec{r}_{ap} with idealized polarization fidelity and telecentricity. The system is enclosed in a black box of temperature T_{stop} , and the focal plane is cooled to T_{fp} .

astronomical sources has been discussed extensively in the literature [9, 29, 28, 227, 226, 228, 3], and in the following sections, we apply these findings to mm-wave telescope design. As overviewed in Chapter 2, modern CMB telescopes employ a wide variety of lens and mirror systems, infrared filter stacks, anti-reflection coatings, and sensing architectures. Despite this variety, we distill a few key instrument characteristics to create a simple yet representative optical system in which to study HBT correlations.

4.2.1 Telescope

The assumed telescope model is depicted in Figure 4.2. It consists of an objective lens with focal length f_{obj} within a blackened enclosure at $T_{\text{stop}} = 4 \text{ K}$. The cold box has a circular aperture of diameter D_{ap} at $z = z_{\text{ap}}$ that truncates incoming radiation from external, field-of-view-filling sources. The objective lens, which is both cold and transparent, focuses the aperture-truncated radiation onto a circular focal plane at $z = z_{\text{fp}}$, the size of which is determined by the telescope's plate scale. The focal plane houses an array of close-packed detector pixels with diameter D_{pix} operating at $T_{\text{fp}} = 0.1 \text{ K}$, and the optical system is

assumed to be diffraction-limited such that the throughput (or etendue) per mode is

$$A\Omega = \lambda^2. \quad (4.24)$$

This model optical system excludes many common features of real telescopes—such as fore-optics, thermal filters, or reimaging lenses—which are needed to form high-fidelity images over a moderate field of view (FOV). Such details are experiment dependent and are therefore beyond the scope of this chapter, but we can capture their effects by imposing several assumptions onto our simple system. These assumptions are not strictly necessary to calculate photon noise correlations, as the scattering matrix formalism in Section 4.1.1 is completely general, but they simplify the correlation calculation significantly while encapsulating the salient features of practical instruments.

Firstly, we assume that all sources—both external and internal to the telescope—are isothermal blackbody emitters large enough to **uniformly illuminate** the aperture across the telescope’s FOV. The assumption of blackbodies allows us to readily evaluate each mode’s occupation number using the Bose-Einstein distribution $n(\nu, T)$, provided each source’s effective brightness temperature T . The assumption that the blackbodies are **FOV-filling** and isothermal generalizes the correlation integrals to follow and is a good approximation for experiments that map extended sources, such as CMB telescopes. In practice, thermal gradients develop across optical elements and atmospheric brightness varies with elevation and cloud structure; however, these variations are typically small and experiment-dependent.

Secondly, we assume a **diffraction-limited, single-moded optical system** that converts stop-truncated plane waves into **spherical waves** converging onto a telecentric focal plane. In other words, any pixel rays with angle $(\theta_{\text{pix}}, \phi_{\text{pix}})$ are mapped via the objective lens onto parallel rays with aperture-plane location \vec{r}_a (see Figure 4.2). In this configuration, the **optical path length** between a detector pixel and a spot on the aperture stop is identical regardless of $(\theta_{\text{pix}}, \phi_{\text{pix}})$ or \vec{r}_a , which simplifies the calculations to follow. We note that when re-imaging optics and a pupil stop (such as Lyot stop) are employed, modes which map onto such spherical waves do not generally correspond to plane waves passing through the aperture. However, such reimaging optics can always be modeled as a simplified equivalent system with a simple aperture stop, provided that their clear-aperture diameters are large enough to pass all modes not truncated by the pupil stop.

Thirdly, we assume an **ideal aperture stop**, such that all detector pixels have the same mapping between ray angle $(\theta_{\text{pix}}, \phi_{\text{pix}})$ and aperture-plane location \vec{r}_a (see Figure 4.2). In other words, each detector pixel illuminates the aperture stop identically regardless of its location on the focal plane.⁹ Practically, this condition is not generally satisfied for a system with a large field of view, as telecentricity and aperture truncation may differ significantly between the central and peripheral regions of the focal plane. However, as we discuss later, photon noise correlations arise predominantly between neighboring pixels where such non-idealities are negligible.

⁹Strictly speaking, the necessary condition for an ideal aperture is for the mapping to be identical only on the stop’s circumference, but when this condition is met, the mapping becomes identical within the aperture stop as well, in practice.

Fourthly, we assume polarization fidelity across the focal plane, such that the objective lens converts a linearly polarized plane wave into a spherical wave with a Ludwig-3 polarization distribution [124]

$$\begin{aligned}\hat{x} &\rightarrow \cos \phi_{\text{pix}} \hat{\theta} - \sin \phi_{\text{pix}} \hat{\phi} \\ \hat{y} &\rightarrow \sin \phi_{\text{pix}} \hat{\theta} + \cos \phi_{\text{pix}} \hat{\phi}.\end{aligned}\tag{4.25}$$

It follows from this assumption that correlations in the following calculations cannot develop between orthogonal polarimeters. In practice, some cross polarization develops within real telescopes, and the degree of polarization mixing may differ between the central and peripheral regions of the focal plane. However, polarization-sensitive optical systems are specifically designed for low cross polarization [138, 47], especially over localized areas of the focal plane for which photon noise correlations are important.

4.2.2 Focal plane

The assumed focal plane model is shown in Figure 4.3, and the focal plane coupling architecture is similar to that described in Section 2.3.3. We assume single-moded, dual-polarization detector pixels with circular apertures and **diffraction-limited Gaussian beams**. Each pixel's angular response function is determined solely by its beam waist w_0 and takes the far-field form (see Equation 2.6)

$$E(\theta) \approx E_0 \exp \left[-\frac{\theta^2}{(\lambda/\pi w_0)^2} \right].\tag{4.26}$$

Each pixel has a diffraction-limited throughput of $A\Omega = \lambda^2$, regardless of its aperture size, and we assume that the beam pattern in Equation 4.26 is symmetric between the pixel's E and H planes. A larger (smaller) pixel will result in a smaller (larger) far-field beam pattern, and we linearly relate the pixel diameter D_{pix} to the beam waist w_0 via a scaling constant w_f

$$w_0 = \frac{D_{\text{pix}}}{w_f}.\tag{4.27}$$

Typical mm-wave detector pixels, such as corrugated feedhorns [145], spline-profiled feedhorns [183], and lenslet-coupled planar antennas [185], achieve $w_f \approx 3$, which we assume for the calculations that follow.

Plugging Equation 4.27 into Equation 4.26 yields a simple relationship between D_{pix} and aperture stop spillover efficiency (see Section 2.3.3)

$$\eta_{\text{apert}} = \frac{\int_0^{\theta_{\text{stop}}} E^2(\theta) d\theta}{\int_0^{\pi/2} E^2(\theta) d\theta} = 1 - \exp \left[\frac{\pi^2}{2} \left(\frac{D_{\text{pix}}}{F\lambda w_f} \right)^2 \right],\tag{4.28}$$

where $F \equiv f_{\text{obj}}/D_{\text{ap}}$ is the f-number at the focal plane and $\theta_{\text{stop}} = \arctan [1/(2F)]$. This Gaussian representation for the spillover efficiency does not hold when $D_{\text{pix}} < \lambda$, as diffraction




Figure 4.3: The assumed layout of detector pixels on the focal plane, shown with $D_{\text{pix}} = 1.2F\lambda$ spacing. Each pixel’s angular response is Gaussian and scales with its diameter, as described in Equation 4.26. Additionally, each pixel has two polarimeters that sense orthogonal polarizations and whose noise outputs do not correlate in our idealized optical system.

at the pixel’s edges will create significant ringing in the far-field beam pattern. However, to remain agnostic to the specifics of the detector’s coupling architecture, we assume that Equations 4.26 and 4.28 remain valid for all values of D_{pix} in the following calculations.

While detector pixels can be arranged in a variety of ways, we choose one common focal plane arrangement to find $|\gamma(\vec{r}_1, \vec{r}_2)|^2$ explicitly. We assume **hex-packed circular pixels**, and we assume that pixel pitch is equal to pixel diameter $p_{\text{pix}} = D_{\text{pix}}$, as shown in Figure 4.3. This assumption allows us to relate pixel packing density to pixel size as $n_{\text{pix}} \propto D_{\text{pix}}^{-2}$, where n_{pix} is the number of pixels per unit focal plane area. In practice, dead space typically exists between pixels that does not scale with pixel size, yielding a more complex relationship between n_{pix} and D_{pix} [149]. However, such details are experiment specific and are therefore beyond the scope of the following discussions. In addition, we assume that each pixel contains two polarimeters and that pixels alternate between “+” and “×” orientations to extract Q and U Stokes polarization, respectively (see Figure 1.9).

4.3 Correlation calculation

We have shown in Section 4.1 that the HBT coefficient $|\gamma_{ij}|^2$ can be calculated using Equations 4.18 and 4.19 given the telescope’s classical scattering matrix $S_{i,\text{kp}}$. In Hill and Kusaka (est. 2021) [78], we describe the model optical system in Section 4.2 using such a scatter-




Figure 4.4: A schematic of the aperture, stop, and pixel radiation models used for the correlation simulation. Radiation within the aperture is decomposed into a basis of plane-waves as in Equation 4.34, and the normally incident mode $(k_x, k_y, k_z) = (0, 0, k)$ is depicted here. The plane waves are assumed to uniformly illuminate a diameter D_{illum} , and each mode's Gaussian-random amplitude is defined by its effective brightness temperature. In turn, the aperture has a diameter $D_{\text{ap}} < D_{\text{illum}}$ and acts as a high-pass filter, cutting off modes below $k_{\text{min}} = 2\pi/D_{\text{ap}}$. Radiation from the stop is generated by a collection of blackbody point sources, as in Equation 4.35, which emit spherical Lambertian wavelets. Both the aperture's plane-wave modes and the stop's wavelet modes couple to the objective lens, which focuses the thermal light onto the focal plane. A detector pixel then senses the incoming spherical wave with the Gaussian angular response in Equation 4.26.

ing matrix, allowing us to readily evaluate $|B_{ij}(\bar{\nu})|^2$. However, in this dissertation section, we instead aim to supplement that compact treatment with a more intuitive description of spatial correlations in our example telescope. A schematic of the radiation model for the central detector pixel is shown in Figure 4.4.

4.3.1 Aperture radiation

Because the model optical system presented in Section 4.2 has finite throughput, it is convenient to decompose incoming radiation from external sources into a basis of plane waves

$$\vec{\psi}_{\text{p}}(\mathbf{m}, \mathbf{n}; x, y, z) = \frac{\hat{\epsilon}_{\text{p}}(\mathbf{m}, \mathbf{n})}{D_{\text{illum}}} e^{i(k_x(\mathbf{m})x + k_y(\mathbf{n})y)} e^{ik_z(\mathbf{m}, \mathbf{n})z}, \quad (4.29)$$

where (m, n) are the mode numbers, D_{illum} is the dimensional extent of the aperture-incident illumination, $\psi_0^{\text{in}}(m, n)$ is the mode's complex amplitude, $\hat{\varepsilon}_p(m, n)$ is the polarization vector, $k_x(m) = 2\pi m/D_{\text{illum}}$, $k_y(n) = 2\pi n/D_{\text{illum}}$, $k_z(m, n) = \sqrt{k_{\text{in}}^2 - k_x^2(m) - k_y^2(n)}$, $k_{\text{in}} = 2\pi/\lambda$, and $\omega_k = ck_{\text{in}}$. The incoming and aperture-truncated electric fields, \vec{E}_p^{in} and \vec{E}_p^{ap} , can then be written as a superposition of these plane waves

$$\vec{E}_p^{\text{in}}(x, y, z) = \sum_{m, n} a_p(m, n) \vec{\psi}_p(m, n; x, y, z) \quad (4.30)$$

$$\vec{E}_q^{\text{ap}}(x, y, z) = \sum_{r, s} b_q(r, s) \vec{\psi}_q(r, s; x, y, z), \quad (4.31)$$

where $a_p(m, n)$ and $b_q(r, s)$ are the incoming and aperture-truncated mode amplitudes, respectively. As discussed in Section 4.1.1, input modes are mapped onto aperture-truncated modes as

$$b_q(r, s) = \sum_p \sum_{m, n} S_{pq}^{\text{ap}}(r, s; m, n) a_p(m, n), \quad (4.32)$$

where the scattering matrix is defined to be the overlap integral of the input and aperture-truncated modes over the aperture σ

$$\begin{aligned} S_{pq}^{\text{ap}}(r, s; m, n) &= \iint_{\sigma} \vec{\psi}_p(m, n; x, y, z) \cdot \vec{\psi}_q^*(r, s; x, y, z) \, dx \, dy \\ &= \hat{\varepsilon}_p(m, n) \cdot \hat{\varepsilon}_q(r, s) e^{i(k_z(m, n) - k_z(r, s))z} \\ &\times \frac{1}{D_{\text{illum}}^2} \iint_{\sigma} e^{i(k_x(m) - k_x(r))x} e^{i(k_y(n) - k_y(s))y} \, dx \, dy. \end{aligned} \quad (4.33)$$

Here, the aperture σ satisfies $\sqrt{x^2 + y^2} \leq D_{\text{ap}}/2$, and the effect of the scattering matrix is to induce **mode overlap**. The scattering matrix is readily evaluated for a *square* aperture [34], but for a circular aperture, the integral in Equation 4.33 is nontrivial. However, to leading order, the impact of the aperture stop is to simply limit the number of modes entering the telescope [225] so that the aperture-truncated field can be approximated as

$$\begin{aligned} \vec{E}_p^{\text{ap}}(u, v, z = z_{\text{ap}}) &\simeq \sum_q \sum_{m, n} \vec{\psi}_q^{\text{in}}(m, n; u, v, z = z_{\text{ap}}) \\ \text{where } \sqrt{m^2 + n^2} &\leq D_{\text{ap}}/\lambda, \end{aligned} \quad (4.34)$$

where (u, v) are the coordinates on the aperture plane. Stated differently, we assert that the aperture stop is large enough to introduce negligible correlations between input modes. This assumption is reasonable because the aperture is much larger than the wavelength $D_{\text{ap}} = 10 \sim 100\lambda$, or the length scale over which thermal light is correlated [127]. This assertion, along with that of uniform illumination noted in Section 4.2, substantially reduces the computation time of the simulations that follow while describing the aperture field with high fidelity [34].

4.3.2 Stop radiation

The stop is located on the aperture plane $z = z_{\text{ap}}$ and is effectively a black, annular source with temperature T_{stop} . Therefore, we model it as a collection of Gaussian-random, uncorrelated thermal emitters that generate **Lambertian spherical wavelets**, as shown in Figure 4.4. These point sources represent atomic thermal motions within the stop's absorbing material, and their wavelets superpose to form incoherent plane waves that the objective lens focuses onto the focal plane (see Figure 4.4). Explicitly, the electric field on the stop $(x, y, z) = (u, v, z_{\text{ap}})$ is written as

$$\vec{E}_p^{\text{st}}(u, v, z_{\text{ap}}) = \psi_0^{\text{st}}(u, v) \hat{e}_p(u, v) \quad (4.35)$$

when $\sqrt{u^2 + v^2} > D_{\text{ap}}/2$, and $\vec{E}_p^{\text{st}}(u, v, z) = 0$ otherwise. Here, the stop's mode amplitude $\psi_0^{\text{st}}(u, v)$ is analogous to the aperture's ψ_0^{ap} in Equation 4.29.

4.3.3 Correlation patterns

Given the electric field both within the aperture (Equation 4.34) and on the stop (Equation 4.35), we now simulate intensity fluctuations and calculate the HBT coefficient at the focal plane. Due to their thermal nature, the amplitude and phase for all modes are drawn from a **normal probability distribution** [136]

$$p[\text{Re}(\psi_0)] = p[\text{Im}(\psi_0)] = \frac{1}{\sqrt{2\pi\langle I \rangle}} \exp\left[\frac{-|\psi_0|^2}{2\langle I \rangle}\right] \quad (4.36)$$

for all (m, n) , where each mode's mean intensity is determined by its Bose-Einstein occupation number

$$\langle I \rangle = h\nu n(T, \nu). \quad (4.37)$$

Once all modes are simulated, they are superposed to form the aperture and stop fields \vec{E}_p^{ap} and \vec{E}_p^{st} in Equations 4.34 and 4.35, respectively. As described in Section 4.2, the objective lens is located directly behind the aperture plane and converts impinging plane waves into diffraction-limited spherical waves focused onto the focal plane. Given this setup, the electric field at focal-plane location $(x, y, z) = (\vec{r}, z_{\text{fp}})$ due to radiation from the aperture/stop is found via the surface integral

$$\vec{e}_p^{\text{ap/st}}(\vec{r}, \nu) = \iint_{\sigma_{\text{ap}}/\sigma_{\text{st}}} \vec{E}_p^{\text{ap/st}}(u, v, z_{\text{ap}}) b(u, v, z_{\text{ap}}) e^{2\pi i \vec{r} \cdot \vec{u} \nu / c} du dv, \quad (4.38)$$

where integration is either over the aperture-illuminated surface σ_{ap} or the stop-illuminated surface σ_{st} of the objective lens, $b(u, v, z_{\text{ap}})$ is the detector's far-field response in Equation 4.26, and the distance between the aperture-field point and the detection point is




Figure 4.5: Simulated correlation at the focal plane due to radiation within the aperture (left panel) and from the stop (right panel). The blue dots show the results of the Monte Carlo simulation, and the red dotted lines show the expectation of the van Cittert-Zernike theorem for an aperture/stop with uniform intensity. Small deviations between the data points and the theory curve are due to a limited number of input modes and field realizations, which are computationally expensive to simulate.

$r = \sqrt{(x - u)^2 + (y - v)^2 + (z_{\text{fp}} - z_{\text{ap}})^2}$. Given the focal plane fields in Equation 4.38, we calculate the HBT correlation coefficient between focal plane points \vec{r}_1 and \vec{r}_2 as

$$|\gamma_{pq}(\vec{r}_1, \vec{r}_2, \nu)|^2 = \frac{\langle |e_p(\vec{r}_1, \nu)|^2 |e_q(\vec{r}_2, \nu)|^2 \rangle - \langle |e_p(\vec{r}_1, \nu)|^2 \rangle \langle |e_q(\vec{r}_2, \nu)|^2 \rangle}{\sqrt{\left[\langle (|e_p(\vec{r}_1, \nu)|^2)^2 \rangle - \langle |e_p(\vec{r}_1, \nu)|^2 \rangle^2 \right] \left[\langle (|e_q(\vec{r}_2, \nu)|^2)^2 \rangle - \langle |e_q(\vec{r}_2, \nu)|^2 \rangle^2 \right]}}, \quad (4.39)$$

where we have noted that $\vec{e}_p(\vec{r}, \nu) \cdot \vec{e}_q^*(\vec{r}, \nu) = |e_p(\vec{r}, \nu)|^2 \delta_{pq}$.

To find the HBT coefficient due to radiation within the aperture $|\gamma_{pq}^{\text{apert}}(\vec{r}_1, \vec{r}_2, \nu)|^2$ and from the stop $|\gamma_{pq}^{\text{stop}}(\vec{r}_1, \vec{r}_2, \nu)|^2$, we simulate 1,000 realizations of the aperture and stop fields with the model telescope described in Section 4.2. We set the aperture diameter to $D_{\text{ap}} = 100\lambda$, the stop's outer diameter to $2D_{\text{ap}}$, and we integrate over a grid with cell size $\Delta\vec{r}_{\text{ap}} = \lambda$. For radiation within the aperture, we simulate 2,500 total modes $(m_{\text{max}}, n_{\text{max}}) = (50, 50)$, which is enough to both resolve the stop's edges and Nyquist sample the integration grid. Figure 4.5 shows the result of the calculation plotted in units of $F\lambda$.

The HBT correlation pattern due to radiation within the aperture resembles an **Airy disk**, which is the Fourier transform of a circular aperture, while that of the stop is tighter

with larger side lobes. These patterns are consistent with the van Cittert-Zernike theorem (VCZT) [35, 238], which states that the correlation pattern of a far-field, incoherent source is the Fourier transform of the source's intensity profile. Written explicitly, the VCZT amplitude correlation (see Appendix B) is

$$\Gamma_{\text{VCZT}}^{(1,1)}(\vec{r}_1, \vec{r}_2, \nu) = \frac{1}{\sqrt{I(\vec{r}_1, \nu)}} \frac{1}{\sqrt{I(\vec{r}_2, \nu)}} \iint_{\sigma} I(u, v, \nu) \frac{e^{i\bar{k}(R_1 - R_2)}}{R_1 R_2} du dv, \quad (4.40)$$

where σ is $\sigma_{\text{ap}}/\sigma_{\text{st}}$, $\vec{r}_i = x_i\hat{x} + y_i\hat{y}$, $R_i = \sqrt{(u - x_i)^2 + (v - y_i)^2 + (z_{\text{ap}} - z_{\text{fp}})^2}$, and $\bar{k} = 2\pi/\bar{\lambda}$.

In its purest form, the VCZT applies to sources that are fully incoherent and in the far field of the detection plane. However, the aperture is *not* in the **Fraunhofer region** $z_{\text{fp}} - z_{\text{ap}} \not\gg D_{\text{ap}}^2/\lambda$ in most mm-wave cameras, and blackbody radiation is not totally incoherent, correlating over the distance of a wavelength [135, 29, 9, 191]. Therefore, it is not immediately obvious that the VCZT should predict the correlation pattern in our model telescope. However, various studies of the VCZT and its applications show that Equation 4.40 applies to **quasihomogeneous sources**, whose spatial intensity variations are slow compared to its coherence length [227, 226, 28, 57]. This classification describes not only our idealized uniform blackbodies but also more realistic thermal sources. Additional studies have shown that the VCZT also applies to sources nearer than the Fraunhofer limit [115, 3, 63, 30], and Gori (2005) [63] showed that its **far-field criterion** is

$$z_{\text{ap}} \gg \frac{D_{\text{ap}} R_c}{\lambda_{\text{obs}}}, \quad (4.41)$$

where D_{ap} is the diameter of the source, R_c is the radiation's correlation length, and λ_{obs} is the observed wavelength, which for CMB experiments is ~ 1 mm.

Mehta and Wolf (1964) [135] showed that blackbody radiation correlates over distances $R_c \approx \hbar c/k_B T$, where T is the radiation's blackbody temperature. In the presence of a $T_{\text{atm}} = 10$ K atmosphere, $R_c \sim \lambda_{\text{obs}}/10$ and the VCZT far-field criterion $z_{\text{ap}} \gg R_{\text{ap}}/10$ is met for all practical cases. For balloon and satellite missions, which observe lower sky intensities, or for experiments with a stop temperature of $T_{\text{stop}} \sim 1$ K, the limit in Equation 4.41 may become marginal, especially at high observation frequencies. However, as we will soon show, the impact of photon correlations is most prominent when observing brighter sources at lower frequencies, making Equation 4.40 a good approximation for most CMB experiments. Given its robust application to thermal light, we use the VCZT to calculate the HBT correlation coefficient for the remainder of this chapter

$$|\gamma_{\text{pq}}(\vec{r}_1, \vec{r}_2, \nu)|^2 = \frac{|\Gamma_{\text{VCZT}}^{(1,1)}(\vec{r}_1, \vec{r}_2, \nu)|^2}{\Gamma_{\text{VCZT}}^{(1,1)}(\vec{r}_1, \vec{r}_1, \nu) \Gamma_{\text{VCZT}}^{(1,1)}(\vec{r}_2, \vec{r}_2, \nu)} \delta_{\text{pq}}, \quad (4.42)$$

where we have noted that orthogonal polarizations cannot correlate [201] given our idealized aperture stop, objective lens, and detector pixels.

4.4 Impact on sensitivity

Using the VCZT-formulated HBT coefficient in Equation 4.42, we are now positioned to investigate the impact of correlations on instrument sensitivity, which is the primary research topic of this chapter. As shown in Figure 4.5, the outputs of detector pixels can correlate if their pitch is $D_{\text{pix}} \lesssim 1.2F\lambda$, and these correlations will slow noise averaging when outputs are coadded. In this section, we introduce a generalized formalism for mapping speed, which measures the total sensitivity of the detector array, and we inspect the impact of HBT correlations on **mapping speed vs. pixel size**, which is an important metric for focal plane design.

4.4.1 Mapping speed

We assume that each detector in the imaging array has three noise components: photon shot noise, photon wave noise, and internal noise. The covariance between outputs (i,j) on the imaging array is

$$\sigma_{ij} = \sigma_{ij}^{\text{shot}} + \sigma_{ij}^{\text{wave}} + \sigma_{ij}^{\text{int}}, \quad (4.43)$$

and the total variance of the detector array is

$$\sigma_{\text{arr}}^2 = \frac{1}{N_{\text{det}}^2} \sum_i \sum_j \sigma_{ij}. \quad (4.44)$$

Shot noise cannot correlate between outputs, and we assert that the internal detector noise cannot either, which allows the covariance to be written as

$$\sigma_{ij} = (\sigma_{ij}^{\text{shot}} + \sigma_{ij}^{\text{int}}) \delta_{ij} + |\gamma_{ij}|^2 \sqrt{\sigma_{ii}^{\text{wave}} \sigma_{jj}^{\text{wave}}}, \quad (4.45)$$

where $|\gamma_{ij}|^2$ is the HBT coefficient. To both simplify and clarify the calculations that follow, we assume that all detectors in the array have the same photon and internal noise

$$\sigma_{ij} = (\sigma_{\text{shot}}^2 + \sigma_{\text{int}}^2) \delta_{ij} + |\gamma_{ij}|^2 \sigma_{\text{wave}}^2, \quad (4.46)$$

where we define σ_{shot} , σ_{wave} , and σ_{int} as each noise component's RMS amplitude. While array uniformity is not true in real experiments, it is common practice to use the median noise expectation when forecasting instrument performance, making a uniform treatment useful for instrument designers.¹⁰ Given the simplification in Equation 4.46, the total variance of the detector array can be written as

$$\sigma_{\text{arr}}^2 = \frac{\sigma_{\text{shot}}^2 + \sigma_{\text{wave}}^2 + \sigma_{\text{int}}^2}{N_{\text{det}}} + \frac{1}{N_{\text{det}}(N_{\text{det}} - 1)} \sum_i \sum_{j \neq i} |\gamma_{ij}|^2 \sigma_{\text{wave}}^2. \quad (4.47)$$

¹⁰Additionally, the details of detector variations are experiment-dependent and are therefore beyond the scope of this dissertation.

The first term represents uncorrelated array noise and averages as $1/N_{\text{det}}$, while the second term quantifies HBT correlations. Let us further define the **array-averaged correlation** as

$$\gamma = \frac{1}{N_{\text{det}} - 1} \sum_i \sum_{j \neq i} |\gamma_{ij}|^2, \quad (4.48)$$

which allows us to write the array sensitivity more compactly as

$$\sigma_{\text{arr}}^2 = \frac{\sigma_{\text{shot}}^2 + (1 + \gamma)\sigma_{\text{wave}}^2 + \sigma_{\text{int}}^2}{N_{\text{det}}}. \quad (4.49)$$

Given Equation 4.49, the impact of correlations on the **detector array noise** is reduced to calculating the array-averaged HBT coefficient γ . In the limit of $|\gamma_{ij}|^2 \rightarrow 1$, $\gamma \rightarrow (N_{\text{det}} - 1)$ and σ_{wave} is not at all suppressed by array averaging.¹¹ In the other limit of $|\gamma_{ij}|^2 \rightarrow 0$, wave noise averages in the familiar way for uncorrelated measurements $\sigma_{\text{wave}}^2/N_{\text{det}}$. It is worth emphasizing that the augmentation of σ_{arr}^2 by γ not only depends on the HBT coefficient $|\gamma_{ij}|^2$ but also on the contribution of wave noise compared to that of other noise terms. As shown in Equation 4.23, $\sigma_{\text{wave}}^2 \propto n^2(\nu, T)$ and therefore becomes more important at lower frequencies and larger brightness temperatures.

To estimate the internal detector noise σ_{int} , we assume that the detectors are bolometers with strong electrothermal feedback, low impedance, and low amplifier-induced noise, as is often the case with transition edge sensors (TESes) (see Section 2.4) [161, 116]. In this paradigm, the detector's internal noise is dominated by fluctuations in heat flow between the bolometer and the bath and is defined to be [129] (see Section 3.5)

$$\sigma_{\text{int}}^2 = 4k_{\text{B}}T_{\text{c}}^2g \propto P_{\text{sat}}, \quad (4.50)$$

where T_{c} is the bolometer's operating temperature, g is its dynamic thermal conduction to the bath, and P_{sat} is its **saturation power**. The specifics of the bolometer's thermal carrier mechanism depend on the details of its design and material composition, and the choice of saturation power depends on the detector's bias parameters and observation conditions (see Section 5.5). For the calculations to follow, we adopt a representative configuration for CMB experiments, which includes a $T_{\text{b}} = 100$ mK bath temperature and a $T_{\text{c}} = 170$ mK detector temperature [32]. These assumptions relate saturation power and thermal noise as

$$\sigma_{\text{int}}^2 = (2.1 \times 10^{-11} \text{ pW/Hz}) P_{\text{sat}}, \quad (4.51)$$

which we adopt for the calculations that follow.¹²

Finally, we quantify the experiment's signal-to-noise using its **mapping speed** (see Section 3.10), which is defined as

$$\text{MS} = \frac{S^2}{\sigma_{\text{arr}}^2} \propto \frac{N_{\text{det}} \eta_{\text{apert}}^2}{\sigma_{\text{shot}}^2 + (1 + \gamma)\sigma_{\text{wave}}^2 + \sigma_{\text{int}}^2}, \quad (4.52)$$

¹¹This fact drives the use of interferometers at low frequencies where $n(\nu, T) \gg 1$, $\sigma_{\text{wave}} \gg \sigma_{\text{shot}}$, and correlation lengths are long.

¹²We assert in this chapter that other noise sources, such as those due to readout, amplifiers, and electromagnetic interference are negligible compared to photon and thermal carrier noise.




Figure 4.6: An example of a “classic” mapping speed calculation—which ignores the impact of HBT correlations—for a 90 GHz instrument with a 4 K stop. The mapping speed (MS) peak arises from the opposing effects of more detectors N_{det} and lower aperture efficiency η_{apert} with smaller pixel size D_{pix} . The optimum when ignoring correlations is $\sim 0.7 \lambda$, but as we later show, the addition of the HBT coefficient modifies this classic curve.

where S is the source signal and η_{apert} is the aperture spillover efficiency defined in Equation 4.28. Mapping speed is a powerful measure of an experiment’s efficacy, as it is $\propto N_{\text{det}}$ and therefore scales with detector yield and observation efficiency.

4.4.2 Pixel size optimization

Using the optical assumptions in Section 4.2.1, the detector assumptions in Section 4.2.2, and the VCZT coefficient in Equation 4.42, we now calculate mapping speed vs. pixel size. Provided a fixed FOV,¹³ decreasing the pixel diameter D_{pix} increases the number of detectors as $N_{\text{det}} \propto D_{\text{pix}}^{-2}$ but decreases aperture spillover efficiency as $\eta_{\text{apert}} \propto \exp[-(\pi D_{\text{pix}}/(F\lambda w_f))^2/2]$. These competing effects combine to form a peak in mapping speed vs. pixel size, revealing the optimal packing density.

An example “classic” mapping speed vs. pixel size curve [64, 5, 197, 41, 39]—one for which $\gamma \equiv 0$ —of a model 90 GHz instrument with a 4 K aperture stop is shown in Figure 4.6. Historically, ground-based CMB experiments have observed at 90, 150, and 220 GHz with $D_{\text{pix}} = 1 \sim 2F\lambda$ [5, 77, 197, 216, 184, 181, 158, 89, 172, 40], where the HBT correlation coefficient is small. However, as more powerful readout architectures mature [45, 46] and as CMB experiments push to lower frequencies for improved synchrotron charac-

¹³The throughput of most modern mm-wave telescopes is limited by cryostat size, which limits the telescope’s FOV.

terization [120, 234, 239], focal planes with $D_{\text{pix}} < F\lambda$ are becoming increasingly practical. Therefore, the impact of HBT correlations on mapping speed is of interest to upcoming mm-wave experiments, such as CMB-S4 [205]. In this section, we calculate HBT-modified mapping speed vs. pixel size curves for the model telescope described in Section 4.2 at several observation frequencies.

Consider a ground-based telescope with cryogenically cooled optics, infrared filters, and mK detectors. The telescope is highly transparent to mm-wave photons and has negligible spillover and scattering losses beyond those of the aperture stop. In this paradigm, we can characterize loading from sources external to the telescope’s cold box, including the CMB, atmosphere, and ambient optics, via their effective brightness temperature

$$T_{\text{b}}(\nu) = \frac{h\nu}{k_{\text{B}} [1 + (h\nu/\eta_{\text{apert}} \eta_{\text{det}} p(\nu))]} ,$$

where $p(\nu)$ is the source’s power spectrum referred to the telescope’s input, η_{det} is the detector’s optical efficiency, and the single-mode throughput is $A\Omega = \lambda^2$. Telescope temperatures can vary considerably depending on the specifics of the mirrors, cryostat window, and anti-reflection coatings, but for the calculations that follow, we assume $T_{\text{tel}} = 10$ K, $T_{\text{stop}} = 4$ K, and $\eta_{\text{det}} = 0.8$. To model power due to atmospheric emission, we assume that the telescope observes from the Chajnantor Plateau in the Atacama Desert of Chile (see Section 2.2), and we use the AM model [151] to generate the atmosphere’s equivalent brightness temperature and transmittance at 1 mm precipitable water vapor (PWV) and 50 deg elevation above the horizon.

Within this telescope + sky model, we consider four top-hat observation bands centered at (35, 95, 150, 220) GHz with bandwidths of (17, 33, 39, 44) GHz. The bands, along with the equivalent brightness temperature of the CMB, atmosphere, and telescope, are shown in Figure 4.7. The channels are chosen to fit within the atmospheric windows and are similar to those of existing Atacama instruments. Additionally, Figure 4.7 shows the bunching fraction for each frequency band vs. sky temperature at $D_{\text{pix}} = 3F\lambda$, for which $\eta_{\text{apert}} > 0.95$. As discussed in Section 4.4.1, lower-frequency modes tend to have larger occupation numbers, but because the sky is brighter at higher frequencies, we expect correlations to impact all four observation bands substantially.

Additionally, we note that detector saturation power is, in general, set after the pixel size is chosen and the expected optical load is estimated. Such a tunable parameter is in contrast to the CMB, telescope, and sky temperatures, which are more-or-less independent of the focal plane optimization.¹⁴ Saturation-power optimization is a rich topic that depends on many experiment specific details, such as observation strategy, telescope and detector optical efficiency, and detector bias technique (see Section 5.5), but we can marginalize these

¹⁴Strictly speaking, the telescope temperature does change with pixel size, which modulates the ratio of aperture-stop loading to that of the other optics. However, most telescope optics are cryogenic and therefore have a similar temperature to that of the stop, and as a result, telescope temperature varies only slowly with aperture spillover efficiency.




Figure 4.7: The assumed observation bands plotted over the assumed CMB, atmosphere, and telescope temperatures (left panel), and the Bose noise fraction vs. sky temperature for each band at $D_{\text{pix}} = 3F\lambda$, where $\eta_{\text{apert}} > 0.95$ (right panel).

factors by asserting that P_{sat} scales with the detected optical power P_{opt} as [197]

$$P_{\text{sat}} = 3P_{\text{opt}}. \quad (4.53)$$

Excess saturation power ensures that the detector can be biased when the sky fluctuates above its median brightness (e.g., when the PWV $\gtrsim 1$ mm and the horizon elevation $\lesssim 50$ deg).

Figure 4.8 shows an example of four pixel-pitch scenarios $D_{\text{pix}} = (0.2, 0.5, 0.7, 1.0) F\lambda$ given the hex packing described in Section 4.2.2, and the color contour shows the HBT coefficient due to both aperture and stop radiation. As expected, the degree of correlation is a 2D analog of the VCZT curves in Figure 4.5. The effect of pushing $D_{\text{pix}} < 1.2F\lambda$ is for the output modes (represented by the pixel apertures) to oversample the input modes (represented by the $|\gamma(\vec{r}_1, \vec{r}_2)|^2$ contour), giving rise to correlated noise.

When calculating mapping speed, we sum the HBT coefficients in Figure 4.8 over a $4F\lambda$ radius to find $\gamma_i = \sum_{j \neq i} |\gamma_{ij}|^2$, noting that because “+” and “ \times ” pixels are orthogonal, $|\gamma_{ij}|^2$ vanishes for half of all (i, j) output pairs. We then assume that this sum applies to all detectors on the focal plane such that $\gamma = \gamma_i$. Such a treatment ignores the fact that edge pixels have fewer neighbors than internal ones, which is a reasonable approximation for a focal plane with a large area. For small or moderately sized detector arrays, the fraction of edge pixels may become important, but such details are experiment-dependent and are therefore beyond the scope of this discussion. Figure 4.9 shows HBT-impacted mapping




Figure 4.8: $|\gamma(\vec{r}_2 - \vec{r}_1)|^2$ given pixel pitches $D_{\text{pix}} = (0.2, 0.5, 0.7, 1.0) F\lambda$ (from top to bottom) for radiation from within the aperture (left panels) and from the stop (right panels). When the pixel pitch is $\lesssim 1.2F\lambda$, the detectors oversample spatially independent modes on the focal plane, giving rise to HBT correlations between neighboring pixels.

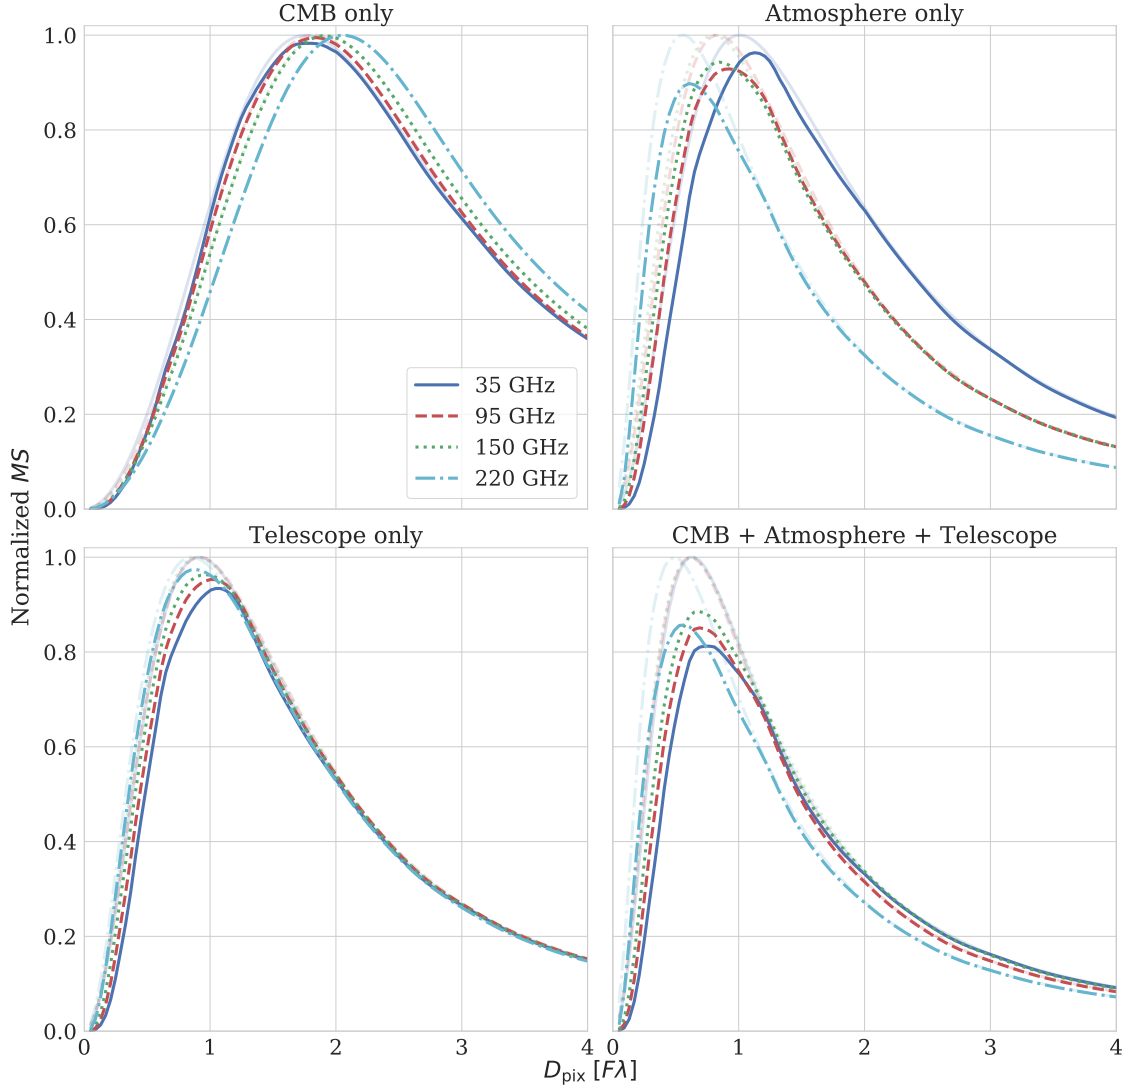


Figure 4.9: Mapping speed vs. pixel size in each frequency band for four sets of external sources: the CMB only, the atmosphere only, telescope emission only, and all three combined. Each band is normalized to the peak of its classic curve ($\gamma \equiv 0$), which is represented by faded lines.

speed vs. pixel size curves for each observation band, normalized to the peak of each band’s “classic” curve for which $|\gamma(\vec{r}_2 - \vec{r}_1)|^2 \equiv 0$. Loading due to the CMB, atmosphere, and telescope are plotted independently to demonstrate the dependence of correlations on source temperature.

There are several features in Figure 4.9 that are worth noting explicitly. Firstly, the impact of HBT correlations depends on source temperature and is most pronounced in the presence of a brightly illuminated aperture. This effect is most clearly seen when contrasting the CMB and atmosphere, especially at 220 GHz where the intensity of the CMB is fading while that of the atmosphere is growing. Secondly, while the HBT impact on curve *shape* is

most pronounced at low frequencies, the *overall suppression* is similar among the frequency bands. This effect arises because the atmosphere is brighter at higher frequencies, increasing the wave-noise fraction (see Figure 4.7) even as the γ decreases. Thirdly, the HBT-impacted mapping speed peak is pushed to slightly larger D_{pix} at 35 and 95 GHz while remaining relatively unmoved at 150 and 220 GHz. This effect arises because the 4 K stop is significantly fainter than the sky at 150 and 220 GHz, and hence $(1 + \gamma)\sigma_{\text{wave}}$ decreases more rapidly than η_{apert} at high frequencies. Lastly, while the impact of correlations on mapping speed is most prominent when $D_{\text{pix}} < 1.2F\lambda$, there are percent-level impacts at larger spacings, which correspond to local maxima in the VCZT patterns shown in Figure 4.5.

4.5 Implications for experiment design

As indicated by Figure 4.9, HBT correlations both modify the optimal pixel packing density and suppress the achievable mapping speed with respect to the “classic” $\gamma \equiv 0$ calculation. However, the severity of the modification depends on a plethora of instrument details, including internal detector noise, amplifier noise, observation site and conditions, stop temperature, detector efficiency, camera optical throughput, and extraneous systematic noise sources, such as electromagnetic interference, vibrational pickup, and detector nonidealities (see Chapter 3). A comprehensive handling of correlations within a more detailed experiment is available via the BoloCalc sensitivity calculator [81] described in Chapter 5, but in this section, we sweep a few parameters in our model telescope to provide a quick reference for focal plane designers. The results of these calculations are shown in Figure 4.10.

The first column of Figure 4.10 shows mapping speed vs. pixel size for various stop temperatures $T_{\text{stop}} = (2, 3, 4, 5)$ K. As stop temperature decreases, so too does photon loading due to stop spillover, favoring smaller pixels. A colder stop also suppresses the relative contribution of σ_{wave} , especially at higher frequencies, which modulates the mapping speed slope below $D_{\text{pix}} = 1.2F\lambda$. The second column of Figure 4.10 shows mapping speed vs. pixel size for various telescope temperatures $T_{\text{tel}} = (10, 20, 30, 40)$ K. As telescope temperature increases, so too does the photon load within the aperture, favoring smaller pixels. In addition, the brighter the aperture radiation, the greater σ_{wave} becomes, further suppressing the mapping speed curves. The third column of Figure 4.10 shows mapping speed vs. pixel size in the presence of a constant internal detector noise $\sigma_{\text{int}}^2 = (0.0, 0.5, 0.7, 1.0) \times \sigma_{\text{ph}}^2$ at $D_{\text{pix}} = 1.2F\lambda$, where $\sigma_{\text{ph}}^2 \equiv (\sigma_{\text{shot}}^2 + \sigma_{\text{wave}}^2)$. This generic noise term can represent a wide variety of phenomena, including excess bolometer thermal carrier noise, amplifier noise, or any other optically-independent white noise source. As non-photon noise increases, a larger pixel size is favored to improve signal strength (see Equation 4.52), and the impact of HBT correlations is reduced both due to the reduced relative contribution of σ_{wave} and due to the optimal $D_{\text{pix}} \gtrsim 1.2F\lambda$ where γ is small.

The examples in Figure 4.10 only graze the rich topic of focal plane optimization, and we leave a more comprehensive discussion of experiment-specific applications to other publications [81]. Nonetheless, regardless of the context, HBT correlations should be considered

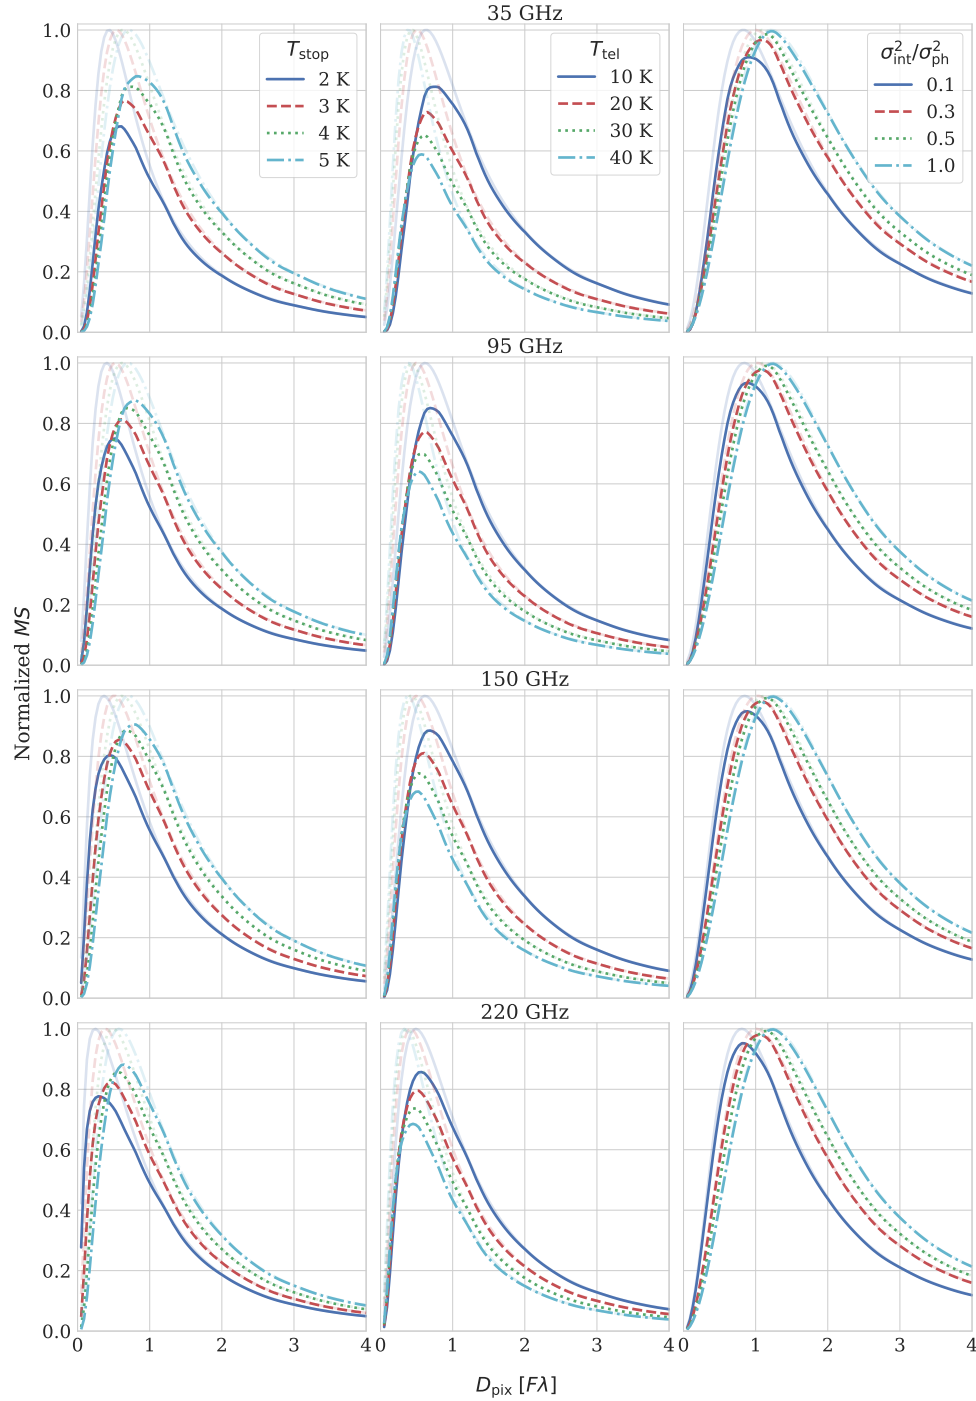


Figure 4.10: The impact of stop temperature T_{stop} , telescope temperature T_{tel} , and detector internal noise σ_{int}^2 on mapping speed vs. pixel size in the presence of HBT correlations. Each band is normalized to the peak of its classic curve ($\gamma \equiv 0$), which is shown by the faded lines. These parameters are among many that vary between experiments, and we include them here as a reference for focal plane designers.

when designing dense focal planes, especially for ground-based experiments for which sky and telescope brightness temperatures are substantially larger than that of the CMB.

4.6 Discussion

In this chapter, we have presented a theoretical formalism for photon noise correlations by extending the quantum optical circuit-based model in Zmuidzinas (2003) [241] to a free-space classical model using the equivalence theorem of Glauber and Sudarshan (1963) [61, 194]. We have used this formalism to simulate the Hanbury Brown-Twiss (HBT) coefficient in a simplified telescope system, and we have shown that these simulations match the expectation of the van Cittert-Zerenike theorem (VCZT). This equivalence allows the HBT coefficient to be calculated with only a knowledge of the intensity profile at the aperture plane.

We uniformly illuminated a model telescope with blackbody radiation from the CMB, atmosphere, and telescope, and we calculated the impact of HBT correlations on mapping speed vs. pixel size within observation bands centered at 35, 95, 150, and 220 GHz. Acknowledging that sensitivity calculations have many inputs and assumptions, we further discussed three useful variations to the simplified instrument—stop temperature, telescope temperature, and internal detector noise—and showed how each parameter modulates the mapping speed curves. This work builds on a piloting discussion by Padin (2010) [149] and formalizes the calculation of photon noise correlations between detector pixels in mm-wave telescopes. The presented formalism is useful to GHz focal plane designers, especially as emerging readout technologies enable the readout of dense detector arrays, and it is implemented as part of the sensitivity calculator discussed in Chapter 5.

Chapter 5

BoloCalc: a sensitivity calculator

This chapter overlaps substantially with Hill, Bruno, Simon, et al. (2018) [81], which we encourage the reader to cite instead of or in addition to this dissertation.

As the CMB research field matures, it is conglomerating the efforts of many institutions into a small number of large projects [215, 204]. This paradigm shift allows for higher-budget observatories that can field more detectors and build more powerful telescopes. Despite the allure of these “mega experiments,” large projects come with a plethora of logistical challenges. One such challenge is to coalesce each institution’s independently-developed sensitivity codes into one versatile, community-wide calculator. The effectiveness of a broadly-usable sensitivity code relies on a few core capabilities:

1. The code must be **general and modular**, which allows it to simulate a wide variety instrument configurations. This capability is especially important during the early stages of observatory design, when many subsystems are being optimized across vast parameter spaces.
2. The code must be **detailed and comprehensive**, which allows the instrument to be modeled accurately. There are dozens of inputs to any sensitivity calculation, and while some may be unimportant for one telescope, those same inputs may be critically important to another.
3. The code must be **easy to use**, which will encourage its utilization throughout the collaboration.

In this chapter, we present **BoloCalc**, a generalized sensitivity calculator for CMB instrument characterization. The tool is publicly available¹ and is being actively used not only within the CMB community but also within IR and radio astronomy.

¹BoloCalc repository: <https://github.com/chill90/BoloCalc>

5.1 Calculator design

BoloCalc is a Python code that (typically) takes $\mathcal{O}(100)$ user-defined inputs and calculates dozens of outputs described in Chapter 3, including optical power P_{opt} , photon, thermal, and readout noise-equivalent power (NEP), noise-equivalent temperature (NET), and mapping speed (MS). While it runs as a standalone command-line tool, BoloCalc also offers a PyQt² graphical user interface (GUI), which eases interaction with input configuration files and adds real-time data visualization for quicker, more comprehensive sensitivity assessments.

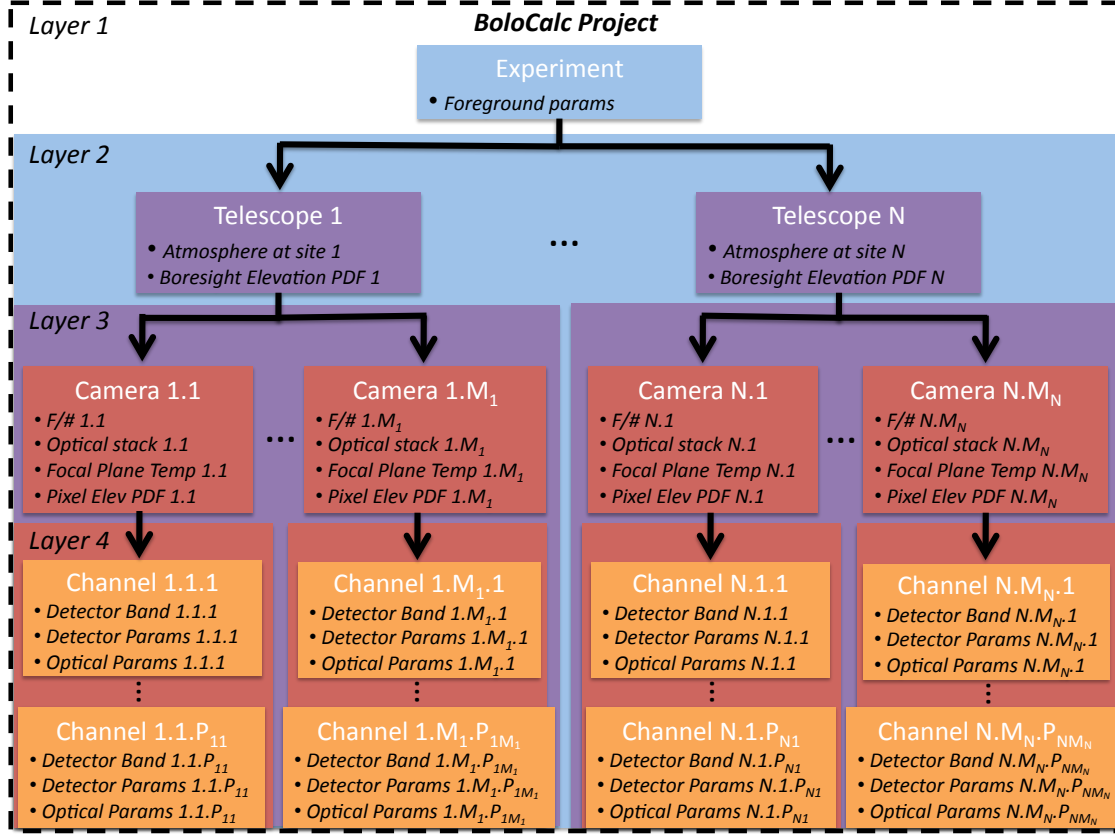
5.1.1 Structure

BoloCalc has a modular **object-oriented** structure, which allows for arbitrary mixtures of sites, telescopes, cameras, optics, focal planes, and detectors. A BoloCalc project has the **parent-child structure** shown in Figure 5.1 and is built with **four layers**: experiments, telescopes, cameras, and channels, which are defined in the table of Figure 5.1. Each experiment can have an arbitrary set of telescopes (at different sites), each telescope an arbitrary set of cameras, and each camera an arbitrary set of channels. Each layer of a BoloCalc project contains various user-defined parameters, and the inheritance structure follows that of parent-child such that each telescope inherits the parameters of its experiment, each camera inherits those of its telescope, and each channel those of its camera. This flexible construction has proven valuable to applications within Simons Observatory (SO) and LiteBIRD, especially during each experiment’s early stages when the numbers of telescopes, cameras, frequencies, and detectors were undecided and when rapid feedback was needed to evaluate various instrument configurations.

The table in Figure 5.1 shows the user-defined parameters for layers 1–4 of a BoloCalc project. Layer 1 defines the **foreground parameters** for each experiment, which determines celestial optical loading on the detector. While foregrounds contribute little in-band power relative to the atmosphere and CMB at ~ 100 GHz, they can become important for satellite experiments that observe at very high and/or low frequencies. Therefore, foregrounds are included in BoloCalc for LiteBIRD and similar missions [173]. Layer 2 defines each telescope’s **atmospheric conditions**, as well as its elevation distribution, observation time and efficiency, and sky fraction. Layer 3 defines each camera’s **optical chain** and magnification, as well as its field of view (FOV) and focal plane temperature. Layer 4 defines the channels within each camera, including the **detector parameters**, bandpasses, and antenna beam properties. This layered organization allows BoloCalc to seamlessly interface with the **experiment directory structure**, which is shown in Figure 5.3.

The BoloCalc source code is composed of 30 classes that form an **inheritance tree**, and each class serves a specific function. This object-oriented structure makes the code easy to modify and computationally efficient. Efficiency is especially important when running

²PyQt: <https://wiki.python.org/moin/PyQt>



Layer	Definition
Experiment	An assemblage of CMB telescopes.
Telescope	A platform that carries and points one or more cameras. It observes at a specified site with a specified observation strategy and can include warm reflectors.
Camera	A cryostat that houses cryogenic optics, filters, and detectors. Multiple cameras can be mounted on the same telescope.
Channel	A frequency band observed by some set of detectors within a camera. A multichroic camera will have multiple channels.

Figure 5.1: The generic layout of a BoloCalc project. There are four layers to a project, each with its own set of parameters: (1) experiments, (2) telescopes, (3) cameras, and (4) channels. There can be an arbitrary set of N telescopes in each experiment, an arbitrary set of M cameras in each telescope, and an arbitrary set of P channels in each camera. Each telescope inherits the parameters of its parent experiment, each camera inherits those of its telescope, and each channel those of its camera. The black bullet points highlight some important parameter definitions in each layer. The table provides definitions for each layer.

Monte Carlo realizations of the experiment and when running parameter sweeps, which we detail in the following sections.

5.1.2 Input and output parameters

BoloCalc accepts 65 unique input parameters that are used to describe foreground emission, the atmosphere and observation strategy, the camera optics, and the detector specifications. These components then compose an instrument model that is used to simulate 12 outputs. Below is a comprehensive list of the inputs, organized by level, along with their corresponding symbol in Chapter 3 and Section 5.2.1.

- **Level 1:** foreground parameters
 - Synchrotron spectral index: n_s in Equation 5.1
 - Synchrotron amplitude: T_s in Equation 5.1
 - Synchrotron scale frequency: ν_s in Equation 5.1
 - Dust temperature: T_d in Equation 5.2
 - Dust spectral index: n_d in Equation 5.2
 - Dust amplitude: A_d in Equation 5.2
 - Dust scale frequency: ν_d in Equation 5.2
- **Level 2:** observation and telescope parameters
 - Atmospheric temperature: T_{ATM} in Equation 2.1
 - Site: site at which the telescope observes (see Section 2.2)
 - Elevation: θ_{elev} in Equation 2.3
 - Precipitable water vapor: PWV in Equation 2.4
 - Observation time: t_{obs} in Equation 3.68
 - Sky fraction: f_{sky} in Equation 3.68
 - Observation efficiency: η_{obs} in Equation 3.68
 - NET margin: factor applied to all NETs within a telescope, which can be treated as contingency
- **Level 3:** camera parameters
 - Boresight elevation: the camera’s boresight elevation with respect to the telescope’s boresight
 - Optical coupling: ξ in Equations 3.59 and 3.60
 - F-number: F in Equation 2.10
 - Bath temperature: T_b in Equations 3.46, 3.47, and 3.50,
- **Level 4:** channel parameters
 - Band center: ν_c in Equation 5.3
 - Fractional bandwidth: f_{BW} in Equation 5.3
 - Pixel size: D_{pix} in Equation 2.10
 - Number of detectors per wafer

- Number of wafers per optics tube
 - Number of optics tubes in this camera
 - Waist factor: w_f in Equation 2.10
 - Detector optical efficiency: η_{det} in Equations 5.3
 - Bolometer saturation power: P_{sat} in Equations 3.50 and 3.54
 - Psat factor: allows the user to define $P_{\text{sat}} = f_{\text{psat}} \times P_{\text{opt}}$
 - Carrier index: n in Equations 3.47 and 3.50
 - Bolometer transition temperature: T_c in Equations 3.46, 3.47, and 3.50
 - Tc fraction: allows the user to define $T_c = f_c \times T_b$
 - Thermal link factor: F_{link} in Equations 3.46 and 3.47
 - Dynamic thermal conductance: g in Equations 3.46 and 3.50
 - Detector yield: Y in Equation 3.62
 - Readout noise: NEI_{read} in Equation 3.54
 - Bolometer operating resistance: R_{bolo} in Equation 3.54
 - Responsivity factor: S_{fact} in Equations 2.19, 3.53, and 3.54
 - Readout noise fraction: allows the user to define $\text{NEP}_{\text{read}} = f_{\text{read}} \sqrt{\text{NEP}_{\text{ph}}^2 + \text{NEP}_{\text{g}}^2}$
- **Level 4:** optics parameters
 - Temperature: T_i in Equation 3.19
 - Emissivity/absorptivity: ϵ_i in Equation 3.19
 - Reflectivity: r_i in Equations 3.25 and 3.40
 - Optical thickness: t in Equation 3.28
 - Refractive index: n in Equation 3.28
 - Loss tangent: $\tan \delta$ in Equation 3.28
 - Conductivity: σ in Equation 3.29
 - Surface roughness: σ_{surf} in Equation 3.36
 - Scatter fraction: δ_i in Equation 3.19
 - Scatter temperature: $T_{\delta;i}$ in Equation 3.19
 - Spillover fraction: β_i in Equation 3.19
 - Spillover temperature: $T_{\beta;i}$ in Equation 3.19

Some optics and channel parameters are **functionally redundant**, offering the user multiple methods for calculating emissivity, efficiency, and scattering. For example, the absorptivity of a refractive optic can be entered explicitly, or it can be derived using the loss tangent, thickness, and refractive index as in Equation 3.28. In a similar manner, there are multiple ways to define certain channel parameters. For example, the dynamic thermal conductance g can be defined either explicitly or using T_c , T_b , and n as in Equation 3.49. As another example, readout noise can be defined using NEI_{read} , R_{bolo} , and P_{sat} , or it can be stated as a fraction of the other noise sources. These redundancies make BoloCalc useful for a wide range of forecasting scenarios, from defining high-level specifications during project conception to evaluating low-level measurements during instrument commissioning.

BoloCalc uses the presented input parameters and the equations in Chapter 3 to calculate **12 output parameters**

- Optical throughput: η_{through} in Equation 3.40
- Optical power: P_{opt} in Equation 3.25
- Telescope temperature: T_{tel} in Equation 3.45
- Sky temperature: T_{sky} in Equation 3.45
- Photon NEP: NEP_{ph} in Equation 3.23
- Bolometer thermal NEP: NEP_{g} in Equation 3.46
- Readout NEP: NEP_{read} in Equation 3.54
- Detector NEP: NEP_{det} in Equation 3.57
- Detector NET (CMB or RJ): NET_{det} in Equation 3.58
- Array NET (CMB or RJ): NET_{arr} in Equation 3.62
- Correlation factor: Γ in Equation 3.67
- Map depth: σ_{map} in Equation 3.68

These parameters are calculated for each channel in each camera, but the NETs and map depths are also combined via inverse-variance weighting (see Equation 3.61) at the telescope and experiment levels to estimate the *total* sensitivity of multiple detector arrays across multiple cameras and telescopes. The outputs’ median and (15, 85)% confidence values are displayed in text tables, but histograms of the output data are also available when running Monte Carlo (MC) instrument realizations (see Section 5.2.3). This capability is important for accurate forecasting, as NET distributions tend to be **skewed** and therefore poorly described using Gaussian distributions, even if the input distributions are Gaussian.

In addition to generating the listed sensitivity outputs, BoloCalc generates tables of optical power and optical throughput as a function of location within the camera’s optics chain. An example of such a table is shown in Figure 5.2. This functionality is quite useful for assessing the propagation of optical power through the telescope system and for identifying areas of improvement when trying to reduce NEP_{ph} .

5.1.3 User interface

There are two ways to interact with BoloCalc. The first way, which has been the primary interface during most of the calculator’s lifetime, is through the **command line**. In this scheme, inputs are set by editing text tables, and outputs are also written to text tables.³ BoloCalc’s `unpack.py` class can then be used to convert the output tables into dictionaries, which are easily loaded, plotted, and manipulated in a more friendly environment.⁴ While the command-prompt scheme is well-tested and self-sufficient, it is cumbersome at times and can be a barrier for new users. Additionally, as is always the case for Python’s runtime interpreter, errors in user inputs are not assessed until execution. While BoloCalc’s logging

³These text tables are set up to be most easily viewed using Emacs or Vim.

⁴The base BoloCalc distribution provides a **Jupyter Notebook** tutorial for output handling.

Element	Power from Sky	Power to Detector	Efficiency	Cumulative Efficiency
Units	[pW]	[nA]	[NA]	[NA]
CMB	0	0.058 +/- (0.011, 0.01)	1	0.12 +/- (0.021, 0.021)
ATM	0.497 +/- (0.026, 0)	0.454 +/- (0.093, 0.086)	0.967 +/- (0.002, 0.001)	0.124 +/- (0.022, 0.021)
Primary	4.255 +/- (0.28, 0.295)	0.179 +/- (0.035, 0.031)	0.988	0.126 +/- (0.022, 0.022)
Mirror	5.64 +/- (0.321, 0.291)	0.03 +/- (0.006, 0.005)	0.998	0.126 +/- (0.022, 0.022)
Window	5.868 +/- (0.328, 0.291)	0.077 +/- (0.015, 0.014)	0.97 +/- (0.002, 0.002)	0.13 +/- (0.023, 0.022)
IRShader1	6.277 +/- (0.33, 0.259)	0.016 +/- (0.003, 0.003)	0.999	0.13 +/- (0.023, 0.022)
IRShader2	6.396 +/- (0.335, 0.26)	0.014 +/- (0.003, 0.002)	0.999	0.13 +/- (0.023, 0.022)
IRShader3	6.497 +/- (0.338, 0.26)	0.011 +/- (0.002, 0.002)	0.999	0.13 +/- (0.023, 0.022)
AbsFilter	6.578 +/- (0.343, 0.259)	0.046 +/- (0.009, 0.008)	0.97	0.134 +/- (0.023, 0.023)
LowPass1	6.724 +/- (0.343, 0.254)	0.024 +/- (0.005, 0.004)	0.941	0.142 +/- (0.025, 0.025)
Lens1	6.491 +/- (0.328, 0.239)	0.008 +/- (0.002, 0.003)	0.938 +/- (0.009, 0.012)	0.151 +/- (0.025, 0.025)
LowPass2	6.138 +/- (0.313, 0.221)	0.002	0.941	0.161 +/- (0.027, 0.027)
Lens2	5.783 +/- (0.295, 0.208)	0	0.99	0.162 +/- (0.027, 0.027)
Aperture	5.727 +/- (0.292, 0.206)	0.092 +/- (0.018, 0.015)	0.309 +/- (0.001, 0.001)	0.525 +/- (0.088, 0.088)
LowPass3	1.87 +/- (0.082, 0.06)	0.001	0.941	0.559 +/- (0.093, 0.093)
Lens3	1.762 +/- (0.077, 0.056)	0.001	0.99	0.564 +/- (0.094, 0.094)
LowPass4	1.745 +/- (0.076, 0.055)	0.001	0.941	0.6 +/- (0.1, 0.1)
Detector	1.643 +/- (0.072, 0.052)	0	0.6 +/- (0.1, 0.1)	1

Figure 5.2: A BoloCalc optical power table for an arbitrarily constructed example experiment. The columns define, from left to right, the incident power on each optic from the sky side, the detected power due to each optic’s emission, spillover, and scattering, the efficiency of each optic, and the cumulative optical efficiency from the detector towards the sky. The central values represent the median, while the +/- values are the (85, 15) percentiles. This table is generated as part of BoloCalc’s Jupyter Notebook interface.

and error tracing are both well developed and comprehensive, the inconvenience of debugging at the command line can be a deterrent for casual users. Therefore, we have developed a graphical user interface (GUI) using PyQt named **BoloCalc GUI (BCG)**.

The GUI substantially improves BoloCalc’s user experience in several ways. First, it eliminates the need to interact with the command line and edit text tables. Because BoloCalc has so many inputs, especially for the optics (see Section 3.4), text files can be cumbersome to maintain. BCG acts as a **wrapper** around the input text files, allowing the user to modify parameters using line edits, combo boxes, and checkboxes, as shown in Figure 5.4. Second, BCG performs **real-time error checking** on user inputs via embedded environments. This scheme transfers input checking to the front end (when data are entered) instead of the back end (during Python execution). In other words, BCG ensures that when the user runs BoloCalc’s executable `calcBolos.py`, the program will complete without interruption. Third, BCG has **embedded descriptions** of input and output parameters in each parameter’s edit window. This feature increases the transparency of the sensitivity calculation and provides immediate information about how each input impacts the outputs. While BoloCalc’s user




Figure 5.3: The file structure of a BoloCalc experiment, which mimics the calculator’s underlying class structure. This layout enables each layer in Figure 5.1 to be modular, making it straightforward to update experiment versions, add telescopes and cameras, and visualize how the various subsystems fit together.

manual⁵ explains how text-file inputs correspond to the equations in Chapter 3, studying the manual while using the command prompt can be cumbersome. BCG’s **in-window definitions**, one of which is shown in Figure 5.5, fix this issue while also serving as a pedagogical tool.

Fourth, BCG provides **real-time data visualization** of both the inputs and, after running `calcBolos.py`, the outputs. For example, BCG allows the user to view plots of input probability distributions and bandpass functions for optics and detectors. In a similar way, outputs can be either downloaded or plotted directly using the GUI’s **analyze** tool. These visualization functionalities allow for a quick assessment of the assumed instrument model and simulation results, helping facilitate the rapid compilation of documentation and presentations. Fourth, BCG effectively divorces BoloCalc’s source code from the user interface, which is enormously useful for development purposes. BoloCalc’s underlying text files and Python classes are interfaced to BCG using **intermediary configuration files**, improving backward compatibility and making program upgrades more seamless.

BCG’s primary goal is to make BoloCalc more accessible to a wider range of potential users. During instrument design and development, many scientists are evaluating disparate instrument systems in parallel, and if the experiment’s sensitivity calculator is not user-friendly, the entire collaboration’s forecasting workload will likely flow through a small num-

⁵The user manual is available as a PDF within the BoloCalc distribution.

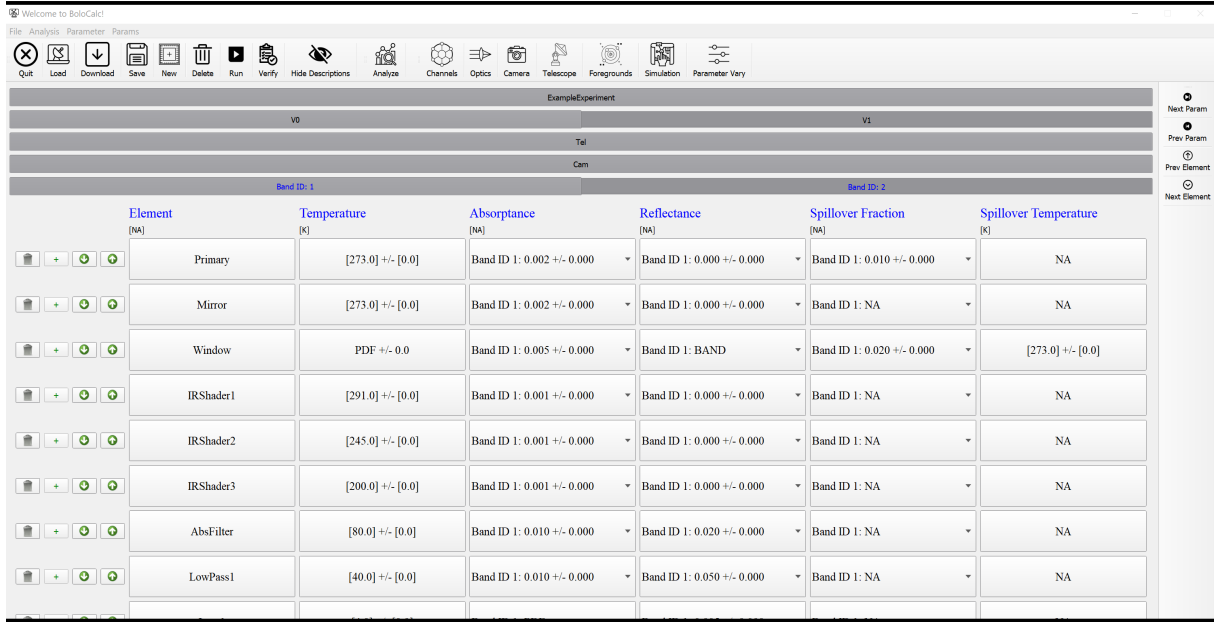


Figure 5.4: A screenshot of BCG’s interface for editing the optics’ input parameters. The dark grey toolbar navigates between telescopes, cameras, and channels, and the interactive table provides an interface for editing optical parameters. The drop-down menus for each table cell allow each channel to be edited independently.

ber of experts. However, if the calculator is user-friendly and easy to learn, sensitivity studies can be more broadly distributed and design feedback can be provided more rapidly. BCG is specifically built to enhance BoloCalc’s usability for ongoing and future experiments, such as SO and CMB-S4, as well as make the calculator more appealing to potential applications outside the CMB community.

5.2 Calculator features

BoloCalc has several features that distinguish it from predecessor sensitivity calculators. Its object-oriented structure makes the code easy to upgrade and modify, and therefore many these features are the result of user requests.

5.2.1 Foregrounds

BoloCalc provides an option to define the power due to dust and synchrotron emission. While foreground intensity is far below that of the CMB at ~ 100 GHz and at high galactic latitudes, satellite experiments often observe at high frequencies where dust becomes bright, while radio telescopes often observe at low frequencies where synchrotron becomes bright. In




Figure 5.5: An example BCG window for editing the “Band Center” parameter. The window on the left shows the parameter’s name, range, unit, and description while allowing the user to define it either as a float or via a probability distribution file (PDF) or band file. In this particular example, the detector band is defined by an uploaded CSV file, and the data points are overplotted onto the assumed (median) atmospheric profile.

order to accommodate such dust and synchrotron mappers, BoloCalc includes the following foreground formalisms, which assume a uniform intensity across the sky.⁶

The **synchrotron power spectral density** for a single polarization in $\text{W Hz}^{-1} \text{sr}^{-1}$ is given by the equation

$$p_{\text{synch}}(\nu) = A\Omega \frac{\nu^2}{c^2} k_B T_s \left(\frac{\nu}{\nu_s} \right)^{n_s} = k_B T_s \left(\frac{\nu}{\nu_s} \right)^{n_s}, \quad (5.1)$$

where T_s is the synchrotron brightness temperature in K_{RJ} at pivot frequency ν_s , n_s is the synchrotron spectral index, and the etendue $A\Omega$ for a diffraction-limited detector is given by Equation 3.21. The **dust power spectral density** in units of $\text{W Hz}^{-1} \text{sr}^{-1}$ is given by the equation

$$p_{\text{dust}}(\nu) = A_d A\Omega \left(\frac{\nu}{\nu_d} \right)^{n_d} \frac{S(T_d, \nu)}{S(T_d, \nu_d)} \times 10^{-20} = A_d \frac{c^2}{\nu^2} \left(\frac{\nu}{\nu_d} \right)^{n_d} \frac{S(T_d, \nu)}{S(T_d, \nu_d)} \times 10^{-20}, \quad (5.2)$$

⁶At the moment, BoloCalc uses the telescope’s boresight elevation above the horizon to monitor atmospheric emission. However, for space telescopes, it may become useful to monitor foreground emission as a function of galactic coordinate, and such a capability can easily be added to BoloCalc in the future.




Figure 5.6: A comparison of the Atacama and the South Pole atmospheric profiles used in BoloCalc, generated using the AM atmospheric simulator. Each profile’s PWV (left) and elevation dependence (right) is also plotted. The deviation between sites at high frequency is due to the Atacama’s being at higher elevation (5,200 m) than the South Pole (2,800 m).

where A_d is the dust amplitude in MJy sr^{-1} at pivot frequency ν_d , n_d is the dust spectral index, $S(T_d, \nu)$ is the blackbody spectral density function defined in Equation 3.19, and T_d is the dust temperature. Note that the conversion from MJy to W is 10^{-20} W/MJy .

5.2.2 Site and atmosphere

BoloCalc uses atmospheric transmissivity and temperature vs. frequency when calculating P_{opt} (see Section 2.2). Such a spectrum is often called an **atmospheric profile** and depends on the details of the observation site, including elevation above sea level, weather patterns, and sky access. There are three methods to define the atmosphere in BoloCalc: via a constant temperature and transmissivity vs. frequency, via a custom profile, and via a simulated profile selected by PWV and horizon elevation. In this section, we highlight the final method, which is typically most useful and accurate.

BoloCalc’s atmospheric profiles are generated by the **am Atmospheric Model**⁷ developed at Harvard [151]. **am** is a C++ command-line tool that processes a configuration file with details of the atmosphere’s molecular composition from the troposphere to the stratosphere. BoloCalc includes atmospheric profiles for three different observation sites: the Atacama of Chile, the South Pole, and the stratosphere, from where balloon-borne experiments typically observe [206, 65, 6]. BoloCalc’s atmospheric configurations are provided by Scott Paine

⁷am Atmospheric Model: <https://www.cfa.harvard.edu/~spaine/am/>

and Denis Barkats, and they are constructed using data from the MERRA-2 meteorological reanalysis.⁸ These atmospheric profiles are well verified, showing excellent agreement with sky-temperature measurements from the South Pole Telescope (SPT) and the Atacama Cosmology Telescope (ACT). A side-by-side comparison of the Atacama and South Pole atmospheric profiles, including their dependence on PWV and elevation, is shown in Figure 5.6.

BoloCalc contains an HDF5 file with atmospheric spectra at [0.1, 0.2, ... 1000.0] GHz for all three sites at PWVs of [0, 0.1, ..., 8.0] mm and elevations of [0, 1, ..., 90] deg above the horizon. When the user defines the “site,” “elevation,” and “PWV” input parameters, BoloCalc accesses the corresponding profile using the `h5py` Python class.⁹ HDF5 files are ideal for atmosphere handling, as they facilitate fast look-up access to large datasets directly from disk. Quick retrieval becomes even more important when running many MC realizations of the sky (see Section 5.2.3).

5.2.3 Monte Carlo simulations

From forecasting an instrument’s scientific impact to measuring its performance in the field, uncertainties are critical to a proper sensitivity estimate. For example, when iterating on high-level design concepts during an experiment’s earliest stages, large error bars will surround back-of-the-envelope calculations, and propagating these uncertainties to NET estimates can inform questions such as resource allocation, contingency scenarios, and even science targets. As another example, when using planet observations to characterize detector responsivity in the field, measurement uncertainties, focal-plane variations, and modeling limitations must be incorporated into accompanying NET simulations to properly estimate the telescope’s capabilities. Given the importance of uncertainty handling, BoloCalc uses **MC simulations** to generate histograms of output parameters.

BoloCalc accepts either a mean and standard deviation or a custom probability distribution for every input parameter. For custom spectral bands (see Section 5.2.5), every data point can be defined with a mean and standard deviation, while PWV and elevation distributions give rise to variations in the selected atmospheric profile (see Section 5.2.2). BoloCalc breaks MC simulations into **three levels**: (1) realizations of the foregrounds, telescope, camera, and optics, (2) realizations of the sky, and (3) realizations of the detector parameters. The number of realizations at each level are defined separately, meaning that for each of N experiments, M skies are generated, and for each of M skies, P detectors are generated. Separating the simulation into these three categories is intended to mimic the instrument’s progression: (1) build and field a telescope, (2) observe with that telescope in a variety of sky conditions, and (3) given the sky conditions, calibrate each detector’s

⁸MERRA-2: <https://gmao.gsfc.nasa.gov/reanalysis/MERRA-2/>

⁹Note that BoloCalc does not make any simplifying assumptions about how the atmospheric loading changes with elevation or PWV, as suggested by Equations 2.4 and 2.3, but instead relies on AM simulations for every PWV-elevation combination.

operating parameters. In addition, this division is **computationally efficient**, allowing the layers in Figure 5.1 to be simulated from the top down.

In its current form, BoloCalc does not offer any functionality for defining correlated errors. Such a capability would be particularly useful for quantities whose fluctuations are obviously tied together, such as the temperatures of optical elements or the waist factors of different channels on the same detector pixels. However, generalized support for error correlations would require interaction with a large correlation matrix, posing a major development project. To counteract this deficiency, BoloCalc offers a comprehensive parameter sweep functionality, which can be used to calculate **input-output covariances** manually.

5.2.4 Parameter sweeps

What most clearly separates BoloCalc from its predecessor calculators is its **parameter sweep** functionality. When defining instrument specifications, it is often critical to understand the derivatives of NET vs. various instrument performance metrics. The scope of useful parameter-sweep studies is quite large, and therefore BoloCalc handles a wide variety of inquiries such as NET vs. elevation, photon NEP vs. primary spillover, and readout NEP vs. saturation power.

BoloCalc enables the user to sweep any set of parameters across any set of channels. One can either calculate all possible permutations of input-parameter lists or iterate over inputs in specified groups.¹⁰ The input lists can be defined either with a linear spacing or via a custom input file. Custom lists enable the user to investigate more specific dependencies using external codes, such as, for example, NEP_{ph} vs. alumina/silicon/sapphire loss tangent (see Figure 7.6). MC simulations can also be carried out for each input-parameter combination, allowing the user to inspect not only how parameter sweeps change the median NET but also how they change the *shape* of the NET distribution. An example MC-generated parameter variation is shown in Figure 5.7.

Parameter variations, especially when combined with MC simulations, can be **computationally expensive**, and therefore BoloCalc’s sweeps must be performed efficiently. To do this, an “initial” set of MC-generated experiment realizations is constructed, and as each parameter is changed, only the affected instrument aspects are resimulated. For example, if the user changes the detector optical efficiency of a 90 GHz bolometer in “Camera 1,” only the MC realizations of the 90 GHz detectors in “Camera 1” are recalculated, while other channels (e.g., 150 GHz), the telescope optics, and all other cameras and telescopes are left untouched. However, if the spillover fraction of the telescope’s primary mirror is varied, then all of the cameras and detectors in that telescope are recalculated, while other telescopes are left untouched. This top-down variation scheme enables efficient iteration over arbitrary parameter combinations, which is especially powerful when simulating thousands of instrument realizations.

¹⁰For example, when investigating the relative impact of cryostat temperature and window load on P_{opt} , one can sweep lens, filter, and stop temperatures together in one group and window emissivity, scattering fraction, and spillover in another group, generating a two-dimensional output.




Figure 5.7: A BoloCalc sweep of detector NEP vs. detector optical efficiency for two P_{sat} values. This simulation includes 3,000 realizations of an arbitrarily constructed example experiment. The blue and orange contours show the distributions’ shapes for each parameter combination, and the central box-and-whisker plot quantifies the combined distribution. This double-sided violin plot demonstrates how MC-generated datasets for each parameter-sweep combination can be extracted and manipulated using BoloCalc’s `output.py` class.

The parameter-sweep output files are constructed in a dedicated directory for user-friendly inspection, and BoloCalc’s `output.py` class loads these files into dictionaries for easy plotting. In addition, BCG’s `analyze` tool allows the user to plot any output parameter vs. an arbitrary combination of input parameters, further facilitating quick and generalized feedback to instrument designers.

5.2.5 Custom bandpasses

BoloCalc offers two methods to handle bands—or input parameters that vary with frequency—such as those of absorptivity, transmissivity, reflectivity, spillover, and scattering fraction. The first method is to define a **trapezoidal band**

$$B(\nu) = \begin{cases} 0 & \text{if } \nu < \nu_{\text{lo}} \\ \alpha(\nu) & \text{if } \nu_{\text{lo}} \leq \nu \leq \nu_{\text{hi}} \\ 0 & \text{if } \nu > \nu_{\text{hi}} \end{cases} \quad (5.3)$$

where $\alpha(\nu)$ is, in general, a function of frequency, and where the band edges ν_{lo} and ν_{hi} mark discontinuities. The most common trapezoidal band is when $\alpha(\nu) = \text{const.}$, defining what is

called a **top-hat band**. The top-hat band definition is convenient because it is defined only by a band center ν_c and fractional bandwidth f_{BW}

$$\begin{aligned}\nu_{\text{lo}} &= \nu_c - \frac{\nu_c f_{\text{BW}}}{2} \\ \nu_{\text{hi}} &= \nu_c + \frac{\nu_c f_{\text{BW}}}{2},\end{aligned}\tag{5.4}$$

making an estimate of band-averaged quantities algebraically trivial.

In reality, bands have nontrivial spectral features that arise due to any number of physical factors. For example, Fabry-Pérot fringes generate peaks/troughs in transmission/reflection due to constructive/destructive interference within a dielectric optical element. As another example, ripples develop in detector bands due to internal reflections within on-chip filter banks or between the antenna and the bolometer (see Section 2.4). Also, realistic band edges have a finite roll-off, and the steepness of this roll-off can be important for understanding the impact of, for example, atmospheric emission on P_{opt} (see Figure 2.1). For these reasons and more, BoloCalc accepts **custom band files** with a mean and standard deviation defined over an array of frequencies (see, for example, Figure 5.5). BoloCalc also allows the spectral locations of these custom bands to be either swept or assigned an error bar during MC simulations. **Band-center uncertainty** is a major consideration during detector fabrication [224, 197], and therefore BoloCalc’s band-placement features have been valuable for detector characterization within Simons Array (SA) and SO.

The **frequency resolution** of any BoloCalc simulation can be tuned to $\Delta f \geq 0.1$ GHz, providing a lever to optimize computation time vs. accuracy. For example, if an optic’s band is measured with high frequency resolution, the user might run a BoloCalc simulation with small Δf to best capture the measured data. In the other limit, if the user inputs minimalist bands (such as top-hats) but wants to run a million MC realizations, they can increase Δf to maintain a reasonable computation time.

5.3 Calculator validation

In 2014, BoloCalc was conceived to forecast the performance of LiteBIRD’s detector arrays. To ensure its accuracy, the sensitivity code was checked against that of LiteBIRD’s Japanese team (Tomotake Matsumura), and all outputs were confirmed to be identical for a variety of input assumptions. During that same period, several sets of parameter sweeps were compared against legacy codes, such as those of POLARBear-2 (Aritoki Suzuki) [197] and SPT-SZ (Nils Halverson), to ensure that BoloCalc’s mapping-speed limits and scalings are consistent with past calculators.

In 2016, BoloCalc was further compared against sensitivity codes from ACT (Suzanne Staggs and Sarah Marie Bruno) to ensure that mapping-speed scalings were consistent across various sets of input parameters. This cross-check was especially critical during SO’s early stages of optimizing mapping speed for a plethora of instrument configurations, some of

which are discussed in Section 5.4. Shortly thereafter, as SO’s instrument design and science forecasts matured [215], BoloCalc models were built for POLARBEAR, ACTpol, and ABS, and each experiment’s published array NETs were compared against BoloCalc’s prediction. This comparison covered a wide variety of detector, optical, and telescope infrastructures, and in each case, the achieved array NET was statistically consistent with BoloCalc’s expectation.¹¹ This final comparison demonstrates that the calculator is detailed and accurate enough to predict and characterize actual instruments.

5.4 Informing the design of SO

One of BoloCalc’s primary applications has been to inform the design and development of SO. During its earliest stages, SO contemplated menagerie of instrument designs, and in order to make progress, it needed an efficient feedback loop between the science team—which uses map depths to set cosmological constraints—and the instrument team—which translates science targets into observatory requirements. BoloCalc was a crucial “glue” that bound these discussions together, and in this section, we discuss some of SO’s most important BoloCalc studies [81].

SO’s design is discussed as part of Chapter 2, but we review the basics here before diving into technical details [59]. SO uses dichroic pixels and distributes its frequency bands between 27/39 GHz “low-frequency” (LF) cameras, 90/150 GHz “mid-frequency” (MF) cameras, and 220/270 GHz “ultra-high-frequency” (UHF) cameras. Lenslet-coupled sinuous antennas are baselined for the LF wafers, while a feedhorn + OMT architecture is baselined for the MF and UHF wafers. The detector wafers are distributed between cameras in the large aperture telescope (LAT), each of which houses three wafers,¹² and small aperture telescopes (SATs), each of which houses seven wafers.

5.4.1 LATR architecture

The LAT is the vehicle for SO’s high- ℓ science, and its goal is to achieve arcminute resolution in 6 frequency bands from 30 GHz to 280 GHz. To achieve this goal, both the telescope and receiver need to be quite large and are therefore major cost drivers for SO. For this reason and others, the LAT system is designed from the *top-down*, with the telescope’s design driving that of the large aperture telescope receiver (LATR), which drives the design of the optics tubes, which drives the design of the detector arrays. Of course, there is iteration and interaction among the teams designing each subsystem, but for the LAT discussions to follow, larger elements are held constant while smaller elements are varied.

The LAT has a 7.8° field of view (FOV) and illuminates a 2.5 m-diameter area at the telescope’s focus, where the LATR then reimages onto cryogenic detector arrays. An early

¹¹We plan to present these findings in a future publication, and hence we avoid detailing the assumptions and results in this dissertation.

¹²The LAT cameras have a FOV that can accommodate up to 4 wafers, allowing for future upgrades.

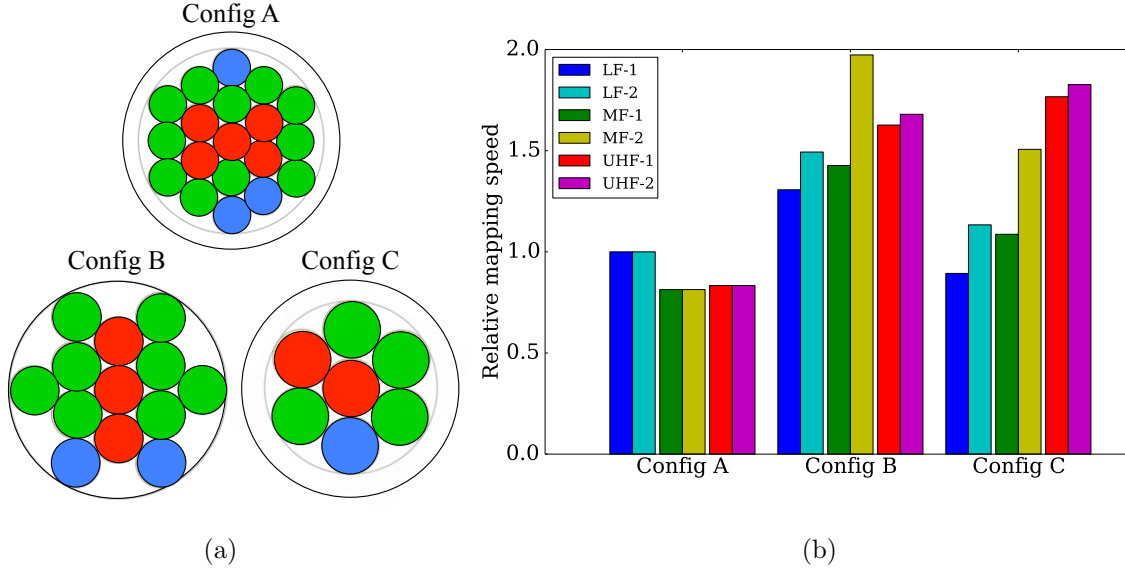


Figure 5.8: Figure 5.8a shows three different optics-tube packings that were investigated for the LATR. Config. A has a maximum of 19 cameras with one wafer per camera, Config. B has 13 cameras with four wafers per camera, and Config. C has 7 cameras with seven wafers per camera. Red circles represent UHF cameras, green MF, and blue LF. The outer-most black circle shows the maximum FOV offered by the telescope. Configs. A and C do not fill the available telescope FOV due to limitations on the magnification and image quality of their cameras’ reimaging optics. Figure 5.8b shows the relative MS of each configuration for all six SO bands. Config. B, when fully filled, is favored in all but the UHF bands.

application of BoloCalc was to help determine how many optics tubes to deploy in the LATR, and while this question involves a complex mixture of telescope optics, cryo-vacuum engineering, mechanical design and assembly, reimaging optics, cost, and upgradability [59, 60, 240, 44], sensitivity was central to the decision making process.

Figure 5.8 shows a study of three different optics-tube configurations that were considered for the LATR. This particular investigation imposes a fixed telescope FOV and a hexagonal camera packing. Config. A has 19 “small-diameter” cameras with a maximum of one wafer per camera, Config. B has 13 “medium-diameter” cameras with a maximum of four wafers per camera, and Config. C has 7 “large-diameter” cameras with a maximum of seven wafers per camera. In this study, pixel size was held constant, and f-number ($F/\#$) was varied with plate scale to maximize each camera’s FOV. Each LATR configuration impacts its cameras’ subsystems (e.g., detectors, refractive optics, cryogenics, etc.) differently, but Figure 5.8b shows the relative maximum-achievable¹³ mapping speed (MS) for each arrangement across all six frequency bands. Config. B has the best overall MS at LF and MF, with only a marginal MS degradation at UHF. Driven in part by this MS calculation, Config. B was chosen as the baseline LATR architecture [240].

¹³Here, the “maximum” is achieved by fully filling each optics tube with detector wafers.




Figure 5.9: Relative MS vs. camera $F/\#$ in each frequency band in the LAT (Figure 5.9a) and the SAT (Figure 5.9b). Given a fixed FOV and pixel size, smaller $F/\#$ improves MS, but the impact is greater at lower frequency and in the SAT. The curves for each frequency channel are individually peak-normalized.

5.4.2 Camera magnification

Another important application of BoloCalc within SO was to assess the impact of camera magnification on MS, which BoloCalc parameterizes using the $F/\#$ at the focal plane. For a fixed FOV and pixel size, a smaller $F/\#$ leads to larger spillover efficiency at the cold aperture stop (Equation 2.10) and thus better sensitivity. However, if the $F/\#$ is *too* small, the **Strehl ratio**¹⁴ will begin to degrade at the edges of the focal plane, leading to fewer detectors with acceptable image quality. Therefore, it is useful to understand how rapidly MS varies with $F/\#$ to inform the sensitivity impact of various reimaging optics designs.

The degree to which a smaller $F/\#$ improves MS depends on frequency and on the optical configuration. Figure 5.9 shows the $F/\#$'s impact on MS in both the LAT and the SAT given a fixed FOV and pixel size for the LF, MF, and UHF cameras. This calculation is combined with **ray-trace simulations** to evaluate the performance of various camera-optics designs [44]. The impact of $F/\#$ on MS is most pronounced at low frequencies, where stop spillover and photon correlations are largest. In contrast, the impact of $F/\#$ on sensitivity is minimal at high frequencies where stop spillover and photon correlations tend to be small. Additionally, increasing per-detector throughput via a smaller $F/\#$ most benefits the SAT, which does not suffer from parasitic loading due to ambient-temperature mirrors. Section 5.4.4 has more details on the impact of warm spillover on the LAT's MS.

¹⁴The Strehl ratio is sometimes called the “phase-front error,” as it measures the degree to which the imaged radiation’s phase varies at the detection point. In a perfect telescope, the image will be fully coherent across focal plane coupling element’s diffraction-limited aperture (see Section 2.3.3).




Figure 5.10: Relative MS in each frequency band in the LAT (Figure 5.10a) and the SAT (Figure 5.10b) vs. pixel pitch, plotted in units of $F\lambda$. Smaller pixels are strongly favored until $D \sim 1.2 F\lambda$, at which point the MS curve flattens due to detector-to-detector photon correlations. The curves for each frequency channel are individually peak-normalized.

5.4.3 Pixel pitch

Another important input to instrument design is the pixel pitch on the focal plane.¹⁵ As described in Section 2.3.4, CMB detectors are single-moded and diffraction-limited, and the aperture stop spillover efficiency can be approximated as a Gaussian function

$$\eta_{\text{apert}} = 1 - \exp \left[-\frac{\pi^2}{2} \left(\frac{D}{\lambda F w_f} \right)^2 \right], \quad (5.5)$$

where D is the pixel pitch (or equivalently, the pixel diameter), F is the $F/\#$ at the focal plane, λ is the observation wavelength, and $w_f = D/w_0$, where w_0 is the beam waist. Smaller pixels have lower efficiency through the stop but allow for denser detector packing. Therefore, for a fixed FOV, there exists a pixel pitch that maximizes MS at each frequency.

Figure 5.10 shows MS vs. pixel pitch, plotted in units of $F\lambda$, for all SO frequencies in the LAT and SAT. In general, smaller pixels give higher MS, as the MS gain due to increased pixel density is faster than the MS degradation due to reduced stop spillover efficiency. The optimal $F\lambda$ spacing is largest in the LF camera, as photon noise correlations (see Chapter 4) are strongest at low frequencies where the photon occupation number is large. Additionally, the optimal $F\lambda$ pixel pitch in the LAT is smaller than that of the SAT, as the LAT has more parasitic ambient-temperature loading (see Section 5.4.4). We also confirmed that the trends found from assuming Gaussian beams (Equation 5.5) were consistent with the integrated η_{apert} from full-beam simulations¹⁶ of prototype feedhorn and lenslet designs at multiple pixel sizes [183, 182].

¹⁵This section describes an SO application of the focal plane optimization discussion in Chapter 4.

¹⁶Ansys HFSS: <https://www.ansys.com/products/electronics/ansys-hfss>




Figure 5.11: The relative optical power (Figure 5.11a) and MS (Figure 5.11b) in each LAT frequency band vs. primary spillover fraction. Lower spillover is always better for sensitivity, but the MS impact is more pronounced at low frequencies. The curves for each frequency channel are individually peak-normalized.

There are many considerations when choosing pixel pitch, including systematic effects that arise when observing with electrically small antennas [38], diffraction artifacts from aggressive aperture truncation [60], readout multiplexing factor [46], wire-bond density [119], achievable saturation power [109], feedhorn and/or lenslet fabrication limitations [182], cost, etc. Therefore, smaller pixel sizes are not always favorable when accounting for these extraneous considerations, but understanding the MS differences between various focal plane layouts was a critical input to the pixel pitch study.

5.4.4 LAT primary spillover

The LAT has a 6 m-diameter primary mirror, a 7.8 degree FOV, and up to 13 cameras in a 2.4 m-diameter receiver that cover more than a decade in frequency [59, 240, 44]. Given the immensity and complexity of the LAT system, ambient-temperature spillover and scattering are major concerns [60], as minimizing optical loading is critical to maintaining low photon NEP.

Figure 5.11 shows relative in-band optical power and MS in each frequency band as a function of LAT primary spillover. The impact of primary spillover on both optical power and MS is largest at low frequencies, where loading due to other sources—such as atmospheric emission and camera thermal emission—is low and the optical efficiency—determined by the absorptivity of the lenses, filters, and on-chip transmission lines—is high. Additionally, the impact of primary spillover on MS is steep, making LAT optical simulations and baffling design a top priority within SO.




Figure 5.12: Relative optical power (Figure 5.12a) and MS (Figure 5.12b) in each frequency band vs. SAT aperture stop temperature. A lower stop temperature is always better for sensitivity, but the impact tends to be larger at low frequencies, where other sources of parasitic loading are small, and in the low band of each dichroic pixel, where the detector beams are largest (see Equation 5.5). The curves for each frequency channel are individually peak-normalized.

5.4.5 SAT stop temperature

Because the SAT does not suffer from loading due to ambient-temperature mirrors, the most important contribution to its photon NEP is its aperture stop temperature. Therefore, understanding how stop temperature impacts MS is critical to setting its cooling requirement [59]. The SAT’s stop is cooled by the 1 K stage of the dilution refrigerator and is connected to that stage via a long heat strap. Sky-side IR radiation is filtered out by a combination of IR shaders, reflective low-pass filters, and alumina absorbing filters (see Section 2.5). Therefore, BoloCalc’s calculation of NET vs. stop temperature helped inform the telescope’s thermal design and heat strapping infrastructure.

Figure 5.12 shows relative in-band optical power and MS as a function of stop temperature for all SAT frequency bands. The impact of stop temperature is largest at low frequencies where other sources of parasitic loading—such as the atmosphere and lens emission—are insignificant. Additionally, the stop’s impact is more dramatic in the low band of each dichroic pixel, which has a lower stop spillover efficiency due to a smaller $D/F\lambda$ pixel size (see Equation 5.5). Finally, the stop’s impact at high frequencies is negligible because the SAT UHF pixels are electrically large and therefore spill little power at the aperture.

5.4.6 Detector saturation power

In addition to characterizing sources of optical loading, BoloCalc can also be used to investigate the impact of detector parameters on sensitivity, including operation temperature, thermal conductivity to the bath, and on-chip optical efficiency. Such calculations can be used to set **detector fabrication tolerances** and provide evaluation criteria for detector




Figure 5.13: Relative MS for each frequency band in the LAT (5.13a) and the SAT (5.13b) vs. saturation power P_{sat} , plotted as a fraction of optical power P_{opt} . P_{sat} impacts MS most at low frequencies, where NEP_{ph} is smallest, and the LAT and SAT impacts are similar. The curves for each frequency channel are individually peak-normalized.

testing and quality control. Figure 5.13 shows relative MS vs. bolometer saturation power P_{sat} , plotted as a fraction of optical power P_{opt} , for both the LAT and SAT in each frequency band. The lowest possible value for $P_{\text{sat}}/P_{\text{opt}}$ is 1, corresponding to zero voltage bias across the bolometer, and the typical range of values that ensure **linear bolometer operation** is $P_{\text{sat}}/P_{\text{opt}} \sim 2\text{--}3$.

Selecting a P_{sat} depends on several considerations, including detector linearity [38], the stability of observing conditions, and uncertainty in the expected telescope-induced loading. However, in the limit of large loop gain (see Equation 2.17), lower saturation power is always better for sensitivity. The impact of $P_{\text{sat}}/P_{\text{opt}}$ is largest in the LF bands where NEP_{ph} is smallest and hence where NEP_{g} makes the largest contribution to NET.¹⁷ As SO has matured, BoloCalc has continued to play a key role in connecting the detector and optical designs, as an accurate calculation of optical power P_{opt} is critical to setting detector parameters. Additionally, as SO detectors begin to undergo evaluation [192] and as uncertainty in the expected optical loading decreases, BoloCalc is increasingly useful for tuning detector performance to maximize MS.

5.4.7 Scan strategy

Another BoloCalc application is to quantify MS trade-offs of observing at different elevations and PWVs and across different sky-patch sizes. Such calculations can be used to optimize the map depth achieved using various **scan strategies**. BoloCalc utilizes atmospheric simulations of the Cerro Toco Atacama observation site that are consistent with measured sky

¹⁷The impact of $P_{\text{sat}}/P_{\text{opt}}$ is similar in the LAT and SAT because $\text{NEP}_{\text{g}} \propto \sqrt{P_{\text{sat}}/P_{\text{opt}}}/\text{NEP}_{\text{ph}}$ (see Equations 3.23 and 3.46), modulated only by wave noise, which is similar in both telescopes.




Figure 5.14: Normalized MS vs. PWV and elevation at the Cerro Toco Atacama observation site in the LAT’s 90 (5.14a) and 150 GHz (5.14b) bands. The impact of elevation is more prominent in the low band, while the impact of water is more prominent in the high band. Additionally, the gradient of MS vs. sky conditions is larger at higher frequency. Such plots are easy to generate using BoloCalc and are useful inputs to scan strategy optimization.

temperatures in existing Atacama experiments (see Section 5.2.2). Figure 5.14 shows the LAT’s normalized MS vs. PWV and elevation in its 90 and 150 GHz bands. The impact of elevation is more prominent in low band, while the impact of water is more prominent in the high band (see Figure 2.1). Additionally, the gradient of MS vs. sky conditions is larger at higher frequency.

5.5 Informing PB-2c Psats

BoloCalc has also been used extensively to inform the design and characterization of SA, which includes three POLARBEAR-2 telescopes (see Section 2.1.2): PB-2a and PB-2b observing at 90 and 150 GHz, and PB-2c observing at 220 and 270 GHz. Unlike SO, SA was quite mature by the time of BoloCalc’s conception. Therefore, SA has relied on the calculator for lower-level analyses, including evaluations of anti-reflection coatings, readout pathologies, and detector bandpasses. In this section, we highlight BoloCalc’s capability to inform detailed hardware considerations by presenting a sensitivity study for PB-2c’s saturation power targets.

As described in Section 2.4, saturation power P_{sat} is the total power conducted from the transition-edge sensor (TES) to the thermal bath to which it is weakly connected. P_{sat} is tuned during detector fabrication (primarily) by tuning the thermal link’s geometry (see Equation 3.48), or specifically its A/l ratio. P_{sat} needs to be large enough to avoid detector saturation but small enough to avoid excess P_{bias} , which leads to excess readout noise (see

Equations 3.54).¹⁸ In other words, P_{sat} should not be too much larger or smaller than the expected $P_{\text{opt}} + P_{\text{bias}}$.

The detector's optical load depends on a plethora of factors, including telescope temperature, sky temperature (see Equation 3.45), telescope optical efficiency, and detector efficiency. While measurements of individual telescope elements can be made in the lab, it is typically difficult to estimate telescope temperature precisely until the instrument is operating in the field. Similarly, detector bands are measurable in the lab, but common systematics in Fourier-transform spectroscopy can make an accurate estimate of on-wafer loss difficult. As a result of these realities, the expected P_{opt} often has large uncertainties that need to be accounted for when selecting a P_{sat} target. In addition, the detectors' optical load also depends on weather conditions, which change as a function of PWV and telescope boresight elevation. The probability of a given sky temperature T_{sky} therefore depends on several factors, including the time of year, observation strategy, fridge cycle duration, and day-to-day weather patterns.

The goal of selecting an optimal P_{sat} is one of minimizing N_ℓ (see Section 3.10), which is related to mapping speed as

$$N_\ell \propto \frac{1}{\text{MS } t_{\text{obs}} \eta_{\text{obs}}}, \quad (5.6)$$

where t_{obs} is observation time, η_{obs} is observation efficiency, and MS is mapping speed. Assuming that P_{sat} does not impact t_{obs} ,¹⁹ minimizing N_ℓ is analogous to maximizing $\text{MS} \times \eta_{\text{obs}}$.

There are several competing factors when optimizing saturation power:

1. Higher P_{sat} leads to higher η_{obs} , as the experiment can observe in the presence of larger T_{sky} and hence larger PWV. Conversely, lower P_{sat} will reduce η_{obs} . A cumulative probability distribution of PWV during the second season of POLARBEAR observations (2013-2014) [212] is shown in Figure 2.2, and it demonstrates that PWV is < 2.2 mm 80% of the time and < 4 mm 90% of the time.
2. Higher P_{sat} leads to higher bolometer thermal carrier noise NEP_g and readout noise NEP_{read} (see Equations 3.51 and 3.54), which reduces MS. Conversely, lower P_{sat} leads to higher MS.
3. Higher T_{sky} leads to higher P_{opt} and hence higher photon noise. Therefore, observations at larger PWV are less sensitive and therefore contribute less *weight* to the final, observation-combined map depth (see Equation 3.68).

Given these factors, the problem of optimizing mapping speed is one of estimating P_{opt} and accurately understanding its scaling with T_{sky} .

¹⁸Excess P_{sat} also leads to excess thermal carrier noise (Equation 3.51), but in SA, readout noise is a more prominent concern.

¹⁹Fabricating a fresh batch of detector wafers takes multiple months. Therefore, any P_{sat} adjustment that induces project delays should improve mapping speed enough to compensate for lost observation time.




Figure 5.15: Probability distributions of sky temperature T_{sky} (left) and telescope temperature T_{tel} (right) in PB-2c's 220 (top) and 270 GHz (bottom) bands. A 50 deg telescope boresight elevation is assumed throughout, and therefore the T_{sky} distributions trace the Atacama's PWV distribution shown in Figure 2.2. Variability in telescope temperature is dominated by uncertainties in optical throughput due to lens absorption, lens reflection, and on-wafer detector loss.

5.5.1 Estimating P_{opt}

According to Equation 3.45, there are two contributions to optical power on the detector P_{opt} : telescope temperature and sky temperature. Telescope temperature has a distribution that is governed by uncertainties in instrument characterization. In other words, once the telescope is observing in the field, its temperature is relatively constant. Sky temperature, on the other hand, follows a distribution that is governed by weather patterns and scan strategy. Throughout this study, we assume a constant elevation of 50 deg above the horizon, allowing us to focus on the impact of PWV variations alone. Probability densities for PB-2c's telescope




Figure 5.16: Optical power probabilities for both PB-2c bands. Figure 5.16a shows a cumulative distribution of the expected optical power at 220 GHz given uncertainties in dewar temperature, sky power, and detector bandpass and assumes a detector efficiency of 60%. Figure 5.16b shows the same distribution but for 270 GHz.

and sky temperatures are shown in Figure 5.15. The sky-temperature distribution traces that of PWV in Figure 2.2b, while telescope temperatures are reasonably Gaussian with 95% confidence intervals of $\pm 5 \text{ K}_{\text{RJ}}$. Uncertainties in telescope temperature are primarily due to uncertainties in the performance of the 220/270 GHz anti-reflection (AR) coatings.

Given probability distributions for T_{sky} and T_{tel} , we can now use BoloCalc to run thousands of MC realizations of PB-2c observing in various weather conditions. Because $P_{\text{opt}} \propto \eta_{\text{det}}$, we consider four cases of detector optical efficiency: 20, 40, 60, and 80%. This delineation also illuminates the interplay between P_{sat} and η_{det} , which are governed by the detector fabrication process and are often assessed during the same suite of lab tests [224]. The resulting P_{opt} distributions are shown in Figure 5.16. The two bands are very similar, which may seem surprising given that atmospheric emission is larger at 270 GHz than at 220 GHz. However, as suggested by Equation 3.28, the refractive optics' absorptivities increase with frequency, reducing the telescope's throughput²⁰ and canceling out the brightening atmosphere.

It is worth emphasizing that the distributions in Figure 5.16 fold *uncertainties* in the instrument configuration and *variations* in sky load together. Therefore, this plot represents P_{opt} when deploying an MC-generated instrument to observe an MC-generated sky. Another approach would be to handle the T_{sky} and T_{tel} separately in order to apply different observation strategies to different telescope realizations. This approach is reasonable, as the observation strategy can be tuned after studying the instrument in the field. However, such

²⁰In fact, the worst in-band absorbers are PB-2c's alumina lenses, which have an increasing loss tangent with frequency. In fact, this quality makes alumina an appealing material for IR absorbing filters, such as in Figure 2.11a.

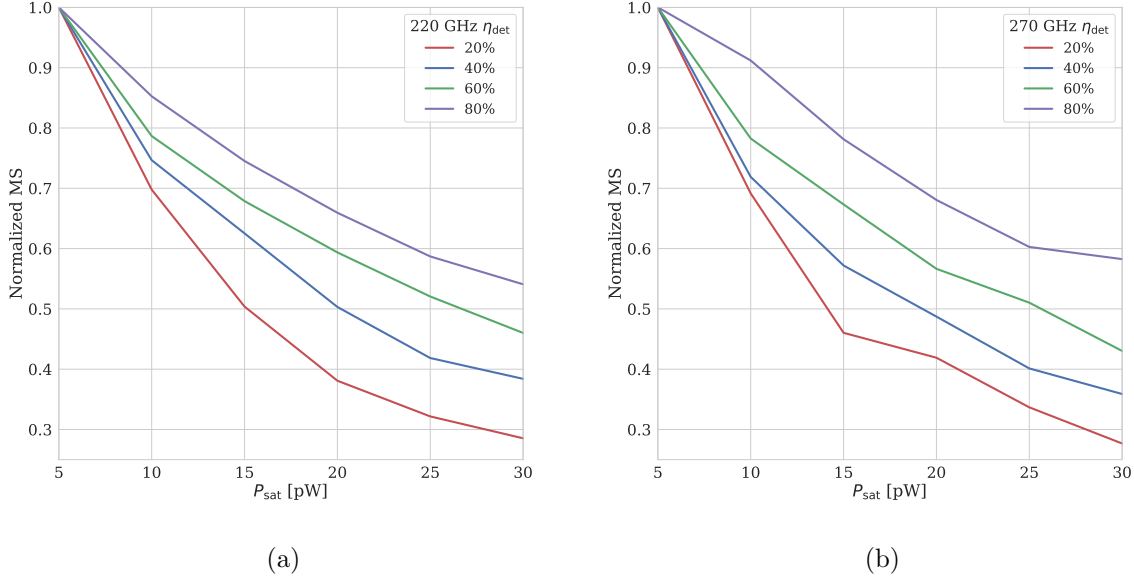


Figure 5.17: Median mapping speed vs. detector saturation power P_{sat} in PB-2c’s 220 (Figure 5.17a) and 270 GHz (Figure 5.17b) bands, normalized to the case of $P_{\text{sat}} = 5$ pW. The impact of P_{sat} is more pronounced when detector efficiency η_{det} is small, as the relative contributions of bolometer thermal carrier noise NEP_{g} and readout noise NEP_{read} to the total detector NEP increases with decreasing photon noise NEP_{ph} (see Equation 3.57).

an analysis involves too many details and independent variables for the PB-2c discussion, which is meant to quickly and straightforwardly feed back to detector fabrication efforts. That said, BoloCalc is fully capable of running arbitrary parameter separations, and a more detailed study may be interesting for other maturing experiments such as SO.

5.5.2 Maximizing $\text{MS} \times \eta_{\text{obs}}$

Armed with the optical-power distributions in Figure 5.16, we now calculate mapping speed as a function of both detector efficiency η_{det} and saturation power P_{sat} . Figure 5.17 shows the peak-normalized *median* mapping speed vs. P_{sat} in both the 220 and 270 GHz bands. As expected, mapping speed decreases with increasing saturation power, and the impact is larger when detector efficiency η_{det} is smaller. Because the P_{opt} distributions in Figure 5.16 are similar, so are each band’s mapping speed vs. saturation power curves.

To quantify the impact of P_{sat} on observation efficiency η_{obs} , we must decide on a **P_{opt} cutoff**, or an optical load above which the telescope will not observe.²¹ The P_{opt} cutoff is

²¹There are practical considerations surrounding weather-driven decisions (e.g., observing in falling snow is ill-advised), but these extraneous factors typically arise at larger PWVs than those considered in this analysis.




Figure 5.18: Peak-normalized mapping speed MS vs. detector saturation power P_{sat} in the PB-2c 220 and 270 GHz bands for various P_{sat} cutoffs P_{cut} (see Equation 5.7) and detector optical efficiencies η_{det} . The impact of deploying bolometers with overly conservative saturation powers is most pronounced at low P_{cut} and high η_{det} .

driven by the detector's **saturation point**, or when $P_{\text{opt}}/P_{\text{bias}}$ is so large that the bolometer becomes non-linear or even unstable. We can define this **cutoff power** in terms of P_{sat} as

$$P_{\text{cut}} \equiv \text{MIN} \left[\frac{P_{\text{sat}}}{P_{\text{opt}}} \right], \quad (5.7)$$

where P_{cut} determines η_{obs} via the cumulative distributions in Figure 5.16. P_{cut} is analogous to a **PWV cutoff**, which corresponds to a sky temperature T_{sky} above which the instrument does not observe.

The cutoff power P_{cut} is usually determined after measuring how detector responsivity scales with P_{bias} in the field. For example, if $dR_{\text{bolo}}/dT_{\text{bolo}}$ is shallow, then the bolometer's loop gain $\mathcal{L} \propto P_{\text{bias}}$, and maintaining a large $P_{\text{bias}}/P_{\text{opt}}$ ratio could be important for **detector linearity**. For these reasons, the following analysis is performed for various P_{cut} values to illuminate possible $MS \times \eta_{\text{obs}}$ scenarios in the field.

Figure 5.18 shows $MS \times \eta_{\text{obs}}$ vs. P_{sat} for various η_{det} and P_{cut} . There are several features of the $MS \times \eta_{\text{obs}}$ optimization that are worth highlighting. First, lower P_{sat} is favored for lower P_{cut} . This behavior is expected given the proportionality relation in Equation 5.7. Second, the optimal P_{sat} is lower for lower η_{det} . This behavior is also expected, as $P_{\text{opt}} \propto \eta_{\text{det}}$. Third, the mapping-speed penalty of fielding a higher P_{sat} than necessary is less pronounced at higher P_{cut} and higher η_{det} . This behavior arises because the total NEP becomes increasingly photon-noise dominated with increasing η_{det} , and when NEP_{ph} is substantially larger than NEP_{g} and NEP_{read} , the importance of P_{sat} on NET is suppressed. Fourth, along a similar vein, a larger P_{cut} allows the experiment to observe in more T_{sky} scenarios, allowing a broader




Figure 5.19: The distribution of bands and sensitivities for an early-stage incarnation (2016) of the LiteBIRD focal plane, as calculated by BoloCalc. The lower 12 bands are located in the low-frequency telescope (LFT), and the top three bands are located in the high-frequency telescope (HFT). Each detector pixel in the LFT is trichroic and is labeled with a distinct color. The y -axis is in units of K_{CMB} , and the frequency dependences of synchrotron and dust emission are plotted with amplitudes that highlight the foreground minimum at ≈ 70 GHz. The carbon-monoxide (CO) emission lines are only slightly polarized but are very bright and therefore are relegated to one band per line. Since this figure was made, the distribution of bands between telescopes and the projected sensitivities have changed, but the importance of this initial study resonates throughout LiteBIRD’s instrument design and mission concept.

range of P_{sat} values to be “viable.” This phenomenon is most clearly seen when contrasting the highest and lowest P_{cut} values. $P_{\text{cut}} = 2$ asserts that the bolometer must satisfy $P_{\text{bias}} \geq P_{\text{opt}}$, and in this case, for all η_{det} , if the $P_{\text{sat}} \leq 10$ pW, the detector is inoperable in *any* sky-loading scenario. However, if the same bolometers are fielded with $P_{\text{cut}} = 1.25$, then bolometers with $P_{\text{sat}} = 10$ pW become optimal in the $\eta_{\text{det}} = (40, 60)\%$ cases.

The presented study was critical to assigning P_{sat} targets for PB-2c’s bolometers and assessing the viability of its fabricated detectors. In addition, this study is a useful demonstration of how BoloCalc can be used to optimize the detector array during late-stage instrument development.

5.6 Other applications

In this chapter, we have presented the design and application of BoloCalc within SA and SO. However, the sensitivity calculator’s impact has been much larger than the presented research. BoloCalc was originally conceived and first used to inform the detector array design for LiteBIRD, a Japanese-funded next-generation CMB satellite experiment whose mission

is to map the entire sky from 40 to 400 GHz [95, 195]. In order to characterize foregrounds with unprecedented precision, LiteBIRD is deploying 15 observation bands within multiple telescopes. Therefore, optimizing the satellite’s focal plane layouts was a major undertaking in the project’s early stages. BoloCalc was at the heart of these efforts, generating a new “sensitivity memo” every few weeks as the instrument design rapidly evolved. Given this experience, BoloCalc has demonstrated applicability to CMB satellite missions and could continue to be useful for upcoming space-based observatories. Figure 5.19 shows a BoloCalc-generated sensitivity plot for the initial US LiteBIRD proposal to NASA in 2016.

BoloCalc’s ultimate goal is to help inform the upcoming Department of Energy-funded experiment **CMB-Stage 4 (CMB-S4)** [204], which will turn hundreds of millions of dollars into the ultimate ground-based measurement of CMB temperature and polarization anisotropies. Similar to SO, CMB-S4 will bring many institutions together, and a standardized, user-friendly sensitivity code will be essential to evaluating and advancing the observatory’s design. Therefore, we anticipate that BoloCalc’s impact will reach into the next generation of CMB experimentation and beyond.

Chapter 6

Half-wave plate polarization modulation

The quest to measure primordial B-modes requires CMB instrumentation not only with unprecedented sensitivity but also with tight control of systematic effects. One such **systematic** for ground-based experiments is **low-frequency noise** (also called **1/f noise** to indicate increasing noise with decreasing frequency) in the detector data caused by fluctuations in atmospheric intensity due to wind, clouds, and other weather effects. Even though these fluctuations are *unpolarized* [71, 189], they are difficult to reject at the level needed to measure B-modes, especially on large angular scales where long-baseline (and therefore long-duration) scans are needed. Section 6.1 discusses atmospheric noise in more detail. Another such systematic is beam distortion caused by telescope optics, focal plane optics, and imperfect detector calibration [179]. These **beam systematics** erroneously convert temperature fluctuations into polarization fluctuations and are also difficult to suppress below the CMB polarization signal. Section 6.2 discusses polarized beam systematics in more detail.

There are a plethora of hardware and analysis remedies to mitigate 1/f noise and beam systematics. One such method is **detector differencing**, which extracts polarization by differencing two orthogonal polarimeters as

$$Q = \langle |E_x|^2 \rangle - \langle |E_y|^2 \rangle . \quad (6.1)$$

Here, $|E_x|^2$ and $|E_y|^2$ are the intensities along the x and y directions, the angle brackets denote statistical averages, and Q is one of the **Stokes parameters** presented in Section 6.3.1. Detector differencing works because *unpolarized* fluctuations modulate $|E_x|^2$ and $|E_y|^2$ together, while *polarized* fluctuations do not. Therefore, Equation 6.1, in theory, nulls the intensity while revealing polarization. An idealized example of detector differencing in both the temporal and frequency domains is shown in Figure 6.1. However, in practice, differencing $|E_x|^2$ and $|E_y|^2$ leaves an intensity residual, contaminating the polarization data with so-called **intensity-to-polarization leakage (I-to-P)** (also called **temperature-to-polarization leakage (T-to-P)**). Because the atmosphere is so much brighter than CMB polarization, even minuscule residuals lead to challenging levels of I-to-P contamination.




Figure 6.1: An idealized simulation showing how detector differencing rejects unpolarized atmospheric fluctuations. The left panels show “dummy” time-ordered data (TOD), and the right panels show the data’s power spectral densities (PSDs). In this toy simulation, the x and y polarimeters measure white noise, which does not correlate between the two, and atmospheric $1/f$ noise, which does. After the two detectors are differenced as $x - y$ to extract polarization, only the white noise is left over, and the effective $1/f$ knee in polarization is reduced from ~ 1 Hz to $\lesssim 10^{-2}$ Hz. While $1/f$ rejection is excellent in this idealized case, in practice, uncorrelated $1/f$ noise develops between detectors. If these drifts are not accounted for during detector differencing, residual $1/f$ noise pushes the $x - y$ spectrum’s $1/f$ knee to higher frequencies.

While detector differencing is a common technique at both the Atacama and South Pole observation sites, the need for tighter control of systematics motivates alternative methods for polarization extraction.

Another popular technique to reject unpolarized atmospheric fluctuations and improve polarization sensitivity on long time scales is to employ **continuous polarization modulation**. A **polarization modulator**’s essential function is to modulate the celestial po-

larization signal while *not* modulating the atmospheric intensity signal. To illustrate this idea, consider a polarization rotator oscillating at $f_{\text{mod}} = \omega_{\text{mod}}/2\pi$, which modifies a single polarimeter’s time-ordered data (TOD) $d_m(t)$ as

$$\begin{aligned} d_m(t) &\propto \langle |E_x(t)|^2 \rangle \cos^2 [\omega_{\text{mod}} t] + \langle |E_y(t)|^2 \rangle \sin^2 [\omega_{\text{mod}} t] \\ &\propto (\langle |E_x(t)|^2 \rangle + \langle |E_y(t)|^2 \rangle) + (\langle |E_x(t)|^2 \rangle - \langle |E_y(t)|^2 \rangle) \cos [2\omega_{\text{mod}} t] \\ &\propto I(t) + Q(t) \cos [2\omega_{\text{mod}} t] . \end{aligned} \quad (6.2)$$

Here, $I(t)$ is the intensity signal, $Q(t)$ is the polarization signal defined in Equation 6.1, and the proportionality operator absorbs any gain or normalization factors. This hardware-based polarization separation technique, quantified by the **modulation function** $\cos [2\omega_{\text{mod}} t]$, has several advantages over analysis-based techniques. First, a single polarimeter can measure both x and y polarization states, and therefore detectors need not be differenced to extract sky polarization. This capability avoids issues associated with **gain mismatch** between orthogonal polarimeters, including mismatched bias parameters, optical properties, and responsivities (see Section 2.4). Second, the polarization signal $\langle |E_x|^2 \rangle - \langle |E_y|^2 \rangle$ is modulated at $2\omega_{\text{mod}} = 4\pi f_{\text{mod}}$, which is a hardware-defined parameter. If f_{mod} is large enough, then the modulated polarization signal $Q(t) \cos [2\omega_{\text{mod}} t]$ will be “faster” than the atmospheric fluctuations $I(t)$, improving signal-to-noise in the polarization channel.

A wide variety of optical polarization modulators have been deployed for CMB observation, and we encourage the curious reader to review the rich modulator literature.¹ In this dissertation, we focus on the **half-wave plate (HWP) polarization modulator**, which is the central topic of Chapters 6-9. HWPs are becoming increasingly popular within the CMB community, and continuously rotating HWPs have been adopted by Simons Array (SA) and Simons Observatory (SO) to enable primordial B-mode measurements at the Atacama observation site. In the following sections, we discuss **sapphire HWPs** for multichroic receivers, formalize their mitigation of $1/f$ noise and beam systematics, and discuss SA’s HWP requirements.

6.1 Atmospheric $1/f$ noise

The first systematic mitigated by a continuously rotating HWP is **atmospheric fluctuation**. As discussed in Section 3.4.7, sky intensity is often expressed as an effective temperature, and as shown in Figure 2.1, a typical sky temperature is $5 \sim 30$ K. In contrast, as shown in Figure 1.12, the CMB’s primordial B-mode polarization is ~ 100 nK, meaning that the target cosmological signal is ~ 50 *million* times fainter than the atmosphere. Even so, ground-based CMB experiments are viable competitors to balloons and satellites because (a) the atmosphere is unpolarized, and (b) the cosmic signal is fixed while the atmosphere varies, allowing atmospheric noise to be **averaged down** over time.

¹ The library of polarization modulators includes half-wave plates [99, 105, 111, 139, 27, 6, 199, 82, 162, 72, 26, 157, 37], variable-delay polarization modulators [33, 73], and Faraday-rotation modulators [141].

There are a few major components to atmospheric noise. First, as discussed in Section 5.2.2, atmospheric mm-wave emission increases photon noise in the detector output. This noise is largely uncorrelated across the focal plane and is best mitigated by observing at a site with dry, thin air. Second, coordinate-dependent fluctuations in atmospheric intensity are generated by spatial variations in water-vapor content, such as those due to clouds. Because these variations change during the course of an observation, they are, in general, averaged down and therefore do not appear as hot and cold spots in the sky maps.² However, because the telescope scans through these atmospheric structures, they generate 1/f noise in the detector data. Third, the atmosphere varies in time, and its fluctuations also follow a 1/f noise spectrum. These final two low-frequency effects are mitigated when using a continuous polarization modulator.

According to Errard *et al.* (2016) [50], atmospheric turbulence at the Atacama site follows a Kolmogorov spectrum $\mathcal{P}_{\text{atm}} \propto 1/k^\alpha$, where k is a three-dimensional wave number and $\alpha = -11/3$. Then, as the telescope scans through the atmosphere’s spatial structure, the detector data follows a spectrum³ of $\mathcal{P}_{\text{det}} \propto f^\alpha$ [50]. In addition to its spectral slope α , atmospheric noise is often quantified by its **1/f knee**, which is the frequency where 1/f noise and white noise are equal and assumes a spectrum

$$N(f) = N^{\text{white}} + N^{\text{red}} \left(\frac{f}{f_{\text{knee}}} \right)^\alpha, \quad (6.3)$$

where f_{knee} is analogous to the ℓ_{knee} Equation 3.71. In fact, the detector’s **audio frequency** f can be roughly related to angular scale ℓ via the telescope’s scan speed f_{scan} (in degrees on the sky⁴ per second) as

$$\ell \sim \frac{f}{f_{\text{scan}}} \times 180^\circ, \quad (6.4)$$

where the \sim acknowledges that the $N(f) \rightarrow N_\ell$ transfer function is complex in practice. Nonetheless, a lower f_{knee} (generally) corresponds to a lower ℓ_{knee} , which in turn corresponds to better low- ℓ sensitivity. Therefore, controlling atmospheric 1/f noise is critical to measuring primordial B-modes between $50 \lesssim \ell \lesssim 150$.

The characteristic length scale of Kolmogorov fluctuations is \sim hundreds of meters, and detectors whose projection through the atmosphere is closer than this distance will experience correlated 1/f noise [50]. The degree to which atmospheric 1/f fluctuations are correlated across the focal plane depends on the telescope’s field of view (FOV) and on whether the

²**Ground pickup**, on the other hand, does have a fixed spatial dependence and therefore is typically subtracted using templates that depend on telescope coordinate.

³Even though this spectrum does not strictly follow a $1/f$ shape with $\alpha = -1$, it is still referred to as “1/f noise.” In other words, the term “1/f noise” is synonymous with “low-frequency noise,” which can have any value for α .

⁴Note here that the scan speed in *sky coordinates* f_{scan} is related to the telescope’s azimuthal rotation frequency in *ground coordinates* f_{tel} as $f_{\text{scan}} \approx f_{\text{tel}} \cos[\theta_{\text{el}}]$, where θ_{el} is the telescope’s boresight elevation above the horizon. As a limiting example, if the telescope is pointing along the z-axis (referred to as zenith) such that $\theta_{\text{el}} = 90$, spinning it in azimuth does not change its sky location.

atmosphere is in the telescope’s far field. The telescope’s **far-field limit** is determined by its **Fraunhofer distance**

$$d_{\text{far}} = \frac{2D_{\text{ap}}^2}{\lambda}, \quad (6.5)$$

where D_{ap} is the telescope’s primary aperture diameter and λ is its observed wavelength. For example, SO’s small aperture telescope (SAT) has a ≈ 0.5 m primary aperture, a $\approx 20^\circ$ FOV, and a far-field limit of ~ 200 m at 90 GHz. In this case, **atmospheric turbulence** appears in the telescope’s far field and has a correlation length smaller than the telescope’s FOV. Therefore, 1/f fluctuations are not correlated across the focal plane and hence average down during detector coaddition. As a counter example, consider an SA telescope, which has a ≈ 2.5 m primary aperture, a $\approx 4^\circ$ FOV, and a far-field limit of $\sim 4,000$ m at 90 GHz. In this case, atmospheric turbulence occurs in the telescope’s near field and is correlated over regions larger than the telescope’s FOV [50]. When atmospheric fluctuations are correlated across the focal plane, they are said to be **common-mode** and do not average down during detector coaddition. Therefore, common-mode atmospheric fluctuations must be suppressed⁵ to maximize low- ℓ sensitivity on SA-style telescopes.

6.2 Polarized beam systematics

The second systematic mitigated by an HWP is I-to-P leakage caused by non-idealities in the telescope optics or by differential optical response between orthogonal polarimeters. These inadvertent conversions from intensity to polarization pose several challenges during data analysis, including CMB temperature fluctuations bleeding into polarization and residual atmospheric 1/f noise after detector differencing (see Section 6.1). Such beam systematics are typically described in terms of their monopole, dipole, and quadrupole moments [179], which are shown in Figure 6.2.

Monopole systematics arise due to **differential gain** or **differential beamwidth** and create a rotationally symmetric residual. Dipole systematics arise from **differential pointing**, where the differenced detectors are not perfectly co-located on the sky. Quadrupole systematics arise from **differential ellipticity**,⁶ where the differenced detectors’ beams have ellipticities⁷ that do not align. While monopole and dipole beam systematics can be removed using beam maps⁸ during analysis, quadrupolar I-to-P is difficult to distinguish from the celestial polarization signal, elevating the importance of beam symmetry during instrument design and evaluation.

⁵*Subtracting* (as opposed to *suppressing*) common-mode fluctuations also removes large-angular-scale CMB signals and therefore degrades low- ℓ sensitivity.

⁶Quadrupolar residuals can also arise from higher-order differential-pointing effects.

⁷CMB detectors are designed to have round beams, but in practice, even the best focal plane optics generate beams with some degree of ellipticity.

⁸Understanding the instrument’s beam is critical to an accurate CMB characterization, and therefore polarized beam mapping is an active research area.

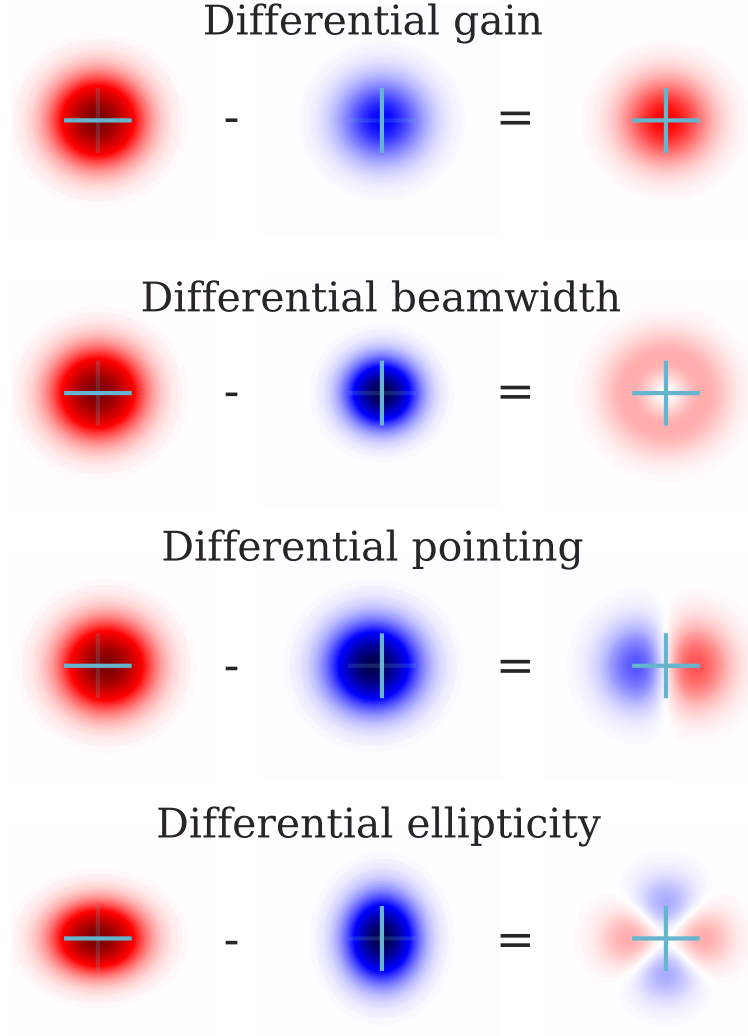


Figure 6.2: A schematic of differential beam systematics. The left (middle) column shows the beam of a horizontally (vertically) oriented polarimeter, and the right column shows their difference. Differential gain and differential beamwidth are monopole effects, while differential pointing and differential ellipticity are dipole and quadrupole effects, respectively.

6.3 Sapphire achromatic HWP

Continuous HWP polarization modulation is a powerful tool to mitigate atmospheric $1/f$ noise and beam systematics. At its most basic level, a continuous HWP spins the sky's polarization vector at a well-monitored frequency, modulating the signal on each polarimeter. This operation allows each detector to measure both polarization states, and because the HWP does not—to leading order—modulate the intensity signal, it eliminates the need for

detector differencing to extract the sky polarization signal.⁹ In addition, if this continuous rotation is fast enough, the HWP up-converts the sky polarization signal to audio frequencies larger than the atmospheric 1/f knee, dramatically suppressing long-time-scale fluctuations in the demodulated detector data. In this section, we review the basics of HWP polarization modulation and discuss its applications to a **sapphire achromatic HWP**. Then, later in the chapter, we discuss the utilities of HWP modulators for CMB polarimetry and how they are being leveraged by SA.

6.3.1 Stokes polarization

Polarization is often decomposed into one of two bases. The first basis, which is perhaps most familiar to physicists, is that of Jones vectors. The **Jones basis** consists of two orthogonal basis vectors (\hat{r}_1, \hat{r}_2) and describes fully polarized light propagating along the \vec{r}_3 direction as

$$\vec{E} = \begin{pmatrix} E_1(t) \\ E_2(t) \end{pmatrix} = \begin{pmatrix} E_{0,1}e^{i\phi_1} \\ E_{0,2}e^{i\phi_2} \end{pmatrix} e^{i(kr_3 - \omega t)}, \quad (6.6)$$

where $E_{0,1}$ and $E_{0,2}$ are wave amplitudes and where ϕ_1 and ϕ_2 are phase factors. In this basis, polarization can be manipulated by a collection of **Jones matrices**, which are used to both rotate linear polarization and convert to/from circular polarization. The subset of these operations most useful to us are

$$\begin{aligned} \begin{pmatrix} E_x \\ 0 \end{pmatrix} &= \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \vec{E} \quad ; \quad \begin{pmatrix} 0 \\ E_y \end{pmatrix} = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} \vec{E} \\ \begin{pmatrix} E_a \\ 0 \end{pmatrix} &= \begin{pmatrix} 1/2 & 1/2 \\ 1/2 & 1/2 \end{pmatrix} \vec{E} \quad ; \quad \begin{pmatrix} 0 \\ E_b \end{pmatrix} = \begin{pmatrix} 1/2 & -1/2 \\ -1/2 & 1/2 \end{pmatrix} \vec{E} \\ \begin{pmatrix} E_l \\ 0 \end{pmatrix} &= \begin{pmatrix} 1/2 & -i/2 \\ i/2 & 1/2 \end{pmatrix} \vec{E} \quad ; \quad \begin{pmatrix} 0 \\ E_r \end{pmatrix} = \begin{pmatrix} 1/2 & i/2 \\ -i/2 & 1/2 \end{pmatrix} \vec{E}, \end{aligned} \quad (6.7)$$

where \vec{E} is 100% linearly polarized in the (\hat{x}, \hat{y}) basis. These three Jones bases are shown as part of Figure 6.3.

While Jones vectors are a familiar, simple description of *fully* polarized light, astronomical sources are generally **partially polarized**, and therefore the bases in Equation 6.7 are inconvenient for CMB applications. Instead, we introduce the **Stokes parameters**, which can be written in terms of the Jones vectors as

$$\begin{aligned} I &\equiv \langle |E_x|^2 \rangle + \langle |E_y|^2 \rangle \\ Q &\equiv \langle |E_x|^2 \rangle - \langle |E_y|^2 \rangle \\ U &\equiv \langle |E_a|^2 \rangle - \langle |E_b|^2 \rangle \\ V &\equiv \langle |E_r|^2 \rangle - \langle |E_l|^2 \rangle \end{aligned} \quad (6.8)$$

⁹The use of an HWP does not preclude detector differencing, however, as analysis techniques can be used to further limit common-mode fluctuations during data processing, if desired.




Figure 6.3: The relationship between Jones vectors and Stokes vectors. Linearly polarized detector pixels for CMB are sensitive to Q and U but not V.

and are shown in Figure 6.3. The Stokes basis has several advantages over the Cartesian basis for CMB measurements. First, the Stokes vectors are written in terms of intensity $|E_i|^2$, which is helpful when using bolometers that detect power (instead of amplitude and phase). Second, they are inherently statistical and measure expectation values, denoted by the angle brackets $\langle \cdot \rangle$, making them well suited to describe noise power. Third, Stokes parameters naturally handle partial polarization, where the **polarization fraction** is defined as

$$P \equiv \frac{\sqrt{Q^2 + U^2 + V^2}}{I}. \quad (6.9)$$

Fourth, the Stokes decomposition is natural for extracting polarization via detector differencing, a common CMB analysis technique (see Figures 6.1 and 6.2). Finally, relating back to Equation 1.36, E-modes and B-modes are constructed in the Stokes basis, and therefore Q/U maps are necessary precursors to E-mode/B-mode maps. While the math of polarization modulators can be equivalently described in either the Jones or Stokes bases, we work in the Stokes paradigm from now on.

6.3.2 Birefringent HWP

A half-wave plate is a **birefringent medium** whose geometry is tuned to introduce a 180° phase shift between light along the medium's **birefringent axes**. The dielectric tensor of such a birefringent material is

$$\varepsilon_{ij} = \begin{pmatrix} \varepsilon_e & 0 & 0 \\ 0 & \varepsilon_o & 0 \\ 0 & 0 & \varepsilon_o \end{pmatrix}, \quad (6.10)$$

where ε_e is medium's dielectric constant¹⁰ along the **extraordinary axis** and where ε_o is its dielectric constant along the **ordinary axes**.

When a wave with an electric field vector $\vec{E} = E_0 \hat{k}$ enters the dielectric medium, it can be decomposed into ordinary and extraordinary waves as

$$\begin{aligned}\vec{E}_1 &= E_1 \hat{e}_1 ; \hat{e}_1 \propto (0, k_z, -k_y) \\ \vec{E}_2 &= E_2 \hat{e}_2 ; \hat{e}_2 \propto (1 - k_x^2, -k_x k_y, -k_x k_z) .\end{aligned}\quad (6.11)$$

The phase velocities of the ordinary and extraordinary waves are

$$\begin{aligned}v_1 &= v_o \\ v_2 &= \sqrt{v_e^2 (1 - k_x^2) + v_o^2 k_x^2} ,\end{aligned}\quad (6.12)$$

where

$$\begin{aligned}v_o &\equiv c/n_o \\ v_e &\equiv c/n_e .\end{aligned}\quad (6.13)$$

In the limit of normal incidence, $\hat{k} = (0, 0, k_z)$, $v_2 = v_e$, and the basis vectors for the ordinary and extraordinary waves defined in Equation 6.11 coincide with the ordinary and extraordinary axes of the birefringent medium $(\vec{e}_1, \vec{e}_2) = (\hat{e}_e, \hat{e}_o)$.

Assume that the birefringent dielectric is cut as in Figure 6.4 such that (\hat{e}_e, \hat{e}_o) are parallel to the surface. Given the plate's thickness d , the generated phase delay between the ordinary and extraordinary waves is

$$\delta\phi = 2\pi \left(\frac{d}{v_e/\nu} - \frac{d}{v_o/\nu} \right) = \frac{2\pi (n_e - n_o) d}{\lambda_0} , \quad (6.14)$$

where λ_0 is the wavelength of the incident light in a vacuum. An ideal HWP has $\delta\phi = \pi$ and therefore has an ideal thickness

$$d_{\text{HWP}} = \frac{\lambda_0}{2(n_e - n_o)} . \quad (6.15)$$

Note that an HWP can only be ideal at a single frequency, and therefore when linearly polarized incident light has a wavelength $\neq \lambda_0$, the HWP does not perfectly preserve its linear polarization.

6.3.3 Mueller matrices

In order to apply the phase-delay discussion in Section 6.3.2 to the Stokes parameters, we utilize the machinery of **Mueller matrices**, which are 4×4 matrix operators that act on the Stokes vector $\vec{S} = (I, Q, U, V)$. We can express a Mueller matrix operation as

$$\vec{S}_{\text{out}} = M \vec{S}_{\text{in}} , \quad (6.16)$$

¹⁰The dielectric constant is also called the “relative permittivity.”




Figure 6.4: A diagram of a perfect HWP, where the phase shift between the ordinary and extraordinary waves is exactly 180° . Note that the optical axis and extraordinary axis are synonymous. This figure is courtesy of Bob Mellish and is taken from the “Waveplate” Wikipedia page [137].

where \vec{S}_{in} and \vec{S}_{out} are the input and output Stokes vectors, respectively. As shown in Figure 6.4, a perfect HWP “flips” the input polarization with respect to the HWP’s extraordinary axis. In terms of the Stokes parameters, this operation is equivalent to inverting the +U polarization component to $-U$ (and $+V$ to $-V$) using the Mueller matrix

$$M_{\text{HWP}}^{\text{ideal}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}, \quad (6.17)$$

where here the extraordinary axis is assumed to be along either the x - or y -axis.

As mentioned earlier, an HWP made from a birefringent substrate is only perfect for a single input frequency: at frequencies increasingly far from λ_0 , incident light is increasingly converted to **elliptical polarization**, which is a mixture of linear and circular polarization. In order to calculate this effect, we introduce a more generalized Mueller matrix formalism for the HWP system. Consider an input Stokes vector

$$\vec{S}_{\text{in}} = I(\nu) \begin{pmatrix} 1 \\ P_{\text{in}} \cos 2\alpha_{\text{in}} \\ P_{\text{in}} \sin 2\alpha_{\text{in}} \\ 0 \end{pmatrix}, \quad (6.18)$$

where $I(\nu)$ is the intensity as a function of microwave frequency ν , $P_{\text{in}} = \sqrt{Q^2 + U^2}/I$ is the input **linear polarization fraction**, and α_{in} is the angle between the input polarization vector and a fixed coordinate system. The effect of the HWP is to introduce a phase delay between the birefringent axes via the operation

$$\vec{S}_{\text{out}} = R(-\chi) \Gamma(\delta) R(\chi) \vec{S}_{\text{in}}, \quad (6.19)$$

where δ is the phase delay defined in Equation 6.14. The first and final operators on S_{in} are rotation matrices

$$R(\Psi) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos 2\Psi & -\sin 2\Psi & 0 \\ 0 & \sin 2\Psi & \cos 2\Psi & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (6.20)$$

and are used to rotate the Stokes vector into and out of the HWP's dielectric coordinate system $\hat{\epsilon}_i = (\hat{\epsilon}_e, \hat{\epsilon}_o, \hat{\epsilon}_o)$. Once in the frame of the ordinary and extraordinary axes, the phase delay operator modifies the Stokes vector as

$$\Gamma(\delta) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \cos \delta & -\sin \delta \\ 0 & 0 & \sin \delta & \cos \delta \end{pmatrix}, \quad (6.21)$$

which reduces to the idealized matrix in Equation 6.17 when $\delta = \pi$.

6.3.4 Achromatic HWP

As noted by Equation 6.15, an HWP is only ideal for a single frequency, and as bandwidth increases, so too does the HWP's conversion of linear polarization into elliptical polarization. Because linear polarimeters are not sensitive to circular polarization,¹¹ this conversion from linear to elliptical polarization results in a polarization efficiency loss. This loss is quantified by the HWP's **linear polarization modulation efficiency**

$$\varepsilon_{\text{mod}} = \frac{P_{\text{out}}}{P_{\text{in}}}, \quad (6.22)$$

where P_{in} and P_{out} are the input and output linear polarization fractions, respectively.

One popular method to achieve broadband polarization modulation efficiency is to use an **achromatic HWP (AHWP)**, or one that is less sensitive to the color of the incident light. There are many ways to realize an AHWP, including via engineered capacitive grids [157] or metamaterial layers [37], but one particularly common solution is to employ a **Pancharatnam AHWP** [152]. The basic concept of a **Pancharatnam stack** is to increase the number

¹¹There are theories of *cosmological birefringence* that predict some level of CMB circular polarization, but none has yet been detected [143, 148].

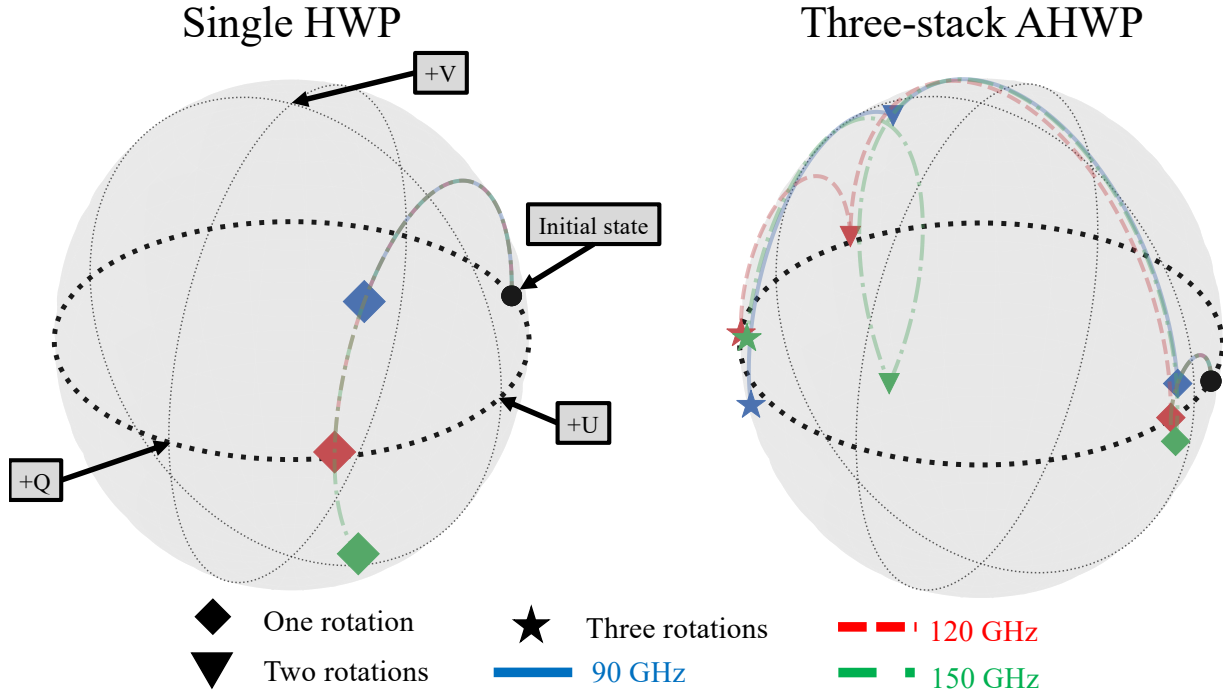


Figure 6.5: A depiction of the Pancharatnam AHPW using the Poincaré sphere. The polarization components (Q, U, V) are the basis vectors of a three-dimensional space, and the sphere's radius $R = \sqrt{Q^2 + U^2 + V^2}$ only considers the polarized component. The HWP manipulates input light by rotating its polarization vector about the sphere. An ideal HWP takes an equatorial (linearly polarized) input vector and outputs another equatorial vector, hence preserving linear polarization fraction, while a non-ideal HWP will convert some fraction of $(Q, U) \rightarrow V$. The left panel shows a single HWP with d_{HWP} tuned to 120 GHz. While linear polarization is well preserved near the ideal frequency, it is increasingly lost for broader bandwidth. The right panel shows a three-stack AHPW also tuned to 120 GHz. Here, three operations keep the output polarization vectors, represented by stars, close to the equator, hence preserving linear polarization between 90 and 150 GHz.

of degrees of freedom that can be used to optimize modulation efficiency¹² by adding more birefringent plates. Equation 6.19 then becomes

$$\vec{S}_{\text{out}} = \prod_{i=1}^{N_{\text{stack}}} [R(-\chi - \theta_i) \Gamma(\delta) R(\chi + \theta_i)] \vec{S}_{\text{in}}, \quad (6.23)$$

where θ_i is the orientation of each plate in the stack of N_{stack} . The product's ordering is from the sky towards the detector, or the order that the plates operate on the input polarization. Figure 6.5 shows how the Pancharatnam stack preserves linear polarization using a **Poincaré sphere**, and Figure 6.6 shows how such a stack is implemented.

¹²This technique is similar in spirit to that used to increase the bandwidth of anti-reflection coatings, a topic detailed in Chapter 10.

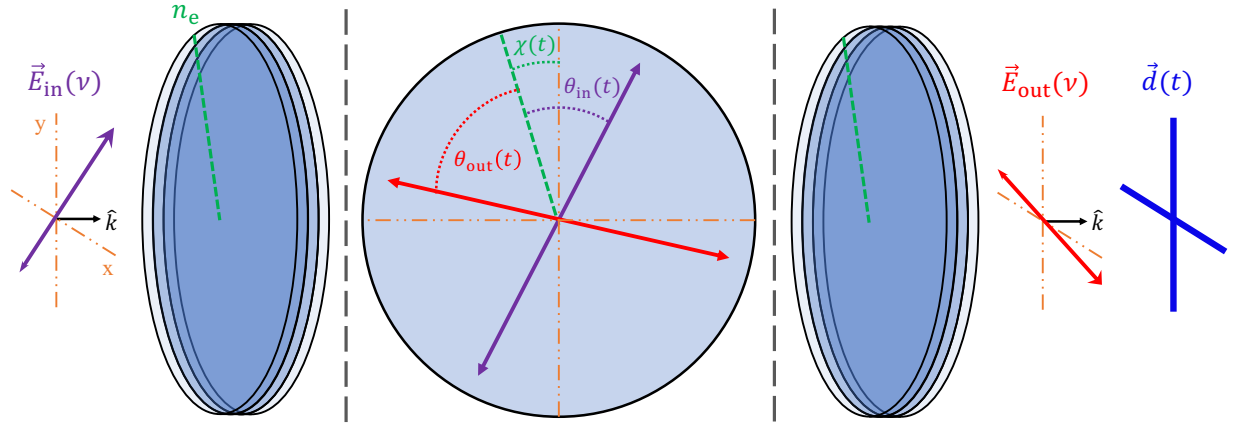


Figure 6.6: A cartoon of the sapphire AHWP. Linearly polarized input light, with wave vector \hat{k} and field vector $\vec{E}_{\text{in}}(\nu)$, is rotated by twice its polarization angle $\theta_{\text{in}}(t)$ with respect to the HWP axis—which coincides with the first sapphire piece’s extraordinary crystal axis n_e —plus a frequency-dependent phase $2\phi(\nu)$. A continuously rotating HWP spins at a constant velocity $d\chi/dt = 2\pi f_{\text{HWP}}$, modulating the output field vector $\vec{E}_{\text{out}}(t)$ at $2f_{\text{HWP}}$ and the detected polarimeter power $\vec{d}(t)$ at $4f_{\text{HWP}}$.

An AHWP can be composed of any birefringent substrate, but an appealing material for mm-wave observation is **sapphire**, which has both a small loss tangent of $\tan \delta \sim 10^{-4}$ (Equation 3.26) and a large differential index $(n_e - n_o) \approx (3.4 - 3.05) \approx 0.35$ at 100 GHz. These two characteristics minimize attenuation and thermal emission in the AHWP, which is important to the telescope’s sensitivity (see Sections 3.4.1 and 3.3). For these reasons and others discussed in Sections 6.1 and 6.2, SA and SO use sapphire AHWPs as polarization modulators.

6.4 AHWP performance

Armed with the formalism for AHWPs given by Equation 6.23, we now evaluate the performance of an AHWP for SA and SO. To keep the discussion focused, we only discuss the 90/150 GHz instruments, noting that much of the following investigation transfers trivially to 34/40 GHz and 220/270 GHz cameras.

Consider linearly polarized light \vec{S}_{in} generated by a narrow-band source that can sweep in frequency. This source illuminates an HWP with N_{stack} plates that rotates the input light to output polarization vector \vec{S}_{out} . The output light is then detected by a polarimeter \vec{d} with




Figure 6.7: Detected intensity vs. HWP angle for a one-, three-, and five-stack HWP (top, middle, and bottom panels) between 50 and 170 GHz. A larger peak-to-peak value indicates a larger modulation efficiency. The frequency-dependent phase shift is a unique quality of the Pancharatnam AHWP.

Mueller matrix

$$G = \frac{1}{2} \begin{pmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad (6.24)$$

such that the detected signal is

$$\vec{S}_{\text{det}} = G \vec{S}_{\text{out}}. \quad (6.25)$$

The effect of the polarimeter matrix G is to admit intensity I_{out} and polarization Q_{out} . In the case of fully polarized output light where $I_{\text{out}} = \sqrt{Q_{\text{out}}^2 + U_{\text{out}}^2}$, the [minimum, maximum] detected power is $[0, I_{\text{out}}]$, while in the case of partially polarized light, it is $[(1 - P_{\text{out}})I_{\text{out}}, I_{\text{out}}]$, where P_{out} is the output polarization fraction. Therefore, we can write the HWP's modula-

tion efficiency in terms of the detected intensity as

$$\varepsilon_{\text{mod}} = \frac{I_{\text{det,max}} - I_{\text{det,min}}}{I_{\text{det,max}} + I_{\text{det,min}}}, \quad (6.26)$$

where $I_{\text{det,max}}$ and $I_{\text{det,min}}$ are the maximum and minimum detected power, respectively. Figure 6.7 shows detected intensity vs. HWP angle for a one-, three-, and five-stack HWP between 50 and 170 GHz. As suggested by the Poincaré sphere in Figure 6.5, modulation efficiency improves away from the central frequency as the number of plates is increased.

In the following subsections, we briefly overview the calculation of modulation efficiency and phase for the AHWP at normal incidence within the context of a 90/150 GHz instrument. However, the topic of AHWP performance is rich, and we encourage the interested reader to learn about other AHWP effects, such as the impacts of increasing detection bandwidth, intensity to polarization leakage, effects at non-normal incidence, and stack orientation optimization [134].

6.4.1 Modulation efficiency

Figure 6.8 shows ε_{mod} as a function of frequency for a one-, three-, and five-stack HWP [134] calculated using Equation 6.26, and the table shows the band-averaged quantities. There are two important findings from this comparison. First, a single HWP provides insufficient bandwidth to cover the 90 and 150 GHz bands, giving rise to a $\approx 25\%$ polarization efficiency loss. This problem is exacerbated by thickness and index tolerances, which will scatter the ε_{mod} contour to higher (lower) frequency and further degrade modulation efficiency in the 90 (150) GHz band. Second, the modulation efficiency gain of a five-stack HWP over a three-stack HWP is only $\approx 0.5\%$, and because a three-stack configuration is 40% thinner and hence 40% less absorptive/emissive (see Equation 3.28), the three-stack HWP is favored for SA's and SO's dichroic receivers. That said, the modulation phase ϕ_ν is also an important consideration of the AHWP's design, and we discuss it in the next section.


6.4.2 Frequency-dependent phase

The I_{det} vs. χ curves in Figure 6.7 can be modeled via the relation

$$I_{\text{det}} = \frac{I_{\text{in}}}{2} [1 + \varepsilon_{\text{mod}} P_{\text{in}} \cos(4\chi - 2\alpha_{\text{in}} - 4\phi_\nu)], \quad (6.27)$$

where I_{in} is the input Stokes intensity, α_{in} is the input polarization angle, P_{in} is the input polarization fraction, χ is the HWP angle, and ϕ_ν is a frequency-dependent phase. We can invert this relationship to write an analytic function for the frequency-dependent phase as

$$\phi_\nu = -\frac{1}{4} \arccos \left[\frac{2I_{\text{det}}/I_{\text{in}} - 1}{\varepsilon_{\text{mod}} P_{\text{in}}} \right] + \chi - \frac{1}{2} \alpha_{\text{in}}. \quad (6.28)$$



N_{stack}	$\varepsilon_{\text{mod}}(90 \text{ GHz})$ [%]	$\varepsilon_{\text{mod}}(150 \text{ GHz})$ [%]	$\Delta\phi_{\nu}(90 \text{ GHz}, 150 \text{ GHz})$ [deg]
1	73.8	76.7	0.0
3	99.5	99.5	0.5
5	99.9	99.9	0.1

Figure 6.8: A comparison of the band-averaged modulation efficiency and phase for one-, three-, and five-stack HWPs. The top panel shows modulation efficiency (solid lines) and phase (dotted lines) vs. frequency, as well as the bandwidth of SA’s 90 and 150 GHz bands. The table shows the band-averaged modulation efficiencies and phase difference $\Delta\phi_{\nu} = \phi(150 \text{ GHz}) - \phi(90 \text{ GHz})$.

Phase vs. frequency for the one-, three-, and five-stack HWPs are plotted in Figure 6.8, and the band-integrated phase differences between the 90 and 150 GHz bands are shown in that figure’s table.

The effect of ϕ_{ν} is to rotate the detected polarization angle as a function of frequency. Such a rotation is often called **cross polarization (cross-pol)** and is of critical importance to an accurate measurement of E-modes and B-modes. The topic of polarization angle calibration in CMB telescopes is intensively studied [101, 110, 7, 144], and an accurate understanding of how E-modes leak into B-modes—a phenomenon called **EB leakage**—requires a detailed instrument model [11, 10] that is beyond the scope of this dissertation. However, we can gain some intuition for the impact of the AHWP’s cross-pol from a relatively simple analysis.

Consider a *measured* polarized signal $[Q \pm iU]'$ that is the *true* sky signal $[Q \pm iU]$

modified by a global polarization angle rotation of $\Delta\alpha$

$$\begin{aligned} [Q + iU]' &= e^{-2i\Delta\alpha} [Q + iU] \\ [Q - iU]' &= e^{-2i\Delta\alpha} [Q - iU] . \end{aligned} \quad (6.29)$$

As discussed in Section 1.5, we can decompose these linear combinations of Q and U into spherical harmonics as

$$\begin{aligned} a_{2,\ell m} &= \int d\Omega Y_{\ell m}^*(\hat{n}) [Q + iU](\hat{n}) \\ a_{-2,\ell m} &= \int d\Omega Y_{\ell m}^*(\hat{n}) [Q - iU](\hat{n}) , \end{aligned} \quad (6.30)$$

where \hat{n} denotes sky location. Then, the B-mode Fourier amplitude is written as

$$a_{B,\ell m} = i [a_{2,\ell m} - a_{-2,\ell m}] , \quad (6.31)$$

and the B-mode power spectrum amplitude is

$$C_\ell^{BB} = \langle a_{B,\ell m} a_{B,\ell m}^* \rangle , \quad (6.32)$$

where the angle brackets denote the statistical average. Using these relationships, we can roughly quantify the impact of rotating E-modes into B-modes on the power spectrum as

$$\text{Bias} [C_\ell^{BB}] \equiv C_\ell^{EE} \frac{C_\ell^{BB} - C_\ell^{BB'}}{C_\ell^{BB}} = \frac{(a'_{2,\ell m} - a'_{-2,\ell m})^2 - (a_{2,\ell m} - a_{-2,\ell m})^2}{(a_{2,\ell m} - a_{-2,\ell m})^2} . \quad (6.33)$$

In general, a proper estimation of the **BB bias** due to polarization angle rotation $\text{Bias} [C_\ell^{BB}]$ requires an end-to-end simulation pipeline that accounts for how the detector data is filtered, the maps are constructed, and the power spectrum is calculated [11, 10]. However, we can make a very simple estimate by plugging $a_{2,\ell m} = [Q + iU]$, $a_{-2,\ell m} = [Q - iU]$, $a'_{2,\ell m} = [Q + iU] \exp[-2i\Delta\alpha]$, and $a'_{-2,\ell m} = [Q - iU] \exp[-2i\Delta\alpha]$ into Equation 6.33, which reduces to


$$\text{Bias} [C_\ell^{BB}] \sim C_\ell^{EE} \sin^2 (2\Delta\alpha) . \quad (6.34)$$

Noting that the inflationary parameter r measures the ratio of B-mode power to E-mode power, we can write

$$\text{Bias} [r] \sim \sin^2 (2\Delta\alpha) . \quad (6.35)$$

Therefore, any miscalibration of the HWP-dependent phase $\Delta\alpha$ needs to be $\lesssim 1$ deg for $\text{Bias} [r] \lesssim 10^{-3}$.

Using the simple BB bias approximation in Equation 6.34, we can estimate the effectiveness of various angle calibration techniques at mitigating EB leakage. The essential issue is that when using an AHWP, if we measure the telescope's polarization angle using



Channel	Tau-A $\Delta\alpha_{\text{inst}}$ [deg]	SWG $\Delta\alpha_{\text{inst}}$ [deg]
90 GHz	-0.2	-0.2
150 GHz	0.3	0.4

Figure 6.9: The estimated angle error induced by the AHWP when calibrated using Tau-A and an SWG. The figure shows peak-normalized spectra for each source as well as the frequency-dependent phase for a three-stack AHWP. The table shows the angle difference between Tau-A/the SWG and the CMB for both the 90 and 150 GHz bands using Equation 6.39.

a source with a different spectrum than that of the CMB, we will bias the BB power spectrum.¹³ In this sense, the problem raised by the AHWP’s frequency-dependent phase ϕ_ν in Equation 6.28 is one of understanding the calibration source’s spectrum.

Consider three example methods for polarization calibration:

1. Self-calibration. This technique assumes no cosmic birefringence and therefore nulls C_ℓ^{EB} to calibrate the polarization angle [101].
2. Calibration using a celestial source. The Crab Nebula (also called “**Tau-A**”) is a polarized supernova remnant that is bright at mm wavelengths [7].
3. Calibration using a terrestrial source. There are many techniques to measure a telescope’s polarization angle from Earth, but we highlight one method that was prototyped on POLARBEAR by Osamu Tajima, which is to reflect ~ 300 K emission from the ground using a **sparse wire grid (SWG)** to generate a small polarized signal [198].

¹³The same principle applies to galactic foregrounds, which have different spectra than the CMB.

The diffraction-limited, single-moded, power spectra of the CMB, Tau-A, and the sparse wire grid are

$$p_{\text{CMB}}(\nu) = \varepsilon_{\text{CMB}} \frac{h\nu}{e^{h\nu/(k_{\text{B}}T_{\text{CMB}})} - 1} \quad (6.36)$$

$$p_{\tau\text{A}}(\nu) = A_{\tau\text{A}} + B_{\tau\text{A}} (\nu/\nu_0)^{\beta_{\tau\text{A}}} \quad (6.37)$$

$$p_{\text{SWG}}(\nu) = \varepsilon_{\text{SWG}} \frac{h\nu}{e^{h\nu/(k_{\text{B}}T_{\text{gnd}})} - 1}. \quad (6.38)$$

The first equation is a $T_{\text{CMB}} = 2.7$ K blackbody spectrum with a frequency-independent polarization fraction ε_{CMB} . In the second equation, $A_{\tau\text{A}}$ and $B_{\tau\text{A}}$ are Tau-A's coefficients in units of W/Hz, and $\beta_{\tau\text{A}}$ is the spectral index, which according to Ritacco *et al.* (2017) is $\beta_{\tau\text{A}} \approx -0.35$ [162]. The final equation is simply that of a blackbody with temperature $T_{\text{gnd}} = 280$ K, modified by the SWG's frequency-independent polarizing efficiency ε_{SWG} . Each peak-normalized spectrum between 60 and 180 GHz is shown in Figure 6.9.

To quantify possible misestimation of the CMB's polarization angle when calibrating the AHWP's phase using either Tau-A or the SWG, we find the band-averaged instrument angle

$$\alpha_{\text{inst}} = \frac{\int_0^\infty \phi(\nu) B(\nu) p(\nu) d\nu}{\int_0^\infty B(\nu) p(\nu) d\nu}, \quad (6.39)$$

where $B(\nu)$ is the detector's band, $p(\nu)$ is the polarized spectrum of the calibration source, and $\phi(\nu)$ is given by Equation 6.28. We then compare α_{inst} for each calibration source to that of the CMB. The values for $\Delta\alpha_{\text{inst}}$ in the 90 and 150 GHz bands (the primary CMB channels) are given in the table of Figure 6.9. Even without any mitigation, the angle bias introduced by the AHWP is $< 1^\circ$, which meets SA's requirements.

There are a few points worth emphasizing about this brief study of the AHWP's frequency-dependent phase. First, the presented analysis is rough, and more rigorous studies have been done for EBEX¹⁴ [11, 10], which employed a five-stack AHWP for balloon-borne, trichroic CMB observation [206]. Nonetheless, EBEX's conclusions are similar to those shown here: even with no mitigation techniques, the C_ℓ^{BB} bias introduced by the AHWP's frequency-dependent rotation is small. Second, the spectral-mismatch problem presented here for CMB calibration is also true for **foreground subtraction**. Because dust and synchrotron emission have different spectra than the CMB (see Section 1.7), ϕ_ν must be accounted for during component separation.

Third, while these frequency-dependent effects have been deemed subdominant for SA and SO, they could become important for future missions with tighter systematic error budgets. Therefore, there is ongoing research into how AHWPs can be constructed to have a frequency-*independent* angle, and the LiteBIRD team has made substantial progress in this area [108]. Finally, we note that the sinuous antenna (see Section 2.4) also has a frequency-dependent **wobble** that generates its own $\Delta\alpha_{\text{inst}}$. However, SA and SO use a technique

¹⁴The impact of the AHWP's frequency-dependent phase is also being actively studied within SO but is not quite published as of December 2020.

called **wobble cancellation** to mitigate the antennas' frequency-dependent angle [49, 197, 39], where the sinuous's handedness is mirrored throughout the focal plane, nulling the wobble when neighboring pixels are coadded. In theory, a similar technique could be used to cancel the AHWP's ϕ_ν by placing a second, mirrored AHWP (immediately) detector-side of the first. However, because ϕ_ν is small and can be corrected during analysis, such a hardware remedy is not considered in this dissertation.

Given its high modulation efficiency and low frequency-dependent phase error, we adopt a three-stack sapphire Pancharatnam AHWP for both SA and SO.

6.5 Continuous polarization modulation

There are two common modes of HWP polarization modulation. One mode is to use a **stepped HWP**, which involves rotating the HWP to discrete angles between observations. This technique allows sky polarization to be observed with various global rotations, providing a powerful tool to separate cosmic polarization from instrumental polarization. It also allows orthogonal polarimeters to measure the same sky polarization, providing a mechanism to separate beam asymmetries. While a stepped half-wave plate mitigates the beam systematics discussed in Section 6.2, it does not mitigate 1/f atmospheric fluctuations, and therefore ground-based experiments instead often employ the second mode: a **continuously rotating HWP**.

SA and SO use continuously rotating sapphire AHWPs as polarization modulators, mitigating both beam systematics and atmospheric 1/f noise. A continuously rotating HWP modulates an input polarized sky signal $[Q_{\text{in}}(t) \pm iU_{\text{in}}(t)]$ onto a detector with output $d_{\text{m}}(t)$ (see Figure 6.6) as

$$d_{\text{m}}(t) = I_{\text{in}}(t) + \varepsilon_{\text{mod}} \text{Re}\{[Q_{\text{in}}(t) \pm iU_{\text{in}}(t)] \exp[\mp i4\chi(t)]\}, \quad (6.40)$$

where χ is the HWP's rotation angle and $m(\chi) = \exp[\mp i4\chi(t)]$ is the **modulation function**. A continuously rotating HWP spins the birefringent stack with a steady velocity f_{HWP} such that

$$\chi(t) = 2\pi f_{\text{HWP}} t, \quad (6.41)$$

where

$$f_{\text{m}} = 4f_{\text{HWP}} \quad (6.42)$$

is the modulation frequency. The detector's output is modulated at $4\times$ the HWP's rotation frequency because both the HWP and the detector have a π symmetry (see Figure 6.6). Equation 6.40 is written in the time domain, but it is often useful to conceptualize the HWP's operation in the Fourier domain. To provide some intuition, consider a sky signal composed of a single Fourier mode

$$[I_{\text{in}} + Q_{\text{in}} \pm U_{\text{in}}](t) = [I_{\text{in}} + Q_{\text{in}} \pm U_{\text{in}}] \exp(i\omega_{\text{sig}} t), \quad (6.43)$$




Figure 6.10: A toy demonstration of amplitude modulation. In this example, the carrier frequency is $f_m = 8$ Hz, and the (top, middle, bottom) panels show the time-domain (left) and Fourier-domain (right) signature of a $f_{\text{sig}} = (0, 1, 5)$ Hz sky signal, where $\omega_{\text{sig}} = 2\pi f_{\text{sig}}$ as in Equation 6.44. The simulated white-noise amplitude is 0.5% and is arbitrary.

which when plugged into Equation 6.40 and Fourier transformed becomes

$$d_m(\omega_{\text{sig}}) = I(\omega_{\text{sig}}) + \frac{1}{2} [Q_{\text{in}} + iU_{\text{in}}] (4\omega_{\text{HWP}} + \omega_{\text{sig}}) + \frac{1}{2} [Q_{\text{in}} - iU_{\text{in}}] (4\omega_{\text{HWP}} - \omega_{\text{sig}}), \quad (6.44)$$

where here $\omega_{\text{HWP}} = 2\pi f_{\text{HWP}}$. Three examples of Equation 6.44 are shown in Figure 6.10, and we discuss a few salient features. First, the polarized sky signals $[Q_{\text{in}} \pm iU_{\text{in}}]$ are “moved” from frequency ω_{sig} to frequencies $4\omega_{\text{HWP}} \pm \omega_{\text{sig}}$, while the intensity sky signal is left at ω_{sig} . This is the signature of **amplitude modulation**, where variations in the polarized sky signal modulate the amplitude of the HWP-generated **carrier signal** at $4\omega_{\text{HWP}}$. Stated mathematically, the modulator multiplies $\exp(i\omega_{\text{sig}}t)$ by $\exp(i4\omega_{\text{HWP}}t)$, which can be simplified as $\sin(\omega_{\text{sig}}t) \sin(4\omega_{\text{HWP}}t) = 1/2[\cos((4\omega_{\text{HWP}} + \omega_{\text{sig}})t) - \cos((4\omega_{\text{HWP}} - \omega_{\text{sig}})t)]$. Second,

because sky signals appear at frequencies $\pm\omega_{\text{sig}}$ around the carrier frequency $4\omega_{\text{HWP}}$, the HWP’s **side bands** $4\omega_{\text{HWP}} \pm \omega_{\text{sig}}^{\text{max}}$ should be in a **white-noise-dominated frequency regime**. Stated another way, the HWP’s rotation frequency f_{HWP} must be large enough such that polarized sky signals are up-converted to frequencies above the atmospheric 1/f noise. We discuss this requirement with more detail in Section 6.7.2.

6.5.1 HWP synchronous signals

Equation 6.40, while instructive, is only a portion of the signal from a real HWP. In addition to the modulated sky signal, the HWP generates **HWP synchronous signals** (HWPSSs), which can be decomposed as

$$\mathcal{P}_{\text{HWPSS}} = \sum_{n=1} \text{Re}\{A_n(t)e^{-in\chi(t)}\}, \quad (6.45)$$

where A_n is the amplitude of the n th HWP harmonic. There are many mechanisms that generate HWPSSs in the detector data, including **differential thermal emission** between the HWP’s crystal axes ($n = 2$), **differential reflection** of incident intensity along the HWP’s crystal axes ($n = 2$), **differential transmission** of the S and P polarizations at non-normal incident angles ($n = [2, 4, 6, \dots]$), imperfections in the HWP’s sapphire stack or AR coatings (any n), **rotation-synchronous** electromagnetic signals generated by the HWP’s drive mechanism (any n), and HWP-rotation-induced vibrations (any n), to name a few. While HWPSSs can appear at any harmonic n , the most prominent are typically those at $n = 2$ and $n = 4$ and are often referred to as the “**2f**” and “**4f**” signals.

The most important HWPSSs to control are those at $4f_{\text{HWP}}$, as $A_4(t)e^{-i4\chi(t)}$ is essentially indistinguishable¹⁵ from the sky signal $\text{Re}\{[Q_{\text{in}}(t) \pm iU_{\text{in}}(t)] \exp[\mp i4\chi(t)]\}$ in Equation 6.40. This dangerous degeneracy motivates many of SA’s HWP requirements, and we highlight two particularly important HWPSS mechanisms here.

The first optical mechanism that induces a $4f_{\text{HWP}}$ HWPSS is instrumental polarization generated by optics that are sky-side of the HWP. In its ideal application, the HWP modulator is the skyward-most optical element [111]. However, in many cases, practical limitations, such as a large primary aperture or the modulator’s being cryogenic, require the HWP to reside “behind” some of the telescope’s optics. As shown in Figure 2.3b, the PB-2b CHWP is located behind the telescope’s mirrors, vacuum window, and IR filters, while Figure 2.4c shows that the SO SAT’s CHWP is located behind its vacuum window and IR filters. Because vacuum windows and IR filters do not have perfect AR coatings and because SA’s off-axis mirrors are not lossless, they generate I-to-P due to differential transmission of the S and P polarizations at non-normal incidence. This sky-side I-to-P signal is modulated at $4f_{\text{HWP}}$ by the HWP, and its amplitude $A_4(t)$ varies with sky intensity. The second optical mechanism that induces a $4f_{\text{HWP}}$ HWPSS is non-idealities in the sapphire stack. As an example,

¹⁵This degeneracy can be broken with an orthogonal data set, such as from thermometers, but it is very difficult to do effectively in practice.

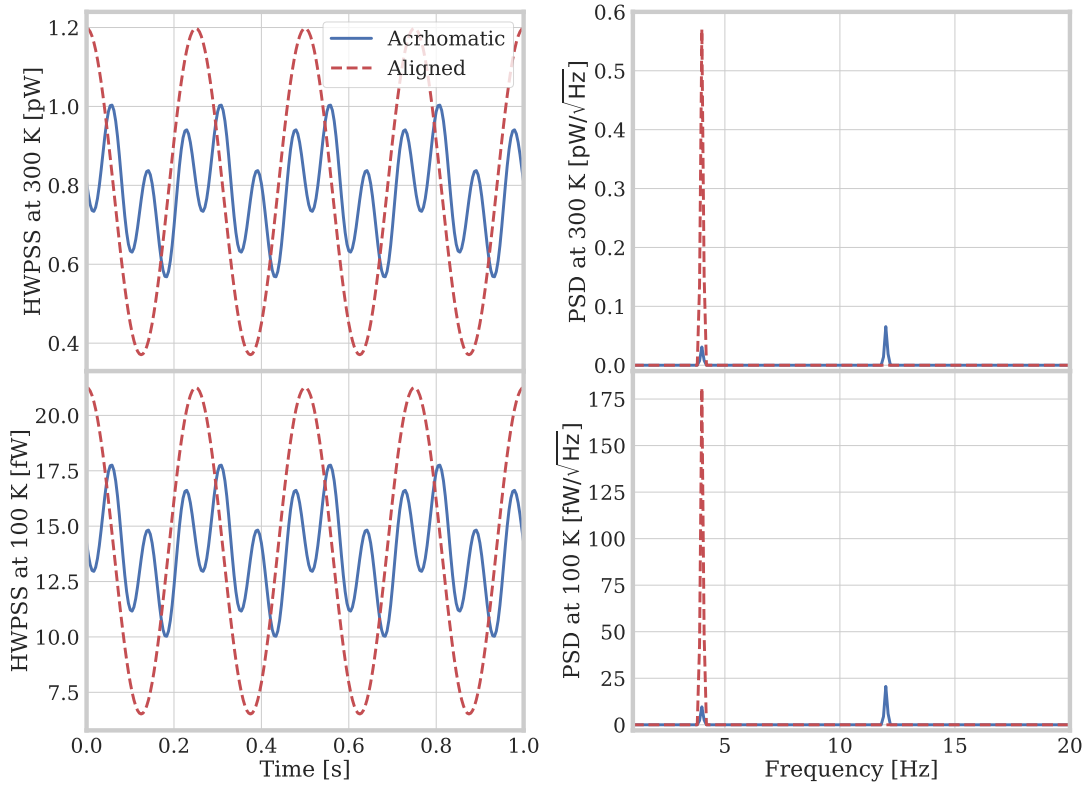


Figure 6.11: Differential thermal emission at normal incidence from a three-stack sapphire AHWP at 300 K and 100 K. The signal with the crystal axes for all three plates aligned is also plotted as a reference. The AHWP’s thermal emission has $2f_{\text{HWP}}$ and $6f_{\text{HWP}}$ components but none at $4f_{\text{HWP}}$.

localized imperfections or position-dependent optical properties can both generate polarized emission that varies with HWP angle. The $4f_{\text{HWP}}$ component of these emissive (reflective) non-idealities will vary with the HWP’s (sky-side optics’) temperature and contaminate the cosmic polarization signal. These two optical examples are only a subset of possible HWPSSs, and we detail some non-optical contaminants when discussing the PB-2b CHWP’s hardware requirements in Chapter 8. For a comprehensive discussion of HWPSSs and how they are handled during analysis, we refer the reader to the work of Satoru Takakura [199, 200].

For a single sapphire HWP, a primary generator of $2f_{\text{HWP}}$ HWPSSs is differential thermal emission between the sapphire’s ordinary and extraordinary crystal axes. Sapphire’s extraordinary axis has both a larger index and a larger loss tangent than the ordinary axis, and therefore the two axes have different emissivities. This differential emission corotates with the HWP and shows up at $2f_{\text{HWP}}$ in the detector data. However, for a three-stack AHWP, the situation is more complicated, as differential emission from sky-side sapphire plates is *rotated* by detector-side plates. An important question for the AHWP is whether this internally modulated differential thermal emission has a $4f_{\text{HWP}}$ component. Figure 6.11

shows the Mueller-matrix-simulated time-ordered data (TOD) and power spectral density (PSD) of the AHWP's thermal emission at 300 K and 100 K. While there are clear signals at $2f_{\text{HWP}}$ and $6f_{\text{HWP}}$, there are no signals at $4f_{\text{HWP}}$, and therefore, to leading order,¹⁶ differential emission from the AHWP does not contaminate the sky signal.

6.6 HWPs for Simons Array

Using the presented mathematical framework to describe and assess AHWPs, we now discuss their implementation within SA. While a similar discussion for SO is beyond the scope of this dissertation, many of the presented principles are shared between SA and SO. Before deriving some SA HWP optical requirements in Section 6.7, we first discuss the general philosophy of the SA modulator as well as the precedent that motivates it.

6.6.1 The POLARBEAR HWP

As described in Section 2.3.1, SA consists of three telescopes with 2.5 m primary mirrors that enable CMB polarization measurements up to $\ell \sim 2,000$. SA's telescope design is functionally identical to that of its predecessor experiment POLARBEAR, and therefore we briefly discuss POLARBEAR's HWP as a precedent for those in SA [199].

The Huan Tran Telescope's (HTT's) primary and secondary mirrors form a compact off-axis Gregorian configuration, which is designed to optimize telescope mobility and hence improve low- ℓ sensitivity. To understand how mobility and low- ℓ sensitivity are related, recall that atmospheric turbulence evolves with time due to wind, cloud movement, and other weather effects. If this atmospheric evolution is faster than the time it takes to scan the sky patch, the cosmic signal is buried beneath atmospheric fluctuations. In order to beat this atmospheric noise, the sky signal must be modulated faster than atmospheric fluctuations, and one way to accomplish this goal is to scan the sky quickly. When *not* using an HWP, the polarized sky signal appears in side bands around the telescope's scan frequency f_{scan} , and therefore the larger the scan speed, the smaller the atmospheric 1/f noise.¹⁷

POLARBEAR has a Goldilocks-type telescope design, having both a large aperture and a small, agile frame, which together enable both low- ℓ and high- ℓ science. For this reason, among others, the initial POLARBEAR telescope design did not include a continuously rotating HWP but instead relied on azimuth motion and detector differencing to mitigate 1/f noise in polarization. Using these techniques, POLARBEAR made the first polarization-only null-hypothesis rejection of lensing B-modes between $600 < \ell < 3,000$ during 2012-2014 [213]. However, as POLARBEAR aimed to measure even lower ℓ ranges in 2014-2016,

¹⁶Higher-order considerations such as internal reflections at non-normal-incidence can bleed the $2f_{\text{HWP}}$ component into $4f_{\text{HWP}}$.

¹⁷As an example, the GroundBIRD CMB experiment is specifically designed to beat atmospheric noise by rotating the telescope at 20 revolutions per minute in azimuth [117].

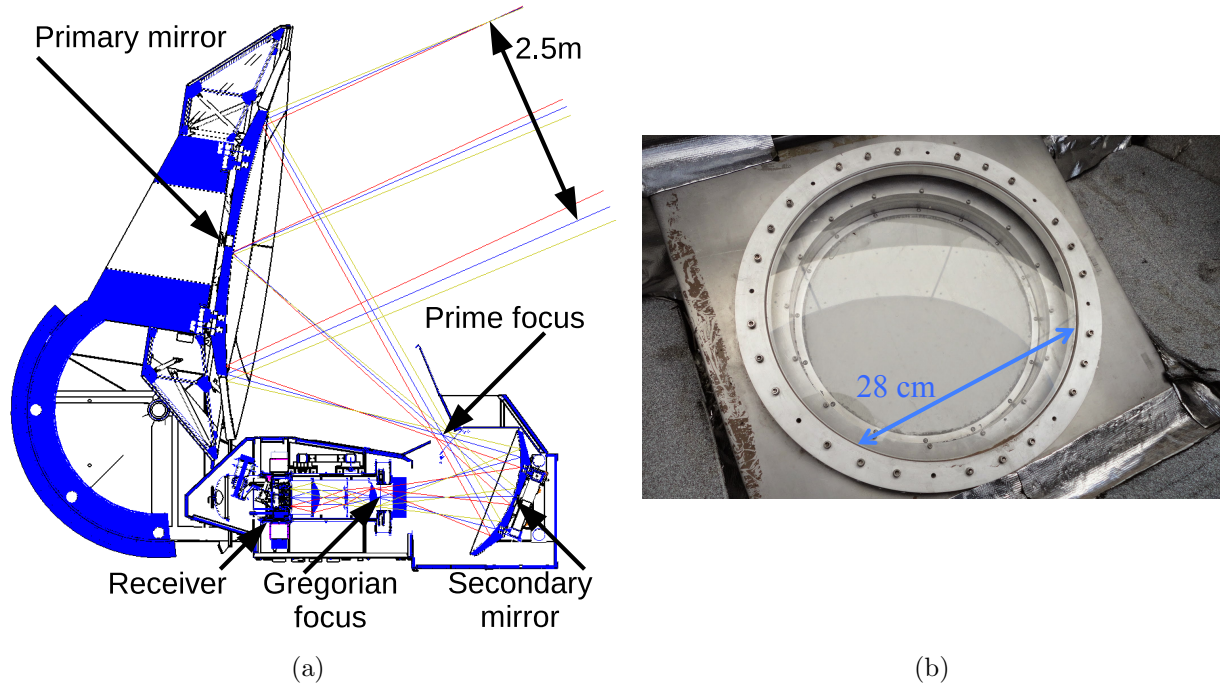


Figure 6.12: An overview of the POLARBEAR HWP’s implementation. Figure 6.12a shows a CAD cross-section and ray trace of the POLARBEAR telescope. The POLARBEAR HWP is located at the prime focus, between the primary and secondary mirrors. The image is taken from Takakura *et al.* (2017) [199]. Figure 6.12b shows a photograph of the POLARBEAR HWP at prime focus, which is weatherproofed by a thin Mylar sheet. The image is courtesy of Satoru Takakura and is taken from Hill, Beckman, *et al.* (2016) [82].

it retrofitted an ambient-temperature continuously rotating HWP at its prime focus, as shown in Figure 6.12 [199].

POLARBEAR’s noise spectrum in both intensity and demodulated polarization, as well as its achieved N_ℓ spectrum (see Equation 3.71), are shown in Figure 6.13. After subtracting I-to-P due to detector non-linearity [199], the POLARBEAR HWP facilitates white-noise-dominated stability down to < 5 mHz, in turn enabling an $\ell_{\text{knee}} \sim 90$. POLARBEAR’s HWP demonstration is the foundation of SA’s HWP development. Because the SA optical design is nearly identical to that of POLARBEAR, the successes of the POLARBEAR HWP acts as an “existence proof” that guides the SA HWP’s design and evaluation.

6.6.2 The SA HWP

Spurred by the success of the POLARBEAR HWP, SA has adopted continuously rotating HWPs for their three dichroic telescopes. Before detailing the SA HWP’s requirements in Section 6.7, we briefly discuss SA’s history and project structure.

When SA was first funded, it comprised three POLARBEAR-2 (PB-2) experiments, one of which was already under laboratory evaluation at the KEK High Energy Research

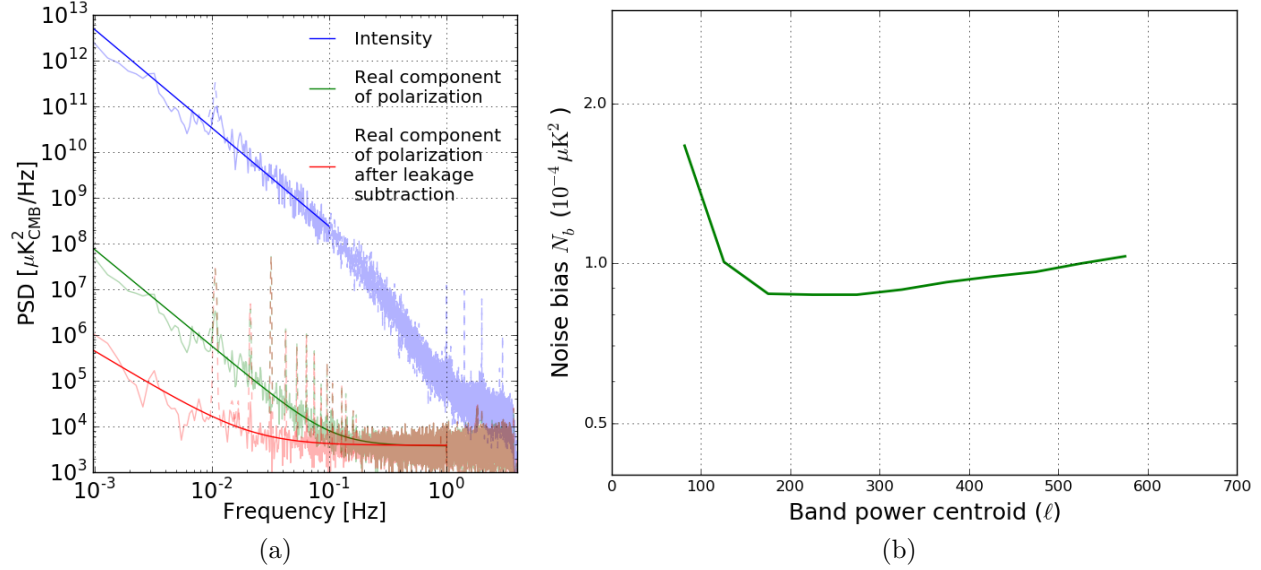


Figure 6.13: The demonstrated performance of the POLARBEAR HWP. Figure 6.13a shows PSDs of time streams for coadded detectors cross the full focal plane. An additional I-to-P leakage subtraction is applied after demodulation to further suppress the $1/f$ knee in polarization. This figure is taken from Takakura *et al.* (2017) [199]. Figure 6.13b shows the N_ℓ curve for a full season of POLARBEAR data with the HWP. The achieved $\ell_{\text{knee}} \sim 90$ is a promising demonstration of the HWP’s ability to enable low- ℓ B-mode science on an SA-style telescope. This figure is taken from POLARBEAR’s third-season C_ℓ^{BB} results paper (2020) [211].

Organization in Japan. PB-2 was originally designed to house a continuously rotating AHWP near the Lyot stop, but as R&D advanced, it became clear that a 4 K modulator would not be ready in time for deployment. Therefore, shortly before PB-2 became PB-2a (the first installment of SA), its HWP was moved outside of the cryostat, decoupling its development from the PB-2a receiver and accelerating the overall project timeline. When SA’s second and third telescopes were funded in 2016, the PB-2b and PB-2c AHWPs moved back into the cryostat to boost instrument sensitivity. However, to avoid the same problems encountered with PB-2a’s 4 K AHWP, the PB-2b/c HWP is instead located on the 50 K stage, relaxing several of its mechanical requirements (see Section 8.2) and enabling its parallel development with the receiver cryostat, just as for PB-2a.

6.7 SA HWP requirements

This dissertation focuses on HWP development for PB-2a and PB-2b, which both observe at 90 and 150 GHz. While we have also worked substantially on the PB-2c HWP, it shares many characteristics with that of PB-2b, and therefore we relegate its detail to future papers. In the following section, we discuss the requirements that are common to the PB-2a and PB-2b HWPs, including rotational velocity, angle encoder noise, and sapphire tolerances. For the historical reasons described in Section 6.6.2, the fielded AHWPs are retrofitted to the




Figure 6.14: PB-2a/b band-averaged modulation efficiency vs. plate orientation for the 3- and 5-stack AH-WPs. While a 5-stack configuration has more bandwidth, it provides no advantage over the 3-stack configuration for PB-2a/b at the percent level. The optimal relative orientations for the 3-stack AHWP are (0, 54, 0) deg.

“original” PB-2 receiver, and therefore much of the modulator’s design effort focuses on maintaining “heritage” system performance.

6.7.1 Modulation efficiency

The first requirement for the SA AHWP is to preserve input linear polarization fraction. As demonstrated in Figure 6.5, an imperfect AHWP converts linear polarization Q/U into circular polarization V, which manifests as a polarization efficiency loss. Polarized mapping speed (see Equation 5.6) is related to polarization efficiency as $MS \propto \epsilon_{\text{mod}}^2$, and therefore we mandate that the AHWP’s modulation efficiency be $> 95\%$.

An AHWP’s modulation efficiency depends on the thickness and relative orientations of its sapphire plates. Assuming that d_{HWP} in Equation 6.15 is optimized for 120 GHz,¹⁸ band-averaged modulation efficiency vs. plate orientation for the 3- and 5-stack configurations are shown in Figure 6.14 [134]. While a 5-stack configuration has more bandwidth, as shown in Figure 6.8, it provides no modulation efficiency advantage at the 1% level¹⁹ for PB-2a/b. The optimal 3-stack configuration is (0, 54, 0) deg.

Given the 3-stack design, we can now estimate the impact of thickness and relative-alignment tolerances on modulation efficiency. We assume uniform distributions between

¹⁸Equal thicknesses for each sapphire plate best ensures even coverage of the 90 and 150 GHz bands [134].

¹⁹While a 5-stack does have better modulation efficiency at the 0.5% level, more sapphire windows increase absorptivity, thermal emission, and assembly complexity, all of which have stronger impacts on experiment performance.




Figure 6.15: The impact of thickness (left) and relative-alignment (right) tolerances on the PB-2a/b AHWP’s modulation efficiency. Each x value is simulated 1,000 times. The solid lines denote the median, and the shaded regions denote the (15, 85) percentiles. The leading-order effect of modulating both thickness and angle is to shift the modulation-efficiency band’s location, and therefore when the 90 (150) GHz band scatters high (low), the 150 (90) GHz band scatters low (high).

each tolerance’s \pm values—as is common practice in manufacturing—and simulate 1,000 Monte Carlo realizations of the band-averaged modulation efficiency at 90 and 150 GHz. The results of the simulation are shown in Figure 6.15. The dominant effect of scattering thickness and alignment is to shift the modulation efficiency band such that one band is favored over the other. Therefore, to keep 90/150 GHz polarization efficiencies well-matched, we set a thickness tolerance of ± 0.2 mm and a relative alignment tolerance of $\pm 2^\circ$.

6.7.2 Rotation velocity

Another common requirement to PB-2a and PB-2b is the HWP’s **rotational velocity**. Recall from Section 6.5 that the polarization modulation frequency f_m is related to the HWP’s rotation frequency f_{HWP} as $f_m = 4f_{\text{HWP}}$. There are two requirements on the minimum modulation frequency.²⁰ First, the HWP must rotate fast enough to suppress $1/f$ noise in the demodulated timestream. To do this, its signal side bands (see Figure 6.10) must be in the white-noise dominated regime of the raw data spectrum. Because faster sky signals appear

²⁰There is technically one requirement on the modulation frequency’s maximum value, which is that the modulated signal’s period cannot be too close to the detector’s time constant. However, TES bolometers (see Section 2.4.1) typically operate with $\tau \lesssim 20$ ms or $f_{\text{bolo}} \gtrsim 50$ Hz, which is much faster than the HWP needs to operate.




Figure 6.16: A notional depiction of the HWP’s modulation band referenced to the PSD of dummy-generated intensity data with an $f_{\text{knee}} = 1$ Hz.

further from the carrier, the **modulation bandwidth** $\pm\Delta f_{\text{sig}}$ is determined by the fastest resolvable sky signal during demodulation.

As discussed in Section 6.1 and shown in Equation 6.4, the maximum resolvable ℓ_{max} is related to the telescope’s scan speed (in deg/s) as

$$\ell_{\text{max}} \sim \frac{f_{\text{max}}}{f_{\text{scan}}} \times 180^\circ. \quad (6.46)$$

PB-2a/b each have an angular resolution of ≈ 3 arcmin at 150 GHz (see Section 2.3.1) that can resolve up to $\ell_{\text{max}} \sim 3,000$, and if we assume the same telescope scan speed $f_{\text{scan}} = 0.4^\circ/\text{sec}$ used by POLARBEAR [211, 214], then the full range of angular scales are resolved within $\pm\Delta f_{\text{sig}} = \pm f_{\text{max}} \approx \pm 3$ Hz of bandwidth. To provide some margin for filtering to higher ℓ , PB-2a/b’s modulation band is chosen to be

$$f_{\text{m}} \pm \Delta f_{\text{sig}} = f_{\text{m}} \pm 4 \text{ Hz}. \quad (6.47)$$

According to POLARBEAR’s intensity PSD in Figure 6.13, the atmospheric $1/f$ knee at the Chile observation site is $1 \sim 2$ Hz, and therefore we set $f_{\text{m}} \geq 8$ Hz and $f_{\text{HWP}} \geq 2$ Hz, placing the signal band comfortably clear of unpolarized $1/f$ noise. The PB-2 HWP modulation band is shown in Figure 6.16 with respect to a representative intensity PSD.

6.7.3 Angle encoding noise

The final requirement that is common to both the PB-2a and PB-2b HWP systems is the accuracy with which the HWP's rotation angle must be reconstructed. An error in the measurement of the **rotation angle** gives rise to an error in the measurement of rotation synchronous signals, preventing such signals from being fully subtracted during demodulation.

Both the PB-2a and PB-2b HWPs are located behind the telescope's off-axis mirrors, which induce I-to-P along the y -direction (see Figure 2.3). POLARBEAR measures up to ≈ 180 mK_{RJ} of I-to-P at the prime focus of an SA-style telescope, largely due to off-axis emission from the primary reflector [199]. Because the PB-2a/b HWP is located behind the primary *and* secondary mirrors, we conservatively²¹ assert an I-to-P amplitude of

$$A_4^0 \sim 400 \text{ mK}_{\text{RJ}}, \quad (6.48)$$

where here K_{RJ} is Rayleigh-Jeans (RJ) temperature. This mirror polarization is modulated at $4f_{\text{HWP}}$, and a mismeasurement $\Delta\chi$ of the HWP's angle χ modifies the $4f_{\text{HWP}}$ amplitude $A_4(\chi)$ as

$$\begin{aligned} A_4(\chi) &= A_4^0 \cos(4[\chi + \Delta\chi]) \\ &\approx A_4^0 \cos(4\chi) - 4 A_4^0 \sin(4\chi) \Delta\chi, \end{aligned} \quad (6.49)$$

such that

$$\Delta A_4 = -4 A_4^0 \sin(4\chi) \Delta\chi. \quad (6.50)$$

Assuming that this **angle jitter** is random, its average effect on the detected noise spectrum is

$$\text{NET}_{\text{CMB}}^{\text{HWP}} = 2\sqrt{2} \frac{dT_{\text{CMB}}}{dT_{\text{RJ}}} A_4^0 \sigma_\chi, \quad (6.51)$$

where σ_χ is the root mean square (RMS) of $\Delta\chi$, $\text{NET}_{\text{CMB}}^{\text{HWP}}$ is the noise equivalent CMB temperature (see Equation 3.58) of the angle jitter, and $dT_{\text{CMB}}/dT_{\text{RJ}} = 1.7$ is the conversion from RJ to CMB temperature at 150 GHz.

All detectors on the focal plane are demodulated using the same angle encoder data, and therefore any noise in the encoder reconstruction is common to all detectors. In order for angle encoder noise to not significantly contribute to the noise of the detector array, $\text{NET}_{\text{CMB}}^{\text{HWP}}$ needs to be subdominant to the array NET (see Equation 3.62), which is forecasted to be $\text{NET}_{\text{CMB}}^{\text{arr}} = 5.8 \mu\text{K}/\sqrt{\text{Hz}}$ for both PB-2a and PB-2b [196]. Given this constraint, the angle noise requirement for the SA HWPs is

$$\sigma_\chi \ll 3 \mu\text{rad}/\sqrt{\text{Hz}}. \quad (6.52)$$

This requirement drives several aspects of the SA HWP designs, including excellent rotation stability, low encoder noise, and sufficient encoder resolution.

²¹In reality, we expect the I-to-P from the secondary reflector to be less than that of the primary due to its being ellipsoidal with a smaller incident angle. However, the secondary mirror's emission has not been studied in detail, and so we assert that each mirror generates similar I-to-P leakage in order to set a conservative angle jitter requirement.

6.8 Discussion

In this chapter, we have discussed the motivation behind polarization modulators, the principles of continuously rotating AHWPs, the precedent for their implementation, and the AHP requirements that are common to both PB-2a and PB-2b. Armed with this background information, we now discuss the mechanical design, construction, and testing of PB-2a's ambient-temperature rotation stage in Chapter 7 and PB-2b's cryogenic rotation stage in Chapters 8-9. Finally, Chapter 10 discusses the cryogenic anti-reflection coating for the PB-2b sapphire, which plays an important role in the modulator's performance.

Chapter 7

A warm HWP for PB-2a

This chapter overlaps substantially with Hill, Beckman, et al. (2016) [82], which we encourage the reader to cite instead of or in addition to this dissertation.

Building on the successes of its predecessor experiment POLARBEAR (PB-1), Simons Array (SA) employs an achromatic half-wave plate (AHWP) polarization modulator on each of its three telescopes. POLARBEAR-2a (PB-2a) fields an ambient-temperature HWP—or a “warm” HWP (WHWP)—operating in front of its receiver cryostat’s vacuum window, while PB-2b and PB-2c employ a cryogenic HWP (CHWP) located inside the cryostat (see Chapter 2.3.1 for more details). In this chapter, we discuss the **WHWP for PB-2a**, including its requirements, design, construction, and laboratory evaluation.

Mimicking the discussion in Section 6.6.2, we briefly review the historical context surrounding the PB-2a experiment, which motivates several of the WHWP’s design trajectories. Before SA came to be in 2015, POLARBEAR-2 (prior to the “a/b/c” labels) [133] was a fully-funded joint venture between researchers at the KEK high energy research organization in Japan¹ and UC Berkeley. The PB-2 receiver cryostat was designed and built in Japan, the PB-2 detectors and readout were developed and fabricated in Berkeley,² and the telescope, which is nearly identical to that of PB-1, was constructed commercially in Italy.³ PB-2’s design principle is to leverage some of PB-1’s most effective technologies—such as the Huan Tran Telescope (HTT), lenslet-coupled antennas, and frequency-domain multiplexing—to build a world-class next-generation instrument. Some of PB-2’s most prominent hardware upgrades over PB-1 are a dichroic receiver, a larger-diameter optics tube and focal plane to enable a larger field of view (FOV), and a higher-frequency readout system to facilitate a larger multiplexing factor.

The original PB-2 design included a continuously rotating AHWP at the 4 K Lyot stop (see Figure 2.3). However, deploying a modulator deep within the cryostat proved to be an enormous technical challenge, and as PB-2’s deployment neared, the HWP subsystem was

¹KEK CMB group: <https://www2.kek.jp/proffice/archives/intra-e/feature/2009/CMB.html>

²Marvell Nanofabrication Laboratory: <https://nanolab.berkeley.edu/>

³General Dynamics Mission Systems: <https://gd-ms.it/>

descope. In 2014, PB-2 moved its AHWP from the Lyot stop to outside of the cryostat, dramatically reducing the technological challenges surrounding the HWP’s development and construction. Therefore, the presented WHWP research was conducted at UC Berkeley from 2014-2016 in parallel with the commissioning of PB-2’s detectors, readout, and optics tube. Shortly following the PB-2 HWP descope, SA was born, and PB-2 became PB-2a.

The PB-2a WHWP is a three-stack Pancharatnam sapphire AHWP designed to modulate linear polarization for 90 and 150 GHz detectors simultaneously. It is located directly in front of the receiver cryostat’s vacuum window and has a clear-aperture diameter of 480 mm. In the following sections, we present the WHWP’s requirements, optical and mechanical design, and laboratory evaluation. The presented system has deployed to Chile and is ready to operate following the commissioning of PB-2a’s detectors and readout.

7.1 Requirements

The PB-2a WHWP shares several design requirements with all SA HWPs (see Section 6.7), including a 2 Hz rotational velocity, $> 95\%$ modulation efficiency, and an angle encoder noise of $\ll 3 \mu\text{K}/\sqrt{\text{Hz}}$. However, there are several requirements specific to the WHWP that we review in this section, including its aperture diameter, reflectivity, and emissivity.

7.1.1 Clear-aperture diameter

There are two possible locations for the WHWP within the PB-2a system, each shown in Figure 7.1. The first is at the **prime focus** between the primary and secondary mirrors, as was used in PB-1 [199]. The advantage of the prime-focus location is that the **beam**⁴ only has a ≈ 330 mm aperture diameter, enabling the reuse of the PB-1 HWP’s rotator design (see Figure 6.12b). That said, there are several disadvantages to operating the WHWP between the telescope’s mirrors. First, because the prime focus plane is at the primary mirror’s focal point,⁵ each detector pixel illuminates a localized region of the HWP’s surface, as shown in the left panel of Figure 7.1b. This optical configuration is prone to half-wave plate synchronous signals (HWPSSs) caused by position-dependent (and, therefore, rotation-angle-dependent) variations in the WHWP’s optical assembly, such as defects in the anti-reflection (AR) coating or sapphire stack. Second, a prime-focus HWP is far from the cryostat, making it vulnerable to ambient-temperature **stray light**, which in turn increases parasitic optical power on the detectors. In addition, while much of this stray light is, in the reverse-time sense (Figure 3.3), expected to “land” on the telescope’s comoving⁶ baffling, some will “escape” the telescope’s structure and “land” on the ground. This **ground pickup** introduces

⁴The telescope’s “beam” is defined as the collection of reverse-time-sense rays from all detectors on the focal plane. See Figure 7.1a for a visual representation.

⁵The telecentric FOV of the primary mirror is small, and therefore the prime focus is only a “true” focus for a small subset of detector pixels. Even so, the beam for all detectors at the prime focus is tightly collimated relative to the rest of the imaging system.

⁶Comoving describes any structure attached to the telescope.

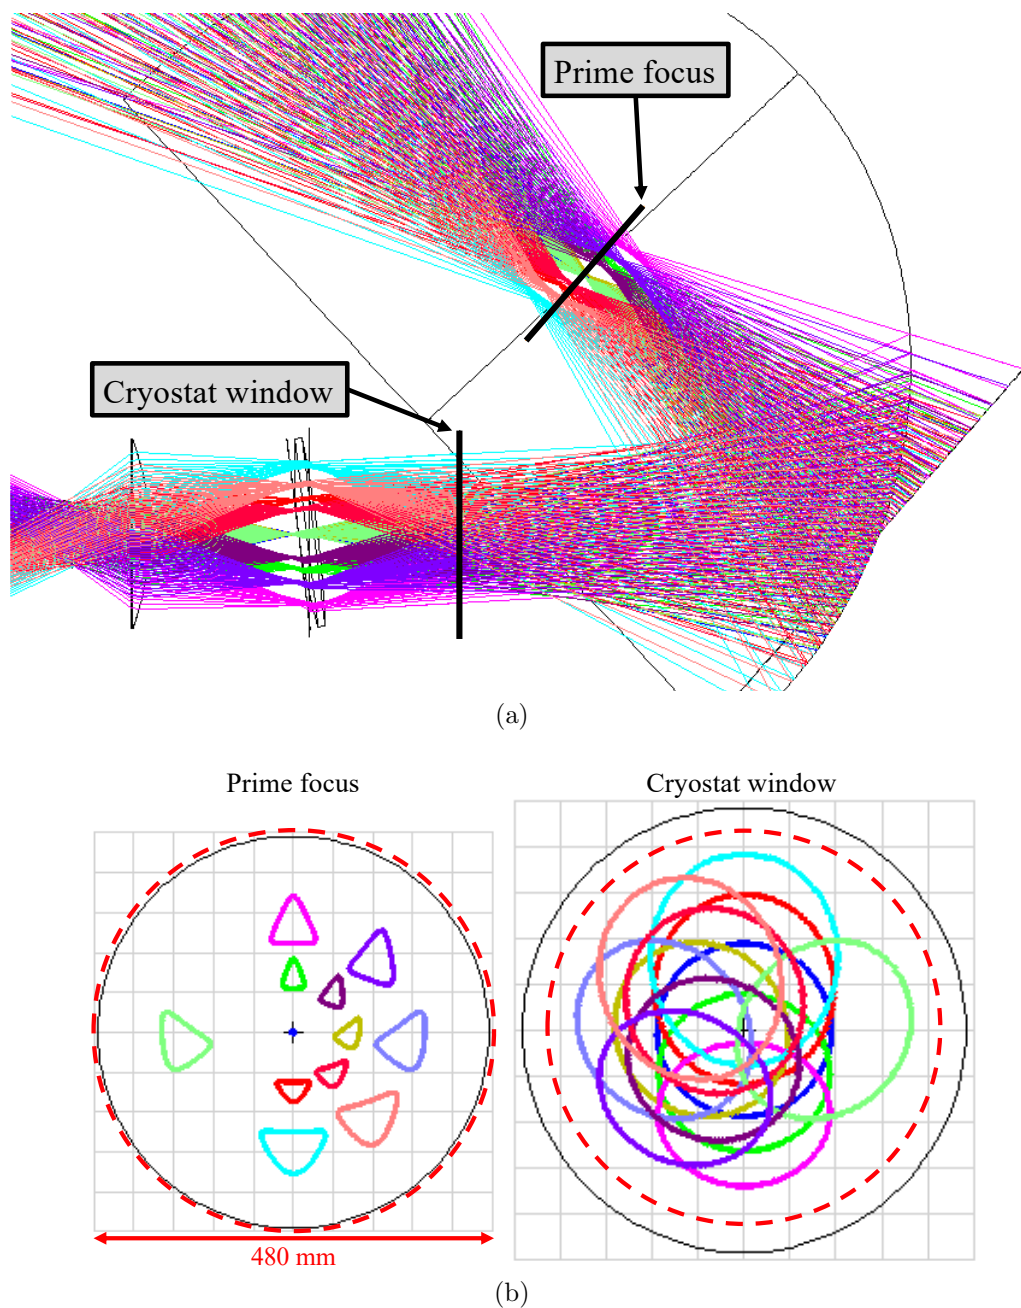


Figure 7.1: A comparison of the prime-focus and cryostat-window WHWP locations. Figure 7.1a is a Zemax (<https://www.zemax.com/>) ray-trace simulation of the PB-2a optical system, including the telescope mirrors and reimaging optics. Each ray color corresponds to a distinct detector pixel on the focal plane. Figure 7.1b shows the **ray footprint** at both the cryostat window and prime focus planes, and the WHWP's 480 mm clear aperture is marked with a red dotted circle. Each pixel's illumination pattern is larger at the cryostat window location, mitigating HWPSSs due to non-uniformities in the HWP's AR coating or sapphire stack. All Zemax outputs are courtesy of Fred Matsuda.

azimuth- and elevation-dependent signals in the detector data that, while removable, degrade signal-to-noise, especially on large angular scales. Third, placing an optical element between the mirrors of a crossed Dragone configuration breaks the **Mizuguchi-Dragone (MD) condition** [138, 47], a condition that enables the secondary mirror to correct cross polarization induced by the primary. This **MD breaking** was studied in detail for the PB-1 system [130] and is expected to be larger in PB-2a due to a larger FOV and the AHWP’s frequency-dependent angle.⁷

For these reasons, among others, the PB-2a WHWP is located at the second possible location directly in front of the receiver cryostat. While this location requires a larger **clear-aperture diameter**,⁸ it has a few key advantages over the prime-focus location. First, it is more broadly illuminated by each detector pixel, as shown in the right panel of Figure 7.1b. This wider illumination pattern better averages each detector’s beam over any WHWP defects, hence reducing HWPSSs. Second, the vacuum-window location is both close to the cryostat and deep within the telescope’s comoving baffling. Therefore, any stray light induced by the WHWP is, in the reverse-time sense, much more likely to land on cold surfaces, hence reducing optical loading on the detectors. Additionally, any light scattered by the WHWP is more likely to terminate on the telescope’s comoving baffling, mitigating ground pickup. Finally, an AHWP at the cryostat window does not break the MD condition and hence induces less cross polarization.

The WHWP’s clear-aperture requirement is determined using the **ray footprint** at the WHWP’s aperture plane, shown in the right panel of Figure 7.1b. The diametrical extent of the rays from detector pixels across the focal plane is 430 mm, and in order to provide margin for diffraction effects and alignment tolerances, we set a clear aperture requirement of ≥ 480 mm.

7.1.2 Emissivity and reflectivity

The most impactful consequence of moving the PB-2a HWP from 4 K to 300 K is the corresponding increase in mm-wave power on the detectors. As discussed in Chapter 3, a CMB instrument’s noise-equivalent temperature (NET) increases sharply with **in-band optical power**,⁹ and therefore limiting WHWP thermal emission is central to mitigating sensitivity degradation. Along a similar vein, limiting reflectivity is critical to the WHWP’s success, not only to maximize throughput but also because WHWP-induced reflections are susceptible to landing (in the reverse-time sense) on 300 K surfaces,¹⁰ further increasing in-band optical loading. While emissivity and reflectivity are independent physical effects, we address

⁷MD breaking with an AHWP in PB-2a/b has not been studied in detail, so the cross polarization in PB-1 [130] can be taken as a lower limit.

⁸The clear-aperture diameter defines a circle that is optically unobstructed.

⁹In-band optical power is the optical power within the detectors’ observation bands, which for PB-2a are centered at 90 and 150 GHz and have bandwidths of ≈ 30 and ≈ 40 GHz, respectively.

¹⁰“Warm reflections” are less of a concern for cryogenic optics, which are deep within the cryostat and therefore more likely to reflect onto cold surfaces.




Figure 7.2: The impact of WHWP emissivity and reflectivity on PB-2a’s mapping speed in the 90 and 150 GHz bands. The cyan stars mark the measured values in the table of Figure 7.13.

them together because the AR coating optimization, presented in Section 7.2.2, considers trade-offs between the two. Figure 7.2 shows PB-2a’s mapping speed (see Section 3.10) vs. WHWP emissivity and reflectivity. As demonstrated by the contour’s orientation, instrument sensitivity depends more strongly on emissivity than on reflectivity.

7.2 Optical design

The PB-2a WHWP’s **optical stack** is designed for high linear polarization modulation efficiency, low emissivity, and low reflectivity. It consists of three 512 mm-diameter, 3.75 mm-thick, α -cut¹¹ sapphire plates manufactured at Guizhou Haotian Optoelectronics Technology (GHTOT).¹² The plates are cut and ground with an α -plane alignment of $\pm 2^\circ$, a surface parallelism of $\pm 100 \mu\text{m}$,¹³ and a surface roughness of $\sim 0.1 \mu\text{m}$ RMS.

The AR coating comprises a top layer of 380 μm -thick **high-density polyethylene (HDPE)** from New Process Fibre (NPF)¹⁴ and a bottom layer of 270 μm -thick **RT/duroid**

¹¹An α -cut sapphire plate’s ordinary axis is normal to its surface so that normally incident light experiences optical birefringence. Another common orientation is c-cut sapphire, whose *extraordinary* axis is normal to its surface. This latter configuration is not birefringent and is often used for sapphire windows, which often aim to preserve incident polarization properties. Because there are two α -axes and only one c-axis, it is easier to grow large-diameter α -cut sapphire windows than c-cut ones.

¹²GHTOT: <http://www.ghot.com/>

¹³Surface parallelism is determined at the manufacturer using a combination of a deep-throat micrometer and a coordinate-measuring machine (CMM).

¹⁴NPF: <http://www.newprocess.com/>




Figure 7.3: An exploded view of the PB-2a WHWP’s optical stack, which is composed of the Pancharatnam sapphire and a two-layer AR coating of HDPE and Rogers RT/duroid RO3006.

RO3006 circuit board laminate from Rogers Corporation.¹⁵ The AR layers are pressed onto the sapphire using ambient pressure, a technique that eliminates thermal emission from glue layers and avoids potential AR delamination during WHWP operation in the field. The table in Figure 7.10 presents the measured parameters of the optical stack. The sapphire and AR indices were obtained using a Fourier-transform spectrometer (FTS) in Figure 7.13, and the loss tangents were obtained using the thermal emission apparatus described in Section 7.2.1. In the following subsections, we discuss the details of the WHWP’s optical stack.


7.2.1 Sapphire

PB-2a selects **sapphire** as its WHWP’s birefringent substrate. Sapphire is single-crystal aluminum oxide (Al_2O_3), and it has a large **differential index** $n_e - n_o \approx 0.35$ (see Equation 6.15) and small loss tangent at ~ 100 GHz, making it ideal for CMB applications. The PB-2a sapphire is fabricated by GHTOT, located in the Guizhou province of China. GHTOT produces among the world’s largest sapphire for an extraordinarily affordable price that is several times less than U.S. competitors.¹⁶ GHTOT’s sapphire boules,¹⁷ shown in the left panel of Figure 7.4a, are grown using the **heat exchanger method (HEM)** over 28 days. After growth, the boule is inspected for “smoke”—or microscopic defects such as bubbles or microcracks—the windows are cut from the boule using a wire saw, the crystal-axis orientation is measured using x-ray diffractometry, and the windows are machined to their final shape via **Blanchard grinding**. One of the most challenging aspects of the PB-2a sapphire manufacturing process is the windows’ high aspect ratio, which makes them vulnerable to cracking during grinding. Therefore, GHTOT carefully approaches the final thickness with several passes of the grinding wheel, giving precise control of the part’s thickness and improving manufacturing yield.

¹⁵RO3006: <https://rogerscorp.com/advanced-connectivity-solutions/ro3000-series-laminates/ro3006-laminates>

¹⁶An example U.S.-based large-diameter sapphire manufacturer is GT Advanced Technologies (<https://gtat.com/>).

¹⁷“Boule” is an industry term to describe the raw sapphire bulk after its growth.



(a)



(b)

Figure 7.4: Photographs of the GHTOT large-diameter sapphire. Figure 7.4a includes photos of a sapphire plate being manufactured in China, courtesy of Carl Johnson. The left panel shows a boule of sapphire after growth, and the right panel shows the Blanchard grinding process. Figure 7.4b shows a photo of the author holding a finished 512 mm-diameter sapphire plate, courtesy of Chris Raum. The window is ground to $\sim 0.1 \mu\text{m}$ surface roughness, giving it a frosty appearance.

Aside from cost and lead time, there are two primary considerations when assessing the usability of GHTOT's sapphire for the PB-2a WHWP. The first consideration is **optical diameter**. PB-2a's sapphire pieces are 512 mm in diameter, allowing the sapphire's fixture (see Section 7.3) to clamp $\Delta R = 5$ mm at the edge and still have a 500 mm clear-aperture diameter. This diameter enables the WHWP to meet its 480 mm clearance specification with margin for baffling and alignment tolerances. The second consideration is **sapphire purity**. To minimize its emissivity, the sapphire should be free of contaminants such as

titanium, chromium, and vanadium.¹⁸ Clean sapphire growth is challenging, and therefore purity was a fixation of PB-2a’s sapphire selection process. GHTOT measures impurity levels using **glow discharge mass spectroscopy (GDMS)**, and the bulk has < 1 ppm of all but tungsten, of which the growth crucible made. A boule of GHTOT sapphire is shown in Figure 7.4, and its neutral color indicates excellent purity.

7.2.2 Anti-reflection coating

In order to meet PB-2a’s performance targets in both its 90 and 150 GHz bands, the WHWP employs a dual-layer AR coating, shown in Figure 7.3. The AR coating’s primary objective is to minimize reflectivity, and because high optical throughput is critical to an experiment’s sensitivity, AR coatings are an active research area within the CMB community. Chapter 10 discusses AR coatings for cryogenic optics and presents a comprehensive overview of techniques to minimize reflection, but in this section, we focus on the WHWP AR coating’s construction and its impact on instrument sensitivity.

The WHWP’s AR coating consists of two layers that gradually impedance match from the index of air $n_{\text{air}} = 1$ to that of the sapphire $(n_e + n_o)/2 \approx 3.21$. The optimal AR indices are $n_{\text{bot}} = 2.31$ and $n_{\text{top}} = 1.40$ (see Figure 10.2),¹⁹ and their ideal thicknesses are $\lambda_c/(4n)$, where λ_c is the AR band’s central frequency. For PB-2a, $\nu_c = c/\lambda_c = 124$ GHz is directly between its 90 and 150 GHz detector bands. In addition to achieving low reflectivity, the AR layers must be low-emissivity to limit parasitic optical loading. According to Figure 7.2, the coating’s emissivity is more important than its reflectivity, and therefore the PB-2a WHWP’s AR coating design prioritizes materials with low mm-wave absorptivity.

The available materials for an ambient mm-wave AR coating are limited. Organic compounds, which are prevalent in commercially available products, absorb at ~ 100 GHz, and therefore our search is restricted to **inorganic materials** such as polyethylenes, fluoropolymers, and ceramics. The table in Figure 7.5 presents a catalog of candidate materials measured using a 50 \sim 300 GHz FTS at UC Berkeley. Upon simulating the performance for all two-layer combinations of the evaluated materials, we select the RT/duroid (**Duroid**) RO3006 circuit board laminate²⁰ as our bottom layer and high-density polyethylene (HDPE) as our top layer. RO3006 is a matrix of PTFE loaded with alumina grains to tune its dielectric constant, while the selected HDPE is manufactured by NPF via a high-precision extrusion process. Both materials have low loss in PB-2a’s frequency range and have reproducible refractive indices, making them suitable AR solutions.

¹⁸Pure single-crystal aluminum oxide is *technically* called **corundum**, but the standard industry term is “sapphire” no matter the purity.

¹⁹A *fully* optimal two-layer WHWP AR coating would itself be birefringent, matching to $n_o = 3.08$ along the ordinary axis and $n_e = 3.35$ along the extraordinary axis. Such birefringent coatings are difficult to fabricate using conventional techniques but are an attractive feature of metamaterial HWP’s [37].

²⁰The Duroid sheets are manufactured with copper cladding into which circuit patterns are typically etched. Therefore, before using the sheets, we strip the copper using large baths of ferric chloride solution.

Material	Measured n	Measured $\tan \delta$ [10^{-4}]
PTFE	1.45 ± 0.01	3.0 ± 1.5
LDPE	1.50 ± 0.01	10.2 ± 2.4
HDPE	1.55 ± 0.01	0.5 ± 1.0
RT5880	1.56 ± 0.01	16.0 ± 1.2
UHMWPE	1.58 ± 0.01	2.6 ± 6.2
RO3003	1.77 ± 0.03	21.1 ± 0.8
RO3006	2.52 ± 0.01	56.5 ± 2.7
RT6006	2.86 ± 0.02	89.9 ± 1.7




Figure 7.5: The table presents measured indices and loss tangents for candidate AR materials using UC Berkeley’s FTS. Because the films are thin, they were pressed onto an alumina substrate, and the spectra with and without the AR film are compared. The selected materials for the WHWP are bolded: HDPE for the top layer, RO3006 for the bottom layer. Figures 7.5a and 7.5b show measured FTS spectra between 80 and 220 GHz for the RO3006 (left panel) and HDPE (right panel). Blue lines show the alumina + film fringe, the green line shows the alumina-only fringe, and the red line shows the fit using transfer-matrix calculus [83]. The “fast” fringe pattern is due to the ~ 6 mm-thick alumina substrate, and the “slow” modulation of the fast fringe is used to extract the AR layer’s refractive index.

As shown in Figure 7.5, HDPE has an index of $n_{\text{HDPE}} = 1.55$ while Duroid has an index of $n_{3006} = 2.52$. Both of these indices are substantially larger than the targets of $(n_{\text{top}}, n_{\text{bot}}) = (1.40, 2.31)$, increasing reflectivity with respect to the optimal case. In addition, because thin plastics are difficult to machine,²¹ PB-2a employs the HDPE and RO3006 at their stock thicknesses, which are both ≈ 25 μm thicker than optimal. While slightly degrading reflection performance, this programmatically driven design feature facilitated

²¹The difficulty of machining the WHWP’s AR coatings is partly because we do not glue them down. If we had affixed the coatings to the rigid sapphire windows, they might have been more easily trimmed using a milling machine.




Figure 7.6: The impact of a potential epoxy glue layer, assumed to have an index of $n_{\text{glue}} = 1.7$ and an ambient-temperature loss tangent of $\tan \delta = 0.01$, on the emissivity and reflectivity of the WHWP. Because the glue layers create index mismatches, reflectivity increases sharply with glue thickness, especially at 150 GHz.

the WHWP’s on-time arrival in Chile.

Because minimizing emissivity and reflectivity is central to the AR coating design, the impact of glue layers is a major concern for the WHWP’s construction. A common practice throughout the CMB community is to affix plastic films using adhesives such as epoxy or rubber cement.²² While adhesives are relatively straightforward to implement during WHWP fabrication, they introduce several complications to the optical assembly. First, the glue layers are **index-mismatched** dielectrics within the WHWP’s optical stack. In order for these layers to have a negligible impact on reflectivity, they must be much thinner than the detected wavelength. For example, generic epoxy²³ typically has a refractive index of $n = 1.7$, which creates index mismatches throughout the WHWP that are especially large between the sapphire windows.²⁴ In addition, adhesives are quite lossy at room temperature ($\tan \delta \sim 10^{-2}$), which also necessitates a very thin layer to avoid increasing the WHWP’s emissivity. Figure 7.6 shows that even a 20 μm glue layer substantially increases both the

²²Low-density polyethylene (LDPE) is also often employed by heat pressing the AR coating onto the optic above the LDPE’s melting temperature. This technique requires substantial R&D for 500 mm-diameter optics and therefore was not pursued for PB-2a.

²³“Generic” describes epoxies without any fillers, such as Stycast 1266.

²⁴Because sapphire is not infinitely rigid, clamping the edge of an unaffixed WHWP causes the first and third pieces in the Pancharatnam stack to bow away from the second, creating $\sim 100 \mu\text{m}$ gaps between the sapphire windows. Therefore, adhesives are needed within the sapphire stack as well as within the AR coating.




Figure 7.7: A cross-sectional view of the AR vacuum module. The HDPE AR layers are glued to aluminum annuli that mate to o-ring seals on the rotor. When the module is evacuated, atmospheric pressure presses the optical stack together, eliminating the need for glue layers.

WHWP’s reflection and absorption. Second, adhesive layers are difficult to make uniform,²⁵ and therefore a glue layer is most likely to generate HWPSSs due to position-dependent emissivities and/or reflectivities. Third, ambient temperatures fluctuate between -20° and 20° C at the Chilean observation site, and every time the WHWP warms (cools), it expands (contracts) and the glue layers become stressed due to differential thermal expansion (contraction). This repeated stress can lead to **progressive AR delamination** that is often difficult to detect during real-time data quality assessments.

In order to avoid adhesive layers, we press the AR coating and sapphire windows together using a **vacuum bagging** technique. A cross-sectional CAD drawing of the AR coating assembly is shown in Figure 7.7, and a photo of the fully-assembled WHWP, including the AR coating, is shown in Figure 7.8a. The bagging method creates a vacuum space using the outermost HDPE layers such that **ambient pressure** forces the sapphire stack and AR layers together. Each HDPE sheet is glued using Stycast 2850FT²⁶ to an aluminum “window clamp” ring, and the ring interfaces to the rotor stage via an o-ring seal. After the vacuum bag is assembled, it is pumped out via a valved vacuum port on the rotor stage, and the WHWP is ready to operate. Maintaining sufficiently low pressure for \sim months of continuous WHWP operation is challenging due to the vacuum chamber’s relatively small volume. Therefore, a tubular chamber was added to increase the chamber’s volume, as shown in Figure 7.8b.

7.3 Mechanical assembly

The WHWP’s mount and rotation stage must spin the 510 mm-diameter optical stack at 2 Hz more than 100 million times throughout PB-2a’s lifetime. Furthermore, the WHWP’s

²⁵This reality contrasts to other layers in the WHWP’s optical stack, which are specifically fabricated for uniformity.

²⁶Stycast 2850FT is often used to seal vacuum spaces, and because its CTE is matched to that of aluminum, the integrity of its seal is robust to variations in ambient temperature.




Figure 7.8: Photographs of the PB-2a WHWP's mechanical assembly. Figure 7.8a shows the fully assembled WHWP on its test stand at UC Berkeley. Key hardware elements include a vacuum module to apply the AR coating, a servo motor coupled to a timing belt for modularity and rotational stability, and vibration-damping mounts to mechanically isolate the HWP from the receiver cryostat. Figure 7.8b shows the WHWP mounted on the PB-2a telescope. The secondary mirror is covered with blue foam, and the receiver cryostat was not yet installed during this test fit. The tubular structure on the rotor adds volume to the AR vacuum module.

assembly must minimize vibrational coupling to the telescope and receiver, maintain adequate rotational stability, accommodate in-field optical alignment, and be robust against mechanical failure. In this section, we describe several key features of the PB-2a WHWP's mechanical assembly, shown in Figure 7.8.

7.3.1 Mount

The WHWP's mount is designed to maximize configurational flexibility and minimize rotation-induced receiver vibrations. The mount attaches to the telescope's boom via a series of 80/20 rails and rotational stages, allowing the HWP's position to be adjusted along three translational and two rotational axes. Additionally, the mount is capable of positioning the HWP either in front of the cryostat window or between the primary and secondary mirrors, as in PB-1. This locational flexibility allows for in-situ evaluation of HWP-related aberrations, instrumental polarization, and beam coverage. In order to ensure that WHWP-induced vibrations are isolated from the receiver cryostat, the assembly is coupled to the telescope using a series of independent rubber sandwich mounts oriented both tangentially and axially to the WHWP's rotation axis. To isolate the sapphire from vibrations in the bearing, the

optical stack is clamped by a thin rubber gasket.

7.3.2 Rotation stage

The WHWP’s rotation stage comprises a **large-bore bearing**, a timing-belt-coupled gear system, and a servo motor. For rotor-stator coupling, we employ a 635 mm-diameter matched-pair **thin-section ball bearing** from SilverThin Bearings.²⁷ This bearing is a scaled-up version of the one used on PB-1 and is chosen for its demonstrated success at the observation site. The bearing is preloaded to meet product specifications and is lubricated with a low-temperature-compatible grease to accommodate on-site weather conditions. The bearing’s races are stainless steel, and the bearing clamps are 7075 aluminum to minimize weight. To validate the bearing’s mechanical compatibility with Chilean conditions, we perform stress and thermal tests to ensure robustness against premature wear and seizing.

The drivetrain is a 400 W, 60 mm **AC servo motor** from Applied Motion Products.²⁸ This motor accommodates our estimated peak- and continuous-torque requirements with a safety factor of ≈ 3 and avoids electrical switching noise present in comparable stepper motors. The servo is tuned to provide loose feedback, keeping the motor from overheating and utilizing the rotor’s large moment of inertia to maintain stable rotation. The drivetrain is connected to the rotor via a Kevlar-reinforced timing belt and a commercial tensioner pulley, allowing for easy installation of replacement parts at the observation site. When on the telescope, the WHWP operates in the receiver cryostat’s **weatherproofed enclosure**, keeping it safe from wind and snow but not from fluctuations in ambient temperature. Therefore, all WHWP components are rated to -20°C .

7.3.3 Angle encoder

To acquire the WHWP’s angle, we utilize a **commercial optical encoder** from RSF Elektronik²⁹ that consists of a steel tape with 10,000 reflective lines and an infrared read head, shown in Figure 7.9a. The encoder tape is mounted on a 636 mm-diameter aluminum ring whose surface maintains a radial distance of $0.75 \pm (0.4, 0.2)$ mm from the read head during WHWP rotation. This system provides 6.5 arcsec resolution via $4\times$ interpolation.

To read out the encoder, we use an **Arduino Leonardo ETH**³⁰ microcontroller. The Arduino samples the encoder signal at 16 kHz to achieve WHWP angle resolution of

$$\sigma_x = \frac{30\ \mu\text{rad}}{\sqrt{12}} \sqrt{\frac{1}{16 \times 10^3\ \text{Hz}}} = 0.07 \frac{\mu\text{rad}}{\sqrt{\text{Hz}}}, \quad (7.1)$$

where $30\ \mu\text{rad} = 6.5\ \text{arcsec}$ is the encoder’s resolution and $1/\sqrt{12}$ is the standard deviation of a uniform distribution, quantifying the Arduino’s sampling randomly between encoder ticks.

²⁷SilverThin Bearings: <http://www.silverthin.com/>

²⁸Applied Motion: <http://www.applied-motion.com/>

²⁹RSF: <http://www.rsf.at/en/>

³⁰Arduino Leonardo: <https://store.arduino.cc/usa/arduino-leonardo-eth>




Figure 7.9: The PB-2a WHWP’s encoder mechanism and performance. Figure 7.9a shows a photograph of the encoder tape and read head. The read-head gap tolerance is $\pm (0.4, 0.2)$ mm, necessitating $< 0.03\%$ ellipticity for the encoder tape’s 635 mm-diameter mounting ring. Figure 7.9b shows a power spectral density of the HWP’s rotation at ≈ 2.2 Hz for 5 min. The drivetrain’s excellent rotational stability demonstrates the effectiveness of the servo + timing belt system.

This angle noise level meets the $\ll 3 \mu\text{rad}/\sqrt{\text{Hz}}$ requirement presented in Section 6.7.3. After collecting a preset number of encoder ticks, the Arduino then compiles a UDP packet containing the WHWP’s angle and a universal timestamp.³¹ It then sends the packet to a central computer where the angle data is verified and paired with matching detector data. The Arduino eliminates the need for a WHWP-dedicated computer and can be replaced quickly and cheaply if it fails at the site.

7.3.4 Mechanical validation

To validate the WHWP’s rotation mechanism, we evaluate vibration damping, bearing robustness, rotational stability, and the vacuum module’s hold time. Figure 7.9b shows a power spectral density (PSD) for five minutes of angle encoder data while the WHWP rotates at ≈ 2.2 Hz. The sharpness of the central peak demonstrates outstanding rotational stability, validating the effectiveness of the servo/timing belt system. We also run the WHWP continuously in the lab for ~ 1 week to confirm that rotation does not impact the AR assembly’s vacuum seal. In addition, after moving it to the Chile site, we rotated the WHWP outdoors to ensure its robustness to ambient conditions, and we test fitted it to the telescope

³¹Simons Array is GPS synchronized using the **Inter Range Instrumentation Group (IRIG-B)** time code.

before the receiver was installed, as shown in Figure 7.8b. WHWP-induced vibrations of the telescope’s structure were found to be negligible.

7.4 Optical evaluation

As described in Section 7.1.2, the primary merit figures for the PB-2a WHWP’s optical performance are transmissivity, emissivity, and linear polarization modulation efficiency. To evaluate optical performance, we use two modes that are cross-checked for consistency. In the first mode, we measure the properties of each optical-stack component. The measured AR layers are presented in Section 7.2.2, which discusses the sapphire’s loss tangent, index, and crystal axis orientation. A compilation of the sapphire’s and AR coating’s optical constants are in the table in Figure 7.10, and these values are used to simulate the performance of the assembled WHWP. In the second mode, we measure both transmissivity vs. frequency and band-averaged modulation efficiency of the assembled optical stack directly, and the results are shown in Figures 7.13 and 7.14.


7.4.1 Sapphire index

In order to both assemble the Pancharatnam sapphire stack (see Figure 6.14) and measure the sapphire’s ordinary and extraordinary refractive indices, we must first identify the orientation of each window’s **crystal axes**. Crystal measurements are most commonly performed using **x-ray diffractometry**, but such techniques are unavailable for PB-2a due to the sapphire’s large diameter. Therefore, we leverage the birefringent window’s mm-wave polarization properties to identify its axes.³²

A schematic of the crystal-axis measurement setup is shown in Figure 7.11a. A **polarized thermal source** is coupled and aligned to a broadband polarimeter, and the sapphire window is rotated between them, modulating the detector’s output. The detected intensity vs. sapphire angle has two local maxima corresponding to where the ordinary/extraordinary axes align with the source/detector’s polarization axes. An example measurement is shown in Figure 7.11b. In this particular test, data is collected over 20° and is fit with a parabola that identifies the maximum to $< 0.5^\circ$. Including uncertainties in the alignment and polarization leakage of the input and output **wire grids**, the accuracy with which the axes are marked, and the precision with which the stack is assembled, the relative alignment between the plates’ crystal axes is $< 1.1^\circ$. This axis-orientation uncertainty drives the uncertainties of the expected modulation efficiency and phase, which are represented as shaded bands in Figure 7.10 and are small enough to maintain $> 97\%$ polarization efficiency across both PB-2a bands.

After identifying the orientation of the sapphire’s crystal axes, we measure each piece’s ordinary and extraordinary refractive indices. We use the FTS system shown in Figure 7.13a,

³²We need only determine the *relative* orientation of each plate’s crystal axes, instead of the absolute orientation.



Stack Element	Thickness [mm]	Index of Refraction	Loss Tangent [10^{-4}]
Top AR Layer: HDPE	0.38 ± 0.04	1.55 ± 0.01	0.5 ± 1.0
Bot AR Layer: RO3006	0.27 ± 0.02	2.52 ± 0.01	56.5 ± 2.7
Ordinary Axis	3.75 ± 0.01	3.05 ± 0.03	0.1 ± 1.3 at 95 GHz
Extraordinary Axis		3.38 ± 0.03	1.1 ± 1.3 at 150 GHz

Figure 7.10: Normal-incidence, Mueller matrix calculations of the PB-2a AHWP (green), compared to a single HWP (red). The shaded green region represents measurement and assembly uncertainties, and the magenta and cyan bands show the approximate PB-2a observation bands. Figure 7.10a shows modulation efficiency as a function of frequency. The dual-band-averaged value is 98%, while that of the single plate is 75%. Figure 7.10b shows the AHWP's modulation phase as a function of frequency, while that of a single HWP is uniformly zero [82]. The table shows the measured ambient-temperature thickness, index of refraction, and loss tangent for each layer in the AHWP's optical stack. Indices are measured using an FTS (Figure 7.13), and loss tangents are measured using a thermal emission apparatus (Figure 7.12). The sapphire's loss tangent is averaged over the ordinary and extraordinary axes.

which includes two aligned wire grid polarizers on either side of the sapphire window to isolate the crystal axis. The measured **Fabry-Pérot fringe pattern** is then fit to extract the indices, which are presented in the table of Figure 7.10. While attenuation is also estimated using the FTS data, the sapphire's loss tangent and thickness are so small that constraining $\tan \delta$ meaningfully is difficult. Nevertheless, a precise estimate of the sapphire's loss is critical to validating the WHWP, and therefore we perform a dedicated measurement of sapphire emissivity.

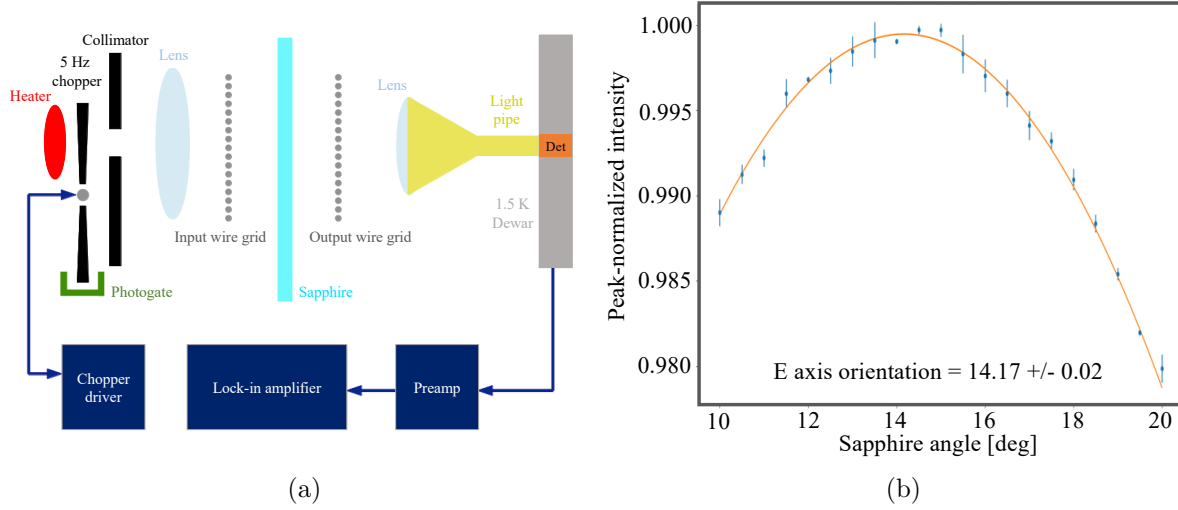


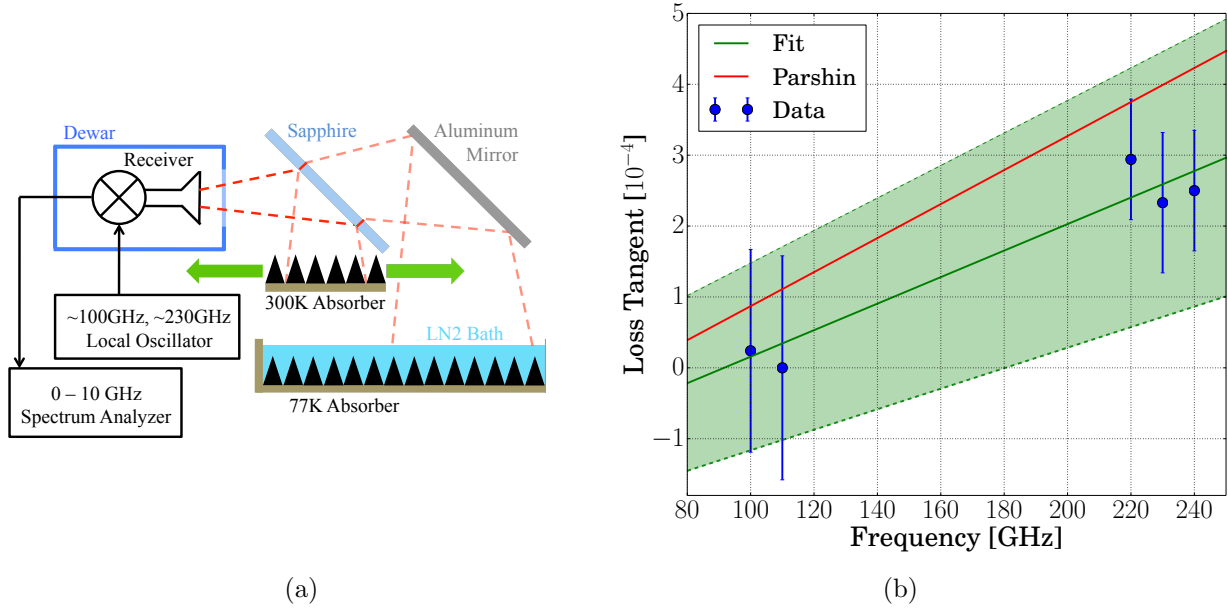
Figure 7.11: A measurement of a sapphire window’s crystal axis orientation. The measurement apparatus is shown in Figure 7.11a. A $\sim 700^\circ\text{C}$ ceramic thermal source is chopped at $\approx 20\text{ Hz}$ and collimated by an absorbing screen. A lens focuses the collimated beam through a wire grid onto the sapphire plate, and the output is polarized using another wire grid, is collimated by a lens and light pipe, and is detected by a neutron transmutation doped (NTD) Germanium bolometer at 1.5 K. The detector’s output vs. sapphire angle is shown in Figure 7.11b, and the maximum corresponds to the extraordinary axis’s location. This orientation is then marked and used to align the Pancharatnam stack. We also measure the ordinary index to gather statistics and better estimate uncertainties.

7.4.2 Sapphire loss

PB-2a’s mapping speed is quite sensitive to the WHWP’s emissivity, as shown in Figure 7.2, and the WHWP’s emissivity is quite sensitive to dielectric loss (see Section 3.4.1) in the sapphire. Therefore, a precise measurement of the sapphire’s loss tangent $\tan\delta$ in PB-2a’s frequency range is critical to the WHWP’s validation. Sapphire is expected to have $\tan\delta \sim 10^{-4}$ at 150 GHz and 300 K [153, 114], making an emissivity measurement of 3.75 mm-thick windows challenging.³³ To meet this challenge, we construct the **thermal emission apparatus (TEA)** shown schematically in Figure 7.12.

The TEA measures the temperature increase over a 77 K background due to thermal emission from a dielectric slab. We use 100 and 300 GHz **heterodyne receivers** previously used on the Combined Array for Research in Millimeter-wave Astronomy (CARMA) [20]. The receiver features horn-coupled superconducting-insulating-superconducting (SIS) mixers cooled to 4 K, fed by a tunable local oscillator (LO) to produce a 1-10 GHz **intermediate-frequency (IF) band** that is digitized with a spectrum analyzer.

³³CMB optical engineers often test thicker, smaller-diameter samples before purchasing the full-sized optic. However, sapphire’s quality is sensitive to the properties of its growth, and it is expensive to purchase small pieces from large-diameter boules. Therefore, we elect to “certify” the 512 mm-diameter, 3.75 mm-thick pieces directly.



Sapphire	Reflected	Transmitted	Purpose
-	-	77 K	Calibrate receiver noise temp
-	-	300 K	Calibrate receiver gain
Y	300 K	77 K	Calibrate reflection
Y	77 K	300 K	Calibrate transmission
Y	77 K	77 K	Measure emission

Figure 7.12: The PB-2a sapphire’s loss tangent at 300 K. Figure 7.12a shows the apparatus used to measure the sapphire’s loss at room temperature. Thermal emission from the sapphire slab is detected over a 77 K background using a heterodyne receiver. A series of configurations involving 300 K absorber are used to calibrate receiver gain, receiver noise, reflection from the sapphire, and transmission through the sapphire. Figure 7.12b shows a measurement of the IF-band-averaged, axis-averaged GHTOT sapphire $\tan \delta$ at five LO frequencies. The plot includes data points with error bars that account for both statistical and systematic uncertainties (blue points), a linear fit to those data points (green band), and data presented in Parshin *et al.* (red line) [153]. The table shows the TEA configurations used to characterize the sapphire’s loss.

Because the sapphire is not AR coated during this evaluation, it is important that radiation reflected (in the reverse-time sense) from the slab’s front surface terminates at 77 K. Thus, as shown in Figure 7.12a, the sapphire is mounted at a 45° angle to the receiver in front of a polished aluminum sheet also mounted at 45°. Both the sapphire and the mirror reflect to pyramidal absorber submerged in a liquid nitrogen bath. The sapphire is on a rotating stage, and the receiver is optimized for P polarization to maximize signal transmission through the slab. Measurements were made at LO frequencies of 100, 110, 220, 230, and 240 GHz. We made a series of five measurements at each frequency, which are listed in the table of Figure 7.12, in order to solve for the receiver’s noise temperature and gain as well as the sapphire’s reflection coefficient, transmission coefficient, and emissivity.

Several factors complicate the calculation of $\tan \delta$. First, the sapphire slab acts as a **Fabry-Pérot etalon**, creating a series of transmission peaks with ≈ 13 GHz spacing. Absorption, and hence thermal emission, in the slab is maximized at these frequencies and minimized at the reflection peaks between them. Second, measuring the Fabry-Pérot pattern is complicated by the **double-sideband receiver**, which folds signals above and below the LO frequency into a single IF band. Finally, the principal axes of the sapphire were not accurately known at the time of these measurements.³⁴ Simulations based on analytic techniques [52, 83] show that averaging the data across the IF band and across a range of sapphire azimuth angles incurs systematic error that is smaller than the noise in our measurement. Therefore, we present the band-averaged, sapphire-angle-averaged $\tan \delta$ values in Figure 7.12b. The measured GHTOT sapphire $\tan \delta$ is consistent with the literature value [153, 114] to within our 1σ uncertainty of $\approx 10^{-4}$.

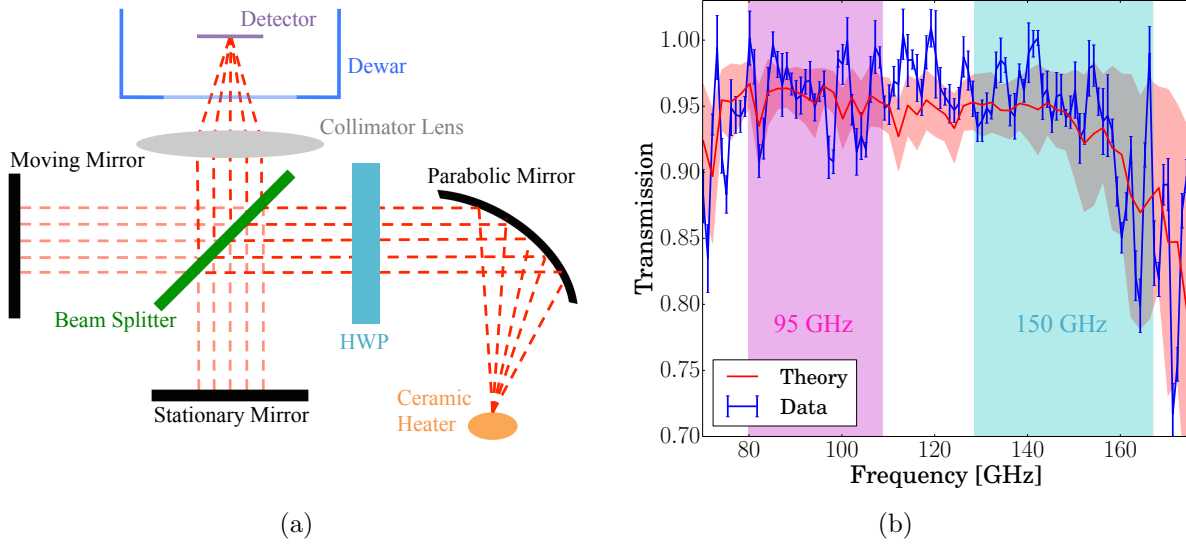
To estimate the sapphire’s $\tan \delta$ in the PB2 frequency bands, we linearly interpolate the measured $\tan \delta$ to 90 and 150 GHz, and the resulting values are given in the table of Figure 7.10. We find that the sapphire’s in-band absorptivity is smaller than that of the RO3006 AR layers.

7.4.3 Transmissivity

To validate the HWP’s AR performance, we utilize an **FTS** coupled to a **broadband detector**, shown schematically in Figure 7.13a. The signal from a temperature-modulated source is collimated by an off-axis parabolic mirror. This beam travels through the WHWP, is split between a pistoned and fixed path by a 250- μm -thick Mylar film, is recombined by the same splitter, and is focused onto a 0.3 K bolometer by an ultra-high-molecular-weight polyethylene (UHMWPE) collimator lens. The detector’s output is sent to a lock-in amplifier whose value is integrated for 0.5 s at each moving-mirror position. The moving mirror’s step size and dynamic range provide 1 GHz resolution over a 300 GHz bandwidth. We repeat the described process with the WHWP removed to **divide out** any spectral effects from the FTS setup and normalize the WHWP’s transmissivity. Additionally, we repeat the measurement at various HWP azimuth positions to average over any polarization induced by the chopper blade or input mirror.

The result of the FTS measurement is shown by the “Data” curve in Figure 7.13b. We integrate the measured transmission across each PB-2a frequency band to obtain the 90/150 GHz HWP transmissivities in the table of Figure 7.13. To estimate the emissivity, we simulate [52] HWP transmission vs. frequency using the dielectric layer parameters in the table of Figure 7.10 and isolate the loss due to absorption. The transmission simulation result, shown by the theory curve in Figure 7.13b, reveals that the simulated transmission is consistent with the measured transmission to within 1σ .

³⁴The crystal-axis measurements presented in Section 7.4.1 were performed after these sapphire loss measurements.



Band	Transmissivity	Emissivity
90 GHz	0.959 ± 0.014	0.020 ± 0.009
150 GHz	0.941 ± 0.015	0.032 ± 0.014

Figure 7.13: A schematic of the FTS used to measure PB-2a’s WHWP and a plot of the resulting transmissivity. Figure 7.13a shows a cartoon of the FTS apparatus. The signal from a temperature-modulated source is collimated, travels through the optical stack into the FTS, and is focused onto a 0.3 K bolometer. Figure 7.13b shows the measured transmission for the PB-2a WHWP between 60 and 180 GHz. The blue points represent the WHWP’s measured angle-averaged transmissivity, while the red line and shaded region represent the theoretical expectation using a transfer matrix calculator [52] and the individual-layer measurements in the table of Figure 7.10. The measured bandpass agrees with the expectation to within 1σ . The table presents the band-integrated transmissivity and emissivity across each PB-2a frequency channel.


7.4.4 Polarization efficiency

To measure the WHWP’s modulation efficiency, we utilize the setup shown schematically in Figure 7.14a. The signal from a chopped thermal source is polarized with an input wire grid, and the grid is tilted to avoid standing waves. The polarized beam travels through the WHWP, is again polarized by an output wire grid,³⁵ and is focused by an UHMWPE collimator lens onto PB-2a-style 90 and 150 GHz detectors. The WHWP’s angle is stepped in 10° increments, and the detector’s output is sent to a lock-in amplifier whose value is integrated for 1.0 s at each orientation.

The result of the modulation efficiency measurement in Figure 7.14b is fit to

$$S_{\text{det}} = \varepsilon \cos^2 \left[2(\chi - \phi) \right] + (1 - \varepsilon), \quad (7.2)$$

³⁵Even though the sinuous antenna is polarized, the output wire grid mitigates its frequency-dependent polarization angle [49].



Band	Modulation Efficiency	Phase Difference
90 GHz	0.989 ± 0.005	1.3 ± 0.1 deg
150 GHz	0.984 ± 0.004	

Figure 7.14: A measurement of the WHWP’s polarization modulation efficiency in PB-2a’s observation bands. 7.14a shows a cartoon of the polarized measurement apparatus. The HWP is mounted on a rotating stage to modulate a polarized thermal source with respect to a polarized detector. 7.14b shows normalized intensity as a function of HWP angle in each PB-2a frequency band. The points are the data, and the dotted lines are the fits using Equation 7.2. The 90 GHz modulation (in red) is slightly “ahead” of the 150 GHz modulation (in green) due to the WHWP’s frequency-dependent phase. The error bars for the measured data points are smaller than the plot’s marker size. The table presents the band-integrated modulation efficiencies across and the phase difference between the PB-2a bands.

where S_{det} is the normalized signal seen by the detector, χ is the WHWP’s angle, ϕ is its modulation phase (see Equation 6.28), and ε is its modulation efficiency (see Equation 6.22). To characterize any **cross polarization** in the setup, we remove the WHWP and step the input wire grid’s azimuthal angle in 10° increments over a 180° range. We fit this data to a model given by

$$S_{\text{det}} = (1 - L) \sin^2 [\chi_{\text{WG}} - \phi_{\text{WG}}] + L, \quad (7.3)$$

where S_{det} is the normalized signal seen by the detector, χ_{WG} is the wire grid’s azimuthal angle, L is its **polarization leakage**, and ϕ_{WG} is some phase. We find the setup’s leakage to be $< 1\%$ in both the 90 and 150 GHz bands.

After accounting for cross polarization, the band-averaged **polarization efficiencies** and **phase differential** are given in the table of Figure 7.14. Our measurement is consistent with the prediction from Figure 7.10a to within the 1σ uncertainty associated with the plates’

relative orientation and the detectors' frequency responses. Figure 7.14b shows that the HWP's polarization angle is controlled to $\approx 1^\circ$ between the PB-2a bands, which is small enough to ease angle calibration in the field.

7.5 Discussion

The presented WHWP has deployed to Chile and will see first light on PB-2a following ongoing commissioning of the detectors, readout, and optics. PB-2a's HWP is the largest deployed on any CMB experiment to date and is one of the first AHWP's to operate on a large-aperture telescope [37]. Therefore, a demonstration of PB-2a's data quality will represent a major step forward for polarization modulators for CMB observation.

While moving the PB-2a HWP from 4 K to 300 K accelerated modulator development, thermal emission from the sapphire and AR coating substantially degrade the instrument's mapping speed (see Figure 7.2). So even though the WHWP improves low- ℓ sensitivity, the experiment's *overall* sensitivity can be improved by cooling the AHWP to cryogenic temperatures. Therefore, PB-2b and PB-2c adopt **cryogenic HWPs**, which are the focus of Chapters 8-10.

Chapter 8

A cryogenic HWP for PB-2b

This chapter overlaps substantially with Hill, Kusaka, Ashton, Barton, et al. (2020) [79] and Hill, Kusaka, Barton, et al. (2018) [80], which we encourage the reader to cite instead of or in addition to this dissertation.

To improve its sensitivity, POLARBEAR-2b (PB-2b), the second installment of Simons Array (SA), implements a cryogenic half-wave plate (CHWP). Like PB-2a, PB-2b observes at 90 and 150 GHz and uses a sapphire Pancharatnam achromatic HWP (AHWP). Therefore, PB-2b is able to leverage much of the optical design and characterization in Chapter 7. However, there are several challenges associated with bringing the AHWP inside the cryostat that have driven substantial hardware advancements.

Similar to the PB-2a warm HWP’s introductory discussion in Chapter 7, we first outline the context surrounding the PB-2b polarization modulator. SA was funded in 2015 to build three POLARBEAR-2-style (PB-2-style) telescopes, at which point PB-2—which was funded in 2012 and was being built in Japan—became PB-2a, and PB-2b and PB-2c were born. PB-2b shares many common characteristics with PB-2a, including an identical telescope design, a nearly identical receiver cryostat design, alumina reimaging optics, lenslet-coupled sinuous antennas, and transition-edge sensor (TES) bolometers. However, there are a few areas where PB-2b improves over PB-2a, including higher-throughput anti-reflection (AR) coatings [97], improved detector wafers and readout [224], and a cryogenic HWP rather than a warm one [80, 79]. The PB-2b backend—which contains the focal plane infrastructure—was built at the University of California, San Diego (UCSD) [85, 84], while the optics tube was built at UC Berkeley. Motivated in part by its being a brand new component, the CHWP was developed at Lawrence Berkeley National Laboratory (LBNL) under the purview of Akito Kusaka. CHWP R&D launched in late 2014, the deployable CHWP was built and tested in 2016-2017, and the modulator was integrated into the PB-2b optics tube in 2018. The optics tube was mated to the backend at UCSD in 2019, and PB-2b was deployed to Chile in early 2020. As of December 2020, PB-2b work in Chile is stunted by the COVID-19 pandemic, but the CHWP is expected to see first light sometime in 2021. Therefore, the presented research represents nearly six years of designing, constructing, and testing and includes

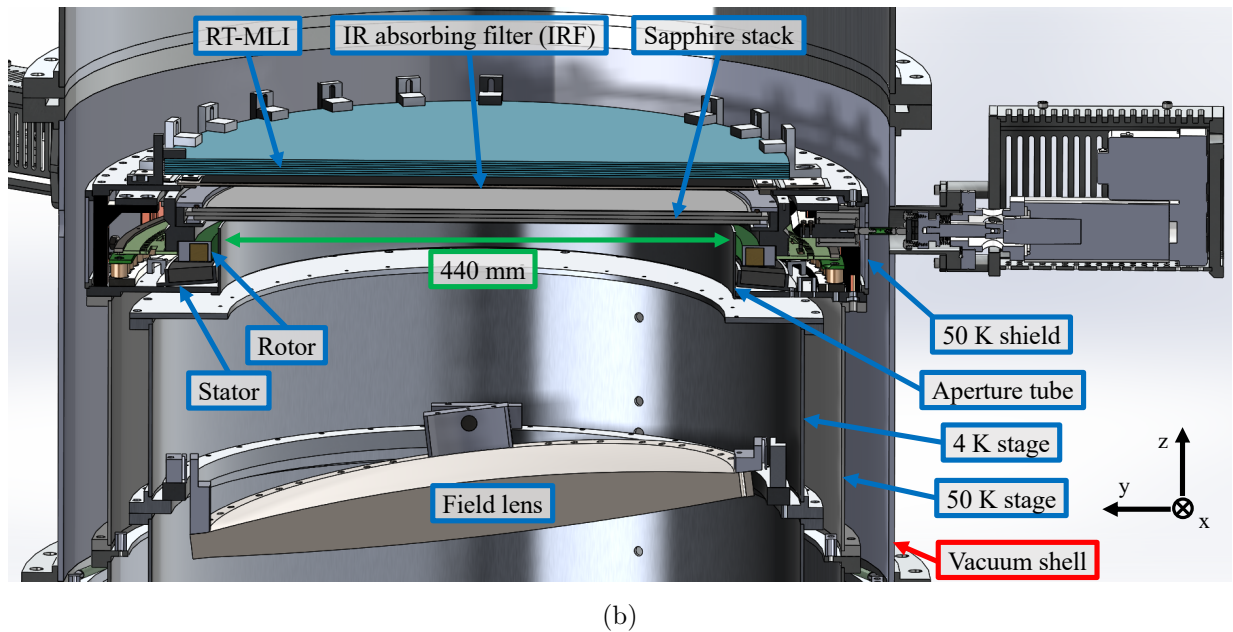
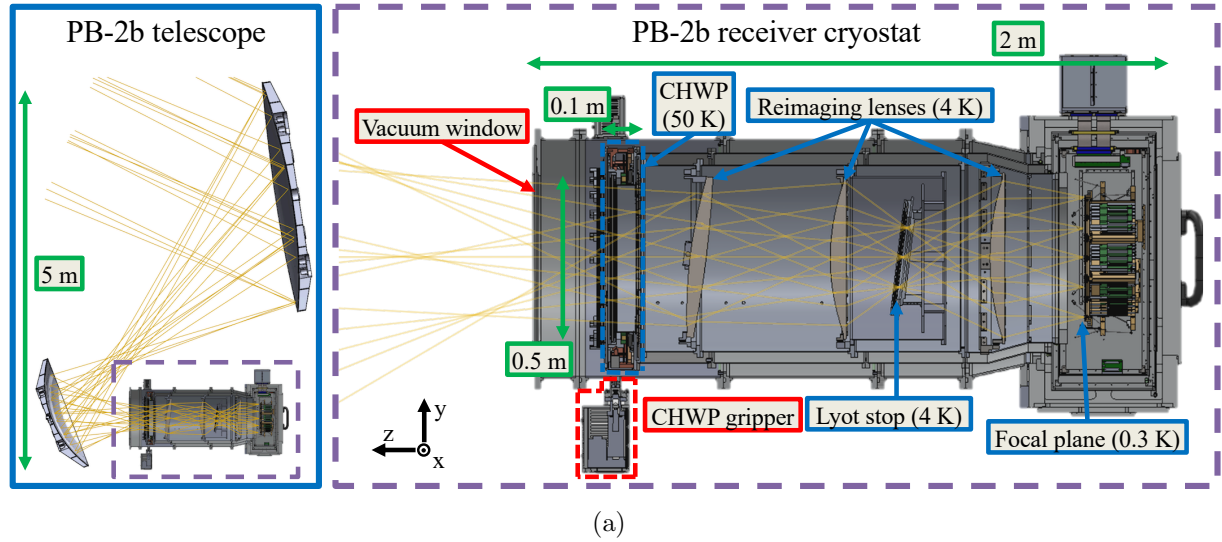


Figure 8.1: Computer-aided-design (CAD) drawings of the CHWP in PB-2b. Figure 8.1a shows an optical ray trace of the PB-2b telescope (left) and receiver cryostat (right), superposed onto CAD cross-sectional views. Red boxes mark warm components, and blue boxes mark cold components. The CHWP is located between the secondary mirror and the field lens and has a 440 mm clear-aperture diameter. Figure 8.1b shows a zoomed CAD cross-section of the CHWP integrated into the PB-2b optics tube, highlighting CHWP-relevant optical and thermal components. Red boxes mark warm parts, and blue boxes mark cold parts. The CHWP's clear aperture is defined by a stationary, HR-10-lined aperture tube (absorber not shown) that hides non-optical rotating components from the telescope's beam.

several advancements in the research area of optical modulators for CMB polarimetry.

While **stepped CHWPs** have deployed on a variety of mm and sub-mm experiments [139, 27, 6, 72, 26], the **continuously rotating CHWP** is an emerging technology [105, 98, 80, 169, 79]. Operating a 500 mm-diameter, tens-of-kilogram, spinning instrument in a cryogenic vacuum space poses a menagerie of challenges, including differential thermal contraction, frictional dissipation, and cryo-mechanical durability. The PB-2b CHWP design addresses these challenges with several hardware advancements.

As shown in Figure 8.1, the PB-2b CHWP is located within the receiver cryostat’s optics tube, which has a 0.8 m-diameter vacuum shell and two cryogenic stages cooled by a pulse tube refrigerator¹ (PTR) to ≈ 50 K and ≈ 4 K. The CHWP operates on the 50 K stage near the telescope’s Gregorian focus in front of three reimaging lenses and behind a vacuum window, radio-transmissive multi-layer insulation (RT-MLI) [31], and an alumina infrared absorbing filter (IRF) [93]. Its drive system consists of a superconducting magnetic bearing and a synchronous electromagnetic motor, and its rotation angle is monitored by a slot-chopped optical encoder. In this chapter, we describe the PB-2b CHWP’s design, and in Chapter 9, we present its evaluation in the laboratory.

8.1 Precedent

CHWPs have been deployed on several experiments prior to SA, and PB-2b’s CHWP design leverages this preceding work. Most notably, the PB-2b CHWP’s design concept relies on aspects of the E and B EXperiment (EBEX) continuously rotating CHWP [105, 132], and we briefly acknowledge this precedent before presenting PB-2b’s modulator.

EBEX was a balloon experiment that flew for 20 days above Antarctica in 2013 and observed at 150, 250, and 410 GHz [206]. They used a five-stack AHWP and a three-layer anti-reflection (AR) coating to accommodate their trichroic, broadband receiver design [134]. The EBEX CHWP is located on the 4 K stage just in front of the aperture stop, rotated on a superconducting magnetic bearing, driven by an external motor coupled to a Kevlar belt, and monitored by a slot-chopped angular encoder [105]. While EBEX’s telescope control system experienced critical failures² during flight [1], the CHWP successfully rotated 584,000 times at its nominal 1.235 Hz rotation speed and achieved a velocity stability of 0.45% per ten-hour observation [206]. Given this outcome, the EBEX CHWP is a success story upon which the PB-2b CHWP can build.

There are a few elements to the PB-2b CHWP implementation that are distinct from those of EBEX. First, the PB-2b CHWP is located on the 50 K stage, which relaxes its power dissipation requirements. Second, it is not located at the aperture stop (see the left panel of Figure 7.1b). Therefore, beams from detectors across the focal plane only partially overlap, heightening the importance of optical-stack uniformity to suppress HWP synchronous signals

¹Cryomech PT415: <https://www.cryomech.com/products/pt415/>

²It is worth noting that balloon experiments are “high-risk, high-reward,” where successful flights are powerful but “showstopping” issues are common. Therefore, EBEX’s outcome is not unique.

Parameter	Requirement	Achieved
Assembly outer diameter	≤ 700 mm	690 mm
Assembly height	≤ 100 mm	95 mm
Clear-aperture diameter	≥ 430 mm	440 mm
Rotor temperature T_{rotor}	≤ 55 K	< 53 K
Thermal dissipation P_{stator}	≤ 2 W	< 1.3 W
Rotor thermalization time	≤ 36 hr	5 hr
Rotation frequency f_{HWP}	≥ 2 Hz	≤ 2.8 Hz
Angle encoder noise σ_{χ}	$\ll 3 \mu\text{rad}/\sqrt{\text{Hz}}$	$0.1 \mu\text{rad}/\sqrt{\text{Hz}}$
B-field interference σ_B @ $4f_{\text{HWP}}$	$\ll 20 \mu\text{G}/\sqrt{\text{Hz}}$	$< 10 \mu\text{G}/\sqrt{\text{Hz}}$
T_{rotor} stability $\sigma_{T_{\text{rotor}}}$ @ 1 mHz	$\ll 1 \text{ mK}/\sqrt{\text{Hz}}$	$< 0.3 \text{ mK}/\sqrt{\text{Hz}}$

Table 8.1: The CHWP’s numerical requirements and achieved values.

(HWPSSs). Third, the PB-2b CHWP must undergo > 100 *million* revolutions as opposed to EBEX’s 100 thousand, elevating the importance of operational durability. And last, SA is projected to be a more powerful instrument than EBEX, and therefore the PB-2b CHWP’s systematic effects must be more tightly controlled to avoid degrading data quality. For these reasons and others, several aspects of the PB-2b CHWP design are fresh and novel.

8.2 Requirements

The CHWP must meet several optical, thermal, noise, and operational requirements to effectively mitigate $1/f$ noise and intensity-to-polarization (I-to-P) leakage while not introducing its own systematic effects. These requirements are central to the CHWP design and depend on the specifics of PB-2b’s telescope, cryogenic, and detector systems. When it was conceived, PB-2b was virtually a copy of PB-2a [133], which does not use a CHWP. As a result, the CHWP is retrofitted to the original PB-2b design and centers around maintaining “heritage” system performance. Table 8.1 summarizes the PB-2b CHWP’s numerical requirements and achieved values. In the following subsections, we detail the CHWP’s design drivers and how they flow into its hardware specifications.

8.2.1 Spatial and optical requirements

The initial set of requirements are on the CHWP’s physical dimensions. As shown in Figure 8.1, the assembly must fit within the confines of the optics tube’s vacuum shell, limiting its **cryogenic diameter** to ≤ 700 mm. Additionally, its height, or dimension along the z-axis, must be ≤ 100 mm so that the window aperture does not extend too far toward




Figure 8.2: The reverse-time-sense, peak-normalized, 90 GHz, x -polarization illumination pattern onto the CHWP aperture plane from detector pixels at the center (left panel) and $-y$ -edge (right panel) of the focal plane. The 430 mm aperture requirement is marked with a dashed cyan circle about the telescope’s chief ray, and the total spillover for the center (edge) pixel is 0.2% (0.1%).

the secondary mirror, where the telescope’s beam³ diverges more rapidly than within the receiver. These dimensional constraints influence much of how the CHWP subsystems are built and arranged, as well as their clearances and alignment tolerances.

While a presentation of sapphire-stack requirements is relegated to Chapters 6 and 10, we do discuss the **clear-aperture diameter**, which is central to the cryo-mechanical design. The CHWP is located sky-side of the field lens (see Figure 8.1a), and its clear aperture must be large enough to admit the beam between the secondary mirror and the Gregorian focus. We use GRASP⁴ to simulate the polarized, 90 GHz response of detector pixels across the focal plane, and Figure 8.2 shows the x -polarization illumination of the central and $-y$ -edge pixels at the CHWP aperture. To set a requirement, we vary the CHWP aperture diameter and compare the telescope’s **far-field beam** to that with the CHWP aperture removed. We evaluate a collection of merit figures, including angular resolution, spill over the primary mirror, side-lobe amplitude, and differential pointing and ellipticity [179]. We find that a 430 mm aperture has no impact on the central pixel and has edge-pixel impacts that are subdominant to preexisting systematic effects in the telescope optics. Therefore, we require that the CHWP aperture encapsulate a 430 mm-diameter circle about the telescope’s chief ray, including a margin for manufacturing and alignment tolerances.

Additionally, we require all rotating, non-optical CHWP components to be hidden from

³As shown in Figure 7.1a, the telescope’s beam is defined as the collection of reverse-time-sense rays from all detectors on the focal plane.

⁴GRASP: <https://www.ticra.com/software/grasp/>

the beam in order to minimize rotation-synchronous signals. To satisfy this requirement, the bearing and sapphire must be sufficiently oversized to encapsulate the **optical baffling** that defines the CHWP aperture, as shown in Figure 8.1b.

8.2.2 Thermal requirements

The next set of requirements are on the CHWP’s thermal impact. Because the PB-2b detectors are expected to be **photon-noise limited**, CHWP-induced optical power can dramatically degrade experiment sensitivity. Figure 8.3 shows the simulated fractional impact of **rotor**⁵ temperature on PB-2b mapping speed (see Section 3.10), compared to that of a PB-2a-style warm HWP (WHWP) (see Figure 7.2). Mapping speed quantifies the number of detector-hours needed to reach a specified CMB map depth, and therefore fractional mapping speed is analogous to detector yield or observation efficiency. In this simulation, the CHWP is assumed to have an epoxy-based AR coating (see Section 10.4), and sensitivity degradation with increasing CHWP temperature is primarily due to increasing **mm-wave emission**, which in turn increases detected photon noise. As shown by the WHWP points, some of the sensitivity loss at ≈ 270 K can be reclaimed by optimizing the AR coating for lower ambient emissivity (see Section 7.2.2), but the biggest sensitivity gain comes from cooling the HWP to cryogenic temperatures. Mapping speed retention is $> 97\%$ when the CHWP is at < 50 K, and the sensitivity loss vs. CHWP temperature is gentle < 100 K. Therefore, the dominant driver of the rotor’s temperature is not the CHWP’s mm-wave emission but instead its **IR emission** onto the field lens.

In the heritage PB-2b design without the CHWP, the 4 K field lens is radiatively loaded⁶ by the IRF (see Figure 8.1b). Therefore, if the CHWP becomes warmer than the IRF or if its dissipation onto the 50 K stage is too large, the CHWP-induced **4 K load** will exceed that of the heritage configuration, which may in turn increase lens and Lyot temperatures and even degrade the performance of the sub-kelvin stage. Section 8.6 discusses the CHWP thermal system in detail, and Figure 8.11 shows the relationship between rotor temperature and field-lens load. We require that the 4 K load not exceed that of the heritage configuration at 85% confidence, requiring the rotor temperature to be

$$T_{\text{rotor}} \leq 55 \text{ K} . \quad (8.1)$$

It is unavoidable that the CHWP’s motor and bearing will dissipate power onto its stator⁷, raising IRF and rotor temperatures. The goal, however, is to keep this temperature rise small enough to not warm 4 K components. Utilizing **load-curve**⁸ measurements of the PB-2b optics tube [84], we mandate that the CHWP warm the 50 K stage < 1 K with respect

⁵The “rotor” collectively refers to the CHWP’s rotating components.

⁶The term “load” is most often used to describe power dissipation in electrical systems but is also used to describe power transfer in thermal systems.

⁷The “stator” collectively refers to the CHWP’s stationary components.

⁸A “load curve” is a plot of an object’s temperature as a function of the load (power) on it.




Figure 8.3: The fractional impact of rotor temperature on CHWP emissivity (Figure 8.3a) and PB-2b mapping speed (Figure 8.3b) with respect to that of a 4 K HWP. For comparison, the emissivity and mapping speed impact of a PB-2a-style warm HWP (WHWP) at 273 K are marked with star points (see Figure 7.2). Note that the mapping speed gain of a 4 K CHWP over a 50 K CHWP is $\approx 1\%$.

to the heritage configuration and therefore set the CHWP’s 50 K dissipation requirement to be

$$P_{\text{stator}} \leq 2 \text{ W} . \quad (8.2)$$

We next require that the CHWP not add to PB-2b’s **cooldown time**. This mandate is important both for restarting observations following a disruption in the field and for rapid-turnaround receiver testing in the lab. The heritage PB-2b system takes about five days to cool, and that duration is limited by the cooldown of the focal plane,⁹ which reaches 4 K about 36 hours after the optics tube [84]. The CHWP’s cooldown time is limited by that of the rotor, which has a large thermal mass and is heatsunked by unfastened interfaces (see Section 8.3.4). Therefore, we require that the rotor reach its base temperature before the focal plane has thermalized, or within 36 hours of the 50 K stage.

Finally, we require that CHWP operation not vibrate the focal plane structure. CHWP-induced vibrations can generate microphonic heating that may modulate the focal plane’s temperature and degrade detector gain stability [84]. To avoid such issues, we require no measurable difference in focal-plane temperature between when the rotor is spinning and when it is not.

⁹The focal plane has a large thermal mass, as it consists of detector wafers, readout modules, and a beefy Invar frame. In addition, the focal plane is well isolated from the 4 K stage in order to reduce the load on the mK fridge during detector operation.

8.2.3 Noise requirements

As discussed in Section 6.7.2, for the CHWP to effectively suppress $1/f$ noise, $f_m = 4f_{\text{HWP}}$ must be large enough such that all modulation sidebands, which are contained within $4f_{\text{HWP}} \pm \Delta f$, are above the Atacama’s $1 \sim 2$ Hz knee frequency [199, 111], or the frequency at which $1/f$ noise power and white noise power are equal. Assuming a $0.4^\circ/\text{s}$ telescope scan speed [214], the sky is fully resolved by ≈ 4 Hz of temporal bandwidth [199]. Therefore, we require the CHWP’s **rotation frequency** to be

$$f_{\text{HWP}} \geq 2 \text{ Hz}, \quad (8.3)$$

which places the **demodulation band** $4f_{\text{HWP}} \pm \Delta f = 8 \pm 4$ Hz comfortably clear of atmospheric fluctuations. PB-1 and the Atacama B-mode Search (ABS) also rotate their HWPs at ≈ 2 Hz in Chile and achieve excellent $1/f$ suppression in their demodulated detector data [199, 111].

While rejecting optical $1/f$ noise, the CHWP introduces other noise sources that can degrade data quality. These CHWP-induced noise sources are largely **common-mode** across the focal plane and hence do not average down during detector coaddition. Therefore, we require the **noise-equivalent CMB temperature** (NET_{CMB}) (see Equation 3.62) of each CHWP noise source to be much less than PB-2b’s forecasted array-averaged NET_{CMB} [196]

$$\text{NET}_{\text{CMB}}^{\text{HWP}} \ll \text{NET}_{\text{CMB}}^{\text{arr}} = 5.8 \mu\text{K}_{\text{CMB}}/\sqrt{\text{Hz}}. \quad (8.4)$$

We use this bound to set requirements on three primary CHWP contaminants: encoder angle jitter, magnetic interference, and rotor temperature stability.

Firstly, **angle jitter** injects $4f_{\text{HWP}}$ noise into the demodulated detector timestreams, as discussed in Section 6.7.3, and the requirement for PB-2b is

$$\sigma_\chi \ll 3 \mu\text{rad}/\sqrt{\text{Hz}}, \quad (8.5)$$

which is identical to that of the PB-2a WHWP.

Secondly, CHWP-induced **magnetic interference** can also inject noise into the demodulated data. The PB-2b detector array consists of aluminum manganese (AlMn) transition-edge sensors (TESes) amplified by superconducting quantum interference devices (SQUIDs) (see Sections 2.4 and 2.4.2). Both AlMn TESes and SQUIDs are sensitive to ambient magnetic fields, and because the PB-2b CHWP consists of a magnetic bearing and an electromagnetic motor (see Section 8.3), magnetic interference at the focal plane must be carefully controlled.

PB-2b uses gradiometric Series SQUID Array Amplifiers (SSAAs) [193] fed back using Digital Active Nulling [67], stabilized by a low-frequency flux-locked loop [14], and shielded via a combination of μ -metal and superconducting niobium [233]. This particular SQUID configuration is robust to time-varying magnetic fields, especially those that are uniform over each SQUID’s $\sim 10 \text{ mm}^2$ area. Therefore, the more considerable CHWP-specific concern is magnetic interference in the detectors themselves.¹⁰

¹⁰Private communication with Tijmen de Haan (KEK High Energy Research Organization) and Darcy Barron (University of New Mexico), May 2020.

Magnetic pickup at $4f_{\text{HWP}}$ mimics modulated sky polarization, especially if the magnetic signal drifts over the course of an observation. The superconducting transition temperature T_c of AlMn varies slightly with ambient magnetic field B [43], and a change in T_c changes the bolometer’s saturation power P_{sat} , which in turn mimics a change in optical power P_{opt} (see Section 2.4.1). We can convert the CHWP-induced **magnetic jitter** σ_B to NET_{CMB} as

$$\text{NET}_{\text{CMB}}^{\text{HWP}} = \frac{dT_{\text{CMB}}}{dP_{\text{opt}}} \frac{dP_{\text{sat}}}{dT_c} \frac{dT_c}{dB} \sigma_B. \quad (8.6)$$

Measurements of the PB-2b AlMn TESes give $dT_c/dB = 0.3 \text{ mK/G}$ [222] and $dP_{\text{sat}}/dT_c = 0.1 \text{ pW/mK}$ [224], while simulations of PB-2b’s optics [81] give $dT_{\text{CMB}}/dP_{\text{opt}} = 12 \text{ } \mu\text{K}_{\text{CMB}}/\text{aW}$. We again require that $\text{NET}_{\text{CMB}}^{\text{HWP}} \ll \text{NET}_{\text{CMB}}^{\text{arr}}$ such that

$$\sigma_B @ 4f_{\text{HWP}} \ll 20 \text{ } \mu\text{G}/\sqrt{\text{Hz}}. \quad (8.7)$$

The final noise requirement is on the CHWP rotor **temperature stability**. Thermal emission from the CHWP can be modulated at $4f_{\text{HWP}}$ due to non-uniformities in the sapphire stack [170]. In turn, variations in CHWP temperature will modulate the amplitude of this thermal signal, mimicking polarized sky fluctuations in the demodulated detector data. At 50 K, the thermal $4f_{\text{HWP}}$ amplitude is expected to arise predominantly due to thickness and index variations in the AR coating, as sapphire is very transparent at mm wavelengths and low temperatures [153, 22]. While a more detailed study of AR-coating uniformity is relegated to Chapter 10, we set an empirical limit by once again invoking the PB-1 HWP’s $4f_{\text{HWP}}$ amplitude (see Equation 6.48), which is $\sim 10\%$ of its thermal emission [199]. We assume that the PB-2b CHWP’s $4f_{\text{HWP}}$ thermal signal will also be 10% of its thermal emission to set a conservative¹¹ requirement on its temperature stability.

Rotor temperature fluctuations $\sigma_{T_{\text{rotor}}}$ can be converted to NET_{CMB} as

$$\text{NET}_{\text{CMB}}^{\text{HWP}} = \frac{0.1\epsilon}{\eta} \frac{dT_{\text{CMB}}}{dT_{\text{RJ}}} \sigma_{T_{\text{rotor}}}, \quad (8.8)$$

where $\epsilon = 0.02$ is the assumed CHWP emissivity (see Section 10.2.2.2), $\eta = 0.9$ is the optical efficiency of the telescope + atmosphere between the CHWP and the CMB, and $dT_{\text{CMB}}/dT_{\text{RJ}} = 1.7$ is evaluated at 150 GHz. Noting that rotor temperature fluctuations follow a **1/f spectrum**, we require $\text{NET}_{\text{CMB}}^{\text{HWP}} \ll \text{NET}_{\text{CMB}}^{\text{arr}}$ on 1,000 s timescales—long enough to encapsulate many long-baseline telescope traversals—which imposes a rotor temperature stability requirement of

$$\sigma_{T_{\text{rotor}}} @ 1 \text{ mHz} \ll 1 \text{ mK}/\sqrt{\text{Hz}}. \quad (8.9)$$

¹¹While a substantial portion of the PB-1 $4f_{\text{HWP}}$ signal is due to polarized emission from the primary mirror (as discussed in Section 6.7.3), we assume here that it is entirely due to HWP non-uniformity in order to set a conservative requirement on CHWP temperature stability.

8.2.4 Operational requirements

As the only moving component inside the receiver cryostat, the CHWP poses unique risks to experiment operations. PB-2b intends to observe for at least three years in Chile, and the CHWP must operate robustly for as long as it is needed. This situational constraint gives rise to several operational requirements, which we highlight below.

First, the CHWP must be robust against **component degradation**. Because warming, opening, and cooling the receiver cryostat cost weeks of observation time, the rotor should be able to undergo $> 100,000,000$ revolutions without being serviced. This robustness requirement drives many aspects of the CHWP design, such as a zero-contact drive and multiple hardware redundancies. Second, standard CHWP operation must require no human intervention. Critical monitoring data, such as encoder data packet drops, rotational velocity, temperatures, and gripper status, are automatically logged. Additionally, the CHWP control software comprises a system of classes, configuration files, and command-line executables designed to interface with the observatory control program in order to fully automate routine operations. Rotor temperature monitoring is particularly important and is discussed in Section 8.6.1.


Third, we mandate an **automated shutdown** procedure to keep the rotor centered if site power is lost. Weather events are not uncommon on Cerro Toco and can cause sudden generator failures. If PTR cooling terminates and site network access is unavailable, the CHWP must automatically stop and stow such that it can reliably restart with minimal performance degradation. Shutdown and recovery testing was central to system evaluation and is discussed in Section 9.3.5. Fourth, in the event of critical equipment failures, an issue with the sapphire stack, or an otherwise unforeseen problem, the CHWP must be simple to service on the telescope. This requirement motivates the CHWP's location near the vacuum window where it can be accessed with minimal disassembly, as discussed in Section 8.7.

8.3 Cryo-mechanical assembly


Given the requirements presented in Section 8.2, we now discuss the design of the CHWP mechanism, shown in Figures 8.1b and 8.4. The CHWP consists of a rotor, stator, motor, encoder, and gripper, all integrated into a compact assembly on the 50 K stage. The following sections detail the designs of each subsystem.

8.3.1 Location

The CHWP is located on the **50 K stage**, which has several advantages over other possible 4 K locations, such as at the Lyot stop. First, the CHWP is skyward of all reimaging optics, mitigating the impact of I-to-P generated at the lens surfaces. Second, at 50 K, the sapphire stack's mm-wave emission is already deeply subdominant to that of other optical elements, limiting the sensitivity gain of cooling further (see Section 8.2.2). Third, the PTR's first stage has substantially more cooling power (~ 10 W/K) than its second (~ 1 W/K),



(a)



(b)

Figure 8.4: The CHWP rotation mechanism without the sapphire stack. Rotating components are labeled with cyan boxes, and non-rotating components are labeled with orange boxes. Figure 8.4a shows a CAD cross-section of the rotation mechanism with the rotor floating 5 mm above the stator. The YBCO ring is buckled due to residual thermal stress between the YBCO tiles and their aluminum fixture. Figure 8.4b shows a photograph of the rotation mechanism with the retaining ring removed and the gripper ports, gripper fingers, and flex circuit added. The magnet and YBCO rings are separated by three 4 mm shims. All internal wiring, including that of the solenoids, photodiodes, LEDs, and thermometers, exit the assembly via four 25-pin Micro-D connectors at the bottom of the photo.

relaxing CHWP dissipation requirements. Fourth, the focal plane structure is anchored to the 4 K stage but not the 50 K, easing CHWP vibrational requirements. Fifth, the 50 K location necessitates minimal optics-tube assembly/disassembly to install/access the CHWP, making it easy to both integrate into the receiver and service on the telescope, if necessary. Finally, the 50 K CHWP is **modular**, with its addition to the PB-2b instrument requiring no adjustments to the heritage lens shapes or positions, thermal filtering, or optics tube assembly procedure. This final feature was particularly powerful for PB-2b integration and testing, allowing the CHWP and optics tube to be validated in parallel prior to full system integration, hence accelerating the receiver commissioning process.

8.3.2 Optical design

As shown in Figure 8.1b, the CHWP sapphire stack consists of three ≈ 505 mm-diameter, 3.8 mm-thick sapphire windows (see Section 7.2.1) with a ≈ 0.7 mm-thick dual-layer AR coating on each of its outermost surfaces. The stack is held in an aluminum cradle, which uses tubular springs¹² to absorb **differential thermal contraction** and keep the sapphire from shifting during cooldown and rotation. The cradle has a 490 mm clear-aperture diameter, is 20 mm thick, and weighs ≈ 10 kg including the sapphire stack.

The CHWP’s clear aperture is defined by a stationary 440 mm-diameter tube lined with HR-10 absorber.¹³ This **aperture tube** is 60 mm tall and forms a 5 mm gap with the back face of the sapphire stack, “hiding” non-optical components from the telescope’s beam and restricting the propagation of any CHWP-induced stray light. The CHWP’s aperture diameter is limited by bearing manufacturing constraints (see Section 8.3.3) yet meets the 430 mm requirement while providing ± 5 mm of radial alignment tolerance.

8.3.3 Bearing

The CHWP bearing needs to be low-friction, have a large bore diameter, and be mechanically robust at low temperatures. Cryogenic ball bearings with a ~ 500 mm bore are not commercially available and are challenging to develop due to thermal contraction, vibration, and durability issues. Therefore, the PB-2b CHWP employs a **superconducting magnetic bearing (SMB)**, as shown in Figure 8.4. The SMB operates by **flux pinning** in an azimuthally symmetric geometry. When suspending a uniformly magnetized ring (the rotor) above a type-II superconducting ring (the stator) and subsequently cooling the superconductor below its transition temperature, the rotor’s permanent magnetic field becomes trapped in the stator’s superconducting bulk, constraining the rotor in the axial and radial directions while allowing it to rotate in azimuth. SMBs are implemented in other CHWP systems for CMB observation [132, 105, 168, 169, 167], and their cryogenic robustness is well demonstrated.

¹²Spira Manufacturing Corporation: <https://www.spira-emi.com/>

¹³Eccosorb: <https://www.laird.com/rfmicrowave-absorbers-dielectrics>

The PB-2b SMB is manufactured by Adelwitz Technologiezentrum GmbH.¹⁴ The magnet ring consists of 16 22.5°, 97 mm × 16 mm annular segments of Neodymium (NdFeB), glued contiguously into an encapsulating G10 fixture to form a highly uniform, $\approx 5,000$ G surface field. The superconducting ring consists of 46 7.8°, 35 mm × 13 mm annular segments of yttrium barium copper oxide (YBCO), which has a $T_c \approx 90$ K, glued into an encapsulating aluminum fixture to form a contiguous type-II superconductor. Both the rotor and stator have a 470 mm inner diameter, which is large enough to fit around the aperture tube with a 5 mm radial clearance (see Figure 8.1b).

The SMB’s effectiveness relies on a small separation between the permanent magnet and YBCO. SMB stiffness is a steep function of rotor-stator separation [91], and therefore controlling the rotor’s axial position near the YBCO’s transition temperature is critical to controlling the SMB’s spring constant. The nominal **rotor-stator separation** in the PB-2b CHWP system is 5 mm, for which we measure a **bearing spring constant** of ≈ 300 N/mm. In addition, SMB friction must be small enough to both achieve the required rotation speed and mitigate heat dissipation. **Eddy losses** are induced by eddy currents in nearby metal and scale as ΔB^2 [12], where ΔB is the rotor’s peak-to-peak magnetic field variation. **Hysteresis losses** arise due to paramagnetic hysteresis in the YBCO and scale as ΔB^3 [237]. Therefore, the uniformity of the magnet ring is critical to minimizing dissipation during continuous rotation. A measurement of rotor friction is presented in Section 9.3.4.

8.3.4 Gripper

Magnetic levitation is a cryogenic phenomenon, and the SMB only engages below the YBCO’s ≈ 90 K superconducting transition temperature. Above this temperature, the YBCO does not flux pin, and the rotor is effectively decoupled from the stator. Therefore, the CHWP employs a **gripper** to support the rotor during cooldown and keep it aligned until levitation initiates. The gripper consists of three subassemblies azimuthally distributed about the rotor (see Figure 9.2): one along $-y$, and the others ± 40 deg¹⁵ about $+y$. Each subassembly has a **linearly actuating vacuum feedthrough** at 300 K, a flexible, thermally isolating “arm” between 300 and 50 K, and a 50 K “finger” that engages the rotor stage. When the **gripper fingers** are extended, the rotor is gripped, and when they are retracted, the rotor is released.


Figure 8.5 shows a CAD rendering and photographs of a single gripper subassembly. The feedthrough¹⁶ consists of a linear actuator¹⁷ with a vacuum-compatible bellows assembly that mounts to a ConFlat port on the vacuum shell. The feedthrough bolts to a rubber spring assembly, which adds compliance both along and about the radial direction. This spring assembly in turn connects to the gripper arm, which is composed of a hinged thermal

¹⁴ATZ: <http://www.atz-gmbh.com/>


¹⁵The opening angle of the top two subassemblies is less than 120° to avoid mechanical interference with the optics tube’s PTR attachment, as shown in Figure 9.2.

¹⁶Huntington Mechanical Laboratories: <https://huntvac.com/>


¹⁷SMC Corporation: <https://www.smcusa.com/>



(a)



(b)



(c)

Figure 8.5: A detailed view of a single gripper subassembly with the retaining ring removed for visual clarity. Red boxes mark warm components, and blue boxes mark cold components. Figure 8.5a shows a CAD cross-section of a gripper subassembly with the rotor un-gripped and floating. Figure 8.5b shows a photograph of a feedthrough actuator attached to a gripper arm. Figure 8.5c shows a photograph of the linear bearing and the copper gripper finger without the gripper arm.

isolator joined to a 6 mm-diameter, 50 mm-long aluminum bearing shaft. The thermal isolator comprises two brass hinges epoxied to either end of a 24 mm-long, 0.8 mm-thick-walled hollow G10 tube. The hinges accommodate differential thermal contraction along the receiver cryostat's axial direction and relax alignment tolerances between the 300 K and 50 K stages.

The bearing shaft enters the 50 K assembly through a **gripper port**, which includes a 10 mm-diameter hole in the 50 K shield and a Frelon-lined linear bearing within which the bearing shaft slides. Because this port introduces a potential 300 K light leak, the outer surface of the linear bearing and the inner surface of the **50 K shield** are blackened with

carbon-loaded Stycast [156] to limit the propagation of any stray light. The innermost end of the bearing shaft bolts to the gripper finger, which is a 6 mm-deep, 90°-angled, oxygen-free high-thermal-conductivity (OFHC) copper wedge, attached to the CHWP stator via two flexible OFHC copper-braid heat straps.¹⁸ When the rotor is gripped, the finger fits into an identically shaped, azimuthally symmetric groove on the rotor stage, heat sinking the rotor while constraining its axial position.

The gripper finger also contains a spring-loaded probe that contacts an azimuthally symmetric **flex circuit** when the rotor is gripped, permitting four-point readout of a silicon diode thermometer¹⁹ on the sapphire stack's **cradle**. The probe is a beryllium copper, four-point battery contact, and the Kapton-based flex circuit²⁰ has four 2 mm-wide, 2.5 mm-pitch, gold-plated copper traces soldered to the thermometer's leads.

The complete gripper assembly is composed of three subassemblies that actuate simultaneously to grip and release the rotor. The motors are driven by a parallel-output controller,²¹ which is commanded using a programmable logic controller (PLC).²² The typical cooldown configuration for the PB-2b receiver is approximately horizontal (see Figure 9.2), with the bottom gripper finger supporting the rotor's 17 kg mass. Therefore, we employ a 450 N-max motor for the bottom subassembly and 200 N-max motors for the top two.²³

The gripper is designed for reliability. Each gripper finger attaches to its bearing shaft with two titanium²⁴ bolts to avoid fastener fatigue and has two PTFE²⁵ stabilizers that slide along the bottom face of the retaining ring (see Figure 8.4a) to constrain rotation about the gripper-port axis. The G10 tubes and hinges are joined with a step to avoid compression failures, and the hinges are screened to withstand more than 4× the rotor's weight. In addition, the rotor's center of mass lies in the plane of its triangular groove, and the bottom finger's copper wedge is slightly V-shaped, forming a 176° angle about the radial direction. These two features allow the rotor to be constrained by the bottom subassembly alone when the receiver is horizontal, providing insurance against possible gripper failure modes such as motor control issues and imperfect subassembly synchronization.

8.4 Motor

The **CHWP motor** needs to drive stable 2 Hz rotation while being low-dissipation, low-magnetic-interference, and mechanically robust at cryogenic temperatures. While EBEX successfully operated a belt drive using an external motor during weeks of CMB observation

¹⁸The OFHC heat straps were machined and welded at the LBNL main machine shop.

¹⁹Lakeshore DT-670: <https://www.lakeshore.com>

²⁰Q-Flex: <https://www.qflexinc.com/>

²¹SMC JXC-831: <https://www.smc Pneumatics.com/JXC831.html>

²²Click C0-00DR-D: <https://automationdirect.com/click-plc/>

²³SMC LEY-32 and LEY-16: <https://www.smcusa.com/>

²⁴Titanium is both strong and non-magnetic.

²⁵PTFE is low-friction and maintains its slipperiness down to cryogenic temperatures.




Figure 8.6: A CAD cross-section of the CHWP’s motor and a photograph of the encoder read heads. Figure 8.6a shows a zoomed CAD cross-section of the CHWP motor. Rotating components are labeled with cyan boxes, and stationary components are labeled with orange boxes. The slotted encoder plate and sprocket ring are semi-transparent for visual clarity. Figure 8.6b shows a photograph of the encoder read heads and nearby solenoids with the slotted encoder plate removed. The collimator holes for the motor LEDs are circled in green and labeled by phase, while those of the angle LEDs are circled in magenta. The solenoids are also labeled by phase and are radially aligned with their corresponding motor LED-PD pairs.

[105, 206], the PB-2b CHWP must run for years, motivating a drive system free of **mechanical fatigue**. Therefore, we utilize a brushless, three-phase, **synchronous electromagnetic motor** driven by custom electronics. We discuss the motor’s cryogenic assembly, driver, and efficiency in the following subsections.

8.4.1 Motor cryogenic assembly

The motor’s cryogenic assembly is shown in Figure 8.6a. Its active component consists of 114 **solenoids**²⁶ with low-carbon-steel magnetic cores²⁷ glued with equal spacing onto a 650 mm-diameter low-carbon-steel ring on the stator. Each coil has an 8.6 mH inductance and a $\approx 3 \Omega$ resistance at 50 K, and the motor’s three phases are driven across three regularly interspersed groups of 38 coils. To reduce the motor’s **equivalent resistance** and hence reduce the voltage needed to operate it, the solenoid array is further divided into four sections—two with 30 total coils, two with 27 total coils—that are driven in parallel. The motor’s passive component consists of 76 1.5 mm-diameter \times 1.5 mm-tall NdFeB **magnet sprockets**, which have a 6,700 G surface field and are glued with equal spacing and alternating polarity onto a 650 mm-diameter low-carbon-steel sprocket ring on the rotor.

The solenoids are energized in three 120°-separated phases with an alternating drive voltage $\pm V_D$, creating an **oscillating magnetic field** that couples to the rotor’s magnet sprockets, driving rotation. The sprockets, solenoids, and cores are carefully chosen to avoid

²⁶ APW Company: <https://apwelectromagnets.com/fc-6035.html>

²⁷ Low-carbon steel has a similar magnetic permeability to that of pure iron but is cheaper and easier to machine. All steel parts are nickel-plated to prevent rusting.

cogging while providing enough torque to attain the required rotational velocity. At 2 Hz rotation, the motor delivers only ≈ 5 N-mm of torque, making it susceptible to **physical touches**. Therefore, the CHWP assembly includes a system of wire harnesses to facilitate clean cable management. Additionally, despite the < 2 mm rotor-stator alignment requirement presented in Section 8.4.3, we provide 5 mm of clearance around the rotor to further limit the possibility of a physical touch impacting rotation.

The coils are driven by custom electronics, described in Section 8.4.2, whose sensing component is an **optical encoder** on the 50 K stage shown in Figure 8.6b. The encoder consists of three 940 nm light-emitting diodes²⁸ (LEDs) shining onto three reverse-biased photodiodes²⁹ (PDs) through a **slotted encoder plate** on the rotor. The LEDs and PDs are soldered to printed wiring boards and are housed in aluminum **read heads** with 1 mm diameter, 3 mm deep collimation holes. As shown in Figures 8.4a and 8.6b, the LED read head is mounted to four precision-ground aluminum standoffs on the stator baseplate, while the PD read head is similarly mounted to the retaining ring on the opposite side of the slotted encoder plate. Each motor LED-PD pair is azimuthally aligned with the solenoid array, and the read heads are aligned to each other using dowel pins during assembly. The gallium aluminum arsenide LEDs and the silicon PDs are cryogenically screened via dunk tests in liquid nitrogen and have been robust throughout hundreds of hours of testing. Even so, the CHWP has two identical pairs of read heads, as shown on the right- and left-hand sides of Figure 8.4, to provide insurance against LED or PD failure. The slotted encoder plate has 38 5 mm-wide motor slots whose edges are azimuthally aligned to the magnet sprockets. During rotation, the slot pattern chops the PD input, and the motor driver converts this photocurrent waveform into the solenoid bias voltage, as described in Section 8.4.2. Both the slotted encoder plate and cap ring are made of G10 to minimize motor-induced eddy currents on both the stator and rotor.

8.4.2 Motor driver

The motor drive electronics must be robust, low-noise, and simple to operate. While commercial drivers for synchronous motors are abundant, most involve auxiliary control software, pulse-width-modulated (PWM) waveforms (which can inject high-frequency noise into the receiver), and awkward interfacing to the PB-2b cryogenic assembly. Therefore, we employ a custom **driver printed circuit board (PCB)** that both meets our requirements and simplifies CHWP operation. The presented PCB is also used to read out the angle encoder signal, as discussed in Section 8.5.

A single PD amplification chain is shown in Figure 8.7, and the complete driver PCB is shown in Figure 8.8. Photocurrent from the PD read head travels through 50 K Manganin ribbon cables,³⁰ a DB-25 vacuum feedthrough,³¹ and 300 K, double-shielded, twisted-pair,

²⁸Vishay VSMB294008G: <https://www.vishay.com/docs/84228/vsmb294008rg.pdf>

²⁹Vishay TEMD1020: <https://www.vishay.com/docs/81564/temd1000.pdf>

³⁰Tekdata Interconnect: <https://www.tekdata-interconnect.com/>

³¹Accu-Glass Products: <https://www.accuglassproducts.com/25d2-450>

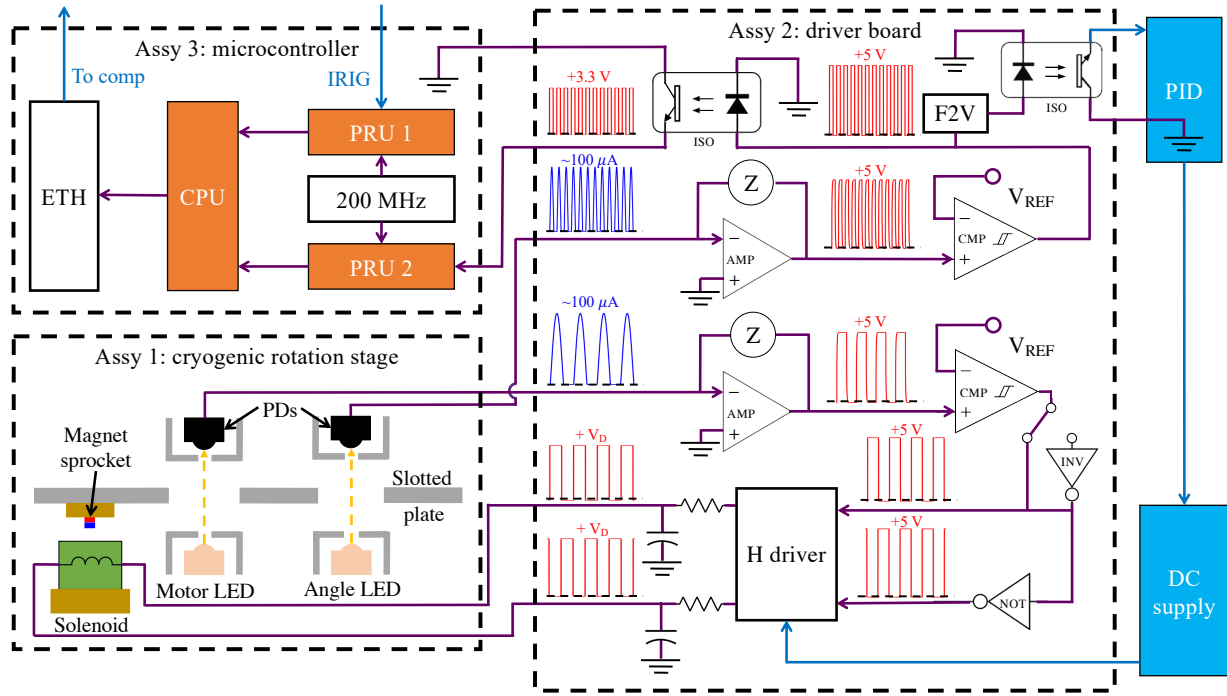


Figure 8.7: A signal diagram for a single phase of the CHWP motor driver and for a single angle encoder output. On the **cryogenic rotation stage**, the slotted encoder plate, which has one lane of slots for motor encoding and another for angle encoding, chops the signals of the LED-PD pairs. The **driver board** converts these photocurrents into voltage signals using transimpedance amplifiers (AMP) with carefully tuned low-pass feedback (Z) and then converts these analog waveforms into 5 V digital signals using comparators (CMP) with potentiometer-tunable reference voltages (V_{REF}). For motor control, both non-inverted and inverted (NOT) signals are passed to an H-bridge driver (H driver), which outputs synchronized, low-pass-filtered, alternating-polarity voltage waveforms $\pm V_D$ to the solenoids. For angle encoder readout, which has a $15\times$ faster signal than that of the motor, the digital signals are opto-isolated (ISO) and sent to a microcontroller unit for processing. Additionally, the angle encoder digital signal is passed to a frequency-to-voltage converter (F2V) whose opto-isolated output is used by a PID controller to provide feedback to the H-driver voltage and stabilize CHWP rotation. To brake or reverse direction, a switched inverter (INV) applies a global 180° phase shift to the H-driver input. The **microcontroller** is a BeagleBone Black, which has two programmable real-time units (PRUs) that share a 200 MHz clock. PRU 1 timestamps and decodes a GPS-synchronized IRIG-B PWM waveform, while PRU 2 timestamps rising and falling edges of the digital angle encoder waveform. These data are written to a shared buffer, which is emptied by the central processing unit (CPU) before being sent to an external computer over Ethernet (ETH).

copper cables³² to the driver board where it is converted to an analog voltage by a transimpedance amplifier. The time constant of the amplifier feedback is chosen to suppress high-frequency noise while outputting a symmetrical waveform. This analog signal is then converted to 5 V digital, and the TTL waveforms from each PD read head are OR-ed so that if one LED-PD pair fails, CHWP operation is unaffected. All 12 solenoid chains (three phases

³²Alpha Wire 6831: <http://www.alphawire.com/>




Figure 8.8: A photograph of the driver PCB, the schematic for which a subset is shown in Figure 8.7. The photocurrents from all ten encoder PDs enter the board through the “PD input” connector before being converted to analog signals by the transimpedance amplifiers “AMPs” and converted to 5 V digital signals by the comparators “CMPs.” At this point, the angle encoder signals are opto-isolated (“ISOs”) and sent to the microcontroller via the “MCU output” connector. The motor signals from both PD read heads are “OR-ed” before being inverted (“INV”) and input to the “H drivers,” which in turn power the solenoids via the “Solenoid output” connector. The motor and encoder LEDs are current biased by potentiometer-adjustable “LED drivers” via the “LED output” connector. DC supply voltages come through the “PWR input” connector.

in four sections) are energized in parallel by **H-bridge drivers** whose output is cleaned by a single-pole ≈ 300 Hz low-pass filter to suppress any high-frequency interference in the cryostat. The PCB’s digital logic and H drivers are powered by low-noise, non-switching DC power supplies,³³ and the PCB layout and ground-plane geometry are specifically designed to avoid contaminating analog signals with digital artifacts.

The CHWP’s rotational velocity is naturally steady due to the rotor’s large rotational inertia and the motor’s small torque. However, the positioning between the motor’s solenoids and rotor’s magnet sprockets changes slightly when the CHWP’s gravitational orientation changes, such as during telescope motion. This modulation in motor coupling slightly modulates the motor’s efficiency (see Section 8.4.3) and causes small, slow drifts in rotational velocity. Therefore, we employ **proportional-integral-derivative (PID) feedback**³⁴ to the H-driver voltage in order to stabilize the CHWP velocity on long timescales.

Finally, in order to both stop the rotor and spin it in the opposite direction, an inverter switch applies a 180° phase shift to both H-driver inputs when toggled by an external digital input. **Braking** is vital during telescope power failure (see Section 9.3.5), and spinning the

³³Kikusui PMX: <https://www.kikusui.co.jp/en/>

³⁴Omega CNi16D52: <https://www.omega.com/en-us/>

CHWP in both directions provides a useful data split during analysis.

8.4.3 Motor efficiency

The motor’s maximum torque is delivered at start-up when the rotor’s sprocket pattern and the solenoid array’s rotating magnetic field are in phase. However, at non-zero rotation frequencies, the solenoids’ inductance creates a phase shift between the H-driver voltage and the solenoid current, degrading motor efficiency. The impact of this **inductive phase shift** is shown in Figure 8.9 and is $\approx 20\%$ at 2 Hz rotation.³⁵

Additionally, motor efficiency relies on concentricity between the rotor and stator. When radially misaligned, the coupling between the solenoids and the magnet sprockets weakens, reducing motor torque. Radial misalignment³⁶ is most likely to occur along the gravitational axis when the cold assembly contracts during cooldown. The simulated impact of rotor radial misalignment on motor efficiency is shown in Figure 8.9 and assumes a 5 mm rotor-stator axial separation. This simulation motivates a concentricity requirement of $\Delta R < 2$ mm, for which the efficiency loss is $< 20\%$.

8.5 Angle encoder

The rotor angle is measured by an **incremental encoder** that uses much of the same infrastructure as the motor encoder (see Section 8.4). The angle encoder’s signal schematic is shown as part of Figure 8.7, and its cryogenic components are shown in Figure 8.6. On each encoder read head, two LED-PD pairs peer through $570 - 1 = 569$ slots on the slotted encoder plate with one missing slot to mark the rotor’s **absolute position**. During 2 Hz continuous rotation, the PD photocurrents are chopped at 1.14 kHz, are digitized on the driver board, are opto-isolated to 3.3 V, and are processed by a **microcontroller unit (MCU)**. The two angle encoder LED-PD pairs are offset by half a slot width, enabling **quadrature readout** to monitor rotation direction.

The MCU is a **BeagleBone Black (BBB)**,³⁷ which houses two **programmable real-time units (PRUs)** and an on-board CPU running Linux. The PRU is a lightweight, low-latency processing unit specifically designed to handle single-threaded inputs, making it ideal for angle encoding. One PRU polls the angle encoder signal and uses a 200 MHz clock³⁸ to timestamp each rising and falling edge. Simultaneously, a second PRU polls and decodes a GPS-synchronous **inter-range instrumentation group B code (IRIG-B)**

³⁵In principle, a microcontroller or equivalent could use the F2V output to time-shift H-driver inputs and recover this loss. However, as shown in later sections, the existing motor scheme meets the CHWP requirements, and therefore the phase delay does not need to be corrected for PB-2b.

³⁶Axial displacement is also important for motor coupling but is easily controlled to ~ 0.1 mm by the gripper’s wedge-and-groove design.

³⁷Beagle Board: <https://beagleboard.org/black>

³⁸Specifically, both PRUs access the BBB’s industrial Ethernet protocol (IEP) timer.




Figure 8.9: CHWP motor efficiency vs. the rotor’s radial misalignment to the motor solenoids and the rotor’s rotation frequency. The PB-2b rotor alignment tolerance is < 2 mm, limiting positional efficiency loss to $< 20\%$. At 2 Hz rotation, the phase-delay efficiency loss is $\approx 20\%$.

PWM waveform³⁹ using the same 200 MHz clock. The encoder and IRIG clock values are written to a shared memory buffer that is periodically emptied by the CPU, which in turn sends IPv4 data packets to an external computer. During post-processing, the rotor angle is reconstructed using the missing reference slot and is interpolated to IRIG time using the MCU clock values. This time-ordered CHWP angle data can then be used to **demodulate** the detectors, which are also synchronized to IRIG.

While the CHWP angle jitter requirement is $\ll 3 \mu\text{rad}/\sqrt{\text{Hz}}$, the encoder only has a resolution of 5 mrad, necessitating **precise interpolation** between ticks on the slotted encoder plate. Such an encoding scheme is feasible because the CHWP’s rotation is very smooth and the encoder is high-signal-to-noise, enabling clean angle reconstruction (see Sections 9.3.3 and 9.4.1).

8.6 Thermal design

The effectiveness of the CHWP system depends centrally on its thermal performance. As shown in Figure 8.1b, the CHWP is located on the 50 K stage, which is cooled by the first stage of the optics tube’s PTR. The rotor is shielded from sky-side radiation by the vacuum window, RT-MLI, and IRF, and it floats between the IRF and the field lens. As discussed

³⁹Spectrum Instruments TM-4: <http://www.spectruminstruments.net/>

in Section 8.2, the CHWP thermal design’s primary objective is to minimally impact both the 4 K and 50 K stage temperatures with respect to the heritage PB-2b configuration.

When the rotor is stationary, loading on the 4 K stage is lower than in the heritage system. The IRF transmits non-negligibly $\lesssim 2$ THz [93], and when present, the CHWP’s sapphire stack absorbs $> 90\%$ of this leaked sky-side power, keeping it from reaching the field lens. Additionally, because the rotor is floating, it acts as **multi-layer insulation (MLI)**, further reducing the 4 K load. When the rotor is spinning, however, the solenoids, LEDs, and rotor generate heat, and if these loads warm the 50 K stage or rotor too much, 4 K improvements could be negated. Therefore, the CHWP’s thermal design focuses on limiting 50 K dissipation and on keeping the rotor cool.

We use an analytic simulation to predict rotor temperature and CHWP-induced power on the 50 K and 4 K stages during continuous operation in the field. Figure 8.10 shows a schematic for the model, and Table 8.2 shows the measured and calculated values for the schematic. The errors on the measured values contain both measurement and configurational uncertainties, while the errors on the calculated values are driven by uncertainties in the assumptions. We find that the most important contributions to the CHWP’s thermal performance (Figures 8.11 and 9.10) are the IRF temperature, 50 K stage temperature, and the IR emissivity of the IRF and sapphire stack.

The CHWP is an 11.4 mm-thick stack of three sapphire windows. At 50 K, the sapphire stack has an emissivity of 0.95 at ~ 3 THz [147], and therefore the CHWP rotor absorbs IR radiation efficiently, even without an AR coating. The dominant sources of radiative transfer to and from the rotor are the IRF, stator, and field lens. The IRF is heat-strapped to its fixture using 24 1 mm-thick, 50 mm-wide, 20 mm-long flexible ribbons of 99.9999% (6N) pure aluminum, which have superior conductivity between 100 \sim 50 K [230], are malleable enough to engage the alumina surface without thermal grease, and are flexible enough to absorb differential thermal contraction between the IRF and its aluminum fixture. The IRF fixture is machined from aluminum 1100 and is bolted to six 90 mm-tall, titanium-helicoiled OFHC copper towers (see Figure 8.4), which provide a thermal path to the stator baseplate. The measured conductivity between the IRF and the stator baseplate is $k_2 = 2.5$ W/K. During laboratory testing with a thick window and no AR coatings (see Section 9.3.4), we see an IRF temperature that is ≈ 5 K warmer than the 50 K stage, suggesting a ≈ 13 W sky-side load, which is slightly less than the model’s median expectation but within its 1σ uncertainty.

The sky-side power leaked onto the rotor depends on the IRF’s AR coating. In this analysis, we assume that the IRF and CHWP sapphire stack are coated with an **epoxy-based AR coating** [163], which has a non-negligible transmissivity $\lesssim 2$ THz [69]. Using the effective temperature of a six-layer RT-MLI stack [31] and the transmission spectra for Stycast-coated alumina [93], we estimate that $\approx 60 \pm 10$ mW of sky-side power leaks through the IRF onto the rotor (or, in the heritage PB-2b configuration without the CHWP, onto the field lens).

The stator baseplate is made from aluminum 1100, houses the motor and encoder, and includes a 3 mm-thick 50 K shield (see Figure 8.5) that encloses the CHWP assembly. The




Figure 8.10: The CHWP's thermal circuit during nominal operation. The squiggly lines represent radiative loads, the straight lines conductive loads, and the dotted lines dissipative loads. Loads labeled by p represent constant power, while loads labeled with k represent power that depends on operating temperatures. Each conductivity has a red-to-blue color gradient from hot to cold. The switched conductivity k_6 represents the gripper connection, which is closed when the rotor is gripped and open when it is not. The lines' thicknesses show the relative magnitudes of the various contributions but are not to scale and are intended only as a visual guide. The IRF (T_{IRF}), rotor (T_{R}), stator (T_{S}), and field lens (T_{FL}) temperatures are the system's merit figures. The measured and calculated values for each symbol are given in Table 8.2.

50 K shield's inner wall is coated with carbon-loaded Stycast [156] and is intended to absorb any 300 K photons that leak into the CHWP cavity through the gripper ports. The 50 K shield is attached to the stator baseplate by only sixteen bolts, and therefore its temperature is typically ≈ 5 K higher than that of the baseplate. The radiative coupling between the stator and the rotor is similar to that between the IRF and the rotor.

The encoder consists of two read heads, each with a set of five LEDs shining through the slotted encoder plate onto five photodiodes. At 50 K, the LEDs are current biased with 30-50 mA at 1.8 V, and therefore the power dissipated by each read head is 300-500 mW. We attach each read head to the stator baseplate using four 6 mm-diameter aluminum 6061 standoffs with polished ends and interfaced with thermal grease,⁴⁰ and the estimated read-head warm-up is < 1 K. The LEDs also shine onto the rotor's slotted encoder plate, whose

⁴⁰Apiezon N Grease: <https://www.apiezon.com/>

Symbol	Description	Value
p_1	Radiative load from 300 K onto the IRF, including that of the sky, window, and RT-MLI	17 ± 5 W
p_2	Radiative load from the encoder LEDs onto the rotor	10 ± 3 mW
p_3	Power dissipated by the encoder LEDs onto the stator	0.8 ± 0.2 W
p_4	Eddy current and hysteresis dissipation onto the rotor due to the motor's oscillating magnetic field	< 1 mW
p_5	Rotor frictional dissipation during 2 Hz rotation	80 ± 10 mW
k_1	Radiative coupling between the IRF and the rotor	7 ± 2 mW/K
k_2	Conductivity between the IRF and the stator	2.5 ± 0.3 W/K
k_3	Radiative coupling between the stator and the rotor	5 ± 2 mW/K
k_4	Radiative coupling between the solenoid array and the rotor	< 1 mW/K
k_5	Conductivity between the motor solenoids and the stator	0.8 ± 0.1 W/K
k_6	Conductivity between the rotor and the stator via the gripper	0.7 ± 0.1 W/K
k_7	Radiative coupling between rotor and the field lens	5 ± 1 mW/K

Table 8.2: Parameter definitions and values for Figure 8.10. The bold values are measured, while the non-bold values are estimates. The error bars represent some combination of uncertainties in the calculation or measurement as well as configurational variations.

slit patterns have 50% duty cycles. The LED emits between 900 and 960 nm, and its beam has a peak intensity of 35 mW/sr at 50 mA bias with a full-width-half-max of $\pm 7^\circ$. In order to limit stray read-head emission, each LED is collimated by a 1 mm wide, 3 mm deep hole that truncates its beam at $\pm 10^\circ$. Upon integrating these LED beams over the collimation holes, the estimated loading on the rotor from both encoder read heads is 10 ± 3 mW.

During 2 Hz rotation, the motor solenoids carry an RMS current of ≈ 20 mA, generating a peak-to-peak magnetic field of ≈ 20 G. The field is generated by a $38f_{\text{HWP}}$, $V_D \approx 12$ V square-wave (at $f_{\text{HWP}} \approx 2$ Hz), which is low-pass filtered above ≈ 300 Hz. To keep the coils cool, we epoxy⁴¹ them to their cores, providing a 7 ± 2 mW/K conductive path to the stator. At 50 K, each coil's resistance is $\approx 3 \Omega$, therefore dissipating ≈ 1 mW per coil during continuous rotation. This dissipation warms the coils < 1 K, and their radiative coupling to the rotor is < 1 mW/K. In order to minimize eddy current dissipation on the rotor due to the motor's oscillating magnetic field, the encoder plate is made of G10, limiting eddy losses on the rotor to the sprocket ring. Electromagnetic dissipation within low-carbon steel at ~ 10 G and ~ 100 Hz frequencies induces < 1 mW of heating on the rotor.

⁴¹Stycast 2850FT: <https://www.henkel-adhesives.com/us/en/>




Figure 8.11: Rotor temperature (Figure 8.11a) and CHWP-induced power on the field lens (Figure 8.11b) as a function of excess, unmodeled power on the rotor. The solid blue line represents the median, while the shaded regions denote one-sigma uncertainties. Each plot’s y -intercept gives the modeled expectation, and the IRF is assumed to be 55 K.

The field lens is ≈ 50 mm thick and is therefore assumed to have an IR emissivity of 1 for all calculations. Its coupling to the rotor is determined purely by the CHWP aperture diameter and the field-lens view factor and is calculated to be 5 ± 1 mW/K.

Assuming the thermal model in Figure 8.10 and Table 8.2, Figure 8.11 shows the rotor temperature and the CHWP-induced load on the field lens as a function of excess power on the rotor, where “excess” is that which is beyond the model’s prediction. The expected rotor temperature is 52 to 54 K, and the expected CHWP-induced field lens load is -70 to -10 mW with respect to the heritage configuration. CHWP-induced power on the field lens is < 0 mW at 85% confidence when the rotor is < 55 K, which motivates the rotor temperature requirement presented in Section 8.2.2. CHWP-induced power on the 50 K stage is < 1.3 W at 85% confidence and is nearly independent of rotor temperature.

8.6.1 Rotor thermometry

Two modes are used to monitor the rotor’s temperature. First, when the rotor is gripped, four spring-loaded contacts on each gripper finger touch four flex-circuit traces on the rotor stage (see Section 8.3.4). These traces are soldered to a four-wire silicon diode thermometer, which is varnished to the sapphire stack’s cradle. While only one contact current biases the diode, all three probe the diode’s voltage; therefore this flex-circuit thermometry is also a



Figure 8.12: A photo of the CHWP being installed into the PB-2b optics tube. The cryostat is constructed vertically on a cart from the focal plane towards the sky. The CHWP is one of the last elements to be installed and is hoisted using a series of handles that attach to the rotor stage. Once above the receiver, the CHWP is slowly lowered into the optics tube and is bolted to the 50 K stage.

touch-sensing system used to verify the fit between the gripper’s copper wedges and the rotor stage’s triangular groove.

Second, during continuous rotation, we use the ring magnet as a thermometer. Neodymium has a **temperature-dependent magnetic field**⁴² that generates a ~ 1 G/K variation ≈ 5 mm from the ring magnet’s face [166]. We varnish a 1 mm thick cryogenic Hall sensor⁴³ to the surface of the YBCO and monitor long-time-scale changes in the ring magnet’s average field. Limited by the degeneracy with rotor position, the sensitivity of this scheme is a few K, and therefore Hall-sensor monitoring is intended to flag thermal events rather than subtract rotor temperature drifts from detector data. If the need to re-calibrate the Hall-sensor arises, the CHWP can be stopped and gripped to cross-check the diode thermometer via the touch probes.

8.7 Integration

The CHWP is designed to be modular and assembles almost entirely on the benchtop (see Figure 8.4), allowing for seamless integration into a nearly-fully-assembled PB-2b receiver. To secure the rotor to the stator and keep the two concentric during installation, the rotor stage is bolted to three of the copper towers in Figure 8.4b using aluminum **installation**

⁴²Neodymium magnetization: http://spontaneousmaterials.com/Papers/TN_0302.pdf

⁴³Lakeshore HGCT-3020: <https://shop.lakeshore.com/>

stanchions. Then, four handles are fastened to the top-most face of the rotor stage, which allows the assembly to be hoisted and maneuvered.

Figure 8.12 shows a photo of the CHWP installation. First, the receiver is tilted with its boresight at zenith. Then, straps are attached to the four handles on the rotor stage, and a hoist lifts the CHWP above the optics tube. At this point, the CHWP is lowered and bolted onto the 50 K stage before the receiver is returned to a horizontal position. Finally, the gripper subassemblies are installed, the retaining ring is added, and the IRF, RT-MLI, and vacuum window close the cryostat.

Because the CHWP integrates into a nearly fully assembled receiver, the two systems were able to be evaluated in parallel prior to full system validation. This facet of the CHWP design was inspired in part by the PB-2a WHWP’s modularity, and it dramatically accelerated the PB-2b commissioning process.⁴⁴ The sapphire stack is also modular, enabling sapphire AR coatings to be easily replaced in the field, if necessary.

⁴⁴As a counterexample, if the CHWP were at the 4 K stage, then modulator development would have been tied to the optics tube, and each subsystem’s cryogenic issues would have been difficult to disentangle.

Chapter 9

PB-2b CHWP evaluation

This chapter overlaps substantially with Hill, Kusaka, Ashton, Barton, et al. (2020) [79], which we encourage the reader to cite instead of or in addition to this dissertation.

The PB-2b CHWP development began as a list of requirements and a concept that vaguely resembled the EBEX modulator. From those earliest stages in 2014 to its Chile deployment in 2020, the CHWP system underwent many iterations of designing, building, testing, and revising. This chapter illuminates the endpoints of that maturation, briefly discussing the CHWP prototype before presenting the laboratory validation of the deployed system. We show that the completed CHWP meets all of the requirements in Table 8.1 and is ready to see first light on PB-2b. This chapter is a continuation of Chapter 8.

9.1 Prototype evaluation

To learn about the superconducting magnetic bearing (SMB), synchronous magnetic motor, encoder system, and gripper, we began the CHWP’s development with ≈ 1.5 years of ~ 120 mm-diameter prototyping. The **CHWP prototype** was designed and built at Lawrence Berkeley National Laboratory (LBNL), and its testing informed nearly every aspect of the PB-2b CHWP presented in Chapter 8. The prototype was cooled eight times in the CAPMAP dewar [19] at LBNL using a Gifford-McMahon (GM) cooler,¹ and each testing cycle provided rapid feedback to each subsystem. In this section, we briefly discuss the most prominent differences between the small-scale and large-scale CHWP assemblies, highlighting lessons learned during the prototyping process.

The biggest difference between the prototype CHWP and the PB-2b CHWP is the gripper. The prototype’s entire gripper mechanism, including its linear actuators and bearings, is within the vacuum space, as shown in Figure 9.1. The gripper assembly is composed of two subassemblies on the $\pm x$ sides of the rotation stage that actuate two **gripper panels** on the $\pm y$ sides of the stage along the $\pm y$ -direction. Each gripper assembly has two 300 K

¹Cryodyne 1020: <https://www.cv.nrao.edu/~rfisher/Cryogenics/CTIdatasheet.pdf>

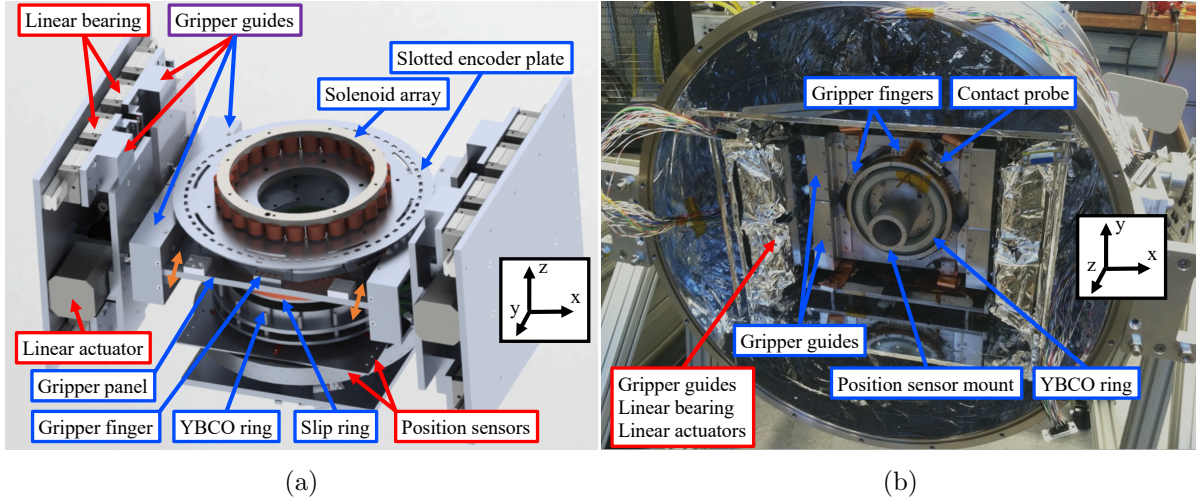


Figure 9.1: The CHWP prototype developed at LBNL prior to constructing the PB-2b system in Chapter 8. The computer-aided-design (CAD) render (Figure 9.1a) shows the fully assembled prototype, and the photo (Figure 9.1b) shows the partially assembled prototype in the CAPMAP dewar. Warm components are marked with red boxes and cold components with blue boxes. The orange arrows in Figure 9.1a denote the $+y$ gripper panel's motion, which is facilitated by vacuum-compatible actuators and linear bearings on the $\pm x$ sides of the rotation stage. The warm gripper hardware in Figure 9.1b is wrapped in aluminized Mylar.

linear bearings, two 300 K linear actuators, and two gripper guides that move along these linear bearings and provide thermal isolation between warm components and the cryogenic gripper panels. Two motors, one on each side of the rotation stage, actuate each gripper panel, and each panel contains three gripper fingers that fit into an azimuthally symmetric triangular groove on the rotor. Similar to the PB-2b CHWP (see Section 8.3.4), a slip ring is used to read out a thermometer on the rotor.

There were several issues with this prototype design that led to an overhauled PB-2b gripper system. First, six gripper fingers on two moving panels **overconstrain** the rotor, making it prone to misalignment. Second, the gripper is itself overconstrained, requiring two motors to actuate each panel. If the motors are not perfectly synchronized, the panel becomes crooked, **jammed**, and difficult to recover. Third, the gripper relies on vacuum-compatible bearings and motors, whose oil-free parts are susceptible to stiction. If a problem with any of these mechanical interfaces arises, the CHWP becomes inoperable until the dewar is opened and the assembly is pulled apart. Fourth, the rectangular assembly was bulky and difficult to interface with CAPMAP's tubular vacuum shell. Furthermore, a scaled-up prototype design would have required a redesign of PB-2b's optics tube, partially negating the CHWP's modularity benefit. For these reasons and others, we abandoned the prototype gripper system and instead adopted the design presented in Section 8.3.4.

The second prominent difference between the prototype CHWP and the PB-2b CHWP is the prototype's system of **position sensors**. Because the rotor is free to move within the cryostat during cooldown, we wished to monitor its position and ensure rotor-stator

alignment. The prototype therefore includes six LEDs shining onto six **one-dimensional photosensors**² through thin slots on the rotor, and the outputs of these sensors are used to reconstruct the rotor’s position. Both the sensors and LEDs are mounted to 300 K surfaces and peer through “windows” in a (complexly shaped) 50 K shell that surrounds the rotation assembly. Despite this linear sensor setup’s demonstration in warm benchtop tests, its utility inside the dewar was limited. Therefore, as we gained confidence in the SMB’s stiffness and the improved gripper design’s robustness (Section 8.3.4), we removed the position sensors to reduce **system complexity**.³

Despite the prototype’s underperforming gripper and position sensing systems, its bearing, motor, encoder, and thermometry designs worked well and were therefore carried to the full-scale design. In this way, the prototyping process provided both affirming and halting lessons, and we emphasize that the deployed design resulted from much ideation, experimentation, and iteration. Provided this historical context, we fast forward to the other endpoint on the CHWP’s development timeline, presenting the modulator’s laboratory evaluation prior to its deployment.

9.2 PB-2b CHWP evaluation

The full-scale PB-2b CHWP is evaluated in two laboratory testing phases. The first phase is conducted in a **non-optical configuration** of the CAPMAP dewar at LBNL and uses an IR-blackened aluminum disk instead of the CHWP’s sapphire stack. The second phase integrates the CHWP into an **optical test configuration** of the PB-2b receiver at the University of California, San Diego (UCSD), shown in Figure 9.2. The UCSD setup employs the full detector array, but the lenses, CHWP, IRF, and vacuum window are not AR coated, and the sapphire stack only includes one 3.8 mm-thick sapphire plate. Though optically incomplete, the UCSD configuration allows us to evaluate the cryo-mechanical, thermal, and magnetic interaction between the CHWP and the rest of the experiment.

We divide a discussion of the CHWP evaluation into two sections. First, we discuss operational testing, including spin-up, continuous rotation, thermal performance, and shutdown. Second, we present the CHWP’s data quality and noise impact, including encoder jitter, magnetic interference, and rotor temperature stability.

9.3 Operational performance

This section reviews CHWP operation, focusing on gripper, motor, bearing, and thermal performance. We evaluate interfacing between various subsystems and show that the CHWP executes its essential functions while meeting the requirements shown in Table 8.1.

²Three one-dimensional sensors monitor the rotor’s radial position, and three other sensors monitor its z -position.

³Nonetheless, a rotor position measurement is appealing, and Simons Observatory is developing a capacitive position sensing scheme for its CHWPs.




Figure 9.2: The CHWP mounted in the PB-2b receiver at UCSD. Before the gripper engages, the rotor is held by three installation stanchions that align it to the stator. The top two “small” gripper subassemblies are only separated by 80° to avoid colliding with the PTR attachment. The single sapphire plate has no AR coating, and its frosty appearance is due to its $\sim 0.1 \mu\text{m}$ RMS surface roughness.

9.3.1 Cooldown

When the YBCO is above its $\approx 90 \text{ K}$ transition temperature, the gripper must keep the rotor centered, and between 300 and 50 K, the rotor contracts $\approx 2 \text{ mm}$ radially with respect to the vacuum shell. Therefore during cooldown, the gripper periodically “re-grips” the rotor, with each gripper finger inching inwards until the grip force is 50% of the rotor’s weight. This routine both keeps the rotor centered until the YBCO goes superconducting and maintains a conductive cooling path from the rotor to the stator baseplate. Figure 9.3b shows gripper-finger position vs. rotor and stator temperature during an LBNL cooldown. In this particular test, the fingers are moved manually at convenient intervals, and their relative positions are maintained to $< 0.2 \text{ mm}$, resulting in a rotor-stator coupling efficiency of $\approx 99\%$ (see Figure 8.9).

Also during cooldown, the rotor’s only cooling mechanisms are conduction through the gripper fingers and radiative coupling to the surrounding environment. Because the rotor’s heat capacity is $\sim 10 \text{ kJ/K}$ at 300 K, good rotor-gripper contact conductance is needed for the CHWP to thermalize within 36 hours of the 50 K stage. The rotor stage’s triangular groove and gripper fingers’ triangular wedges are precision-cut to maximize thermal contact, and the measured rotor-to-gripper conductance is 0.7 W/K at 100 K. Figure 9.3a shows rotor and stator temperatures during a cooldown in the LBNL cryostat. The rotor lags behind the stator by only five hours, which is well within the 36-hour requirement.




Figure 9.3: The results of CHWP cooldown testing at LBNL. Figure 9.3a shows CHWP temperatures during a thermal cycle to 25 K in LBNL’s dark cryostat. While the rotor lags behind the stator by ≈ 10 K, it thermalizes within five hours of its surroundings. Figure 9.3b shows radial displacement of each gripper finger vs. rotor and stator temperature. The gripper-finger labels refer to their locations in Figure 9.2. The motors are manually commanded at convenient points during the ≈ 100 -hour cooldown, and their relative positions are maintained to ≤ 0.2 mm throughout.

9.3.2 Start-up

After the YBCO becomes superconducting and the CHWP assembly thermalizes, the rotor is held in place for a few hours to let the bearing “relax.” Flux pinning needs time to find its lowest-energy configuration [159], and if the rotor moves during this relaxation, pinning sites are more likely to escape the superconducting bulk, hence reducing the bearing’s spring constant. This **flux migration** is logarithmic in time, with the vast majority of equilibration occurring within hours of the YBCO transition. After the bearing has relaxed, the gripper fingers retract, and the stator supports the rotor’s weight. When the receiver is horizontal, as in Figure 9.2, the rotor is observed to “sag” away from its gripped position by ≈ 0.5 mm, and because the bearing’s spring force is elastic, this displacement decreases with increasing receiver inclination.

When floating and stationary, forces between the magnet sprockets and the solenoids’ ferromagnetic cores keep the rotor azimuthally constrained. In order to overcome this **stiction** during start-up, the motor energizes with ≈ 0.4 A (or $\sim 5\times$ its current draw during continuous operation). Once rotation commences, **cogging** quickly diminishes, the rotor transitions into smooth rotation, and the solenoid bias is reduced. Rotation frequency vs.




Figure 9.4: CHWP rotation frequency vs. time during start-up for various H-driver voltages in the LBNL cryostat. The motor attains rotation frequencies up to ≈ 2.8 Hz, at which point it becomes limited by motor-phase-delay efficiency (see Figure 8.9) and rotor friction (see Figure 9.6).

time for various H-driver voltages during spin-up are shown in Figure 9.4. The achievable rotation frequencies are 1.8-2.8 Hz, and the equilibration time is tens of minutes. While PID control will reduce the start-up time somewhat, the CHWP is intended to run without interruption throughout each day’s telescope observations. Therefore, this level of latency meets PB-2b’s needs.

9.3.3 Continuous rotation

Once spun up, the CHWP enters **constant-velocity mode**. To measure the CHWP’s open-loop stability, we collect one hour of continuous rotation data in the LBNL test cryostat without PID control. Such a test helps identify any pathologies in the drive system that may manifest as modulations in rotational velocity. While the PB-2b CHWP does not have a velocity stability requirement, steady rotation helps assess encoding accuracy, as discussed in Section 9.4.1.

Figure 9.5 shows velocity drift vs. time at $f_{\text{HWP}} = 2.15$ Hz. The total $\Delta f_{\text{HWP}} = 0.8$ mHz corresponds to a rotational stability of 0.04%/hr, which is much better than that of ABS [111] and PB-1. While changes in telescope inclination will slightly modulate rotor-stator concentricity, which in turn modulates motor efficiency (see Section 8.4.3), we anticipate even better stability when using PID control in the field (see Section 8.4).




Figure 9.5: Δf_{HWP} during one hour of continuous rotation in the LBNL cryostat, sampled once per rotation and averaged over one-minute intervals. The mean velocity is $f_{\text{HWP}} = 2.15$ Hz, and therefore the fractional stability is 0.04%/hr.

9.3.4 Thermal impact

The CHWP's thermal impact is evaluated in three stages. First, we measure rotor friction and motor dissipation in the LBNL cryostat. **Rotor friction** is determined by measuring deceleration vs. velocity with the motor powered off, and **motor dissipation** is determined by measuring the current through and resistance of the solenoids. Dissipation as a function of rotation frequency is shown in Figure 9.6a, and at 2 Hz, rotor and motor heating are 80 mW and < 200 mW, respectively. A polynomial fit to rotor dissipation vs. f_{HWP} shows that at 2 Hz, most friction is due to eddy currents ($P_{\text{eddy}} \propto f_{\text{HWP}}^2$) as opposed to hysteresis loss ($P_{\text{hyst}} \propto f_{\text{HWP}}$). The motor dissipation data points represent measured values with excellent alignment at $T_{\text{rotor}} = 60$ K, and the bands quantify possible variations in motor torque due to variations in rotor-stator concentricity and solenoid temperature. The motor dissipation's steep slope is due to both decreasing motor efficiency and increasing rotor friction with increasing f_{HWP} .

Second, we evaluate **cryo-optical performance** in the UCSD test setup shown in Figure 9.2. Even though this configuration has no AR coatings on the IRF, sapphire stack, or lenses, bare alumina and sapphire are largely IR absorptive, allowing us to evaluate the thermal model presented in Section 8.6. During UCSD testing, the CHWP's rotor thermometer disconnected, leaving us to infer the rotor's temperature rather than monitor it directly. Using a load curve of the field lens [84] and comparing field-lens and 50 K temperatures to a dark run without the CHWP [84], we measure a rotor temperature of < 50 K and a field-lens temperature of 6.4 K, both of which are $\approx 1\sigma$ better than the thermal model's




Figure 9.6: Measured 50 K dissipation and mK heating during continuous rotation. Figure 9.6a shows f_{HWP} -dependent dissipation on the 50 K stage during continuous rotation. As described in Section 9.3.4, a polynomial fit shows that rotor friction at 2 Hz is predominantly due to eddy-current losses, while motor dissipation is due to resistive heating in the solenoid array. The motor dissipation’s error bars represent possible variations due to rotor positioning, solenoid temperature, and cryostat angle. Figure 9.6b shows mK-fridge and focal-plane temperatures as the CHWP is turned on and off. There is no measurable CHWP-induced heating to within a ≈ 0.3 mK/hr background drift.

expectation after accounting for no AR coatings.⁴ On the other hand, we measure a 50 K load of ≈ 2.5 W, which is larger than the 2 W requirement. This test was performed before adopting the **solenoid heat sinking** scheme described in Section 8.6, which has been shown to reduce solenoid heating by $\sim 5\times$ in auxiliary tests. Therefore, we anticipate a stator load that is consistent with the thermal model’s < 1.3 W expectation when the CHWP operates in Chile.

Third, we search for any CHWP-induced heating of the focal plane. We spin the rotor up and down at various drive voltages over several hours in the UCSD setup, and as shown in Figure 9.6b, we see no changes in mK temperatures to within a ≈ 0.3 mK/hr background drift,⁵ implying negligible CHWP-induced focal-plane vibrations. This null result is expected, as the CHWP’s resonant frequency $\sqrt{k/m} \approx 130$ Hz is much higher than its ≈ 2 Hz rotation frequency.

⁴With one piece of 3.8 mm-thick sapphire and with no AR coatings, the rotor temperature and field-lens power in the UCSD test are expected to be lower and higher, respectively, than shown in Figure 8.11.

⁵No focal plane temperature regulation was employed during this test, and therefore this background drift is larger than that which is achievable.




Figure 9.7: Spin-down tests from various rotation frequencies using various braking voltages. These tests did not employ the PID controller, which will shorten the stopping time for continuous rotation frequencies < 2.8 Hz.

9.3.5 Shutdown and recovery

A critical capability for long-term field operation is the **clean recovery** of the rotor following a power disruption. When the optics tube’s PTR shuts off, the CHWP warms with the rest of the experiment, and in the absence of a robust recovery procedure, the rotor may become difficult to recenter without opening the cryostat. Therefore, we employ two redundant procedures to recapture the CHWP in the event of a cryostat warmup.

The primary procedure is to stop and grip the rotor before the 50 K stage warms appreciably. The CHWP electronics are powered via an **uninterrupted power supply (UPS)** that provides a ~ 30 min window after site-wide power loss during which the CHWP must be stilled, re-gripped, and stowed. Necessitated by its low friction, the rotor is actively braked by globally inverting the motor’s H-driver outputs (see Section 8.4.2). Figure 9.7 shows rotation frequency vs. time for various braking voltages, and the measured spin-down time is ≈ 5 min, which is sufficiently shorter than the UPS duration. After rotation stops, the rotor is gripped loosely to provide a margin for thermal expansion, the gripper motors’ brakes are applied, and the CHWP electronics are shut down. When site power is later restored, the cooldown procedure in Section 9.3.2 commences, and nominal CHWP operation resumes. We have tested this emergency stop and **re-grip procedure** multiple times in the LBNL cryostat and have shown that it keeps the rotor centered to within 0.5 mm.

In the event of an issue when stopping and re-gripping, we employ a **backup procedure** to recover the rotor after it has “fallen.” As shown in Figure 8.4, the retaining ring includes six PTFE crash pads designed to gently “catch” the rotor and limit its misalignment after

the bearing disengages. Once site power is restored, we orient the receiver horizontally, as shown in Figure 9.2, lift the rotor along the $+y$ direction using the bottom gripper finger (see Section 8.3.4), and grab it using the top two fingers. This after-the-fact technique is less than ideal, as it requires the rotor to be secured and realigned from an indeterminate position. Nonetheless, it provides necessary insurance against opening the cryostat, which is an expensive operation. We have tested this backup procedure in the LBNL and UCSD cryostats and have shown it to recenter the rotor to within 1.5 mm of its original position, which is good enough to reattain 2 Hz rotation.

9.4 Data quality

In this section, we review the impact of the CHWP on experiment **data quality** within the context of the noise requirements presented in Section 8.2.3. Specifically, we discuss encoder jitter, magnetic interference, and rotor temperature stability, highlighting the achieved values in Table 8.1.

9.4.1 Angle encoder performance

We use the same hour-long data set presented in Section 9.3.3 to evaluate the angle encoder described in Section 8.5. For purely historical reasons, this test is performed using an Arduino MCU,⁶ which employs two separate 16 MHz clocks, instead of the BBB, which employs one shared 200 MHz clock. Additionally, this test did not utilize PID control and therefore relied on open-loop motor stability to maintain steady rotation. Our primary goal in this section is to measure the angle encoder’s white noise level and compare it to the $\ll 3 \mu\text{rad}/\sqrt{\text{Hz}}$ requirement in Table 8.1.

As described in Section 8.5, the **rotor angle** $\chi(t)$ is reconstructed by linearly interpolating angle encoder ticks to IRIG time t as

$$\chi(\tau_{\text{Enc}}) \longleftrightarrow t(\tau_{\text{IRIG}}), \quad (9.1)$$

where τ_{Enc} and τ_{IRIG} are the MCU clocks for the encoder and IRIG signals, respectively. Deviations between τ_{Enc} and τ_{IRIG} on the Arduino lead to additional noise that will not exist when using the BBB in the field. Therefore, we analyze both $\chi(t)$ and $\tau_{\text{IRIG}}(t)$ to distinguish MCU drifts from those of the rotation mechanism and encoder system.

To facilitate the following discussion, we consider two illustrative angle jitter definitions: the residual after subtracting a quadratic fit

$$\delta\chi_{\text{HWP}}^{\text{poly}}(t) = \chi(t) - P_2^x(t), \quad (9.2)$$

and the residual after subtracting a quadratic fit plus an **angle-dependent function** $A(\chi)$

$$\delta\chi_{\text{HWP}}^{\text{temp}}(t) = \chi(t) - [P_2^x(t) + A(\chi)]. \quad (9.3)$$

⁶Arduino Leonardo: <https://store.arduino.cc/usa/arduino-leonardo-eth>




Figure 9.8: A measurement of encoder performance during 1 hour of testing in the LBNL setup. Figure 9.8a shows a three-second sample of the angle and IRIG jitters. The rotor is spinning at ≈ 2.15 Hz and its angle is sampled 1,140 times per revolution. Figure 9.8b shows PSDs of the angle and IRIG jitters, including a zoomed inset of the $1f_{\text{HWP}}$ peak with finer binning. We multiply $\delta\tau_{\text{IRIG}}^{\text{poly}}$ by $2\pi f_{\text{HWP}}$ to convert it from seconds to radians.

For this test data, we select a simple 2π -repeating template

$$A(\chi) = A(\chi - 2\pi n) ; n \in \{0, 1, \dots, N_{\text{rev}}\} , \quad (9.4)$$

where N_{rev} is the number of completed revolutions in the data set and where $A(\chi - 2\pi n)$ is constructed by taking a window-weighted average of the encoder pattern over all N_{rev} . In practice, a more complex function can be used to remove HWP-synchronous signals in the field [53], but a **uniform template** works well for lab characterization. Additionally, we define **IRIG clock jitter** as

$$\delta\tau_{\text{IRIG}}^{\text{poly}}(t) = \tau_{\text{IRIG}}(t) - P_2^{\tau}(t) , \quad (9.5)$$

which we use to evaluate MCU clock drifts.

A three-second segment of angle jitters $\delta\chi_{\text{HWP}}^{\text{poly}}$ and $\delta\chi_{\text{HWP}}^{\text{temp}}$ are shown in Figure 9.8a. The polynomial-subtracted spectrum $\delta\chi_{\text{HWP}}^{\text{poly}}$ shows a distinct peak-to-peak variation of ≈ 1 mrad at f_{HWP} in addition to higher-order structures, and the template-subtracted spectrum $\delta\chi_{\text{HWP}}^{\text{temp}}$ shows that most of this jitter is removed by $A(\chi)$. A unique property of the SMB is that it does not wobble, even if it is gravitationally imbalanced. Therefore the observed $\delta\chi_{\text{HWP}}^{\text{poly}}$ signal is caused by a slight misalignment between the slotted encoder plate and magnet ring, while higher-frequency structures are caused by non-uniform patterning of the encoder slots. Three seconds of IRIG clock jitter $\delta\tau_{\text{IRIG}}^{\text{poly}}$ is also shown in Figure 9.8a and is effectively

a measurement of Arduino clock noise. Because the IRIG signal is GPS synchronized to ≤ 40 ns, the observed drift is due to the instability of the Arduino’s crystal oscillators, which we discuss further below.

Power spectral densities (PSDs) of the angle and clock jitters are shown in Figure 9.8b. The sharpness of the peaks in the $\delta\chi_{\text{HWP}}^{\text{poly}}$ spectrum demonstrates excellent rotational stability, and the suppression of those peaks in the $\delta\chi_{\text{HWP}}^{\text{temp}}$ spectrum confirms that $A(\chi)$ measures all but a few $\delta\chi_{\text{HWP}}$ features with high signal-to-noise. The $A(\chi)$ subtraction also suppresses side-band power in the f_{HWP} harmonics, better revealing a **white-noise level** of $\approx 0.1 \mu\text{rad}/\sqrt{\text{Hz}}$.

Both $\delta\chi_{\text{HWP}}$ spectra have a 1/f knee of ≈ 1 Hz, which is larger than expected from the Δf_{HWP} measurement in Figure 9.5. Therefore, we compare the angle and IRIG spectra to analyze the contribution of MCU clock drifts to the observed 1/f noise. The two MCU clocks, τ_{IRIG} and τ_{Enc} (see Equation 9.1), fluctuate with $\approx 90\%$ coherence, and because common-mode clock drifts are subtracted during angle-time interpolation, $\delta\chi_{\text{HWP}}$ 1/f power dips beneath that of $2\pi f_{\text{HWP}} \delta\tau_{\text{IRIG}}^{\text{poly}}$ between ≈ 0.05 Hz and the IRIG’s 0.5 Hz Nyquist frequency. Below ≈ 0.05 Hz, drifts in rotational velocity begin to contribute and the $\delta\chi_{\text{HWP}}$ spectra steepen, but even so, MCU-induced 1/f noise remains dominant down to ~ 0.01 Hz. This finding motivated us to replace the Arduino with the BBB, which uses a single shared clock ($\tau_{\text{Enc}} = \tau_{\text{IRIG}}$) and has a measured 1/f knee of < 0.1 Hz. This configuration allows the 1 Hz IRIG signal to fully subtract MCU clock drifts during angle-time interpolation, and therefore we expect significantly improved low-frequency noise when measuring $\chi(t)$ in the field.

9.4.2 Magnetic interference

The CHWP’s motor and rotor both introduce **magnetic interference** that can affect the detectors and their SQUID amplifiers. In particular, interference at $4f_{\text{HWP}}$ mimics a sky signal and must be especially well controlled. The motor energizes three phases across 114 solenoids at $38f_{\text{HWP}}$, and its large multipole number causes its ≈ 20 G field to decay quickly with distance. The magnet ring, on the other hand, has only 16 segments—resulting in a lower multipole number—and a surface magnetization of $\approx 5,000$ G, posing a greater $4f_{\text{HWP}}$ interference concern.

Magnetic field testing is performed at LBNL using a room-temperature magnetometer⁷ placed 1.5 m behind the CHWP assembly at the approximate location of a detector near the focal plane’s edge⁸ (see Figure 2.3). The CHWP rotates steadily at ≈ 2.15 Hz, and the ambient magnetic field is measured at 150 samples per second for 100 s. The results are presented in Figure 9.9. The time-ordered data show a DC field of 926 mG⁹ and a ≈ 5 mG oscillation at $1f_{\text{HWP}}$, while the spectrum shows peaks at $1f_{\text{HWP}}$, $2f_{\text{HWP}}$, and at harmonics

⁷Honeywell HMR2300: <https://aerospace.honeywell.com/>

⁸We calculate that CHWP-induced magnetic fields are larger near the edge of the focal plane than near the center. Therefore, this measurement represents a worst-case scenario for potential detector interference.

⁹This average field applies a ≈ -0.3 mK DC shift to the TES transition temperature T_c , which is less than T_c variation across the focal plane [224] and is calibrated out during detector tuning.




Figure 9.9: A measurement of the magnetic field 1.5 m behind the CHWP assembly during 100 s of continuous rotation at 2.15 Hz in the LBNL cryostat. Figure 9.9a shows 3 s of time-ordered magnetometer data, showing a clear ≈ 5 mG peak-to-peak variation at $1f_{\text{HWP}}$. Figure 9.9b shows the B-field’s power spectral density. Red arrows mark rotation-synchronous peaks, while the magenta arrow marks that of the GM cooler. The unmarked peaks are environmental and are predominantly GM harmonics.

of the GM cooler’s cycle frequency. However, the $4f_{\text{HWP}}$ component lies beneath the noise floor, bounding its amplitude to $< 10 \mu\text{G}/\sqrt{\text{Hz}}$.

Full-system interference testing is performed at UCSD by monitoring detector outputs when the CHWP is spinning. During two minutes of rotation at 2 Hz, data is collected from both optical and **non-optical bolometers**¹⁰ across the full focal plane. While $2f_{\text{HWP}}$ and $4f_{\text{HWP}}$ signals are visible in the optical detectors, as expected, no CHWP-induced signals are detected in *any* of the dark detectors. In addition, because $4f_{\text{HWP}}$ magnetic interference is $\leq 10^{-6}$ that at $1f_{\text{HWP}}$, as shown in Figure 9.9, this test also suggests that no $4f_{\text{HWP}}$ magnetic features will arise even when $\sim 5 \times 10^3$ detectors are coadded during data analysis.

9.4.3 Rotor temperature stability

To evaluate **thermal stability**, we simulate the rotor temperature using the thermal model presented in Section 8.6. Because it is floating, the rotor’s temperature variations are dominated by those of the 50 K stage to which it is radiatively tied. These CHWP fluctuations are expected to be very slow, as the rotor has a ≈ 900 J/K heat capacity at base temperature and a ≈ 10 mW/K coupling to its 50 K surroundings. Due to **diurnal variations** in ambient temperature and changes in telescope elevation—which impact PTR performance—50 K temperature drifts in Chile are substantially larger than those in the lab. Therefore, we use

¹⁰Non-optical bolometers are often called “dark” detectors.




Figure 9.10: Power spectra of measured 50 K temperature variations from PB-2a in Chile and of the resulting simulated rotor temperature variations. The rotor’s large thermal mass and small radiative coupling make it an effective low-pass filter, suppressing IRF fluctuations by $> 100\times$ at 1 mHz.

24 hours of 50 K temperature data from PB-2a, which is operating in the field, to simulate the expected rotor stability. The selected PB-2a data includes telescope slew testing, and therefore the presented 50 K-stage variations represent an upper bound on those expected during science observations.

The measured 50 K-stage and simulated rotor PSDs are shown in Figure 9.10. The rotor acts as a **low-pass filter**, and its simulated drift is $\approx 0.3 \text{ mK}/\sqrt{\text{Hz}}$ at 1 mHz, which is below the $1 \text{ mK}/\sqrt{\text{Hz}}$ requirement in Table 8.1. We note that rotor temperature stability can be substantially improved by regulating the 50 K-stage temperature, which may become necessary for future experiments with tighter noise requirements.

9.5 Discussion

Chapters 8 and 9 have presented the design and evaluation of the PB-2b CHWP. This research advances the area of cryogenic polarization modulators for mm-wave astronomical observations, introducing a new motor, encoder, and grip-and-release mechanism while pushing the superconducting magnetic bearing and sapphire diameters to the largest ever deployed. The CHWP has been evaluated in both a non-optical standalone cryostat and an optical test configuration of the PB-2b receiver. Its performance satisfies all PB-2b-specific requirements and includes a 430 mm clear aperture, a rotor temperature of $< 53 \text{ K}$, $< 1.3 \text{ W}$ of dissipation during continuous operation, rotation frequencies up to 2.8 Hz, $0.1 \mu\text{rad}/\sqrt{\text{Hz}}$ of encoder noise, and $< 10 \mu\text{G}/\sqrt{\text{Hz}}$ of magnetic interference at $4f_{\text{HWP}}$. Additionally, an

emphasis has been placed on system robustness to ensure years of continuous operation in the field. The presented CHWP instrument has deployed to the Chilean observation site and is expected to see first light in 2021. However, one important component is yet to be completed: the sapphire AR coating, which we discuss in Chapter 10.

Chapter 10

Sapphire anti-reflection coating

The final component of the cryogenic half-wave plate (CHWP) system is the sapphire’s **anti-reflection (AR) coating**. Sapphire has a crystal-axis-averaged **refractive index** of $n_{\text{sapphire}} \approx 3.2$, and therefore its **Fresnel reflection coefficient** at normal incidence¹ is

$$R = \frac{(n_{\text{sapphire,avg}} - n_{\text{vacuum}})^2}{(n_{\text{sapphire,avg}} + n_{\text{vacuum}})^2} \approx 0.3, \quad (10.1)$$

where $n_{\text{vacuum}} = 1$ is the refractive index of vacuum. This level of reflection at any refractive surface in the experiment has several problematic ramifications, including decreased optical throughput, beam distortions, and multiple reflections within the cryostat that generate **ghosting**, or the formation of “ghost” images. In addition, sapphire’s birefringence induces **differential transmission** between its ordinary axis with $n_o \approx 3.08$ and extraordinary axis with $n_e \approx 3.35$, generating $2f_{\text{HWP}}$ and $4f_{\text{HWP}}$ HWP synchronous signals (HWPSSs)² in the detector data (see Section 6.5.1). This **intensity-to-polarization (I-to-P)** is difficult to remove during data analysis and can induce detector pathologies such as non-linear responsivity [199]. AR coatings limit these systematic effects by minimizing **reflectivity** (Equation 3.37), and developing high-performance sapphire coatings is an imperative of the CHWP’s development.³

There are many methods to AR coat **cryogenic mm-wave optics**, and given the recent rise of polarization modulators, many of these methods are being investigated on sapphire. In this chapter, we review several coating technologies developed for the Simons Array (SA) CHWP. For a discussion of the CHWP’s sapphire—which is identical to that of the PB-2a warm HWP (WHWP)—see Section 7.2.1, and for a related discussion of the WHWP’s AR coating, see Section 7.2.2.

¹Reflectivity increases at non-normal incidence, and therefore Equation 10.1 represents a lower limit.

²Sapphire’s **differential reflectivity** is $\Delta R = R_e - R_o \approx 3\%$, which for 10 K of atmospheric loading (see Figure 2.1) corresponds to a $2f_{\text{HWP}}$ signal of ≈ 300 mK with a harmonic component at $4f_{\text{HWP}}$ for light at non-normal incidence [170].

³Note that we cannot adopt PB-2a’s *warm* HWP AR coating in Section 7.2.2 on PB-2b’s *cryogenic* HWP for reasons discussed in Section 10.1.




Figure 10.1: A schematic for a quarter-wave AR coating. The electric field oscillates with a free-space wavelength λ_0 and is decomposed into the transmitted component E_{trans} and the reflected component E_{refl} . The two AR layers have indices $n_{\text{air}} < n_1 < n_2 < n_{\text{sapphire}}$ such that the index gradually transforms from that of air/vacuum to that of the sapphire substrate. Each layer has a thickness $d_{\lambda/4}$ according to Equation 10.2 such that the transmitted and reflected waves are 180° out of phase, destructively interfering the reflected wave and maximizing transmission.

10.1 AR coating design

An AR coating suppresses reflectivity by gradually transforming the refractive index from that of air/vacuum $n_{\text{air}} = n_{\text{vacuum}} = 1$ to that of the refractive optic. The more gradual the transformation, the less light is reflected, and therefore in the absence of practical considerations, more AR layers generally corresponds to lower reflectivity.⁴ In addition, reflection can be further suppressed by tuning each layer's thickness to

$$d_{\lambda/4} = \lambda_0 / (4n) \quad (10.2)$$

to maximize **destructive interference** between forward- and backward-traveling light. Here, λ_0 is the wavelength of the incident light in vacuum, and n is the AR layer's refractive index. This thickness construction is called a **quarter-wave coating**, and if all layers are $d_{\lambda/4}$, it achieves zero reflection for incident light of wavelength λ_0 . A schematic of the quarter-wave coating is shown in Figure 10.1.

⁴This statement is not *always* true, as the wavelength, layer thicknesses, and indices matter, but if the coating is properly optimized, more layers indeed perform better.




Figure 10.2: A comparison of one-, two-, and three-layer AR coatings optimized to maximize average transmissivity across the 90 GHz (magenta) and 150 GHz (cyan) detector bands. The sapphire substrate is assumed to have a thickness of 3.75 mm and an axis-averaged refractive index of $n_{\text{sapphire,avg}} = 3.23$. The optimal AR indices, thicknesses, and resulting reflectivities in each band are shown in the table. The largest transmissivity improvement comes between the one- and two-layer coatings.

From the perspective of PB-2b, which has a fixed observation bandwidth, adding more layers to the AR coating adds more **degrees of freedom** over which the coating can be optimized. Figure 10.2 shows fully⁵ optimized [83] one-, two-, and three-layer coatings and their integrated reflectivities across PB-2b’s 90 and 150 GHz bands. A single-layer coating has insufficient bandwidth to achieve reasonable reflectivity in both bands, and therefore the PB-2b CHWP must employ a **multi-layer coating**. A two-layer coating does substantially better, achieving 1.1% reflection in each band, while a three-layer coating does even better, achieving (0.3, 0.1)% reflection at (90, 150) GHz.

Multi-layer AR coatings are a relatively new technology in the field of CMB instrumentation. While multichroic CMB receivers have become increasingly prominent, the problem of achieving high optical throughput across broad bandwidths remains a considerable chal-

⁵“Fully” means that all layer thicknesses and indices are allowed to vary during the optimization.

lenge. Therefore, even though the three-layer coating in Figure 10.2 performs better than the two-layer coating, the PB-2b CHWP’s development focuses primarily on two-layer coatings to minimize technical complexity.⁶

Just as for the PB-2a WHWP’s AR coating (Section 7.2.2), the CHWP’s AR materials must have close to optimal indices and must be transparent at 100 GHz. Because 100 GHz loss tangents (Equation 3.28) are generally lower at cryogenic temperatures, and because thermal radiation is $\sim 10\times$ fainter at 50 K than at 300 K, the catalog of optically viable AR materials is larger for the CHWP than for the WHWP (Figure 7.5). However, diminished concerns about finding index-matched, low-loss dielectrics are replaced with concerns about **cryo-mechanical reliability**, which is a major driver of this chapter’s AR investigations.

10.1.1 Tolerances

The contours and values in Figure 10.2 represent idealized coatings, but in reality, there will be variations and uncertainties in each layer’s thickness and refractive index. Figure 10.3 shows the impact of index and thickness tolerances on reflectivities in the 90 and 150 GHz bands. As we discuss in Section 10.2.1.2, the index of each layer is usually measured before it is machined to thickness, and typical index measurement uncertainties are $< 1\%$. Therefore, assuming the coating’s indices are close to optimal, **thickness tolerances** are the primary performance driver for this chapter’s AR technologies. “Standard” machining uncertainties are $\pm 100\ \mu\text{m}$, which is far too large to ensure adequate AR transmissivity. Therefore, thickness control is an important R&D item for the CHWP coating’s fabrication process.

To set a thickness tolerance, we consider not only the performance impact shown in Figure 10.3b but also the practical limitations of the machining process. While we use a numerical mill to machine the AR layers (see Section 10.2.1.3), several challenging factors impact the milling process, including temperature control, bit wear, and fixturing repeatability. Taking these practical considerations into account, we impose a thickness tolerance of $< 25\ \mu\text{m}$, which maintains an absolute reflectivity of $< 2\%$ and a reflection uniformity of $\pm 0.5\%$.⁷

10.1.2 Objectives

There are several figures of merit when evaluating an AR coating technology. The first merit figure is reflectivity. Because the instrument’s noise-equivalent temperature (NET) is a strong function of the telescope’s **optical throughput** (see Section 3.8), achieving **percent-**

⁶That said, three-layer coatings can be used to combat challenges associated with material selection and fabrication limitations, and for this reason, we consider three-layer coatings with lower priority in parallel. One such consideration is described in Section 10.4.2.

⁷As discussed in Section 6.5.1, non-idealities in the CHWP’s optical stack can lead to HWPSSs, and therefore AR uniformity is an especially important consideration for the sapphire coatings in this chapter.

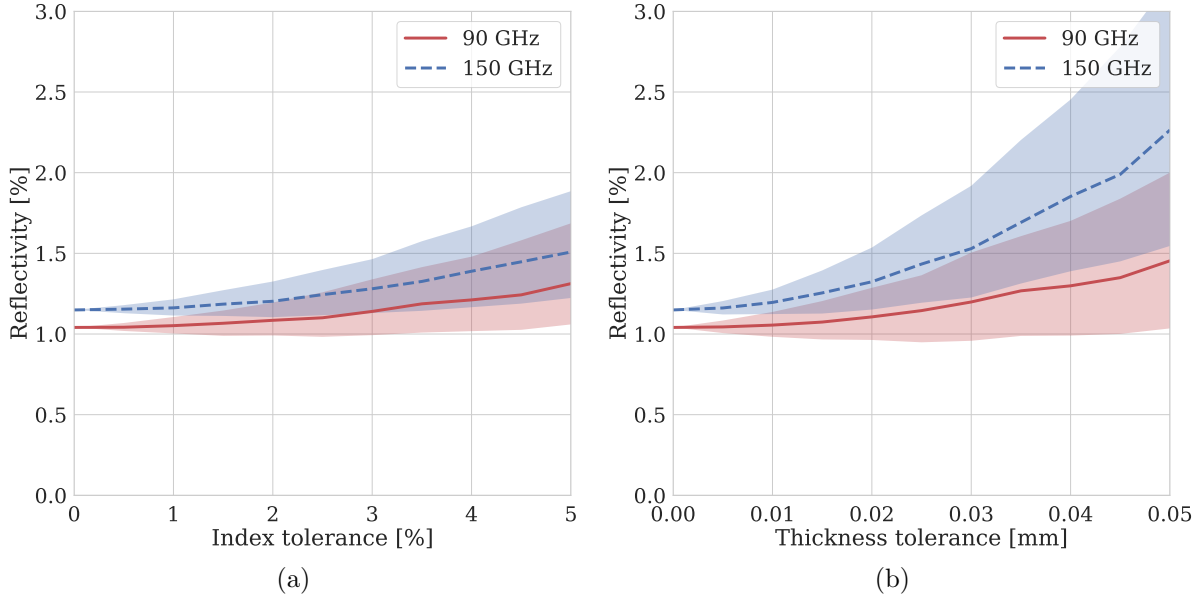


Figure 10.3: The impact of index (Figure 10.3a) and thickness (Figure 10.3b) tolerances on the two-layer AR coating’s reflectivity in PB-2b’s bands. Each tolerance denotes the \pm spread of a uniform distribution centered on the ideal index and thickness values in Figure 10.2, and each case is Monte Carlo simulated 1,000 times. The solid lines denote the medians, while the shaded areas denote the (15, 85)% confidence regions.

level reflectivity⁸ on *each* optic, including the CHWP, is critical to fielding a cutting-edge instrument. The second merit figure is **cryo-mechanical robustness**. The CHWP is cooled to ≈ 50 K, introducing challenges such as **differential thermal contraction**, cryogenic adhesion, low-temperature material properties, and vacuum compatibility. The third merit figure is **manufacturability**. For example, there is a finite number of readily obtainable materials available for mm-wave AR coatings (see, for example, Section 7.2.2). In addition, the coating must be applied with excellent thickness and index control, the fabrication process must be repeatable and reliable, and the application procedure must be of reasonable cost and lead time. The fourth merit figure is **implementation risk**. To reduce project risk both during technology development and after deployment, we aim to leverage the successes of existing technologies. For this reason, many of the AR coatings considered for the PB-2b CHWP utilize materials and techniques used for past experiments.

In Sections 10.2 and 10.3, we review two sapphire AR coating technologies developed for the PB-2b CHWP, recount their development paths, describe their pros and cons, and assess their viability. Each section is titled “A + B AR,” indicating that the bottom AR layer is composed of “A” and the top AR layer is composed of “B.”

⁸There is no clear-cut requirement on the PB-2b AR coating’s reflection or absorption, as the question of deploying an “acceptable” optics tube is a complicated convolution of robustness, project schedule, technology risk, and other factors.




Figure 10.4: A cartoon schematic of the epoxy + epoxy AR coating. The bottom layer is Stycast 2850FT, and the top layer is Stycast 1090. The coating is applied to the outward face of the outermost sapphire pieces in the Pancharatnam stack.

10.2 Epoxy + epoxy AR

Early AR coating development for POLARBEAR-2 demonstrated a two-layer **epoxy + epoxy AR coating** comprising a bottom layer of **Stycast 2850FT**⁹ + Catalyst 23LV and a top layer of **Stycast 1090**¹⁰ + Catalyst 9 [163]. Stycast 2850FT is an epoxy resin loaded with microscopic alumina grains that both raise its mm-wave refractive index and reduce its coefficient of thermal expansion (CTE) to match that of aluminum. Because it is $\approx 30\%$ alumina by weight, 2850FT is an excellent thermal conductor and is therefore a common adhesive in cryogenic engineering. Catalyst 23LV, the chosen hardening agent, offers a stable refractive index vs. mm-wave frequency, gives the pre-cured mixture a low viscosity, and provides a ≈ 2 -hour pot life, which is comfortably long enough for most AR application processes. Stycast 1090 is an epoxy resin loaded with $\sim 50\ \mu\text{m}$ -diameter, hollow silica, nitrogen-filled **microspheres** that lower the epoxy's density, and Catalyst 9 is chosen to minimize the 1090's cured refractive index. During POLARBEAR-2's initial R&D in 2013 [163], the measured refractive indices of each epoxy layer were $n_{2850\text{FT}} = 2.28$ and $n_{1090} = 1.42$, and the resulting transmissivity vs. mm-wave frequency on an alumina substrate is shown in Figure 10.5.

Because the epoxy + epoxy coating is **moldable** and shows excellent performance on a 50 mm-diameter sample, PB-2a selected the technology for its lenses' curved surfaces. The process of applying 2850FT + 1090 to full-scale alumina optics was adopted from a procedure developed at Stanford for BICEP Array [90], and the PB-2a lenses were epoxy coated at UC Berkeley between 2014 and 2016. The lenses' flat surfaces were coated at the High Energy Research Organization (KEK) in Japan using a bottom layer of **thermal-sprayed mullite** and a top layer of an expanded polyimide foam called **Skybond**.¹¹ For more information

⁹Stycast 2850FT: https://www.henkel-adhesives.com/us/en/product/potting-compounds/loctite_stycast_2850ft.html

¹⁰Stycast 1090: https://www.henkel-adhesives.com/us/en/product/potting-compounds/loctite_stycast_1090bk.html

¹¹Skybond foam: https://www.istcorp.jp/en/industrial_material/skybond-foam/




Figure 10.5: Measured performance of the epoxy + epoxy AR coating on a 50 mm-diameter, 3 mm-thick alumina sample at both 300 K and 140 K. While the epoxy coating is lossy at room temperature, its emissivity decreases substantially at cryogenic temperatures. The measured fractional bandwidth is $\approx 65\%$, which is wide enough to cover PB-2b’s 90 and 150 GHz bands. This figure is taken from Rosen, Suzuki, *et al.* (2013) [163].

about the mullite + plastic AR combination,¹² see Section 10.4.2.

The AR-coated PB-2a lenses were not measured directly, as taking the spectrum of a large optic at cryogenic temperatures requires dedicated infrastructure that is time-consuming and expensive to develop. Instead, they were installed in the PB-2a receiver cryostat, and the band-averaged optical throughput of the full instrument was measured at KEK. In addition, the receiver’s near-field beam was mapped in the lab to assess its Gaussianity and angular extent. These tests’ results were satisfactory for deployment, and the PB-2a epoxy coatings were deemed a success [100]. Leveraging the work done for PB-2a, PB-2b adopted the epoxy + epoxy coating for *all* of its alumina and sapphire surfaces, both curved and flat. In this section, we review the epoxy coating process and present the results of its application to sapphire.

10.2.1 Fabrication

The process for applying the epoxy + epoxy coating is:

1. Prepare the substrate.
2. Mold Stycast 2850FT onto the substrate’s bare surface.
3. Mill the 2850FT layer to its target thickness.
4. Mold Stycast 1090 onto the machined 2850FT layer.

¹²Note that Skybond foam was discontinued and is therefore not considered as a possible AR layer in this dissertation.




Figure 10.6: A microscope comparison of a sapphire surface after 24-grit sandblasting at 50 PSI (left) and of a different sapphire surface after abrading with 60-grit diamond sandpaper (right). Sandblasting cracks the sapphire’s surface, making it vulnerable to fracture, while sanding roughens the surface without any visible damage.

5. Mill the Stycast 1090 layer to its target thickness.
6. Strain relieve the layers.
7. Thermal cycle to cryogenic temperatures and inspect for mechanical degradation.

We describe each of these steps in the following subsections, highlighting advancements made as part of PB-2b’s CHWP development.

10.2.1.1 Surface preparation

The first step in Section 10.2.1 is to prepare the sapphire’s surface. Aluminum oxide, of which both sapphire (single-crystal) and alumina (polycrystalline) are composed, is inert and does not easily bond to organic compounds. Therefore, a **mechanical bond**, one in which the epoxy “grabs” microscopic crevices in the sapphire’s surface, is critical to robust adhesion. As manufactured, the sapphire’s surface has an RMS roughness of $\sim 0.1 \mu\text{m}$ (see Figure 7.4), which auxiliary tests have shown to be insufficient for a strong epoxy bond. Therefore, we roughen the surface before coating it.

While PB-2a found **sandblasting**¹³ to be adequate for robust epoxy-to-alumina adhesion, sapphire requires a more nuanced abrasion approach. Figure 10.6 shows a comparison of a sapphire surface that has been sandblasted with 24-grit silicon carbide vs. a surface that has been sanded with 60-grit diamond sandpaper. As is evident in the left photo, sandblasting introduces **microcracks** at the sapphire’s surface, and these fractures can propagate when stressed upon cooling. For this reason, we sand the surface using diamond sandpaper, making sure to roughen both evenly and gently.

¹³Metal Fusion: <https://www.metalfusioninc.com/>




Figure 10.7: Photographs of the AP-134 application technique used for the PB-2b sapphire. The photo on the left shows a bottle of AP-134 (right) next to the felt applicator tip (left) used to apply it. The photo on the right is a snapshot of the application process. The bottle is turned upside-down and squeezed, and the solution soaks the felt tip. The AP-134 is then “painted” onto the sapphire, and its dispensation rate is controlled by how hard the tip is pressed against the surface.

After the surface is roughened, we clean it thoroughly using Scotch-Brite¹⁴ and solvents, and we apply a thin film of **Lord AP-134 adhesion promoter**.¹⁵ AP-134 is a ceramic-compatible primer that allows the epoxy to *chemically* bond to the sapphire and consists of an organosilane dissolved in solvents. When the primer is applied to the sapphire substrate, the solvents quickly evaporate, the agent **hydrolyzes**, and the organosilane bonds to dangling hydroxyl groups on the surface. This curing process sets up a stout, crosslinked silane network with free-radical groups that the epoxide molecules can bond to. The adhesion promoter’s efficacy relies on the AP-134 layer being thin and on the silane network being sufficiently hydrolyzed. During PB-2a’s epoxy development, the AP-134 was wiped across the optic’s surface with a Kimwipe. Because this process has poor thickness control, PB-2b uses **felt applicator tips** from Designetics¹⁶ to provide thinner, more even coverage. The AP-134 layer is left to hydrolyze for 2-3 hours at 50-80% humidity¹⁷ before the coating process advances. Figure 10.7 shows an example applicator tip as well as a photo of the adhesion promoter being applied to the sapphire’s surface. After hydrolyzation is complete, the cured AP-134 leaves a hazy finish whose uniformity can be verified by eye.

10.2.1.2 Coating

After the sapphire surface is prepared, the Stycast 2850FT is mixed and applied. We note that the 1090 application follows a nearly identical procedure, and therefore this section

¹⁴Scotch-Brite: https://www.scotch-brite.com/3M/en_US/scotch-brite/

¹⁵Lord AP-134: <https://www.lord.com/products-and-solutions/chemlok-ap-134-primer>

¹⁶Designetics: <https://designetics.com/>

¹⁷If the ambient humidity is less than 50%, as is sometimes the case during summertime in Berkeley, the cure time is extended appropriately.



Figure 10.8: Photographs of the epoxy AR application and cure. The left photo is of the epoxy resin being mixed in the can to evenly distribute the fillers. The right photo is of the epoxy mixture on the sapphire’s surface right after application. It is coated all the way to the edge, and a barrier of masking tape keeps the epoxy from leaking over the edge. A leveling table keeps the layer flat as it cures.

describes both steps 2 and 4 of the fabrication process in Section 10.2.1. The Stycast application’s foremost goal is to achieve a thin, uniform layer free of **air bubbles**.¹⁸

First, to ensure that the fillers are evenly distributed throughout the epoxy, we warm the resin to $\approx 40^\circ\text{C}$ to reduce its viscosity and mix it with a hand drill and paddle for ≈ 5 min, as shown in the left panel of Figure 10.8. This step is especially important when opening a fresh can of resin, as the fillers may sink (float) to the bottom (top) during the epoxy’s shelf life. After the resin cools to $< 30^\circ$, we extract ≈ 700 g of 2850FT and mix it with ≈ 50 g of Catalyst 23LV (by hand) for ≈ 5 min or until the mixture’s color is even. We then evacuate the mixture to ≈ 3 Torr over 8 min using a high-throughput roughing pump, removing trapped air.¹⁹

After the epoxy is mixed and evacuated, we monitor its temperature with an IR thermometer until it is $< 30^\circ\text{C}$. It is important to control the epoxy’s temperature throughout the mixing process,²⁰ as epoxide catalysis is an **exothermal reaction** that can cause thermal runaway and hence make the epoxy difficult to handle. Once the epoxy is cool, it is slowly poured into a blob at the center of the sapphire’s surface, and the plate is rotated so that gravity spreads the epoxy to the sapphire’s edges. After the surface is fully coated, it is stored on a leveling table to cure for 24 hours in ambient conditions.

For each 2850FT and 1090 layer, we fabricate and measure a 50 mm-diameter, ≈ 20 mm-thick witness sample to determine the layer’s refractive index. Fluctuations in index arise due to fluctuations in filler fraction, and we find variations in both $n_{2850\text{FT}}$ and n_{1090} as large

¹⁸Voids act as mm-wave scatterers, which can substantially reduce the AR coating’s transparency (see Equations 3.36 and 3.35). We do not have an air-bubble requirement for PB-2b, but a good “rule of thumb” at 150 GHz is for bubbles to be $< 100\ \mu\text{m}$ and sparsely populated.

¹⁹The evacuation step has been tuned carefully, as pumping for too long can cause the epoxy to overheat, and pumping too little results in populations of > 0.1 mm air bubbles throughout the AR layer.

²⁰This is especially true during evacuation, as the epoxy cannot cool in a vacuum space.




Figure 10.9: Photographs of epoxy-AR machining and metrology. The left photo shows a roughing cut on a cured 2850FT layer. The roughing cut is followed by a pre-finish cut to check for any z -offsets in the spindle, and the finish cut trims the layer to its target thickness. The photo on the right shows the Renishaw contact probe used to measure both the bare optic surface before the epoxy is cured and the epoxy surface after machining.

as 5% between batches.²¹ Because the ideal thickness for the coating is tied to its refractive index (see Equation 10.2), this index verification is important to minimizing reflectivity.

10.2.1.3 Machining

After the epoxy layer has cured, we machine it to its target thickness. We use a vacuum chuck to affix the piece in a **Haas VM3**²² computer numerical controlled (CNC) mill, and we use a diamond-coated end mill to remove the epoxy. The diamond-coated bit is necessary because the fillers in both epoxies abrade tooled steel and silicon carbide. Stycast particles are hazardous to breathe, and therefore machining dust is evacuated using a vacuum nozzle coupled to a high-efficiency particulate air (HEPA) filter.

Because the tolerance on each layer's thickness is tight, the spindle height must be carefully controlled throughout the machining process, in turn relying on control of the mill's temperature. The **UC Berkeley physics machine shop** is not climate controlled as of December 2020; therefore, we schedule machining cycles during times of stable ambient temperature²³ to avoid thickness gradients. In addition, we find that the mill's **spindle length** increases by $\sim 100 \mu\text{m}$ during the first \sim hour after it is started. Therefore, we run an hour-

²¹The manufacturing batch number is printed on the front of each resin can.

²²Haas VM3: <https://www.haascnc.com/machines/vertical-mills/mold-machines/models/vm-3.html>

²³For example, we typically perform finish cuts during the afternoon, and we are wary about working on especially hot days.

long “warm-up” cycle prior to each epoxy cut, after which we use gauge blocks to calibrate the spindle’s elongation. To minimize the impact of any additional spindle growth during the machining cycle, we employ an aggressive stepover to keep the process time to $\lesssim 1$ hour.

After the warm-up cycle, we machine the epoxy layer in three steps: a rough cut to remove most excess material, a pre-finish cut to cross-check and calibrate the cutter’s z -dimension, and a finish cut. After the pre-finish and finish cycles, the surface is probed using a Renishaw **contact probe**²⁴ with a 1 mm-diameter ruby tip. We probe four points at three diameters as well as one point at the center, and these thirteen measurements are differenced with an identical measurement of the bare sapphire surface (or in the case of the 1090 top layer, the 2850FT’s machined surface) to extract the layer’s thickness and uniformity. This machining and metrology procedure has reliably met our $\pm 25 \mu\text{m}$ tolerances.

Voids within the cured epoxy layer sometimes appear during the machining process, and sometimes they need to be filled. Therefore, after the pre-finish cut, which is typically $\approx 100 \mu\text{m}$ over the target thickness, we mix a small amount of epoxy, roughen the cavities using a steel brush, and fill them in place on the mill. These patches then cure overnight before proceeding to the finish cut.²⁵

10.2.1.4 Strain relieving

After both the 2850FT and 1090 layers have been applied and machined, the sapphire is ‘officially’ AR coated, and if the HWP operated at ambient temperature, we would at this point be finished. However, there are substantial CTE differences between the epoxy layers and the sapphire substrate, and therefore we must **strain relieve** the coating to avoid delamination when cooling the CHWP to 50 K.

As was done for the PB-2a optics, we **laser dice** the epoxy at Laserod.²⁶ Laserod uses a 355 nm, ~ 10 W, nanosecond-pulsed laser to separate the AR coating into detached **islands**. The ultraviolet wavelength is necessary to achieve a $\lesssim 30 \mu\text{m}$ kerf, which is required to avoid polarized diffraction induced by the strain-relief cuts. Also following the findings of PB-2a, we dice with a 1 cm pitch along two orthogonal directions, which discretizes the epoxy layer into a grid of $1 \text{ cm} \times 1 \text{ cm}$ **squares**. To calibrate the cut’s depth, the machinist looks for laser light to shine through the sapphire’s semi-transparent backside and adds $\approx 10\%$ additional passes to ensure complete epoxy separation. Verifying complete separation is critical, as PB-2a found that **incomplete dicing** leads to widespread delamination upon cooling.

Additional investigations of the laser dicing process were performed to better understand how the epoxy is ablated, how the heat-affected zone changes with pulse duration, and how

²⁴Renishaw OMP400: <https://www.renishaw.com/en/omp400-high-accuracy-machine-probe--6089>

²⁵This filling process does not remedy any voids *below* the (250, 450) μm -thick (bottom, top) layer’s surface. However, we have machined a few AR layers completely off and have empirically seen that bubbles tend not to form at material interfaces. When embedded voids do appear, they tend to reside near the optic’s edge, where their optical impact is diminished.

²⁶Laserod: <https://laserod.com/>




Figure 10.10: Photographs of the epoxy-coated sapphire before thermal cycling. The photo on the left shows the fully fabricated epoxy + epoxy AR coating on a 500 mm-diameter sapphire window mounted in the CAPMAP dewar prior to thermal cycling. The faint 1 cm \times 1 cm grid lines are laser cuts for strain relief. The photo on the right is that same piece held with its bare, semi-transparent backside to the window, showing that the dicing lines cut all the way through the AR coating.

overdicing impacts the sapphire’s mechanical strength. These investigations were part of the epoxy + Duroid AR campaign and are therefore discussed in Section 10.3.1.3.

10.2.1.5 Thermal cycling

The final step of the AR fabrication process is to validate cryo-mechanical performance. As demonstrated by Figure 10.5, the refractive indices of 2850FT and 1090 are largely consistent with temperature,²⁷ and the coating’s transmissivity is excellent (on a 50 mm-diameter sample) at ~ 100 K. Therefore, this final step needs only to assess the AR coating’s (full-scale) cryogenic adhesion.

The fully coated sapphire is cooled to ≈ 20 K and warmed back to ≈ 300 K over ≈ 48 hours using a Gifford-McMahon (GM) cooler in the CAPMAP dewar at Lawrence Berkeley National Laboratory (LBNL), shown in Figure 10.10. After the piece is returned to ambient conditions, its surface is inspected for any delaminated epoxy islands, and the adhesion layer is inspected through the sapphire’s semi-transparent backside. If substantial damage is found, then the fabrication process is, in most cases, deemed unsuccessful. On the other hand, if the adhesion is unaffected by thermal cycling, the piece is cooled several more times to assess the coating’s long-term viability. If the piece survives this gauntlet of cryogenic testing, then it is qualified for optical use.²⁸

²⁷This statement has also been verified by independent measurements at Tohoku University in Japan (Makoto Hattori) and at Stanford (Keith Thompson).

²⁸Note that we do not optically measure each full-scale sapphire piece, as doing so at cold temperatures requires dedicated, expensive infrastructure. Instead, we use index measurements of the witness pieces and thickness measurements of the machined layers to calculate each part’s expected performance. This technique is a valid approximation of the realized performance in the absence of any mechanical degradation (see Figure 7.13b).

10.2.2 Performance

We consider two performance metrics for the epoxy + epoxy AR coating: reflectivity at low temperatures and cryo-mechanical robustness. Because a comparison of the epoxy coating's warm and cold transmissivity has already been published on a 50 mm-diameter, un-diced piece, this section focuses on the performance of larger-diameter, diced parts, which are representative of deployable optics. We will discuss the cryo-mechanical performance first, which informs a discussion of optical performance.

10.2.2.1 Cryo-mechanical performance

The primary differences between small-scale epoxy coatings, which have been characterized and published, and full-scale epoxy coatings, which have not, are the diameter and the laser strain relieving. The dominant mechanism for full-scale performance degradation is **cryogenic delamination**, or separation of the AR coating from the optical substrate at low temperatures. Therefore, cryogenic robustness is the most important assessment of the performance evaluation process. If full-scale cryogenic robustness is adequate, then the performance of small-scale pieces will likely represent that of full-scale pieces.

Figure 10.11 shows the results of thermal cycling epoxy-coated 500 mm-diameter and 95 mm-diameter sapphire pieces. Cryogenic delamination is evident, and the epoxy separates by pulling “chunks” from the substrate,²⁹ leaving the surface severely damaged. Even when this delamination occurs only in areas, it progresses with each cooldown and therefore affects a larger fraction of the surface over time.

10.2.2.2 Optical performance


As discussed in the previous section, the epoxy AR coating breaks the sapphire substrate upon cooling. While this result is undesirable, we use an optical measurement to quantify its impact on reflectivity. To eliminate polarization effects in the measurement, we coat a 150 mm-diameter, 3 mm-thick piece of **alumina**, which has nearly identical optical properties to sapphire, and measure it in a **reflectometer** at the University of Michigan.

Figure 10.12a is a photo of the alumina sample, which is coated on both sides, suspended over a bath of liquid nitrogen (LN2), and cooled to ≈ 140 K. It is subtle to see, but the square islands are raised at the dicing lines, and this **partial delamination** creates a non-uniform air gap between the bottom 2850FT epoxy layer and the alumina substrate. Measurements of reflectance between 140-170 GHz at various temperatures are shown in Figure 10.12b, and the air gap clearly increases reflectance with decreasing temperature. This measurement was performed several times, and the result was shown to be repeatable.

²⁹This surface damage occurs even sanding, which is gentler than sandblasting.



(a)



(b)

Figure 10.11: Photographs of cryogenically delaminated Stycast 2850FT + Stycast 1090 on sapphire. The top panel shows the 500 mm-diameter piece in Figure 10.10 after one cooldown. Though nearly every epoxy square delaminates and “curls” at 20 K, most squares flatten again at room temperature, making damaged AR difficult to detect by eye. The bottom panel shows the epoxy’s adhesion layer through the transparent backside of a 95 mm-diameter, polished sapphire sample. The left (right) panels present the sapphire before (after) cooling. Greyed areas indicated AR separation, while shimmering patches indicate fractured sapphire.

10.2.3 Assessment

The epoxy + epoxy AR coating of Stycast 2850FT + Stycast 1090 has several issues. Despite advancements in the coating technique, including gentler roughening and better-controlled adhesion-promoter application, the AR coating separates from the sapphire when cooled to cryogenic temperatures. In fact, the epoxy pulls chunks out of the substrate itself, best shown in Figure 10.11b, indicating that the AR coating’s **large CTE** and **large tensile modulus** may be more than the sapphire can withstand. There are two clear potential remedies to this problem. The first remedy is to laser dice with a finer pitch, shrinking the isolated islands and reducing the stress due to differential contraction. While this modification is

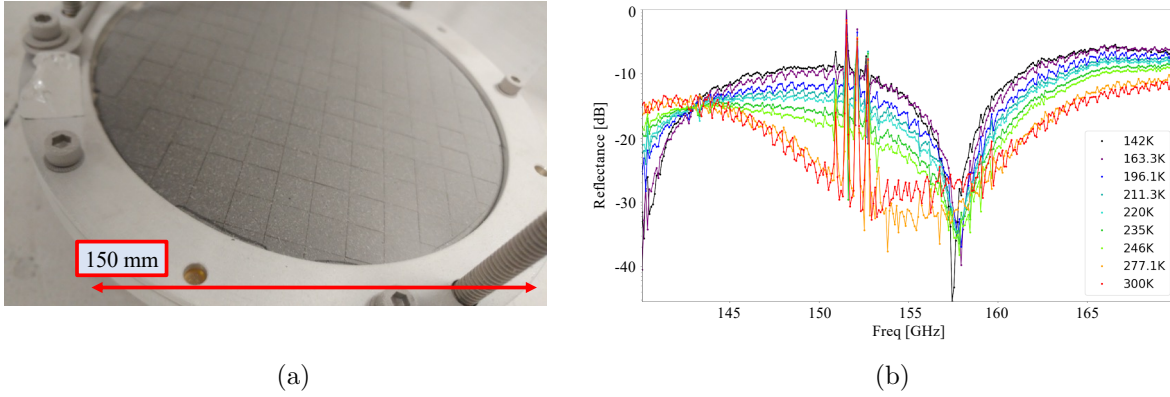


Figure 10.12: Partial delamination of epoxy from 150 mm-diameter alumina coated on both sides, and the resulting impact on reflectance vs. temperature. Figure 10.12a is a photo of the sample at ≈ 140 K suspended above a liquid nitrogen bath. Figure 10.12b shows the resulting reflectance vs. temperature between 140~300 K. The increased reflectance is well modeled by a non-uniform air gap of 50~100 μm between the sapphire and bottom 2850FT epoxy layer. Both figures are courtesy of Shreya Sutariya.

straightforward to implement, it is expensive³⁰ and would quickly make the epoxy + epoxy AR coating unaffordable. In addition, as discussed in Section 10.3.1.3, laser dicing introduces its own risks to the coating's adhesion.

The second remedy is to change the AR coating's top layer. Successful cryo-mechanical tests of the 2850FT layer alone on 95 mm samples suggest that the 1090 layer is to blame for the observed delamination. Stycast 1090 both shrinks more quickly than 2850FT and is stiff at low temperatures, therefore attempting to *peel* the bottom layer away from the substrate.³¹ In addition, Stycast 1090's index has drifted over time, presumably due to changes in the epoxy's manufacturing process,³² with its index in 2019 of $n_{1090} = 1.53$ being higher than $n_{1090} = 1.42$ in 2014. This shift increases the coating's idealized band-averaged reflectivity from $< 1\%$ to $\approx 2\%$, while also tightening tolerances. For these reasons and others, we deem the epoxy + epoxy AR technology to be unfit for the PB-2b CHWP, and we look to other AR solutions in the Sections 10.3 and 10.4.

³⁰Machining time on Laserod's 355 nm laser is expensive, and process time scales with the square of the dicing pitch.

³¹This peeling issue is less problematic for single-layer AR coatings, where (nearly) all of the force due to differential CTE is tangential to the surface, where adhesion is strongest. For this reason, among others, single-layer coatings are substantially simpler to implement for cryogenic applications, and the presented cryo-mechanical challenges are largely unique to multichroic coatings.

³²We have reached out to Loctite several times about whether the filler concentration in Stycast 1090 may have decreased over the years, but the manufacturer assures us that the formula remains unchanged. Keith Thompson at Stanford has also seen 1090's index increase since 2014, and therefore it is unlikely that the detected drift is due to measurement error alone.




Figure 10.13: A cartoon schematic of the epoxy + Duroid AR coating. The bottom layer is Stycast 2850FT, and the top layer is Duroid 5880LZ. The coating is applied to the outward face of the outermost sapphire pieces in the Pancharatnam stack.

10.3 Epoxy + Duroid AR

Section 10.2 demonstrates that Stycast 1090’s CTE, cryogenic tensile modulus, and refractive index are too large to be a viable sapphire AR candidate. Therefore, we replace the 1090 top layer with a circuit board laminate from Rogers Corporation called **RT/duroid 5880LZ**.³³ Duroid 5880LZ is a matrix of PTFE densely loaded with a filler of $\sim 50\ \mu\text{m}$ hollow, nitrogen-filled, aluminosilicate microspheres.³⁴ It is specifically engineered to have a low CTE, and its material composition is more flexible than epoxy’s at cryogenic temperatures.³⁵ Furthermore, Duroid has a refractive index of $n_{5880\text{LZ}} = 1.41$, which is nearly identical to that of the idealized two-layer coating shown in Figure 10.2, and its manufacturing tolerances are remarkably tight, nullifying concerns about batch-to-batch variability. Lastly, Duroid is readily available at low cost and with short lead times, making it amenable to rapid prototyping and development. Motivated by these advantages, we investigate the viability of a Stycast 2850FT + Duroid 5880LZ AR coating.

10.3.1 Fabrication

The fabrication process for epoxy + Duroid is nearly identical to that of epoxy + epoxy, except for a few modifications that are marked with an asterisk:

1. Prepare the substrate.
2. Mold Stycast 2850FT onto the bare substrate.
3. Mill the 2850FT layer to its target thickness.
4. * Glue the Duroid to the machined 2850FT layer.
5. * Machine the Duroid to its target thickness.

³³Rogers Duroid 5880LZ laminate: <https://www.rogerscorp.com/Advanced-Connectivity-Solutions/RT-duroid-Laminates/RT-duroid-5880LZ-Laminates>

³⁴While the microspheres found in 5880LZ are similar to those in Stycast 1090, 1090’s spheres are silica while 5880LZ’s are aluminosilicate. Duroid’s fillers are specifically engineered to improve thermal stability.

³⁵Duroid’s cryogenic modulus has not been measured, but 5880LZ is certainly flimsier than 1090 at room temperature, and PTFE is known to maintain its properties to low temperatures.




Figure 10.14: A comparison of the epoxy + epoxy and epoxy + Duroid AR coatings overplotted onto PB-2b's 90 (magenta) and 150 GHz (cyan) detector bands. The sapphire substrate is assumed to have a thickness of 3.75 mm and an axis-averaged refractive index of $n_{\text{sapphire}} = 3.23$. The band-integrated reflectivity for each band is shown in the table. Using Duroid instead of Stycast 1090 for the top AR layer substantially improves reflectivity.

6. * Strain relieve the layers.
7. * Anneal the AR coating.
8. Thermal cycle to cryogenic temperatures and inspect for any mechanical degradation.

We cover the starred steps in the following subsections and refer the reader to Section 10.2.1 for details about surface preparation, 2850FT application, and thermal cycling.

10.3.1.1 Duroid application

Duroid 5880LZ is procured from Rogers Corporation as a 600 mm \times 600 mm \times 0.5 mm panel plated with 30 μm of copper. We send the Duroid panels to Westak Circuits,³⁶ who uses conveyORIZED ammonium chloride baths to strip the copper and water rinses to wash

³⁶Westak Circuits: <https://www.westak.com/>




Figure 10.15: Photographs of the EpoTek 301-2 application for bonding the Duroid 5880LZ top AR layer to the machined Stycast 2850FT bottom AR layer. The left panel shows the author spraying the transparent 301-2 onto the 2850FT layer in a fume hood. The right panel shows the EpoTek-covered 2850FT layer after spraying is complete.

away etchant residue.³⁷ After the copper is removed, one side of the Duroid sheet is gently roughened using 140 grit aluminum oxide sandpaper to improve adhesion to the bottom 2850FT layer. Note that we do not roughen the machined Stycast, as the mill's stepover size leaves a large enough residual roughness to facilitate strong bonding.³⁸

To affix the Duroid to the 2850FT, we use **EpoTek 301-2**, an optical epoxy with a viscosity of ~ 300 cPs, a pot life of 8 hours, and a low storage modulus of ~ 300 kpsi that persists to cryogenic temperatures. These properties are favorable for the Duroid-on-epoxy application, as a low viscosity enables a $< 10 \mu\text{m}$ bondline,³⁹ a long pot life provides flexibility to the fabrication process, and a **low storage modulus** absorbs differential contraction between the Duroid and 2850FT. EpoTek 301-2 was recommended to us by Ed Wollack at NASA Goddard and was used to bond the AdvACT metamaterial ambient HWP at the University of Michigan [36]. Therefore, we have imported aspects of ACT's procedure into our own.

We investigate two methods to apply the EpoTek 301-2. The first application method is a variation of the adhesion-promoter application process described in Section 10.2.1.1. Designetics makes custom applicator tips for a wide range of fluid viscosities, and they designed a specialized tip to paint a $\approx 5 \mu\text{m}$ -thick EpoTek layer onto the 2850FT. The second application method is to spray EpoTek using a paint gun. This spraying can be done

³⁷In a circuit board's typical application, a circuit pattern is etched into the copper, but for AR coating, we etch away the conductor entirely.

³⁸It is worth noting that we use an epoxy to glue the Duroid to the 2850FT, and because the interfacing epoxy chemically bonds to 2850FT epoxy, surface roughness is less important than when bonding to the 2850FT than when bonding to the sapphire substrate.

³⁹See Section 7.2.2 for a discussion of how glue-layer thickness impacts HWP reflectivity. Note that the glue's emissivity is much less important at cryogenic temperatures than in ambient conditions.

without a thinning agent,⁴⁰ as the EpoTek has a low enough viscosity to **aerosolize** under $\gtrsim 20$ psi of pressure. In this dissertation, we pursue the second method, which consistently achieves a thinner bondline than the painting method.

Figure 10.15 shows the EpoTek application process. When spraying, close attention is paid to the discharge rate, air pressure, raster scanning speed, and the nozzle-to-surface distance to ensure even coverage. After the 2850FT is sufficiently coated with EpoTek, the Duroid is laid onto the bondline, and the assembly is **vacuum bagged** for 48 hours.⁴¹ Vacuum pressing the Duroid and Stycast-covered sapphire together provides a large amount of even and repeatable pressure.

10.3.1.2 Duroid machining

To machine the Duroid to its target thickness, we adopt the procedure in Section 10.2.1.3, with a few modifications. The Duroid’s manufactured thickness is $500\ \mu\text{m}$, and the AR coating’s target thickness is $430\ \mu\text{m}$; therefore, the machining process only removes $\approx 15\%$ from the Duroid stock. In order to skim the Duroid effectively, we run a 70 mm-diameter **fly cutter**, whose tip is angled at 15° to the surface, at 5,000 RPM. We do a pre-finish cut $\approx 25\ \mu\text{m}$ above the target thickness to calibrate the spindle’s z -offset before proceeding to the final cut. Because the Duroid’s softness causes it to recoil when pressed, we use both gauge blocks and the Renishaw contact probe to check the layer’s dimension.

10.3.1.3 Strain relieving

One possible explanation for AR delamination (see Section 10.2.2.1) is that laser cutting seeds separation between the bottom 2850FT layer and the sapphire substrate. Therefore, during the epoxy + Duroid development, we performed several tests to better understand the impact of laser ablation on adhesion. As described in Section 10.2.1.4, we use a $\sim 10\ \text{W}$, 355 nm, nanosecond-pulse laser at Laserod to cut strain relief lines, and we study three laser parameters: pulse duration, the number of passes on each cut, and pulse repetition rate.

The first investigated parameter is **pulse duration**. When laser machining, a shorter pulse duration reduces the laser’s **heat-affected zone (HAZ)**—or the area around the laser’s beam that is “heated” rather than “ablated”—leading to a cleaner cut. Minimizing the HAZ is important for AR strain relieving, as a large HAZ can melt the epoxy and weaken adhesion at the dicing lines.⁴² The AR squares’ corners are particularly vulnerable, as they are ablated along both orthogonal directions *and* they experience more contraction along the

⁴⁰Avoiding thinners, such as toluene, is important for a maximally strong bond, as solvents weaken the epoxy’s polymer cross-linking.

⁴¹EpoTek’s cure time at room temperature is two full days.

⁴²Because the AR coating has two layers, this weakened edge allows the top layer to more easily peel the bottom layer away from the sapphire.




Figure 10.16: Microscope photos of the laser dicing lines for various pulse durations and pass numbers. The left panel shows scanning electron microscope (SEM) images of cuts with nanosecond and femtosecond pulses. The right panel shows optical microscope images through the fully-transparent backside of two AR-coated sapphire samples (similar to those in Figure 10.11b): one with 100% of the passes needed to fully separate the epoxy AR islands, and another with 150% of the needed passes.

island’s diagonal direction.⁴³ For these reasons, we abandon a nanosecond pulse duration for a **femtosecond pulse**, and the difference is shown in the left panel of Figure 10.16. As is evident in the scanning electron microscope (SEM) images, a shorter pulse results in a cleaner cut with less melted epoxy.

The second investigated parameter is the **number of passes** made by the laser on each dicing line. PB-2a found that if the Stycast 2850FT + 1090 AR coating’s $1\text{ cm} \times 1\text{ cm}$ islands are not completely independent, the AR coating delaminates in sheets at cryogenic temperatures. Therefore, during the early stages of PB-2b’s AR development, we found the number of passes needed to sever the AR coating,⁴⁴ and then multiplied by 150%.⁴⁵ These “excess” passes ensured that the diced islands were truly detached from one another. However, as epoxy + Duroid development progressed, we became concerned about the effect of this **overcutting** on the sapphire substrate. The right panel of Figure 10.16 compares a sapphire sample with 100% of the needed passes to one with 150%. As seen under the microscope through the sapphire’s transparent backside, overcutting cracks the sapphire, and these weaknesses seed separation between the coating and the substrate. To avoid overcutting *and* ensure complete separation, we calibrate the number of passes by cutting a test line at the optic’s edge⁴⁶ and looking for laser light to peek through the sapphire’s

⁴³We investigated the possibility of “skipping” over existing lines when cutting along the second orthogonal direction. While we think this technique can be implemented using Laserod’s equipment and expertise, it falls just beyond the scope of this dissertation’s research.

⁴⁴We calibrate the number of passes on free-standing epoxy squares that had delaminated from earlier optics.

⁴⁵While 50% extra cuts might seem excessive, cut depth is not linear with the number of passes. Most material is ablated during the first $\sim 20\%$ of passes, but as the groove deepens, it acts as an aperture that decreases laser power at the cut’s base.

⁴⁶We oversize the CHWP’s AR coated area such that these calibration cuts lie outside the optically active region.

transparent backside.⁴⁷ Using this technique, we avoid sapphire cracking that may seed delamination.

The third parameter is **pulse repetition rate**, which, when holding the feed rate constant, determines the **pulse overlap**. In PB-2a, the laser was run at 100 kHz, which, given a 40 mm/s feed rate and a $\sim 30\ \mu\text{m}$ laser spot size, corresponds to a pulse overlap of $\sim 75\%$. This level of overlap is often desired in typical laser machining applications, as it ensures a clean-cut edge. However, because we are concerned about HAZ-induced melting, such an overlap is undesirable for epoxy AR dicing. Therefore, we reduce the repetition rate to 25 kHz.⁴⁸

These three laser dicing modifications together improved the strain relief cuts' integrity and therefore improved adhesion performance on 95 mm samples.. Given the dicing's 1 cm \times 1 cm length scale, we expect these laser-cutting improvements to transfer to 500 mm pieces.

10.3.1.4 AR annealing

The final step before thermal cycling the epoxy + Duroid AR coating is to **relax** the Stycast 2850FT layer. As the bottom layer cures, epoxide molecules in the resin are joined by amine molecules in the catalyst to form a three-dimensional, cross-linked web of long-range polymer chains. However, as it **phase transitions** from a liquid to a solid, the epoxy shrinks, and because cross-linking is a stochastic process, complex stress profiles develop within the epoxy's bulk. Such effects are typically not a concern when using Stycast 2850FT for potting, but because PB-2b's epoxy coatings are large-volume ($\sim 700\ \text{g}$ of resin) and high-aspect-ratio (500 mm in diameter but only $250\ \mu\text{m}$ in thickness), these stresses can compound CTE-induced stresses and increase the probability of cryogenic delamination.

To demonstrate the phenomenology of **cure-induced stress**, Figure 10.17 shows the impact of annealing a 50 mm-diameter, $300\ \mu\text{m}$ -thick, freestanding epoxy sample milled on a vacuum chuck, just as for a cured epoxy layer on sapphire. Without annealing, the epoxy wafer is demonstrably warped, indicating substantial cure-induced and machining-induced⁴⁹ stresses. Following the advice of epoxy experts at Ellsworth Adhesives,⁵⁰ we alleviate these stresses via **annealing**.

The goal of the annealing process is to relax the epoxy while not thermally shocking it. We use a large Despatch oven⁵¹ at LBNL, which has an in-chamber thermocouple and can

⁴⁷Even though 355 nm is technically outside the visible spectrum, the human eye has a logarithmic roll-off into the UV, and because the laser is so bright, our eyes can see a faint blue hue when the laser's beam is scattered.

⁴⁸Power conservation suggests that decreasing the repetition rate by $4\times$ should mandate more $4\times$ more passes to achieve the same cut depth. However, this relationship was found to be less severe in practice, with the decreased pulse rate only requiring a moderate increase in overall machining time.

⁴⁹We notice that three samples similar to that in Figure 10.17 bow *toward* the end mill, a trend that implicates some amount of machine-induced stress.

⁵⁰Ellsworth: <https://www.ellsworth.com/>

⁵¹Despatch V-series oven: <https://www.despatch.com/>

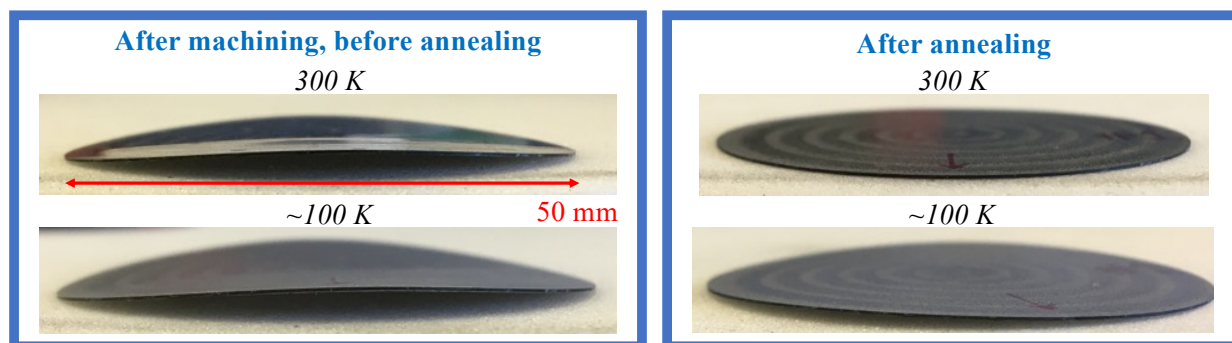


Figure 10.17: The impact of annealing on epoxy stress. The left panel shows a 50 mm-diameter, $300\text{ }\mu\text{m}$ -thick epoxy sample at both 300 K and $\sim 100\text{ K}$ after dunking in LN₂. Warping indicates cure-induced and machining-induced stresses. The right panel shows the same piece after annealing. Post-annealing low-temperature bowing is likely due to stochastic variations in the epoxy’s cross-linked structure, which appears to be sub-dominant to the cure-induced stresses alleviated by annealing.

simultaneously house several full-scale optics. The **annealing profile** is preprogrammed and includes the following steps:

1. Warm to 60° C over 3 hours.
2. Soak at 60° C for 3 hours.
3. Cool to 40° C over 1 hour.
4. Warm to 65° C over 2 hours.
5. Cool to 25° C over 2 hours.

When cured at room temperature, the **glass transition temperature** of Stycast 2850FT is $\approx 40^\circ\text{ C}$, and therefore steps 1 and 2 soften the epoxy and allow it to relax.⁵² After soaking, steps 3-5 again soften the epoxy and slowly cool it to harden. This annealing profile is similar to those used for other plastics, such as polyethylenes, and is well established in industrial settings.

10.3.2 Performance

Similar to with the epoxy + epoxy AR coating, we consider both the optical and cryomechanical performance of the Stycast 2850FT + Duroid 5880LZ coating. Guided by our experiences with epoxy + epoxy, we focus particularly on cryogenic adhesion, knowing that if the coating survives thermal cycling, its optical performance will likely match expectations.⁵³

⁵²Raising the epoxy above its glass transition temperature makes it slightly malleable and rubbery. However, it does not reflow, and therefore the annealing process has no impact on the bond between the 2850FT and the sapphire.

⁵³As noted in Section 7.4.3, the expectation is set by the simulated AR performance assuming measured indices and thicknesses for the individual layers.

10.3.2.1 Optical performance

The first sample fabricated for optical testing is a 150 mm-diameter, 3 mm-thick piece of alumina coated on both sides, shown in Figure 10.18a. This piece is first measured at ambient temperature using two **coherent-source setups**: one that measures reflectivity at the University of Michigan and one that measures transmissivity at UC Berkeley.⁵⁴ Both setups use narrow-band, tunable sources and coherent detectors to measure the frequency-dependent amplitude and phase of the sample’s reflected/transmitted signal. The results of both measurements are shown in Figure 10.18b, and the measured reflectivity at (90, 150) GHz is (0.6, 1.1)%. Attenuation in the transmissivity measurement arises because both Stycast 2850FT and Duroid 5880LZ are lossy at room temperature. However, measurements of both materials at ~ 100 K show substantial loss reductions at cryogenic temperatures,⁵⁵ and therefore the expected band-averaged transmissivity is $\approx 98\%$ at the CHWP’s operating temperature.⁵⁶

Encouraged by the outstanding reflectivity of the warm sample in Figure 10.18 and by the index stability of 2850FT and Duroid down to cryogenic temperatures, we next measure the sample’s reflectivity at ≈ 140 K. We employ the same reflectivity setup used for the warm measurement, except with the sample suspended over an **LN2 bath**. Careful analysis models the impact of adding LN2⁵⁷ to the optical path, and heaters prevent ice from forming on the Styrofoam⁵⁸ cryogen bucket. Several calibration measurements are made with both a fully-reflective plate and a slab of high-density polyethylene (HDPE) to cross-check for systematic effects in the setup. The results of the cryogenic reflection measurement are shown in Figure 10.19.

Unfortunately, a similar effect to that seen with the epoxy + epoxy AR coating is also seen with epoxy + Duroid. Upon cooling, the epoxy squares delaminate at the dicing lines and **curl** to form an air gap between the bottom 2850FT AR layer and the alumina substrate. While the resulting reflectivity increase is not as severe as for epoxy + epoxy (Figure 10.11), epoxy + Duroid’s cryogenic reflectivity is $3\text{--}4\times$ larger than nominal, necessitating improved adhesion.

10.3.2.2 Cryo-mechanical performance

While epoxy + Duroid’s *cryogenic* reflectivity is unacceptable, its *warm* reflectivity in Figure 10.18b is nearly ideal. Therefore, we launched a dedicated R&D campaign to “fix” the epoxy + Duroid’s cryogenic adhesion and “save” the technology’s viability. This adhesion R&D was largely performed on 95 mm-diameter sapphire samples, which are quick to coat,


⁵⁴The measurement apparatuses at Michigan and Berkeley are courtesy of Jeff McMahon and Dick Plambeck, respectively.

⁵⁵Cryogenic loss measurements of Stycast and Duroid are courtesy of Makoto Hattori at Tohoku University.

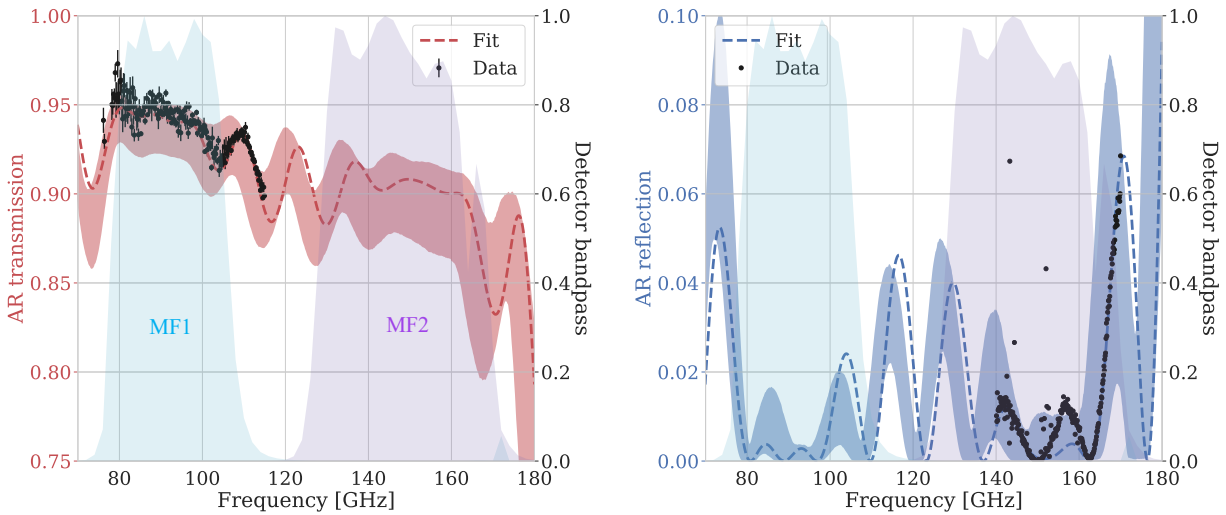
⁵⁶The quoted transmissivity is actually across Simons Observatory’s mid-frequency bands, which are nearly identical to those of PB-2b.

⁵⁷LN2 has a refractive index of $n_{\text{LN2}} \approx 1.2$ at millimeter wavelengths.

⁵⁸Styrofoam, which is expanded polystyrene, is essentially transparent at ~ 100 GHz.



(a)



(b)

Figure 10.18: An ambient measurement of epoxy + Duroid on a 150 mm-diameter alumina sample. Figure 10.18a shows a photo of the sample at room temperature before any thermal cycles. The sample is coated on both sides, and the dicing pattern is $1\text{ cm} \times 1\text{ cm}$, just as for the epoxy + epoxy coatings. Figure 10.18b shows the results of the transmissivity (left) and reflectivity (right) measurements. The dotted lines are a simultaneous fit to both datasets, and the shaded bands represent the expectation from measurements of the individual-layer indices, loss tangents, and thicknesses. The fit and expectation are largely consistent. The observed frequency shift is likely due to a slight mismeasurement of the alumina substrate, which was measured before coating. The shaded cyan and magenta regions are Simons Observatory’s 90 and 150 GHz detector bands, which are nearly identical to those of PB-2b. The transmission and reflection data are courtesy of Dick Plambeck and Shreya Sutariya, respectively.

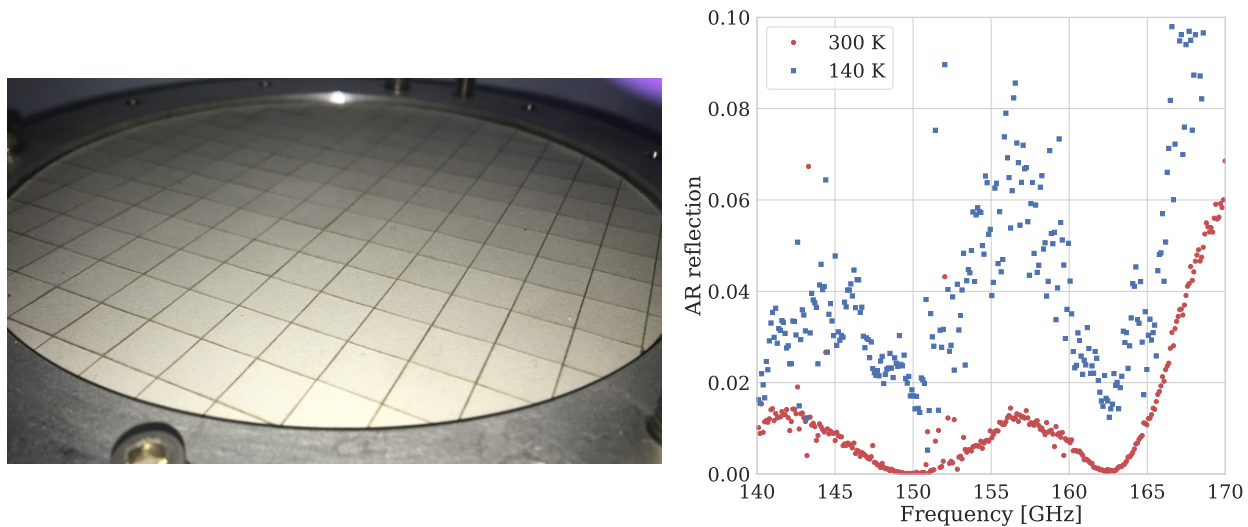


Figure 10.19: A cryogenic measurement of epoxy + Duroid on the 150 mm-diameter alumina sample in Figure 10.18a. The left panel shows the square islands curling at ≈ 140 K, and the right panel shows the resulting reflectivity increase between 300 K and 140 K. The reflectometry data is courtesy of Shreya Sutariya.

cheap to laser dice, and easy to cool. A dozen or so rapid-turnaround experiments led to many of the fabrication advancements presented in Section 10.3.1, and the combination of these process improvements eliminated delamination on small samples. Because the length scale of differential-CTE-induced stress is set by the dicing's 1 cm pitch, improvements on 95 mm sapphire should, in theory, transfer to 500 mm sapphire. Therefore, once the epoxy + Duroid fabrication process was demonstrated on small pieces, we coated the PB-2b CHWP's sapphire.⁵⁹

The result of thermal cycling the PB-2b sapphire to ≈ 20 K in a vacuum chamber over ≈ 48 hours is shown in Figure 10.20. As is most clearly seen through the sapphire's semi-transparent backside, the AR coating delaminated over $> 90\%$ of the surface, an unexpected result given the success on 95 mm pieces.

10.3.3 Assessment

Although Duroid 5880LZ's CTE and tensile modulus are lower than those of Stycast 1090, the epoxy + Duroid coating does not solve the defunct epoxy + epoxy technology's delamination problems. Therefore, the sapphire pieces in Figure 10.20 were deemed unfit for deployment on PB-2b. The degree of cryogenic delamination witnessed on the PB-2b sapphire was surprising. As mentioned in Section 10.3.2.2, the robustness of epoxy + Duroid AR on 95 mm sapphire should transfer to 500 mm, as the length scale of differential-CTE-induced

⁵⁹The CHWP has coatings on one side of its two outermost sapphire pieces. See Section 6.4 for more details about achromatic HWPs.




Figure 10.20: The results of thermal cycling epoxy + Duroid on 500 mm-diameter sapphire to ≈ 20 K. The top panel shows the coated (left) and uncoated (right) sides of a PB-2b piece before cooling, and the bottom panel shows the same piece after cooling. The color contrast observed through the sapphire's semi-transparent backside before and after cooling demonstrates widespread delamination.

stress is 1 cm for both diameters. However, this reasoning proved to be flawed, and in this section, we pose a few theories that are being examined by ongoing R&D.

The first theory for the epoxy + Duroid's large-piece delamination is that adhesion degrades with increasing diameter. During the epoxy's cure, the formation of tight-knit polymer chains shrinks the Stycast 2850FT bulk by 0.3%.⁶⁰ This shrinkage occurs slowly during the cure⁶¹ and amounts to a maximum $\Delta L = (25, 150) \mu\text{m}$ for the (95, 500) mm-diameter samples. Because the coating is diced, this stress is relieved before the coating is cooled; however, the epoxy's cure shrinkage might compromise its adhesion to the sapphire or even introduce microfractures at the sapphire's surface.⁶² This effect would be more impactful at larger diameters and could therefore explain the discrepant performance between 500 mm

⁶⁰Stycast 2850FT datasheet: <https://webaps.ellsworth.com/edl/Actions/GetLibraryFile.aspx?document=7673&language=en>

⁶¹According to conversations with consultants at Ellsworth Adhesives, the shrinkage is not linear in time, but instead, the vast majority of cross-linking occurs during the cure's first ~ 12 hr.

⁶²After cooling the AR coated 500 mm-diameter sapphire, we indeed find widespread surface damage similar to that shown in Figure 10.11b.

and 95 mm pieces. The second theory is that the sapphire plates themselves may have varying degrees of cryogenic robustness due to differences in the sapphire’s growth or machining practices. Therefore, we are investigating the sapphire’s quality, finish, and manufacturing technique more carefully as AR development continues.⁶³ The final theory is that the base temperature of the thermal cycle impacts robustness, as the 500 mm pieces were cooled to 20 K while the 95 mm pieces were cooled to 80 K. Most materials—including sapphire and plastics—experience nearly all cryogenic contraction above ~ 100 K. Therefore, tests using LN2 are typically indicative of cryo-mechanical behavior at lower temperatures. However, it could be that an unexpected physical mechanism causes the coating to fail between 80 K and 20 K, and this hypothesis is testable by cooling small samples to 20 K.

In summary, replacing the Stycast 1090 top AR layer with Duroid 5880LZ improves ambient optical performance and improves cryo-mechanical performance on 95 mm samples. However, widespread adhesive failure on 500 mm samples requires further R&D to validate the epoxy + Duroid technology for use on PB-2b’s CHWP.

10.4 Ongoing work

As of December 2020, the PB-2b CHWP remains without an AR coating. Coating the sapphire is an urgent outstanding item for PB-2b, and several research thrusts are being pursued in parallel. We briefly review each pursuit in this section, highlighting their strengths and challenges.

10.4.1 Epoxy + Duroid

The primary problem with the Stycast 2850FT + Duroid 5880LZ coating is adhesion between the bottom layer and the sapphire. One hypothesis for this issue is that the Stycast’s shrinkage when curing on the sapphire’s surface prevents a sufficiently robust bond. To solve this problem, we aim to cure the epoxy as a **standalone wafer**, glue this cured wafer to the sapphire using a thin layer of EpoTek 301-2, machine the glued wafer to its target thickness, and proceed with the remaining AR fabrication as presented in Sections 10.2.1 and 10.3.1. The primary challenge with this technique is to achieve an EpoTek layer that is $< 10 \mu\text{m}$ (see Figure 7.6) and free of air pockets. The Stycast wafer technique is currently being cryo-mechanically demonstrated on both small and large samples, and an optical measurement will soon follow.

10.4.2 Mullite + Duroid

The problematic player in the Stycast 2850FT + Duroid 5880LZ AR coating appears to be the bottom Stycast layer. One strategy is to replace that layer with a material that better adheres to the sapphire, and we look to PB-2a’s experience for guidance. PB-2a

⁶³For example, annealing the sapphire after grinding could improve its surface’s robustness.

used two different coatings on its alumina optics: epoxy + epoxy on all curved surfaces (see Section 10.2) and **mullite + Skybond** [92] on all flat surfaces. Mullite ($3\text{Al}_2\text{O}_3\cdot 2\text{SiO}_2$) is a silicate with a refractive index of $n_{\text{mullite}} = 2.52$ and a **small CTE** that is well matched to sapphire. Skybond is an expanded polyimide foam whose density is tuned for an index of $n_{\text{Skybond}} = 1.43$. The mullite + Skybond coating was cryogenically robust on PB-2a’s flat surfaces and achieved reflectivities of (2.1, 1.1)% in the (90, 150) GHz bands [92]. However, mullite + Skybond is not considered for the PB-2b CHWP because shortly after PB-2a’s AR fabrication, Skybond manufacturing was discontinued. Nonetheless, we seek to leverage PB-2a’s success by replacing the Skybond layer with Duroid 5880LZ, which has an index of $n_{\text{Duroid}} = 1.41$ that is well matched to mullite.

Mullite is applied via **thermal spraying** (sometimes called **plasma spraying**). During thermal spraying, mullite powder is heated into a plasma, is accelerated using an electric field, and is discharged through a nozzle. The mullite plasma then sticks to the sapphire and cools upon contact, and the spray gun is raster scanned across the surface to achieve even coverage. Thermal spraying is an attractive technology for mm-wave AR coatings because the spray parameters—such as deposition rate, raster rate, nozzle distance, etc.—can be tuned to control each layer’s density, uniformity, and thickness [97]. Mullite spraying is conducted commercially in Japan⁶⁴ and is facilitated by our collaborators at the KEK High Energy Research Organization and the University of Tokyo. After the mullite is applied, we use aerosolized EpoTek 302-1 (Section 10.3.1.1) to apply the Duroid 5880LZ layer, which is pre-machined to the target thickness. After the EpoTek has cured, the Duroid is diced using a razor blade into $\sim 5 \times 5$ cm square islands while the mullite layer is not diced.

The mullite + Duroid coating has shown promising optical and cryo-mechanical performance on small samples, and large-scale testing on both alumina and sapphire is ongoing. Inspired by mullite’s successes and by EpoTek 301-2’s cryo compatibility, we are also investigating a **three-layer AR coating** of mullite + Duroid 6002 ($n_{6002} \approx 1.7$) + porous PTFE ($n_{\text{PTFE}} \approx 1.2$), which would achieve $< 1\%$ reflectivity in both PB-2b bands. This technology leverages many lessons learned during previous sapphire AR coating campaigns and could become an attractive option for both Simons Observatory and CMB Stage-4.

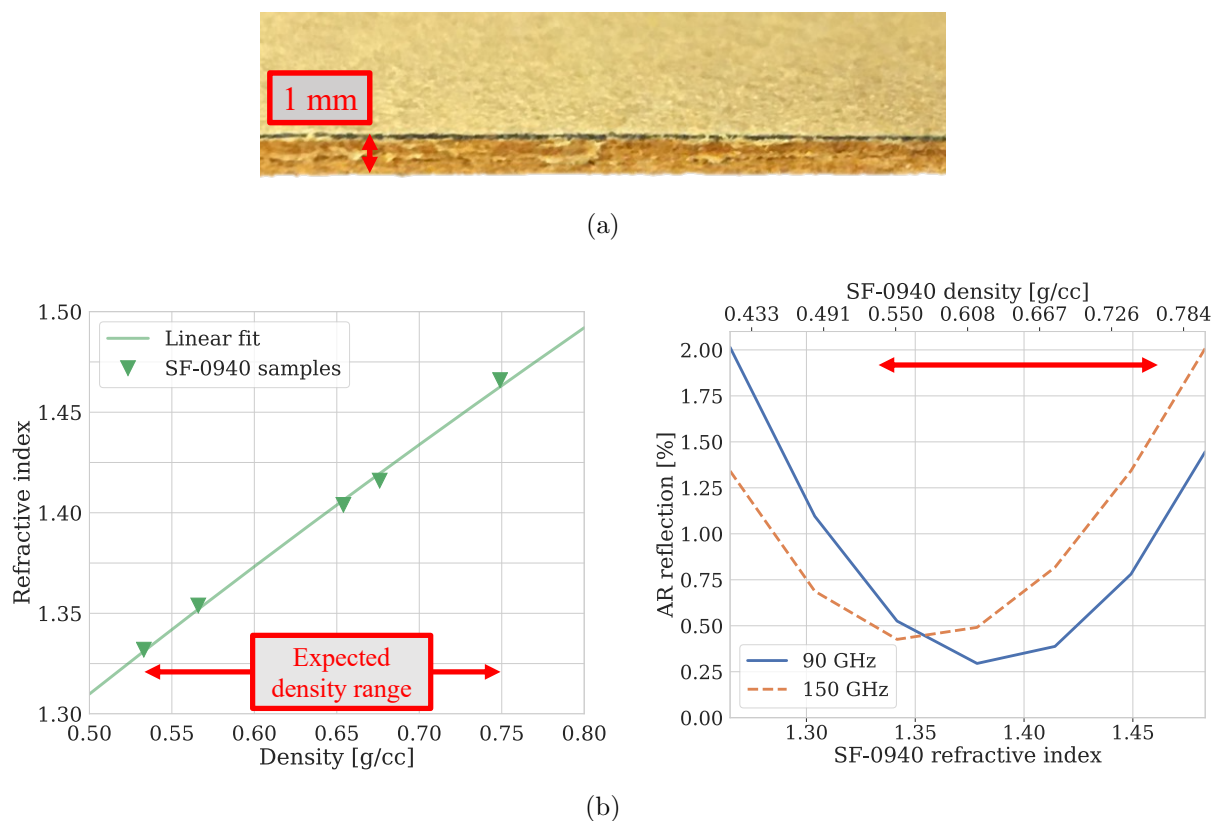
10.4.3 Vespel SF-0940

The fundamental cryo-mechanical challenge with any two-layer AR coating is that the top layer shrinks more quickly than the bottom layer⁶⁵ and therefore peels the coating away from the substrate. Therefore, we investigate replacing the top Duroid layer with a material that has an even lower CTE and lower tensile modulus. Inspired by the Skybond polyimide foam in PB-2a, we consider the **Vespel SF-0940** polyimide foam from DuPont,⁶⁶ shown in

⁶⁴TOCALO: <https://www.tocalo.co.jp/english/index.html>

⁶⁵CTE, density, and index trend together in all conventional (non-exotic) materials.

⁶⁶DuPont SF-0940: <https://www.dupont.com/knowledge/vespel-sf-polyimide-foam.html>



Parameter	Duroid	Vespel
Index	1.40 ~ 1.42	1.32 ~ 1.47
300 → 77 K contraction [%]	0.67	0.59
Tensile modulus at 300 K [kpsi]	240	76

Figure 10.21: A photograph and performance of the Vespel SF-0940 foam. Figure 10.21a shows a cross-sectional photo of the foam, highlighting how density and thickness are tuned by heat pressing intermediate layers. The left panel of Figure 10.21b shows how density variations in a single 500 mm × 500 mm × 1 mm piece correspond to index variations. The right panel of Figure 10.21b shows the resulting reflectivity variation in PB-2b's 90 and 150 GHz bands. The table shows the refractive index, thermal contraction between 300 and 77 K measured using liquid nitrogen, and tensile modulus at room temperature.

Figure 10.21a. Vespel polyimide is thermally stable,⁶⁷ and SF-0940's open-cell foam structure gives it a low composite density and a low refractive index. Every Vespel SF part is made to order, and the foam's density is tunable via a **heat-induced densification** process. The fabrication techniques are proprietary, but DuPont's manufacturing capabilities are quite

⁶⁷To evidence polyimide's thermal stability, Kapton tape is a polyimide film also from DuPont that is widely used for applications in extreme temperatures.

flexible⁶⁸ to customer-defined specifications. A summary of Vespel’s optical and mechanical properties is shown in the table of Figure 10.21, and its low modulus is particularly appealing for a cryo-mechanically robust AR coating.

DuPont manufactured several 500 mm \times 500 mm \times 1 mm Vespel sheets for both SA and Simons Observatory (SO) in mid-2019. Due to the limitations of the heat-pressing infrastructure, each piece has a positive **radial density gradient**, such that the edges are at a higher density than the center. These density variations modulate the foam’s void density and hence modulate its refractive index. To measure the relationship between density and index, we take one Vespel sheet, cut it into many small pieces, measure each piece’s density,⁶⁹ and measure each piece’s index (at room temperature).⁷⁰ The results of these measurements are shown in the left panel of Figure 10.21b, and the relationship between index and density is well-described by a linear fit. The index is $1.32 \sim 1.47$ and increases away from the sheet’s center. In addition, we measure the index of a single piece at 77 K and find that Vespel’s index is stable with temperature.⁷¹

The primary concern with a position-dependent Vespel index is that the CHWP’s AR coating will then have a position-dependent reflectivity, which may introduce HWPSSs. Because different detectors on the focal plane “see” different regions of the CHWP (as shown in Figure 7.1b), the reflectivity “seen” by each detector depends on the HWP’s angle. To quantify these possible HWPSSs,⁷² the right panel of Figure 10.21b shows reflectivity vs. Vespel density, and the resulting range of reflectivities in the (90, 150) GHz bands is $(0.5, 0.7) \sim (1.0, 1.8)\%$. These HWPSSs are mitigated by the density’s radial gradient but are nonetheless a substantial concern for the epoxy + Vespel AR coating.

Since mid-2019, DuPont has upgraded its heat-press infrastructure to improve the density uniformity over 500 mm diameters,⁷³ and therefore we continue to pursue Vespel foam as a potential CHWP AR material.

⁶⁸Vespel is developed primarily for aerospace applications and is often used as an interface material in turbines, which come in a wide variety of geometries.

⁶⁹As expected for the heat-pressing process, density correlates nearly perfectly with thickness.

⁷⁰Measurements performed by Dick Plambeck and UC Berkeley.

⁷¹Cryogenic index measurements performed by Makoto Hattori at Tohoku University. PB-2a also found Skybond’s index to be stable with temperature [93].

⁷²We have not yet modeled AR-induced HWPSSs in detail, but if Vespel is further considered for mm-wave AR coatings, such investigations will become important.

⁷³This effort is useful not only for two-layer AR coatings but also for three-layer coatings, which need a low-density top-most material. SF-0940 is the highest-density Vespel offering, but Dupont also manufactures SF-0930, which is \sim half the density of that shown in Figure 10.21. Therefore, improving Vespel uniformity has multiple application pathways.

Chapter 11

Research impact

This dissertation advances the research areas of sensitivity forecasting and hardware-based systematic error mitigation for cosmic microwave background (CMB) polarization measurements. These advancements address needs within Simons Array (SA) and Simons Observatory (SO), but their applications are broad, and they are influencing developments across the field. In this chapter, we briefly discuss the research impacts of the BoloCalc sensitivity calculator and SA’s half-wave plate (HWP) polarization modulators.

11.1 BoloCalc

We have written a unified, broadly applicable noise-equivalent temperature (NET) simulation code called BoloCalc [81], whose comprehensive feature set has helped quantify a wide variety of optimization questions. For example, SA, SO, and LiteBIRD have used BoloCalc to quantify requirements on pixel packing, optic temperatures, manufacturing tolerances, and validation metrics. Chapter 5 demonstrates how BoloCalc can inform high-level telescope design, such as the arrangement of optics tubes in SO’s large-aperture telescope receiver (LATR), as well as low-level fabrication targets and tolerances, such as the saturation power for SA’s detectors. BoloCalc’s ability to assemble arbitrary combinations of telescopes, optics tubes, and detector arrays, to import measured histograms and spectra, to sweep over arbitrary combinations of input parameters, and to generate layered Monte Carlo simulations provides value at every stage of a CMB project, from conceptualization to in-field characterization.

In addition, BoloCalc is the first sensitivity calculator, to our knowledge, that includes the impact of Hanbury Brown-Twiss (HBT) correlations on the array-averaged NET of densely packed detector arrays. Chapter 4 explores the implications of these photon noise correlations for observatory design, and we have studied their effects both on a generic telescope system and on the performance of SO’s focal planes. The importance of accurate HBT correlation calculations will grow as larger and denser detector arrays deploy on future experiments, such as CMB-S4.




Figure 11.1: A standalone photo of the SO CHWP (similar to Figure 8.4b), taken at LBNL during laboratory evaluation. The sapphire, 50 K shield, and retaining rings are removed for visual clarity. Many aspects of the SO CHWP design are functionally identical to those of the PB-2b CHWP, including a superconducting magnetic bearing, a magnetic motor, an optical encoder, and a three-pronged gripper with ambient actuators. This photo is courtesy of Peter Ashton.

BoloCalc is already being used across the CMB community, not only within SO and SA but also for LiteBIRD detector modeling [131], CMB-S4 band selection [205], and independent cosmology forecasting studies [112]. We anticipate that BoloCalc will become increasingly prolific over the coming decade as CMB projects conglomerate and need unified simulation tools. In addition, BoloCalc’s modular, object-oriented structure makes it amenable to feature addition and functionality expansion in the future. Finally, BoloCalc’s graphical user interface (GUI) has boosted its accessibility, and as the calculator continues to mature, we foresee a broad, diverse userbase for years to come.

11.2 HWPs

We have designed, built, and validated two HWP polarization modulators for SA: a warm HWP (WHWP) for POLARBEAR-2a (PB-2a) and a cryogenic HWP (CHWP) for POLARBEAR-2b. The PB-2a WHWP (Chapter 7) is the largest-diameter HWP ever deployed for CMB observation and introduces innovative techniques for sapphire characterization, dichroic anti-reflection (AR) coatings, and HWP laboratory validation. The PB-2b CHWP (Chapters 8-10) advances the research area of large-diameter cryogenic polarization modulators in several areas, including low-dissipation rotation, cryo-mechanical robustness, high-precision angle encoding, and experiment-integrated performance. While the PB-2b CHWP’s design is similar in spirit to that of EBEX [105], it is distinctively the first of its kind and improves the feasibility and lucrativeness of polarization modulation for CMB observation.

The PB-2b CHWP design is already making its way into other experiments. PB-2c, the third installment of SA, is building an identical CHWP for observation at 220 and 270 GHz, while SO has adopted the presented system, with some experiment-dependent modifications,




Figure 11.2: CAD of LiteBIRD’s polarization modulator unit (PMU) for its CHWP system, taken from Sakurai *et al.* (2018) [169].

for cryogenic polarization modulation on their small aperture telescopes (SATs) [59, 4]. Figure 11.1 shows an SO CHWP for the first SAT, which will observe at 90 and 150 GHz. SO’s modulator has an even larger diameter than PB-2b’s and therefore has a different set of dimensional constraints. However, its core systems—including the motor and driver, gripper, angle encoder, and sapphire—are effectively identical to the CHWP described in this dissertation.¹ The first SO CHWP was built at LBNL, and a second SO CHWP is under construction (as of December 2020) at the University of Tokyo.² In addition, a similar CHWP is being developed for LiteBIRD [169]—a Japanese CMB satellite mission [131]—and computer-aided-design (CAD) images of LiteBIRD’s modulator are shown in Figure 11.2. Several design aspects are shared between the LiteBIRD and PB-2b modulator systems, and their simultaneous development has resulted in a fruitful collaboration between the R&D teams. That said, the LiteBIRD CHWP faces additional challenges and stricter requirements, as the stakes are higher for a satellite experiment than for a ground-based telescope. Nonetheless, PB-2b’s CHWP serves as a pathfinder for LiteBIRD’s polarization modulation efforts.

The PB-2b CHWP has also been designed with portability in mind so that its hardware may be seamlessly adopted for non-SA telescopes. For example, the motor driver is easy to implement and includes universal built-in features like noise suppression, ground-isolation, solenoid-signal cleaning, and simple operation. These design choices will ease the adaptation of PB-2b’s CHWP system into other experiments, and their utility has in fact helped accelerate SO’s CHWP development. Furthermore, once the scientific power of the PB-2b CHWP is solidified, we foresee variants of the presented system making their way into a wider range of telescopes within mm-wave and sub-mm astronomy.

¹SO uses a different manufacturer for its superconducting magnetic bearing and achieves an inner diameter of ≈ 500 mm, larger than SA’s 470 mm [167].

²Kusaka Lab at the University of Tokyo: <http://www.cmb.phys.s.u-tokyo.ac.jp/>

The major outstanding item for the PB-2b CHWP development is its sapphire anti-reflection (AR) coating, described in Chapter 10. Despite the epoxy + Duroid coating’s lucrative (warm) performance and an advancing understanding of its cryo-mechanical properties, we have not yet demonstrated a robust sapphire AR solution. We are pursuing several technologies in parallel, including an improved epoxy + Duroid application process and a mullite + Duroid hybrid. Given the non-overlapping nature of each technology’s strengths and weaknesses, we anticipate one or more of these options to converge in time for CHWP operation on PB-2b.

11.3 Conclusion

The search for primordial B-modes is at the center of cutting-edge CMB science, and this dissertation advances the efforts of SA and SO to discover inflation through improved sensitivity and systematic error mitigation. The presented research is only a fragment of a much greater effort, and exciting instrumentation developments throughout the CMB community indicate a bright future for mm-wave telescope technology. So while this dissertation’s work comes to a close, the hunt for inflation is far from over.

Bibliography

- [1] A. Aboobaker *et al.* “The EBEX balloon-borne experiment - gondola, attitude control, and control software.” *The Astrophysical Journal Supplement Series* 239.1 (Feb. 2017), p. 9. DOI: 10.3847/1538-4365/aae435.
- [2] P. A. R. Ade, G. Pisano, C. Tucker, and S. Weaver. “A review of metal mesh filters.” *Millimeter and Submillimeter Detectors and Instrumentation for Astronomy III*. Vol. 6275. SPIE, June 2006, 62750U. DOI: 10.1117/12.673162.
- [3] G. S. Agarwal, G. Gbur, and E. Wolf. “Coherence properties of sunlight.” *Optics Letters* 29.5 (Mar. 2004), pp. 459–461. DOI: 10.1364/OL.29.000459.
- [4] A. M. Ali and The Simons Observatory Collaboration. “Small Aperture Telescopes for the Simons Observatory.” *Journal of Low Temperature Physics* 2020 (Jan. 2020), pp. 1–11. DOI: 10.1007/s10909-020-02430-5.
- [5] K. S. Arnold. “Design and Deployment of the POLARBEAR Cosmic Microwave Background Polarization Experiment.” PhD thesis. UC Berkeley, 2010. URL: <https://escholarship.org/uc/item/99m8b32x>.
- [6] P. C. Ashton. “Interstellar Dust Polarimetry with BLASTPol and BLAST-TNG.” PhD thesis. Northwestern University, 2017. URL: <https://arch.library.northwestern.edu/downloads/qb98mf67m?locale=en>.
- [7] J. Aumont, J. F. Macías-Pérez, A. Ritacco, N. Ponthieu, and A. Mangilli. “Absolute calibration of the polarisation angle for future CMB B -mode experiments from current and future measurements of the Crab nebula.” *Astronomy & Astrophysics* 634 (Feb. 2020), A100. DOI: 10.1051/0004-6361/201833504.
- [8] N. A. Bahcall. “Hubble Law and the expanding universe.” *Proceedings of the National Academy of Sciences* 112.11 (Mar. 2015), pp. 3173–3175. DOI: 10.1073/pnas.1424299112.
- [9] H. P. Baltes, B. Steinle, and G. Antes. “Spectral coherence and the radiant intensity from statistically homogeneous and isotropic planar sources.” *Optics Communications* 18.3 (Aug. 1976), pp. 242–246. DOI: 10.1016/0030-4018(76)90122-X.

- [10] C. Bao, C. Baccigalupi, B. Gold, S. Hanany, A. Jaffe, and R. Stompor. “Maximum likelihood foreground cleaning for cosmic microwave background polarimeters in the presence of systematic effects.” *The Astrophysical Journal* 819.1 (Feb. 2016), p. 12. DOI: 10.3847/0004-637x/819/1/12.
- [11] C. Bao, B. Gold, C. Baccigalupi, J. Didier, S. Hanany, A. Jaffe, B. R. Johnson, S. Leach, T. Matsumura, A. Miller, and D. O’Dea. “The impact of the spectral response of an achromatic half-wave plate on the measurement of the cosmic microwave background polarization.” *Astrophysical Journal* 747.2 (Mar. 2012), p. 97. DOI: 10.1088/0004-637X/747/2/97.
- [12] C. P. Bean. “Magnetization of High-Field Superconductors.” *Reviews of Modern Physics* 36.1 (Jan. 1964), pp. 31–39. DOI: 10.1103/RevModPhys.36.31.
- [13] A. N. Bender and The South Pole Telescope Collaboration. “On-Sky Performance of the SPT-3G Frequency-Domain Multiplexed Readout.” *Journal of Low Temperature Physics* 199.1-2 (Apr. 2020), pp. 182–191. DOI: 10.1007/s10909-019-02280-w.
- [14] A. N. Bender, J.-F. Cliche, T. de Haan, *et al.* “Digital frequency domain multiplexing readout electronics for the next generation of millimeter telescopes.” *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII*. Vol. 9153. SPIE, July 2014, 91531A. DOI: 10.1117/12.2054949.
- [15] D. J. Benford, T. R. Hunter, and T. G. Phillips. “Noise Equivalent Power of Background Limited Thermal Detectors at Submillimeter Wavelengths.” *International Journal of Infrared and Millimeter Waves* 19.7 (July 1998), pp. 931–938. DOI: 10.1023/A:1022671223858.
- [16] C. L. Bennett, D. Larson, J. L. Weiland, *et al.* “Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results.” *The Astrophysical Journal Supplement Series* 208.2 (Sept. 2013), p. 20. DOI: 10.1088/0067-0049/208/2/20.
- [17] F. Beutler, C. Blake, M. Colless, D. H. Jones, L. Staveley-Smith, L. Campbell, Q. Parker, W. Saunders, and F. Watson. “The 6dF Galaxy Survey: Baryon acoustic oscillations and the local Hubble constant.” *Monthly Notices of the Royal Astronomical Society* 416.4 (Oct. 2011), pp. 3017–3032. DOI: 10.1111/j.1365-2966.2011.19250.x.
- [18] R. S. Bhatia, S. T. Chase, S. F. Edgington, J. Glenn, W. C. Jones, A. E. Lange, B. Maffei, A. K. Mainzer, P. D. Mauskopf, B. J. Philhour, and B. K. Rownd. “A three-stage helium sorption refrigerator for cooling of infrared detectors to 280 mK.” *Cryogenics* 40.11 (Jan. 2000), pp. 685–691. DOI: 10.1016/S0011-2275(00)00072-2.
- [19] C. Bischoff, L. Hyatt, J. J. McMahon, G. W. Nixon, D. Samtleben, K. M. Smith, K. Vanderlinde, D. Barkats, P. Farese, T. Gaier, J. O. Gundersen, M. M. Hedman, S. T. Staggs, and B. Winstein. “New Measurements of Fine-Scale CMB Polarization Power Spectra from CAPMAP at Both 40 and 90 GHz.” *The Astrophysical Journal* 684.2 (Sept. 2008), pp. 771–789. DOI: 10.1086/590487.

- [20] D. C.-J. Bock, A. D. Bolatto, D. W. Hawkins, A. J. Kemball, J. W. Lamb, R. L. Plambeck, M. W. Pound, S. L. Scott, D. P. Woody, and M. C. H. Wright. “First results from CARMA: the combined array for research in millimeter-wave astronomy.” *Ground-based and Airborne Telescopes*. Vol. 6267. SPIE, June 2006, p. 626713. DOI: 10.1117/12.674051.
- [21] N. W. Boggess, J. C. Mather, R. Weiss, *et al.* “The COBE mission - Its design and performance two years after launch.” *The Astrophysical Journal* 397 (Oct. 1992), p. 420. DOI: 10.1086/171797.
- [22] V. B. Braginsky, V. S. Ilchenko, and K. S. Bagdassarov. “Experimental observation of fundamental microwave absorption in high-quality dielectric crystals.” *Physics Letters A* 120.6 (Mar. 1987), pp. 300–305. DOI: 10.1016/0375-9601(87)90676-1.
- [23] R. H. Brown, R. C. Jennison, and M. K. D. Gupta. “Apparent Angular Sizes of Discrete Radio Sources: Observations at Jodrell Bank, Manchester.” *Nature* 170.4338 (Dec. 1952), pp. 1061–1063. DOI: 10.1038/1701061a0.
- [24] R. H. Brown and R. Q. Twiss. “Correlation between Photons in two Coherent Beams of Light.” *Nature* 177.4497 (Jan. 1956), pp. 27–29. DOI: 10.1038/177027a0.
- [25] R. H. Brown and R. Q. Twiss. “LXXIV. A new type of interferometer for use in radio astronomy.” *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 45.366 (July 1954), pp. 663–682. DOI: 10.1080/14786440708520475.
- [26] S. Bryan *et al.* “Optical design of the TolTEC millimeter-wave camera.” *Millimeter, Submillimeter, and Far Infrared Detectors and Instrumentation for Astronomy IX*. Vol. 10708. International Society for Optics and Photonics, July 2018, 107080J. DOI: 10.1117/12.2314130.
- [27] S. Bryan and The SPIDER Collaboration. “A cryogenic rotation stage with a large clear aperture for the half-wave plates in the Spider instrument.” *Review of Scientific Instruments* 87.1 (Jan. 2016), p. 14501. DOI: 10.1063/1.4939435.
- [28] W. H. Carter and E. Wolf. “Coherence and radiometry with quasihomogeneous planar sources.” *JOSA* 67.6 (June 1977), pp. 785–796. DOI: 10.1364/JOSA.67.000785.
- [29] W. H. Carter and E. Wolf. “Coherence properties of lambertian and non-lambertian sources.” *JOSA* 65.9 (Sept. 1975), pp. 1067–1071. DOI: 10.1364/JOSA.65.001067.
- [30] R. Cerbino. “Correlations of light in the deep Fresnel region: An extended Van Cittert and Zernike theorem.” *Physical Review A* 75.5 (May 2007), p. 53815. DOI: 10.1103/PhysRevA.75.053815.
- [31] J. Choi *et al.* “Radio-transparent multi-layer insulation for radiowave receivers.” *Review of Scientific Instruments* 84.11 (Nov. 2013), p. 114502. DOI: 10.1063/1.4827081.
- [32] S. K. Choi *et al.* “Characterization of the Mid-Frequency Arrays for Advanced ACT-Pol.” *Journal of Low Temperature Physics* 193.3-4 (Nov. 2018), pp. 267–275. DOI: 10.1007/s10909-018-1982-4.

- [33] D. T. Chuss, J. R. Eimer, D. J. Fixsen, J. Hinderks, A. J. Kogut, J. Lazear, P. Mirel, E. Switzer, G. M. Voellmer, and E. J. Wollack. “Variable-delay polarization modulators for cryogenic millimeter-wave applications.” *Review of Scientific Instruments* 85.6 (June 2014), p. 064501. DOI: 10.1063/1.4879499.
- [34] D. T. Chuss, E. J. Wollack, S. H. Moseley, S. Withington, and G. Saklatvala. “Diffraction Considerations for Planar Detectors in the Few-Mode Limit.” *Publications of the Astronomical Society of the Pacific* 120.866 (Apr. 2008), pp. 430–438. DOI: 10.1086/587494.
- [35] P. H. van Cittert. “Die Wahrscheinliche Schwingungsverteilung in Einer von Einer Lichtquelle Direkt Oder Mittels Einer Linse Beleuchteten Ebene.” *Physica* 1.1-6 (Jan. 1934), pp. 201–210. DOI: 10.1016/S0031-8914(34)90026-4.
- [36] K. P. Coughlin, J. J. McMahon, *et al.* “Pushing the Limits of Broadband and High-Frequency Metamaterial Silicon Antireflection Coatings.” *Journal of Low Temperature Physics* 193.5-6 (Dec. 2018), pp. 876–885. DOI: 10.1007/s10909-018-1955-7.
- [37] K. Coughlin. “Metamaterial Optics for Precision Cosmic Microwave Background Observation.” PhD thesis. University of Michigan, 2018. URL: <https://deepblue.lib.umich.edu/handle/2027.42/146072>.
- [38] K. T. Crowley *et al.* “Advanced ACTPol TES Device Parameters and Noise Performance in Fielded Arrays.” *Journal of Low Temperature Physics* 193.3-4 (Nov. 2018), pp. 328–336. DOI: 10.1007/s10909-018-2028-7.
- [39] A. Cukierman, A. T. Lee, C. Raum, A. Suzuki, and B. Westbrook. “Hierarchical sinuous-antenna phased array for millimeter wavelengths.” *Applied Physics Letters* 112.13 (Mar. 2018), p. 132601. DOI: 10.1063/1.5021962.
- [40] S. Dahal *et al.* “The CLASS 150/220 GHz Polarimeter Array: Design, Assembly, and Characterization.” *Journal of Low Temperature Physics* 199.1-2 (Apr. 2020), pp. 289–297. DOI: 10.1007/s10909-019-02317-0.
- [41] R. Datta, J. Hubmayr, C. Munson, *et al.* “Horn coupled multichroic polarimeters for the atacama cosmology telescope polarization experiment.” *Journal of Low Temperature Physics* 176.5-6 (Feb. 2014), pp. 670–676. DOI: 10.1007/s10909-014-1134-4.
- [42] R. Datta, C. D. Munson, M. D. Niemack, J. J. McMahon, *et al.* “Large-aperture wide-bandwidth antireflection-coated silicon lenses for millimeter wavelengths.” *Applied Optics* 52.36 (Dec. 2013), p. 8747. DOI: 10.1364/ao.52.008747.
- [43] S. W. Deiker, W. Doriese, G. C. Hilton, *et al.* “Superconducting transition edge sensor using dilute AlMn alloys.” *Applied Physics Letters* 85.11 (Sept. 2004), pp. 2137–2139. DOI: 10.1063/1.1789575.
- [44] S. Dicker, P. Gallardo, J. Gudmudsson, P. Mauskopf, *et al.* “Cold optical design for the large aperture Simons’ Observatory telescope.” *Proceedings of SPIE - The International Society for Optical Engineering*. Vol. 10700. 2018. DOI: 10.1117/12.2313444.

- [45] M. A. Dobbs, M. Lueker, *et al.* “Frequency multiplexed superconducting quantum interference device readout of large bolometer arrays for cosmic microwave background measurements.” *Review of Scientific Instruments* 83.7 (July 2012), p. 73113. DOI: 10.1063/1.4737629.
- [46] B. Dober *et al.* “Microwave SQUID multiplexer demonstration for cosmic microwave background imagers.” *Applied Physics Letters* 111.24 (Dec. 2017), p. 243510. DOI: 10.1063/1.5008527.
- [47] C. Dragone. “Offset Multireflector Antennas with Perfect Pattern Symmetry and Polarization Discrimination.” *Bell System Technical Journal* 57.7 (1978), pp. 2663–2684. DOI: 10.1002/j.1538-7305.1978.tb02171.x.
- [48] S. M. Duff *et al.* “Advanced ACTPol Multichroic Polarimeter Array Fabrication Process for 150 mm Wafers.” *Journal of Low Temperature Physics* 184.3-4 (Aug. 2016), pp. 634–641. DOI: 10.1007/s10909-016-1576-y.
- [49] J. M. Edwards, R. O’Brien, A. T. Lee, and G. M. Rebeiz. “Dual-polarized sinuous antennas on extended hemispherical silicon lenses.” *IEEE Transactions on Antennas and Propagation* 60.9 (2012), pp. 4082–4091. DOI: 10.1109/TAP.2012.2207048.
- [50] J. Errard and The POLARBEAR Collaboration. “Modeling Atmospheric Emission For CMB Ground-based Observations.” *The Astrophysical Journal* 809.1 (Aug. 2015), p. 63. DOI: 10.1088/0004-637X/809/1/63.
- [51] T. Essinger-Hileman *et al.* “The Atacama B-mode Search: CMB polarimetry with transition-edge-sensor bolometers.” *AIP Conference Proceedings*. Vol. 1185. 1. American Institute of Physics AIP, Dec. 2009, pp. 494–497. DOI: 10.1063/1.3292387.
- [52] T. Essinger-Hileman. “Transfer matrix for treating stratified media including birefringent crystals.” *Applied Optics* 52.2 (Jan. 2013), p. 212. DOI: 10.1364/AO.52.000212.
- [53] T. Essinger-Hileman, A. Kusaka, *et al.* “Systematic effects from an ambient-temperature, continuously rotating half-wave plate.” *Review of Scientific Instruments* 87.9 (Sept. 2016), p. 94503. DOI: 10.1063/1.4962023.
- [54] U. Fano. “Quantum Theory of Interference Effects in the Mixing of Light from Phase-Independent Sources.” *American Journal of Physics* 29.8 (Aug. 1961), pp. 539–545. DOI: 10.1119/1.1937827.
- [55] D. J. Fixsen. “The temperature of the cosmic microwave background.” *Astrophysical Journal* 707.2 (Nov. 2009), pp. 916–920. DOI: 10.1088/0004-637X/707/2/916.
- [56] D. J. Fixsen, E. S. Cheng, J. M. Gales, J. C. Mather, R. A. Shafer, and E. L. Wright. “The Cosmic Microwave Background Spectrum from the Full COBE FIRAS Data Set.” *The Astrophysical Journal* 473.2 (Dec. 1996), pp. 576–587. DOI: 10.1086/178173.
- [57] A. T. Friberg. “Radiation From Partially Coherent Sources.” *Optical Engineering* 21.2 (Apr. 1982), p. 212362. DOI: 10.1117/12.7972911.

- [58] M. Galeazzi. “Fundamental Noise Processes in TES Devices.” *IEEE Transactions on Applied Superconductivity* 21.3 (June 2011), pp. 267–271. DOI: 10.1109/TASC.2010.2091243.
- [59] N. Galitzki and The Simons Observatory Collaboration. “The Simons Observatory: Instrument overview.” *Proceedings of SPIE - The International Society for Optical Engineering*. Vol. 10708. 2018. DOI: 10.1117/12.2312985.
- [60] P. Gallardo, J. Gudmundsson, B. Koopman, F. Matsuda, S. Simon, *et al.* “Studies of Systematic uncertainties in the Simons Observatory: Optical effects and sensitivity considerations.” *Proceedings of SPIE - The International Society for Optical Engineering*. Vol. 10708. 2018. DOI: 10.1117/12.2312971.
- [61] R. J. Glauber. “Coherent and Incoherent States of the Radiation Field.” *Physical Review* 131.6 (Sept. 1963), pp. 2766–2788. DOI: 10.1103/PhysRev.131.2766.
- [62] R. J. Glauber. “Photon Correlations.” *Physical Review Letters* 10.3 (Feb. 1963), pp. 84–86. DOI: 10.1103/PhysRevLett.10.84.
- [63] F. Gori. “Far-zone approximation for partially coherent sources.” *Optics Letters* 30.21 (Nov. 2005), pp. 2840–2842. DOI: 10.1364/OL.30.002840.
- [64] M. J. Griffin, J. J. Bock, and W. K. Gear. “The Relative Performance of Filled and Feedhorn-Coupled Focal-plane Architectures.” *Applied Optics* 41.31 (Nov. 2002), p. 6543. DOI: 10.1364/AO.41.006543.
- [65] R. Gualtieri, J. P. Filippini, *et al.* “SPIDER: CMB Polarimetry from the Edge of Space.” *Journal of Low Temperature Physics* 193.5-6 (Dec. 2018), pp. 1112–1121. DOI: 10.1007/s10909-018-2078-x.
- [66] A. H. Guth. “Inflationary universe: A possible solution to the horizon and flatness problems.” *Physical Review D* 23 (Jan. 1981), pp. 347–356. DOI: 10.1103/PhysRevD.23.347.
- [67] T. de Haan, G. Smecher, and M. Dobbs. “Improved performance of TES bolometers using digital feedback.” *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VI*. Vol. 8452. Sept. 2012, 84520E. DOI: 10.1117/12.925658.
- [68] E. Hagen and H. Rubens. “Ueber Beziehungen des Reflexions- und Emissionsvermoegens der Metalle zu ihrem elektrischen Leitvermoegen.” *Annalen der Physik* 316.8 (Jan. 1903), pp. 873–901. DOI: 10.1002/andp.19033160811.
- [69] M. Halpern, H. P. Gush, E. Wishnow, and V. D. Cosmo. “Far infrared transmission of dielectrics at cryogenic and room temperatures: glass, Fluorogold, Eccosorb, Stycast, and various plastics.” *Applied Optics* 25.4 (Feb. 1986), pp. 565–570. DOI: 10.1364/AO.25.000565.

- [70] J.-C. Hamilton, R. Charlassier, C. Cressiot, J. Kaplan, M. Piat, and C. Rosset. “Sensitivity of a bolometric interferometer to the cosmic microwave background power spectrum.” *Astronomy & Astrophysics* 491.3 (Dec. 2008), pp. 923–927. DOI: 10.1051/0004-6361:200810504.
- [71] S. Hanany and P. Rosenkranz. “Polarization of the atmosphere as a foreground for cosmic microwave background polarization experiments.” *New Astronomy Reviews* 47.11-12 (Dec. 2003), pp. 1159–1165. DOI: 10.1016/j.newar.2003.09.017.
- [72] D. A. Harper, M. C. Runyan, C. D. Dowell, C. J. Wirth, *et al.* “HAWC+, the Far-Infrared Camera and Polarimeter for SOFIA.” *Journal of Astronomical Instrumentation* 07.04 (July 2018), p. 1840008. DOI: 10.1142/S2251171718400081.
- [73] K. Harrington *et al.* “Variable-delay polarization modulators for the CLASS telescopes.” *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*. Vol. 10708. SPIE, July 2018, p. 92. DOI: 10.1117/12.2313614.
- [74] K. Harrington, T. Marriage, *et al.* “The Cosmology Large Angular Scale Surveyor.” *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VIII*. Vol. 9914. SPIE, July 2016, 99141K. DOI: 10.1117/12.2233125.
- [75] K. Hattori *et al.* “Development of Readout Electronics for POLARBEAR-2 Cosmic Microwave Background Experiment.” *Journal of Low Temperature Physics* 184.1-2 (July 2016), pp. 512–518. DOI: 10.1007/s10909-015-1448-x.
- [76] S. W. Henderson *et al.* “Highly-multiplexed microwave SQUID readout using the SLAC Microresonator Radio Frequency (SMuRF) electronics for future CMB and sub-millimeter surveys.” *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*. Vol. 10708. SPIE, July 2018, p. 43. DOI: 10.1117/12.2314435.
- [77] J. W. Henning *et al.* “Feedhorn-coupled TES polarimeter camera modules at 150 GHz for CMB polarization measurements with SPTpol.” *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VI*. Vol. 8452. 5. SPIE, Oct. 2012, 84523A. DOI: 10.1117/12.927172.
- [78] C. A. Hill and A. Kusaka. “Photon noise correlations in millimeter-wave cameras.” *In preparation* (2020).
- [79] C. A. Hill, A. Kusaka, P. Ashton, P. Barton, *et al.* “A cryogenic continuously rotating half-wave plate mechanism for the POLARBEAR-2b cosmic microwave background receiver.” *Review of Scientific Instruments* 91.12 (Dec. 2020), p. 124503. DOI: 10.1063/5.0029006.
- [80] C. A. Hill, A. Kusaka, P. Barton, *et al.* “A Large-Diameter Cryogenic Rotation Stage for Half-Wave Plate Polarization Modulation on the POLARBEAR-2 Experiment.” *Journal of Low Temperature Physics* 193.5-6 (2018). DOI: 10.1007/s10909-018-1980-6.

- [81] C. A. Hill, S. M. Bruno, S. Simon, *et al.* “BoloCalc: A sensitivity calculator for the design of Simons Observatory.” *Proceedings of SPIE - The International Society for Optical Engineering*. Vol. 10708. 2018. DOI: 10.1117/12.2313916.
- [82] C. A. Hill, S. Beckman, *et al.* “Design and development of an ambient-temperature continuously-rotating achromatic half-wave plate for CMB polarization modulation on the POLARBEAR-2 experiment.” *Proceedings of SPIE - The International Society for Optical Engineering*. Vol. 9914. 2016. DOI: 10.1117/12.2232280.
- [83] H. S. Hou. “Method for Optimized Design of Dielectric Multilayer Filters.” *Applied Optics* 13.8 (Aug. 1974), p. 1863. DOI: 10.1364/ao.13.001863.
- [84] L. Howe. “The POLARBEAR-2 Cryogenic Receiver for Cosmic Microwave Background Polarization Science.” PhD thesis. UC San Diego, 2019. URL: <https://escholarship.org/uc/item/5897v3wz>.
- [85] L. Howe, C. Tsai, L. Lowry, *et al.* “Design and characterization of the POLARBEAR-2b and POLARBEAR-2c cosmic microwave background cryogenic receivers.” *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*. Vol. 10708. July 2018, 107083W. DOI: 10.1117/12.2312821.
- [86] W. Hu and T. Okamoto. “Mass Reconstruction with CMB Polarization.” *The Astrophysical Journal* 574.2 (Aug. 2002), pp. 566–574. DOI: 10.1086/341110.
- [87] W. Hu and M. White. “A CMB polarization primer.” *New Astronomy* 2.4 (Oct. 1997), pp. 323–344. DOI: 10.1016/S1384-1076(97)00022-5.
- [88] E. Hubble. “A relation between distance and radial velocity among extra-galactic nebulae.” *Proceedings of the National Academy of Sciences* 15.3 (Mar. 1929), pp. 168–173. DOI: 10.1073/pnas.15.3.168.
- [89] H. Hui *et al.* “BICEP3 focal plane design and detector performance.” *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VIII*. Vol. 9914. 27. SPIE, July 2016, 99140T. DOI: 10.1117/12.2232986.
- [90] H. Hui *et al.* “BICEP Array: a multi-frequency degree-scale CMB polarimeter.” *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*. Vol. 10708. 9. SPIE, July 2018, p. 49. DOI: 10.1117/12.2311725.
- [91] J. R. Hull. “Effect of permanent-magnet irregularities in levitation force measurements.” *Superconductor Science and Technology* 13.6 (May 2000), pp. 854–856. DOI: 10.1088/0953-2048/13/6/347.
- [92] Y. Inoue *et al.* “POLARBEAR-2: An instrument for CMB polarization measurements.” *Proceedings of SPIE - The International Society for Optical Engineering*. Vol. 9914. 2016. DOI: 10.1117/12.2231961.
- [93] Y. Inoue, T. Matsumura, *et al.* “Cryogenic infrared filter made of alumina for use at millimeter wavelength.” *Applied Optics* 53.9 (Mar. 2014), pp. 1727–1733. DOI: 10.1364/AO.53.001727.

- [94] Y. Inoue, N. Stebor, *et al.* “Thermal and optical characterization for POLARBEAR-2 optical system.” *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII*. Vol. 9153. SPIE, Aug. 2014, 91533A. DOI: 10.1117/12.2055572.
- [95] H. Ishino *et al.* “LiteBIRD: Lite satellite for the study of B-mode polarization and inflation from cosmic microwave background radiation detection.” *Proceedings of SPIE - The International Society for Optical Engineering*. Vol. 9904. 2016. DOI: 10.1117/12.2231995.
- [96] N. Jarosik *et al.* “Design, Implementation, and Testing of the Microwave Anisotropy Probe Radiometers.” *The Astrophysical Journal Supplement Series* 145.2 (Apr. 2003), pp. 413–436. DOI: 10.1086/346080.
- [97] O. Jeong, A. Lee, C. Raum, and A. Suzuki. “Broadband Plasma-Sprayed Anti-reflection Coating for Millimeter-Wave Astrophysics Experiments.” *Journal of Low Temperature Physics* 184.3-4 (Aug. 2016), pp. 621–626. DOI: 10.1007/s10909-015-1442-3.
- [98] B. R. Johnson, F. Columbro, D. Araujo, *et al.* “A large-diameter hollow-shaft cryogenic motor based on a superconducting magnetic bearing for millimeter-wave polarimetry.” *Review of Scientific Instruments* 88.10 (Oct. 2017), p. 105102. DOI: 10.1063/1.4990884.
- [99] B. R. Johnson, J. Collins, *et al.* “MAXIPOL: Cosmic Microwave Background Polarimetry Using a Rotating Half-Wave Plate.” *The Astrophysical Journal* 665.1 (Aug. 2007), p. 42. DOI: 10.1086/518105.
- [100] D. Kaneko and The POLARBEAR Collaboration. “Deployment of Polarbear-2A.” *Journal of Low Temperature Physics* 199.3-4 (2020). DOI: 10.1007/s10909-020-02366-w.
- [101] B. G. Keating, M. Shimon, and A. P. Yadav. “Self-calibration of cosmic microwave background polarization experiments.” *Astrophysical Journal Letters* 762.2 (Jan. 2013), p. 23. DOI: 10.1088/2041-8205/762/2/L23.
- [102] Z. D. Kermish. “The POLARBEAR Experiment: Design and Characterization.” PhD thesis. UC Berkeley, 2012. URL: <https://escholarship.org/uc/item/6wk67864>.
- [103] S. A. Kernasovskiy, S. E. Kuenstner, E. Karpel, Z. Ahmed, D. D. Van Winkle, S. Smith, J. Dusatko, J. C. Frisch, *et al.* “SLAC Microresonator Radio Frequency (SMuRF) Electronics for Read Out of Frequency-Division-Multiplexed Cryogenic Sensors.” *Journal of Low Temperature Physics* 193.3-4 (Nov. 2018), pp. 570–577. DOI: 10.1007/s10909-018-1981-5.
- [104] M. Kesden, A. Cooray, and M. Kamionkowski. “Separation of Gravitational-Wave and Cosmic-Shear Contributions to Cosmic Microwave Background Polarization.” *Physical Review Letters* 89.1 (June 2002), p. 11304. DOI: 10.1103/PhysRevLett.89.011304.

- [105] J. Klein and The EBEX Collaboration. “A cryogenic half-wave plate polarimeter using a superconducting magnetic bearing.” *Cryogenic Optical Systems and Instruments XIII*. Vol. 8150. Sept. 2011, p. 815004. DOI: 10.1117/12.893669.
- [106] L. Knox. “Cosmic Microwave Background Anisotropy Observing Strategy Assessment.” *The Astrophysical Journal* 480.1 (May 1997), pp. 72–78. DOI: 10.1086/303959.
- [107] L. Knox and Y.-S. Song. “Limit on the Detectability of the Energy Scale of Inflation.” *Physical Review Letters* 89.1 (June 2002), p. 11303. DOI: 10.1103/PhysRevLett.89.011303.
- [108] K. Komatsu, T. Matsumura, *et al.* “Demonstration of the broadband half-wave plate using the nine-layer sapphire for the cosmic microwave background polarization experiment.” *Journal of Astronomical Telescopes, Instruments, and Systems* 5.4 (Nov. 2019), p. 44008. DOI: 10.1117/1.JATIS.5.4.044008.
- [109] B. J. Koopman *et al.* “Advanced ACTPol Low-Frequency Array: Readout and Characterization of Prototype 27 and 39 GHz Transition Edge Sensors.” *Journal of Low Temperature Physics* (May 2018), pp. 1–9. DOI: 10.1007/s10909-018-1957-5.
- [110] B. J. Koopman *et al.* “Optical modeling and polarization calibration for CMB measurements with ACTPol and Advanced ACTPol.” *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VIII*. Vol. 9914. SPIE, July 2016, 99142T. DOI: 10.1117/12.2231912.
- [111] A. Kusaka, T. Essinger-Hileman, *et al.* “Modulation of cosmic microwave background polarization with a warm rapidly rotating half-wave plate on the Atacama B-Mode Search instrument.” *Review of Scientific Instruments* 85.2 (Feb. 2014), p. 24501. DOI: 10.1063/1.4862058.
- [112] P. La Plante, A. Lidz, J. Aguirre, and S. Kohn. “The 21 cm kSZ–kSZ Bispectrum during the Epoch of Reionization.” *The Astrophysical Journal* 899.1 (Aug. 2020), p. 40. DOI: 10.3847/1538-4357/aba2ed.
- [113] J. M. Lamarre. “Photon noise in photometric instruments at far-infrared and submillimeter wavelengths.” *Applied Optics* 25.6 (Mar. 1986), pp. 870–876. DOI: 10.1364/AO.25.000870.
- [114] J. W. Lamb. “Miscellaneous data on materials for millimetre and submillimetre optics.” *International Journal of Infrared and Millimeter Waves* 17.12 (1996), pp. 1997–2034. DOI: 10.1007/BF02069487.
- [115] J. C. Leader. “Far-zone range criteria for quasi-homogeneous partially coherent sources.” *JOSA* 68.10 (Oct. 1978), pp. 1332–1338. DOI: 10.1364/JOSA.68.001332.
- [116] A. T. Lee, J. M. Gildemeister, S.-F. Lee, and P. L. Richards. “Voltage-biased high-T/sub c/ superconducting infrared bolometers with strong electrothermal feedback.” *IEEE Transactions on Applied Superconductivity* 7.2 (June 1997), pp. 2378–2381. DOI: 10.1109/77.621718.

- [117] K. Lee *et al.* “GroundBIRD: A CMB Polarization Experiment with MKID Arrays.” *Journal of Low Temperature Physics* 200.5-6 (Sept. 2020), pp. 384–391. DOI: 10.1007/s10909-020-02511-5.
- [118] A. Lewis and A. Challinor. “Weak gravitational lensing of the CMB.” *Physics Reports* 429.1 (June 2006), pp. 1–65. DOI: 10.1016/j.physrep.2006.03.002.
- [119] Y. Li *et al.* “Assembly and Integration Process of the High-Density Detector Array Readout Modules for the Simons Observatory.” *Journal of Low Temperature Physics* 199.3-4 (May 2020), pp. 985–993. DOI: 10.1007/s10909-020-02386-6.
- [120] Y. Li *et al.* “Performance of the advanced ACTPol low frequency array.” *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*. Vol. 10708. SPIE, July 2018, p. 8. DOI: 10.1117/12.2313942.
- [121] D. J. Lockwood. “Rayleigh and Mie Scattering.” *Encyclopedia of Color Science and Technology*. Springer New York, 2016, pp. 1097–1107. DOI: 10.1007/978-1-4419-8071-7_218.
- [122] E. V. Loewenstein, D. R. Smith, and R. L. Morgan. “Optical Constants of Far Infrared Materials 2: Crystalline Solids.” *Applied Optics* 12.2 (Feb. 1973), p. 398. DOI: 10.1364/ao.12.000398.
- [123] R. Loudon. *The Quantum Theory of Light*. Clarendon Press, 1983, pp. 8–13. ISBN: 978-0198501763.
- [124] A. Ludwig. “The definition of cross polarization.” *IEEE Transactions on Antennas and Propagation* 21.1 (Jan. 1973), pp. 116–119. DOI: 10.1109/TAP.1973.1140406.
- [125] L. Mandel. “Fluctuations of Photon Beams and their Correlations.” *Proceedings of the Physical Society* 72.6 (Dec. 1958), pp. 1037–1048. DOI: 10.1088/0370-1328/72/6/312.
- [126] L. Mandel. “Fluctuations of photon beams: The distribution of the photo-electrons.” *Proceedings of the Physical Society* 74.3 (Sept. 1959), pp. 233–243. DOI: 10.1088/0370-1328/74/3/301.
- [127] L. Mandel and E. Wolf. “Photon Correlations.” *Physical Review Letters* 10.7 (Apr. 1963), pp. 276–277. DOI: 10.1103/PhysRevLett.10.276.
- [128] J. C. Mather, E. S. Cheng, R. E. Eplee Jr., *et al.* “A preliminary measurement of the cosmic microwave background spectrum by the Cosmic Background Explorer (COBE) satellite.” *The Astrophysical Journal Letters* 354 (May 1990), pp. L37–L40. DOI: 10.1086/185717.
- [129] J. C. Mather. “Bolometer noise: nonequilibrium theory.” *Applied Optics* 21.6 (Mar. 1982), p. 1125. DOI: 10.1364/ao.21.001125.

- [130] F. T. Matsuda, S. Takakura, *et al.* “Cross-polarization systematics due to Mizuguchi-Dragone condition breaking by a continuously rotating half-wave plate at prime focus in the Huan Tran telescope.” *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*. Vol. 1070849. SPIE, July 2018, p. 144. DOI: 10.1117/12.2313177.
- [131] T. Matsumura and The LiteBIRD Collaboration. “LiteBIRD: Mission Overview and Focal Plane Layout.” *Journal of Low Temperature Physics* 184.3-4 (2016). DOI: 10.1007/s10909-016-1542-8.
- [132] T. Matsumura. “A cosmic microwave background radiation polarimeter using superconducting magnetic bearings.” PhD thesis. University of Minnesota, 2006. URL: https://www.researchgate.net/publication/259670185_A_cosmic_microwave_background_radiation_polarimeter_using_superconducting_magnetic_bearings.
- [133] T. Matsumura and The POLARBEAR Collaboration. “POLARBEAR-2 optical and polarimeter designs.” *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VI*. Vol. 8452. Oct. 2012, 84523E. DOI: 10.1117/12.926770.
- [134] T. Matsumura, S. Hanany, *et al.* “Performance of three- and five-stack achromatic half-wave plates at millimeter wavelengths.” *Applied Optics* 48.19 (July 2009), pp. 3614–3625. DOI: 10.1364/AO.48.003614.
- [135] C. L. Mehta and E. Wolf. “Coherence Properties of Blackbody Radiation. I. Correlation Tensors of the Classical Field.” *Physical Review* 134.5A (June 1964), A1143–A1149. DOI: 10.1103/PhysRev.134.A1143.
- [136] C. Mehta. “Coherence and statistics of radiation.” *Lect.Theor.Phys.c* 7 (1965), pp. 345–401. URL: <https://inspirehep.net/literature/1365348>.
- [137] B. Mellish. *Half-wave plate*. 2020. URL: https://en.wikipedia.org/wiki/Waveplate#Half-wave_plate.
- [138] Y. Mizuguchi, M. Akagawa, and H. Yokoi. “Offset Dual Reflector Antenna.” *IEEE International Symposium on Antennas and Propagation Digest* (1976), pp. 2–5. URL: <https://ui.adsabs.harvard.edu/abs/1976isap.conf....2M/abstract>.
- [139] L. Moncelsi, P. A. R. Ade, F. E. Angilè, *et al.* “Empirical modelling of the BLASTPol achromatic half-wave plate for precision submillimetre polarimetry.” *Monthly Notices of the Royal Astronomical Society* 437.3 (Jan. 2014), pp. 2772–2789. DOI: 10.1093/mnras/stt2090.
- [140] B. L. Morgan and L. Mandel. “Measurement of Photon Bunching in a Thermal Light Beam.” *Physical Review Letters* 16.22 (May 1966), pp. 1012–1015. DOI: 10.1103/PhysRevLett.16.1012.

- [141] S. Moyerman, E. Bierman, *et al.* “Scientific verification of faraday rotation modulators: Detection of diffuse polarized galactic emission.” *Astrophysical Journal* 765.1 (Mar. 2013), p. 64. DOI: 10.1088/0004-637X/765/1/64.
- [142] A. Nadolski, J. D. Vieira, J. A. Sobrin, A. M. Kofman, *et al.* “Broadband, millimeter-wave antireflection coatings for large-format, cryogenic aluminum oxide optics.” *Applied Optics* 59.10 (Apr. 2020), p. 3285. DOI: 10.1364/ao.383921.
- [143] J. M. Nagy *et al.* “A New Limit on CMB Circular Polarization from SPIDER.” *The Astrophysical Journal* 844.2 (Aug. 2017), p. 151. DOI: 10.3847/1538-4357/aa7cfd.
- [144] M. F. Navaroli. “Precise Astronomical Polarization Angle Calibration and its Impact on Studying Lorentz and Parity Violation in the Cosmic Microwave Background.” PhD thesis. UC San Diego, 2020. URL: <https://escholarship.org/uc/item/1380c9mn>.
- [145] J. P. Nibarger, J. A. Beall, D. Becker, H. M. Cho, A. Fox, G. C. Hilton, J. Hubmayr, D. Li, M. D. Niemack, K. D. Irwin, J. Van Lanen, J. Britton, J. McMahon, and K. W. Yoon. “An 84 pixel all-silicon corrugated feedhorn for CMB measurements.” *Journal of Low Temperature Physics*. Vol. 167. 3-4. Springer, June 2012, pp. 522–527. DOI: 10.1007/s10909-011-0428-z.
- [146] S. Oguri *et al.* “GroundBIRD: Observing Cosmic Microwave Polarization at Large Angular Scale with Kinetic Inductance Detectors and High-Speed Rotating Telescope.” *Journal of Low Temperature Physics* 184.3-4 (Aug. 2016), pp. 786–792. DOI: 10.1007/s10909-015-1420-9.
- [147] Oxford Instruments. *Windows for cryogenic environments*. 2015. URL: <https://nanoscience.oxinst.com/assets/uploads/products/nanoscience/documents/Spectroscopy-Windows.pdf>.
- [148] I. L. Padilla, J. R. Eimer, Y. Li, *et al.* “Two-year Cosmology Large Angular Scale Surveyor (CLASS) Observations: A Measurement of Circular Polarization at 40 GHz.” *The Astrophysical Journal* 889.2 (Jan. 2020), p. 105. DOI: 10.3847/1538-4357/ab61f8.
- [149] S. Padin. “Mapping speed for an array of corrugated horns.” *Applied Optics* 49.3 (Jan. 2010), pp. 479–483. DOI: 10.1364/AO.49.000479.
- [150] L. Page, C. Barnes, G. Hinshaw, D. N. Spergel, J. L. Weiland, E. Wollack, C. L. Bennett, M. Halpern, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, G. S. Tucker, and E. L. Wright. “First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Beam Profiles and Window Functions.” *The Astrophysical Journal Supplement Series* 148.1 (Sept. 2003), pp. 39–50. DOI: 10.1086/377223.
- [151] S. Paine. “The am atmospheric model” (Sept. 2019). DOI: 10.5281/ZENODO.3406496.
- [152] S. Pancharatnam. “Achromatic combinations of birefringent plates.” *Proceedings of the Indian Academy of Sciences - Section A* 41.4 (Apr. 1955), pp. 137–144. DOI: 10.1007/BF03047098.

- [153] V. V. Parshin, R. Heidinger, *et al.* “Silicon as an advanced window material for high power gyrotrons.” *International Journal of Infrared and Millimeter Waves* 16.5 (May 1995), pp. 863–877. DOI: 10.1007/BF02066662.
- [154] S. C. Parshley, M. D. Niemack, R. Hills, *et al.* “The optical design of the six-meter CCAT-prime and Simons Observatory telescopes.” *Ground-based and Airborne Telescopes VII*. Vol. 10700. SPIE, July 2018, p. 145. DOI: 10.1117/12.2314073.
- [155] A. A. Penzias and R. W. Wilson. “A Measurement of Excess Antenna Temperature at 4080 Mc/s.” *The Astrophysical Journal* 142 (July 1965), p. 419. DOI: 10.1086/148307.
- [156] M. J. Persky. “Review of black surfaces for space-borne infrared systems.” *Review of Scientific Instruments* 70.5 (May 1999), pp. 2193–2217. DOI: 10.1063/1.1149739.
- [157] G. Pisano, G. Savini, P. A. R. Ade, V. Haynes, and W. K. Gear. “Achromatic half-wave plate for submillimeter instruments in cosmic microwave background astronomy: experimental characterization.” *Applied Optics* 45.27 (Sept. 2006), pp. 6982–6989. DOI: 10.1364/AO.45.006982.
- [158] C. M. Posada *et al.* “Fabrication of Detector Arrays for the SPT-3G Receiver.” *Journal of Low Temperature Physics* 193.5-6 (Dec. 2018), pp. 703–711. DOI: 10.1007/s10909-018-1924-1.
- [159] E. Postrekhin, K. B. Ma, H. Ye, and W.-K. Chu. “Dynamics and relaxation of magnetic stress between magnet and superconductor in a levitation system.” *IEEE Transactions on Applied Superconductivity* 11.1 (Mar. 2001), pp. 1984–1987. DOI: 10.1109/77.920242.
- [160] B. Rabii, C. D. Winant, J. S. Collins, A. T. Lee, P. L. Richards, M. E. Abroe, S. Hanany, B. R. Johnson, *et al.* “MAXIMA: A balloon-borne cosmic microwave background anisotropy experiment.” *Review of Scientific Instruments* 77.7 (July 2006), p. 071101. DOI: 10.1063/1.2219723.
- [161] P. L. Richards. “Bolometers for infrared and millimeter waves.” *Journal of Applied Physics* 76.1 (July 1994), pp. 1–24. DOI: 10.1063/1.357128.
- [162] A. Ritacco, N. Ponthieu, A. Catalano, *et al.* “Polarimetry at millimeter wavelengths with the NIKA camera: calibration and performance.” *Astronomy & Astrophysics* 599 (Mar. 2017), A34. DOI: 10.1051/0004-6361/201629666.
- [163] D. Rosen, A. Suzuki, *et al.* “Epoxy-based broadband antireflection coating for millimeter-wave optics.” *Applied Optics* 52.33 (Nov. 2013), pp. 8102–8105. DOI: 10.1364/AO.52.008102.
- [164] A. J. Ross, L. Samushia, C. Howlett, W. J. Percival, A. Burden, and M. Manera. “The clustering of the SDSS DR7 main Galaxy sample - I: A 4 per cent distance measure at $z = 0.15$.” *Monthly Notices of the Royal Astronomical Society* 449.1 (Feb. 2015), pp. 835–847. DOI: 10.1093/mnras/stv154.

- [165] K. Rotermond, B. Barch, S. Chapman, K. Hattori, A. Lee, N. Palaio, I. Shirley, A. Suzuki, and C. Tran. “Planar Lithographed Superconducting LC Resonators for Frequency-Domain Multiplexed Readout Systems.” *Journal of Low Temperature Physics* 184.1-2 (July 2016), pp. 486–491. DOI: 10.1007/s10909-016-1554-4.
- [166] Y. Sakurai, T. Matsumura, N. Ktayama, H. Kanai, and T. Iida. “Development of a Cryogenic Remote Sensing Thermometer for CMB Polarization Experiment.” *29th IEEE International Symposium on Space THz Technology*. Pasadena, CA, 2018. URL: <https://www.nrao.edu/meetings/isstt/papers/2018/2018169174.pdf>.
- [167] Y. Sakurai, P. Ashton, A. Kusaka, C. A. Hill, K. Kiuchi, N. Katayama, and O. Tajima. “Half-meter Scale Superconducting Magnetic Bearing for Cosmic Microwave Background Polarization Experiments.” *Journal of Physics: Conference Series* 1590.1 (July 2020), p. 012060. DOI: 10.1088/1742-6596/1590/1/012060.
- [168] Y. Sakurai, T. Matsumura, H. Kataza, S. Utsunomiya, and R. Yamamoto. “Estimation of the Heat Dissipation and the Rotor Temperature of Superconducting Magnetic Bearing Below 10 K.” *IEEE Transactions on Applied Superconductivity* 27.4 (June 2017), pp. 1–4. DOI: 10.1109/TASC.2017.2672688.
- [169] Y. Sakurai, T. Matsumura, N. Katayama, *et al.* “Design and development of a polarization modulator unit based on a continuous rotating half-wave plate for LiteBIRD.” *Proceedings of SPIE - The International Society for Optical Engineering*. Vol. 10708. 2018. DOI: 10.1117/12.2312391.
- [170] M. Salatino, J. Lashner, M. Gerbino, S. Simon, J. Didier, *et al.* “Studies of systematic uncertainties for Simons Observatory: Polarization modulator related effects.” *Proceedings of SPIE - The International Society for Optical Engineering*. Vol. 10708. 2018. DOI: 10.1117/12.2312993.
- [171] M. Sathyanarayana Rao, M. Silva-Feaver, *et al.* “Simons Observatory Microwave SQUID Multiplexing Readout: Cryogenic RF Amplifier and Coaxial Chain Design.” *Journal of Low Temperature Physics* 199.3-4 (May 2020), pp. 807–816. DOI: 10.1007/s10909-020-02429-y.
- [172] A. Schillaci *et al.* “Design and Performance of the First BICEP Array Receiver.” *Journal of Low Temperature Physics* 199.3-4 (May 2020), pp. 976–984. DOI: 10.1007/s10909-020-02394-6.
- [173] Y. Sekimoto *et al.* “Concept design of the LiteBIRD satellite for CMB B-mode polarization.” *Proceedings of SPIE - The International Society for Optical Engineering*. Vol. 10698. 2018. DOI: 10.1117/12.2313432.
- [174] U. Seljak and M. Zaldarriaga. “Direct Signature of Evolving Gravitational Potential from Cosmic Microwave Background.” *Physical Review D* 60.4 (July 1999). DOI: 10.1103/PhysRevD.60.043504.

- [175] U. Seljak and M. Zaldarriaga. “Signature of Gravity Waves in Polarization of the Microwave Background.” *Physical Review Letters* 78.11 (Mar. 1997), pp. 2054–2057. DOI: 10.1103/PhysRevLett.78.2054.
- [176] A. L. Shands. *Mean Precipitable Water in the United States*. Tech. rep. United States Department of Commerce, 1949. URL: https://www.nws.noaa.gov/oh/hdsc/Technical_papers/TP10.pdf.
- [177] J. A. Shaw. “Radiometry and the Friis transmission equation.” *American Journal of Physics* 81.1 (Jan. 2013), pp. 33–37. DOI: 10.1119/1.4755780.
- [178] B. D. Sherwin and M. Schmittfull. “Delensing the CMB with the cosmic infrared background.” *Physical Review D* 92.4 (Aug. 2015), p. 43005. DOI: 10.1103/PhysRevD.92.043005.
- [179] M. Shimon, B. Keating, N. Ponthieu, and E. Hivon. “CMB polarization systematics due to beam asymmetry: Impact on inflationary science.” *Physical Review D* 77.8 (Apr. 2008), p. 83003. DOI: 10.1103/PhysRevD.77.083003.
- [180] F. E. Silveira and S. M. Kurbart. “Hagen-Rubens relation beyond far-infrared region.” *EPL* 90.4 (June 2010), p. 44004. DOI: 10.1209/0295-5075/90/44004.
- [181] S. M. Simon *et al.* “The Advanced ACTPol 27/39 GHz Array.” *Journal of Low Temperature Physics* 193.5-6 (Dec. 2018), pp. 1041–1047. DOI: 10.1007/s10909-018-1963-7.
- [182] S. M. Simon, J. Golec, *et al.* “Feedhorn development and scalability for Simons Observatory and beyond.” *Proceedings of SPIE - The International Society for Optical Engineering*. Vol. 10708. 2018. DOI: 10.1117/12.2313405.
- [183] S. M. Simon *et al.* “The design and characterization of wideband spline-profiled feedhorns for Advanced ACTPol.” *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VIII*. Vol. 9914. SPIE, July 2016, p. 991416. DOI: 10.1117/12.2233603.
- [184] S. M. Simon *et al.* “The design and characterization of wideband spline-profiled feedhorns for Advanced ACTPol.” *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VIII*. Vol. 9914. SPIE, July 2016, p. 991416. DOI: 10.1117/12.2233603.
- [185] P. Siritanasak *et al.* “The Broadband Anti-reflection Coated Extended Hemispherical Silicon Lenses for POLARBEAR-2 Experiment.” *Journal of Low Temperature Physics* 184.3-4 (Aug. 2016), pp. 553–558. DOI: 10.1007/s10909-015-1386-7.
- [186] G. F. Smoot, C. L. Bennett, A. Kogut, *et al.* “First results of the COBE satellite measurement of the anisotropy of the cosmic microwave background radiation.” *Advances in Space Research* 11.2 (Jan. 1991), pp. 193–205. DOI: 10.1016/0273-1177(91)90490-B.

- [187] G. F. Smoot, C. L. Bennett, A. Kogut, E. L. Wright, *et al.* “Structure in the COBE differential microwave radiometer first-year maps.” *The Astrophysical Journal* 396 (Sept. 1992), p. L1. DOI: 10.1086/186504.
- [188] J. A. Sobrin *et al.* “Design and characterization of the SPT-3G receiver.” *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*. Vol. 10708. SPIE, Aug. 2018, p. 51. DOI: 10.1117/12.2314366.
- [189] S. Spinelli, G. Fabbian, A. Tartari, M. Zannoni, and M. Gervasi. “A template of atmospheric O₂ circularly polarized emission for cosmic microwave background experiments.” *Monthly Notices of the Royal Astronomical Society* 414.4 (July 2011), pp. 3272–3280. DOI: 10.1111/j.1365-2966.2011.18625.x.
- [190] A. A. Starobinski. “Spectrum of relict gravitational radiation and the early state of the universe.” *Soviet Journal of Experimental and Theoretical Physics Letters* 30 (Dec. 1979), p. 682. URL: <https://inspirehep.net/literature/147727>.
- [191] B. Steinle and H. P. Baltes. “Radiant intensity and spatial coherence for finite planar sources.” *JOSA* 67.2 (Feb. 1977), pp. 241–247. DOI: 10.1364/JOSA.67.000241.
- [192] J. R. Stevens, N. F. Cothard, E. M. Vavagiakis, *et al.* “Characterization of Transition Edge Sensors for the Simons Observatory.” *Journal of Low Temperature Physics* 199.3-4 (May 2020), pp. 672–680. DOI: 10.1007/s10909-020-02375-9.
- [193] G. M. Stiehl *et al.* “Time-Division SQUID Multiplexers With Reduced Sensitivity to External Magnetic Fields.” *IEEE Transactions on Applied Superconductivity* 21.3 (June 2011), pp. 298–301. DOI: 10.1109/TASC.2010.2091483.
- [194] E. C. G. Sudarshan. “Equivalence of Semiclassical and Quantum Mechanical Descriptions of Statistical Light Beams.” *Physical Review Letters* 10.7 (Apr. 1963), pp. 277–279. DOI: 10.1103/PhysRevLett.10.277.
- [195] H. Sugai *et al.* “Updated Design of the CMB Polarization Experiment Satellite Lite-BIRD.” *Journal of Low Temperature Physics* 199.3-4 (2020). DOI: 10.1007/s10909-019-02329-w.
- [196] A. Suzuki and The POLARBEAR Collaboration. “The POLARBEAR-2 and the Simons Array Experiments.” *Journal of Low Temperature Physics* 184.3-4 (2016). DOI: 10.1007/s10909-015-1425-4.
- [197] A. Suzuki. “Multichroic Bolometric Detector Architecture for Cosmic Microwave Background Polarimetry Experiments.” PhD thesis. UC Berkeley, 2013. URL: https://digitalassets.lib.berkeley.edu/etd/ucb/text/Suzuki_berkeley_0028E_13878.pdf.
- [198] O. Tajima, H. Nguyen, C. Bischoff, A. Brizius, I. Buder, and A. Kusaka. “Novel Calibration System with Sparse Wires for CMB Polarization Receivers.” *Journal of Low Temperature Physics* 167.5 (2012), pp. 936–942. DOI: 10.1007/s10909-012-0545-3.

- [199] S. Takakura *et al.* “Performance of a continuously rotating half-wave plate on the POLARBEAR telescope.” *Journal of Cosmology and Astroparticle Physics* 2017.5 (2017). DOI: 10.1088/1475-7516/2017/05/008.
- [200] S. Takakura. “Characterization of a continuous polarization modulator using a half-wave plate for measurements of degree-scale cosmic microwave background with the POLARBEAR experiment.” PhD thesis. Osaka University, 2017. DOI: <https://doi.org/10.18910/67054>.
- [201] J. Tervo, T. Setälä, J. Turunen, and A. T. Friberg. “Van Cittert-Zernike theorem with Stokes parameters.” *Optics Letters* 38.13 (July 2013), pp. 2301–2303. DOI: 10.1364/OL.38.002301.
- [202] The BICEP-2 and Keck Collaborations. “BICEP2 / Keck Array V: Measurements of B-mode Polarization at Degree Angular Scales and 150 GHz by the Keck Array.” *The Astrophysical Journal* 811.2 (Sept. 2015), p. 126. DOI: 10.1088/0004-637X/811/2/126.
- [203] The BICEP-2 and Keck Collaborations. “Constraints on Primordial Gravitational Waves Using Planck, WMAP, and New BICEP2/Keck Observations through the 2015 Season.” *Physical Review Letters* 121.22 (Nov. 2018), p. 221301. DOI: 10.1103/PhysRevLett.121.221301.
- [204] The CMB-S4 Collaboration. “CMB-S4 Science Book, First Edition” (Oct. 2016). URL: <http://arxiv.org/abs/1610.02743>.
- [205] The CMB-S4 Collaboration. “CMB-S4 Technology Book, First Edition” (June 2017). URL: <http://arxiv.org/abs/1706.02464>.
- [206] The EBEX Collaboration. “The EBEX Balloon-borne Experiment—Optics, Receiver, and Polarimetry.” *The Astrophysical Journal Supplement Series* 239.1 (Nov. 2018), p. 7. DOI: 10.3847/1538-4365/aae434.
- [207] The Planck Collaboration. “Planck 2018 results: I. Overview and the cosmological legacy of Planck.” *Astronomy and Astrophysics* 641 (Sept. 2020), A1. DOI: 10.1051/0004-6361/201833880.
- [208] The Planck Collaboration. “Planck 2018 results. IV. Diffuse component separation.” *Astronomy & Astrophysics* (Jan. 2019). DOI: 10.1051/0004-6361/201833881.
- [209] The Planck Collaboration. “Planck 2018 results. VI. Cosmological parameters.” *Astronomy and Astrophysics* 641 (Sept. 2020), A6. DOI: 10.1051/0004-6361/201833910.
- [210] The Planck Collaboration. “Planck 2018 results. XI. Polarized dust foregrounds.” *Astronomy & Astrophysics* (Nov. 2018). DOI: 10.1051/0004-6361/201832618.

- [211] The POLARBEAR Collaboration. “A measurement of the CMB E-mode angular power spectrum at subdegree scales from 670 square degrees of POLARBEAR data.” *The Astrophysical Journal* 904.1 (May 2020), p. 65. DOI: 10.3847/1538-4357/abbacd.
- [212] The POLARBEAR Collaboration. “A Measurement of the Cosmic Microwave Background B-Mode Polarization Power Spectrum at Sub-Degree Scales from 2 years of POLARBEAR Data.” *The Astrophysical Journal* 848.2 (Oct. 2017), p. 121. DOI: 10.3847/1538-4357/aa8e9f.
- [213] The POLARBEAR Collaboration. “A measurement of the cosmic microwave background B-mode polarization power spectrum at sub-degree scales with POLARBEAR.” *Astrophysical Journal* 794.2 (Oct. 2014), p. 171. DOI: 10.1088/0004-637X/794/2/171.
- [214] The POLARBEAR Collaboration. “A Measurement of the Degree-scale CMB B-mode Angular Power Spectrum with POLARBEAR.” *The Astrophysical Journal* 897.1 (July 2020), p. 55. DOI: 10.3847/1538-4357/ab8f24.
- [215] The Simons Observatory Collaboration. “The Simons Observatory: Science goals and forecasts.” *Journal of Cosmology and Astroparticle Physics* 2019.2 (2019). DOI: 10.1088/1475-7516/2019/02/056.
- [216] R. J. Thornton *et al.* “The Atacama Cosmology Telescope: the Polarization-Sensitive ACTPol Instrument.” *The Astrophysical Journal Supplement Series* 227.2 (Dec. 2016), p. 21. DOI: 10.3847/1538-4365/227/2/21.
- [217] H. Tran, A. Lee, S. Hanany, M. Milligan, and T. Renbarger. “Comparison of the crossed and the Gregorian Mizuguchi-Dragone for wide-field millimeter-wave astronomy.” *Applied Optics* 47.2 (Jan. 2008), pp. 103–109. DOI: 10.1364/AO.47.000103.
- [218] R. J. Trumpler. “Absorption of Light in the Galactic System.” *Publications of the Astronomical Society of the Pacific* 42.248 (Aug. 1930), p. 214. DOI: 10.1086/124039.
- [219] N. Turok. “Ekpyrotic universe: Colliding branes and the origin of the hot big bang.” *Physical Review D - Particles, Fields, Gravitation and Cosmology* 64.12 (Nov. 2001), p. 24. DOI: 10.1103/PhysRevD.64.123522.
- [220] E. D. Valentino and L. Mersini-Houghton. “Testing predictions of the quantum landscape multiverse 1: The Starobinsky inflationary potential.” *Journal of Cosmology and Astroparticle Physics* 2017.3 (Mar. 2017), p. 002. DOI: 10.1088/1475-7516/2017/03/002.
- [221] J. H. Van Vleck. “The absorption of microwaves by oxygen.” *Physical Review* 71.7 (Apr. 1947), pp. 413–424. DOI: 10.1103/PhysRev.71.413.
- [222] E. M. Vavagiakis, S. W. Henderson, K. Zheng, *et al.* “Magnetic Sensitivity of AlMn TESes and Shielding Considerations for Next-Generation CMB Surveys.” *Journal of Low Temperature Physics* 193.3 (Nov. 2018), pp. 288–297. DOI: 10.1007/s10909-018-1920-5.

- [223] J. W. Waters. “2.3. Absorption and Emission by Atmospheric Gases.” *Methods in Experimental Physics* 12.PB (Jan. 1976), pp. 142–176. DOI: 10.1016/S0076-695X(08)60684-5.
- [224] B. Westbrook and The POLARBEAR Collaboration. “The POLARBEAR-2 and Simons Array Focal Plane Fabrication Status.” *Journal of Low Temperature Physics* 193.5-6 (2018). DOI: 10.1007/s10909-018-2059-0.
- [225] S. Withington, C. Tham, and G. Yassin. “Theoretical analysis of planar bolometric arrays for THz imaging systems.” *Millimeter and Submillimeter Detectors for Astronomy*. Vol. 4855. SPIE, Feb. 2003, p. 49. DOI: 10.1117/12.459155.
- [226] E. Wolf and W. H. Carter. “A radiometric generalization of the Van Cittert-Zernike theorem for fields generated by sources of arbitrary state of coherence.” *Optics Communications* 16.3 (Mar. 1976), pp. 297–302. DOI: 10.1016/0030-4018(76)90001-8.
- [227] E. Wolf and W. H. Carter. “Angular distribution of radiant intensity from sources of different degrees of spatial coherence.” *Optics Communications* 13.3 (Mar. 1975), pp. 205–209. DOI: 10.1016/0030-4018(75)90081-4.
- [228] E. Wolf. “Coherence and radiometry.” *JOSA* 68.1 (Jan. 1978), pp. 6–17. DOI: 10.1364/JOSA.68.000006.
- [229] E. Wolf. *Introduction to the Theory of Coherence and Polarization of Light*. Cambridge University Press, 2007, p. 130. ISBN: 9780521822114.
- [230] A. L. Woodcraft. “Predicting the thermal conductivity of aluminium alloys in the cryogenic to room temperature range.” *Cryogenics* 45.6 (June 2005), pp. 421–431. DOI: 10.1016/j.cryogenics.2005.02.003.
- [231] W. L. K. Wu, L. M. Mocanu, and The South Pole Telescope Collaboration. “A Measurement of the Cosmic Microwave Background Lensing Potential and Power Spectrum from 500 deg² of SPTpol Temperature and Polarization Data.” *The Astrophysical Journal* 884.1 (Oct. 2019), p. 70. DOI: 10.3847/1538-4357/ab4186.
- [232] Z. Xie, W. Xue, H. Chen, and Y. Huang. “Mechanical and thermal properties of 99% and 92% alumina at cryogenic temperatures.” *Ceramics International* 37.7 (Sept. 2011), pp. 2165–2168. DOI: 10.1016/j.ceramint.2011.03.066.
- [233] B. Xu and W. O. Hamilton. “Combined mu-metal and niobium superconductor shielding for dc SQUID operation.” *Review of Scientific Instruments* 58.2 (Feb. 1987), pp. 311–312. DOI: 10.1063/1.1139278.
- [234] Z. Xu, M. K. Brewer, P. F. Rojas, Y. Li, K. Osumi, B. Pradenas, *et al.* “Two-year Cosmology Large Angular Scale Surveyor (CLASS) Observations: 40 GHz Telescope Pointing, Beam Profile, Window Function, and Polarization Performance.” *The Astrophysical Journal* 891.2 (Mar. 2020), p. 134. DOI: 10.3847/1538-4357/ab76c2.

- [235] M. Zaldarriaga and U. Seljak. “An All-Sky Analysis of Polarization in the Microwave Background.” *Physical Review D* 55.4 (Feb. 1997), pp. 1830–1840. DOI: 10.1103/PhysRevD.55.1830.
- [236] M. Zaldarriaga and U. Seljak. “Gravitational lensing effect on cosmic microwave background polarization.” *Physical Review D* 58.2 (June 1998), p. 23003. DOI: 10.1103/PhysRevD.58.023003.
- [237] M. Zeisberger and W. Gawalek. “Losses in magnetic bearings.” *Materials Science and Engineering: B* 53.1 (May 1998), pp. 193–197. DOI: 10.1016/S0921-5107(97)00326-7.
- [238] F. Zernike. “The concept of degree of coherence and its application to optical problems.” *Physica* 5.8 (Aug. 1938), pp. 785–795. DOI: 10.1016/S0031-8914(38)80203-2.
- [239] C. Zhang *et al.* “Characterizing the Sensitivity of 40 GHz TES Bolometers for BICEP Array.” *Journal of Low Temperature Physics* 199.3-4 (May 2020), pp. 968–975. DOI: 10.1007/s10909-020-02411-8.
- [240] N. Zhu, J. L. Orlowski-Scherer, Z. Xu, *et al.* “Simons Observatory large aperture telescope receiver design overview.” *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*. Vol. 10708. 9. SPIE, July 2018, p. 79. DOI: 10.1117/12.2312871.
- [241] J. Zmuidzinas. “Thermal noise and correlations in photon detection.” *Applied Optics* 42.25 (Sept. 2003), pp. 4989–5008. DOI: 10.1364/AO.42.004989.

Appendix A

Quantum and classical expectation values

In this appendix, we discuss the relation between the quantum and classical expectation values in Equation 4.19. Given the coherent-state eigenbasis, the expectation value of the first-order coherence operator $\hat{E}_i^{(+)}(t)\hat{E}_j^{(-)}(t)$ can be written as

$$\begin{aligned} & \left\langle \{\alpha_{\vec{k}_p}\} \left| \hat{E}_i^{(+)}(t)\hat{E}_j^{(-)}(t) \right| \{\alpha_{\vec{k}_p}\} \right\rangle \\ &= \sum_{\vec{k}, \vec{k}'} \sum_{p, p'} \frac{\hbar \sqrt{\omega_k \omega_{k'}}}{2\epsilon_0 v} (\hat{\epsilon}_{\vec{k}_p} \cdot \hat{\epsilon}_{\vec{k}'p'}) \alpha_{\vec{k}'p'}^* \alpha_{\vec{k}_p} e^{i(\vec{k}'\vec{r}_2 - \vec{k}\vec{r}_1) - i(\omega_{k'}t_2 - \omega_k t_1)} \\ &= \vec{E}^*(\vec{r}_1, t_1) \cdot \vec{E}(\vec{r}_2, t_2), \end{aligned} \quad (\text{A.1})$$

where $|\{\alpha_{\vec{k}_p}\}\rangle = \prod_{\vec{k}} |\alpha_{\vec{k}_p}\rangle$ and where the *classical* amplitude $\vec{E}(\vec{r}, t)$, which is the observable quantity, is defined as the Fourier transform of the coherent state amplitude $\alpha_{\vec{k}_p}$

$$\vec{E}(\vec{r}, t) \equiv \sum_{\vec{k}, p} \sqrt{\frac{\hbar \omega_k}{2\epsilon_0 v}} \hat{\epsilon}_{\vec{k}_p} \alpha_{\vec{k}_p} e^{i(\vec{k}\vec{r} - \omega_k t)}. \quad (\text{A.2})$$

For a generalized mixed state, we can define the density matrix in the coherent-state basis as

$$\hat{\rho} = \int \cdots \int \{d^2\alpha_{\vec{k}_p}\} p(\{\alpha_{\vec{k}_p}\}) \left| \{\alpha_{\vec{k}_p}\} \right\rangle \left\langle \{\alpha_{\vec{k}_p}\} \right|. \quad (\text{A.3})$$

Therefore, the expectation value of the first-order quantum coherence for the mixed state is given by

$$\begin{aligned}
& \left\langle \hat{\mathbf{E}}^{(-)}(\vec{r}_1, t_1) \hat{\mathbf{E}}^{(+)}(\vec{r}_2, t_2) \right\rangle \\
& \equiv \text{Tr} \left[\hat{\rho} \hat{\mathbf{E}}^{(-)}(\vec{r}_1, t_1) \hat{\mathbf{E}}^{(+)}(\vec{r}_2, t_2) \right] \\
& = \sum_{\{n_{\vec{k}_p}\}} \int \cdots \int \{d^2 \alpha_{\vec{k}_p}\} p(\{\alpha_{\vec{k}_p}\}) \left\langle \{\alpha_{\vec{k}_p}\} \left| \hat{\mathbf{E}}^{(-)}(\vec{r}_1, t_1) \right| \{n_{\vec{k}_p}\} \right\rangle \cdot \left\langle \{n_{\vec{k}_p}\} \left| \hat{\mathbf{E}}^{(+)}(\vec{r}_2, t_2) \right| \{\alpha_{\vec{k}_p}\} \right\rangle \\
& = \int \cdots \int \{d^2 \alpha_{\vec{k}_p}\} p(\{\alpha_{\vec{k}_p}\}) \vec{E}^*(\vec{r}_1, t_1) \cdot \vec{E}(\vec{r}_2, t_2), \tag{A.4}
\end{aligned}$$

where $\{\alpha_{\vec{k}_p}\}$ are the coherent state eigenvalues and $|\{n_{\vec{k}_p}\}\rangle$ are the Fock states, which form a complete orthogonal basis on which we take the trace Tr.

On the other hand, the *classical* probability distribution of $\vec{E}(\vec{r}_1, t_1)$ and $\vec{E}(\vec{r}_2, t_2)$ can be written as a convolution of $p(\{\alpha_{\vec{k}_p}\})$ following Equation A.2:

$$\begin{aligned}
p(\vec{E}_1, \vec{E}_2; \vec{r}_1, t_1, \vec{r}_2, t_2) & \equiv \\
& \int \cdots \int \{d^2 \alpha_{\vec{k}_p}\} p(\{\alpha_{\vec{k}_p}\}) \delta(\vec{E}_1 - \vec{E}(\vec{r}_1, t_1)) \delta(\vec{E}_2 - \vec{E}(\vec{r}_2, t_2)). \tag{A.5}
\end{aligned}$$

This leads to the fact that the first-order coherence is the same between the quantum and classical calculations

$$\begin{aligned}
\left\langle \vec{E}^*(\vec{r}_1, t_1) \cdot \vec{E}(\vec{r}_2, t_2) \right\rangle_c & \equiv \int d\vec{E}_1 \int d\vec{E}_2 p(\vec{E}_1, \vec{E}_2; \vec{r}_1, t_1, \vec{r}_2, t_2) \vec{E}_1^* \cdot \vec{E}_2 \\
& = 2\epsilon_0 v \left\langle \hat{\mathbf{E}}^{(-)}(\vec{r}_1, t_1) \hat{\mathbf{E}}^{(+)}(\vec{r}_2, t_2) \right\rangle, \tag{A.6}
\end{aligned}$$

where $\langle \cdots \rangle_c$ denotes the *classical* expectation value. The equality between the mean of the classical field and the expectation value of the quantum field operator is a result of the **optical equivalence theorem**, which was first noted by Glauber and Sudarshan (1963) [61, 194].

There are two useful outcomes of the quantum-classical equivalence for coherence calculations. Firstly, according to Equation 4.1, the number of photons detected at output i can be decomposed into a summation over input modes a_k^\dagger , modified by the optical scattering matrices, and according to Equation 4.10, the electric field operator is a modal decomposition of photon operators. Therefore, we can calculate the quantum mutual intensity in the basis of the electric field operators

$$B_{ij}(\nu) = \langle b(\vec{r}, \nu) b^\dagger(\vec{r}, \nu) \rangle = 2\epsilon_0 v \left\langle \hat{\mathbf{E}}^{(-)}(\vec{r}_1, t_1) \hat{\mathbf{E}}^{(+)}(\vec{r}_2, t_2) \right\rangle = \left\langle \vec{E}^*(\vec{r}_1, t_1) \cdot \vec{E}(\vec{r}_2, t_2) \right\rangle_c, \tag{A.7}$$

where we have defined the output modes $\{b_i, b_j\}$ by their space-time coordinates (\vec{r}, t) and the input modes by the Fourier complement (\vec{k}, t) . Secondly, according to Equation A.6, the

ensemble average of the mutual intensity for the quantum field is described by the statistics of the classical field. Therefore, we can utilize the machinery of classical coherence and superposition to describe the coherence between quantum detectors.

A.1 Thermal photon density matrix

The statistical state, or the mixed state, of thermal photons can be described using the density matrix $\hat{\rho}$. We first consider a state that is single-moded in both the spatial and frequency domains as well as in the polarization state. Using the creation and annihilation operators a^\dagger and a , the density matrix can be written as that of a Bose-Einstein distribution

$$\hat{\rho} = \frac{e^{-\gamma a^\dagger a}}{\text{Tr}(e^{-\gamma a^\dagger a})} = \sum_n p(n; \bar{n}) |n\rangle \langle n|, \quad (\text{A.8})$$

with

$$\gamma \equiv \frac{h\nu}{k_B T} \quad \text{and} \quad p(n; \bar{n}) \equiv \frac{1}{1 + \bar{n}} \left(\frac{\bar{n}}{1 + \bar{n}} \right)^n. \quad (\text{A.9})$$

Here, $\{|n\rangle\}$ is the Fock state and \bar{n} is the mean occupation number. We consider a detection process whose integration time τ is significantly longer than the coherence time $1/\Delta\nu$, where $\Delta\nu$ is the detection bandwidth. In this case of $\tau \gg 1/\Delta\nu$, the occupation number \bar{n} can be written as

$$\bar{n} = n(T, \nu) = \frac{1}{e^\gamma - 1}. \quad (\text{A.10})$$

We can rewrite the density matrix in terms of Glauber's coherent state $|\alpha\rangle$ as

$$|\alpha\rangle \equiv e^{\alpha a^\dagger - \alpha^* a} |0\rangle, \quad (\text{A.11})$$

which satisfies

$$a |\alpha\rangle = \alpha |\alpha\rangle. \quad (\text{A.12})$$

Then, the density matrix can be rewritten using the **Glauber-Sudarshan P representation** [194, 61] as

$$\hat{\rho} = \int d^2\alpha p_g(\alpha; \bar{n}) |\alpha\rangle \langle \alpha| \quad \text{with} \quad p_g(\alpha; \bar{n}) \equiv \frac{1}{\pi \bar{n}} \exp\left(-\frac{|\alpha|^2}{\bar{n}}\right), \quad (\text{A.13})$$

where the integral is over the entire complex plane. Here, the complex amplitude α follows a Gaussian distribution $p_g(\alpha; \bar{n})$, in agreement with the expectation in the classical limit. The photon counting of a coherent state follows a Poisson distribution as

$$\left| \langle n | \alpha \rangle \right|^2 = \exp(-|\alpha|^2) \frac{|\alpha|^{2n}}{n!}. \quad (\text{A.14})$$

Equations A.13 and A.14 immediately lead to a special case of Mandel's formula [125, 126] as

$$\langle n | \hat{\rho} | n \rangle = \int_0^\infty dW e^{-W} \frac{W^n}{n!} \frac{e^{-W/\bar{n}}}{\bar{n}} = p(n; \bar{n}). \quad (\text{A.15})$$

Appendix B

Classical intensity correlations

As described in Section 4.1.2, the equivalence between the statistics of the quantum and classical electric fields allows us to calculate the photon bunching term $|B_{ij}(\nu)|^2$ for thermal photons in Equation 4.8 using classical intensity fluctuations, and in this appendix, we show this equivalence and highlight its key assumptions. To simplify the calculation, we assume scalar fields, noting that fluctuations between orthogonal polarizations do not correlate. While effects within the optical system, such as scattering and diffraction, can cause mode mixing that induces polarized correlations, we reserve a discussion of these effects to Section 4.2, which investigates photon correlations within a model telescope system.

To begin, we define the classical mutual intensity of the scalar field $E(\vec{r}, \nu)$ at location \vec{r} and frequency ν as

$$\Gamma^{(1,1)}(\vec{r}_i, \vec{r}_j, \nu) = \langle E(\vec{r}_i, \nu) E^*(\vec{r}_j, \nu) \rangle, \quad (\text{B.1})$$

where the superscripts $(1, 1)$ indicate the moment and symmetry of the ensemble average. Using this notation, the intensity at location \vec{r} and frequency ν is $I(\vec{r}, \nu) = \Gamma^{(1,1)}(\vec{r}, \vec{r}, \nu)$. Intensity fluctuations at a given location and frequency are given by

$$\Delta I(\vec{r}, \nu) = I(\vec{r}, \nu) - \langle I(\vec{r}, \nu) \rangle, \quad (\text{B.2})$$

and intensity correlations take the form

$$\begin{aligned} \langle \Delta I(\vec{r}_i, \nu_i) \Delta I(\vec{r}_j, \nu_j) \rangle &= \langle [I(\vec{r}_i, \nu_i) - \langle I(\vec{r}_i, \nu_i) \rangle] [I(\vec{r}_j, \nu_j) - \langle I(\vec{r}_j, \nu_j) \rangle] \rangle \delta(\nu_i - \nu_j) \\ &= \langle I(\vec{r}_i, \nu) I(\vec{r}_j, \nu) \rangle - \langle I(\vec{r}_i, \nu) \rangle \langle I(\vec{r}_j, \nu) \rangle \\ &= \Gamma^{(2,2)}(\vec{r}_i, \vec{r}_j, \nu) - \Gamma^{(1,1)}(\vec{r}_i, \vec{r}_i, \nu) \Gamma^{(1,1)}(\vec{r}_j, \vec{r}_j, \nu), \end{aligned} \quad (\text{B.3})$$

where we note that different frequencies do not correlate such that $\nu_i = \nu_j = \nu$ and where we have defined the symmetric fourth probability moment of the field to be

$$\Gamma^{(2,2)}(\vec{r}_i, \vec{r}_j, \nu) = \langle E(\vec{r}_i, \nu) E^*(\vec{r}_i, \nu) E(\vec{r}_j, \nu) E^*(\vec{r}_j, \nu) \rangle. \quad (\text{B.4})$$

In the case of thermal light, the field amplitude is drawn from a **Gaussian distribution** with a variance equal to its average intensity $\sigma^2 = \langle I \rangle$ [136], allowing us to factorize $\Gamma^{(2,2)}(\vec{r}_i, \vec{r}_j, \nu)$

as [229]

$$\begin{aligned}\Gamma^{(2,2)}(\vec{r}_i, \vec{r}_j, \nu) &= \langle E(\vec{r}_i, \nu) E^*(\vec{r}_i, \nu) \rangle \langle E(\vec{r}_j, \nu) E^*(\vec{r}_j, \nu) \rangle \\ &\quad + \langle E(\vec{r}_i, \nu) E^*(\vec{r}_j, \nu) \rangle \langle E(\vec{r}_j, \nu) E^*(\vec{r}_i, \nu) \rangle \\ &= \Gamma^{(1,1)}(\vec{r}_i, \vec{r}_i, \nu) \Gamma^{(1,1)}(\vec{r}_j, \vec{r}_j, \nu) + \left| \Gamma^{(1,1)}(\vec{r}_i, \vec{r}_j, \nu) \right|^2.\end{aligned}\quad (\text{B.5})$$

Plugging Equation B.5 into Equation B.3 yields the following form for intensity correlations

$$\langle \Delta I(\vec{r}_i, \nu) \Delta I(\vec{r}_j, \nu) \rangle = \left| \Gamma^{(1,1)}(\vec{r}_i, \vec{r}_j, \nu) \right|^2, \quad (\text{B.6})$$

which allows us to equate the photon bunching term $|B_{ij}(\nu)|^2$ in Equation 4.8 to classical intensity correlations via Equation 4.20.

B.1 Polarization correlations

The relationship in Equation B.6 is derived for a scalar field but can be straightforwardly generalized to polarization using classical Stokes vectors (see Section 6.3.1), which are commonly used to describe polarization in mm-wave instruments (see Section 2.4). We specifically focus on linear polarization states Q and U , which are central to CMB characterization, but we note that the following formalism generalizes to circular polarization V as well. In a similar way to Equation B.3, consider the linear Stokes polarization at location \vec{r}_i and frequency ν

$$S_i(\vec{r}_i, \nu) = [Q_i(\vec{r}_i, \nu) + iU_i(\vec{r}_i, \nu)]. \quad (\text{B.7})$$

The polarized intensity correlation between Stokes polarimeters¹ at locations \vec{r}_i and \vec{r}_j is then

$$\begin{aligned}\langle \Delta S_i(\vec{r}_i, \nu) \Delta S_j(\vec{r}_j, \nu) \rangle &= \langle \Delta Q_i(\vec{r}_i, \nu) \Delta Q_j(\vec{r}_j, \nu) \rangle + \langle \Delta U_i(\vec{r}_i, \nu) \Delta U_j(\vec{r}_j, \nu) \rangle \\ &\quad + i [\langle \Delta Q_i(\vec{r}_i, \nu) \Delta U_j(\vec{r}_j, \nu) \rangle + \langle \Delta U_i(\vec{r}_i, \nu) \Delta Q_j(\vec{r}_j, \nu) \rangle].\end{aligned}\quad (\text{B.8})$$

Assuming idealized polarization properties of the instrument, as described in Section 4.2, we fix a global polarization basis such that $Q_i = Q_j = Q$ and $U_i = U_j = U$. In addition, Appendix B.2 shows that $\langle \Delta Q \Delta U \rangle = \langle \Delta U \Delta Q \rangle = 0$ such that

$$\langle \Delta S(\vec{r}_i, \nu) \Delta S(\vec{r}_j, \nu) \rangle = \langle \Delta Q(\vec{r}_i, \nu) \Delta Q(\vec{r}_j, \nu) \rangle + \langle \Delta U(\vec{r}_i, \nu) \Delta U(\vec{r}_j, \nu) \rangle. \quad (\text{B.9})$$

Given the globally-applied polarization basis, the relative sensitivities of $S(\vec{r}_i, \nu)$ and $S(\vec{r}_j, \nu)$ to Q/U are determined only by the Stokes pixels' relative orientation ϕ_{ij} ,² and the correlation relation can be written as

$$\begin{aligned}\langle \Delta S(\vec{r}_i, \nu) \Delta S(\vec{r}_j, \nu) \rangle &= \langle \Delta I_p(\vec{r}_i, \nu) \Delta I_p(\vec{r}_j, \nu) \rangle \cos(2\phi_{ij}) \\ &= \left| \Gamma^{(1,1)}(\vec{r}_i, \vec{r}_j, \nu) \right|^2 \cos(2\phi_{ij}),\end{aligned}\quad (\text{B.10})$$

¹A Stokes polarimeter is defined to be two orthogonal polarimeters whose outputs are subtracted to extract Q or U , as described in Section 6.3.1.

²For example, a pixel oriented as a “+” and another oriented as an “×” have $\phi_{ij} = \pi/4$ and are therefore orthogonal.

where $I_p(\vec{r}_i, \nu) = \sqrt{Q^2(\vec{r}_i, \nu) + U^2(\vec{r}_i, \nu)}$ is the total *polarized* intensity at Stokes polarimeter $S(\vec{r}_i, \nu)$, and where the second equality follows from Equation B.6. Given the above relation, it is advantageous to orient otherwise correlated Stokes pixels with $\phi_{ij} = \pi/4$, which orthogonalizes their outputs and nullifies their photon noise correlation.

B.2 Stokes orthogonality

In this appendix, we show that correlations between Stokes vectors Q and U are zero under the assumption that fluctuations between the field vectors E_x and E_y are uncorrelated. The correlation between the two pixels Q and U is

$$\langle \Delta Q \Delta U \rangle = \langle QU \rangle - \langle Q \rangle \langle U \rangle. \quad (\text{B.11})$$

We can write out this correlation coefficient as

$$\begin{aligned} \langle \Delta Q \Delta U \rangle &= \langle E_x E_x^* E_x E_y^* \rangle + \langle E_x E_x^* E_x^* E_y \rangle \\ &\quad - \langle E_y E_y^* E_x E_y^* \rangle - \langle E_y E_y^* E_x^* E_y \rangle \\ &\quad - \langle E_x E_x^* \rangle \langle E_x E_y^* \rangle - \langle E_x E_x^* \rangle \langle E_x^* E_y \rangle \\ &\quad + \langle E_y E_y^* \rangle \langle E_x E_y^* \rangle + \langle E_y E_y^* \rangle \langle E_x^* E_y \rangle. \end{aligned} \quad (\text{B.12})$$

Noting that

$$\Gamma^{(2,2)}(a, b, c, d) = \langle E_a^* E_b^* E_c E_d \rangle \quad (\text{B.13})$$

and that

$$\Gamma^{(1,1)}(a, b) = \delta_{ab} \langle E_a^* E_b \rangle \quad (\text{B.14})$$

due to the mode independence of orthogonal polarizations (x, y) , we can rewrite

$$\Gamma(Q, U) = \Gamma^{(2,2)}(x, y, x, x) + \Gamma^{(2,2)}(x, x, x, y) + \Gamma^{(2,2)}(y, y, x, y) + \Gamma^{(2,2)}(y, x, y, y). \quad (\text{B.15})$$

Because the correlations for electromagnetic wave noise are inherently thermal and hence fluctuate with a Gaussian random distribution (see Equation B.5), we can write the fourth-order correlation function in terms of the second order correlation function as

$$\Gamma^{(2,2)}(a, b, c, d) = \Gamma^{(1,1)}(a, c) \Gamma^{(1,1)}(b, d) + \Gamma^{(1,1)}(a, d) \Gamma^{(1,1)}(b, c), \quad (\text{B.16})$$

which allows us to simplify the expression in Equation B.15 as

$$\begin{aligned} \Gamma(Q, U) &= 2\Gamma^{(1,1)}(x, x) \Gamma^{(1,1)}(y, x) + 2\Gamma^{(1,1)}(x, x) \Gamma^{(1,1)}(x, y) \\ &\quad + 2\Gamma^{(1,1)}(y, y) \Gamma^{(1,1)}(y, x) + 2\Gamma^{(1,1)}(y, y) \Gamma^{(1,1)}(x, y). \end{aligned} \quad (\text{B.17})$$

Noting that $\Gamma^{(1,1)}(a, b) = \Gamma^{(1,1)*}(b, a)$, the above can be simplified to

$$\Gamma(Q, U) = 4\text{Re}\{\Gamma^{(1,1)}(y, x)\} (\Gamma^{(1,1)}(x, x) + \Gamma^{(1,1)}(y, y)). \quad (\text{B.18})$$

Finally, due to orthogonality between the x and y polarizations, $\Gamma^{(1,1)}(y, x) = 0$ and therefore

$$\Gamma(Q, U) = 0. \quad (\text{B.19})$$