

CONSTRAINING TIME-DEPENDENT PARITY
VIOLATION WITH ACT AND DEVELOPING
NEXT-GENERATION MICROWAVE OBSERVATORIES

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Zachary B. Huber

August 2025

© 2025 Zachary B. Huber

ALL RIGHTS RESERVED

CONSTRAINING TIME-DEPENDENT PARITY VIOLATION WITH ACT AND DEVELOPING NEXT-GENERATION MICROWAVE OBSERVATORIES

Zachary B. Huber, Ph.D.

Cornell University 2025

Precision measurements of the intensity and polarization of the cosmic microwave background (CMB) contain a wealth of information about the early universe and its development over time. Over the last twenty years, telescopes with apertures around six meters in diameter observing from the ground at microwave frequencies, including the Atacama Cosmology Telescope (ACT) in the Atacama Desert in Chile, have begun to contribute heavily to CMB studies on medium and small scales. Their success has driven the development of next-generation millimeter and submillimeter observatories like the Simons Observatory (SO) and the CCAT Observatory.

This dissertation outlines the development of new optical designs for the cryogenic receivers and data acquisition hardware, software, and procedures that will enable the Prime-Cam instrument on the CCAT Observatory's Fred Young Submillimeter Telescope to become the premier mid-scale observatory at submillimeter wavelengths. On the optical design side, exploration of the use of biconic lenses to correct the residual astigmatism of the telescope optics led to the adoption of a biconic lens for an upgrade to SO but concluded that they were not necessary for the 350 GHz module for CCAT. A new design for the Epoch of Reionization Spectrometer, which will use a Fabry-Perot Interferometer (FPI) for line intensity mapping of ionized carbon near the end of the epoch of reionization, was also developed to account for the unique needs of the FPI.

On the data acquisition side, the designs for hardware and software needed to operate Prime-Cam remotely and efficiently are presented along with the current status of their implementation. In addition to adapting much of SO's software to work with Prime-Cam, new software was developed for capturing and storing data from Prime-Cam's RFSoc warm readout system with the largest number of polarization-sensitive microwave kinetic inductance detectors ever deployed for astronomy.

This work concludes with a summary of the first stage of a search for time-dependent, isotropic birefringence in short-duration ACT polarization maps of the kind that could be caused by axion-like dark matter particles (ALPs). Using a power spectrum based estimator, a polarization angle has been calculated for nearly 13,000 ACT depth-1 maps covering five years of observations, and work continues to convert that angle timestream into constraints on ALP properties. Measurements of polarized millimeter-wave emission by ACT and future observatories holds great promise for discoveries in cosmology and astrophysics in the coming decade, including further studies of parity violating effects in galaxies and in the CMB which could be caused by dark matter or new physics.

BIOGRAPHICAL SKETCH

Originally from Lebanon, OH, Zachary Huber graduated from the University of Notre Dame in 2019 with a BA in the Program of Liberal Studies and Physics. While at Notre Dame, he worked on a variety of research projects including a study of r-process nucleosynthesis in simulations of binary neutron star mergers with Dr. Rebecca Surman, analysis of Mars Reconnaissance Orbiter data with Dr. Huiqun Wang at the Smithsonian Astrophysical Observatory, and STM and LEED measurements of Weyl semimetal candidates with Dr. Weida Wu during an REU at Rutgers. For PLS and the Glynn Family Honors Program, he wrote a senior thesis on the role of medievalism in the writings of J.R.R. Tolkien.

After beginning doctoral studies at Cornell University in the fall of 2019, he joined the observational cosmology lab led by Dr. Michael Niemack that winter. Over the last six years, he worked on an array of projects related to the design, testing, and implementation of optics, cryogenics, control software, detectors, and readout for microwave telescopes including the Simons Observatory and CCAT Observatory as well as analysis of Atacama Cosmology Telescope maps. His work was supported by a NASA Space Technology Graduate Research Opportunity fellowship for the development and deployment of CCAT's large-format MKID arrays and next-generation readout, which are slated to be deployed to the field in early 2026 after tremendous efforts by the CCAT team to overcome many challenges.

Outside the lab, Zach was widely engaged in science outreach and education programs, philosophical and theological groups on campus, and the local Catholic parish. Following his graduation, he will be entering the Order of Preachers (colloquially known as the Dominicans), a religious order within the Catholic Church dedicated to study and teaching, where he hopes to combine his many interests in the service of the Church and the world.

Ad maiorem Dei gloriam

ACKNOWLEDGEMENTS

It has long been a source of wonder for me that our telescope projects work at all, just as it is currently a source of wonder for me that my doctoral studies have come to an end. In both endeavors, many people made these things possible.

I could not have asked for a better advisor for my doctorate than Mike Niemack. Mike, I have always been impressed by how much you know about every facet of our telescopes and science and how much work you are able to accomplish, but I have come to appreciate over time that your positive approach to leading people is far more impressive and powerful. You have been an outstanding example to me of how to be a researcher, teacher, mentor, father, and friend, and I am deeply grateful for all the lessons you have taught me. Thank you.

When I was trying to decide what field of physics to specialize in early in my time at Cornell, I decided to join the Niemack group in part because the science was exciting and wide-ranging and in part because I found the people so friendly and welcoming. The latter part was just as important, if not more important, than the former. To the graduate students who came before me, Pato, Eve, Nick, Jason, and Cody: thank you for patiently teaching me so much and for building a lab culture that made it a joy to build our (occasionally rather frustrating) telescopes. Though I've missed those of you who have left, it has only increased my joy to see how each of you have continued to grow as scientists and people.

I also learned so much from our postdoctoral researchers: Yaqiong, Thuong, and Steve. In particular, Yaqiong helped me to become interested in the world of detectors, and Steve challenged me to grow as a researcher by asking good questions and encouraging my interests in new areas of our research. When I wanted to jump into analyzing telescope data, he graciously helped me to get started on the project that became chapter five of this dissertation, and he guided me through the many

winding turns, frustrations, and successes of that project. Along the way, I learned a great deal about analyzing CMB data, but I learned more from his careful way of approaching our work and seeking to do the best science possible. Steve, you have been like a second advisor for me, and I am excited to see where your career takes you. Thank you for everything.

One of the strange things about doing a Ph.D. is that you eventually become the old person in the group. The process of teaching new graduate students and working with new postdocs has helped me to grow as a person and researcher - the group members who arrived after me have taught me more than I have taught them. Ben, Lawrence, Rodrigo, Alicia, Ema, Darshan, Sam, Yuhan, and Victoria: thank you for walking with me as we've dealt with cranky DRs and wirebonders, brought the MKIDs in Mod-Cam to life, attended conferences and research trips, and so much more. I have also been very fortunate to work with so many wonderful undergraduates during my time here: it has always been fun to see you tackle new projects and grow as researchers.

We do not build our telescopes in a vacuum. Our network of collaborators for ACT, CCAT, and SO has always been so supportive. Here at Cornell, I've had the great pleasure of working with Abby Crites, Thomas Nikola, Bugao Zou, and Gordon Stacey. I would like to thank Abby Crites and Eanna Flanagan in a special way for serving on my dissertation committee. Ronan Higgins, Mike Nolte, Steve Parshley, Reinhold Schaaf, Erik Rosolowsky, Graham Gibson, Sourav Sarkar, Yoko Okada, Christof Buchbender, and so many others in the CCAT collaboration have helped me on projects as we brought telescope and instrument subsystems from ideas on paper to reality. There are too many SO and ACT collaborators who have been helpful along this journey to list here, but I will always have a special place in my heart for Brian Koopman and Sanah Bhimani - our time together in

Chile is something that I will never forget.

I have been very fortunate to be supported in my research by a NASA Space Technology Graduate Research Opportunity fellowship. So much of my work was made possible by NASA (and you, federal taxpayer - thank you!), but the best benefit of this grant was the people with whom it brought me into contact. I am especially grateful for my NASA mentor, Peter Day. Peter, I know that my research interests and the contingencies of my projects led me away from your particular field of expertise over time, but I greatly appreciate your advice early in my Ph.D. and the hospitality of you and your team at JPL. That visit helped to broaden my understanding of how to operate MKIDs and where they fit into the future of astronomy. Thank you.

My NASA grant also allowed me to spend a great deal of time working on CCAT MKIDs at NIST Boulder. The experts there, especially Jordan, Jay, and Anna, welcomed me when I knew almost nothing about detectors and taught me more about cryogenics, nanofabrication, and RF engineering than I thought was possible, all while having a lot of fun along the way. Thank you all.

Outside of the lab, I have been blessed with friends who have made my life at Cornell rich beyond measure. Dante, Carmen, Dan, Thomas, Luisa, so many wonderful friends past and present from Grad Group and the Thomistic Institute, Fr. Chumo, Mary Therese, Rich, Jen, the Immaculate Conception parish community, my friends from Notre Dame, and Eric G.: thank you so much for everything. This space is inadequate to recall you all, but I am grateful for the ways in which you have touched my life and allowed me to be a part of yours.

Finally, to my family: Mom, Dad, Jacob, Michelle, Charlotte, and Dani: I could have never reached this point in my life without your love and support. In a way, everything good that I do is from you and for you. I love you.

And now, it is time for you, dear reader, to enter into the world of cosmology and microwave telescopes. I hope that you will find it entertaining and informative. While this work has been done in collaboration with many people, any mistakes or oversights here are mine. Enjoy!

TABLE OF CONTENTS

Biographical Sketch	iii
Dedication	iv
Acknowledgements	v
Table of Contents	ix
List of Tables	xi
List of Figures	xii
1 Introduction	1
1.1 The Cosmic Microwave Background	2
1.2 The Λ CDM Model of the Universe	10
1.3 Beyond Λ CDM	16
1.3.1 Initial Conditions	16
1.3.2 Dark Energy and Dark Matter	17
1.4 Observatories	19
1.4.1 The Atacama Cosmology Telescope	20
1.4.2 The Simons Observatory	21
1.4.3 The CCAT Observatory	24
1.5 Outline	25
2 Observatory Instrumentation	27
2.1 Optics	27
2.2 Mechanical Design and Cryogenics	31
2.3 Detectors and Cold Readout	33
2.3.1 Transition Edge Sensors	34
2.3.2 Microwave Kinetic Inductance Detectors	35
2.3.3 Detector Packaging and Magnetic Shielding	40
2.4 Warm Readout Electronics	41
2.5 Operations Software	45
2.6 Mapmaking and Analysis	48
3 Optical Designs for Microwave Telescopes	51
3.1 Design and Optimization	51
3.2 Broadband Imaging Modules	63
3.2.1 350 GHz Module	64
3.2.2 Advanced Simons Observatory	74
3.3 Spectrometer Modules: The EoR-Spec Module	75
3.4 Future Designs	83
4 Prime-Cam Operations and Data Acquisition	86
4.1 Data Acquisition Equipment and Layout	87
4.1.1 Data Acquisition Hardware	89
4.1.2 Computers and Networking	93
4.2 PCS Development and Implementation	97

4.3	Operations and Data Management	100
4.3.1	Initiating Observations	101
4.3.2	Detector Operations	106
4.3.3	Data Storage and Management	117
4.3.4	Remote Observing Tools	122
4.4	Conclusion	126
5	Constraining Axion-Like Dark Matter with ACT	129
5.1	Axion-Like Dark Matter	129
5.2	ACT Data Release 6	133
5.3	Method	135
5.3.1	Polarization Rotation	135
5.3.2	Angle Estimator	136
5.3.3	Power Spectra Calculation	138
5.3.4	Angle Extraction	145
5.4	Timestream Results	149
5.5	Per Observation Systematic Tests	151
5.6	Next Steps	154
5.6.1	Timestream Fitting	154
5.6.2	Timestream Null Tests	157
5.6.3	Amplitude Upper Limits	159
5.6.4	ALP-Photon Coupling Upper Limits	160
5.7	Conclusion	165
5.A	Appendix: Derivation of Estimator and Covariance	168
5.A.1	Estimator Derivation	168
5.A.2	Covariance Derivation	170
6	Conclusion	172
A	Acronym List	175
B	Zemax Tips and Tricks	181
B.1	Optical Parameters to Check	182
B.2	Zemax Settings	185
B.3	Merit Function and Optimization Settings Considerations	190
B.4	Tolerancing Procedure	193
	Bibliography	197

LIST OF TABLES

1.1	The best fit Λ CDM parameters from the 2018 Planck data release from combined TT, TE, EE, low-E, and lensing data [1].	15
3.1	The lens parameters for the final, as-machined version of the EoR-Spec optical design. Diameters and radii are in millimeters. Higher order coefficients have units that are powers of millimeters and can be obtained from looking at the relevant powers of equation 3.3. The clear diameter for lens 1 is 378 mm, and its center thickness from vertex to vertex is 18.7 mm. For the other lenses, those same two parameters are 336 mm and 21.69 mm (L2), 336 mm and 25.38 mm (L3), and 310 mm and 14.12 mm (L4). Note that these thicknesses are the ones from Zemax and do not include the AR coating thickness, so they differ slightly from the lens drawings.	83
4.1	Summary of DAQ hardware for initial deployment of Prime-Cam with expected quantities and locations.	90
4.2	Summary of key specifications of Prime-Cam computing resources. There will likely be at least five of the servers and two of the NUCs for initial deployment. All of the Supermicro servers will be located in the ES, while all of the NUCs will be located in IS1. These specifications are up to date as of May 2025, but we are considering upgrading the NUCs and may do so in the future.	94

LIST OF FIGURES

1.1	Early COBE data plotted over a 2.735 K blackbody intensity curve from Fig. 2 in Ref. [2]. Later analyses would refine the temperature to 2.725 K and would reduce the errorbars to smaller than the thickness of the line used to plot the best fit blackbody!	3
1.2	Map of the CMB anisotropies as mapped by the Planck satellite from Fig. 6 in Ref. [3]. This map was generated using a component separating technique to remove foregrounds such as emission from the Milky Way and other sources, which were concentrated in the region outlined in gray.	4
1.3	The angular power spectra of the CMB plotted as a function of spherical harmonics multipole ℓ . Measurements are taken from a variety of different experiments. From Fig. 25 in Ref. [4].	8
1.4	The Atacama Cosmology Telescope (center) within its ground screen. The top half of the primary mirror can be seen on the right. Photo credit: Mark Devlin.	20
1.5	<i>Left:</i> The Simons Observatory Large Aperture Telescope after installation on Cerro Toco. Many members of SO are shown for scale. Photo credit: Rolando Dunner. <i>Right:</i> One of the SO Small Aperture Telescopes inside its ground screen. The main structure that is visible is the forebaffle that blocks out stray light; the SAT itself is located in the center of the structure mounted above the center of the platform. Photo credit: Nicholas Galitzki.	22
1.6	The Fred Young Submillimeter Telescope, assembled in Xanten, Germany prior to shipping to Chile. As is clear from a comparison with Fig. 1.5, it is nearly identical to the SO LAT. Photo from this news article as of May 2025.	23
2.1	<i>Left:</i> The optical design of FYST's primary and secondary mirrors showing the large focal plane on the right and large possible field of view at the top. <i>Right:</i> The mechanical design for FYST. From Fig. 1 in Ref. [5].	28
2.2	Ray trace of many different fields on the sky (the different colors) for the 280 GHz module and SO LATR optics design. Various filters, lenses, the cryogenic window, and the Lyot stop are labeled. The focal plane where the detector arrays absorb the light is shown on the right. From Fig. 3 in Ref. [5].	30

2.3	<i>Left:</i> CAD drawing of Prime-Cam in partial cross-section view so that the interior can be seen. The instrument modules occupy the front, and the DR can be seen in the rear behind them. The thermal connections from the DR to the modules are not shown. The 4 K stage (purple) is where the instrument stages mount; the 4 K, 40 K (blue), and 80 K (green) shells are cooled by multiple pulse tubes. One of the readout harness for cabling can be seen on the far right (red and pink). Figure credit: Lawrence Lin. <i>Right:</i> Prime-Cam in the lab at Cornell with the outer shells assembled for vacuum testing. Photo credit: Mike Niemack.	32
2.4	A cutaway view of the 280 GHz module CAD design showing key optical, mechanical, and electronic components. The feedhorns and detector arrays, while not labeled directly, sit directly to the right of the LPE filter. From Fig. 5 in Ref. [6].	33
2.5	<i>Left:</i> a picture of two SO TES test pixels under a microscope. The top one is an MF band (90/150 GHz) dichroic, polarization-sensitive pixel, and the lower one is a UHF band (220/280 GHz) pixel. The OMT antenna is smaller for UHF because the relevant wavelengths are smaller. The four TESs in each pixel (one per polarization at each frequency) appear as small silver rectangles surrounding the central absorber. <i>Right:</i> resistance vs. temperature measurements for these TESs. At low temperatures, the measured resistance fluctuates around zero while the TES is superconducting, but at some critical temperature (~ 180 mK here), the resistance rises to its designed normal resistance of $6\text{ m}\Omega$	35
2.6	<i>Left:</i> The CCAT 280 GHz TiN array in its packaging without the feedhorn array on top. The insert in the bottom right shows a zoomed in view of the MKIDs themselves. From Fig. 1 in Ref. [7]. <i>Right:</i> A comparison of the pixels for the Al and TiN 280 GHz arrays showing the different absorber structures. From Fig. 2 in Ref. [8].	37
2.7	<i>Top Left:</i> two layers of anodized black aluminum collimators mounted on the feedhorn array for the first 280 GHz Al array. <i>Right:</i> the full LED mapping setup mounted in a DR. The LEDs are on the other side of the green PCBs. <i>Bottom Left:</i> The data for a single resonator. The plot with the green background in the bottom right shows a VNA sweep with the corresponding LED on and off - the clear shift in frequency is due to the photons from the LED.	38

2.8	<i>Left:</i> The fractional frequency shift of the fabricated resonators compared to the design for the first CCAT Al 280 GHz MKID array, as determined by LED mapping. <i>Right:</i> The fractional frequency shift for the same array after post-fabrication capacitor trimming. The scale is the same as the first plot, showing that the scatter from the new design frequencies is much smaller than before. Figure credit: Jordan Wheeler.	39
2.9	<i>Left:</i> Magnetic sensitivity of an SO UMM with aluminum outer shell and with copper outer shell. From Fig. 3 of Ref. [9]. <i>Right:</i> Magnetic sensitivity of SO UMM with and without aluminum feed-horn block added. Both tests were conducted without a tin-plated shield around the device. From Fig. 4 of Ref. [9].	41
2.10	<i>Left:</i> An example RFSoc in site-deployable packaging from ASU with auxiliary networking, power, and heat sinking components. Photo credit: Chris Groppi. <i>Right:</i> An example of four VNA sweeps taken simultaneously with an RFSoc for four RF networks on one of the CCAT 280 GHz MKID arrays. Each line below the baseline is the resonance for an MKID. Figure credit: Darshan Patel.	43
2.11	Schematic of a simple SO-OCS network in which one hardware agent is controlled by a client to acquire data that is saved to disk by the aggregator agent. From Fig. 1 in Ref. [10].	47
3.1	<i>Left:</i> Image of Prime-Cam from Fig. 1 of Ref. [11] showing the likely plan for where each module will go in Prime-Cam along with the corresponding Zemax configuration number. The arrow indicates the effective direction that the telescope is pointing in the Zemax design. <i>Right:</i> The six outer optics tubes in Prime-Cam from Zemax, as viewed from the image side behind the detector arrays (looking out towards the telescope mirrors). The colored lines depict rays from different fields on the sky, while the various circles and the hexagon are the different lenses, filters, image surface, and window. Each configuration is labeled with an identifying number and the coordinate system in Zemax as a reference for other figures. This figure is rotated 60 deg relative to Fig. 5 in Ref. [12] to match the convention for module locations in the left side of this figure. The arrow indicates the effective direction that the telescope is pointing in the Zemax design.	54

3.2	<i>Top Left:</i> An example linear-scale contour plot of a PSF for the CCAT 280 GHz design, sampled at a single field for a single configuration at an input frequency of 410 GHz during tests of this design for a possible CCAT 410 GHz module. The central beam is fit with a 2D Gaussian to extract the FWHM and ellipticity. The first of the peaks outside the main beam is visible as a faint pink ellipse. <i>Top Right:</i> The ideal 1D Fraunhofer diffraction intensity pattern with parameters roughly matched to the cross-section in the y-direction of the example PSF. <i>Bottom Left:</i> The log-scale x-direction cross-section of the example PSF along its center. Unlike the linear-scale plot in the top left, many secondary peaks are visible at much lower amplitude. <i>Bottom Right:</i> The log-scale y-direction cross-section of the example PSF along its center.	57
3.3	The Strehl ratio at 280 GHz across the field of view for all six possible Prime-Cam locations for the 280 GHz module. The hexagons represent the three detector arrays in a module, and the caption gives the percentage of the area of those hexagons that is above a Strehl ratio of 0.8.	58
3.4	A colormap (<i>left</i>) and histogram (<i>right</i>) of the chief ray angle for one configuration in the SO UHF optical design. For the colormap, the x and y axes are normalized field coordinates. The pattern is identical for all inner ring configurations, so it will be very similar for any 280 GHz module position in Prime-Cam. All chief ray angles are below 2 deg.	60
3.5	The final lens and image plane in the 280 GHz optical design, annotated to show the angle used for calculating the F/# in one direction from the two extreme pupil rays.	61
3.6	A colormap (<i>left</i>) and histogram (<i>right</i>) of the F/# at the image surface for one configuration in the SO UHF optical design. For the colormap, the x and y axes are normalized field coordinates. The pattern is identical for all inner ring configurations, so it will be very similar for any 280 GHz module position in Prime-Cam. The mean F/# is around 2 for these optics, and all but the most extreme edges of the focal plane where there are few detectors have an F/# within 0.1 of the mean.	62
3.7	Strehl ratio as a function of frequency for one configuration of the 280 GHz module design at 280 GHz (1100 μm), 350 GHz (860 μm), and 410 GHz (730 μm). Without additional optimization of the design at each frequency, the Strehl ratio decreases at higher frequencies.	63

3.8	The RMS wavefront error in units of number of waves for one configuration of the 280 GHz module optical design as a function of angle in the y direction away from the center of the module. In addition to the wavefront error increasing for fields further off-axis, the RMS wavefront error is a greater proportion of a wavelength at 0.86 mm compared to 1.1 mm. The overall wavefront error is effectively constant with frequency in units of μm , so the difference in number of waves only comes from the difference in wavelengths.	65
3.9	The sag along the x and y axis cross-sections for the back of lens two in a CCAT 350 GHz test design that used a biconic lens (<i>left</i>) and one that used an aspheric lens (<i>right</i>). The difference in shape along two orthogonal directions is characteristic of a biconic lens.	67
3.10	<i>Left:</i> A comparison of Strehl ratios for two possible Prime-Cam locations of the 350 GHz module between the same design as the 280 GHz module (top row) and the best single biconic 350 GHz design (bottom row). The Strehl is indeed better for the 350 GHz biconic design. <i>Right:</i> Comparisons of the ellipticity and FWHM of the PSFs for the same two designs across all possible locations in Prime-Cam (though note that these plots are for PSFs calculated in the SO locations; they should be similar in distribution to the Prime-Cam ones). The higher Strehl for the biconic design is driven by the slightly better FWHM, but the higher ellipticity contributed to our choice to stick with the 280 GHz design. This is an updated version of Fig. 4 in Ref. [12] to add more bins and make the bins match across designs for easier comparison.	69
3.11	Comparison of the 50th (lowest bank of data points), 80th (middle bank), and 95th (top bank) percentile encircled energy radii for 25 fields in each possible module location for a 350 GHz aspheric design (left) and a 350 GHz biconic design (right). Each colored line represents the fit for a different module. Generally, fields with higher Strehl ratio have lower encircled energy radii (a decreasing slope for the linear fit), but the biconic design often bucks this trend and has greater scatter in the encircled energy for the fields. This is likely due to its distorted central beam shapes. Figure credit: Isla Steinman.	70
3.12	Ray trace for the best single biconic 350 GHz optical design. The back surface of lens two is the biconic surface. The back of lens three took on an unusual curvature in all of these designs, likely due to the constraints on the lens location and sag.	71
3.13	The Strehl ratio at 350 GHz (from the same design as the 280 GHz module) across the field of view for all six possible Prime-Cam locations for the 350 GHz module. The 350 GHz module is nominally slated to go into configuration four, but it could be switched with the 280 GHz module in configuration three to improve performance.	73

3.14	Comparisons of the Strehl ratio (<i>left</i>) and PSF distributions (<i>right</i>) for the ASO biconic design (“ASO-5”, bottom row) and the original all-aspheric SO design (“ID12”, top row) in the outermost ring of the LATR. The Strehl plots have a different colorbar scale – the ASO biconic design is considerably better than the original SO design for this outer ring module location. The same is true for the PSF ellipticities and FWHMs, showing the great utility of biconic lenses for correcting astigmatism far from the optical axis of the receiver.	75
3.15	Ray trace for the final EoR-Spec optical design showing the four lenses and the highly collimated beam at the Lyot stop for each field. Every color shows a series of rays from the same field on the sky and focused at the detector array.	78
3.16	The Strehl ratio at 280 GHz across the field of view for all six possible Prime-Cam locations for the EoR-Spec design. EoR-Spec is nominally planned to go into configuration five. For most locations in Prime-Cam, all three arrays are diffraction-limited.	79
3.17	The Strehl ratio at 350 GHz across the field of view for all six possible Prime-Cam locations for the EoR-Spec design. EoR-Spec is nominally planned to go into configuration five. At least one array is diffraction-limited for most locations in Prime-Cam.	80
3.18	The $F/\#$ at the Lyot stop across the field of view for all six possible Prime-Cam locations for the EoR-Spec design. This can be a tricky measurement to visualize. From each field on the sky (which maps to fields on the focal plane), rays are traced to the extreme locations of the Lyot stop. These four rays are used to calculate the $F/\#$ using the same procedure mentioned above, and this serves as a proxy for the collimation of the beam at the Lyot stop. We targeted values above 100, which show up as yellow on the truncated scale used here. EoR-Spec is nominally planned to go into configuration five.	81
3.19	The chief ray angle across the field of view for the planned location in Prime-Cam for the EoR-Spec module. The results are similar for other locations. All angles fall below the target of 2 deg except a small number at the edges where there are almost no detectors.	82
3.20	The $F/\#$ at the image across the field of view for the planned location in Prime-Cam for the EoR-Spec module. The results are similar for other locations. The median is slightly higher than the 280 GHz module, but the scatter across the focal plane is reduced, so the overall illumination should be more uniform.	82

4.1	A cross-sectional view of FYST with the two compartments used by Prime-Cam DAQ hardware labeled. Instrument space 1 (IS1, blue) is where Prime-Cam itself will be located along with two electronics racks for DAQ hardware, and the electronics space (ES, orange) is where Prime-Cam computing resources will go.	88
4.2	The Prime-Cam network, showing all computers and connection types as well as relevant HK equipment for the initial deployment of Prime-Cam with four modules. Computers are shown in red with servers in light red and NUCs in darker red. Readout equipment is in green, while other HK devices are in pink. The connection to the site network switch in the ES also connects the subnet to the gateway for outside connections that remote observers will use to authenticate when accessing Prime-Cam computers.	96
4.3	Diagram of the configuration of the PCS agents and other associated software containers for Prime-Cam. The container groups refer to separate docker-compose files that allow that group to be altered without affecting the operation of the rest of the system. It is quite likely that this diagram will evolve over time as new equipment is brought online and as the exact operation of the RFSocS is fine-tuned.	99
4.4	Flow chart of a typical observation with Prime-Cam. Prime-Cam hardware and software are shown in blue. Interfaces with the observatory software (gray), remote observers via a webpage (light blue), and low-level and quick look analysis tools (orange) are shown where they are essential for understanding the observational workflow.	100
4.5	The steps and interfaces involved in initiating an observation with Prime-Cam.	102
4.6	Example of the Nextline web interface for one of the SO SATs. The currently running line is highlighted. Other tabs show the past schedules, any queued schedules, and the real-time scheduler interface, which is under development at this time.	105
4.7	The steps and software packages involved in making measurements with Prime-Cam’s MKIDs during an observation.	107
4.8	Contents of an example g3 file from the testing of the <code>rfsoc-streamer</code> showing the frame structure and contents of the first few frames. Each g3 file begins with an Observation frame, and the data from the RFSoc is stored in Scan frames. The PTP clock on the RFSoc was not calibrated when these data were taken, which is why the PTP timestamps do not match the receiving timestamp.	113
4.9	The steps and software packages involved in processing and transferring the raw g3 files off-site.	117

4.10	The tools involved in remotely monitoring the status of an observation.	121
4.11	The main Grafana dashboard for monitoring Mod-Cam cryogenic performance. The most recent value for the main temperatures and pressures are shown for ease of viewing, and time series of the HK data can be found by scrolling down. In this particular time window, Mod-Cam was warming up to room temperature. Many other types of dashboards and panels can be created.	121
4.12	The Campana web interface for allowing users to configure alerting methods for the SO site. Users can easily toggle alert methods on and off, and the master switch at the top for the user allows all alerts to be paused when the observer is not on duty.	123
4.13	Example of the <code>ocs-web</code> interface for the SO LAT. This particular panel shows the controls for one of the compressor agents. The left sidebar shows all agents available on the crossbar network. The middle panel shows the current status of the agent and a lock to prevent users from inadvertently changing the state of the agent, and the right panels show various agent functions that can be executed via the interface.	124
4.14	A screen capture of the TOD live viewing tool in action. In this example, a Mod-Cam test timestream is loaded in and displayed for all 18 RF networks of the 280 GHz module. The left hand side shows the current value of the timestream in each pixel, while the right hand side shows the timestream and controls. Photo credit: Graham Gibson.	125
5.1	Several different f150 depth-1 Stokes Q maps to provide a sense of the wide variety of shapes and sizes. The depth-1 maps can range from a single pass back and forth at constant elevation to many passes over hours. Some of these examples also show evidence of data cuts during the mapmaking process or missing data. The dominant visual effect in these maps is correlated noise along the scan direction.	134
5.2	Examples of depth-1 maps that had missing pixels that led to unusual apodization patterns.	139
5.3	An example of the tapered mask (top), depth-1 Q (middle), and depth-1 U (bottom) for a single depth-1 map.	142
5.4	An example of the likelihood (top), measured estimator with comparison to the theoretical estimator evaluated at the best fit angle (left), and covariance (right) for the depth-1 map in Fig. 5.3. Each spectrum is plotted as D_ℓ , which is equal to $\frac{\ell(\ell+1)}{2\pi}C_\ell$. The x-axis is ℓ for the spectra.	143

5.5	An example of the depth-1 E cross reference B spectrum (top left), reference E cross depth-1 B spectrum (top right), depth-1 EE autospectrum (middle left), reference EE autospectrum (middle right), depth-1 BB autospectrum (bottom left), and reference BB autospectrum (bottom right) measured for the depth-1 map in Fig. 5.3. Each spectrum is plotted as D_ℓ , which is equal to $\frac{\ell(\ell+1)}{2\pi}C_\ell$. The x-axis is ℓ	144
5.6	Examples of likelihoods for high signal-to-noise maps. The sampled likelihood is shown in blue. These maps form nicely Gaussian likelihoods with peaks and widths that are fit equally well by Gaussian (orange) and skew Gaussian (green) curves.	146
5.7	Examples of likelihoods for very low signal-to-noise maps. The sampled likelihood is shown in blue. These maps do not isolate a peak well and are cut from our analysis.	147
5.8	Examples of likelihoods for borderline signal-to-noise maps. The sampled likelihood is shown in blue. These maps do seem to have a well-defined peak (almost always at an angle far from zero), but they exhibit non-Gaussian behavior near the boundary of the principal value domain of the likelihood. The peak values of these likelihoods, in particular, are better fit by the skew Gaussian curve (green) than the Gaussian (orange).	148
5.9	The timestream of angles for f090 and f150 from ACT DR6 depth-1 maps. 5414 (62.5%) f090 and 7347 (52.8%) f150 maps passed our likelihood cuts. The error bar comes from approximating the FWHM of each likelihood as a Gaussian.	150
5.10	This snippet of the angle timestream shows all f090 and f150 angles labeled by array over a 46.3 day period near the end of 2017. . . .	150
5.11	PTE distribution (dof=4) for the difference between the estimator and the theoretical estimator evaluated at the best fit angle for f150 and f090 for all arrays. A flat PTE distribution indicates that the model is generally fitting the data well. The black line shows the average per bin for a perfectly flat distribution, while the gray lines show the average plus or minus the square root of the average per bin to give a sense of the expected random variations from uniformity. We see a slight excess at low PTEs in the PA5 data for both frequencies, which could be interpreted as an underestimate of the errors by roughly 5–10%, but this is small enough to proceed with the analysis.	152
5.12	PTE distributions (dof=5) for the difference of estimators for measurements of the same area of the sky made by different arrays at the same time for f150 and f090 for all array pairs. The black and gray lines are the same as described in Fig. 5.11.	153

5.13	Constraints on the ALP-photon coupling constant as a function of ALP mass for similar measurements to ours. We will eventually add a line on this plot for our constraints once our null tests are complete, assuming that we do not make a detection. The various colored dashed lines show constraints from AC oscillation searches undertaken by other CMB experiments. The green exclusion region comes from the constraints due to the washout effect explored in Ref. [13] with Planck data. Gray lines and exclusion regions refer to other astrophysical limits (vertical lines) [14] [15] or direct detection limits (horizontal line) [16]. Not shown are time-varying polarization rotation methods that use a different source than the CMB.	163
5.14	A forecast of the best-case 95% upper limits that we could place with our data set obtained by using a timestream with all angles set to zero but the real error bars and timing information from our final timestream. At low frequencies (below roughly 0.2 inverse days), the limits are approximately constant, though there is some scatter at the lowest frequencies that have only a small number of periods in the full length of the timestream. The limits degrade by up to 13% at the highest frequency considered in this test due to the finite duration of our observations. The step-like nature of the limits comes from the sampling resolution of the amplitude in this test.	164
B.1	The machining drawing for lens two of the EoR-Spec optical design. Some of the relevant mechanical constraints shown on the drawing are the optically-active diameter (336 mm, smaller than the total diameter of 352 mm which includes the clamping region), the central thickness of the lens including the sag of both sides and the thickness of the clamp (section A-A), and the shape of the clamp (detail B). Figure credit: Rodrigo Freundt.	185
B.2	The Solidworks CAD drawing of the flange for lens two of the EoR-Spec optical design. The flange width of 4.5 mm (D4) is fixed by ensuring that the black dashed-dotted line, the tangent line that the metamaterial antireflection cutting blade makes to the lens surface, has 0.2 mm clearance on each side of the flange (D6 and D9). The 8 mm length of the flange (D3) is fixed by the clamp size in the mechanical design. Figure credit: Rodrigo Freundt.	186
B.3	The twenty-five fields used for optimizing and analyzing CCAT designs	188
B.4	Settings that I commonly used for the default sequential merit function.	190
B.5	Sag constraints in the merit function	192

B.6	The tolerance data editor for the EoR-Spec tolerancing. The parameters in gray were not used in tolerancing, so the distances between lenses and relative tilts and decenters were the only things toleranced. Units are in mm. The tightest tolerance is the distance between L4 and the array (0.564 mm, moving L4 closer to the array), but this is also the most easily constrained parameter by our carbon fiber strut, so we felt comfortable proceeding.	196
-----	---	-----

CHAPTER 1

INTRODUCTION

The foundation of the modern scientific enterprise is the ability to make precise measurements of the natural world. Such measurements test mathematical models and hypothesized physical principles, sifting our explanations of natural phenomena to separate those consistent with our tests of nature from those that fail to describe the world. The history of science shows time and again that improvements in the accuracy and precision of our measurements leads directly to new observations, insights, and physical theories.

This historical conclusion is undoubtedly true in the field of observational cosmology. Theorizing about the cosmos, the universe or the whole of nature, is an ancient intellectual exercise, but the empirical and theoretical foundation of modern cosmology begins to form in the first three decades of the twentieth century with two key developments: observational evidence for the existence of other galaxies moving away from our own and Einstein's theory of general relativity. (The basic story presented here is given in much greater detail in the engaging history of the field in Ref. [17]. The reader can also find references to original papers there.) The measurements that determined that certain nebulae existed at extragalactic distances and high recession velocities were enabled by significant advances in the size and purity of optical elements at places like Mount Wilson Observatory in southern California.

As the field continued to grow, observational cosmologists increasingly sought empirical tests for competing models of the early universe like the Steady State model of Hoyle, Bondi, and Gold and the hot Big Bang model developed from the work of Lemaître, Gamow, Alpher, and others. Advances in radio frequency and

microwave electronics in the 1940s and 1950s enabled the serendipitous discovery of the cosmic microwave background (CMB) by Wilson and Penzias in 1964 (published in 1965). This ubiquitous and isotropic background radiation at microwave wavelengths is one of the key observables of modern cosmology, and unlocking its secrets motivates the work described below.

1.1 The Cosmic Microwave Background

One of the core predictions of the hot Big Bang theory is that there should be a bath of thermal radiation left over from the time when the universe was much hotter than it is today. Around 13.8 billion years ago, the universe was so hot and dense that neutral atoms could not form without being ionized immediately by highly energetic radiation. Likewise, photons could not travel far without being scattered by charged particles. The universe in this period was opaque, much like a fog is opaque to headlight beams on a car. As the universe expanded and cooled, the photons were redshifted to lower energies. When the effective temperature of this radiation dropped well below the equivalent temperature of the energy required to ionize hydrogen atoms, the universe rapidly switched from a largely ionized state to a largely neutral state. Since neutral atoms only interact with specific frequencies of light, the universe became transparent to photons. This period is often referred to as the epoch of recombination (of electrons with protons) or decoupling (of the CMB from matter). Some of these photons still reach us from the edge of the observable universe today, and we measure these as the cosmic microwave background. The following survey of the CMB and the Λ CDM model is drawn largely from the reviews and textbooks in Refs. [18], [19], [20], [21], and [22], and interested readers should refer to them for greater detail.




Figure 1.1: Early COBE data plotted over a 2.735 K blackbody intensity curve from Fig. 2 in Ref. [2]. Later analyses would refine the temperature to 2.725 K and would reduce the errorbars to smaller than the thickness of the line used to plot the best fit blackbody!

As mentioned above, this radiation is found to be highly isotropic - it is the same in every direction once one removes emission from our own galaxy. The isotropy is even more exact when the leading dipole anisotropy due to Earth's motion through space is also measured and subtracted. This sameness consists largely in its uniform spectral energy density, which is described nearly perfectly by a thermal blackbody spectrum with $T_{CMB} = 2.725$ K. This blackbody spectrum, shown in Fig. 1.1, peaks around a frequency of 160 GHz, which is equivalent to a wavelength around 1.9 mm.

Though the CMB is nearly perfectly uniform, it fluctuates on the level of hundreds of parts per million in intensity or, equivalently, in temperature. Fig. 1.2 shows a full-sky map made by the Planck satellite of these tiny deviations away from the average blackbody temperature. These fluctuations indicate places in the




Figure 1.2: Map of the CMB anisotropies as mapped by the Planck satellite from Fig. 6 in Ref. [3]. This map was generated using a component separating technique to remove foregrounds such as emission from the Milky Way and other sources, which were concentrated in the region outlined in gray.

early universe of slightly higher and lower density that will eventually evolve into the large-scale structure of galaxy clusters that we see today. In addition to these temperature anisotropies, the light from the CMB is also partially polarized by Thompson scattering off of free electrons near decoupling.

At the largest scales, the CMB anisotropies are best measured with satellites like COBE (observed 1989–1993), WMAP (2001–2010), and Planck (2012–2015) because water, carbon dioxide, and other molecules in Earth’s atmosphere scatter light near the peak frequencies of the CMB. Balloon-borne missions also provide excellent observing conditions at a fraction of the cost, albeit for a much shorter time. For fluctuations on smaller scales, it can be quite difficult to send telescopes with sufficiently large apertures to space, and it is easier to devise strategies to mitigate atmospheric noise at these scales for ground-based telescopes observing in millimeter and submillimeter wavelength bands at high and dry sites. Accordingly,

observational cosmologists have devoted great effort in the past two decades to the development and operation of terrestrial observatories.

To the limits of current measurements across all scales, the statistical properties of the CMB temperature and polarization anisotropies are described by Gaussian random fields. In this case, all of the information in the CMB is contained in its temperature and the power spectrum of the anisotropies, the Fourier transform of the two-point correlation function in real space. Since the measurements are made on the curved sphere of the sky, it is convenient to decompose the pattern of CMB anisotropies into coefficients for the various spherical harmonics, $Y_{\ell m}$. For a map of the temperature fluctuations away from the mean value parameterized by a direction on the sky, $T(\hat{n})$, the coefficients are defined by

$$T(\hat{n}) = \sum_{\ell m} a_{\ell m}^T Y_{\ell m} \quad (1.1)$$

where ℓ and m follow the usual conventions for spherical harmonics. The power spectrum is defined by the spatial average

$$C_{\ell}^{TT} = \langle a_{\ell m}^{*T} a_{\ell m}^T \rangle, \quad (1.2)$$

but our measurements do not constitute a true average over all positions because we can only measure from one vantage point. The observed power spectrum, therefore, is an average over the azimuthal coordinate m :

$$C_{\ell}^{TT,obs} = \frac{1}{2\ell + 1} \sum_m a_{\ell m}^{*T} a_{\ell m}^T. \quad (1.3)$$

By making a map of the CMB temperature anisotropies, we enable a direct measurement of the components of the power spectrum. These components can be related to the physics of the early universe using general relativity and statistical mechanics, allowing us to use our measurements to test different models of the universe's history and development.

Similar mathematical tools can be used to describe the polarization of the CMB. The polarization state of light can be fully described by a set of four numbers known as the Stokes parameters; see section 7.2 of Ref. [23] or section 7.4 of Ref. [24] for the mathematical details and their application to cosmology. These parameters are easily extracted from our measurements. Commonly denoted T (or I), Q, U, and V, they effectively are measurements of the total intensity (T/I), the linear polarization orientation and polarization fraction (combinations of Q and U), and amount of circular polarization (V). The CMB is observed to have only linear polarization, so we directly measure Q and U in addition to T. These can be decomposed into power spectra using the coefficients of second-order spin-weighted spherical harmonics:

$$(Q \pm iU)(\hat{n}) = \sum_{\ell m} a_{\pm 2, \ell m} \times {}_{\pm 2}Y_{\ell m}(\hat{n}). \quad (1.4)$$

The spin-weighted spherical harmonics correctly account for the rotational symmetry of the Q and U Stokes parameters.

It is most common, however, to represent the polarization power spectrum in terms of nonlocal linear combinations of these spin-weighted coefficients, usually called E modes and B modes.

$$a_{E, \ell m} = -\frac{1}{2}(a_{+2, \ell m} + a_{-2, \ell m}) \quad (1.5)$$

$$a_{B, \ell m} = \frac{i}{2}(a_{+2, \ell m} - a_{-2, \ell m}) \quad (1.6)$$

Dividing linear polarization up into E and B modes is a long-standing convention in cosmology to avoid issues with defining local polarization coordinates on a curved sky. This choice conveniently also separates signals from different physical sources based on the parity of the polarization patterns they create (see section 1.4 of Ref. [18] or the original paper, Ref. [25]). E modes are curl-free fields, somewhat

like static electric fields; B modes are divergence-free fields, like magnetic fields (hence the names).

Fig. 1.3 compiles fairly recent power spectrum results from major CMB telescopes. The power spectrum of the temperature anisotropies is denoted TT and is the dominant signal. Power spectra can also be constructed for E-mode (EE) and B-mode (BB) polarization in analogy to Eq. 1.3, as well as the cross-spectrum between temperature and E-mode polarization (TE). In standard cosmology, the two polarization modes are not correlated: the TB and EB cross-spectra are expected to be zero. Any evidence for parity violating interactions that generate such correlations would point towards something beyond the standard cosmological model.

The EE signal is largely produced by Thomson scattering due to quadrupolar temperature fluctuations near recombination, while the B-mode signal is produced primarily by gravitational lensing of E-modes as the polarized CMB photons travel to us today. As is evident from Fig. 1.3, the temperature and E-mode polarization spectra are far better constrained than the B-mode signal, though current experiments aim to measure B-mode polarization to much greater precision.

Several features appear in these spectra, each of which tells us about some aspect of the physics of the universe. The most prominent is the series of harmonic peaks in the TT and EE spectra. These peaks result from baryon acoustic oscillations (BAO), standing pressure (i.e. sound) waves produced by the pull of gravity and the push of radiation pressure around regions of slight higher and lower density than the mean, often referred to as overdensities and underdensities, in the tightly-coupled plasma of photons and baryons that existed prior to decoupling. The oscillations in the temperature spectrum and the polarization spectra result from different physical aspects of the same phenomenon: driven oscillations in the




Figure 1.3: The angular power spectra of the CMB plotted as a function of spherical harmonics multipole ℓ . Measurements are taken from a variety of different experiments. From Fig. 25 in Ref. [4].

photon-baryon fluid and quadrupolar temperature distributions from the velocity pattern of the fluid, respectively. While out of phase, these patterns probe the same acoustic scale, giving two independent measurements of the dynamics of this early plasma.

At high multipoles, which corresponds to small angular scales, the power spectra exhibit a long damping tail. The damping tail of the CMB above $\ell \approx 1000$ is driven by diffusion of photons from hotter regions to colder regions prior to

decoupling. The mean free path of the photons is related to their scattering off of charged particles, so it is sensitive to the total number of baryons in the universe.

The horizon over which light could have traveled at decoupling has an angular scale which roughly corresponds to $\ell = 100$. As such, any mode with a lower multipole number (and so a larger angular scale) is only affected by the initial conditions and physical processes in place after decoupling, particularly scattering during the epoch of reionization (EoR) in the later universe and the effects of time varying gravitational potentials in the late universe (the integrated Sachs-Wolfe effect).

Breaking the degeneracy between early-time and late-time effects at low multipole requires additional information. From the CMB itself, this comes from two directions. The scattering of CMB photons during the EoR produces new E-mode polarization by the same Thomson scattering mechanisms that produced it prior to recombination. While polarization prior to recombination is at medium and small scales, this late stage polarization is on large scales due to the larger horizon size at the EoR compared to reionization and produces a characteristic “reionization bump” at low ℓ in the TE and EE spectra (visible as a spike on the far left in the EE power spectrum in Fig. 1.3 - the scale of the TE spectrum makes it difficult to see by eye). Precise measurements of the gravitational lensing of high multipole CMB temperature anisotropies also helps disentangle these effects.

In the six decades since the discovery of the CMB, developments in telescope design and location, detector and readout technologies, and data analysis equipment and techniques have enabled high-precision measurements of the intensity and polarization spectra of the CMB on all scales. When combined with measurements of the large-scale structure (LSS) of the universe at more recent times, studies of the

abundance of light nuclei formed in the epoch prior to the formation of the CMB, and the expansion history of universe as traced by supernovae and other standard candles, these measurements led to the development of a concordance model of the history of the universe, its contents, and its dynamical evolution over time: the Λ CDM model of cosmology.

1.2 The Λ CDM Model of the Universe

While the universe might appear to be endlessly complex at first glance, on the largest scales it is very well described by the laws of general relativity and statistical mechanics, a few basic assumptions about the initial conditions of the universe, the contents of the universe, and their distribution, and a set of no more than ten numbers, small enough to count on two hands. A few of those parameters are constrained by assumption or precise measurement, reducing the final number of parameters in the standard model of cosmology to six. The construction of this model, known as the Λ CDM model of cosmology for its inclusion of cold dark matter (CDM) and a cosmological constant (Λ), is one of the most significant developments of the last fifty years of physics. One of the chief achievements of the Λ CDM model is that it provides a detailed picture of the development of the universe from a given set of initial conditions through early-universe nucleosynthesis to recombination and the formation of large-scale structure in the late universe that passes many independent tests from a range of measurements.

In the Λ CDM model, the universe is well described by the Friedmann-Lemaitre-Robertson-Walker (FLRW) metric, which postulates a homogeneous and isotropic expanding universe that is filled with several different sources of energy density.

Small perturbations to the distribution of those particles evolve from the early universe into the structure we see today. The major categories of particles and energy sources in this model are ordinary “baryonic” matter (including both baryons and leptons), non-interacting (“cold”) dark matter, photon radiation, neutrinos,¹ and dark energy.

From observations that the universe is expanding and becoming less dense, the Λ CDM model extrapolates backwards to periods when the universe was more dense and significantly hotter, a state commonly referred to as the “hot Big Bang” scenario. The hot early state of the universe permitted the formation of the conventional sources of energy density that populate the FLRW metric including almost all of the hydrogen and helium in the universe and the sea of thermal radiation that we measure as the CMB today.

A universe with an FLRW metric need not have Euclidean intrinsic spatial geometry; it could indeed be flat, but it could also be open or closed with negative or positive intrinsic curvature, respectively. General relativity dictates that the curvature of a space is related to the energy density it contains, so the relative amount of the density of the constituents of the universe to the critical density reveals the curvature of the universe. For this reason, the energy densities of each of the components described is often given as a fraction of this critical density:

$$\Omega_i = \frac{\rho_i}{\rho_{crit}}. \quad (1.7)$$

ρ_{crit} can be calculated from the Einstein equations of general relativity and depends on the Hubble parameter, H_0 , the relationship between the distance to an object and the apparent speed with which the expansion of the universe carries it

¹The initial construction of Λ CDM assumed that neutrinos are massless, so the inclusion of massive neutrinos is already a small, well understood addition to the model.

away from an observer (see Eq. 22.11 of Ref. [26] for the full equation). Physically, it represents the density for which the universe is spatially flat. There are currently disagreements between different astrophysical measurements over the value of the Hubble parameter, so it is common to parameterize the value of H_0 as a fixed number times an unknown parameter h . This parameter h is usually combined with Ω_i as $\Omega_i h^2$ to create a parameter that is agnostic to the true value of H_0 . In this form, these density parameters ($\Omega_b h^2$, $\Omega_c h^2$, $\Omega_\gamma h^2$, $\Omega_\nu h^2$, $\Omega_\Lambda h^2$, and $\Omega_k h^2$ - representing the physical density of baryons, cold dark matter, photons, neutrinos, dark energy, and the curvature of the universe, respectively) are the primary set of parameters in the Λ CDM model.

Measurements of the CMB and distance-redshift relation of the universe show that the universe is flat to very good accuracy (see, for instance, the final conclusion of section 7.3 in Ref. [1]), and this is generally taken as an assumption of the Λ CDM model. This is equivalent to setting $\Omega_k = 0$. It also implies that the sum of all of the Ω_i parameters equals one, providing a constraint relationship that limits the number of independent density parameters by one. Cosmologists frequently choose to calculate Ω_Λ from the others and assume that dark energy is a cosmological constant with equation of state $w = -1$.

Two more of these density parameters are also constrained by assumption or measurement. The density of photons, Ω_γ is generally taken as fixed due to the precision of the measurement of the temperature of the CMB, which is directly related to Ω_γ . Neutrinos are usually assumed to be relativistic, in which case their energy density can be related to that of photons, or are taken to have a mass sum of $\Sigma m_\nu = 0.06$ eV, the lowest mass sum based on current neutrino oscillation measurement constraints [22]. Removing these parameters from our

list of independent density parameters leaves just the physical energy density of baryons and that of cold dark matter as measured today; these are the first two independent parameters of the Λ CDM model.

Since the relative proportions of cold dark matter and baryonic matter govern the dynamics of the plasma prior to decoupling, both of these parameters can be constrained by the relative heights and locations of the peaks of the CMB power spectra. Measurements of the power spectrum slope in the diffusion tail provide additional constraints on Ω_b since this parameter sets the diffusion length for photons in the early universe. These two parameters provide us with information about the constituents of the homogeneous universe. Along with the assumptions and physical laws detailed above, they can be used to trace the global development of the universe from very early times to the present day.

An accurate description of all the features of the universe today, however, must account for perturbations to the homogeneous background and the development of structure over cosmic time. Four additional independent parameters provide this information in the simplest Λ CDM model. Two parameters, the amplitude A_s and the spectral index n_s , describe a model of the initial conditions of the universe’s matter distribution. The source of the initial fluctuations remains an open question to which we will return. A power law distribution is given by many models for the spectrum of these primordial fluctuations such that

$$P(k) = A_s \left(\frac{k}{k_0} \right)^{n_s - 1}, \quad (1.8)$$

where the pivot scale, k_0 , is usually taken to be 0.002 Mpc^{-1} (though occasionally 0.05 Mpc^{-1} is used). It is possible to extend this model by including higher order terms in an expansion of the spectral index as a function of wavenumber (a non-zero “running” of the spectral index), but thus far observations of this additional

parameter are consistent with zero (see, for example, section 8.2 of Ref. [27]). It is also possible to add additional parameters r and n_t to describe the power spectrum of tensor perturbations due to inflation in an analogous way. These are generally not taken as part of the Λ CDM model, but attempts to measure whether r deviates from zero form a significant part of modern cosmology (see, for example, section 3 of Ref. [28]).

The next parameter, θ_* , measures the angular size of the acoustic scale at recombination. This angle provides information about the maximum size of features in the CMB due to the oscillations of the photon-baryon fluid prior to recombination, which in turn acts as a sort of standard ruler to measure the redshift of recombination. Formally, it is defined as the ratio of the comoving sound horizon of the photon-baryon fluid (the distance that a perturbation can influence) to the comoving angular distance to the CMB

$$\theta_* = \frac{r_s(a_*)}{D_A(a_*)} \quad (1.9)$$

(see Ref. [29] for full derivation). This parameter is given by Planck [1], ACT [27], and SPT-3G [30] in terms of the CosmoMC software package’s approximation to the full numerical value, θ_{MC} [31]. It is measured straightforwardly from the multipole of the CMB power spectrum peaks (see section 3.1 of Ref. [1]), particularly the location of the first peak.

The final free parameter in the Λ CDM model provides information about the period of reionization in the late universe: the optical depth of reionization, τ . Physically, this number corresponds to the fraction of CMB photons that have been Thomson scattered since the EoR. While the optical depth appears in degenerate combinations with amplitude of initial fluctuations (A_s) in models of the CMB power spectra, it can be measured directly from the EE reionization bump

$\Omega_b h^2$	0.02237 ± 0.00015
$\Omega_c h^2$	0.1200 ± 0.0012
$100 \theta_{MC}$	1.04092 ± 0.00031
τ	0.0544 ± 0.0073
$\ln(10^{10} A_s)$	3.044 ± 0.014
n_s	0.9649 ± 0.0042

Table 1.1: The best fit Λ CDM parameters from the 2018 Planck data release from combined TT, TE, EE, low-E, and lensing data [1].

mentioned above to separate the two. Information about the structure of matter in the late universe, usually quantified in the parameters Ω_m and σ_8 , where the latter is the root mean squared matter fluctuations averaged over spheres of size $8h^{-1}$ Mpc in linear perturbation theory, also helps. This parameter can be thought of as quantifying the amplitude of structure formation at a given time. In particular, the amplitude of fluctuations today, σ_8 , is related to the amplitude of the primordial fluctuations, A_s , in a way that does not depend on τ , so the degeneracy can also be broken by precise measurements of the lensing potential through the high- ℓ temperature anisotropies [21].

It is worth noting that alternative parameterizations can be chosen (and often are in the early CMB anisotropy literature). In particular, θ_{MC} and H_0 can both be used for calculations of the redshift of recombination and are sometimes interchanged. CosmoMC recommends the use of θ_{MC} since it has fewer degeneracies with other parameters in the model than H_0 and can be more easily calculated [32]. In any case, the six free parameters chosen are then used to calculate many other cosmological parameters of interest.

Table 1.1 shows the highest precision measurements of the six Λ CDM parameters from the Planck satellite mission. Together, along with the equations of general relativity, nuclear physics, and statistical mechanics that describe the evo-

lution of the universe (see Ref. [20] for a thorough introduction), these represent the culmination of decades of theoretical and observational labor and a momentous accomplishment in the history of physics.

1.3 Beyond Λ CDM

Yet, for all of the success of the Λ CDM model in explaining observational features from the clustering of large-scale structure to the relative abundances of the elements in the universe, several outstanding questions about the Λ CDM model persist. Each remains an intensely active area of research today.

1.3.1 Initial Conditions

One such area is nature of the initial conditions of the universe. As we see it today, the universe is thoroughly homogeneous on the largest scales. The CMB on one side of the sky is the same temperature as the diametrically opposed side to several digits of precision, yet these two regions of space have been too far apart from one another to exchange information or come into thermal equilibrium throughout the entire expansion history of the universe. The leading hypothesis to explain this perplexing state of affairs is that the universe underwent a dramatic expansion in a tiny fraction of a second far back in its history. Starting from a tremendously hot and dense state with physics dominated by one or more scalar fields, it is possible that the universe doubled its size many times over, leaving regions that were initially close together far apart. This scenario, known as inflation, could also help to explain the spatial flatness of the universe's geometry, the nearly scale-

invariant spectrum of initial fluctuations given in Eq. 1.8, and more.

Most inflationary models predict the production of primordial gravitational waves, which would in turn generate primordial B mode polarization in CMB. Disentangling such primordial B modes, if they exist, from the B modes produced by gravitational lensing of E modes is a major goal of contemporary cosmology and is one avenue to constrain various models of inflation. Ref. [33] presents a good review of the motivation for and models of inflation.

1.3.2 Dark Energy and Dark Matter

Two further mysteries lie hidden in the very name of the model itself: dark energy and dark matter. The Λ in Λ CDM stands for the contribution of a cosmological constant in the Einstein field equations of general relativity, often referred to as dark energy. Observational evidence of the redshift-distance relationship at high redshifts from supernovae in distant galaxies suggests that the universe is accelerating in its expansion [34] [35]. The simplest explanation for this state of affairs is an additional, constant energy density, and multiple observational probes including the CMB support the existence of such a term (see, e. g. section 7.4 in Ref. [1] for a discussion of Planck constraints on dark energy or Ref. [36] for a general review). This energy density could be related to the quantum zero-point energy of the vacuum in quantum field theory, but the expected QFT energy density and the measured amount of dark energy differ by many orders of magnitude. At this time, it is not clear what dark energy is, and it seems quite likely that additional physics beyond the Λ CDM model or the Standard Model of particle physics will be needed to elucidate its nature.

The CDM in Λ CDM stands for cold dark matter. Much like dark energy, the existence of dark matter is well-supported by observational evidence, but its nature remains elusive. From the scales of individual galaxies to the distributions of the largest-scale structures probed with the CMB, measurements consistently indicate that there is more gravitational attraction, and therefore more mass, present in massive structures than can be accounted for by baryonic matter found in ordinary places like stars, planetary bodies, nebula, and interstellar media. Astronomers began to find hints along these lines in the 1920s and 1930s, culminating in measurements by Fritz Zwicky of the velocities of galaxies within the Coma cluster that suggested the galaxies were moving too fast to be accounted for by the gravitational attraction of luminous mass alone. Careful measurements of galaxy rotation curves by Vera Rubin and W. Kent Ford in the 1970s added evidence of similar excessive speeds for stars within galaxies. Studies of Big Bang Nucleosynthesis (BBN), the large-scale clustering of galaxies, and the CMB also reached the same conclusion that the amount of matter in the universe exceeds the amount of baryonic matter. In order to be consistent with the amount of clustering of the universe today, they also suggest that dark matter is relatively cold, meaning that its velocity distribution is much less than that of the speed of light. Using tools like weak gravitational lensing, the deflection of the CMB or light from distant sources, we can map the location of this dark matter quite precisely, but it remains uncertain what exactly dark matter is.

Many different models of dark matter exist, often proposing some new particle. Since we have not found these particles in conventional particle physics experiments, they must interact with normal matter weakly, if at all. A form of weakly interacting massive particle (WIMP), like those naturally provided by models of supersymmetry, is a popular dark matter model. Another set of candidates

come from axions or axion-like particles (ALPs), which form the subject of Ch. 5. While the theoretical and experimental parameter space is vast, a wide variety of astrophysical measurements and direct detection experiments in particle physics have not yet discovered any of the best developed theoretical ideas. Ref. [37] and Ref. [38] provide an older and newer review, respectively, of the evidence for dark matter, various models, and types of measurements that offer insight into it.

Modifications to our theory of gravity might also be able to account for this extra gravitational attraction. Like particle descriptions of dark matter, many theoretical concepts have been advanced, but none have yet proved able to account for all the observed data. Ref. [39] reviews existing tests of general relativity and concludes that there is currently no clear evidence for deviations from its description of gravity.

As such, the search to understand the nature of dark matter continues. All of these puzzles in the Λ CDM model point towards the need for deeper, more integrated theories of cosmology and particle physics. It seems likely at this juncture that improved measurements of the CMB might provide hints about the way towards such theories, particularly in the wealth of information that could be contained in its faint polarization signals.

1.4 Observatories

The last three decades have seen the construction and operation of a large number of excellent observatories dedicated to improving our understanding of cosmology and resolving the outstanding questions of the Λ CDM model by making precise measurements of the CMB. I have had the good fortune of working on three such




Figure 1.4: The Atacama Cosmology Telescope (center) within its ground screen. The top half of the primary mirror can be seen on the right. Photo credit: Mark Devlin.

experiments at different stages of their life cycles: the Atacama Cosmology Telescope, the Simons Observatory, and the CCAT Observatory. Each of these facilities is located on a mountain in the Atacama Desert in northern Chile. As one of the highest and driest locations in the world, this desert offers excellent observing conditions for ground-based experiments at millimeter and submillimeter wavelengths where the CMB and associated foregrounds are brightest.

1.4.1 The Atacama Cosmology Telescope

The Atacama Cosmology Telescope (ACT) was a 6-m off-axis Gregorian telescope [40] (shown in Fig. 1.4) that observed between 2007 and 2022 from Cerro Toco at 5200 m in the Atacama. ACT went through three major receivers: the MBAC (2007–2011) [41], ACTPol (2013–2017) [42], and AdvACT (2017–2022) [43]. Each receiver contained three different arrays observing at different frequencies. Over

its lifetime, ACT covered frequency bands centered around 30, 40, 90, 150, 220, and 280 GHz, but the total sensitivity of the observatory is dominated by the monochroic or dichroic polarization-sensitive transition edge sensor (TES) arrays at 90 and 150 GHz in ACTPol and AdvACT.

ACT tackled a variety of science cases including measurements of the CMB anisotropies and CMB lensing [44] [4] [45], surveys of galaxy clusters and cluster science with the Sunyaev-Zel’dovich (SZ) effects [46] [47], studies of millimeter transients [48], and tests of alternative or extended cosmological models beyond Λ CDM [49], to cite just a few of the many papers written with ACT data. To support these science cases, ACT observed several small fields (as small as ~ 100 s sq. deg.) to greater depth as well as a wide field survey ($\sim 19,000$ sq. deg) [50] [51].

1.4.2 The Simons Observatory

In many respects, the Simons Observatory (SO) is the direct successor to ACT as the largest CMB experiment on Cerro Toco, but it also aims at a wider horizon of science goals than ACT. SO is comprised of a suite of telescopes, beginning with one large aperture telescope (LAT) and three small aperture telescopes (SATs). The LAT and one of the SATs are shown in Fig. 1.5. The LAT is a 6-m aperture crossed Dragone telescope optimized for high-resolution science across a wide field of view [5] [52] and a survey area of roughly 40% of the sky. It hosts the Large Aperture Telescope Receiver (LATR), a modular cryogenic camera deploying tens of thousands of dichroic, polarization-sensitive TESs across low frequency (LF; 30 and 40 GHz), mid frequency (MF: 90 and 150 GHz), and ultrahigh frequency (UHF: 220 and 280 GHz) bands with three arrays of two thousand detectors per module [53] [54]. The SATs make use of a ~ 0.5 m refractive aperture and cryogenic



Figure 1.5: *Left:* The Simons Observatory Large Aperture Telescope after installation on Cerro Toco. Many members of SO are shown for scale. Photo credit: Rolando Dunner. *Right:* One of the SO Small Aperture Telescopes inside its ground screen. The main structure that is visible is the forebaffle that blocks out stray light; the SAT itself is located in the center of the structure mounted above the center of the platform. Photo credit: Nicholas Galitzki.

half wave plate to maximize polarization sensitivity for a focal plane of seven TES arrays at a given frequency band [55] [56] [57] across a survey area of roughly 10% of the sky. Both the LAT and the SATs have a higher number of total detectors in the MF band since the CMB is easiest to observe at these frequencies.

The main science goals of SO are outlined in Ref. [28]. One of the chief goals of the observatory is to measure or put tighter upper limits on the tensor-to-scalar ratio, r , through careful measurements of B-mode polarization at the large scales probed by the SATs. Many inflationary models predict a faint B-mode signal at large scales caused by gravitational waves generated during inflation, so this measurement could provide a rare quantitative window into the epoch of inflation. The large scale B-mode signal is masked by foreground contamination and the B-mode signal generated by weak gravitational lensing of E-mode polarization. The LAT survey will allow these confounding signals to be measured independently and removed. The LAT will also be able to study large-scale structure through weak lensing of the CMB, galactic structure, galaxy clusters via the SZ effects,



Figure 1.6: The Fred Young Submillimeter Telescope, assembled in Xanten, Germany prior to shipping to Chile. As is clear from a comparison with Fig. 1.5, it is nearly identical to the SO LAT. Photo from [this news article](#) as of May 2025.

the number and mass of neutrinos via their effect on the CMB damping tail at small scales, and other extragalactic sources like active galactic nuclei and various transients.

For its initial deployment, the SO LATR will only have seven of a possible thirteen instrument modules. Funding has been secured for an extension to SO called Advanced Simons Observatory (ASO), which will fill out the remainder of the LAT and fund a solar array to power the observatory. ASO will enable further small-scale and millimeter transient science, as detailed in Ref. [58].

1.4.3 The CCAT Observatory

The CCAT Observatory² will provide complementary measurements to ACT and SO to bolster their science goals, particularly via the removal of thermal dust emission foregrounds. It will also pursue a host of astrophysical goals of its own. CCAT’s Fred Young Submillimeter Telescope (FYST) is nearly identical to the SO LAT, with the main differences being changing the structural material of the roughly forty ton elevation structure from steel to invar to maintain more precise alignment between the mirrors by reducing thermal expansion and the precision to which the mirrors are polished. The lower surface roughness and better alignment of the FYST mirrors allows observations at higher frequencies, which enables FYST to take better advantage of its 5600 m site on Cerro Chajnantor in the Atacama Desert for measurements between 220 and 850 GHz (1.4 mm to 0.35 mm). FYST will initially observe with Mod-Cam, a single instrument module cryostat for first-light, commissioning, and early science [6]. The first module will be the 280 GHz module, deploying three arrays of monochroic, polarization-sensitive microwave kinetic inductance detectors (MKIDs; sometimes shortened to KIDs for broader applications) [59] [7]. Eventually, Mod-Cam will be replaced by Prime-Cam, a receiver similar to the LATR capable of hosting seven instrument modules. These modules will include both imagers and spectrometers. Currently funded modules include the 280 GHz (1.1 mm), 350 GHz (0.86 mm), and 850 GHz (0.35 mm) [60] broadband imaging modules and the Epoch of Reionization Spectrometer (EoR-Spec) module [61]. About 75% of FYST’s observing time will be devoted to Mod-Cam/Prime-Cam, while the remainder will go to the CCAT Heterodyne Array

²The name carries a significant history that should be understood for parsing the names of relevant papers. The original CCAT proposal was a 25-m submillimeter telescope located on Cerro Chajnantor. After funding fell through, the alternative 6-m mm and submm telescope that became FYST was initially called CCAT-Prime. This was shortened to CCAT again in 2023.

Instrument (CHAI) [62] (see also section 3.4 of Ref. [63]).

The higher frequency bands for Prime-Cam will enable a varied cosmological and astrophysical science program described in Ref. [64]. The science program for CHAI is described in part in section 5.5.2 of Ref. [63]. CCAT will enable detailed studies of the epoch of reionization, the period during which light from early stars and galaxies reionized intergalactic and interstellar gas, via line intensity mapping with EoR-Spec and future spectrometers. It will also enable studies of the dynamics of galaxy formation and substructure through mapping of galactic polarization and dust with the 280 and 350 GHz modules, dusty star-forming galaxies with the 850 GHz module, and the distribution of various types of gas in galaxies with CHAI. Studies of the distribution and properties of thermal dust will be useful for removing this foreground from SO's measurements of the CMB power spectra. Furthermore, it is expected that CCAT could eventually contribute to mm and submm transient science in a similar manner to ACT and SO.

1.5 Outline

All three of these observatories, ACT, SO, and CCAT, have a role to play in the development of cosmology in the coming decade, and each is represented in the work presented in this dissertation. Generally speaking, the remainder of this text will follow CMB photons through an observatory from measurement to scientific results. After the overview of our contemporary understanding of cosmology and the observatories on which I have worked in this chapter, Ch. 2 provides a more thorough introduction to the different facets of an observatory with CCAT as a model. Ch. 3 introduces the design of cold refracting optics for EoR-Spec and

the CCAT 350 GHz module as well as some considerations for the design of the ASO modules. Ch. 4 summarizes the implementation of hardware and software for data acquisition and operations of Mod-Cam and Prime-Cam on FYST. Ch. 5 analyzes data from ACT to put constraints on time-dependent polarization angle oscillations that could arise from the interaction of axion-like dark matter particles with CMB photons. Finally, Ch. 6 summarizes my work and looks forward to how future studies might improve upon it.

CHAPTER 2

OBSERVATORY INSTRUMENTATION

Modern astronomical observatories often entail great complexity. In order to make measurements at the highest sensitivity possible, many details must be attended to in the design, testing, deployment, and operational phases of the experiment. Despite this significant complexity, every contemporary millimeter and submillimeter wavelength observatory shares a great deal in common with other members of its kind. This chapter introduces the core subsystems of our observatories with CCAT and Prime-Cam as the common example.

2.1 Optics

The goal of the optical elements on the telescope is to ensure, as far as possible, that all of the light in a given wavelength band which reaches the telescope is focused onto the detector arrays so that it can be measured. In any real optical system, there will be losses and inefficiencies at each element due to scattering, reflection, absorption, diffraction, and reduction in focus due to the design or machining of the parts. As such, the design and materials must be chosen carefully to maximize the optical efficiency, the fraction of light from the sky that is absorbed by the detectors.

There are two major parts to the focusing system for FYST and other telescopes of similar size at these wavelengths: the reflecting telescope optics and the refracting receiver¹ optics.

¹I will use 'receiver' and 'camera' interchangeably to refer to the cryostat housing the refracting optics, filters, and detectors. Prime-Cam and the LATR are examples of receivers.




Figure 2.1: *Left:* The optical design of FYST’s primary and secondary mirrors showing the large focal plane on the right and large possible field of view at the top. *Right:* The mechanical design for FYST. From Fig. 1 in Ref. [5].

Mirrors can be designed to have excellent reflectivity and low loss across a wide range of wavelengths provided that the size of the mirror is much larger than the wavelength (to minimize diffraction effects) and that the surface roughness of the mirror is much smaller than the wavelength (to minimize scattering effects) (see, e.g., section 5.5.3 of Ref. [65]). FYST will observe at a range of wavelengths from 1.4 mm down to 0.35 mm using two mirrors of diameter ~ 6 m and surface roughness below 7 microns rms. The optical design of the telescope is presented in detail in Ref. [66], building on the design study in Ref. [67]. Fig. 2.1 shows the optical and mechanical design of FYST (which is identical to the SO LAT).

These two mirrors are arranged in a crossed Dragone design, meaning that their relative orientation satisfies the Mizuguchi-Dragone condition (see footnote 15 of Ref. [68] for citations of all the relevant papers and Eq. 13 of Ref. [69] for one statement of the condition). This arrangement has a variety of useful properties: it cancels on-axis astigmatism and coma (an introduction to the basics of aberration theory can be found in section 6.3 of Ref. [70]) and preserves the polarization of incoming light well while allowing for a large diffraction-limited field of view. It also provides a more compact design for ease of construction compared to alternatives

with a similar field of view. The FYST design introduces small perturbations to the shape of the mirrors to suppress off-axis coma, leaving off-axis astigmatism as the dominant optical aberration.

The large field of view due to the mirror size and arrangement allows more of the sky to be monitored at one time by a suitable camera. Due to this design, FYST has an excellent mapping speed, which means it takes less observing time to reach a certain map-level noise threshold. Since the minimum angular distance between two distinguishable sources varies as the inverse of the aperture diameter, the relatively large primary aperture diameter (~ 6 m) compared to previous generations of telescopes like Planck (~ 1.5 m) means that FYST is sensitive to smaller scales (higher ℓ in the power spectra) and point sources than Planck. The angular resolution depends on wavelength, but for FYST's target wavelengths it is around 1 arcminute or smaller.

In order to ensure good optical coupling to the detectors, it is often necessary to use additional optics to reimage the telescope focus to another focal surface. The use of superconducting detectors (described in Sec. 2.3) inside a large cryostat also provides additional space for reimaging optics. In Prime-Cam, this reimaging is accomplished within a series of independent instrument modules by a room temperature ultra-high molecular weight polyethylene vacuum window, cryogenic metal-mesh, alumina, or foam filters to absorb unwanted frequencies of light, metamaterial absorber baffling to absorb scattered and reflected light [71], and cryogenic silicon lenses to maintain a high image quality across a flat focal plane. Each of the optical elements use a metamaterial (MMA) antireflection (AR) layer [72] to reduce reflections.

A Lyot stop is also present in the system. The size of this stop controls how

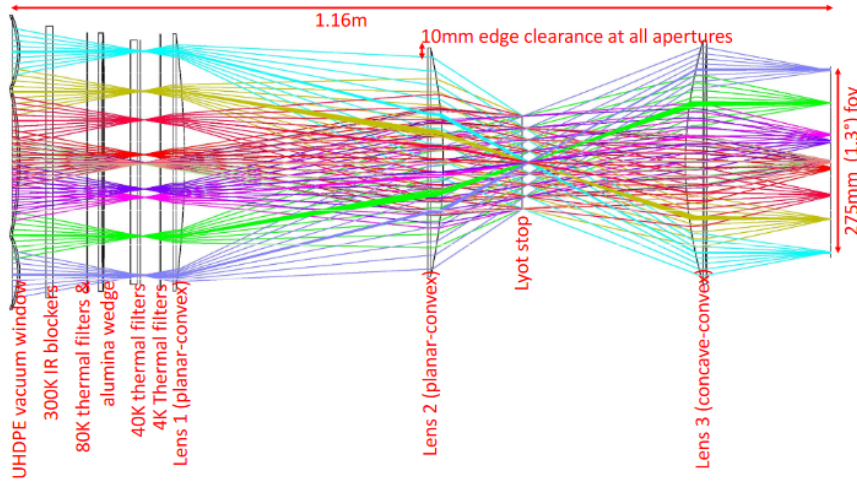


Figure 2.2: Ray trace of many different fields on the sky (the different colors) for the 280 GHz module and SO LATR optics design. Various filters, lenses, the cryogenic window, and the Lyot stop are labeled. The focal plane where the detector arrays absorb the light is shown on the right. From Fig. 3 in Ref. [5].

much of the light from the primary mirror makes it to the detectors, allowing for the absorption of unwanted stray light from the telescope structure or ground before it reaches the detectors. Thinking in a time-reversed sense in which photons move from the detectors to the sky, the Lyot stop controls the illumination of the primary mirror and the field of view on the sky.

After the lenses bring the light to a new focus, metal or silicon feedhorns attached to the detector packaging for SO [73] and Prime-Cam [74] [60] help to couple the focused light to the absorbing elements of the detectors. These feedhorns are small, approximately conical waveguides that are optimized to allow for proper spacing between the detectors (see Ref. [75] and section 2.3 of Ref. [68]), minimizing reflections off of the focal plane, and maximizing absorption by matching the beam of the telescope optics.

The 280 GHz module will use the same optical design as the SO LATR UHF

modules [6]. A ray trace of that optics design is shown in Fig. 2.2, and details of the lens design can be found in Ref. [5]. Within Prime-Cam, the 850 GHz module will be located in the central tube since it has the most stringent optical performance requirements and that location has the highest optical quality. The remaining modules will be located in the one of the outer tubes. Studies of designs for other Prime-Cam modules can be found in Ch. 3.

2.2 Mechanical Design and Cryogenics

In order to measure the tiny temperature fluctuations of the CMB at high sensitivity in a reasonable amount of time, the CMB community often uses large cryostats to cool parts of the optical system and the detectors to temperature below 4 K. For CCAT, we will accomplish this with Prime-Cam, a 1.8 m diameter, 2.8 m long cryostat that will be able to cool seven instrument modules to 4 K with a series of Cryomech pulse tube cryocoolers (PTs). The focal planes of those instrument modules are cooled down further to 100 mK or colder with a Bluefors dilution refrigerator (DR).

Prime-Cam consists of a series of stages cooled to different temperatures, each of which hosts different optical components (see Ref. [76], but note that some development in the design has happened since 2018. Section 2 of Ref. [77] covers the design as of 2021). The assembled Prime-Cam vessel is shown in Fig. 2.3. The room temperature (300 K) outer shell is the main pressure vessel and hosts the AR-coated windows. In Prime-Cam itself, there are 80 K, 40 K, and 4 K shells that host some of the filters to block out infrared light. There are also thermal BUS structures at 1 K and 100 mK (described in Ref. [12], though these may




Figure 2.3: *Left:* CAD drawing of Prime-Cam in partial cross-section view so that the interior can be seen. The instrument modules occupy the front, and the DR can be seen in the rear behind them. The thermal connections from the DR to the modules are not shown. The 4 K stage (purple) is where the instrument stages mount; the 4 K, 40 K (blue), and 80 K (green) shells are cooled by multiple pulse tubes. One of the readout harness for cabling can be seen on the far right (red and pink). Figure credit: Lawrence Lin. *Right:* Prime-Cam in the lab at Cornell with the outer shells assembled for vacuum testing. Photo credit: Mike Niemack.

eventually be replaced by more flexible thermal straps if sufficient cooling power can be obtained with the straps). The instrument modules mount to Prime-Cam at the 4 K stage and have further 1 K and 100 mK stages that are connected directly to the thermal BUS structure. All stages are thermally isolated from one another by G10 or carbon fiber structures with low thermal conductivity. The designs for Mod-Cam [6] and the SO LATR [53] are similar but with some modifications due to the difference in size and complexity of those cryostats.

The elements of the 280 GHz instrument module are shown in Fig. 2.4. The lenses, filters, magnetic shielding, and baffling are all cooled to either 4 K or 1 K, while the detector arrays and the final low pass filter are mounted at 100 mK




Figure 2.4: A cutaway view of the 280 GHz module CAD design showing key optical, mechanical, and electronic components. The feedhorns and detector arrays, while not labeled directly, sit directly to the right of the LPE filter. From Fig. 5 in Ref. [6].

and thermally separated by a carbon fiber truss structure. The 1 K and 100 mK cold fingers can be seen sticking out of the back of the module, where they would then connect to the Prime-Cam thermal BUS structure. In addition to the optical and mechanical elements, the RF and DC cables needed to communicate with thermometers, heaters, amplifiers, and the detector arrays interface with the back of the instrument module after entering Prime-Cam through one of its readout harnesses.

2.3 Detectors and Cold Readout

The purpose of the optical, mechanical, and cryogenic design of the receiver is to allow the light to be measured effectively by the detectors that sit deep in the heart of the cryostat. While there are many types of detectors available, two kinds dominate the current and upcoming generation of experiments: transition edge

sensors and microwave kinetic inductance detectors.

2.3.1 Transition Edge Sensors

For precision broadband measurements at mm and submm frequencies, the gold standard for detectors in the field over the last two decades has been the transition edge sensor (TES). These superconducting bolometers act as highly sensitive thermometers to determine how much energy is absorbed by the detector in a given time. In designs standard in our field (see left of Fig. 2.5), incident light is absorbed by a polarization-sensitive antenna, often a sinuous design or an orthomode transducer (OMT), and is conveyed through a series of on-chip filters to the TES itself. The TES consists of a superconducting element biased by an applied voltage to sit near its critical temperature, T_c . This transition region is characterized by a sharp rise in resistance from zero to the normal resistance of the thin film, meaning that small changes in the temperature of the device correspond to large changes in the resistance (and thus the current flowing through the device). The right side of Fig. 2.5 shows a measurement of the critical temperature for two SO test pixels.

These TESs are suspended on a thin membrane that thermally isolates them from the rest of the system. Thin legs allow wiring to reach the TES and control the thermal conductivity of the system. By carefully tuning the thermal conductivity and the heat capacity of the TES, the system can operate across a wide range of loading conditions with excellent noise performance. The TES bias circuit is then usually coupled to a readout circuit involving superconducting quantum interference devices (SQUIDs) for reading out many TESs with a small number of wires using time domain multiplexing (TDM) or frequency domain multiplexing (FDM). More information about the SO TES design can be found in Ref. [78] (especially

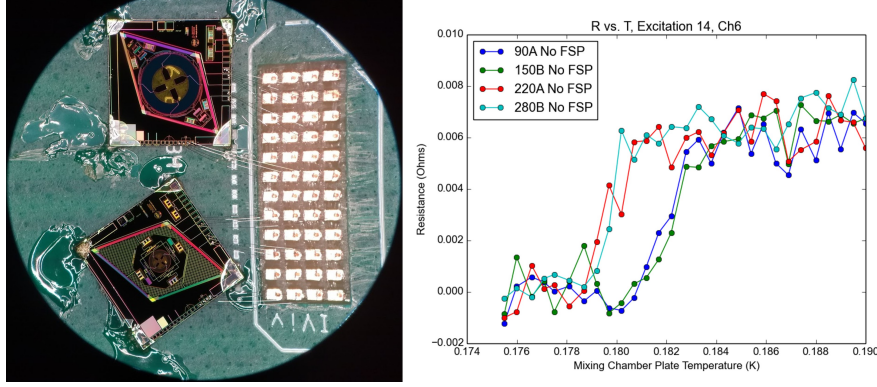


Figure 2.5: *Left:* a picture of two SO TES test pixels under a microscope. The top one is an MF band (90/150 GHz) dichroic, polarization-sensitive pixel, and the lower one is a UHF band (220/280 GHz) pixel. The OMT antenna is smaller for UHF because the relevant wavelengths are smaller. The four TESs in each pixel (one per polarization at each frequency) appear as small silver rectangles surrounding the central absorber. *Right:* resistance vs. temperature measurements for these TESs. At low temperatures, the measured resistance fluctuates around zero while the TES is superconducting, but at some critical temperature (~ 180 mK here), the resistance rises to its designed normal resistance of 6 m Ω .

section 2 and its references). The SO universal focal-plane module (UFM) structure for combining the feedhorns, detectors, and microwave multiplexing (μ mux) readout (a type of FDM) is described in Refs. [79], [80], [81], and [82] (among others). A more thorough summary of the operational principles and mathematics of TESs can be found in Ref. [83].

2.3.2 Microwave Kinetic Inductance Detectors

Despite the success of TESs, their relatively bulky pixel size and complex readout circuit can make it difficult and costly to produce large arrays with a high detector count. As a result, significant development work has gone into developing alternative detector technologies. Microwave kinetic inductance detectors (MKIDs or

KIDs) are one of the most promising of this new group of detectors [84]. Like TESs, MKIDs rely on the superconducting properties of a thin layer of metal, often an aluminum (Al) or a titanium nitride (TiN) film. These films can be patterned into an inductor and a capacitor, forming an LC circuit with a characteristic resonant frequency. The impedance of the superconductor has contributions from the ordinary geometrical inductance of the material and from the kinetic energy stored in the supercurrent of Cooper pairs. This second kind of inductance changes as the Cooper pairs are broken up by photons incident on the detector, which changes the resonant frequency and quality factor of the circuit. By regularly sending a microwave signal into the resonator at its resonant frequency, MKIDs reduce the problem of measuring the number of incident photons to the problem of measuring small changes in the amplitude and phase of the microwave signal. The devices themselves are highly tunable by altering the capacitive properties of the MKID during fabrication, allowing for a wide range of possible resonant frequencies. This enables considerable frequency multiplexing and increases in readout speed by tuning each resonator to a slightly different frequency and using a comb of frequencies to acquire the phase and amplitude data from many resonators at once. Such a design also significantly reduces detector crosstalk and is easier to fabricate than most other superconducting sensors since there is no need for separate resonators and detectors. A detailed review of MKID operation can be found in Ref. [85].

Building on the heritage of experiments like TolTEC [74] and BLAST-TNG [86], CCAT will deploy up to 100,000 MKIDs across a fully populated Prime-Cam. The 280 GHz module will deploy over 10,000 MKIDs spread across three detector arrays: one TiN array that is similar to the TolTEC array in the same band [7] and two Al arrays that should have better noise properties at the expense of more complicated detector tuning [8]. The structure of the array and individual MKIDs




Figure 2.6: *Left:* The CCAT 280 GHz TiN array in its packaging without the feedhorn array on top. The insert in the bottom right shows a zoomed in view of the MKIDs themselves. From Fig. 1 in Ref. [7]. *Right:* A comparison of the pixels for the Al and TiN 280 GHz arrays showing the different absorber structures. From Fig. 2 in Ref. [8].

for the Al and TiN arrays are shown in Fig. 2.6.

The arrays are split into six RF networks with roughly 600 detectors with different resonator frequencies on each network. In the picture of the TiN array in Fig. 2.6, there are two RF networks located side by side in each of the three large rhombi on the array. Each of these networks is monitored by a separate readout chain involving attenuators, coaxial stripline, a cryogenic amplifier, and the warm readout electronics described in Sec. 2.4.

Each MKID on a given RF network is designed to have a different readout frequency within a 500 MHz bandwidth. In practice, due to minute variations in the amount of material etched away during the fabrication process, the inductance and capacitance of each resonator deviates slightly from the designed values. This variation causes shifts in the resonant frequency that can cause two resonators to switch places, making it difficult to know how many resonators are working and which resonators in frequency space go with which optical pixel. Resonators may also collide, meaning that their resonances overlap in frequency space. These res-

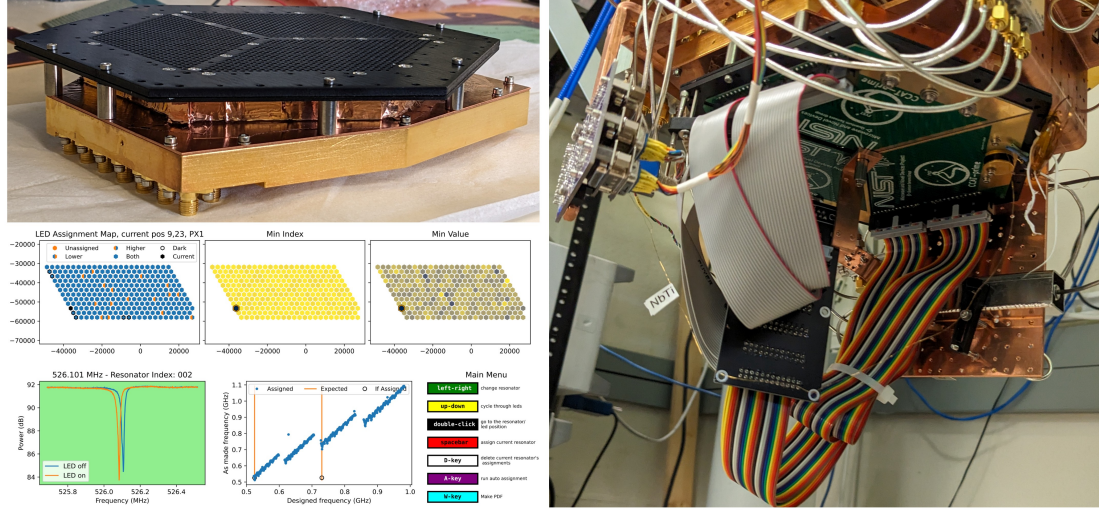


Figure 2.7: *Top Left:* two layers of anodized black aluminum collimators mounted on the feedhorn array for the first 280 GHz Al array. *Right:* the full LED mapping setup mounted in a DR. The LEDs are on the other side of the green PCBs. *Bottom Left:* The data for a single resonator. The plot with the green background in the bottom right shows a VNA sweep with the corresponding LED on and off - the clear shift in frequency is due to the photons from the LED.

onators are often in different physical locations on the array, so collided resonators generally are not useful for observations because they have a high level of crosstalk with another resonator that is measuring a different part of the sky. With observations of point sources on the telescope, it is possible to identify which physical pixel goes with which resonator and which resonators are collided, but if we figure them out before the array is deployed, we can do a post-fabrication trimming procedure to alter all of the resonators to restore the original order and spacing.

To accomplish this, we use an array of LEDs that are mounted above the feedhorns as the array is cooled down. By flashing the LEDs in sequence and monitoring which resonators respond, we can map each resonator to a physical detector. Fig. 2.7 shows the black collimators used to minimize optical crosstalk, the green PCBs with the LEDs mounted on the first Al array in a dilution refrigerator,

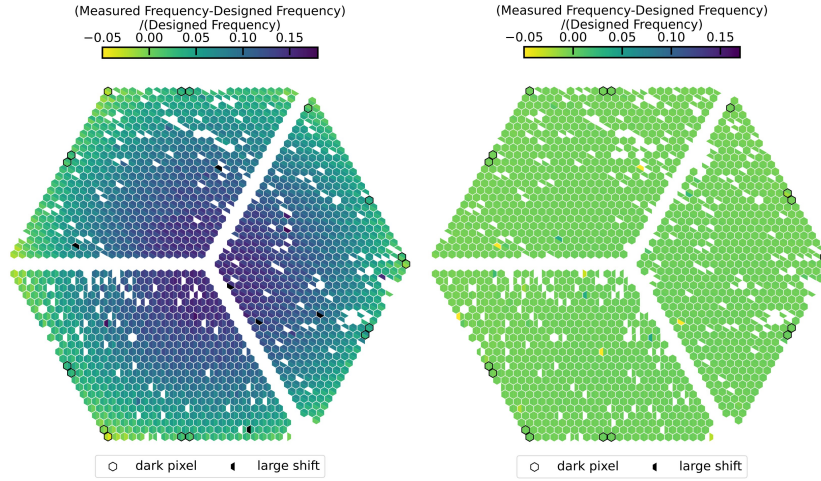


Figure 2.8: *Left:* The fractional frequency shift of the fabricated resonators compared to the design for the first CCAT Al 280 GHz MKID array, as determined by LED mapping. *Right:* The fractional frequency shift for the same array after post-fabrication capacitor trimming. The scale is the same as the first plot, showing that the scatter from the new design frequencies is much smaller than before. Figure credit: Jordan Wheeler.

and the data for one of the resonators. With this configuration, we have been able to map several of our arrays between NIST and Cornell, and the post-processing trimming procedure has been successfully completed for two arrays. Fig. 2.8 shows the fractional frequency shift of each resonator in the array before and after LED trimming. The resonators are spaced much more uniformly after trimming, and the number of collided resonators dropped from $\sim 35\%$ to under 10%.

Results for the 280 GHz TiN array are given in Ref. [87], while results for a 280 GHz Al array are mentioned at the end of section 2 of Ref. [88]. Details about the trimming method in general and the CCAT trimming process are in Ref. [89] and the presentation at Ref. [90], respectively.

Significant progress has been made in characterizing the MKID arrays for the 280 GHz module and in designing and testing early arrays for EoR-Spec, the 350

GHz, and the 850 GHz modules at NIST, Cornell, and the University of British Columbia. Over 10,000 MKIDs are being tested with the 280 GHz module in Mod-Cam at the time of writing and will be deployed on the sky in early 2026.

2.3.3 Detector Packaging and Magnetic Shielding

In addition to the fabrication and operation of the detectors themselves, a variety of other considerations are important to the successful operation of large-format arrays. The cryostat must guarantee enough cooling power to maintain a stable and sufficiently cold operating temperature, and the detector packaging must enable good thermal conductivity for the array. By integrating the feedhorns with the detectors, it must also minimize leakage of light from one detector to another and maximize absorption of the light by the detectors. Furthermore, the packaging must provide the correct electrical ground for the RF networks on the array to minimize reflections, standing waves, and other RF transmission issues.

Another key role of the packaging is providing shielding from stray RF signals and magnetic fields. Both TESs and MKIDs can have their electrical properties altered by time-variable magnetic fluxes generated inside of the cryostat or by Earth’s magnetic field, which can simulate optical signals due to pickup in the detector or its cold readout (for effects in SO TESs and μ mux readout, see Ref. [91] and Ref. [92]; for the effect of ambient magnetic fields on MKID quality factor and noise, see Ref. [93]). Both SO and CCAT address this by adding a large magnetic shield around the back half of the instrument module, but the array packaging itself provides additional shielding when composed of a superconducting material like aluminum that expels magnetic fields below 1.2 K.




Figure 2.9: *Left:* Magnetic sensitivity of an SO UMM with aluminum outer shell and with copper outer shell. From Fig. 3 of Ref. [9]. *Right:* Magnetic sensitivity of SO UMM with and without aluminum feedhorn block added. Both tests were conducted without a tin-plated shield around the device. From Fig. 4 of Ref. [9].

Fig. 2.9 shows an example of the importance of the detector packaging for the SO universal multiplexing module (UMM), which comprises the readout wafers and multiplexing chips (full details in Ref. [9]). These results primarily probe the magnetic pickup of the SQUIDs on the multiplexing chips. With only copper packaging, the fractional phase shift of the resonators is nearly an order of magnitude higher than with the dark aluminum packaging. Adding the aluminum feedhorn block reduces the pickup by another order of magnitude. Tests such as these helped to inform design decisions about the packaging for SO.

2.4 Warm Readout Electronics

In order to make measurements of the amount of light absorbed by the detector over time, readout electronics must be used. For Prime-Cam, this system is the Xilinx ZCU111 Radio Frequency System on a Chip (RFSoc) with custom firmware written by the collaboration [94]. The RFSoc firmware implements a type of frequency domain multiplexing, generating a superposition of sine waves with frequencies corresponding to the resonant frequencies of the resonators on a single

RF network. Each sine wave is commonly referred to as the probe tone for the resonator. This superposition is boosted from a lower frequency baseband to the relevant frequency range of the resonators by an on-board mixing circuit and local oscillator, converted from a digital signal to an analog one by on-board digital-to-analog converters (DACs), and sent into the cryostat via coaxial cables. The phase and amplitude of the tones are altered by interacting with the MKID in ways that depend on the MKID’s optical loading (i.e. absorbed power from photons incident on the detector) and the amount of electrical loading delivered to the MKID by the probe tone. After the signal passes through the remainder of the cold readout and returns to the warm readout system, analog-to-digital converters (ADCs) and a host of electronics for downconverting, filtering and channelization measure the complex transmission (often written S_{21}) in various frequency bins. The system then reports the results by saving output files on the RFSoc or streaming data via UDP packets over the network.

The S_{21} of the system measures the ratio of the input voltage to the output voltage of the system, and it can be decomposed into its real and imaginary parts. The real, or in-phase, component is often denoted ‘I’ and carries the primary signal. The imaginary, or quadrature, component is often denoted ‘Q’. It mostly carries information about the noise of the system, though in resonators driven near the nonlinearity of the kinetic inductance by the electrical power of the probe tone, some of the signal information may be contained in the quadrature component. For Prime-Cam, we record both I and Q for each sample of a resonator to extract the state of the system more accurately. A detailed introduction of MKID operation relevant for Prime-Cam can be found in chapter two of Ref. [95].

The key benefit of the RFSoc that marks advances over previous generations of

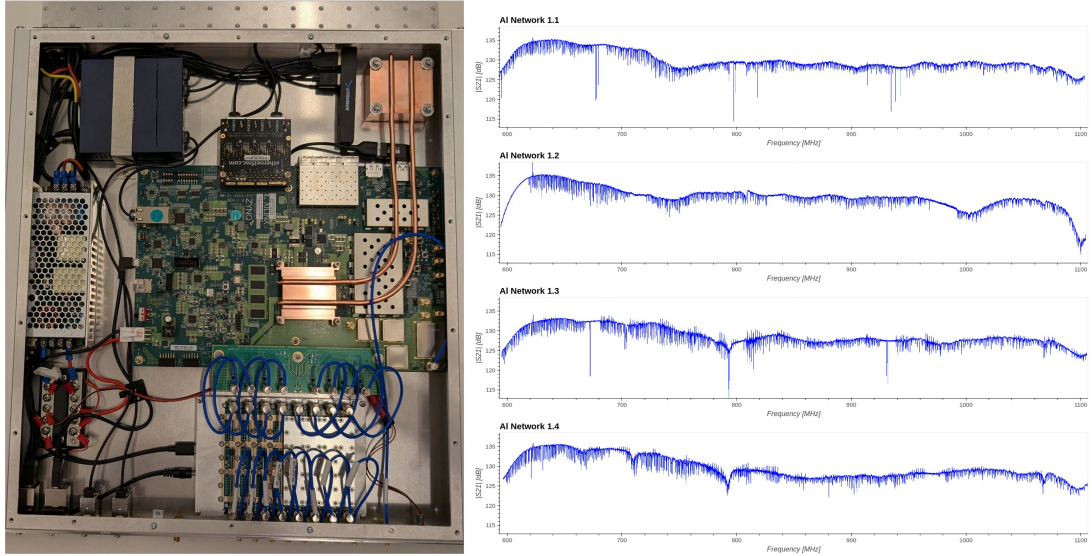


Figure 2.10: *Left:* An example RFSoc in site-deployable packaging from ASU with auxiliary networking, power, and heat sinking components. Photo credit: Chris Groppi. *Right:* An example of four VNA sweeps taken simultaneously with an RFSoc for four RF networks on one of the CCAT 280 GHz MKID arrays. Each line below the baseline is the resonance for an MKID. Figure credit: Darshan Patel.

FDM systems like the ROACH-2 is that it puts all of the DACs, ADCs, and mixers on the board along with a field programmable gate array (FPGA) that does the heavy computational lifting. Combining everything into one board with upgraded modern technology makes for a lighter, smaller, more power efficient, higher bandwidth, and higher multiplexing factor device than previous warm readout systems. The Prime-Cam RFSocs currently can read out up to 1024 resonators across a 512 MHz bandwidth, but expansion to a larger bandwidth is planned for the higher detector counts of the 850 GHz module [96] [97]. Each RFSoc has sufficient components to read out up to four RF networks simultaneously. Fig. 2.10 shows the packaged RFSoc with additional networking and heat sinking components.

The system is able to perform a variety of functions useful for the process of operating MKIDs. Given that changes to the grounding environment, mag-

netic pickup, operational temperature, and more can affect the frequencies of the MKIDs, the first step in operating the system is to sweep quickly through the complex transmission of the system in frequency space to find the resonator frequencies at low precision. This procedure is often called a “VNA sweep” since it mimics the behavior of a vector network analyzer (VNA). The results of some VNA sweeps with the RFSoc can be seen in Fig. 2.10. After identifying the resonator frequencies, a finer-grained sweep called a “target sweep” is performed in the region around each resonator, and this information is used to find the optimal location (i.e. frequency) for the probe tone. The signal-to-noise ratio is highest for the resonator when the resonator is driven near its bifurcation power, which is where the nonlinear part of the kinetic inductance begins to dominate, so it is often necessary to do a series of target sweeps at different probe tone amplitudes to identify the optimal tone location and amplitude. The attenuation of the system can also affect this process and could be optimized as well. Finally, once the frequency comb of probe tone locations and amplitudes is set, the RFSoc places the probe tones and streams a timeseries of data over the network. Many of these low level functions are built into the RFSoc firmware and controlled by operational software called `primecam_readout` [98]. The broader project of tuning MKIDs and operating them in the context of acquiring data from Prime-Cam will be discussed further in Ch. 4.

Beyond what has already been implemented for the Prime-Cam RFSoc, additional features may eventually be added. One useful feature implemented by the SLAC Microresonator Radio Frequency (SMuRF) electronics [99] [100] used by SO which is not yet present in the RFSoc firmware is tone tracking. The optical loading of our detectors is dominated by absorption of light from the atmosphere, even at our high-altitude sites. As we change the telescope’s elevation or scan

over a bright source, the optical loading on our detectors can change dramatically, shifting our resonant frequencies. If they shift far enough, they can shift away from our static probe tones, significantly reducing our signal-to-noise ratio or eliminating the signal entirely. Tone tracking is an algorithm that periodically adjusts the probe tone to account for changes in the resonator frequency and quality factor. While it can still fail to track the resonator if it moves far enough, tone tracking generally keeps the resonators well tuned at the cost of greater complexity in the firmware. We hope to add a tone tracking algorithm to the RFSoc firmware for Prime-Cam eventually, which could greatly simplify our tuning procedures and ensure higher quality data.

2.5 Operations Software

While operating the detectors is one of the most important aspects of the instrument, it is only one piece of a larger data acquisition (DAQ) and operations framework. For operating our cryogenic systems, we must be able to monitor and remotely control systems to read out our thermometers, pressure gauges, and helium compressors for the pulse tubes. There are systems to control power outlets and backup batteries for emergency power, the dilution refrigerator, and specialized hardware for subsystems like the EoR-Spec motor controls. These systems must be controlled by a series of computers plugged in to the site network and integrated with the broader telescope systems. The detector data and the house-keeping (HK) data, which includes measurements from all non-detector subsystems for monitoring operational health, must be packaged and transported off the computers at the telescope site for long-term storage and analysis. All of these tools should be remotely controllable through various command line and web-based ap-

plications so that remote observers can control all aspects of the observations and operations.

To accomplish these tasks, Prime-Cam will make use of much of the same hardware as the SO LATR. Its software will largely be taken from the SO Observatory Control System (SO-OCS) [10] and associated operations software packages and tools for running schedules (Nextline), alarms (Campana) [101], and remote operations (`ocs-web`, Grafana, etc.) [102]. For convenience, I will often refer to SO-OCS to include both those software tools that are formally in either the `ocs`² or `socs`³ repositories and the broader suite of SO operations tools that interface with SO-OCS.

Within the SO-OCS framework, agents are pieces of software that control individual hardware components. They often, but not always, are wrappers of an underlying driver code that communicates with the device directly via TCP, SNMP, or some other protocol. The agent exposes certain functions to the crossbar router, the open-source tool that coordinates the actions of all agents on the network. Most of the agents in the SO-OCS framework run in Docker containers for ease of deployment. The core agents for SO-OCS functionality live in the `ocs` repository on GitHub linked in a footnote above;⁴ this system is designed to be agnostic to SO and can be used for any observatory or instrument that wishes to use its framework. The agents for SO-specific hardware can be found in the `socs` repository on GitHub linked above. Since much of this hardware is the same as that used by Prime-Cam, we use many `socs` agents.

²<https://github.com/simonsobs/ocs>

³<https://github.com/simonsobs/socs>

⁴The documentation for SO-OCS is generally quite up-to-date and can be found at <https://ocs.readthedocs.io/en/main/index.html>. It contains more extensive explanations of the details of SO-OCS.

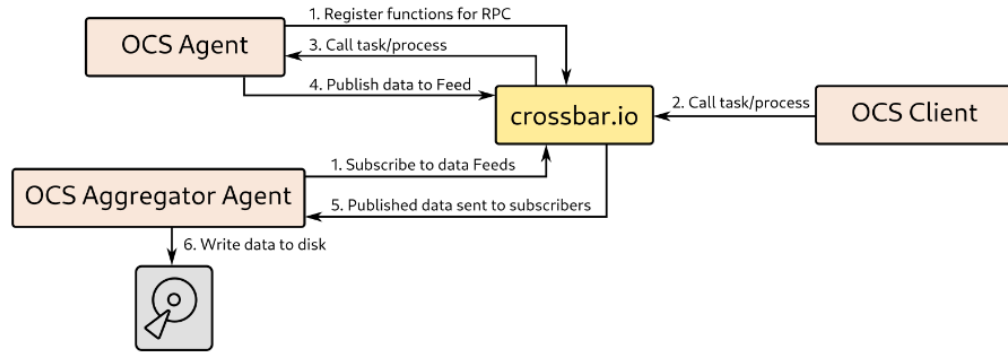


Figure 2.11: Schematic of a simple SO-OCS network in which one hardware agent is controlled by a client to acquire data that is saved to disk by the aggregator agent. From Fig. 1 in Ref. [10].

Agents are controlled by clients. Clients are often scripts run from the command line or embedded into a tool like `ocs-web` that calls the script from a web page, and they call the various functions registered by agents to coordinate the activities of one or more agents on the crossbar network. Fig. 2.11 shows an example of a simple SO-OCS network. In the example, one agent controlling one piece of hardware is started up and registers itself on the crossbar network. A client script commands it to do something like beginning the acquisition of data from the device, and the agent begins to pass that information to the crossbar router. Other agents can register themselves to receive any data passed to the crossbar router. In this example, the aggregator agent, a core SO-OCS agent that writes housekeeping data to g3 files, is listening on the system. It receives the hardware agent's data from the crossbar router and saves it to disk.

Prime-Cam will also make use of agents written for specific hardware that SO does not use. These agents live in the Prime-Cam Control System (PCS) repository on GitHub.⁵ It will also interface with the CCAT Observatory Control System

⁵<https://github.com/ccatobs/pcs>

(which I will refer to as OCS), a suite of software designed to run the core telescope systems that are shared by Prime-Cam and CHAI such as telescope control and monitoring of site power, for example. While SO has only one instrument on its LAT and therefore uses the SO-OCS framework to control the entire telescope and instrument, CCAT takes a distributed approach due to its two instruments which share observing time over the course of an observing season. Each instrument is responsible for its own operational software and for integrating that software with the telescope OCS. Though PCS strictly speaking refers to the agent-based software described here, I will often use it as a shorthand to refer to the broader set of Prime-Cam DAQ and operations software. The full suite of hardware and software for managing an observation with Prime-Cam is described in detail in [Ch. 4](#).

2.6 Mapmaking and Analysis

Once the data are acquired, the raw streams of time-ordered data (TODs) are usually filtered to remove noise and to flag high noise or incomplete data. Some analyses, like short duration transient searches, require the use of the TODs themselves, but most of the analyses for ACT, SO, and CCAT take place at the level of maps. There are a variety of procedures for combining detector pointing, the location on the sky of a given detector at a given time, with the TODs to make a map. The simplest procedure is known as the “filter and bin” method. There are various forms of this method, but they all apply some sort of filtering or template removal to mitigate noise and then average all the samples that fall into a given pixel in the map. Filter and bin mapmaking biases the results depending on your filter; while this can be mitigated with simulations, this method is not always op-

timal. An example of this method in action for the BICEP2 experiment can be found in section IV (and bias correction in section VI) of Ref. [103].

A much more computationally intensive, but unbiased and optimal, method is maximum likelihood (ML) mapmaking. See Ref. [104] for a proof of optimality and Ref. [51] for the most recent example of applying this method to ACT data, which is also a good example of the complex nature of modeling realistic telescope data. This method linearly maps the sky signal to the measured TOD in the presence of noise using some model that can account for a wide range of possible effects on the data. Solving for the “true” map of the sky then involves inverting a matrix, though the numerical methods and details are often quite complicated. If your model of the data is perfectly correct, this method returns an unbiased and optimal estimate in the sense of maximizing the likelihood of obtaining this map given the input data.

The particular science case and computational resources often dictate the best mapmaking strategy to use. For the wide field cosmology surveys of ACT, SO, and CCAT, ML maps are the standard. For smaller field surveys with Prime-Cam or the SO SATs, faster filter and bin maps are often sufficient.

In addition to cosmology-grade maps that combine all of the data for a given array in a given year, other mapmaking products for calibration or specialized analyses can be produced. For example, ACT generates shorter timescale maps for transient analyses. These are often maps containing three days worth of data (“three-day maps”) or maps containing a single continuous constant-elevation scan over a region of the sky (“depth-1 maps” - read “depth one”). The depth-1 maps are used for the time-variable polarization rotation studies in ACT Data Release 6 described in Ch. 5. Figures 8, 9, and 12 in Ref. [51] show examples of portions of

the full ACT maximum likelihood maps, while Fig. 14 shows an example depth-1 map in total intensity.

Once maps are produced, they can be analyzed with a range of techniques including power spectra analyses like those described in Sec. 1.1. Ch. 5 provides an example of such an analysis with the ACT DR6 depth-1 maps.

From the initial design through the final analysis of the data, each subsystem in the observatory must be optimized for the science goals and available resources of the project. In each of the chapters to come, we will explore that process of design and implementation for one or more of the subsystems presented here, beginning with the optical designs for CCAT and ASO.

CHAPTER 3

OPTICAL DESIGNS FOR MICROWAVE TELESCOPES

The journey of a CMB photon through our experiments begins with the warm reflective and cold refractive optics. As described in section 2.1, the crossed Dragone telescope design for the SO LAT and FYST forms its own image while correcting for coma and on-axis astigmatism, but the cryogenic lenses in the LATR or Prime-Cam are needed to reimage the telescope focus to the detector feedhorns. They also can improve the image quality along the way by compensating for other aberrations, mostly notably off-axis astigmatism.

In this chapter, we will summarize the design and optimization process for the cold optics of the CCAT 350 GHz (section 3.2.1) and EoR-Spec (section 3.3) modules and the ASO UHF modules (section 3.2.2) after an introduction to the design process and key optical and mechanical parameters in section 3.1. Section 3.4 outlines possibilities for development work in the future.

3.1 Design and Optimization

Each type of module aims to accomplish particular scientific goals, and the optical and mechanical design of the module must be tailored to those goals to maximize the potential of the module. For each of the modules presented in this chapter, we began the design process by modifying the SO UHF module design, which is identical to that of the CCAT 280 GHz module.

The design and optimization were produced with a numerical ray tracing software, Ansys Zemax OpticStudio.¹ Among its many features, Zemax allows the

¹<https://www.ansys.com/products/optics/ansys-zemax-opticstudio> provides more in-

user to define a merit function with different optical metrics as inputs and design parameters as variables. Many of the optical metrics which we use in the merit function will be defined below, and the design parameters which we allow to vary are most frequently the shapes of the lenses (radius of curvature, conic constant, and higher order polynomial corrections to shape) and the relative distances between the lenses, the Lyot stop, and the image. During this process, the position and shape of lens one is fixed because lens one reimages the primary mirror onto a common Lyot stop to control the illumination of the primary mirror for all field points in a module. The remaining two or three lenses in a given design are allowed to vary.

Once the merit function and variables are defined, Zemax can run an optimization process that adjusts the variables to minimize this merit function. This procedure can be done simultaneously for multiple configurations, which are effectively separate optical systems in the same design file. Since the optical performance of the telescope mirrors is not perfectly rotationally symmetric due to differential astigmatism from the mirror shapes (which is the dominant remaining aberration after coma correction - see the discussion in section 2.2 and Figure 6 in Ref. [66]), the performance of the optical design for a module varies depending on its location in the cryostat. We use the different configurations in Zemax to track the performance of a design in different locations to make decisions about the best place to put each module in Prime-Cam. Additionally, the telescope mirrors tilt in elevation independently of the cryostat, which means that the field observed by a given location in the cryostat varies with elevation. By measuring the optical performance in different configurations, we can compare neighboring module locations to get a sense of how performance will change as the elevation changes. In

formation about this software (as of April 2025).

theory, a design could be optimized separately for each location, producing several different designs. We did not take this path; instead, we optimized the design to work well across as many locations as possible to increase the interchangeability of our modules.

The language of configurations will appear in many of the subsequent figures and discussion, and it is worth paying attention to the reality that the apparent labeling of the configurations depends on whether you are looking at the receiver from the front (i.e. as if you were riding along with the photons as they come from the sky and enter the receiver) or from the back (i.e. as if you were tracing out rays from the receiver to the sky as if time were flowing backwards). For the Prime-Cam designs presented here, the configurations are labeled as shown in Fig. 3.1. The left side shows the numbering in the time-forward sense and the right side shows the numbering when viewing the design in the time-reversed sense (i.e. from behind the cryostat looking out towards the sky). Though the locations could still change, the left side of the figure shows where we expect each module to be placed within Prime-Cam.

There are a few other things to be aware of when interpreting results from different designs. Prime-Cam has its module locations and window orientations rotated by 30 deg relative to the SO LATR. As a result, I made a new version of the SO Zemax file containing only the six inner ring tubes (i.e. matching Fig. 3.1 and excluding the central tube and the additional outer tubes that are present in the SO LATR but not in Prime-Cam) that includes these rotations. The plots for the 280 GHz and 350 GHz modules are taken from this design, which is why they do not exactly match the patterns reported in Ref. [5] and Ref. [105]. Zemax requires all of its designs to be pointed towards zenith. In reality, we will observe at some lower




Figure 3.1: *Left:* Image of Prime-Cam from Fig. 1 of Ref. [11] showing the likely plan for where each module will go in Prime-Cam along with the corresponding Zemax configuration number. The arrow indicates the effective direction that the telescope is pointing in the Zemax design. *Right:* The six outer optics tubes in Prime-Cam from Zemax, as viewed from the image side behind the detector arrays (looking out towards the telescope mirrors). The colored lines depict rays from different fields on the sky, while the various circles and the hexagon are the different lenses, filters, image surface, and window. Each configuration is labeled with an identifying number and the coordinate system in Zemax as a reference for other figures. This figure is rotated 60 deg relative to Fig. 5 in Ref. [12] to match the convention for module locations in the left side of this figure. The arrow indicates the effective direction that the telescope is pointing in the Zemax design.

elevation, so the design files are all constructed to simulate observations at 60 deg elevation by rotating the cryostat 30 deg. Fig. 3.1 shows the pointing direction of the telescope for the design. When I display plots of optical metrics, I will display them in their real locations within Prime-Cam instead of this adjusted location in the Zemax file, which is why quantities that ought to be radially symmetric in the receiver may not look so in the way I arrange my plots. These quantities are indeed radially symmetric in reality.

The LATR has a co-rotator that can be used to adjust the clocking of the

receiver as the telescope changes elevation; Prime-Cam, however, does not. For Prime-Cam, this means that the effective location of each instrument module will change as the telescope’s elevation changes, as mentioned above. In all of the figures below showing optical performance in each of the six module locations, observations at elevations other than 60 degrees will cause the performance to be a combination of the performance between neighboring modules.

More details specific to Zemax about the optimization process and settings used for these design files can be found in Appendix B. The remainder of this section will be devoted to describing the main scientific drivers of our designs with the 280 GHz module design as an example.

The 280 GHz design shown in Fig. 2.2 was previously optimized by our SO collaborators to provide good image quality. From the perspective of geometric optics, image quality essentially refers to how well-focused a bundle of rays is when traced through the system for each point in the focal plane. Viewed from the perspective of physical optics and the theory of diffraction, image quality is usually framed in terms of deviations from perfectly spherical wavefronts. In practice, there are many ways that the image quality can be measured, including the Strehl ratio, the point spread function (PSF), the RMS wavefront error, and more. Each provides different information about the optical properties of the design.

A common metric among designers is the Strehl ratio, which folds a complex set of effects that alter image quality into a single number. The Strehl ratio is defined as the peak intensity of the point spread function divided by the peak intensity of the point spread function in the absence of aberrations

$$S = \frac{I_{PSF}^{max}}{I_{PSF,unaberrated}^{max}}. \quad (3.1)$$

See Ch. 9 of Born and Wolf’s textbook on optics [106] for a detailed study of the

mathematics behind the Strehl ratio, but it may be summarized thus: calculating the PSF resulting from diffraction and aberrations in a particular optical system involves integrating a complex exponential that takes the geometry of the situation and the wavefront deviation due to aberrations as its arguments. The numerator of this Strehl ratio is the value of this full integral at the central point of the beam, while the denominator is the same integral but with the aberration function set to zero in the argument. Zemax calculates these numerically by tracing a grid of rays and performing a series of calculations to assess the aberration function and diffraction.

Physically (and perhaps more intuitively), the PSF is a measure of the response of a system to a point source. Due to diffraction, a point source observed through a circular aperture will never form a true point image. It will always form an extended image with a shape given by a characteristic Fraunhofer pattern of a central peak followed by lower amplitude peaks and minima in intensity for increasing distance from the center of focus. In the presence of aberrations, wavefront errors will accumulate and will alter the shape of this pattern at the image plane, generally increasing the full-width half-maximum (FWHM) or the ellipticity of the central peak and possibly introducing noticeable sidelobe structures to the lower-order fringes outside the central beam. Fig. 3.2 shows an example Zemax PSF from the CCAT 280 GHz design for a single off-axis field (i.e. a single location in the focal plane or, equivalently, a single point source location on the sky) with cross-sections of the central beam, which is well approximated by a 2D Gaussian. The top right shows the theoretical diffraction pattern for one of these cross-sections with parameters chosen roughly to match the y-direction cross-section. The effect of aberrations on the focus is clear from the distortions in the x-direction cross-section that cause it to differ from the theoretical ideal. In this example, the PSFs

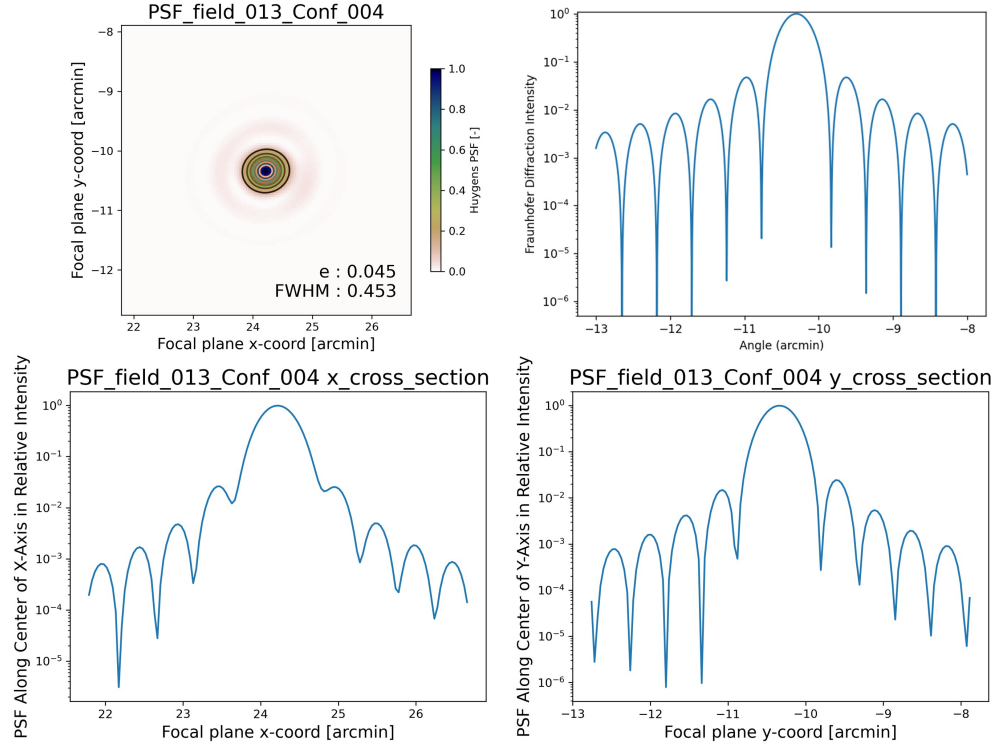


Figure 3.2: *Top Left:* An example linear-scale contour plot of a PSF for the CCAT 280 GHz design, sampled at a single field for a single configuration at an input frequency of 410 GHz during tests of this design for a possible CCAT 410 GHz module. The central beam is fit with a 2D Gaussian to extract the FWHM and ellipticity. The first of the peaks outside the main beam is visible as a faint pink ellipse. *Top Right:* The ideal 1D Fraunhofer diffraction intensity pattern with parameters roughly matched to the cross-section in the y-direction of the example PSF. *Bottom Left:* The log-scale x-direction cross-section of the example PSF along its center. Unlike the linear-scale plot in the top left, many secondary peaks are visible at much lower amplitude. *Bottom Right:* The log-scale y-direction cross-section of the example PSF along its center.

have been normalized, but if the maximum value of the PSF for this field was instead divided by the maximum value of the PSF that would result from the ideal diffraction case without aberrations, the peak value of the PSF would be the Strehl ratio for this field.

While we occasionally make use of the more complete information provided

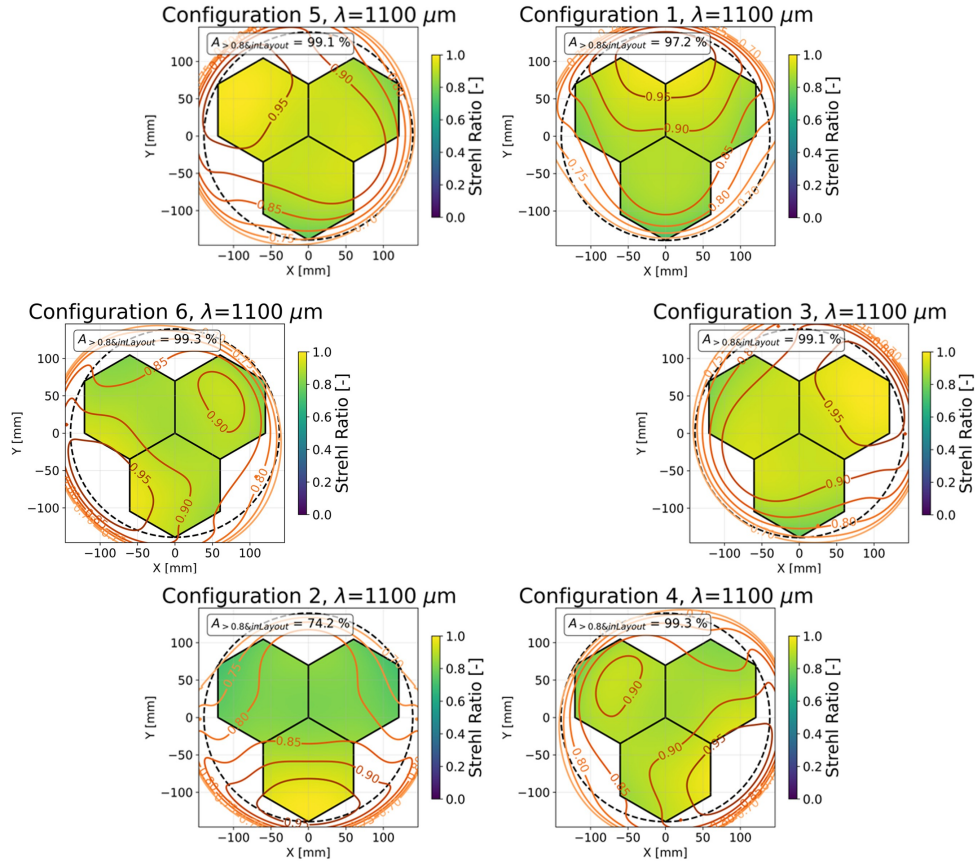


Figure 3.3: The Strehl ratio at 280 GHz across the field of view for all six possible Prime-Cam locations for the 280 GHz module. The hexagons represent the three detector arrays in a module, and the caption gives the percentage of the area of those hexagons that is above a Strehl ratio of 0.8.

by PSF analyses like this one to inform our design decisions (see, especially, the section on the 350 GHz module below), we have not studied in detail how the ray tracing methods of Zemax compare to a full treatment of the PSF using dedicated physical optics software like GRASP.² The Strehl ratio offers a way of summarizing these complex effects on image quality without getting too involved in the details. Fig. 3.3 shows the Strehl ratio at 280 GHz for the 280 GHz module optical design in the six outer tubes of Prime-Cam. The 850 GHz module is

²More information about GRASP can be found at <https://www.ticra.com/software/grasp/> (as of April 2025).

earmarked for the center of Prime-Cam, so all of the other modules were designed to function in the ring of locations surrounding the center. The three hexagons display one possible orientation for the locations of the three detector arrays in the module, and the percentage reported in the top corner shows the fraction of those three detector arrays that is above a Strehl ratio of 0.8. While higher Strehl ratio always helps improve image quality, beam size, and optical efficiency, a Strehl ratio of 0.8 is generally considered to be diffraction limited by the community,³ meaning that the optical quality is dominated by the effects of diffraction instead of aberrations. Lower Strehl values typically increase beam FWHM and ellipticity, leading to reduced resolution, more complicated beams corrections, and reduced mapping speed for point sources. For most locations in Prime-Cam at 280 GHz, the design presented above ensures that practically the entire active focal plane is diffraction-limited.

In order to interpret these and future Strehl plots correctly, there are two important notes. One is that the design generally has the highest image quality along the radius running from the center of Prime-Cam to the center of each module. In this figure, the reason that the Strehl plots do not seem to have the right line of symmetry through the center of the figure is because of the rotation to optimize the design at 60 deg elevation described above. The second is that the design also usually has the highest image quality near the center of the cryostat. The cryogenic optics invert the image on the sky, however, and a mirroring in x and y has been applied to these plots to undo that effect so that these plots show effectively “what we would see on the sky.” This inversion causes the highest Strehl regions to appear on the part of the figure furthest from the center of the cryostat.

³In section 9.3 of Ref. [106], Born and Wolf trace the origin of this tolerancing guideline back to Rayleigh (1879) and Maréchal (1947). Rayleigh implemented a rule of using systems with less than 20% degradation due to primary spherical aberration, and Maréchal studied expansions of this rule to other aberrations.




Figure 3.4: A colormap (*left*) and histogram (*right*) of the chief ray angle for one configuration in the SO UHF optical design. For the colormap, the x and y axes are normalized field coordinates. The pattern is identical for all inner ring configurations, so it will be very similar for any 280 GHz module position in Prime-Cam. All chief ray angles are below 2 deg.

Beyond high image quality, the 280 GHz module design was also optimized to produce a relatively flat focal plane. Naturally, the optics of the camera would like to converge on a curved focal surface that would require all the instrument modules to be angled slightly away from the center of the cryostat and each another. Such a design would be highly inconvenient for the cryogenic design of the camera. A combination of a tilted alumina filter that redirects the chief rays of the outer modules in Prime-Cam to run parallel to the central, on-axis module and constraints on the chief ray angle for fields across the focal plane during the lens optimization ensures that all modules share a flat focal plane. The optimization process enforces this constraint by keeping the chief rays nearly normal to the flat image surface in the design for all fields across the focal plane. Fig. 3.4 shows these results for the 280 GHz module design in the SO tube locations, though the results are effectively rotationally symmetric around the center of the receiver and apply to all Prime-Cam locations as well. The chief rays all are less than 2 deg away from perfectly normal, which is our target to minimize asymmetric illumination of

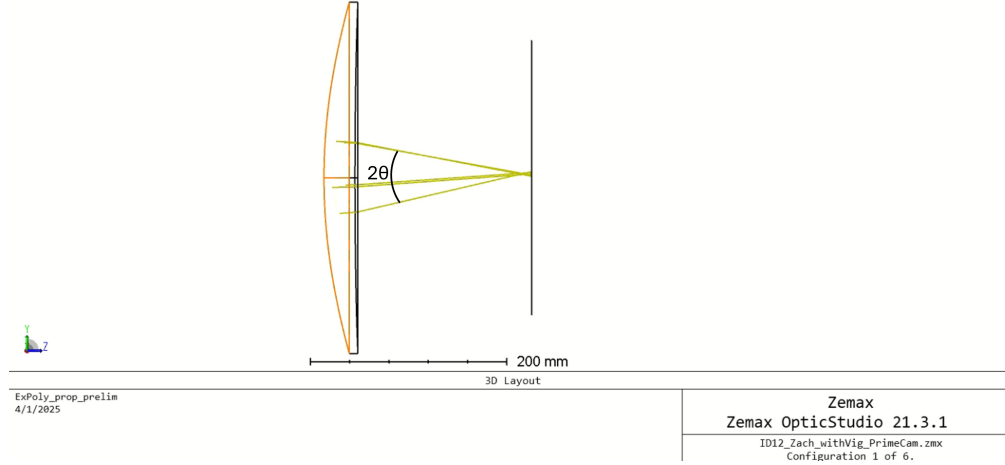


Figure 3.5: The final lens and image plane in the 280 GHz optical design, annotated to show the angle used for calculating the $F/\#$ in one direction from the two extreme pupil rays.

the Lyot stop by the reflections off of the feedhorn arrays.

A third major optical property of the design is maintaining nearly constant illumination of the Lyot stop across the focal plane. The illumination of the stop translates to the illumination of the primary mirror. By ensuring that the each field sees roughly the same illumination, the design minimizes beam and mapping speed variations across the focal plane. One way to quantify this illumination is by measuring the working $F/\#$ of the system at the image plane. The $F/\#$ is the ratio of the focal length to the diameter of the effective aperture, and we calculate it in Zemax from the angle of convergence of the extreme (sagittal or tangential) rays of the pupil at the surface in question using the formula

$$F/\# = \frac{f}{D} = \frac{1}{2 \tan(\theta)}, \quad (3.2)$$

where θ is defined as half the angle between the two extreme rays in one pupil direction (e.g. the rays that pass through the max and min x coordinate of the pupil with $y=0$) as shown for a single field in our Zemax design in Fig. 3.5. The overall $F/\#$ is calculated by averaging the results for the x and y directions. Fig. 3.6




Figure 3.6: A colormap (*left*) and histogram (*right*) of the $F/\#$ at the image surface for one configuration in the SO UHF optical design. For the colormap, the x and y axes are normalized field coordinates. The pattern is identical for all inner ring configurations, so it will be very similar for any 280 GHz module position in Prime-Cam. The mean $F/\#$ is around 2 for these optics, and all but the most extreme edges of the focal plane where there are few detectors have an $F/\#$ within 0.1 of the mean.

shows an example of this quantity calculated across the focal plane for the 280 GHz module design in the SO tube locations, though the results are effectively rotationally symmetric around the center of the receiver and apply to all Prime-Cam locations as well. The values cluster around the mean of two for most of the focal plane, a value fixed by our desired detector spacing and feedhorn beam coupling. See Ref. [75] and section 2.3 of Ref. [68] for further discussion of the tradeoffs between different detector packing schemes and telescope $F/\#$ for various science cases. There is higher variance at the edges, but there are also very few detectors located at the edges because of our hexagonal array design.

In addition to ensuring that each design has good image quality, low chief ray angle, and consistent $F/\#$, there are also a series of mechanical constraints enforced during our optimizations. The most important of these are ensuring that the optical path was short enough to fit inside of the cryostat and that there is no vignetting or overflow of rays past the bounds of the optical and mechanical

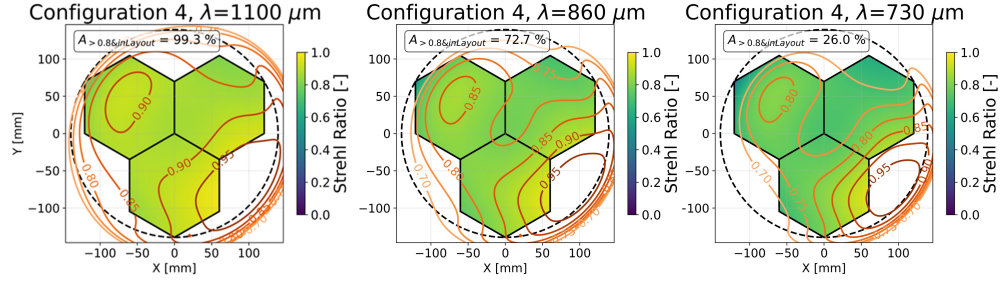


Figure 3.7: Strehl ratio as a function of frequency for one configuration of the 280 GHz module design at 280 GHz ($1100 \mu\text{m}$), 350 GHz ($860 \mu\text{m}$), and 410 GHz ($730 \mu\text{m}$). Without additional optimization of the design at each frequency, the Strehl ratio decreases at higher frequencies.

elements. In the designs presented below, we also enforced an additional constraint that the sag on each surface of the lens should be no more than 14 mm to make cutting the metamaterial AR coating easier. These optical and mechanical constraints comprise the most important tools for evaluating the success of the design optimizations that will be presented next.

3.2 Broadband Imaging Modules

There are two major classes of modules for Prime-Cam: broadband imaging modules, like the 280 GHz module, and spectrometer modules like EoR-Spec. Depending on the spectrometer technology, the optical design considerations may be considerably different for the spectrometer modules, so we will finish considering all of the broadband modules first before turning to look at the design process for the EoR-Spec optics.

3.2.1 350 GHz Module

The broadband imaging modules are all rather similar in operation, and their optical designs are fairly broadband in the sense that the optical performance across many of our main SO and CCAT frequency windows (i.e. 90 GHz, 150 GHz, 220 GHz, 280 GHz, 350 GHz, etc.) is relatively similar for modules in similar locations in the cryostat. The performance is not identical, though, and it diminishes at higher frequencies for a given optical design as shown in Fig. 3.7. Higher frequencies correspond to smaller diffraction-limited beam sizes, which means that wavefront errors are closer in magnitude to the relevant wavelength and beam size. As the wavefront errors make up a larger fraction of a wavelength at smaller wavelengths (see Fig. 3.8), they contribute more to interrupting the ideal diffraction pattern and reduce the Strehl ratio, especially for off-axis field points.

Put concretely for the 280 GHz module optical design, the Strehl ratio exceeds our target of 0.8 for more than 90% of the detectors in the focal plane in the 280 GHz band. In the 350 GHz band, the same design only covers 70-80% of the detectors with a Strehl ratio above 0.8 for the locations in Prime-Cam in which we are considering placing the 350 GHz module, though more than 95% of the detectors would still see a Strehl ratio above 0.7.

In order to try bringing the optical quality back up to diffraction-limited across the full focal plane, we attempted to reoptimize the 280 GHz module optical design to yield better performance at 350 GHz. To save cost and time by accommodating previously machined parts and filters, we operated under the significant constraint that the mechanical design should match that of the 280 GHz design. This mechanical restriction meant that we could only alter the shapes of the lenses without changing their relative location. We also sought to accommodate the request of

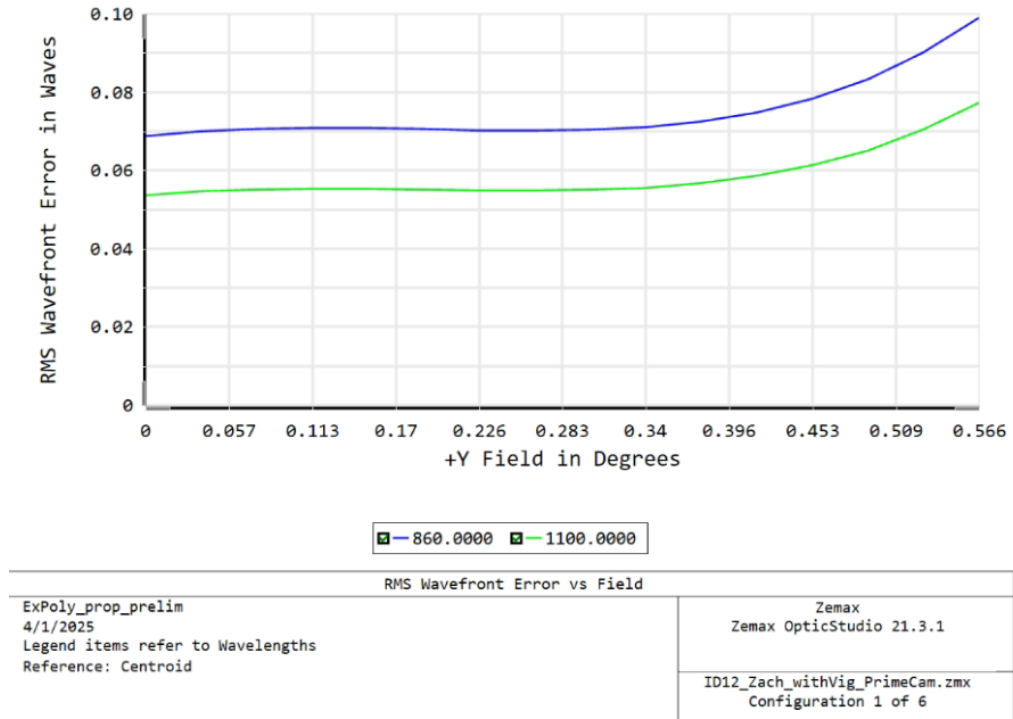


Figure 3.8: The RMS wavefront error in units of number of waves for one configuration of the 280 GHz module optical design as a function of angle in the y direction away from the center of the module. In addition to the wavefront error increasing for fields further off-axis, the RMS wavefront error is a greater proportion of a wavelength at 0.86 mm compared to 1.1 mm. The overall wavefront error is effectively constant with frequency in units of μm , so the difference in number of waves only comes from the difference in wavelengths.

the team that machines the AR coatings mentioned above that the sags of each surface be kept under 14 mm. This constraint was not enforced for previous designs like that of the 280 GHz module. The other design metrics were largely the same; the chief difference in the optimization process was the choices for the lens shapes for lenses two and three as well as using 0.86 mm as the main optimization wavelength in Zemax’s optimization wizard.

In the 280 GHz module optical design, all lens surfaces are either flat or aspheric. Aspheric lenses include all radially symmetric lenses that are not spherical,

which translates to having a surface defined by a radius of curvature and a nonzero conic constant. We also allow for small polynomial perturbations to the surface of the lenses to help cancel aberrations, though we only make use of the even polynomial terms to preserve radial symmetry. The full equation for the shape of our aspheric lenses is

$$z(r) = \frac{cr^2}{1 + \sqrt{1 - (1 + k)c^2r^2}} + a_2r^2 + a_4r^4 + a_6r^6 + a_8r^8, \quad (3.3)$$

where $c = 1/R$ is related to the radius of curvature, R , and k is the conic constant. In the development of the 280 GHz design, it was found that there was little benefit to going to more than three polynomial terms (see section 3.2.2 in Ref. [5]). We experimented with both three and four terms and found similarly that it did not make a large difference; some of our designs have three and some have four. For the Zemax optimization tool, it may be advantageous to use one fewer optimization variable per curved lens surface for speed and convergence to a good solution, but this impression was not rigorously tested.

For the modifications to lenses two and three for possible 350 GHz module designs, the plano-convex and convex-concave lenses of the 280 GHz module design were converted to convex-convex lenses before beginning the optimization process, in part to reduce the sag on any given surface below 14 mm by transferring some of the total sag from one side of the lens to the other. Since Zemax measures distances between elements from the back of one surface to the front of the next, additional lines were added to the merit function to ensure that the mechanical clamping location for the lens remained the same as the sag on each side of the lens changed.

In addition to considering all-aspheric designs, we also considered more significant lens changes, namely swapping one or both of the aspheric lenses for a

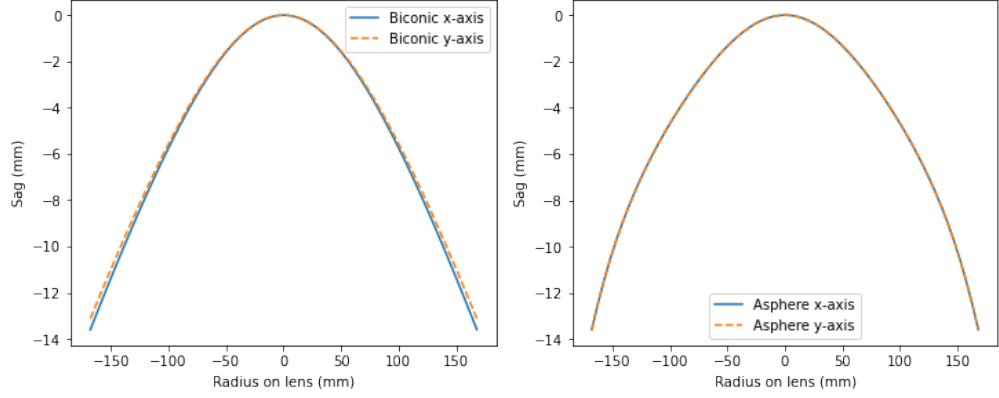


Figure 3.9: The sag along the x and y axis cross-sections for the back of lens two in a CCAT 350 GHz test design that used a biconic lens (*left*) and one that used an aspheric lens (*right*). The difference in shape along two orthogonal directions is characteristic of a biconic lens.

biconic lens. As the name suggests, biconic lenses can have different conic sections for shapes along two orthogonal principal axes. This leads to an equation for the surface that looks like

$$z(x, y) = \frac{c_x x^2 + c_y y^2}{1 + \sqrt{1 - (1 + k_x)c_x^2 x^2 - (1 + k_y)c_y^2 y^2}} + a_{2,x}x^2 + a_{2,y}y^2 + a_{4,x}x^4 + a_{4,y}y^4 + a_{6,x}x^6 + a_{6,y}y^6 \quad (3.4)$$

with separate radii of curvature, conic constants, and polynomial corrections for each principal direction. In our designs, we always orient one of these axes to point towards the center of the receiver with the other perpendicular to it. For the ASO design presented below, up to the sixth order polynomial terms were used. These were found to be rather small corrections, so many of the CCAT designs that involved biconic lenses did not use the polynomial terms at all, relying instead on the core biconic shape alone. Fig. 3.9 shows how a biconic lens has a different cross-sectional shape in two perpendicular directions while an aspheric lens has an identical shape.

Biconic lenses correct astigmatism well since astigmatism involves a difference

in focus in two perpendicular planes of rays. Biconic lenses naturally can focus two perpendicular planes of rays in different locations to eliminate the mismatch. Despite this appealing feature, biconic lenses have not been used on previous CMB experiments due to the extra complexity for machining and AR coating them as well as the additional alignment challenge to maintain the correct clocking of the cryogenically-cooled lens, a problem that does not exist for the radially symmetric aspheric lenses. Given the tight mechanical constraints for the CCAT 350 GHz design and that ASO was developing procedures for using biconic lenses for their far off-axis UHF modules as described in section 3.2.2, we decided to consider these options. The choice to use the same mechanical design ruled out using a fourth aspheric lens, another straightforward way to improve the optical performance (albeit with greater cost and complexity).

We considered over a dozen different designs for the CCAT 350 GHz module. Tests included all-aspheric, lens two biconic,⁴ and lens two and three biconic options. Among the best of each category, further variations were tested with different combinations of lens parameters allowed to vary including with and without higher order polynomial corrections and with varying levels of complexity for the shapes of the surfaces that had previously been planar in the 280 GHz module design. We often started with simpler surfaces for fewer optimization variables to optimize the basic design, then added additional variables to see if a basic design solution could be marginally improved. We also tried using different combinations of initial values for variables and for the merit function targets to probe different parts of the highly multidimensional optimization space.

⁴Generally, we expected that making the surface closest to the Lyot stop the biconic surface would be best for fixing astigmatism. This is the back surface of lens two in this design. Quick tests with making lens three biconic and lens two aspheric did not seem to be an improvement, so I mostly considered designs in which only lens two is biconic.

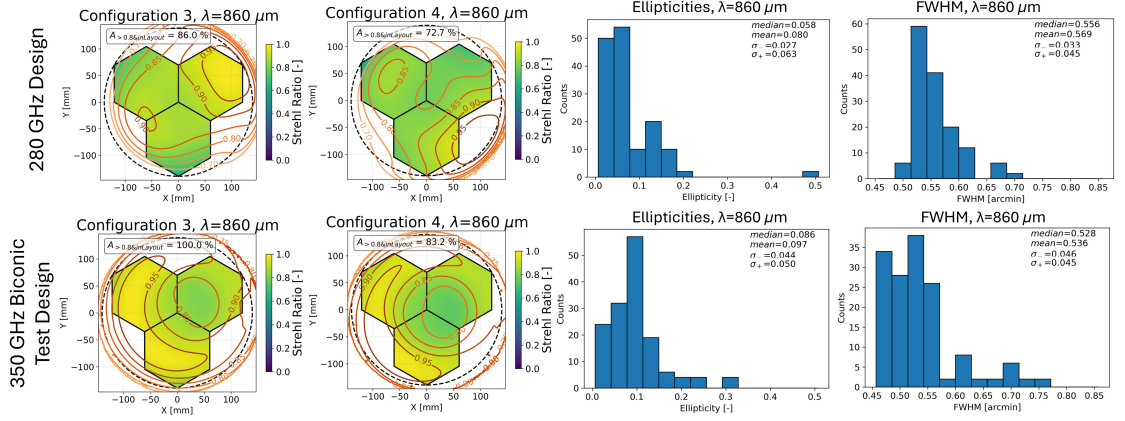


Figure 3.10: *Left:* A comparison of Strehl ratios for two possible Prime-Cam locations of the 350 GHz module between the same design as the 280 GHz module (top row) and the best single biconic 350 GHz design (bottom row). The Strehl is indeed better for the 350 GHz biconic design. *Right:* Comparisons of the ellipticity and FWHM of the PSFs for the same two designs across all possible locations in Prime-Cam (though note that these plots are for PSFs calculated in the SO locations; they should be similar in distribution to the Prime-Cam ones). The higher Strehl for the biconic design is driven by the slightly better FWHM, but the higher ellipticity contributed to our choice to stick with the 280 GHz design. This is an updated version of Fig. 4 in Ref. [12] to add more bins and make the bins match across designs for easier comparison.

The best overall performance of any design that we considered resulted from changing the back surface of lens two to a biconic lens. As presented in Ref. [12], this design would have resulted in a roughly 10% increase in the number of diffraction-limited detectors compared to the best all-aspheric lens design, but an analysis of the PSFs in Zemax revealed that this improvement in Strehl ratio was accompanied by larger scatter in the ellipticity of the PSF. Example PSFs showing this effect can be found in Fig. 4 of Ref. [12]. While the PSF size and ellipticity improved for the biconic design in the center of the focal plane, the middle and edges of the focal plane became much more elliptical.

The right side of Fig. 3.10 shows histograms of the full width at half maximum




Figure 3.11: Comparison of the 50th (lowest bank of data points), 80th (middle bank), and 95th (top bank) percentile encircled energy radii for 25 fields in each possible module location for a 350 GHz aspheric design (left) and a 350 GHz biconic design (right). Each colored line represents the fit for a different module. Generally, fields with higher Strehl ratio have lower encircled energy radii (a decreasing slope for the linear fit), but the biconic design often bucks this trend and has greater scatter in the encircled energy for the fields. This is likely due to its distorted central beam shapes. Figure credit: Isla Steinman.

(FWHM) and ellipticity for each design. These histograms use data from all six outer positions in Prime-Cam with the PSFs sampled at twenty-five regularly spaced locations across the focal plane at each position. While the biconic design did reduce the worst outliers in ellipticity, it also caused the median ellipticity to increase while reducing the median FWHM compared to the all-aspheric design. This reduction in FWHM is what caused the Strehl ratio coverage to increase, though the extra ellipticity limited the optimization from achieving even higher Strehl ratios across the focal plane. These effects were true in the aggregate for all six possible module locations and for each possible module location individually.

Studies of the encircled energy, the radius that encloses a fixed fraction of the beam, for these biconic designs undertaken by an undergraduate researcher in our group revealed that the additional ellipticity in these designs led to counterintuitive behavior. Usually, fields that have a high Strehl ratio also have a lower encircled

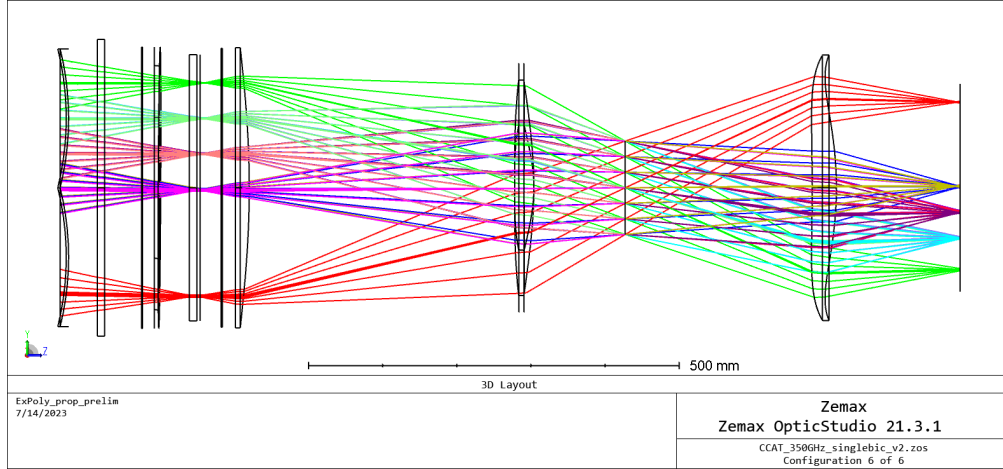


Figure 3.12: Ray trace for the best single biconic 350 GHz optical design. The back surface of lens two is the biconic surface. The back of lens three took on an unusual curvature in all of these designs, likely due to the constraints on the lens location and sag.

energy radius than those with lower Strehl ratios, regardless of the energy cutoff, which naturally follows from fields with higher Strehl ratios having central beam lobes closer to the diffraction-limited Airy disk. For the biconic designs, though, some of the highest Strehl ratio fields also had the highest encircled energy radii, indicating that the main beam has been distorted while minimizing aberrations.

The increase in the median ellipticity and the variation of the ellipticity across the focal plane would have made it much harder to correct for the effects of the beam in the maps made with this module. The shape of the third lens in these designs also preferred to take on a funky shape as seen in Fig. 3.12; it could be forced to maintain a more conventional convex-convex shape, but only at the expense of image quality near the outside of the focal plane. This design could have been produced, but the cost and tolerance of the machining may have been worse.

No single design perfectly accomplished everything that we wanted. Roughly

similar increases in ellipticity were seen in the designs with two biconic lenses. These designs had marginally better Strehl than the single biconic designs, but the extra complexity of appropriately machining and aligning two biconic lenses without a large and clear gain in performance deterred us from pursuing this option further at this time. Likewise, it was possible to generate a higher Strehl and lower ellipticity in a single biconic lens design but only if the radius of the illuminated focal plane increased beyond the size of our detector wafers. This would have generated unacceptable reflections and loss of optical efficiency.

In the end, we decided to reduce the risk of additional systematic effects on our data and the risk of schedule delay from developing new metamaterial AR coatings for the biconic lenses by adopting the 280 GHz module optical design for the 350 GHz module as well. The lenses are now being fabricated for this module. Given that we were optimizing new designs for this module under tight constraints from using the same mechanical design as the 280 GHz module, it is quite possible that improvements could be made to the optical design for a future upgrade to this module if necessary. In the short term, however, this choice reduces the complexity of this module and mitigates a substantial amount of risk in the design, module assembly, and data analysis.

By keeping the same design as the 280 GHz module at 350 GHz, a larger percentage of our detectors will have a Strehl ratio below our target of 0.8, as can be seen in Fig. 3.13. Generically, the main effect of this is a wider main beam for the detectors in areas with lower Strehl, though our beam shapes should still be relatively uniform. Lower Strehls can also mean that we see less total power on our feedhorns, leading to a lower optical efficiency, or that more of our total power is shifted into our sidelobes instead of the main beam. However, the consistency of

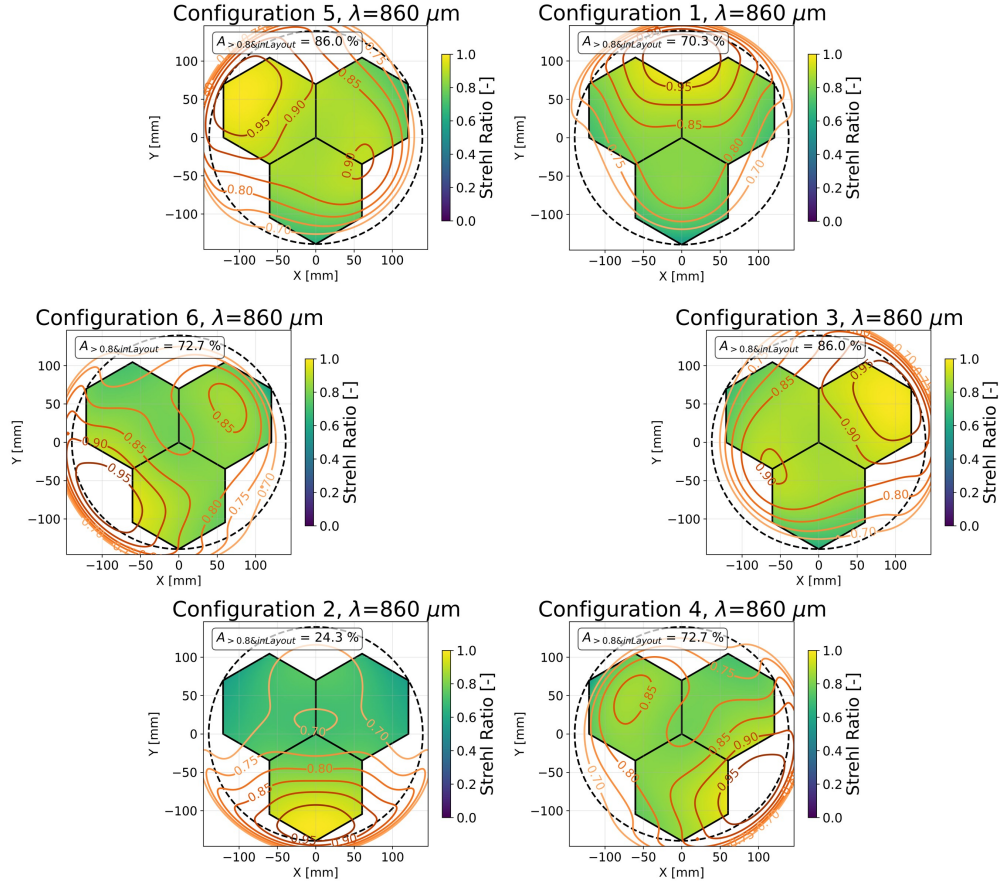


Figure 3.13: The Strehl ratio at 350 GHz (from the same design as the 280 GHz module) across the field of view for all six possible Prime-Cam locations for the 350 GHz module. The 350 GHz module is nominally slated to go into configuration four, but it could be switched with the 280 GHz module in configuration three to improve performance.

the encircled energy analysis suggests that the dominant effect is only the widening of the beam instead of these other effects, which should not significantly affect the mapping speed for diffuse sources like the CMB and foregrounds. It will degrade our point source mapping speed, however, and it will be important to characterize the effect of reduced Strehl on performance when this module is deployed for scientific observations in Prime-Cam.

3.2.2 Advanced Simons Observatory

The analysis of PSFs in Zemax for the CCAT 350 GHz designs also contributed to the study of designs for the Advanced Simons Observatory’s UHF (220 and 280 GHz) modules. The initial deployment of the LATR for SO only had funding to fill seven of the thirteen optics tubes. Six of these were slotted into the central tube or the inner ring of six tubes because the optical performance in the initial SO design is better there. Most of the aberrations that remain in the telescope optical design scale with the distance off-axis, so the image quality decreases considerably for the outer parts of the LATR, especially at higher frequencies. Fig. 4 of Ref. [5] shows a drop-off in the inner ring Strehl ratios themselves at 280 GHz, and it does not even bother to show the Strehl ratio at 280 GHz for the outermost six tubes (showing 90 GHz instead) because the optical quality is so low there at the higher frequency. There will be at least one UHF tube in the outer ring for ASO, so a new design was needed to improve the image quality.

In addition to a single biconic design built upon the previous SO design work with similar mechanical constraints, we also considered the possibility of using one of the biconic designs that I was producing for CCAT, which led to comparisons of the PSFs between the designs. We found that the biconic design produced by our SO collaborators (mostly Simon Dicker, but with input from Patricio Gallardo and others) greatly improved the image quality, PSF FWHM, and PSF ellipticity in the outer ring compared to the baseline SO design (see Fig. 3.14) and performed equally well or better to the CCAT 350 GHz biconic design (including slightly better ellipticity) for the remaining inner ring UHF tube as well. This process showed that it was still possible to produce biconic designs that significantly improve the image quality for off-axis modules in receivers like the LATR and Prime-Cam.

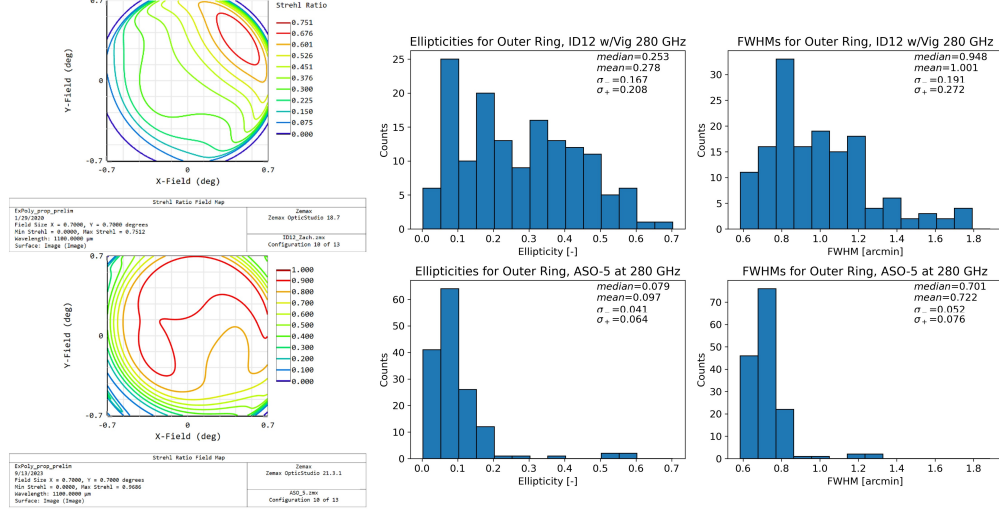


Figure 3.14: Comparisons of the Strehl ratio (*left*) and PSF distributions (*right*) for the ASO biconic design (“ASO-5”, bottom row) and the original all-aspheric SO design (“ID12”, top row) in the outermost ring of the LATR. The Strehl plots have a different colorbar scale – the ASO biconic design is considerably better than the original SO design for this outer ring module location. The same is true for the PSF ellipticities and FWHMs, showing the great utility of biconic lenses for correcting astigmatism far from the optical axis of the receiver.

It matters substantially more to use biconic lenses to correct the main telescope design for the outer ring of the LATR, and the PSF analyses showed an acceptable level of beam ellipticity for the ASO biconic design. The returns are more marginal for the inner ring, which is why a biconic design was not as helpful for the 350 GHz module in Prime-Cam. ASO is proceeding with this biconic design for its higher frequency modules.

3.3 Spectrometer Modules: The EoR-Spec Module

The Epoch of Reionization Spectrometer module necessarily departs significantly from the broadband mechanical and optical designs presented above in order to accommodate the Fabry-Perot interferometer (FPI) located at the Lyot stop of the

system [107] [61] [108]. This moderate resolving power ($R \sim 100$) spectrometer will enable line intensity mapping studies of a variety of spectral lines that trace star formation in the late epoch of reionization ($z \sim 3.5\text{--}8$). Fundamentally, the FPI consists of two reflective surfaces set parallel to one another, forming an optical cavity called an etalon with resonant frequency set by the distance between the mirrors. All rays entering the cavity at a fixed angle, regardless of their origin, will form a single interference annulus on the focal plane. For monochromatic light, as the angle of incidence varies, higher order peaks in the interference pattern (called “fringes”) become visible. Similarly, if the angle of incidence and etalon distance are kept constant, higher order fringes will appear at different radii on the focal plane for different frequencies of light. EoR-Spec will take advantage of this property to measure two bands (roughly 210–315 GHz and 315–420 GHz) simultaneously with the second and third order fringes of the FPI. A small set of motors allow the distance between the mirrors to be altered, which shifts the narrow frequency band that detectors at a given radius will see. Normal observations will involve mapping the target at one etalon distance, then stepping to the next distance and mapping again. By mapping at many different etalon distances, the entire frequency bandwidth is built up over time for the entire field. See section 9.6.1 of Ref. [70] for more information about the basic operational principles of an FPI.

The operation of the FPI constrains the optical design. Because we are going to observe two fringes at a range of etalon distances, we need to have good optical quality over the range 210–420 GHz. In practice, we optimized and assessed the design at the central frequencies for the two bands (280 GHz and 350 GHz), though the optical quality decreases as we go to the highest frequencies in the upper band as we saw above.

The other major constraint on the optical design is that it is important to control the dispersion of the input angle to the FPI to preserve the finesse and resolving power of the instrument. For our broadband module designs, the rays for a single field generally are converging as they pass through the Lyot stop (as, for example, in Fig. 2.2 or Fig. 3.12). This phenomenon is undesirable for the FPI. Converging angles of incidence could cause the radius of the fringes on the focal plane to be too large, missing the relevant detectors. Perhaps more importantly, scatter in the angles of incidence from different fields on the sky can alter the interference pattern of each fringe, reducing its frequency resolution. As a result, FPI performance requires all the rays from a given field to be as close to parallel as possible when entering the Lyot stop.

To account for these constraints, we needed to make a new optical design that has well-collimated rays for all fields at the Lyot stop while still maintaining good optical quality at the focal plane. Beginning from the 280 GHz module design, early design work showed that shifting lens two closer to lens one improved the collimation of all fields. It also showed that adding a fourth lens was necessary to refocus the beam sufficiently in roughly the same mechanical length, a constraint imposed by the cryostat size.

During the optimization process of this initial design, we had thought we would need a biconic lens to achieve diffraction-limited image quality across the detector footprint for our full bandwidth, and such a design was presented in Ref. [109]. Compared to the optimization process for the broadband modules, additional constraints were added to the Zemax optimization merit function to enforce a highly collimated beam at the Lyot stop. Specifically, we calculate the $F/\#$ as described in section 3.1 at the Lyot stop. By Eq. 3.2, the $F/\#$ is related to the angle of inci-




Figure 3.15: Ray trace for the final EoR-Spec optical design showing the four lenses and the highly collimated beam at the Lyot stop for each field. Every color shows a series of rays from the same field on the sky and focused at the detector array.

dence of the extreme rays at a surface, meaning that the $F/\#$ scales inversely (in the small angle limit) with the incidence angle. For a perfectly collimated beam, the $F/\#$ would go to infinity. In our design, the constraints aimed to achieve an $F/\#$ greater than 100 for fields across the focal plane, a conservative target based on calculations of the level of collimation needed to preserve the FPI frequency resolution done by the team developing the FPI. This target corresponds to a half-angle at the stop around 0.3 degrees.

After learning from the optimization of the 350 GHz module and ASO UHF designs that biconic lenses may not be as helpful as expected for the focal plane area of Prime-Cam, especially given that we have a fourth lens to improve optical performance in the EoR-Spec design, we revisited simpler all-aspheric designs for EoR-Spec. In part by using Zemax’s real ray aiming feature to calculate the collimation of the beam near the Lyot stop more accurately and by using a more hands-on approach to help the optimization process avoid suboptimal local minima in the merit function, we found an all-aspheric design that outperformed the single

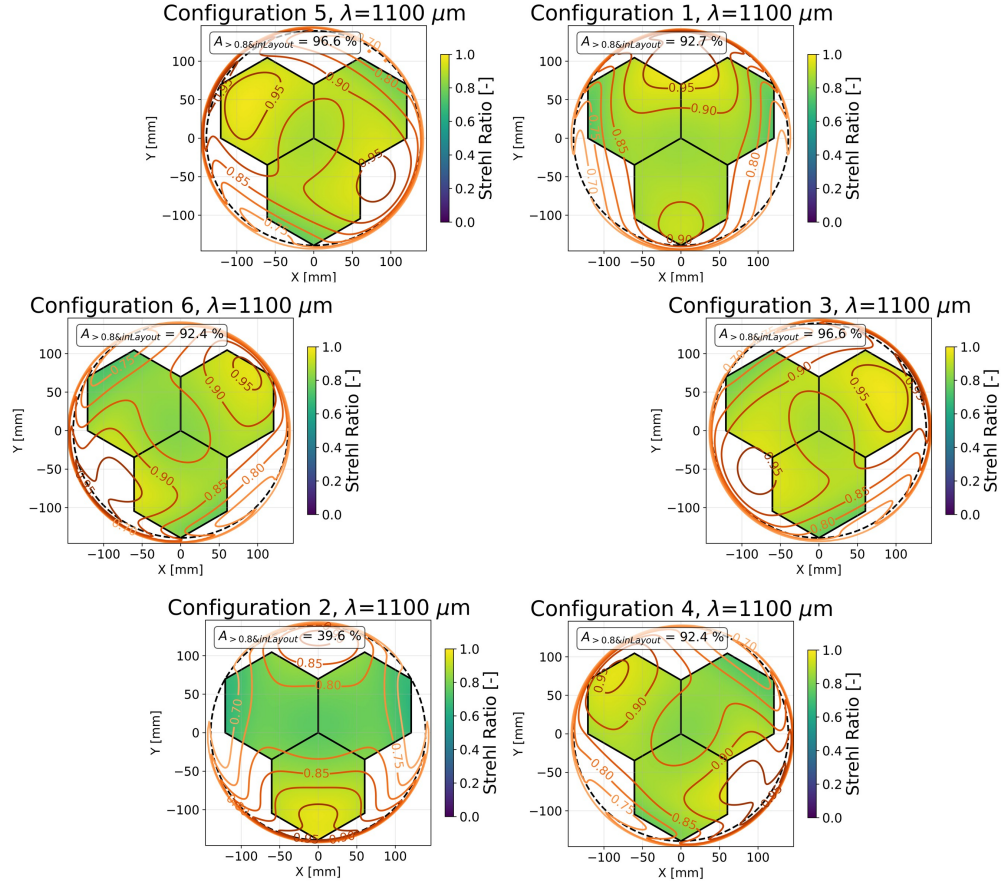


Figure 3.16: The Strehl ratio at 280 GHz across the field of view for all six possible Prime-Cam locations for the EoR-Spec design. EoR-Spec is nominally planned to go into configuration five. For most locations in Prime-Cam, all three arrays are diffraction-limited.

biconic design. The all-aspheric design is presented in Ref. [12], and the raytrace is shown in Fig. 3.15. The Strehl results at 280 GHz (Fig. 3.16) and 350 GHz (Fig. 3.17) are re-presented here (with adjusted module locations relative to Fig. 5 of Ref. [12] to be consistent with the planned module locations presented above) along with the maps of the $F/\#$ at the Lyot stop (Fig. 3.18). The planned location for EoR-Spec is configuration five. Two of the three arrays will have frequency bands centered at 280 GHz and one will be centered at 350 GHz. Sufficient focal plane area is diffraction-limited at the relevant frequencies for nearly all detectors, and a majority of the focal plane area that is filled with detectors has an $F/\#$ at the




Figure 3.17: The Strehl ratio at 350 GHz across the field of view for all six possible Prime-Cam locations for the EoR-Spec design. EoR-Spec is nominally planned to go into configuration five. At least one array is diffraction-limited for most locations in Prime-Cam.

stop greater than the target of 100. The regions that do not reach that conservative target generally still have values greater than 50, which would correspond to a half-angle of 0.6 deg and should still be quite good for our purposes. Fig. 3.19 shows that the chief ray angles still fall below our target of two degrees except for a small number of pixels at the very edge of the focal plane where there are almost no detectors. Fig. 3.20 shows that the working $F/\#$ at the image plane is approximately 10% higher than 2.0, the average working $F/\#$ for the 280 GHz module design, but with a smaller spread, thus achieving more uniform illumination




Figure 3.18: The $F/\#$ at the Lyot stop across the field of view for all six possible Prime-Cam locations for the EoR-Spec design. This can be a tricky measurement to visualize. From each field on the sky (which maps to fields on the focal plane), rays are traced to the extreme locations of the Lyot stop. These four rays are used to calculate the $F/\#$ using the same procedure mentioned above, and this serves as a proxy for the collimation of the beam at the Lyot stop. We targeted values above 100, which show up as yellow on the truncated scale used here. EoR-Spec is nominally planned to go into configuration five.

of EoR-Spec’s detectors.

Once the design met our image quality and other optical and mechanical specifications, slight modifications were then made to the central thicknesses of the lenses (but not the optically active curved surfaces) to match requirements for the machining of the mounting structures and the AR coating. This process should not affect the optical performance, and the details are described in section B.1 of




Figure 3.19: The chief ray angle across the field of view for the planned location in Prime-Cam for the EoR-Spec module. The results are similar for other locations. All angles fall below the target of 2 deg except a small number at the edges where there are almost no detectors.




Figure 3.20: The $F/\#$ at the image across the field of view for the planned location in Prime-Cam for the EoR-Spec module. The results are similar for other locations. The median is slightly higher than the 280 GHz module, but the scatter across the focal plane is reduced, so the overall illumination should be more uniform.

the Zemax appendix. These lenses are currently being machined and AR coated for deployment in the coming years in Prime-Cam. The final lens parameters can be found in Table 3.1.

	Type	$r = 1/c$	k	a_2	a_4	a_6	a_8
L1 front	Flat	-	-	-	-	-	-
L1 back	Convex	-1510.610	12.097	-5.877e-5	1.769e-9	-1.310e-15	1.995e-19
L2 front	Convex	2227.425	-95.015	8.029e-7	2.695e-11	-1.422e-14	1.823e-19
L2 back	Convex	-1753.243	-0.863	-4.925e-5	2.310e-11	-2.316e-14	5.544e-20
L3 front	Convex	1164.043	1.861	-7.222e-7	-7.891e-10	-2.054e-14	2.196e-19
L3 back	Convex	-1975.868	17.968	6.735e-6	7.305e-10	-2.880e-14	3.531e-19
L4 front	Convex	1335.318	68.672	2.994e-7	-7.368e-9	-2.309e-14	-8.241e-18
L4 back	Flat	-	-	-	-	-	-

Table 3.1: The lens parameters for the final, as-machined version of the EoR-Spec optical design. Diameters and radii are in millimeters. Higher order coefficients have units that are powers of millimeters and can be obtained from looking at the relevant powers of equation 3.3. The clear diameter for lens 1 is 378 mm, and its center thickness from vertex to vertex is 18.7 mm. For the other lenses, those same two parameters are 336 mm and 21.69 mm (L2), 336 mm and 25.38 mm (L3), and 310 mm and 14.12 mm (L4). Note that these thicknesses are the ones from Zemax and do not include the AR coating thickness, so they differ slightly from the lens drawings.

3.4 Future Designs

With the designs presented here and the 850 GHz module design presented in Ref. [110], the optical design for four of the seven initial modules of Prime-Cam have been finalized and sent for fabrication. Of the remaining modules, one is likely to be another spectrometer, whether a second EoR-Spec module or a spectrometer-on-a-chip module. Either way, this spectrometer will likely reuse an existing design: the EoR-Spec design (for a second EoR-Spec module) or the 280 GHz design (for a spectrometer-on-a-chip module). The baseline plan for Prime-Cam also involved a 220 GHz module, which could reuse the 280 GHz module design, and a 410 GHz module, though the final decisions about these modules depend on funding and which frequencies could maximize the scientific output of CCAT at the time that funding is procured for those modules.

If a 410 GHz module is eventually produced, its image quality would suffer substantially more than the 350 GHz module if it uses the same design as the 280 and 350 GHz modules, as highlighted in Fig. 3.7. An appealing alternative that would require little additional work would be to consider adopting the ASO biconic design, an improvement that could also be valuable for any upgrades to the 350 GHz module or a spectrometer-on-a-chip module. A more detailed study of the ASO biconic design at 410 GHz would be required to verify that it meets all of the desired optical specifications at this frequency.

When considering optical designs beyond the initial deployment of Prime-Cam, a whole host of interesting ideas for alternative cameras in a FYST-style crossed Dragone telescope have been explored in CMB-S4 design studies (for example, the 85-module camera presented in Ref. [111]), and it seems entirely possible that Prime-Cam could eventually be upgraded to such a camera for a future generation of CCAT. Another appealing alternative would be to populate a larger camera with many spectrometers, given that CCAT’s outstanding site conditions reduce interloper atmospheric lines and that this receiver could study astrophysics that could not be examined by ASO. Such a camera would likely want to explore updated designs involving biconic lenses for improved off-axis performance over a wider range of frequencies, including possibly pushing to higher frequencies than EoR-Spec’s 420 GHz.

In this chapter and the associated appendix, we have presented the optical designs and optimization process for the 280 GHz, 350 GHz, and EoR-Spec modules for CCAT, all of which will operate in diffraction-limited conditions for the large majority of detectors. Studies of designs involving biconic lenses for the 350 GHz module and ASO UHF outer ring modules found that biconic lenses were necessary

at 280 GHz and higher frequencies for the far off-axis locations in the SO LATR, but these designs were not necessarily a significant improvement in Prime-Cam's footprint due to higher beam ellipticities when trying to maintain identical or similar mechanical designs to the 280 GHz module. Each of these designs have been or will soon be fabricated, laying the first stone in the foundation for making Prime-Cam and the ASO LATR successful millimeter wave receivers.

CHAPTER 4

PRIME-CAM OPERATIONS AND DATA ACQUISITION

The hardware and software to operate the instrument and acquire data remotely will be critical for the success of Prime-Cam on FYST. While a team based in the nearby town of San Pedro de Atacama will be available to perform maintenance and recover site systems when there are issues, the difficulty of accessing and working at the site require that the instrument team must be able to make observations without being at the telescope. This chapter will explain the various operations and data acquisition (which I will sometimes abbreviate as DAQ) subsystems for Prime-Cam. By data acquisition, I include all of the hardware and software needed to take detector and housekeeping data and store it to disk, and by operations I include all of the software and procedures for controlling the telescope and Prime-Cam and managing the data acquired by DAQ tools until it is transferred off-site to our long-term storage and analysis data center in Cologne, Germany. As much as possible, I will distinguish data acquisition and operations from low-level timestream processing and mapmaking. The latter constitute the next steps in the pipeline for producing science-ready data products, which are usually maps but could also include calibrated timestreams for certain types of specialized mapmaking procedures or time-domain analyses. In practice, the data pipeline is the joint responsibility of members of the data acquisition and science working groups for CCAT, and some aspects of this processing (e.g. quicklook mapmaking) will be described in brief here.

Prime-Cam's DAQ software is largely built upon the SO-OCS framework of agents and supporting packages for remote observing described in section 2.5. I refer the reader there for an introduction to many of the terms and acronyms

that will be used in this chapter, but I will remind the reader that I will use SO-OCS to refer to SO’s Observatory Control System [10] and OCS to refer to the Observatory Control System developed by the CCAT software team [112] for control of the instrument-agnostic parts of FYST. Wherever possible, we use SO software, but it often must be modified, expanded, or augmented to fit in with OCS and technology specific to Prime-Cam.

To begin, I provide an overview in section 4.1 of the electronics and computing resources that will be deployed with Mod-Cam and Prime-Cam. Section 4.2 summarizes the development of PCS agents and the how agents will be deployed at the telescope site, while section 4.3 then walks step by step through the typical pattern of a remote observation with Prime-Cam, explaining the required software tools and interfaces along the way. Many of these steps have been tested with Mod-Cam in the lab at Cornell or by SO at their site, but some of them are plans that have yet to be implemented. The framework described here is up-to-date as of the time of the publication of this thesis, but we assuredly will make changes during the coming commissioning and early operation of Mod-Cam. The reader ought to be aware of this reality and consult updated references as they become available.

4.1 Data Acquisition Equipment and Layout

The Prime-Cam team is responsible for acquiring, configuring, and operating all the data acquisition electronics and computers needed to run Prime-Cam, but we will also need to interface with existing site hardware and the site network. Most of our equipment will be located in two different compartments of FYST: instrument space 1 (IS1) and the electronics space (ES). The locations of these areas of the




Figure 4.1: A cross-sectional view of FYST with the two compartments used by Prime-Cam DAQ hardware labeled. Instrument space 1 (IS1, blue) is where Prime-Cam itself will be located along with two electronics racks for DAQ hardware, and the electronics space (ES, orange) is where Prime-Cam computing resources will go.

telescope are shown in Fig. 4.1. IS1 will hold Prime-Cam itself and any of the hardware that must be able to connect directly to Prime-Cam. There will be two electronics racks without doors or cooling power in IS1 to hold this equipment, and more equipment may be mounted in various enclosures on Prime-Cam or its mounting raft if needed. The gas handling system (GHS) for the Prime-Cam dilution refrigerator will also be located in this space. The ES will hold computing resources that can afford to communicate over a network connection. There are several racks located in the ES; two of them have been earmarked for Prime-Cam's use while the others are for CHAI and the observatory more broadly. Since our ES racks will be near the observatory racks, any connections to the broader observatory

network or resources will be made here, including to the pulse tube compressors (PTCs) that are being managed for Prime-Cam by the observatory team.

The observatory management team also will store some network resources in shipping containers near the telescope. While Prime-Cam will have its own storage containers which could be used to house some of our equipment, including one that will be an assembly and staging area for hardware, we currently plan on keeping all data acquisition resources in the designated areas on the telescope.

4.1.1 Data Acquisition Hardware

Setting aside the computers and networking hardware for the moment, the main data acquisition hardware for Prime-Cam falls into three categories: cryogenics, instrument readout, and power. Any equipment for measuring temperatures, pressures, or the performance of the cryostat itself falls under the first category. The instrument readout category includes the RFSocCs described in section 2.4 as well as tools for biasing the warm amplifiers and any module-specific hardware. The power category includes power distribution units (PDUs) and uninterruptible power supplies (UPSs). Table 4.1 summarizes all of the equipment that will be deployed by Prime-Cam.

The thermometers on the 80 K, 40 K, 4 K, 1 K, and 100 mK stages of Prime-Cam’s DR, PTs, and modules will be monitored with Lakeshore 372s and Lakeshore 240s. There will be 71 thermometers for the initial deployment of Prime-Cam with four modules, and this number will grow to around 89 for all seven modules. The Lakeshore 372s can monitor up to 16 channels by cycling through them one at a time, and it can also control heaters as necessary. The Lakeshore 240s can measure

Category	Device	Quantity	Description
Cryo	Lakeshore 372	2 (for four modules)	AC bridge readout for cryogenic thermometry inside of Prime-Cam. 16 channels. Located in IS1 racks.
	Lakeshore 240	6 (for four modules)	4-lead readout for cryogenic thermometry. 8 channels. Located in IS1 racks.
	Bluefors GHS. May also include Bluefors TC	1	Monitoring and control for Bluefors DR system. GHS is standalone item in IS1.
	Cryomech PTC	4	Monitoring and control for Cryomech PTs. Located in compressor room.
	Teledyne pressure gauge	1	Pressure monitor for needle valve manifold for pumping out and venting Prime-Cam. Located on top of Prime-Cam in IS1.
	Labjack	1	Versatile readout board; used for monitoring ambient temperature of Prime-Cam and racks. Likely mounted in IS1.
Readout	RFSoc	25 (max for four modules)	Warm readout electronics in 1U enclosures. Located in IS1 racks.
	Amplifier biasing crate	1	Power supplies for up to 144 warm amplifiers. Located in IS1 racks.
	EoR-Spec capacitive bridge readout and motor controller	1 each	Sensors for measuring parallelism of FPI mirrors and homemade controller for stepper motor. Located in IS1 racks or mounted on Prime-Cam.
Power	Raritan PDU PX4-53A8C-C8E7A0	3	24-outlet vertically-mounted remotely switchable PDU. One in each IS1 rack and one in main ES rack.
	JOVYTECH PME3000K UPS	3	UPSs (and battery packs) to permit controlled shutdown during outages. One in each IS1 rack and one in main ES rack.

Table 4.1: Summary of DAQ hardware for initial deployment of Prime-Cam with expected quantities and locations.

up to eight channels simultaneously. Since the Lakeshore 240s have lower measurement accuracy than the Lakeshore 372s, we generally use the Lakeshore 240s for the higher temperature stages and the Lakeshore 372s for the lower temperature stages. Ambient pressures around the cryostat may be monitored with a Labjack temperature readout device.

For operation of the DR and PTs, the Bluefors GHS and Cryomech PTCs control the flow of gas through the system. The software that comes with the Bluefors DR records critical pressures for assessing the health of the DR and can be used to control the operation of the DR. We may also make use of the Bluefors Temperature Controller (BFTC or Bluefors TC) for operating heat switches and reading out a limited number of heaters and thermometers. Additional pressure gauges may be used to monitor the status of the vacuum in Prime-Cam itself, especially when pumping down or venting to local atmospheric pressure. The compressors will be located in a compressor room inside the base of the telescope and can be remotely accessed on the telescope network for control and monitoring of key internal temperatures and pressures.

The main warm readout data acquisition hardware is the RFSoc itself. Our Xilinx ZCU111 RFSocs are packaged in an enclosure by our collaborators at Arizona State University that includes the RFSoc, a power supply for the RFSoc that provides an external C13 connection and on/off switch, heat sinking for the RFSoc FPGA, an unmanaged switch to route the outputs of the RFSoc to a control RJ-45 port and a data RJ-45 port on the exterior of the enclosure (plus some other diagnostic ports), and the equipment to expose SMA coaxial connectors for connecting each network's input and output coax. The RFSocs will be located in the electronics racks in IS1 with short coax cables running to an interface panel

mounted somewhere on the rack to reduce strain on the SMA connections on the front of the RFSoc package. Longer coax cables will run from this interface panel to the readout harnesses on Mod-Cam or Prime-Cam. There are eighteen RF networks (five RFSocs with four networks per RFSoc) needed to read out the three arrays for the 280 GHz module; the 350 GHz module needs the same number, while EoR-Spec needs only twelve RF networks (three RFSocs). The 850 GHz module could need as many as forty-five RF networks (twelve RFSocs), though studies are underway to implement a wider bandwidth version of the RFSoc firmware that could reduce this by perhaps a factor of two.¹ In the current default scenario for the initial deployment of Prime-Cam with these four modules, as many as 25 RFSocs could be required, though that number could be reduced to 23 if some of the drones on a given RFSoc are shared between two modules.

The other core warm readout electronics equipment is a 4U rack-mounted amplifier biasing power supply from our collaborators at Arizona State University. It is based on a model used for powering the amplifiers for the ALPACA project, though it will be modified slightly for Prime-Cam’s needs. When fully populated, it can provide 1.5 to 2.5 V to 144 amplifiers, and it contains a control unit that can be used for remote control of the biasing. We plan on developing a PCS agent to operate it remotely.

Some modules may have additional hardware in IS1. EoR-Spec, for example, will have a controller for adjusting the stepper motors that govern the FPI mirror separation and a capacitive bridge monitoring device for the relative tilt of the mirrors to ensure that they remain parallel within required tolerances. The de-

¹Even if this firmware were implemented, it would not cause reductions in the needed RF network count for the other three modules because the detector arrays have already been planned and/or fabricated for the other three modules. It would reduce the needs for future modules, however.

velopment and testing of these tools are still in the preliminary stage. They may operate under a single PCS agent, or they may require two separate ones.

The power subsystem of Prime-Cam DAQ equipment includes both PDUs and UPSs designed to supply power at 220-240 V and 50 Hz. We will be using Raritan 24-outlet PDUs that mount vertically in the back of the rack instead of in the rack itself. These devices can remotely control or power cycle each C13 outlet, and they monitor power flow and a variety of other electrical parameters through each outlet. These PDUs are similar, but not identical, to other Raritan PDUs being used by the observatory software team.

Each rack in IS1 and the main Prime-Cam rack in the ES will have a UPS with a battery (each 2U) that can last up to 5-7 minutes at full power to enable computers and readout electronics to shut down gently during a site-wide power outage. Especially during early deployment, we will not operate at the full load of the batteries, so their charge should last longer. We intend to test this in the lab with Prime-Cam, and it is easy to add a second battery pack if needed. The UPSs will power the PDUs and provide a guarantee of power quality and surge protection for the racks as well. These UPSs are identical to those being used by the observatory software team for other electronics racks on the telescope.

4.1.2 Computers and Networking

In order to communicate effectively with the devices in the previous section, the Prime-Cam team will deploy a series of computers and network switches to provide connectivity, on-site data acquisition and storage, and resources for analyzing data during commissioning and early science operations.

Category	Servers	Small form-factor NUCs
Supplier	Supermicro	OnLogic
Model	SYS-510P-MR	ML100G-52
Processor	Intel Xeon Gold 6312U	i5 Intel Whiskey Lake
RAM	8x16 GB DDR4	1x16 GB DDR4
Storage	(internal) 960 GB + (ext.) 4x 3.84 TB	256 GB
Networking	2x 1 Gbit RJ45 + 2x 10 Gbit SFP+	2x 1 Gbit RJ45
Mount	1U Rack-mount	DIN-rail mounted

Table 4.2: Summary of key specifications of Prime-Cam computing resources. There will likely be at least five of the servers and two of the NUCs for initial deployment. All of the Supermicro servers will be located in the ES, while all of the NUCs will be located in IS1. These specifications are up to date as of May 2025, but we are considering upgrading the NUCs and may do so in the future.

The plan for computing resources could easily change in the future as additional needs are identified, but the nominal scheme involves one main DAQ node (known as `daq-primecam`) that will run most of the PCS agents and observing software. `daq-primecam` may also host the web interface for Prime-Cam observers. Each module in Prime-Cam will have its own data collection server with naming convention `data-<module tag>` (e.g. `data-280GHz`, `data-350GHz`, `data-eorspec`, etc.). These data collection servers run the software that saves detector data to disk and can be used for low-level data management and processing, as described later in this chapter. Each of these identical servers is outfitted with four solid-state drives (SSDs) in front-mounted external bay drives. These are currently around 3.8 TB each and are configured in a software (`mdadm`) RAID 10 setup so that two of the drives serve as mirror backups for the primary two drives. Each server runs Ubuntu 22.04 LTS at the moment, though this will be upgraded over the lifetime of the project. Additional specifications can be found in Table 4.2.

Two additional sets of data servers might eventually be acquired for site use: one or more dedicated analysis servers for data quality assessments and quicklook

mapmaking and a large central RAID server for easier access to the data being generated on the data collection servers. For the initial deployment of Mod-Cam, it is quite likely that a data server for one of the later modules will be used as an early analysis server. Similarly, the initial deployment of Mod-Cam may use network file system (NFS) tools to mount the large external bay drive storage of each data collection server on `daq-primecam` or a dedicated analysis machine for ease of accessing the data generated by each module. A centralized RAID server would take the place of this system, though it would still be mounted on shared computers with NFS.

In addition to these larger servers, a pair of small form-factor Intel NUC (next unit of computing) computers will be used in IS1 to accomplish specific tasks requiring dedicated interfaces. The Bluefors GHS requires software that can only run on Windows and requires a single USB connection for control of the GHS. We meet these requirements by installing Windows 11 on one of the NUCs, named `nuc-bluefors`, and locating it on the IS1 rack near the GHS. Similarly, the Lakeshore 240s require a Linux computer with a USB connection to be located in the IS1 racks to run their SO-OCS drivers, so we use a NUC with Ubuntu 22.04 LTS (`nuc-daq`) and a ten-port StarTech USB hub to handle these connections and any other data acquisition hardware that requires a USB connection instead of Ethernet.

The communication between the computers is mediated by network switches for the Prime-Cam subnet within the telescope. We expect that two switches per IS1 rack and one switch in the main Prime-Cam ES rack will be needed for the initial deployment of Prime-Cam, though we will start with one switch per rack for the deployment of Mod-Cam. Each of the switches is a PlanetTech IGS-6325-

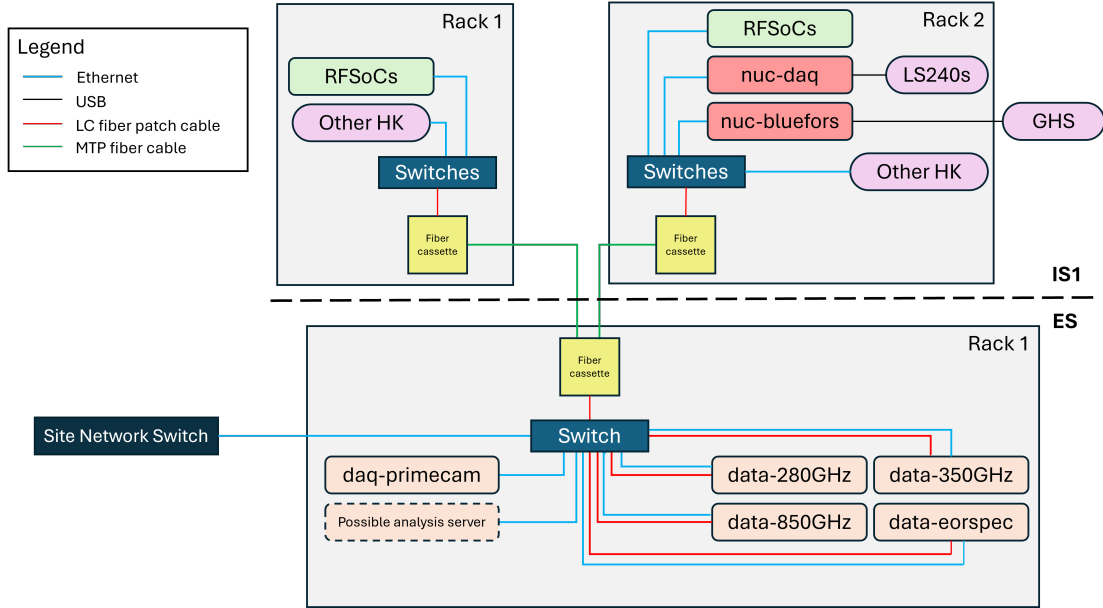


Figure 4.2: The Prime-Cam network, showing all computers and connection types as well as relevant HK equipment for the initial deployment of Prime-Cam with four modules. Computers are shown in red with servers in light red and NUCs in darker red. Readout equipment is in green, while other HK devices are in pink. The connection to the site network switch in the ES also connects the subnet to the gateway for outside connections that remote observers will use to authenticate when accessing Prime-Cam computers.

20T4C4X switch with sixteen 1 Gbit RJ45 ports, four 10 Gbit SFP+ optical fiber ports, and four ports that can be either additional 1 Gbit RJ45 ports or 1 Gbit SFP optical fiber ports.

Generally, devices and computers will connect to the switches with 1 Gbit copper (i.e. regular Ethernet) cables. The connections that carry the RFSoc data packets, however, will be 10 Gbit fiber connections wherever possible. This includes the network path between the switches in IS1 and those in the ES and from the ES switches to the data collection computers for each module. At Prime-Cam’s full data rate with seven modules, a 10 Gbit bandwidth should allow for a bandwidth safety factor of 2-4 at any potential network bottlenecks.

The two 15-m fiber MTP fiber cables that will run between IS1 and the ES are each a ruggedized, 24-fiber bundle that will feed into a breakout cassette in the electronics racks in each compartment. From these cassettes, patch LC cables will go to transceivers in the switches (and from the switches to the transceivers in the computers). Fig. 4.2 summarizes the setup of the Prime-Cam subnet.

An absolute PTP timing reference signal from a Meinberg microSync RX200 clock synchronized to GPS will be distributed across the telescope network. This signal will be received by `daq-primecam` via a connection to the site network on the ES switch and redistributed across the Prime-Cam network to the other computers and RFSOCs using `ptp4l` for Linux for the most time-sensitive data collection. Regular system time for computers and the timing for most housekeeping equipment can be provided by NTP. The PlanetTech switches are PTP compatible and will be set up as boundary clocks in the network. We expect there to be some degradation of the level of synchronization for the RFSOCs due to the unmanaged switch in the RFSOC enclosure, but initial tests indicate that synchronization to the level of hundreds of microseconds or better is possible, which is more than an order of magnitude below our sampling rate. Further testing will be necessary to understand whether this is sufficient or if additional steps must be taken to improve the timing precision delivered to the RFSOCs.

4.2 PCS Development and Implementation

Within the SO-OCS framework, the Prime-Cam Control System (PCS) refers to any agents developed for communicating with Prime-Cam data acquisition hardware that is unique to Prime-Cam (i.e. not being used by SO) or CCAT lab testing.

Progress of development and deployment can be tracked on the PCS GitHub repo, which also builds a common Docker image of all PCS agents for ease of deployment.² As of the time of writing, agents are in use for the Bluefors Temperature Controller, our Raritan PDUs, the Lakeshore 325, and a software-only interface agent for communicating with the Telescope Control System (TCS) with CCAT OCS for directing the Antenna Control Unit (ACU) to move the telescope when Prime-Cam is observing. An agent has also been developed for controlling the RFSocCs.³

Informal agents have been developed locally (i.e. not yet on GitHub) for other testing equipment with Mod-Cam, including our beam mapper and polarized source X-Y stage and a Teledyne pressure gauge for monitoring the needle valve manifold that slowly vents or pumps down Mod-Cam. In the future, some of these agents will be added to the PCS repository, and agents will be developed for other hardware like the UPSs and module-specific hardware. The agents for most of the rest of our HK equipment are shared in common with SO, and we will be using the relevant and well-tested `socs` agents for these tools.

Many of the agents operate passively to monitor the state of their respective hardware at all times, but some of the agents are used more actively during observations, such as the ACU interface agent. The use of these agents will be described in the sections below. For commanding the agents which are usually passive during laboratory testing or commissioning, the full agent API is available via the PCS documentation.⁴

²<https://github.com/ccatobs/pcs>

³This agent is not in the main PCS repo but the `primecam_readout` repo: https://github.com/TheJabur/primecam_readout/blob/develop/src/queen_agent.py

⁴<https://pcs.readthedocs.io/en/latest/> (and https://github.com/TheJabur/primecam_readout/blob/develop/docs/docs_primecam_readout.md for the RFSoc agent)

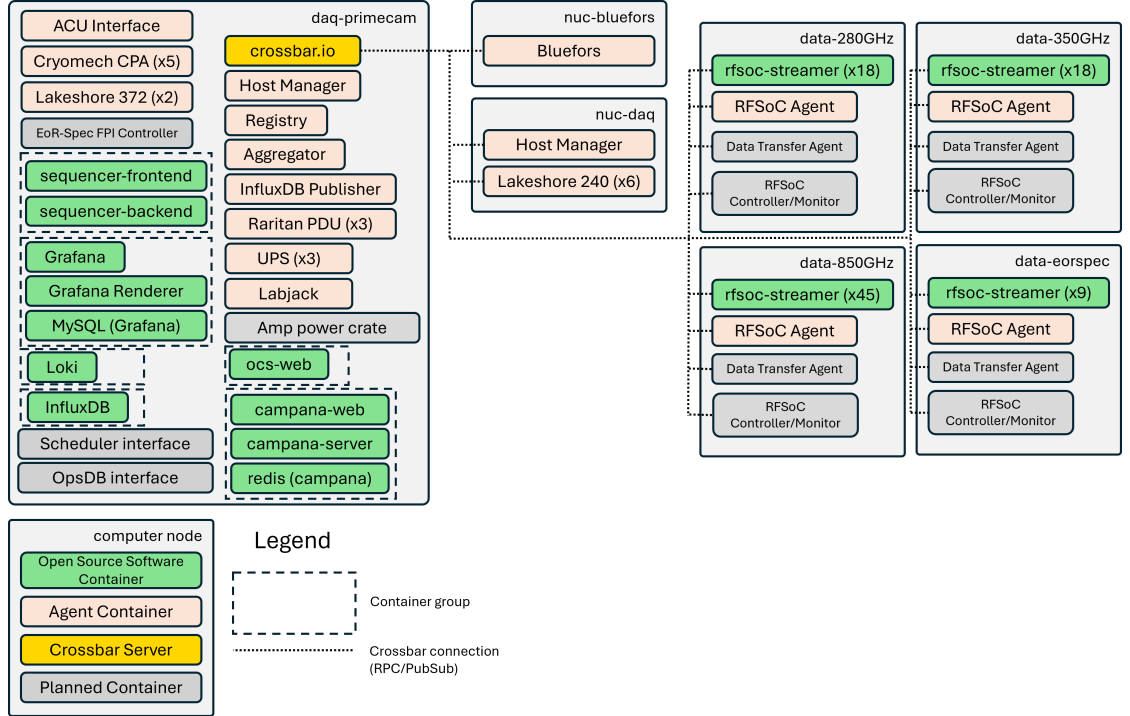


Figure 4.3: Diagram of the configuration of the PCS agents and other associated software containers for Prime-Cam. The container groups refer to separate docker-compose files that allow that group to be altered without affecting the operation of the rest of the system. It is quite likely that this diagram will evolve over time as new equipment is brought online and as the exact operation of the RFSocCs is fine-tuned.

The agents are deployed on a single crossbar network for communication, but each agent runs in a Docker container on one of the computers spread throughout the Prime-Cam network depending on what is needed for communicating with the relevant hardware. Fig. 4.3 shows where each agent runs and how the agents communicate over the network. It also shows Docker containers for other relevant open source software tools, some of which have been developed by the Prime-Cam team, some of which come from SO, and some of which is developed by the open source software community. These will be described in greater detail in the coming sections.

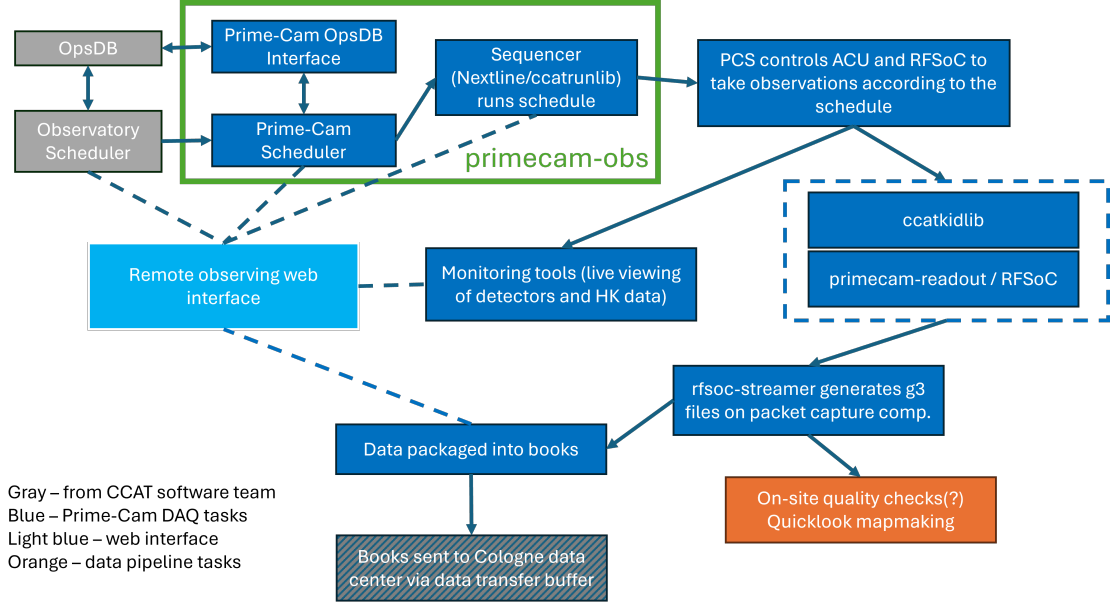


Figure 4.4: Flow chart of a typical observation with Prime-Cam. Prime-Cam hardware and software are shown in blue. Interfaces with the observatory software (gray), remote observers via a webpage (light blue), and low-level and quick look analysis tools (orange) are shown where they are essential for understanding the observational workflow.

4.3 Operations and Data Management

In addition to the hardware described in Sec. 4.1 and the PCS software described in Sec. 4.2, there are a variety of other software tools and interfaces required to facilitate observations with Prime-Cam. The following sections will step through the plan for a standard observation with Prime-Cam; that is, an observation after the commissioning phase is over and the instrument is operating semi-autonomously under the supervision of a remote observing coordinator (ROC). Along the way, we will have occasion to discuss many different software tools, their status, and possible extensions or alternatives to this plan.

Fig. 4.4 depicts the steps of a Prime-Cam observation schematically. There are three major movements. First, the remote observer interacts with the telescope

scheduler and operations database (OpsDB) to identify targets for one or more observations. This information allows the ROC to generate a PCS client schedule for running the observation. Second, the client schedule is run, executing the observations. Finally, the data that is generated by the observation is packaged and safely copied off-site to the CCAT allocation at the data center at the University of Cologne. Throughout the process, information about the status of the observation and data is tracked and made available to the ROC and other CCAT team members for viewing via the web interface.

4.3.1 Initiating Observations

Fig. 4.5 shows the initial steps in a Prime-Cam observation. Here, there are interfaces with two observatory software tools and the remote observer web interface, which will be facilitated by a collection of software packages loosely referred to as `primecam-obs` for convenience: the Prime-Cam scheduler, Prime-Cam OpsDB interface, and Prime-Cam sequencer.

The first step in an observation with Prime-Cam is identifying the target for the observation. Our CCAT collaborators, especially Yoko Okada at the University of Cologne, have developed a scheduling tool which I will refer to as the observatory scheduler. The scheduler is a Python script that generates a list of the five best targets available for the next twelve hours for either Prime-Cam or CHAI, depending on which is observing. It takes into account the various science goals for the observatory and the time that has been allocated to each for a given season, ensuring that sufficiently deep observations are made in the requested fields to accomplish the science proposed by the CCAT science working groups. The scheduler also takes into account which fields are available in pre-defined elevation

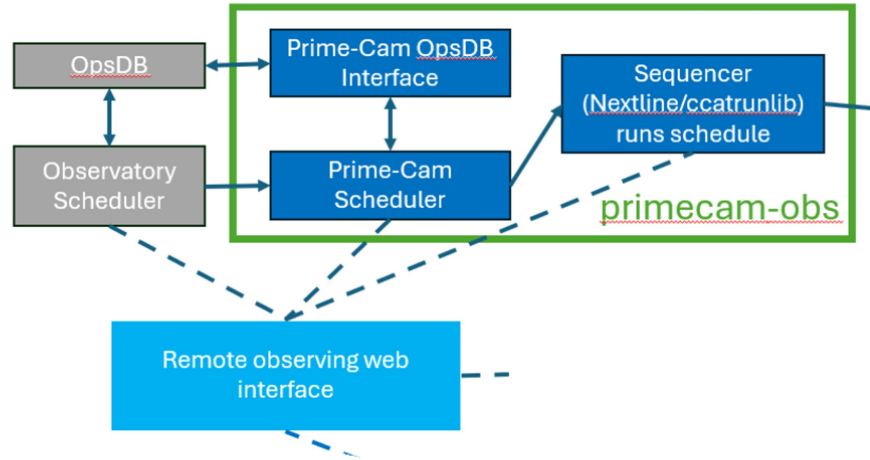


Figure 4.5: The steps and interfaces involved in initiating an observation with Prime-Cam.

ranges, and it incorporates information about sun avoidance when proposing the optimal targets.

In order to process the targets efficiently, all of the possible observational targets are tracked in the CCAT Operations Database (OpsDB), which will likely be hosted at the data center in Cologne or on the telescope site. It was developed largely by Christof Buchbender and Yoko Okada. Within the OpsDB, individual observations are stored as ObsUnit objects. Each observation constitutes a roughly one hour scan with information about the field, scan pattern, and any special constraints on the observation (e.g. if the observation must be made at a certain elevation or while the field is rising on the night sky). There is a single ObsUnit for every class of observation. Since the full science cases require longer total observations than one hour, the same ObsUnit may be scheduled many times. The observatory scheduler aims to fulfill as many of the ObsUnits for different science cases as possible in an equitable way.

The observatory scheduler will update its results at a set cadence and publish the results to an API endpoint. The Prime-Cam team will query this endpoint as part of its own scheduling software package. The Prime-Cam scheduler will receive the list of target ObsUnits from the Observatory Scheduler, parse them with reference to the OpsDB, and translate them into a PCS client script that can command the ACU interface agent to execute the appropriate scan and the RFSoc control agent to take data when needed. The Prime-Cam scheduler has not been developed yet as of the time of writing.

Initially, schedules will need to be generated by the remote observer regularly, perhaps as often as every few hours for more complex observations. For longer wide field survey scans, schedules could be generated more infrequently, stringing together multiple ObsUnits within a single schedule. Utilizing the observatory scheduler's ability to project optimal targets over a twelve hour window, the long-term goal would be to produce schedules once a night (or once at night and once during the day if the instrument performs well enough to add daytime observing). Optimizing the Prime-Cam schedules to minimize downtime between observations within a schedule while waiting for targets to rise, tuning, and moving the telescope will be a goal of the commissioning phase for the instrument.

The remote observer will be able to access information about scheduling through the remote observer web interface. This interface is still in the initial design phase, so it could end up changing considerably. It will likely be hosted by the observatory software team and would accommodate the observatory Grafana dashboards and observatory scheduler interface as well as the web viewing tools for Prime-Cam described in greater detail in the next section under the **primecam** subdomain,⁵ though the site has not yet been developed or tested. A tool is

⁵For example, the Prime-Cam Grafana instance might be available at something like

planned for parsing and visualizing the five targets recommended by the observatory scheduler so that the ROC can select which one will be used next. In practice, Prime-Cam will likely always select the observatory scheduler’s top target unless there is a special reason not to do so. We would also like to have a tool to use the Prime-Cam scheduler interactively to generate schedules. This tool would enable the quick production of standard schedules for emergency procedures like stowing the telescope for bad weather. For initial usage, this tool may be as simple as a form that could be submitted to turn a single ObsUnit into a PCS client schedule.

As schedules are produced and run, the information about which ObsUnits have been selected and run needs to be fed back into the OpsDB. Communicating with the OpsDB will be handled by a small Prime-Cam OpsDB Interface package that will provide Python tools for communicating with the OpsDB. This interface package will be called during the schedules to generate a new ExecutedObsUnit object in the database as soon as an observation begins, which is important so that the scheduler knows what is currently running and updates its outputs accordingly. When an observation finishes in the Prime-Cam schedule, another OpsDB interface function will be called to update the status of the ExecutedObsUnit to “finished” and add additional information as required. If the observation is interrupted, the ExecutedObsUnit will be updated to reflect that the observation was canceled. This procedure has not yet been implemented or tested for Prime-Cam, but it will make significant use of the existing implementation for CHAI.

The schedules generated by the Prime-Cam scheduler are simply PCS client scripts that communicate with the ACU interface agent to direct the telescope to execute one of several pre-programmed scan patterns and the RFSoc agent to stream data during the scan. The ACU interface agent (and the underlying

site.ccatobs.org/primecam/grafana.

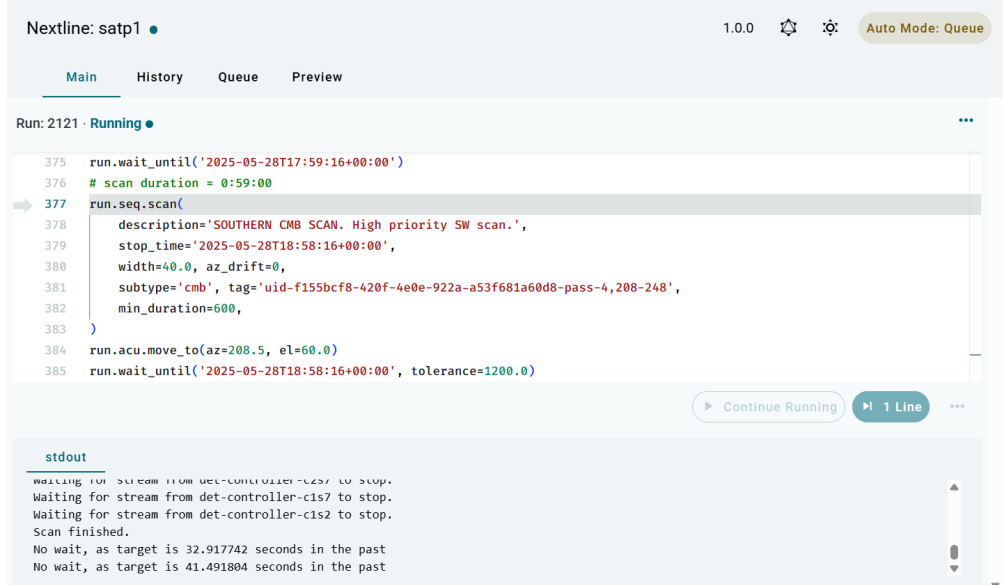


Figure 4.6: Example of the Nextline web interface for one of the SO SATs. The currently running line is highlighted. Other tabs show the past schedules, any queued schedules, and the real-time scheduler interface, which is under development at this time.

TCS code) currently permits moving the telescope to a single location, running a constant elevation scan, and running an arbitrary scan from a list of azimuth and elevation points; additional scan patterns will be added to both the TCS and the ACU interface agent before and during the commissioning period. The ACU interface agent also records pointing information (azimuth and elevation, plus additional ACU flags) from UDP packets generated by the TCS at 200 Hz. These data, along with all of the other housekeeping data generated by PCS agents, are stored in HK g3 files on **daq-primecam**.

While these schedules could be run as ordinary Python scripts, we will run them with the Prime-Cam sequencer Docker containers. Adopted from the SO sequencer described in section 4.4 of Ref. [113], this tool provides a web-based frontend called Nextline [114] that allows remote observers to see where they are in the schedule, to interrupt schedules easily when needed, and to store previous

schedules and their final result (succeeded or stopped). Fig. 4.6 shows what this web interface will look like; we expect it to be hosted as a page on the remote observer web interface. SO has used Nextline extensively for their site operations already. The sequencer backend will incorporate a package called `ccatrunlib`, modeled off of SO’s `sorunlib`, that lightly wraps some of the PCS client code to make the schedule more human-readable. This package has not been adapted and tested yet, but SO’s version has extensive use in the field.

4.3.2 Detector Operations

RFSoc Control

Many of the commands in the Prime-Cam schedules direct the RFSoc readout to measure the state of the MKIDs in each module. The stages of the Prime-Cam observation flowchart related to detector operations are shown in Fig. 4.7. The first, and most fundamental, layer of remote control of the RFSocs is the `primecam_readout` package,⁶ written by our Canadian CCAT collaborators (particularly James Burgoyne). This software uses a central program, known as the “queen” and runs as a PCS agent, a command line interface (CLI), or a graphic user interface (GUI) to control a group of subordinate programs called “drone” instances. Each drone controls a single RF network. With the current firmware, a single RFSoc board supports four RF networks, so there are commonly four drones running on each RFSoc CPU as Linux services.

To measure the response of our MKIDs on a single network, probe tones must be placed at appropriate frequencies with the right amplitudes (or, equivalently,

⁶https://github.com/TheJabur/primecam_readout/

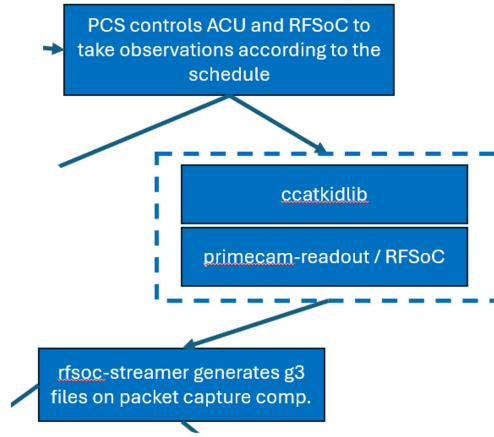


Figure 4.7: The steps and software packages involved in making measurements with Prime-Cam’s MKIDs during an observation.

the total tone power) to maximize the response to optical power while minimizing readout noise. The process of setting the probe tones is known as “tuning.” While the precise procedure that we will use to tune our MKIDs is still being developed and will take into account differences between our TiN and Al MKIDs, the basic procedure is as follows:

1. Once the detectors have been cooled to the desired operational temperature (generally 100 mK for Prime-Cam), each drone takes a VNA sweep, a relatively coarse measurement of the transmission through the RF network as a function of frequency.
2. From the VNA sweep, a set of algorithms are run to identify resonators as sharp dips in the transmission in a narrow frequency range and to weed out false positives.
3. Once the resonator locations have been roughly identified, a target sweep is run with finer frequency bins around each resonator to map the precise shape

of the resonator.

4. The target sweep is fit to a model for the resonator response to identify the minimum of the resonator and other parameters.
5. A frequency comb (frequency, amplitude, and phase) with a peak at each resonator location is generated and sent into the RF network. In temporal space, this is equivalent to generating a waveform that is the superposition of sine waves corresponding to the tone frequencies, digitizing it, and sending it into the RF network.

The VNA sweep and target sweep data are saved to the board and can be copied from there to the control computer. Custom comb files are generated on the control computer and moved to the appropriate location on the RFSoc in order to be parsed by the board. Once the tuning is finished and the probe tones have been written, the drone begins streaming data packets at a rate of $512 \times 10^6 / 1024^2 \simeq 488$ Hz. Each packet reflects a measurement of the complex S_{21} (or forward transmission) for each tone. The documentation for the full list of `primecam_readout` functions can be found at https://github.com/TheJabur/primecam_readout/blob/develop/docs/docs_primecam_readout.md#board-commands-board-commands (as of June 2025).

Each of the steps in this process can be accomplished by one or more functions within the queen agent. For the purposes of grouping some of these low-level steps together into higher-level functions that map more cleanly onto operational steps, integrating additional code for more complex fitting and tuning (especially for the more complicated AI detectors), and handling input and output files for configuring the entire Prime-Cam detector system, another software package called

`ccatkidlib` has been developed. Ultimately, the functions called by `ccatrunlib` in the Prime-Cam schedules will be `ccatkidlib` functions with descriptive names like `find_resonators()`, `tune_all_resonators()`, etc. These are only sketches of the final operations functions; at the moment, `ccatkidlib` is in a developmental phase that is optimized for laboratory testing of MKIDs in Mod-Cam. While the RFSoc agent in `primecam-readout` can be used directly in scripts run in Nextline, another PCS agent along the lines of the SO `pysmurf-controller` agent might be needed to use `ccatkidlib` functions in our scripts. Integrating `ccatkidlib` into the broader observatory control system is a current priority for our team.

As part of the development of the functions that will perform the tuning procedures, a key outstanding question for Prime-Cam that must be resolved during the commissioning phase is how often the detectors need to be retuned and the precise procedure for doing so. As the atmosphere changes and as the telescope changes its elevation (and thus the effective line-of-sight thickness of the atmosphere), the optical loading on the MKIDs will change. If the loading changes enough relative to the loading when the resonator was tuned, then the resonator may shift far enough that the probe tone is no longer located in the resonator. The quality factor will also degrade at higher loading, which affects the optimal probe tone power for the measurements to remain dominated by the signal from the sky instead of readout noise. In this case, the probe tone locations and amplitudes need to be reset by running the entire tuning procedure more or less from the beginning, a time-consuming process. It may be the case, however, that simpler and more frequent tuning procedures could be developed that make smaller adjustments to the probe tones without the need for a full set of VNA and target sweeps. This is an active area of research with Mod-Cam, and the commissioning period will permit more realistic tests of the arrays in Mod-Cam for different scan strategies and site

operating conditions.

Capturing Data Packets

Once tuning is done and data streaming is turned on in `primecam_readout` or `ccatkidlib`, the RFSoc firmware generates and sends a UDP packet from a different (configurable) IP address for each drone. The packet structure (as of June 2025) is defined at https://github.com/TheJabur/primecam_readout/blob/develop/docs/rfsoc_datagram.csv. A maximum of 1024 channels can be read out for a single drone. Every packet contains a single measurement of the complex S_{21} broken up into two 32 bit fields, one for the real and one for the imaginary part. These are usually referred to as the in-phase (I) and quadrature (Q) components, respectively. Every packet contains 2048 of these fields, regardless of whether there are 1024 tones placed or not; only the number of fields corresponding to twice the number of active tones provide meaningful data, beginning with the first channel. In addition to the main data fields, each packet contains two single bit flags that can be set by the user via a `primecam_readout` function, the packet number since the RFSoc was last initialized (32 bit field), the number of active tones (16 bit field), and a 12 bit PTP timestamp produced by the RFSoc when the packet is generated.

The packets are broadcast over the network to a common IP address that can be set in the board configuration files of `primecam_readout` on each RFSoc. Within our computing setup, this will be an IP address on the data subnet (a separate subnet from the regular Prime-Cam subnet) for the data collection computers. To take a concrete example, there will be eighteen drones streaming data for the 280 GHz module if all RF chains are operational. Each one will have an IP address

that could look like 10.10.101.x (the exact subnet number has yet to be assigned), and they would stream data to the fiber interface on **data-280GHz** that could have an IP address like 10.10.101.2. Since the sampling rate is roughly 488 Hz, this corresponds to 18×488 or 8784 packets per second. Each packet is 8253 bytes (including the UDP header), so the total bandwidth for this module on the network prior to capturing the packets would be at least 69.1 MB/s. When multiplying by seven and a safety factor of two because some of the later modules could have many more detectors per network and would require larger packets, a rough estimate for the total Prime-Cam network bandwidth required for the transfer of packets from IS1 to the ES alone is 0.95 GB/s. This bandwidth will be distributed across multiple 1 Gbit/s copper cables to go from each RFSoc to the IS1 switches and will be transferred from IS1 to the data collection computers in the ES by two 10 Gbit/s fiber connections, so there should be significant safety factors for network bandwidth throughout the relevant interfaces for data collection.

Once the packets reach the relevant data collection computer, they are captured by an instance of the **rfsoc-streamer** package.⁷ Modeled extensively off of the **smurf-streamer** package used by SO [115], the **rfsoc-streamer** makes use of software from the South Pole Telescope Third Generation (SPT-3G) to generate g3 files.⁸ This custom file format, used by SPT-3G and SO for their data, is a serial file format comprised of a stack of frames. Each frame contains objects defined in **spt3g-software**⁹ or **so3g**¹⁰ for wrapping and storing data types. Each frame is effectively a dictionary containing objects that have been serialized with

⁷<https://github.com/ccatobs/rfsoc-streamer/>

⁸See the File IO section of the docs at https://cmb-s4.github.io/spt3g_software/index.html for more information.

⁹https://github.com/CMB-S4/spt3g_software

¹⁰The SO extension for working with g3 files. As of June 2025, it is being reworked to incorporate some of its code into **spt3g-software**, so our use of it may change in the future. <http://github.com/simonsobs/so3g>

`spt3g-software`, and the tools in `spt3g-software` must be used to deserialize the objects correctly to load them back into memory.

The `rfsoc-streamer` comprises a set of C++ and Python modules for a G3Pipeline, the `spt3g-software` framework for processing data by running a series of programs that operate on the data. The `RfsocTransmitter` class begins the pipeline by capturing new UDP packets from a particular IP address via a socket. The system is designed so that there is a separate instance of the `rfsoc-streamer` running for each drone, which permits the streamer to listen on only a single IP address and ensures that all of the data from different drones will be stored in different files for clarity. The `RfsocTransmitter` wraps the raw packet in an `RfsocSample` object with a timestamp denoting when the packet was received (which can be used in place of the packet PTP timestamp as the data timestamp if only low timing precision is necessary) and passes this object to the `RfsocBuilder` module. The builder sets up a pair of queues that aggregate data for a configurable period of time (default: three seconds). At the end of this period, all of the packets received during that time are passed to the `FrameFromSamples()` method. This method parses the packets, extracts the I and Q samples for each active channel (each channel from index 0 to the channel count listed in the packet), and saves the timestamps, channel names, and aggregated data as arrays in the `G3SuperTimestream` object provided by `so3g`. This object is stored in a g3 Scan (or data) frame along with a timestamp denoting when the frame was generated, the active channel count, and the number of packets aggregated into the frame.

Each frame is then passed to two further modules. The `SessionManager` module handles the starting and stopping of the g3 files and adds metadata like the software version number, a stream ID that is the ten digit Linux ctime when the new data

```

Frame (Observation) [
  "ccatstream_id" (G3String) => "rfsoc02_drone1"
  "ccatstream_version" (G3Int) => 1
  "frame_num" (G3Int) => 0
  "session_id" (G3Int) => 1739818198
  "stream_placement" (G3String) => "start"
  "time" (G3Time) => 17-Feb-2025:18:49:58.649623000
]
Frame (Scan) [
  "ccatstream_id" (G3String) => "rfsoc02_drone1"
  "ccatstream_version" (G3Int) => 1
  "channel_count" (G3Int) => 1
  "data" (G3SuperTimestream) => G3SuperTimestream(2, 1465)
  "frame_num" (G3Int) => 1
  "num_samples" (G3Int) => 1465
  "packet_flag_1" (G3Int) => 0
  "packet_flag_2" (G3Int) => 0
  "session_id" (G3Int) => 1739818198
  "time" (G3Time) => 17-Feb-2025:18:50:01.649447000
  "utc_from_ptp" (G3Time) => 04-Nov-2026:21:20:27.189869000
]
Frame (Scan) [
  "ccatstream_id" (G3String) => "rfsoc02_drone1"
  "ccatstream_version" (G3Int) => 1
  "channel_count" (G3Int) => 1
  "data" (G3SuperTimestream) => G3SuperTimestream(2, 1465)
  "frame_num" (G3Int) => 2
  "num_samples" (G3Int) => 1465
  "packet_flag_1" (G3Int) => 0
  "packet_flag_2" (G3Int) => 0
  "session_id" (G3Int) => 1739818198
  "time" (G3Time) => 17-Feb-2025:18:50:04.649933000
  "utc_from_ptp" (G3Time) => 04-Nov-2026:21:20:30.190087000
]

```

Figure 4.8: Contents of an example g3 file from the testing of the `rfsoc-streamer` showing the frame structure and contents of the first few frames. Each g3 file begins with an Observation frame, and the data from the RF-SoC is stored in Scan frames. The PTP clock on the RFSoc was not calibrated when these data were taken, which is why the PTP timestamps do not match the receiving timestamp.

stream began, and the frame number in the stream. As of June 2025, new files are generated by listening for the first Scan frame to enter the pipeline. While no packets are being received by the `RfsocTransmitter`, an empty `None` frame is passed through the rest of the pipeline by the `RfsocBuilder` every aggregation period. The `SessionManager` simply discards these. When new packets arrive and the first Scan frame is built, the `SessionManager` initializes a new session, generating an Observation frame that marks the timestamp of the beginning of a

new stream and passing both this Observation frame and the Scan frame on to the FileWriter module. Fig. 4.8 shows an example of the fields stored in Observation and Scan frames. Similarly, when there is a transition from a steady stream of Scan frames to a None frame, the SessionManager ends the session, writing out an end Observation frame to tell the FileWriter to close the current file. Future development may change this system to make use of the flags in the packets to know when to close files so that the session is ended actively instead of passively.

The final module is the FileWriter module. This module writes any g3 frame passed to it by the pipeline to a g3 file using an instance of the G3Writer class from `spt3g-software`. It also sets up the filesystem structure of the output files and rotates the files after a preset time (default: ten minutes) to prevent any single file from getting too large. Within the top-level data directory, a folder is created with the first five digits of the session ID, which roughly ensures that all timestamps from a single day will end up in the same folder. Within this folder, there are folders for each drone, specified by a board ID and a drone ID that is set in the `stream_config.yaml` file used to set up all the `rfsoc-streamer` instances running on the network. Each RFSoc should be assigned a different board ID, and this mapping needs to be tracked over time. Within each folder, the files are labeled by the session ID, board and drone number, and an index that is incremented as files are rotated. See Listing 4.1 for an example of this filesystem structure for a mock series of thirty minute observations using three drones across two RFSocs.

Listing 4.1: Mock rfsoc-streamer output file structure for observations made with different drones over a period of two days. On the first day, corresponding to timestamps beginning with 17389, a single 30 minute stream was commanded for three drones across two RFSocCs, generating three files each. Due to latency in commanding the system, the timestamps in the files of rfsoc02_drone1 are incremented by one second relative to the other two, which can happen in our real data. On the second day, corresponding to timestamps beginning 17390, a 30 minute stream was taken with two drones. A second 10-20 minute stream was then taken on rfsoc02_drone04 alone.

```

/data/g3/
├── 17389/
│   ├── rfsoc01_drone1/
│   │   ├── r01d1_1738956095_000.g3
│   │   ├── r01d1_1738956095_001.g3
│   │   └── r01d1_1738956095_002.g3
│   ├── rfsoc01_drone2/
│   │   ├── r01d2_1738956095_000.g3
│   │   ├── r01d2_1738956095_001.g3
│   │   └── r01d2_1738956095_002.g3
│   └── rfsoc02_drone1/
│       ├── r02d1_1738956096_000.g3
│       ├── r02d1_1738956096_001.g3
│       └── r02d1_1738956096_002.g3
└── 17390/
    └── rfsoc01_drone1/

```

```

|
|   — r01d1_1739069374_000.g3
|   — r01d1_1739069374_001.g3
|   — r01d1_1739069374_002.g3
|
|   rfsoc02_drone4/
|   |
|   |   — r02d4_1739069374_000.g3
|   |   — r02d4_1739069374_001.g3
|   |   — r02d4_1739069374_002.g3
|   |   — r02d4_1739087125_000.g3
|   |   — r02d4_1739087125_001.g3

```

One area for future development in `ccatkidlib` or the `rfsoc-streamer` is a plan for storing the status of the RFSoc. SO stores information about the low-level configuration of the SMuRF readout system in status frames in its g3 files because the SMuRF system supplies status UDP packets that can be handled in a similar fashion to the data packets. The RFSoc does not do something similar, but there is a function in `primecam-readout` that collects the status of user-configurable options on the RFSoc so that it can be saved. One possibility is that this function could be called every time that a stream is started in `ccatkidlib`. The output could be dumped into a timestamped text file with the board and drone ID in the name, and these status files could be stored alongside any other `ccatkidlib` auxiliary files. Effort needs to be devoted for deciding the exact content of these files and the cadence for producing them.




Figure 4.9: The steps and software packages involved in processing and transferring the raw g3 files off-site.

4.3.3 Data Storage and Management

After the detector and housekeeping data have been written to g3 files on various computers, there are two final steps of data processing and management before the data acquisition process is fully complete. This subset of the observation flowchart can be found in Fig. 4.9.

The first of these steps is adapted from SO and is known as the bookbinding process. In the SO paradigm, Level 1 (L1) data refers to the data as it is being generated in the readout and data packets. Data in the initial g3 files is referred to as Level 2 (L2) data. These files do not guarantee that there are common starting points or frame boundaries across files from different drones, and they may be missing frames or have dropped samples within frames. The detector data for a given observation are located in g3 files within one part of the site filesystem, while the pointing data needed to make maps from the detector time-ordered data (TOD) resides in HK g3 files on another computer. All of this makes accessing and using the L2 data unwieldy.

To resolve these issues, SO runs all their detector and pointing data through a

program called the bookbinder. The most complete treatment of the bookbinder can be found in Ch. 4 of Ref. [116] and the references to SO documentation therein. The code itself can be found in the `sotodlib` library.¹¹ Put simply, the bookbinder processes the relevant data from an observation into a book, which is a directory containing several files in the final format that will be used for archiving and analysis. The g3 files produced by the bookbinder are coterminous, cosampled, and complete. They are coterminous in the sense that all of the files for a given module have been aligned such that they have the same frame boundaries. The pointing data is sampled at 200 Hz while the detector data is sampled around 488 Hz, so the pointing data is interpolated to match the sampling rate of the detector data; the final pointing and detector timestreams are cosampled. The bookbinder also completes the timestream by using linear interpolation to fill in any missing frames or samples within frames to mitigate artifacts from irregular sampling in the Fourier transforms that are used to process the TODs for mapmaking and science analyses. A full observation book contains a small number of g3 files with the pointing data, many g3 files containing the detector data, and a handful of metadata files that provide high level information about the book for ease of sorting through books.

Work is underway (especially by Sourav Sarkar at the University of Alberta) to adapt the SO bookbinding code for use with Prime-Cam data. Our g3 files and associated auxiliary files have a slightly different structure than SO's that needs to be taken into account. The SO bookbinder assumes a constant elevation scan pattern, but Prime-Cam will observe in other scan patterns as well, requiring modifications to the “scanification” feature of the SO bookbinder. We expect to make one book for each module in Prime-Cam that is active during a given

¹¹<https://github.com/simonsobs/sotodlib/blob/master/sotodlib/io/bookbinder.py>

ObsUnit. Work remains to be done to set up the database structure for tracking new files that the bookbinder uses and testing this code with Mod-Cam. It is likely that we will need an equivalent to the SO `pysmurf-monitor` agent or some other solution to monitoring output files within `ccatkidlib` for filling out these databases.

In addition to “obs” books that contain the main detector and pointing data, the SO bookbinder makes books for housekeeping data, operations books that contain detector tuning and other calibration information, and miscellaneous books to handle edge cases. Similar books will need to be implemented to track our tuning outputs, planet scan data, and any other `ccatkidlib` information or configuration files that are needed to interpret the raw data. None of these types of books have been tested with CCAT data yet. The process of integrating our detector operation software with the data handling and remote operation tools that we hope to borrow from SO and our CCAT collaborators remains one of the most important areas for development if we wish to follow the framework outlined in this chapter.

On the SO site, a workflow management system called Prefect regularly runs the bookbinder on new observations. It is likely that we will eventually adopt this system or a similar workflow management tool (the simplest of which would be a cronjob) to run the bookbinder regularly on new measurements. This has not yet been implemented for Prime-Cam.

Once the books are produced, the data must be transferred off-site. CCAT has an underground fiber link through REUNA that allows the data to be copied to the collaboration’s storage space at the data center in Cologne. Additional storage sites in North America are also being considered. For data transfer, we expect to zip the

books and transfer them using the bbcp open-source file copying application¹² for handling the data transfer and checksum verification. Our German collaborators have begun testing this application for data transfers from Chile to Cologne. The path to each book should also be registered with the appropriate ExecutedObsUnit in the OpsDB so that there is an easily retrievable record of which data belong to which observation.

Once the data have been backed up off-site, the data can be removed from the site computers. Each data collection computer has about 7.8 TB of storage space, which should be enough for 4-7 days of data storage for each module depending on the telescope's observing efficiency. It is possible that a central RAID server will eventually be added to the Prime-Cam network to permit easier sharing of data and provide a much larger storage space. Until that happens, the bookbinding process will likely take place on each data collection computer and may result in a second copy of the data being made temporarily, reducing the amount of time available before the storage is full. Another possibility that could be used with either the central RAID server or individual computers is that the SO SupRsync agent¹³ could be used to handle internal data transfer before or after bookbinding and to remove old data that has been backed up locally or remotely. Procedures to remove the L2 data once the L3 books are bound and to remove the site L3 books once they are backed up to Cologne need to be developed.

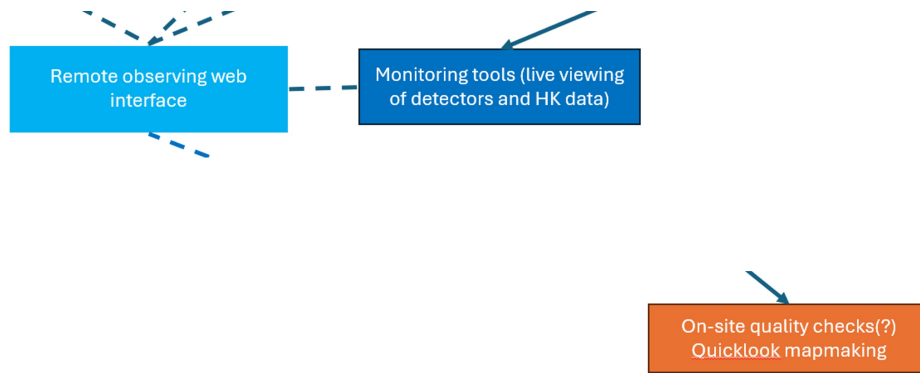


Figure 4.10: The tools involved in remotely monitoring the status of an observation.

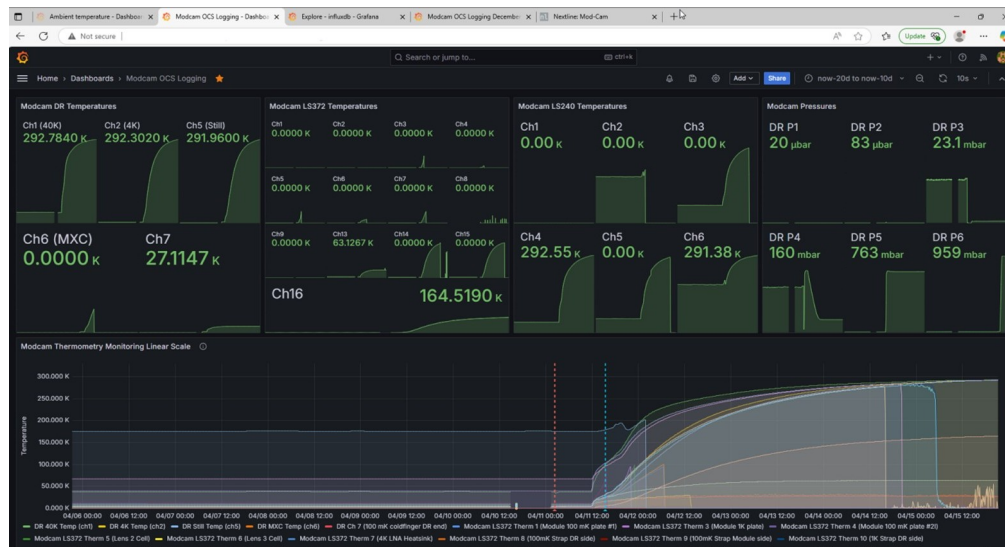


Figure 4.11: The main Grafana dashboard for monitoring Mod-Cam cryogenic performance. The most recent value for the main temperatures and pressures are shown for ease of viewing, and time series of the HK data can be found by scrolling down. In this particular time window, Mod-Cam was warming up to room temperature. Many other types of dashboards and panels can be created.

4.3.4 Remote Observing Tools

The remote observer has a variety of tools to assess the status of the telescope and to monitor for issues once an observation has been put into action by all the steps described above. Fig. 4.10 highlights where these tools fit into the observation flowchart. Since the schedules are PCS client scripts that command PCS agents, they automatically take advantage of the full set of tools built into SO-OCS for monitoring and controlling the system. All HK information can be viewed via an InfluxDB backend through the remote web interface using an open source data visualization tool called Grafana. An example Grafana dashboard is shown in Fig. 4.11 with Mod-Cam temperatures and pressures. Dashboards are planned for monitoring all major DAQ subsystems: cryogenic health, detector status, telescope pointing, computer resource usage via a tool called Telegraf, and power distribution for Prime-Cam equipment. The logs for the Docker containers used to operate many of the PCS agents can also be tracked and viewed with a Docker plugin called Loki. All of these tools are described in greater detail for SO-OCS in Refs. [10], [113], and [102]. These Grafana dashboards are separate from the ones for viewing FYST housekeeping data, which are provided by the telescope software team.

In addition to the ability to view the data in Grafana, the raw HK data can always be extracted from the HK g3 files using tools in the `so3g` library.

Grafana also can trigger alarms on a wide array of conditions. For Mod-Cam, we have alerts configured for the oil temperature of the compressors, the temperature of the module focal plane, and other critical cryogenic metrics. Alerts

¹²<https://www.slac.stanford.edu/~abh/bbcp/>

¹³<https://socs.readthedocs.io/en/latest/agents/suprsync.html>

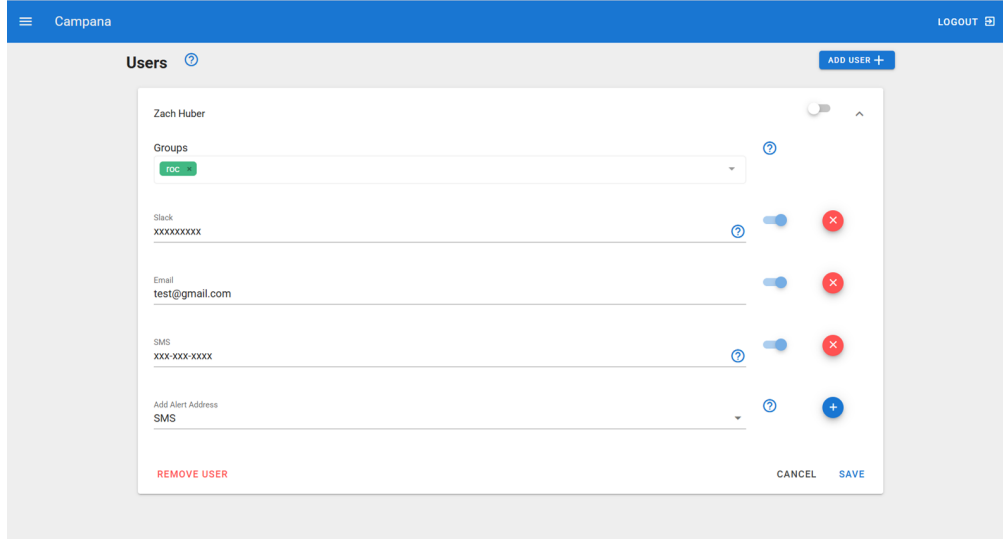


Figure 4.12: The Campana web interface for allowing users to configure alerting methods for the SO site. Users can easily toggle alert methods on and off, and the master switch at the top for the user allows all alerts to be paused when the observer is not on duty.

can also be triggered through the SO-OCS registry agent when individual agents crash. We expect to continue expanding the number and kind of alerts throughout the remainder of the laboratory testing phase and the commissioning phase for Prime-Cam.

We will adopt the SO Campana tool [101] for distributing alerts to ROCs. Campana can send emails, Slack messages, text messages, and phone calls to observers for critical alerts. The method of delivery can be set up in an accessible web interface shown in Fig. 4.12, and observers can toggle alerts on and off so that they do not receive messages when they are not on duty. An alert system will also be in place for observatory-wide alerts from the telescope OCS [112]; the method for the delivery of these alerts to the instrument teams has not been finalized, but it would be possible to store these alerts to the Campana alert database and trigger alerts through Campana for the Prime-Cam team if desired.

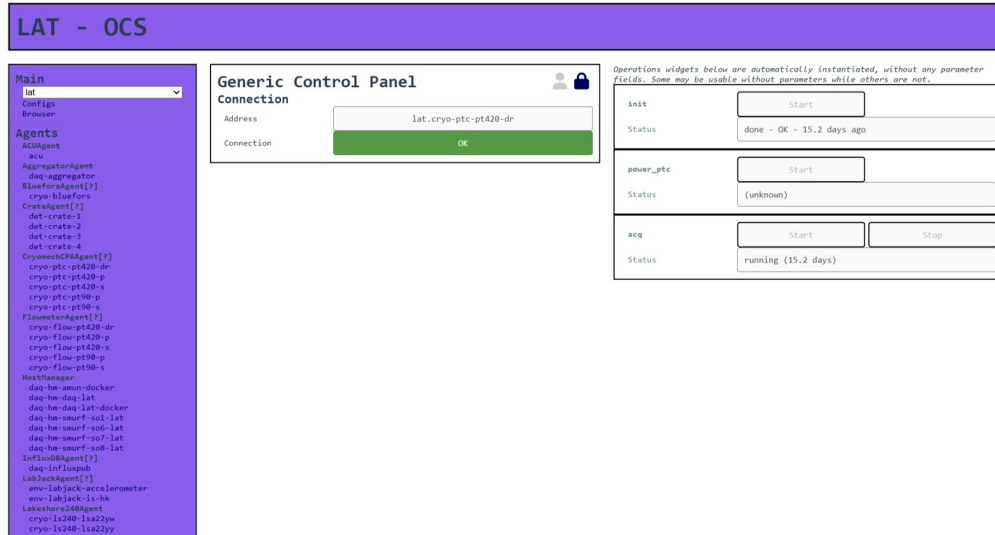


Figure 4.13: Example of the `ocs-web` interface for the SO LAT. This particular panel shows the controls for one of the compressor agents. The left sidebar shows all agents available on the crossbar network. The middle panel shows the current status of the agent and a lock to prevent users from inadvertently changing the state of the agent, and the right panels show various agent functions that can be executed via the interface.

Another tool adapted from SO that gives ROCs easier access to controlling PCS agents is `ocs-web`. This program runs in a Docker container and permits users to view the status of PCS agents and execute short client scripts to call individual agent functions from within a convenient web interface. An example for the SO LAT is shown in Fig. 4.13. Panels already exist for all agents that we share in common with SO, but further development remains to be done to add `ocs-web` panels for PCS-specific agents. This tool complements Nextline – the latter permits more complex client scripts to be run, while the former allows for a simplified approach to specific tasks that is well suited for remote observers who are less familiar with the underlying system. Eventually, we intend to develop a remote observing manual that includes standard operating procedures to try in the event of outages, and these procedures will make use of `ocs-web`.

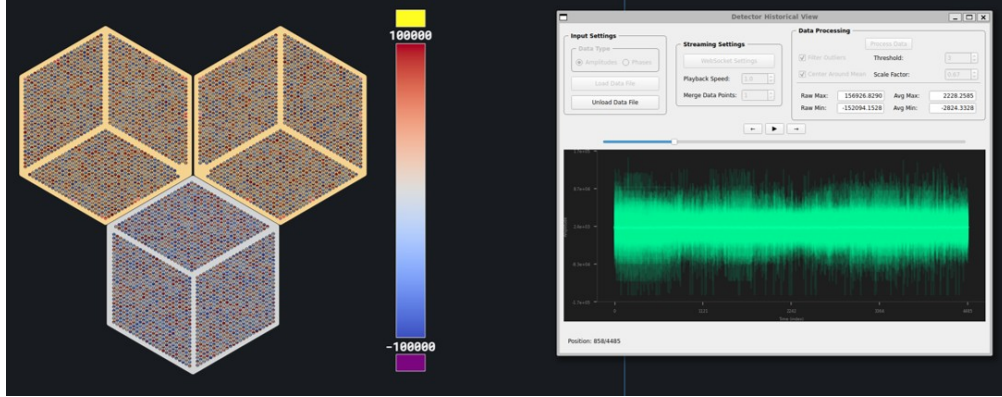


Figure 4.14: A screen capture of the TOD live viewing tool in action. In this example, a Mod-Cam test timestream is loaded in and displayed for all 18 RF networks of the 280 GHz module. The left hand side shows the current value of the timestream in each pixel, while the right hand side shows the timestream and controls. Photo credit: Graham Gibson.

Data Quicklook Tools

In addition to the ability to monitor HK data and the status of the instrument and telescope, it is important to be able to assess the quality of the detector data as it is acquired. Such tools allow subtle issues with data collection to be caught and fixed before a great deal of time is wasted on taking suboptimal data (or no real data at all!).

It is likely that several different tools will eventually be developed to assess data quality for Prime-Cam’s MKIDs. Some of these tools could target TODs. A timestream viewer has already been developed for lab testing that can effectively play a movie of a TOD from a file over the layout of the array as shown in Fig. 4.14. There is interest in integrating this tool with the `rfsoc-streamer` to view the focal plane response via data streaming over the network instead of loading the TOD from a file.

Tools are also being developed in `cctaskidlib` for assessing the health of de-

tector tuning and making plots that summarize key metrics like median detector noise. Some of these tools can be viewed in Grafana, but SO also found it useful to develop a web viewing tool for summary plots saved to the site computers that could be adapted for use with Prime-Cam.

Beyond the TOD level, there is also interest within the CCAT Collaboration in developing quicklook mapmaking tools that would be used to make simple maps on site. These maps would be especially valuable for assessing the quality of small-field observations, but an automated pipeline for generating these maps from each book could prove valuable for monitoring detector health and time-domain science. As of the time of writing, multiple groups are exploring existing tools and possible extensions that could be useful for Prime-Cam. It is also quite possible that a dedicated site analysis server would prove beneficial for the computational requirements of any (semi-)automated quicklook mapmaking done at the site.

4.4 Conclusion

Effective data acquisition hardware and software systems enable every aspect of Prime-Cam’s science; inversely, failures in these systems could cripple the instrument. Over the past two years, a tremendous amount of progress has been made on procuring hardware, adapting or writing software, and integrating these subsystems with Mod-Cam. Wherever possible, we have sought to benefit from the expertise, open-source software, and similarities between the Simons Observatory data acquisitions systems, particularly those of the LATR, and Prime-Cam by adapting much of the same computing resources and HK hardware, the SO-OCS framework for the Prime-Cam Control System, and the data file format and pro-

cessing software. Different detector and warm readout technologies have led to some differences in our systems, as has the need to integrate with the broader CCAT OCS and operational software.

Despite this considerable progress, several key software packages remain in development, including `ccatkidlib`. Integrating `ccatkidlib` operations of the RFSOCs with the sequencer tools and other PCS agents for remote operation and bookbinding, adapting the SO bookbinder, interfacing with the observatory scheduler and OpsDB, and generating PCS client schedules for Prime-Cam all remain high priorities for development and testing. Aspects of the detector hardware might still be changed, especially around the question of an on-site RAID array for expanded storage capabilities, or are still being procured. Questions remain about networking, timing, and consistency of operations. The general plan sketched here will continue to be developed and implemented during the remaining several months of laboratory testing with Mod-Cam and throughout the initial deployment and commissioning of Mod-Cam at the telescope site.

The end products of Prime-Cam’s data acquisition system are neatly packaged books which have been safely transferred off-site. This stage is only the beginning of the journey for the data, however. Whether on-site or at the data center in Cologne, the timestreams will need to be processed to remove artifacts, filter out low-frequency drifts from instrumental and atmospheric stability, and more, likely following the preprocessing pipeline from SO in the `sotodlib` library.¹⁴ The processed timestreams will need to be coupled with calibration information from scanning planets and point sources to convert the raw frequency shift of the resonators into a measurement of absorbed optical power. Along with the raw pointing information and a pointing model for the telescope, the calibrated timestreams will

¹⁴<https://github.com/simonsobs/sotodlib>

then serve as inputs to various mapmaking processes for different science cases. Many of our collaborators have been thinking about these further steps already, but much work remains to be done.

Solid data acquisition procedures provide the foundation for the challenging work of generating good measurements of the cosmos with Prime-Cam. By working consistently and efficiently, this important subsystem enables quick and accurate scientific results for all of Prime-Cam’s science cases. As data collection, calibration, and mapmaking become more efficient, it also could make possible various short-duration time-domain science analyses that are not part of the initial science goals of Prime-Cam but are of great interest to many in the collaboration. Continuing to implement and improve our data acquisition tools, then, is not just an investment in the initial success of CCAT, but also in its long-term growth and legacy.

CHAPTER 5

CONSTRAINING AXION-LIKE DARK MATTER WITH ACT

The development and implementation of instrumentation for our telescopes is fascinating science in its own right and important work, but it is always in service of the observations that we make with our telescopes and the analysis of those data. Much of the work described in the previous two chapters paves the way for observations that will set a new standard for measurements of polarized light from the CMB and other sources over several millimeter and submillimeter wavelength bands, and it was informed throughout by the analysis of ACT polarization data presented here. Understanding the tools we use to analyze our data and how they are impacted by instrumental systematics like beam uniformity, cross-polarization in the optics and detectors, overall detector noise, consistency of observations, and more sheds much light on the need for precise optical characterization, beams, calibration procedures, data storage, and remote operations. The search for axion-like forms of dark matter with ACT data is but one example of the diverse cosmological and astrophysical science that is enabled by excellent measurements of the polarized sky at these wavelengths, and CCAT and SO will continue to drive forward the precision of such data and the insight of new analyses.

5.1 Axion-Like Dark Matter

As detailed in section 1.3, understanding the nature of dark matter is a major priority for modern cosmology and particle physics. Among the many theories attempting to explain dark matter in terms of a new particle beyond the Standard Model of particle physics, much effort has gone into articulating and testing theories that include an ultralight axion-like particle (ALP). The term axion was

initially associated with the QCD axion [117] [118], a proposal to resolve the strong CP problem. While the QCD axion would require specific properties and couplings to the Standard Model, more generic ALP models with similar behavior but a wide range of possible masses and coupling strengths appear in theories beyond the Standard Model like certain string theories [119] [120]. With sufficiently large occupation numbers, even ALPs with masses many orders of magnitude below an eV could account for all the dark matter in the universe (see Ref. [121] and its references for a review of models and cosmological implications).

The standard QCD axion would have a mass larger than approximately 10^{-12} eV, and a wide variety of experiments test such models through their coupling to photons or other Standard Model particles. Select examples include light-shining-through-walls experiments [122] [123], microwave cavity searches [124], and vacuum birefringence searches [125].

If the ALP mass is much smaller than the QCD axion limit, the de Broglie wavelength of the ALP becomes astrophysical in size, leading to a variety of effects. For masses less than roughly 10^{-27} eV, the particle wavelength spans cosmological scales; such ALPs would function as dark energy instead of dark matter, but they are strongly disfavored by searches for their imprint on the CMB power spectrum [126] [127] (see also section 6.3 of Ref. [49]). The parameter space of masses several orders of magnitude heavier than this regime has received intense interest because ALPs with masses of 10^{-23} – 10^{-21} eV could function as fuzzy dark matter (FDM) [128] [129]. FDM proposals seek to resolve possible tensions between the predictions of standard cold dark matter (CDM) models with observations of small-scale structure (roughly less than 10 kpc) in the Milky Way and other galaxies via the small-scale wave-like effects native to ALPs with de Broglie wavelengths around

this length.

One way of constraining ultralight ALP models, therefore, involves measuring small-scale structure formation over time. Example observables include the effect of ALPs on Lyman- α emission [15], the distributions of Milky Way satellite galaxies [14], and galaxy rotation curves [130].

Axion-photon interactions can convert axions into photons or vice versa. Direct detection experiments like CAST search for conversions between solar axions and photons [16] [131], while various astrophysical measurements constrain this effect including searches for excess gamma rays from ALPs produced in supernovae [132] or alterations to the x-ray spectra of active galactic nuclei (AGN) [133] or quasars [134], to name only a few.

Another common method in this mass range is to search for ALP-induced birefringence effects across cosmic scales. Any generic Chern-Simons coupling between a light pseudoscalar boson and electromagnetic fields generates rotation in linearly polarized light as it interacts with the ALP field [135]. The absence of the rotation of light from polarized sources in regions rich in dark matter limits how strongly ALPs could couple to photons. Example studies in this category include pulsar timing and polarization [136] [137] [138], the polarization of AGN jets [139], and the polarization variability of polarized millimeter-wave sources like Tau A [140]. Many of the methods cited here involve extensive modeling of the source environment and Milky Way, which complicate the interpretation of the results despite being done with great care.

The CMB provides a cleaner source for studies of the interaction between ALPs and photons, as pointed out by Fedderke, Graham, and Rajendran [13] (hereafter

referred to as F19). The authors discuss two main effects of ALP-photon interactions on measurements of CMB polarization: the “washout” effect and the “AC oscillation” effect. In the former, variations in the ALP field strength over the period of the decoupling of the CMB cause a range of small rotation angles that sum incoherently on the detectors of CMB telescopes to reduce the overall amplitude of the measured Stokes Q and U parameters. This small amplitude reduction would appear in both the TE and EE spectra, and F19 uses Planck spectra to constrain this effect beyond what ground-based CMB experiments are currently capable of doing with the washout effect. Their final constraints for this method are a reduction of no more than 0.58% in the amplitude in Q and U, which translates to a constraint on the ALP-photon coupling constant of (Eq. 73 of F19):

$$g_{\phi\gamma} \lesssim 9.6 \times 10^{-13} \text{ GeV}^{-1} \times \left(\frac{m_\phi}{10^{-21} \text{ eV}} \right) \times \left(\kappa \times \frac{\Omega_c^0 h^2}{0.11933} \right)^{-1/2} \quad (5.1)$$

where m_ϕ is the ALP mass, κ is the fraction of dark matter that is made up by ALPs, and $\Omega_c^0 h^2$ is the present-day fractional energy density in CDM.

The AC oscillation effect refers to a rotation of the QU or EB rotation angle as a function of time due to the oscillation of the ALP field strength in the Milky Way. Equation 54 of F19 gives the full expression for both effects as:

$$(Q \pm iU)(\hat{n}) = J_0[g_{\phi\gamma}\langle\phi_*\rangle(\hat{n})] \exp \left[\pm 2i \left(\frac{g_{\phi\gamma}}{2} \phi_0 \cos(m_\phi t + \alpha) \right) \right] (Q \pm iU)_0(\hat{n}). \quad (5.2)$$

The AC oscillation part is the oscillation in the argument of the complex exponential. Since the CMB is isotropic and the effect depends only on the local ALP field, this polarization angle oscillation should be both time-dependent and isotropic, distinguishing this measurement from ACT searches for constant and anisotropic cosmic birefringence as described in Refs. [141] and [142]. It should also be coherent across the timescales of CMB experiments and in phase across the sky. The frequency of the oscillation is directly related to the ALP mass, and the period

of oscillation falls between a few hours to a few years for ALP masses between 5×10^{-18} and 1×10^{-23} eV. These timescales align well with regular measurements of CMB polarization by active ground-based experiments, and constraints on ALP-photon coupling in these mass ranges have been placed by BICEP/Keck [143] [144] (hereafter referred to as BXII and BXIV, respectively), SPT-3G [145] (hereafter referred to as S22), and POLARBEAR [146] (hereafter referred to as P23) using this method.

This chapter presents the first major search for an ALP-induced AC oscillation effect in daily maps of CMB polarization made by the Atacama Cosmology Telescope. The data are described in section 5.2. Our estimator based on flat sky power spectra is presented in section 5.3, and the resulting timestream of angles over five years of ACT data is the subject of section 5.4. Tests for per-observation systematics can be found in section 5.5. The remaining steps for putting constraints on ALP models are outlined in section 5.6, though these will be developed and the results presented in detail in a future paper. The entire process is summarized with an eye towards future improvements in section 5.7.

5.2 ACT Data Release 6

The intensity and polarization measurements made with the Advanced ACTPol receiver on ACT from 2017–2022 cover over 19,000 square degrees on the sky at roughly arcminute resolution and were recently released as ACT Data Release 6 (DR6). The data collection, selection, and mapmaking procedure for DR6 is described in detail in Ref. [51] (hereafter referred to as N25).

The DR6 data were collected with three modular optics tubes known as

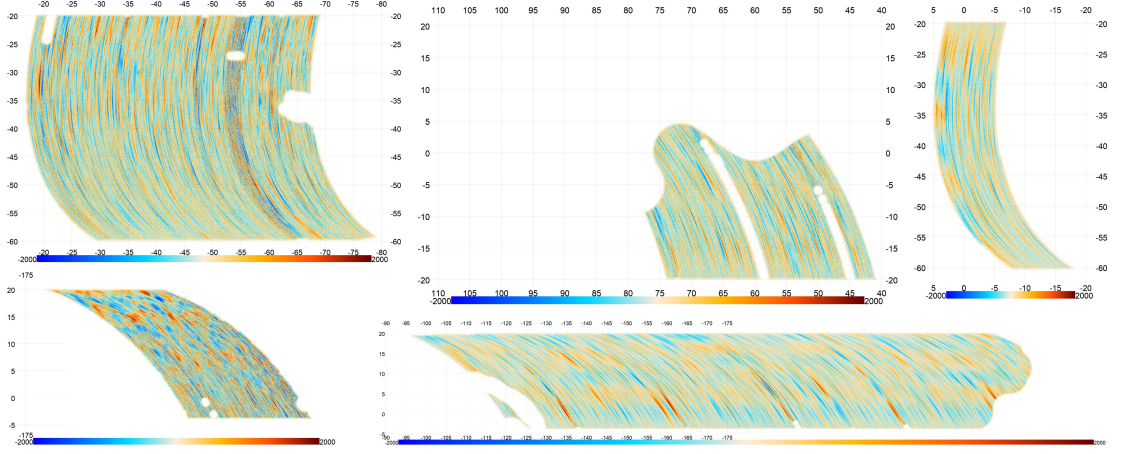


Figure 5.1: Several different f150 depth-1 Stokes Q maps to provide a sense of the wide variety of shapes and sizes. The depth-1 maps can range from a single pass back and forth at constant elevation to many passes over hours. Some of these examples also show evidence of data cuts during the mapmaking process or missing data. The dominant visual effect in these maps is correlated noise along the scan direction.

polarization-sensitive arrays (PAs): PA4, PA5, and PA6.¹ Each PA contained dichroic transition-edge sensors with bandpasses spread across three frequency bands centered at 90, 150, and 220 GHz. We will refer to these as the f090, f150, and f220 bands. PA4 covered f150 and f220, and PA5 and PA6 covered f090 and f150. While PA4 and PA5 observed for the full duration of DR6, PA6 was replaced at the beginning of 2020 with PA7, which observed in bands centered at 30 and 40 GHz.

In order to search for time-varying birefringence due to the impact of ALPs on CMB polarization, we need to group the data into measurements over short timescales. While such a measurement could be made by examining the time-ordered data directly (as originally envisioned by F19), we found it convenient to work instead with the short duration maps known as “depth-1 maps” (read “depth one”) already produced for other astrophysical time-domain analyses with

¹PA1, PA2, and PA3 refer to the optics tubes for the earlier ACTPol receiver. The results for ACTPol were released in DR4 and DR5.

DR6. As described in section 3.7 of N25, these maps constitute the data from a single continuous scan at constant elevation while the sky drifts over the telescope. Depending on the details of the observation and data cuts, the depth-1 maps range in their duration and sky area from as little as 15 minutes of data covering around 100 sq. deg. to several hours of data covering around 2900 sq. deg. The maps cover varying portions of the larger ACT DR6 footprint given in Fig. 1 of N25, effectively sampling a wide range of locations and shapes on the sky. An example depth-1 total intensity map is shown in Fig. 14 of N25, and a series of example depth-1 Stokes Q maps of various shapes and sizes is shown in Fig. 5.1.

Since we expect the ALP polarization rotation signal to be isotropic and coherent for the mass range we will consider, the reality that the depth-1 maps cover different shapes and sizes across a wide sky area is not a problem for our analysis. We expect this signal to appear primarily in CMB-dominated bands, so we restrict our analysis to f090 and f150 depth-1 maps. We refer interested readers to N25 for further details about the depth-1 mapmaking and data selection.

5.3 Method

5.3.1 Polarization Rotation

Equation 5.2 gives the full form of the AC oscillation effect for the CMB in the presence of an ALP field that couples to photons via a Chern-Simons term in the Lagrangian. The time-varying polarization rotation angle has the form:

$$\theta(t) = \frac{g_{\phi\gamma}}{2} \phi_0 \cos(m_\phi t + \alpha), \quad (5.3)$$

where $g_{\phi\gamma}$ is the ALP-photon coupling constant, ϕ_0 is the ALP field value local to the observations, m_ϕ is the mass of the ALP field, and α is the phase of field oscillations. We shall parameterize this rotation in terms of an amplitude A , frequency f , and phase α :

$$\theta = A \sin(2\pi ft + \alpha). \quad (5.4)$$

For each depth-1 map, we measure a rotation angle, generating a timestream covering over five years of data. Sine waves of the form of Eq. 5.4 are fit to this timestream to extract the best fit amplitude and phase for many different frequencies corresponding to different ALP masses. These results then determine the ALP-photon coupling constant if an oscillation is detected or set an upper limit on the coupling constant if no detection is made.

5.3.2 Angle Estimator

Because the depth-1 maps vary considerably in size, shape, and location on the sky, we calculate the rotation angle for each depth-1 map² using an EB power spectrum estimator of the form

$$c_\ell = C'_\ell{}^{E_1 \times B_2} - C'_\ell{}^{E_2 \times B_1}, \quad (5.5)$$

where primed quantities denote rotated quantities; these are the spectra we obtain from the maps as we actually have them, as opposed to ‘primordial’ or theoretical spectra that do not include any rotation effects. For our analysis, index 1 is a depth-1 map, and index 2 is the reference map relative to which we are calculating the angle. We use per-frequency coadded Q and U maps of all night-time DR6

²For convenience, we will speak of each depth-1 observation that generates a T, Q, and U map as a “depth-1 map” except when we need to refer to one of the individual components directly. The same convention will be used for the reference maps.

data as our reference maps.³ These reference maps are masked to cover the same footprint as each depth-1 map for calculating the cross-spectra.

This estimator is proportional to the difference between the rotation angles in the two maps:

$$C'_\ell{}^{E_1 \times B_2} - C'_\ell{}^{E_2 \times B_1} = C_\ell^{EE} \sin[2(\theta_1 - \theta_2)], \quad (5.6)$$

where θ_1 and θ_2 are the rotation angles in maps 1 and 2, respectively, and C_ℓ^{EE} is the best-fit Λ CDM EE spectrum of the CMB. This relation holds under the assumptions that $C_\ell^{EE} \gg C_\ell^{BB}$ and that $C_\ell^{EB} = 0$. The full derivation can be found in Appendix 5.A.1.

This estimator can easily accommodate the diversity in the depth-1 map shapes and can be used to calculate an analytic covariance matrix to set an error bar on each angle measurement without the computational expense of simulations to ascertain the uncertainty. The diagonal elements of the covariance matrix for this estimator are

$$\begin{aligned} \Theta_b = \frac{1}{\nu_b} [& (\langle C'_\ell{}^{E_1 \times E_1} \rangle \langle C'_\ell{}^{B_2 \times B_2} \rangle + \langle C'_\ell{}^{E_1 \times B_2} \rangle^2) \\ & + (\langle C'_\ell{}^{E_2 \times E_2} \rangle \langle C'_\ell{}^{B_1 \times B_1} \rangle + \langle C'_\ell{}^{E_2 \times B_1} \rangle^2) \\ & - 2(\langle C'_\ell{}^{E_1 \times E_2} \rangle \langle C'_\ell{}^{B_2 \times B_1} \rangle + \langle C'_\ell{}^{E_1 \times B_1} \rangle \langle C'_\ell{}^{E_2 \times B_2} \rangle)], \end{aligned} \quad (5.7)$$

which is derived in Appendix 5.A.2. All spectra in this expression can be measured directly from the maps.

Since each depth-1 map contains only a single scan, there is no way to split

³Specifically, we used `act_dr6.02.coadd_AA_night_f150_map_srcfree.fits` from the DR6 release and the equivalent f090 map as the reference maps. These maps come from a later mapmaking run than the one that produced the depth-1 maps, and the pixellization convention of the maps was shifted by half a degree between those runs. In order to process the maps together, we reprojected these maps into the same pixellization as the depth-1 maps. This procedure can introduce a transfer function that reduces power slightly at high ℓ , but this effect should be well less than 1% for the multipole range we use in our analysis.

our data to provide an unbiased estimate of C_ℓ^{EB} if we were calculating a standard EB rotation angle from each depth-1 map alone. Adding a reference map allows us to take cross-spectra that eliminate this noise bias at the expense of calculating the angle relative to the reference map instead of an absolute rotation angle for each depth-1 map. The AC oscillation effect described in F19 is not sensitive to a constant offset in the polarization rotation angle, however, so measuring a relative angle is sufficient for our needs. It also allows us to ignore any systematic constant polarization angle offsets from the instrument that remain in the maps after the mapmaking procedure. Any constant offset from foregrounds in different regions of the sky likewise do not affect this method. Simulations both with and without realistic noise from our depth-1 maps confirm that this method can recover an injected angle within the uncertainty of the measurement.

5.3.3 Power Spectra Calculation

We estimate the various power spectra used in the estimator or its covariance using a flat sky, pseudo- C_ℓ method. Since we are using maps with a relatively small area on the sky, the flat sky approximation should not bias our measurement significantly, especially at the scales we are considering. Any constant offsets introduced by this choice will not affect our search for a time-varying signal.

Each depth-1 Q and U map is loaded into memory and tapered with a mask that goes gently to zero in a 0.5 degree radius around any missing pixels and the edge of the map to prevent ringing in the Fourier transform. The taper follows the equation $w(d) = \frac{d}{r} - \frac{1}{2\pi} \sin(\frac{2\pi d}{r})$ where r is the radius of the taper and d is the distance from the edge of the mask. When $d = 0$ at the boundary of the mask, the taper goes to zero. Likewise, when $d = r$ at the boundary between the taper




Figure 5.2: Examples of depth-1 maps that had missing pixels that led to unusual apodization patterns.

and the interior of mask, the taper goes to one, as expected.

We frequently observe that there are pixels in some of the depth-1 maps that were not observed or that were cut during the mapmaking process. Missing data can result in large portions of the map being apodized; two examples are shown in Fig. 5.2. After considering different options to fill in these holes or otherwise mitigate this issue, we did not find a satisfactory solution and decided to leave the maps as they are. Using a consistent apodization procedure ensures unbiased results even if it is suboptimal by reducing our amount of usable data. From studying a small subset of f150 maps in detail, these issues affect 10% or fewer of the maps, and many of those maps were small or noisy enough that they were unlikely to pass our data cuts (described below), so any effect on our results is small. Of all the arrays, PA6 was affected the most.

Prior to tapering, a mask for the galactic plane is applied as well. Our mask reprojects the Planck PR3 HFI galaxy mask that masks 70% of the galaxy (the same mask used as the starting point in section C.2 of Ref. [141], though without the extension to mask additional dust clouds described in that reference) into the pixellization used by the depth-1 maps. The coadded reference maps have the brightest point sources removed during the mapmaking process, and we do not

apply a further point source mask to these maps (or any point source mask to the depth-1 maps at all) because bright point sources are relatively infrequent in polarization. Any depth-1 map containing a bright point source was likely eliminated by the data cuts described below. Some depth-1 maps fall entirely in the region that is masked by the galaxy mask, and spectra cannot be calculated for these maps. The maps that are not entirely eliminated by the mask and tapering are converted to Fourier space E and B maps and filtered. The filter is the same one applied in ACT DR4 and ACT DR6 (see section 3.3.1 of Ref. [141]) power spectra analyses to eliminate ground pickup by removing Fourier modes with $|\ell_x| < 90$ and $|\ell_y| < 50$. For each observation, the Q and U reference maps are trimmed to the same shape as the depth-1 maps and masked with the same mask before undergoing the same filtering process.

Once the maps are filtered, they are Fourier transformed back to real-space Q and U maps, multiplied by the inverse variance maps for each map inside of the mask footprint, and tapered by 0.5 deg again. The inverse variance map is normalized prior to being applied. Occasionally, the inverse variance maps contained large, unphysical outliers, so the the normalization was done to the 95th percentile of the inverse variance map – any pixel above the 95th percentile was set to one.

Finally, these maps were converted once more to Fourier space. The power spectra were calculated for the relevant pairs of maps and binned radially in bins of width $\Delta\ell = 400$ over the ell range $\ell \in [1000, 3000]$. The minimum ℓ in this range was chosen to match the most conservative baseline polarization range for all arrays used in the main ACT DR6 papers (see Table 1 in Ref. [141]), which ensures that our measurements are minimally contaminated by atmospheric noise at lower multipoles. The bin size was determined by a series of tests such that it

reduced anticorrelations between bins in the measured EB spectra. The maximum ℓ was chosen by balancing the need for at least several bins with the observed residual white noise level at high multipoles in our measured EE spectra.

Following the calculation of the power spectra, several corrections were applied to the binned spectra to account for instrumental effects and biases introduced by our process. We divide each spectrum by the same analytic filtering transfer function, t_b^2 , described in section 6.6 of Ref. [4] to account for the bias introduced by our k-space filtering. It has the form

$$t_b^2(\ell) = 1 - \frac{4(\ell_x + \ell_y)}{2\pi\ell}, \quad (5.8)$$

where ℓ_x and ℓ_y are the cutoff modes in the Fourier-space filter. We use $\ell_x = 90$ and $\ell_y = 50$. This transfer function is calculated for each ℓ and then binned with the same binning scheme as our spectra. We also divide each spectrum by a window factor, w_2 , to correct approximately for the effects of masking and inverse variance weighting (cf. equation A32 in Ref. [147]). This factor is the mean of the product of the inverse variance weighted mask for the two maps that go into the power spectrum calculation. The geometric contribution from the shape of the mask is the same for every spectrum, but the inverse variance weighting contribution varies depending on whether it is a depth-1 autospectrum (in which case it comes from two copies of the depth-1 inverse variance map), a cross spectrum between the depth-1 map and the reference map (one copy of each inverse variance map), or a reference map autospectrum (two copies of the inverse variance map for the reference map). This correction is a single number between 0.1 and 1.0 for each spectrum. Finally, we also divide by a copy of the appropriate beam transfer function for each map, again binned with the same binning scheme as our spectra. These beam transfer functions are publicly-available ACT DR6 data products for

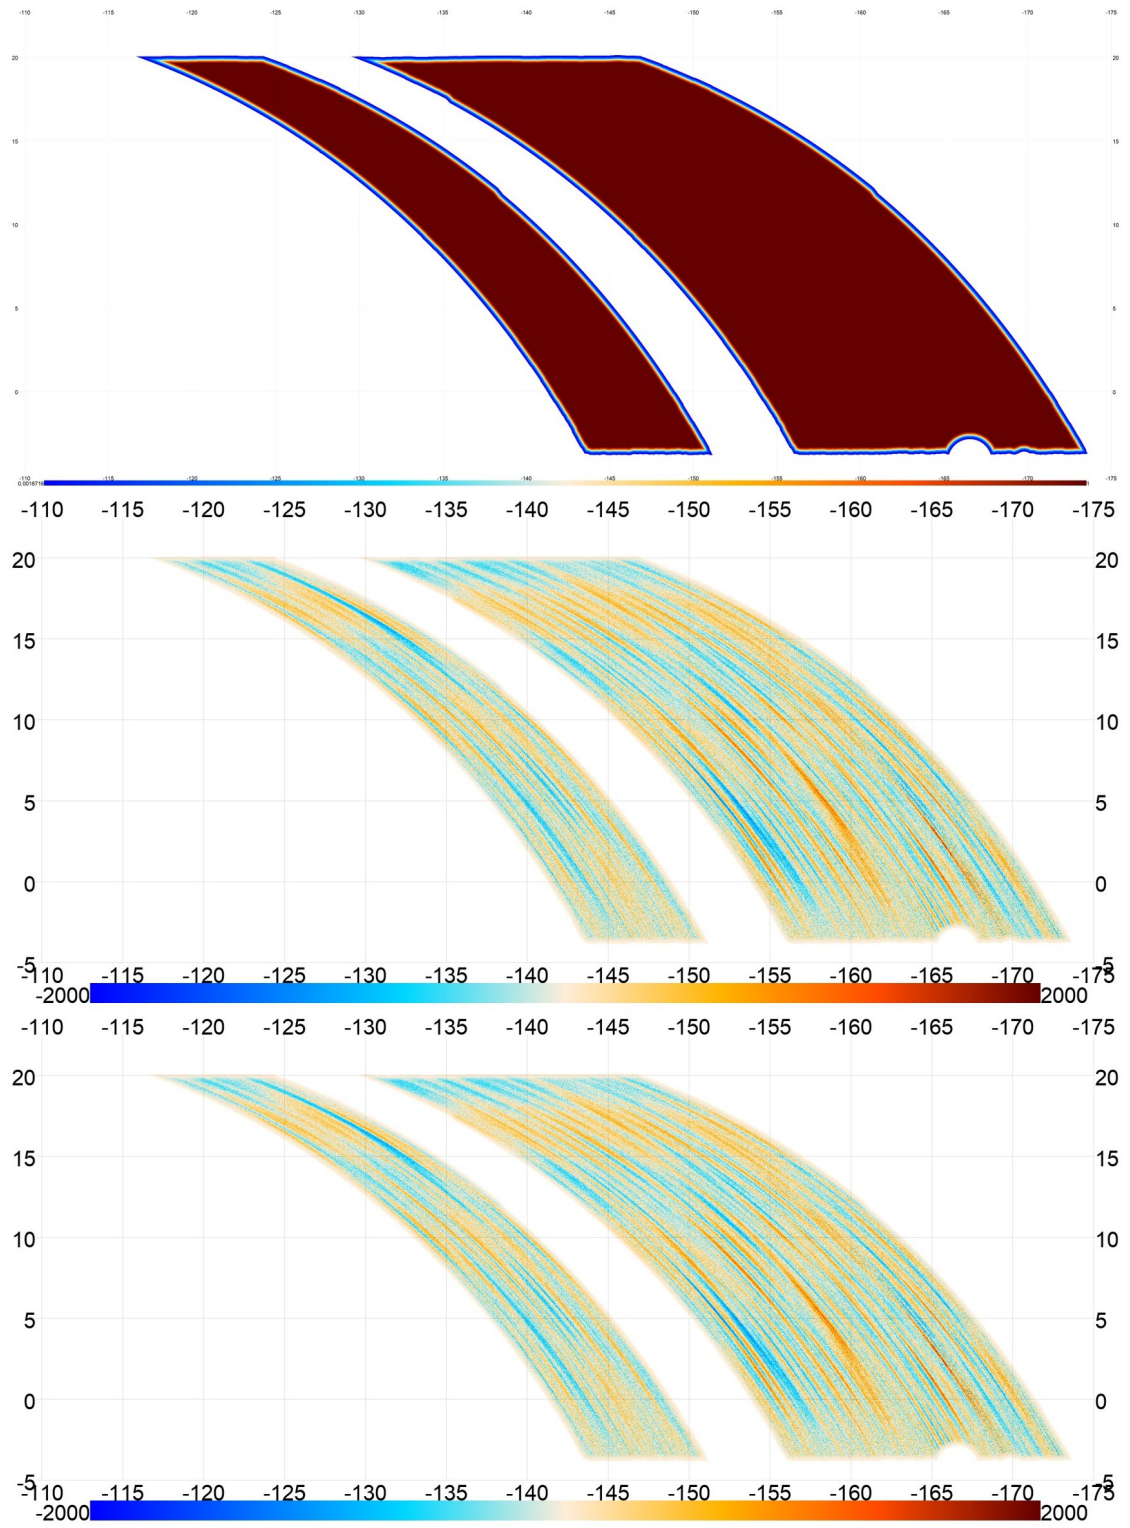


Figure 5.3: An example of the tapered mask (top), depth-1 Q (middle), and depth-1 U (bottom) for a single depth-1 map.

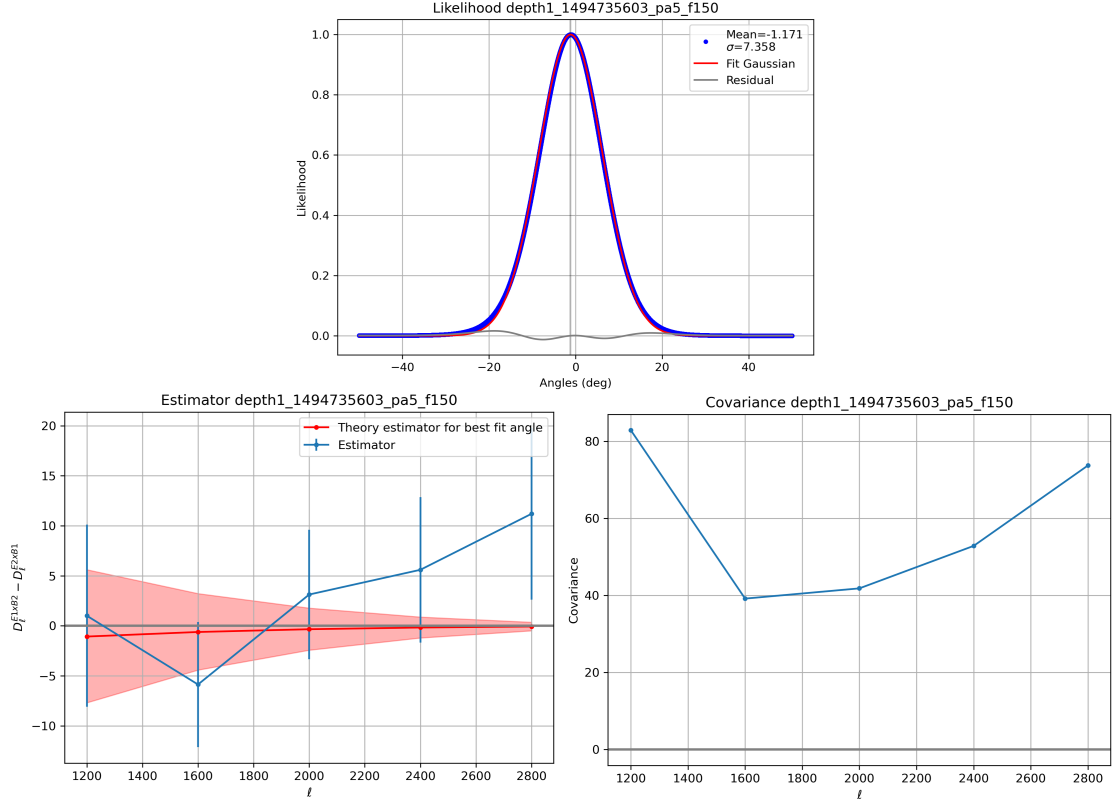


Figure 5.4: An example of the likelihood (top), measured estimator with comparison to the theoretical estimator evaluated at the best fit angle (left), and covariance (right) for the depth-1 map in Fig. 5.3. Each spectrum is plotted as D_ℓ , which is equal to $\frac{\ell(\ell+1)}{2\pi}C_\ell$. The x-axis is ℓ for the spectra.

the reference maps.⁴ The depth-1 maps were produced in a earlier mapmaking process than these publicly-available products. For the depth-1 map beams, we use the equivalent data product from this earlier DR6 mapmaking run that produced the depth-1 maps instead of the publicly-available beam products.

Figs. 5.3, 5.4, and 5.5 show examples for a single depth-1 map of the tapered mask, the original Q and U maps, the likelihood, the estimator and covariance, and many of the individual spectra that go into their calculation.

⁴Specifically, the reference maps were both designed to match the PA5 coadded beam from the data release.

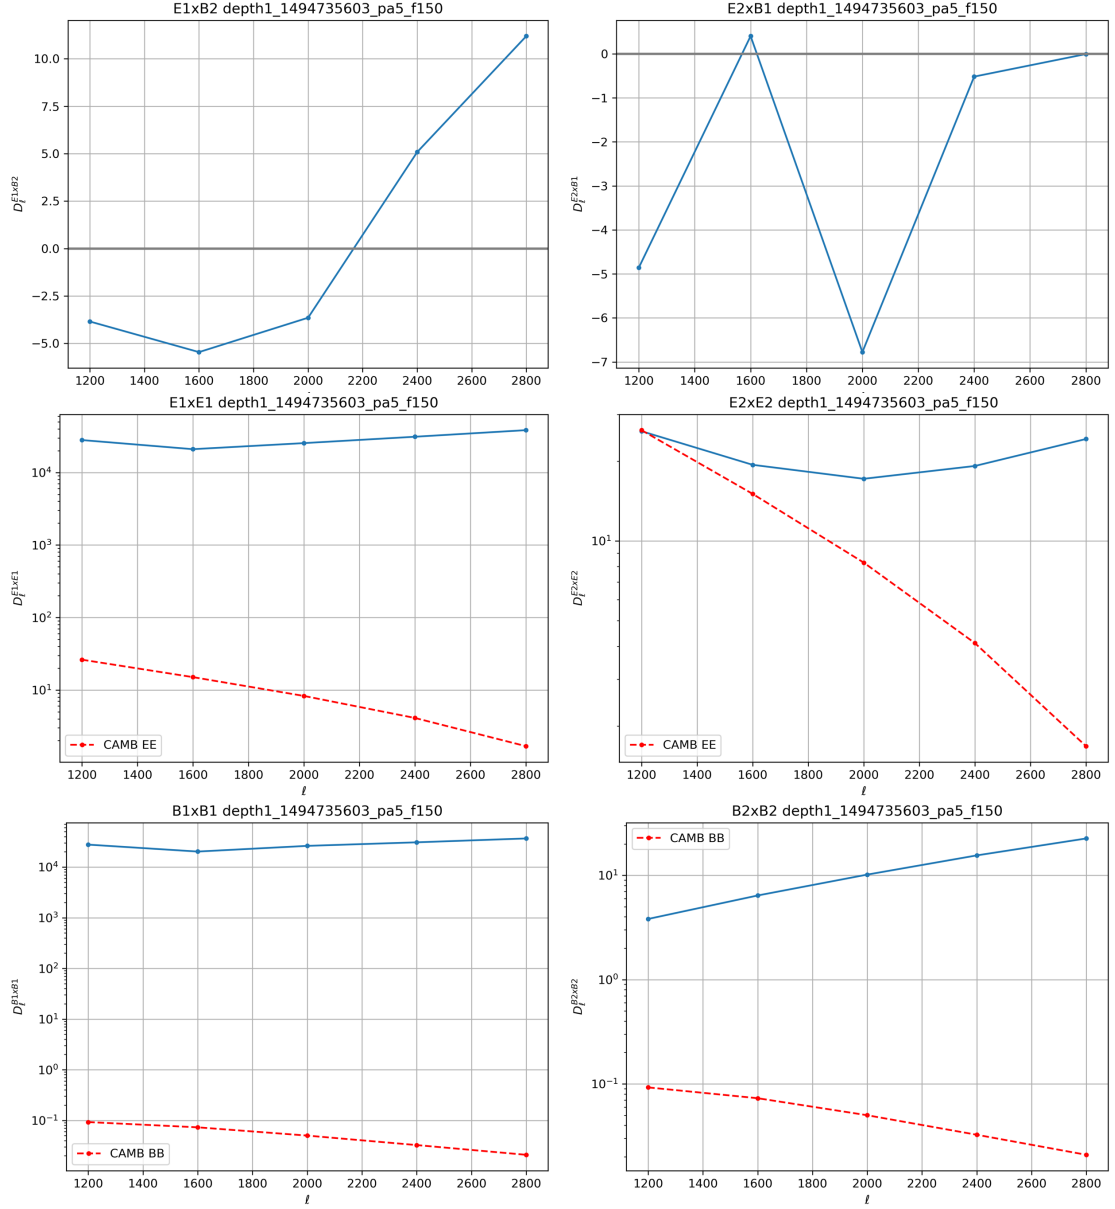


Figure 5.5: An example of the depth-1 E cross reference B spectrum (top left), reference E cross depth-1 B spectrum (top right), depth-1 EE autospectrum (middle left), reference EE autospectrum (middle right), depth-1 BB autospectrum (bottom left), and reference BB autospectrum (bottom right) measured for the depth-1 map in Fig. 5.3. Each spectrum is plotted as D_ℓ , which is equal to $\frac{\ell(\ell+1)}{2\pi}C_\ell$. The x-axis is ℓ .

5.3.4 Angle Extraction

Once the spectra are calculated, the estimator and diagonal elements of the covariance matrix are constructed from the relevant spectra. The covariance term has a mode counting factor, ν_b , that also needs to be corrected for the reduction in total modes due to the filtering transfer function, masking, and inverse variance weighting. This factor begins with the total number of Fourier-space pixels that go into each bin and is multiplied by one copy of the filtering transfer function, t_b , and w_2^2/w_4 , a slightly different window function than the one used to correct the spectrum amplitude (cf. equation 17 in Ref. [147]). Simulations confirmed that this factor better matched the theoretical number of modes for a masked sky. This window function accounts both for the effect of the geometry of the mask and the inverse variance weighting on the number of modes.

These terms are used to construct a likelihood for the best fit angle:

$$\mathcal{L}_{angle} \propto e^{-\frac{1}{2}\chi_{angle}^2}, \quad (5.9)$$

with

$$\chi_{angle}^2 = \sum_b \frac{(c_b - C_b^{EE} \sin(2\phi))^2}{2\Theta_b}, \quad (5.10)$$

where c_b denotes bin b of the binned estimator, C_b^{EE} is a best-fit Planck 2018 EE spectrum from CAMB,⁵ ϕ is the angle we seek to measure, and Θ_b is the relevant bin of the diagonal element of the covariance matrix for this estimator. We will only ever be concerned with relative likelihoods, so we drop the normalization factor in front of the likelihood.

We evaluate this likelihood at many points between -45° and 45° , the principal value domain of the $\sin(2\theta)$ in our estimator, and normalize it so that the peak value

⁵Calculated using the `planck_2018_acc.ini` file on the CAMB GitHub repository for the baseline Planck 2018 LCDM model using `base_plikHM.TTTEEE_lowl_lowE.lensing`.

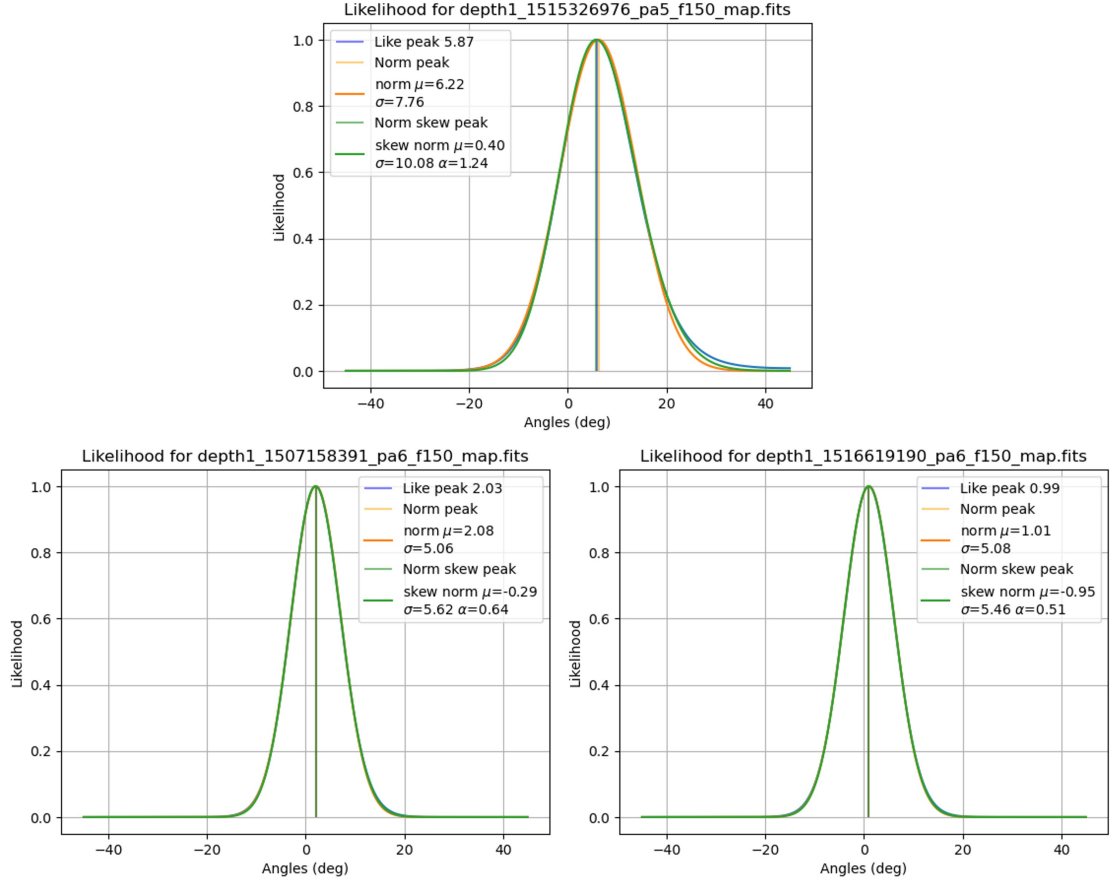


Figure 5.6: Examples of likelihoods for high signal-to-noise maps. The sampled likelihood is shown in blue. These maps form nicely Gaussian like-likelihoods with peaks and widths that are fit equally well by Gaussian (orange) and skew Gaussian (green) curves.

is one. For maps with high signal-to-noise ratio characterized by a relatively small covariance, this likelihood becomes a Gaussian with the best fit angle and error bar corresponding to the mean and standard deviation of the Gaussian. Examples of this group of likelihoods are shown in Fig. 5.6. Many of our depth-1 maps produce such likelihoods, but there are also many depth-1 maps that produce a lower signal-to-noise ratio. As noted in N25, the depth-1 maps are less processed than the large maps used in the ACT DR6 cosmology analyses, so some of the depth-1 maps have beam, pointing, and calibration issues. Given the highly variable nature of the maps, some of them are also quite small or noisy relative to the others. For

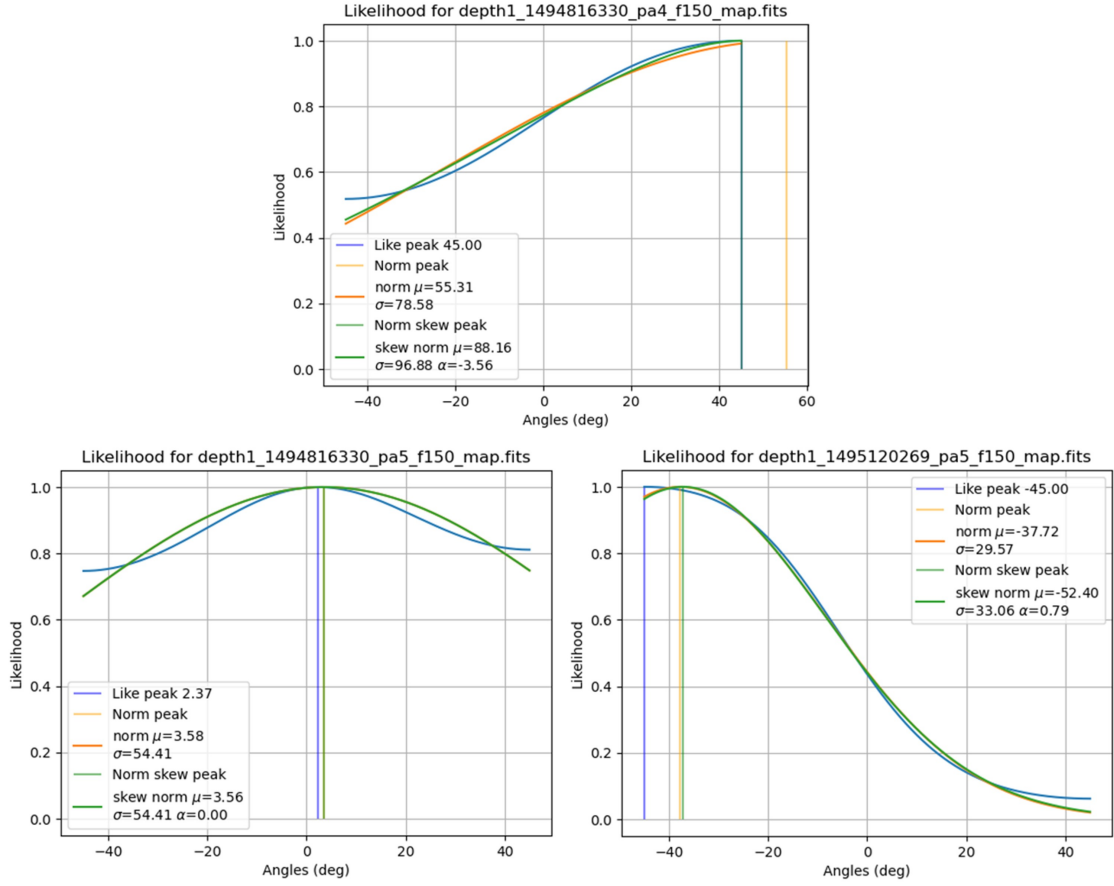


Figure 5.7: Examples of likelihoods for very low signal-to-noise maps. The sampled likelihood is shown in blue. These maps do not isolate a peak well and are cut from our analysis.

many of these maps, the likelihood does not produce a peak at all, and we cut such maps from our analysis. Examples are shown in Fig. 5.7. Some depth-1 maps fall between these extremes: they still peak at some value, but the shape of the likelihood itself is no longer Gaussian. The shape of these likelihoods most frequently takes the form of a central peak with a tail on one side of the peak, as shown in Fig. 5.8. In order to accommodate as much of the data set as possible, we fit all the likelihoods with a skew normal (also called skew Gaussian) distribution and use the peak location as our best fit angle for that map. For the error bar, we directly calculate the FWHM of the skew normal distribution and convert it into

the equivalent Gaussian standard deviation. For the higher signal-to-noise ratio maps, this is equivalent to a regular Gaussian fit. For the lower signal-to-noise ratio maps, this approximates the true FWHM with a Gaussian, but it avoids issues with inaccuracies in finding the peak location that occurred when we fit a Gaussian to these likelihoods directly.




Figure 5.8: Examples of likelihoods for borderline signal-to-noise maps. The sampled likelihood is shown in blue. These maps do seem to have a well-defined peak (almost always at an angle far from zero), but they exhibit non-Gaussian behavior near the boundary of the principal value domain of the likelihood. The peak values of these likelihoods, in particular, are better fit by the skew Gaussian curve (green) than the Gaussian (orange).

After calculating an angle and uncertainty for each map, we use five cuts to eliminate maps with poor quality or other issues from our final timestream:

1. Some of the maps are entirely cut by our galaxy mask. No spectra can be calculated for these, so they are not included in the analysis.
2. For a very small number of maps (less than ten), the fitting tools cannot find a satisfactory fit for a normal or a skew normal distribution to the likelihood. Upon examination, these were usually likelihoods with unusual shapes that do not constrain a peak, so we exclude these.

3. To remove most of the maps that did not form likelihoods with good peaks, we cut all maps with a FWHM greater than 50 degrees for the skew normal distribution. This cut discriminated well between the likelihoods that had no peak and those that did without eliminating any in the latter category.
4. After implementing the previous cuts, inspection of the likelihoods revealed that a small number of maps did not have good likelihoods but had features near the boundary of our fitting range that caused the fitted skew normal distribution to have a lower FWHM. These could be easily removed by cutting all maps with peak angles greater than 30° or less than -30° , as no good likelihoods were found to peak in this range.
5. Finally, among the maps that passed these quality checks, we found that a population of maps from the final observing season in DR6 were much longer in duration than the rest of the depth-1 maps. All of the rest of our maps are 13 hours or less in duration, while some of these were a day or longer in duration with large gaps between data. This longer duration washes out the signal for most frequencies, and the large gaps in these maps raised questions about their data quality and noise level. Since these maps constitute less than half a percent of the total length of the timestream and may have undesirable properties that we do not fully understand, we cut all maps with durations longer than 14 hours.

5.4 Timestream Results

The final timestreams for f090 and f150 are shown in Fig. 5.9. Of the 8664 f090 depth-1 maps in DR6, 5414 (62.5%) pass these cuts. Of the 13920 f150 depth-1

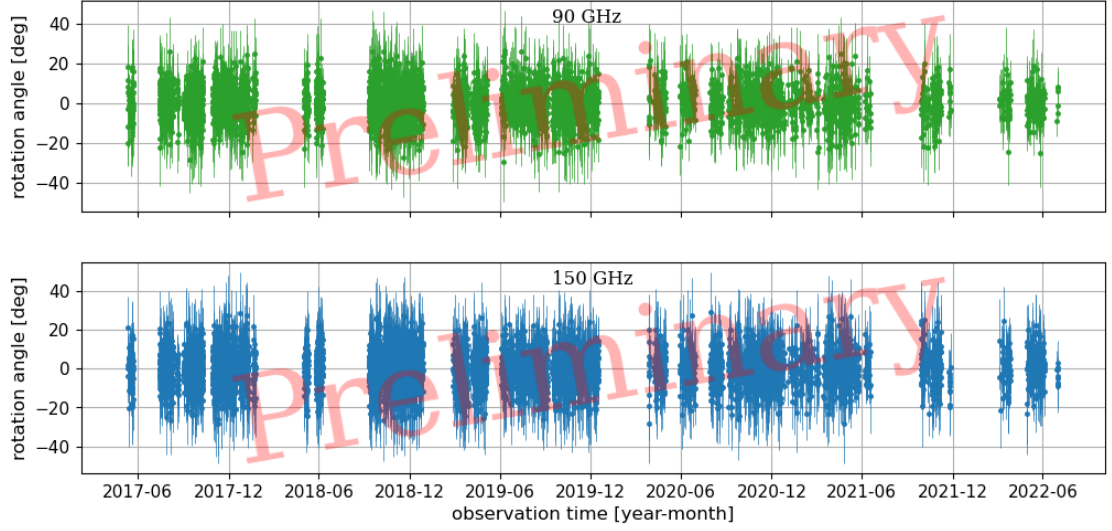


Figure 5.9: The timestream of angles for f090 and f150 from ACT DR6 depth-1 maps. 5414 (62.5%) f090 and 7347 (52.8%) f150 maps passed our likelihood cuts. The error bar comes from approximating the FWHM of each likelihood as a Gaussian.




Figure 5.10: This snippet of the angle timestream shows all f090 and f150 angles labeled by array over a 46.3 day period near the end of 2017.

maps in DR6, 7347 (52.8%) pass these cuts. The large gaps in the timestream correspond to periods when the telescope was not operating due to adverse weather, telescope repairs, or upgrades to the telescope and camera between observing seasons. Fig. 5.10 shows a small portion of these timestreams over nearly seven weeks in late 2017 to make individual data points and error bars clearer.

5.5 Per Observation Systematic Tests

In order to validate our method for calculating angles from each map, we ran a pair of tests for each observation that passed our cuts: a goodness of fit test and an array differencing null test. These tests provided valuable information while identifying bugs in our code and issues with the various correction factors in our method.

For each map, we can assess the goodness of fit of the model (the right hand side of Eq. 5.6) to the estimator by constructing a chi-squared variable that is the difference between the estimator and the model evaluated at the best-fit angle. The covariance for this chi-squared variable is the estimator covariance. For each map that passes our data cuts, we calculate the probability to exceed (PTE) value of this chi-squared variable for $n_{bins} - n_{params} = 5 - 1 = 4$ degrees of freedom. The PTE distributions are shown in Fig. 5.11. Since we expect each angle to be uncorrelated from the others, this chi-squared variable should be randomly distributed. Therefore, we expect a uniform PTE distribution. There is an excess at low PTEs for PA5 at both frequencies, which could result from a subset of PA5 maps that have some unidentified systematic issue or from some underestimate of the errors for all PA5 maps. We have not yet found any clear issue along the lines

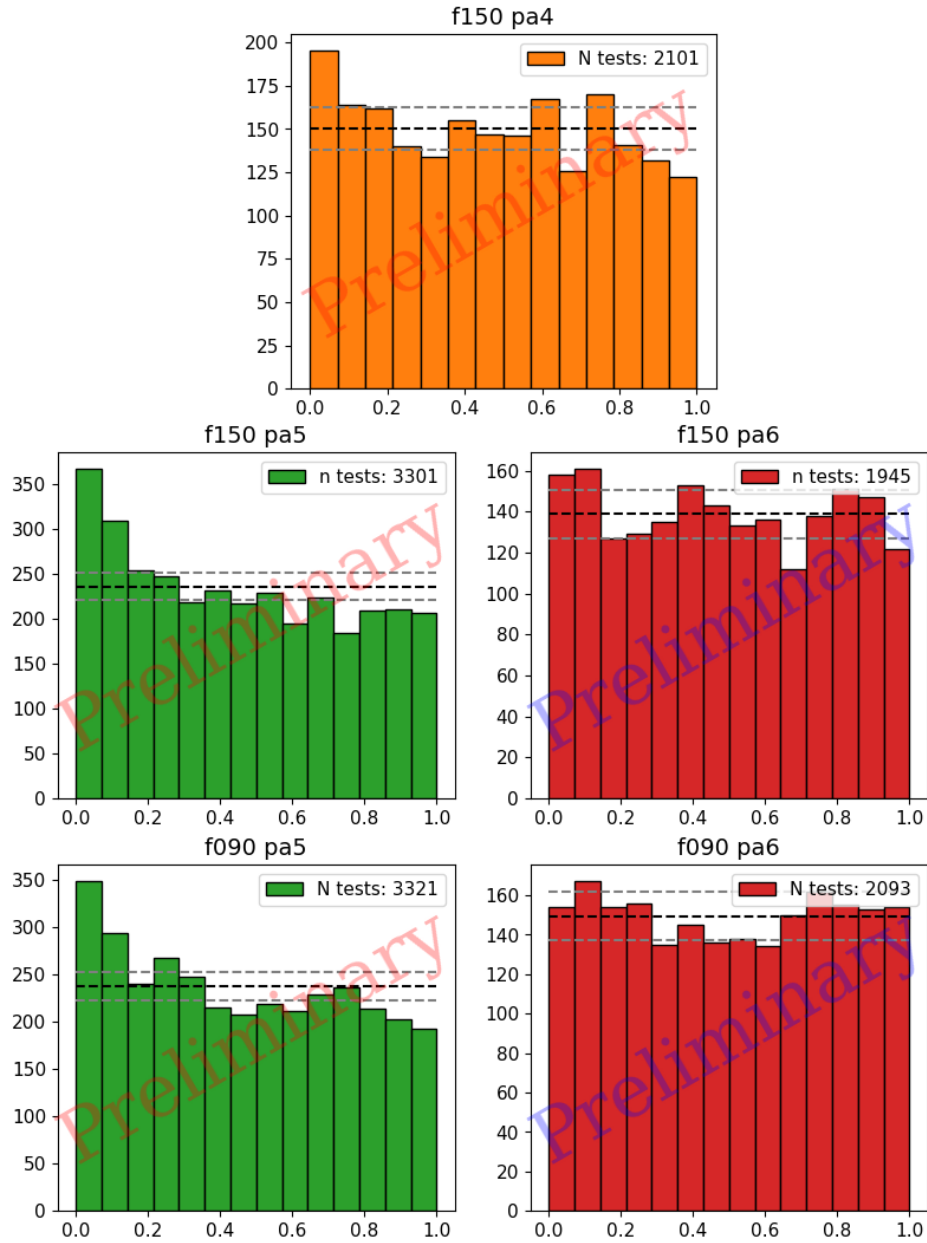


Figure 5.11: PTE distribution (dof=4) for the difference between the estimator and the theoretical estimator evaluated at the best fit angle for f150 and f090 for all arrays. A flat PTE distribution indicates that the model is generally fitting the data well. The black line shows the average per bin for a perfectly flat distribution, while the gray lines show the average plus or minus the square root of the average per bin to give a sense of the expected random variations from uniformity. We see a slight excess at low PTEs in the PA5 data for both frequencies, which could be interpreted as an underestimate of the errors by roughly 5–10%, but this is small enough to proceed with the analysis.




Figure 5.12: PTE distributions (dof=5) for the difference of estimators for measurements of the same area of the sky made by different arrays at the same time for f150 and f090 for all array pairs. The black and gray lines are the same as described in Fig. 5.11.

of the former. If it is the latter, this would correspond to an underestimate of the errors by roughly 5–10%. We consider this to be a small enough effect to proceed with the analysis.

For a given depth-1 map footprint, there are often a pair or group of maps from different arrays that all share similar footprints and are taken at the same time. These maps should yield similar results, so we expect that subtracting their estimators should cancel out any signal. For each pair of such maps from different arrays and frequencies that pass our data cuts, we construct a chi-squared

variable with this null estimator and the sum of the covariances, then calculate the probability to exceed that chi-squared value for $n_{bins} = 5$ degrees of freedom. The PTE distributions are shown in Fig. 5.12 and are consistent with a flat distribution as expected.

5.6 Next Steps

After calculating an angle for each map, several steps remain to translate the angle timestream into constraints on ALP models including fitting the data with sinusoids of variable frequency, studying these results for possible systematic issues, and converting amplitude fits into information about ALP-photon coupling. The remainder of the process and current status of the work is summarized in this section, though final details and results will be deferred to a future paper. Since the work described in this section is either underway or has yet to begin, the plans outlined here should be considered provisional.

5.6.1 Timestream Fitting

Once the timestream has been generated, we use a maximum likelihood estimator to calculate the best fit amplitude for sine waves of different frequencies. Following P23, we incorporate the effect of the finite duration of the measurement for each depth-1 map by integrating Eq. 5.4 over the length of the observation,

$$\begin{aligned}\langle\theta_i\rangle &= \frac{1}{t_i^{end} - t_i^{start}} \int_{t_i^{start}}^{t_i^{end}} dt A \sin(2\pi f t + \alpha) \\ &= A \text{sinc}(\pi f \Delta t_i) \sin(2\pi f \langle t_i \rangle + \alpha).\end{aligned}\tag{5.11}$$

Here, $\Delta t_i = t_i^{end} - t_i^{start}$ is the duration of the depth-1 observation inside of our non-apodized region used for the power spectra calculations, and $\langle t_i \rangle = \frac{t_i^{end} + t_i^{start}}{2}$ is the mean time of the same region.

By construction, the uncertainties on each angle are Gaussian or approximated as Gaussian, so we construct the following chi-squared variable and likelihood:

$$\chi_{ts}^2(A, f, \alpha) = \sum_i \frac{(\phi_i - \langle \theta_i \rangle)^2}{\sigma_i^2} \quad (5.12)$$

and

$$\mathcal{L}_{ts}(A, f, \alpha) \propto e^{-\frac{1}{2}\chi_{ts}^2}, \quad (5.13)$$

where ϕ_i is the measured angle for a single depth-1 map and σ_i is the uncertainty for the angle. The sum is over all f090 and f150 observations that pass cuts. We will only ever be concerned with relative likelihoods, so we drop the normalization factor in front of the likelihood.

In order to evaluate χ_{ts}^2 quickly for a range of amplitudes, frequencies, and phases, we follow BXII in writing this variable as six sums by expanding the square in the numerator and using the sine sum formula. Each of the resulting sums contains only quantities from the depth-1 maps: the angles, error bars, durations, and mean times. These sums can be calculated for each frequency and saved to disk. The full chi-squared variable is a linear combination of the amplitude, various trigonometric functions of the phase, and these sums. We rapidly sample χ_{ts}^2 for amplitudes between $[0.001, 3.0]$ deg in steps of 0.001 deg and phases in the range $[0, 2\pi)$ rad in steps of 1 deg at each frequency. These exact values are subject to change as the analysis proceeds.

For the AC rotation effect described in F19, the phase should be coherent across the sky on the timescales of our experiment, but it should be random and contain

no physical information about the ALP field. We marginalize our likelihood over the phase by integrating with a uniform prior $P(\alpha) = \frac{1}{2\pi}$:

$$\mathcal{L}_{ts}(A, f) = \int_0^{2\pi} P(\alpha) \mathcal{L}_{ts}(A, f, \alpha) d\alpha. \quad (5.14)$$

The maximum value of this marginalized likelihood is the maximum likelihood estimate for the amplitude at a given frequency. We use this marginalized likelihood for the timestream-level null tests described in section 5.6.2 and both the assessment of whether a detection has been made and amplitude constraints in the absence of a detection in section 5.6.3.

Frequency Range

A key consideration in this analysis is determining which oscillation frequencies to sample. We are not sensitive to frequencies corresponding to periods shorter than the duration of our depth-1 maps, which are generally less than 10 hours of data and no greater than 14 hours. The irregular cadence of our observations effectively convolves any regularly sampled oscillation with a complicated window function with structure at periods less than 1 day. To avoid introducing possible oscillatory systematics related to this cadence, we restrict our minimum period to 24 hours; this sets the maximum frequency for our range at 1 inverse day or 0.042 inverse hours.

The minimum frequency is bounded by the length of our timestream, which is nearly 1900 days. We are still exploring whether we can accurately recover unbiased information up to this full range or if we need to limit the range we use.

If our timestream were regularly sampled, as for a standard discrete Fourier transform, the frequency spacing required to extract all available frequency infor-

mation would be the the inverse of the total length of time of the timestream. For an irregularly sampled timestream, the window function introduces correlations between neighboring frequencies that could cause oscillatory signals to be missed if the frequency grid is the same as the discrete Fourier transform (see section 7.1 of Ref. [148] for a discussion of such considerations). To mitigate this issue, we oversample by a factor of 5: $\Delta f = f_{min}/5$.

While the final set of frequencies has yet to be decided, we will likely sample between 4500 and 9000 frequencies from roughly $5 - 10 \times 10^{-4} \text{ days}^{-1}$ to 1.0 days^{-1} . This would correspond to ALP masses from roughly $4 - 8 \times 10^{-24} \text{ eV}$ to $8 \times 10^{-21} \text{ eV}$ with spacing around $1 \times 10^{-24} \text{ eV}$.

5.6.2 Timestream Null Tests

If there are periodic changes in the measurement of polarization due to astrophysical sources or aspects of telescope operations, these could mimic the ALP signal that we are trying to measure. To test for such systematics, we are performing a set of null tests by comparing the best fit amplitudes for splits of the full timestream including both frequencies. The tentative set of these splits is:

1. A temporal split into two roughly equal segments to check for issues with apparent low-frequency oscillations as the receiver is taken out and replaced,
2. Daytime vs nighttime splits to check for systematics related to the diurnal cycle,
3. Moon up vs moon down splits to check for systematics related to the lunar cycle,

4. Elevation splits comparing scans taken at our three standard elevations (40, 45, or 47 deg elevation, with roughly a third of measurements at each elevation) to check for ground pickup systematics,
5. Splits along periodic behavior in the scan strategy since there could be scan strategy related cycles such as always following a rising scan of a field by a setting scan of the same field 24 hours later,
6. And, for completeness, a timestream level null test of f150 vs f090. This would verify the per-observation null test and make sure there is nothing weird happening with the maps that do not have a paired observation to go into the per-observation null test.

For each test, we will perform the method described in section 5.6.1 to identify the best fit amplitude of the marginalized posterior for each split. We then form a test statistic like:

$$T_{null}(f) = \frac{(A_1 - A_2)^2}{[\sigma(A_1^{sims} - A_2^{sims})]^2}, \quad (5.15)$$

where the 1 and 2 refer to the two splits. The frequency-dependent normalization factor is the variance of the simulations described below for each test, though this normalization is subject to change. This test statistic will be large if there is an oscillatory systematic in one of the splits but not the other.

We then compare these T_{null} values at each frequency to T_{null} values at each frequency for 1000 simulated timestreams. Each simulated timestream is drawn from the real data by drawing a new angle from the Gaussian distribution represented by the angle and the uncertainty for every point in the real timestream. We continue to use the same uncertainty value for that new point, and the simulated timestreams are split in the exact same way as the real data. These timestreams have similar statistical properties to the original, but any periodic signal will be

removed. The number of simulated T_{null} values that exceed the real value at a frequency provides a PTE value for that frequency.

For each test, we take the minimum PTE value from all of the frequencies and compare it to the testing threshold of $0.05/n_{tests}$. Due to correlations between frequencies, the distribution is not a random variable, so only consider the minimum PTE instead of the full shape. Each test must have a minimum PTE above this value, and the collection of minimum PTEs should be uniformly distributed. If both of these conditions are met, our timestream-level systematic tests pass.

5.6.3 Amplitude Upper Limits

Once the null tests are complete and show no evidence for large systematics, the method of section 5.6.1 will be applied to the full timestream to identify the best fit amplitude at each frequency. The values of these best fit amplitudes may vary widely over the high number of frequencies sampled. Since we sample so many frequencies, we expect at least some random outliers; even large amplitudes may not be indicative of a detection of an oscillation caused by ALPs. We will form the chi-squared distributed statistic

$$\Delta\chi_{ts}^2(f) = \chi_{ts}^2(A = 0, f) - \chi_{ts}^2(A_{mle}, f) \quad (5.16)$$

to assess whether a detection is made relative to a background-only model ($A = 0$), and we will calibrate this statistic relative to 5000 simulations formed by drawing new angles from the likelihood for each map in the timestream just as in section 5.6.2. We will form a PTE for each frequency and compare the minimum value out of all of the frequencies, which corresponds to the frequency that deviates most from the background, to a detection threshold of 0.05 to assess how

consistent the data are with the background-only model.

If the data are not consistent with the background-only model, we can use the best fit amplitude to quantify the ALP-photon coupling constant at the mass corresponding to the frequency of the oscillation. If they are consistent with the background-only model, we instead will integrate the marginalized likelihood to set Bayesian 95% upper limits on the amplitude of the oscillation at each frequency.

To set the upper limits, we will use a uniform prior for the amplitude in the region sampled by our grid,

$$P(A) = \begin{cases} \frac{1}{A_{max}} & 0 \leq A \leq A_{max} \\ 0 & \text{otherwise.} \end{cases} \quad (5.17)$$

The posterior probability distribution is the normalized product of the prior and the marginalized likelihood given in Eq. 5.14, and this can be integrated to form the cumulative density function

$$\text{CDF}(A, f) = \frac{1}{N} \int_0^A P(A') \mathcal{L}_{ts}(A', f) dA', \quad (5.18)$$

where N is the total area under the posterior probability distribution curve. The 95% upper limit is the value of A for which $\text{CDF}(A, f) = 0.95$ at each frequency.

Another option would be to set frequentist upper limits following the Neyman procedure, as in P23.

5.6.4 ALP-Photon Coupling Upper Limits

Early papers searching for the AC oscillation effect in CMB polarization measurements (BXII, BXIV, and S22) directly converted the upper limits on oscillation

amplitude into upper limits on the ALP-photon coupling constant, $g_{\gamma\phi}$ following equations 48 and 81 of F19. For our case, the relevant equations are

$$A = \frac{g_{\gamma\phi}\phi_0}{2}$$

and

$$\rho_0 = \frac{1}{2}\phi_0^2 m_\phi^2,$$

where ρ_0 is the local dark matter density (taken to be 0.3 GeV/cm^3), ϕ_0 is the local ALP field strength, and m_ϕ is the ALP mass. Solving these for the coupling constant and benchmarking to standard values gives

$$g_{\gamma\phi} = (9.3 \times 10^{-10} \text{ GeV}^{-1}) \times A \times \left(\frac{m_\phi}{10^{-21} \text{ eV}} \right) \times \left(\frac{\kappa \rho_0}{0.3 \text{ GeV/cm}^3} \right)^{-1/2}, \quad (5.19)$$

where κ is the fraction of dark matter made up by the ALP. Analyses of this kind usually assume κ is equal to one.

Following Ref. [149], P23 points out, however, that this deterministic approach is inappropriate for this kind of measurement. An ALP field of fixed mass would oscillate only at one frequency if the motion of the dark matter is neglected. When including the virial velocity distribution of dark matter in the Milky Way, velocity dispersion would cause different frequency modes of the ALP field with random phases to appear and interfere with one another, causing the overall field amplitude to be a stochastic variable with a characteristic coherence time $\tau_c \sim (f v_{MW}^2 / c^2)^{-1}$.

When the ALP coherence time is much shorter than the duration of observation, then the measured amplitude would sample this variable well and would equal the average local dark matter density as in the deterministic case. Direct detection experiments for ALPs heavier than the peV scale often operate in this regime.

For the frequencies we consider (assuming $v_{MW} \sim 10^{-3}c$ and using $f_{max} = 1 \text{ day}^{-1}$), the minimum coherence time is around 2700 years. We clearly fall into

the opposite regime in which the ALP coherence time is much longer than our observations, so we only sample one amplitude from a Rayleigh distribution of stochastic amplitudes for the ALP field. The distribution for the field amplitude takes the form

$$P(\phi_0) = \frac{2\phi_0}{\phi_{DM}^2} e^{-\phi_0^2/\phi_{DM}^2} \quad (5.20)$$

where ϕ_{DM} is the field strength in the deterministic case for which all of the dark matter is in the ALP field.

In this case, the true value for the ALP field amplitude becomes another parameter to marginalize over. A likelihood for the coupling constant as a function of frequency can be constructed, and either frequentist or Bayesian 95% confidence intervals can be set on the ALP-photon coupling constant for our range of masses.

Eventually, assuming there is no detection, we will add a constraint to an exclusion plot like Fig. 5.13, which shows the results from BXIV, S22, and P23 along with select limits from other astrophysical measurements. Though the SPT and BICEP measurements do not take into account the stochastic nature of the ALP field (and thus likely overestimate their bounds), SPT’s higher resolution and larger number of measured polarization modes lead it to the best constraint so far. Since we have a similar resolution, we expect to provide similar constraints over a wider range of ALP masses than SPT using a different method with different systematics.

Fig. 5.14 provides an estimate of the amplitude sensitivity of our data set for comparison to SPT’s stated limit on the oscillation amplitude. To obtain this estimate, the Bayesian methods of sections 5.6.1 and 5.6.3 are applied to a timestream with all the angles set to zero but using the real timing information and error bars from our data set. Such a timestream corresponds to the best possible




Figure 5.13: Constraints on the ALP-photon coupling constant as a function of ALP mass for similar measurements to ours. We will eventually add a line on this plot for our constraints once our null tests are complete, assuming that we do not make a detection. The various colored dashed lines show constraints from AC oscillation searches undertaken by other CMB experiments. The green exclusion region comes from the constraints due to the washout effect explored in Ref. [13] with Planck data. Gray lines and exclusion regions refer to other astrophysical limits (vertical lines) [14] [15] or direct detection limits (horizontal line) [16]. Not shown are time-varying polarization rotation methods that use a different source than the CMB.

limits from our data in the case where there is exactly no oscillation. Since we do not use the real angles, this test does not constitute an unblinding of our real dataset. For this estimate, the frequency range spans the full range of possibilities from the inverse of the full length of the timestream to one inverse day; we may choose a different frequency range from the one used here for the analysis of our real timestream.




Figure 5.14: A forecast of the best-case 95% upper limits that we could place with our data set obtained by using a timestream with all angles set to zero but the real error bars and timing information from our final timestream. At low frequencies (below roughly 0.2 inverse days), the limits are approximately constant, though there is some scatter at the lowest frequencies that have only a small number of periods in the full length of the timestream. The limits degrade by up to 13% at the highest frequency considered in this test due to the finite duration of our observations. The step-like nature of the limits comes from the sampling resolution of the amplitude in this test.

Three things stand out in our best-case upper limits. First, the step-like nature of the limits here comes from the resolution of our amplitude sampling. We plan on sampling at 0.001 deg, though that is subject to change in the final analysis. Here, I switched to 0.0005 deg to show the differences between frequencies more clearly. Second, there is some numerical fluctuation in the limits at very low frequencies that could be related to having only a few periods of oscillation for these frequencies or large gaps in the timestream due to weather events and other interruptions in the observations. Before the scatter at the lowest frequencies, though, the limits

are effectively constant, as expected. Third, there is a degradation in the limits as we move to higher frequencies because the finite duration of the observations averages out the signal at high frequencies. This degradation is no more than 13% at the highest frequency that we consider here.

SPT’s median upper limit after accounting for the finite duration of their observations is 0.151 deg (see section IV of S22). This quick test indicates that our median limits from a significantly different data set and analysis method could be within about 10% of SPT’s at lower frequencies, assuming that we do not make a real detection, but we will have to complete the rest of our null tests and unblind before we can place our limits on Fig. 5.13 and know exactly what part of parameter space our real results rule out.

5.7 Conclusion

If parity violating ultralight ALPs constitute most of the local dark matter and couple appreciably to light, then we should be able to detect their faint, periodic imprints on the polarization of the CMB. Thus far, our search for such oscillations in ACT data has yielded a timestream of angles from all f090 and f150 depth-1 maps over the five years covered by DR6. Timestream-level null tests to search for systematic effects that could mimic the ALP signal are underway, and a detection or constraint on oscillation amplitudes across a wide range of ALP masses will follow soon after those null tests are complete.

In addition to searching for ALPs, this timestream may also be useful for assessing the stability of the polarization angle of the instrument over time or identifying possible sources of systematic uncertainty in the polarization angle for

each array, which could prove useful in broader searches for constant or anisotropic birefringence in ACT data. The possibility of a small cosmic birefringence angle from WMAP and Planck data raised by Ref. [150] has led to great interest in recent years in understanding the polarization calibration of other CMB instruments. As discussed in section 4.4 of Ref. [141], some amount of miscalibration may remain in the ACT polarization maps after our optical modeling. Any tool that could help us to better understand whether the source of this polarization angle is from the instrument or an astrophysical (or cosmological) source would be valuable.

The SO LAT also could produce depth-1 maps, and many of the same methods and considerations map directly onto studies of polarization angles in the LATR. With many more arrays and a greater survey depth, the constraints on time-variable parity violating effects from the LATR will improve significantly on our results from ACT. In the first five-year survey with the LATR, there will be twelve dichroic arrays operating at f090 and f150 for a total of twenty-four independent measurements of the sky. These arrays will have more detectors with comparable or better noise performance to the ACT arrays, so we expect the sensitivity of each depth-1 map to improve. Even if the sensitivity of each measurement were exactly the same as ACT, having twenty-four depth-1 maps for each measurement instead of at most five in this analysis should yield roughly five times as many depth-1 maps. If there is no ALP signal, the timestream is primarily limited by random noise in the angle measurement, and we could expect improvements on our ACT constraints by at least a factor of $\sqrt{24/5} = 2.2$. In reality, improvements to the noise of each measurement, observing efficiency, depth-1 mapmaking, and our analysis method likely will make this factor even larger.

For future analyses using this method, it will be worth considering how to

improve the regularity and yield of the maps that pass cuts. A significant amount of time and effort was devoted to trying various methods to use more data, including the skew normal fitting of the likelihood and testing several reference maps. The dominant term in the covariance is the product of the depth-1 map BB spectrum with the reference map EE spectrum. The depth-1 BB spectrum is primarily noise, so more regular observations and mapmaking techniques that avoid excessively small or noisy depth-1 maps would improve yield. The reference map EE spectrum was composed of a mixture of signal and noise. Reducing the noise in this map significantly reduced the number of cut maps and was the driving factor behind using the full coadd of all ACT DR6 night-time data instead of using separate reference maps made of all the night-time data for each array. It would be worth considering whether there is a deeper reference map that could be used or if there are other ways to improve this estimator to make it more robust.

We used the depth-1 maps because they provided the shortest duration maps available from the regular ACT DR6 mapmaking. For time-domain studies, some three-day maps were also made. At the expense of being able to sample a lower maximum frequency, it could be interesting to carry out a preliminary study using this method with those maps to see the effect of using slightly larger and deeper maps on cuts and results. This approach might produce results that are better behaved and easier to work with, especially if the SO LATR produces them as well. Any other short duration maps produced on a regular cadence by the SO LATR could be used for a similar analysis.

While this method could not be used directly for CMB polarization studies with CCAT because CCAT does not observe at f090 and f150, studies of polarized foregrounds at higher frequencies could inform masking choices for similar studies

with SO. Measurements with Prime-Cam should also prove useful for a broader suite of studies of parity violation and magnetic fields in galaxy clusters and our own Milky Way.

5.A Appendix: Derivation of Estimator and Covariance

5.A.1 Estimator Derivation

The derivation of our power spectrum angle estimator (Eq. 5.6) is as follows: The rotation of the E and B coefficients by a nonzero EB mixing angle is

$$a'_\ell{}^E = \cos(2\theta)a_\ell{}^E + \sin(2\theta)a_\ell{}^B \quad (5.21)$$

$$a'_\ell{}^B = -\sin(2\theta)a_\ell{}^E + \cos(2\theta)a_\ell{}^B \quad (5.22)$$

and the conversion to the power spectrum is

$$C_\ell^{EE} = \langle a'_\ell{}^E a'_\ell{}^E \rangle, \quad (5.23)$$

where E can be replaced by T, B, Q, or U as needed.

Expanding out the two power spectra from our estimator in terms of the unrotated (unprimed) coefficients,

$$\begin{aligned} C_\ell'^{E_1 \times B_2} &= \langle a_\ell'^{E_1} a_\ell'^{B_2} \rangle \\ &= \langle [\cos(2\theta_1)a_\ell^{E_1} + \sin(2\theta_1)a_\ell^{B_1}] [-\sin(2\theta_2)a_\ell^{E_2} + \cos(2\theta_2)a_\ell^{B_2}] \rangle \\ &= \langle -\sin(2\theta_2)\cos(2\theta_1)a_\ell^{E_1}a_\ell^{E_2} - \sin(2\theta_1)\sin(2\theta_2)a_\ell^{E_2}a_\ell^{B_1} + \\ &\quad \cos(2\theta_1)\cos(2\theta_2)a_\ell^{E_1}a_\ell^{B_2} + \sin(2\theta_1)\cos(2\theta_2)a_\ell^{B_1}a_\ell^{B_2} \rangle \end{aligned}$$

and

$$\begin{aligned}
C_\ell'^{E_2 \times B_1} &= \langle a_\ell'^{E_2} a_\ell'^{B_1} \rangle \\
&= \langle [\cos(2\theta_2) a_\ell^{E_2} + \sin(2\theta_2) a_\ell^{B_2}] [-\sin(2\theta_1) a_\ell^{E_1} + \cos(2\theta_1) a_\ell^{B_1}] \rangle \\
&= \langle -\sin(2\theta_1) \cos(2\theta_2) a_\ell^{E_2} a_\ell^{E_1} - \sin(2\theta_2) \sin(2\theta_1) a_\ell^{E_1} a_\ell^{B_2} + \\
&\quad \cos(2\theta_2) \cos(2\theta_1) a_\ell^{E_2} a_\ell^{B_1} + \sin(2\theta_2) \cos(2\theta_1) a_\ell^{B_2} a_\ell^{B_1} \rangle.
\end{aligned}$$

Taking the difference between these two expressions and allowing the angle brackets to convert products of coefficients back into power spectra gives

$$\begin{aligned}
C_\ell'^{E_1 \times B_2} - C_\ell'^{E_2 \times B_1} &= -\sin(2\theta_2) \cos(2\theta_1) C_\ell^{E_1 \times E_2} - \sin(2\theta_1) \sin(2\theta_2) C_\ell^{E_2 \times B_1} \\
&\quad + \cos(2\theta_1) \cos(2\theta_2) C_\ell^{E_1 \times B_2} + \sin(2\theta_1) \cos(2\theta_2) C_\ell^{B_1 \times B_2} \\
&\quad + \sin(2\theta_1) \cos(2\theta_2) C_\ell^{E_1 \times E_2} + \sin(2\theta_1) \sin(2\theta_2) C_\ell^{E_1 \times B_2} \\
&\quad - \cos(2\theta_1) \cos(2\theta_2) C_\ell^{E_2 \times B_1} - \sin(2\theta_2) \cos(2\theta_1) C_\ell^{B_1 \times B_2}.
\end{aligned}$$

Combining like terms,

$$\begin{aligned}
C_\ell'^{E_1 \times B_2} - C_\ell'^{E_2 \times B_1} &= [\sin(2\theta_1) \cos(2\theta_2) - \cos(2\theta_1) \sin(2\theta_2)] C_\ell^{E_1 \times E_2} \\
&\quad + [\sin(2\theta_1) \cos(2\theta_2) - \cos(2\theta_1) \sin(2\theta_2)] C_\ell^{B_1 \times B_2} \\
&\quad + [\sin(2\theta_1) \sin(2\theta_2) + \cos(2\theta_1) \cos(2\theta_2)] C_\ell^{E_1 \times B_2} \\
&\quad - [\sin(2\theta_1) \sin(2\theta_2) + \cos(2\theta_1) \cos(2\theta_2)] C_\ell^{E_2 \times B_1}.
\end{aligned}$$

Using the difference formulas for sine and cosine, these can be rewritten as

$$\begin{aligned}
C_\ell'^{E_1 \times B_2} - C_\ell'^{E_2 \times B_1} &= \sin[2(\theta_1 - \theta_2)] C_\ell^{E_1 \times E_2} \\
&\quad + \sin[2(\theta_1 - \theta_2)] C_\ell^{B_1 \times B_2} \\
&\quad + \cos[2(\theta_1 - \theta_2)] C_\ell^{E_1 \times B_2} \\
&\quad - \cos[2(\theta_1 - \theta_2)] C_\ell^{E_2 \times B_1}.
\end{aligned}$$

These spectra are now unrotated “primordial” quantities, so we expect $C_\ell^{E_1 \times B_2} = C_\ell^{E_2 \times B_1} = 0$. In Λ CDM, $C_\ell^{E_1 \times E_2} \gg C_\ell^{B_1 \times B_2}$, so we will neglect $C_\ell^{B_1 \times B_2}$.

In this case, our estimate reduces to Eq. 5.6 above. Both quantities on the left hand side of the equation can be estimated directly from the depth-1 maps and corresponding full-season coadds, while the right hand side is either the angle that we are trying to measure or the Λ CDM prediction for C_ℓ^{EE} .

5.A.2 Covariance Derivation

To estimate the uncertainty in this estimator, we construct the covariance matrix as follows.

$$\Theta = \langle [C_\ell'^{E_1 \times B_2} - C_\ell'^{E_2 \times B_1} - \langle C_\ell'^{E_1 \times B_2} - C_\ell'^{E_2 \times B_1} \rangle] \times [C_\ell'^{E_1 \times B_2} - C_\ell'^{E_2 \times B_1} - \langle C_\ell'^{E_1 \times B_2} - C_\ell'^{E_2 \times B_1} \rangle] \rangle. \quad (5.24)$$

Expanding out the product and rearranging terms, this longer covariance term can be written as the sum of three simpler ones. The steps are as follows:

$$\begin{aligned} \Theta &= \langle C_\ell'^{E_1 \times B_2} C_\ell'^{E_1 \times B_2} - 2C_\ell'^{E_1 \times B_2} C_\ell'^{E_2 \times B_1} + C_\ell'^{E_2 \times B_1} C_\ell'^{E_2 \times B_1} \\ &\quad + \langle C_\ell'^{E_1 \times B_2} \rangle^2 - 2\langle C_\ell'^{E_1 \times B_2} \rangle \langle C_\ell'^{E_2 \times B_1} \rangle + \langle C_\ell'^{E_2 \times B_1} \rangle^2 \\ &\quad - 2C_\ell'^{E_1 \times B_2} \langle C_\ell'^{E_1 \times B_2} \rangle + 2C_\ell'^{E_1 \times B_2} \langle C_\ell'^{E_2 \times B_1} \rangle \\ &\quad + 2C_\ell'^{E_2 \times B_1} \langle C_\ell'^{E_1 \times B_2} \rangle - 2C_\ell'^{E_2 \times B_1} \langle C_\ell'^{E_2 \times B_1} \rangle \rangle, \\ \Theta &= \langle C_\ell'^{E_1 \times B_2} C_\ell'^{E_1 \times B_2} - 2C_\ell'^{E_1 \times B_2} \langle C_\ell'^{E_1 \times B_2} \rangle + \langle C_\ell'^{E_1 \times B_2} \rangle^2 \rangle \\ &\quad + \langle C_\ell'^{E_2 \times B_1} C_\ell'^{E_2 \times B_1} - 2C_\ell'^{E_2 \times B_1} \langle C_\ell'^{E_2 \times B_1} \rangle + \langle C_\ell'^{E_2 \times B_1} \rangle^2 \rangle \\ &\quad - 2\langle C_\ell'^{E_1 \times B_2} C_\ell'^{E_2 \times B_1} - C_\ell'^{E_1 \times B_2} \langle C_\ell'^{E_2 \times B_1} \rangle \\ &\quad - C_\ell'^{E_2 \times B_1} \langle C_\ell'^{E_1 \times B_2} \rangle + \langle C_\ell'^{E_1 \times B_2} \rangle \langle C_\ell'^{E_2 \times B_1} \rangle \rangle, \end{aligned}$$

and

$$\begin{aligned}\Theta = & \langle [C'_\ell{}^{E_1 \times B_2} - \langle C'_\ell{}^{E_1 \times B_2} \rangle] \times [C'_\ell{}^{E_1 \times B_2} - \langle C'_\ell{}^{E_1 \times B_2} \rangle] \rangle \\ & + \langle [C'_\ell{}^{E_2 \times B_1} - \langle C'_\ell{}^{E_2 \times B_1} \rangle] \times [C'_\ell{}^{E_2 \times B_1} - \langle C'_\ell{}^{E_2 \times B_1} \rangle] \rangle \\ & - 2 \langle [C'_\ell{}^{E_1 \times B_2} - \langle C'_\ell{}^{E_1 \times B_2} \rangle] \times [C'_\ell{}^{E_2 \times B_1} - \langle C'_\ell{}^{E_2 \times B_1} \rangle] \rangle.\end{aligned}$$

Using the notation of Appendix A (the unlabeled equation between A4 and A5) in Ref. [151], each row in the last equation above can be written as a simpler covariance matrix:

$$\Theta = \Theta^{E_1 \times B_2; E_1 \times B_2} + \Theta^{E_2 \times B_1; E_2 \times B_1} - 2\Theta^{E_1 \times B_2; E_2 \times B_1}. \quad (5.25)$$

Each of these covariance matrices can be simplified by using equation A7 from Ref. [151], yielding the final result:

$$\begin{aligned}\Theta = & \frac{1}{\nu_b} [(\langle C'_\ell{}^{E_1 \times E_1} \rangle \langle C'_\ell{}^{B_2 \times B_2} \rangle + \langle C'_\ell{}^{E_1 \times B_2} \rangle^2) \\ & + (\langle C'_\ell{}^{E_2 \times E_2} \rangle \langle C'_\ell{}^{B_1 \times B_1} \rangle + \langle C'_\ell{}^{E_2 \times B_1} \rangle^2) \\ & - 2(\langle C'_\ell{}^{E_1 \times E_2} \rangle \langle C'_\ell{}^{B_2 \times B_1} \rangle + \langle C'_\ell{}^{E_1 \times B_1} \rangle \langle C'_\ell{}^{E_2 \times B_2} \rangle)],\end{aligned} \quad (5.26)$$

where ν_b is the effective number of pixels in the annulus of Fourier k-space for bin b. The ν_b term needs to be corrected for masking and filtering by multiplying by w_2^2/w_4 (cf. equation 17 in Ref. [147]), which accounts for both f_{sky} and the inverse variance weighting, and the filtering transfer function, t_b , as in the text after Eq. 1 in Ref. [4]. Symbolically, the final ν_b is given by $\nu_b = \nu_{b,i} t_b \frac{w_2^2}{w_4}$. All of the quantities here are spectra involving only the rotated maps, so they can all be directly calculated from our maps using the flat-sky psuedo- C_ℓ method and corrected for the filtering transfer function, beams, and inverse variance weighting as described in the main body of the chapter.

CHAPTER 6

CONCLUSION

Modern cosmology and astronomy are rich with powerful instruments and precise data like never before in history. ACT has contributed to this wealth with an unmatched depth for a field covering a large fraction of the sky from a ground-based microwave observatory, and its CMB temperature, polarization, and gravitational lensing maps continue to provide new insights into the past, present, and future of the universe that we call home. On the whole, ACT DR6 results match the Λ CDM model to excellent precision, though we continue to use ACT data for analyses like the ALP analysis presented in the previous chapter that probe the remaining mysteries beyond our best fit cosmological model. Cross-correlations with measurements of the LSS and other cosmological observables across the entire electromagnetic spectrum also continue to drive fruitful inquiry into cosmology and astrophysics with ACT data.

The next decade, however, promises only to increase the deluge of high precision cosmological data delivered by powerful ground-based experiments. In standard CMB studies alone, SO, ASO, and upgrades to the various telescopes at the South Pole promise to surpass the noise levels and sky coverage of the previous generation of telescopes within only a few years. CCAT will fill a hole in the submillimeter regime with its outstanding site and ability to probe small cosmological scales over relatively wide fields, enabling improved studies of dust foregrounds, the Milky Way, early dusty star-forming galaxies, the epoch of reionization, and more. We will also deploy the largest number of MKIDs and RFSoc FDM readout chains ever used for astronomical instrumentation, helping to pave the way for future observatories with even larger detector counts and fields of view.

The work presented in this dissertation helps enable this exciting vision of the near future across every aspect of these projects. The optical design work of Ch. 3 sets the first group of Prime-Cam instrument modules up for success by ensuring good image and beam quality for the broadband modules and tailored optical performance for the FPI in EoR-Spec. It also helps our community better understand the use and trade-offs associated with biconic lenses in our current generation of receivers to inform choices about how to maximize the performance of future CCAT and ASO modules.

The record numbers of CCAT MKIDs presented in Ch. 2 would not be able to operate at their full potential without the data acquisition plans for Mod-Cam and Prime-Cam presented in Ch. 4. While the science goals, optics, and detector technologies are all crucial to the success of the observatory, the hardware and software to operate the instrument remotely with little downtime over a long period are essential and often overlooked. Much work remains to be done to implement the outstanding pieces of this puzzle for Prime-Cam, but much of the core structure has been built.

Finally, the search for time-variable parity violating photon-DM interactions in ACT polarized depth-1 maps presented in Ch. 5 informs our desire to limit polarization systematics in SO and CCAT. The new methods of this analysis set the stage for further innovative searches for physics beyond the Λ CDM model with upcoming observatories. Microwave time-domain astrophysics across many scales, wavelengths, and phenomena is poised to grow dramatically in this decade. Similarly, interest in cosmic birefringence and parity violating physics like magnetic fields in galaxies and dust characteristics on small scales drives demand for the high-precision polarization measurements that SO and CCAT will be able to

provide. With techniques similar to those presented here, our next generation of microwave observatories will contribute significantly to both of these areas while also enabling searches for dark matter and new physics beyond the standard cosmological and particle physics models.

By developing new millimeter-wave instrumentation, we improve our understanding of the polarized microwave sky. In turn, this understanding may yet teach us more about the most profound realities of the natural world. Like so much of science, better measurements lead to clearer ideas and new projects to explore. Along the way, they also provide those of us who make them with the joy of discovering even a hint of the grandeur and beauty that the universe is waiting to show us.

APPENDIX A

ACRONYM LIST

This appendix lists all (at least in theory) acronyms used in this dissertation. Since some acronyms appear in multiple chapters, the list is arranged alphabetically instead of order or location of appearance.

- **AC** - alternating current (but used occasionally by me to indicate any oscillating effect, as in the AC rotation effect in the ALP chapter)
- **ACT** - Atacama Cosmology Telescope
- **ACU** - antenna control unit, the control system for the SO LAT and FYST
- **ADC** - analog-to-digital converter
- **AGN** - active galactic nuclei
- **ALP** - axion-like particle
- **ASO** - Advanced Simons Observatory
- **BBN** - Big Bang Nucleosynthesis
- **BFTC or Bluefors TC** - Bluefors Temperature Controller
- **BICEP** - the Background Imaging of Cosmic Extragalactic Polarization suite of telescopes; located at the South Pole along with the Keck Array
- **CCAT** - (backronym) Cerro Chajnantor Atacama Telescope
- **CDM** - cold dark matter
- **CHAI** - the CCAT Heterodyne Array Instrument; along with Prime-Cam, the other main science instrument for CCAT
- **CLI** - command line interface

- **CMB** - cosmic microwave background
- **CP** - in the context of chapter five, a charge-parity symmetry
- **CPU** - central processing unit
- **DAC** - digital-to-analog converter
- **DAQ** - data acquisition
- **DC** - direct current
- **DR** - dilution refrigerator
- **DR6** - Atacama Cosmology Telescope Data Release 6
- **EoR** - epoch of reionization
- **EoR-Spec** - the Epoch of Reionization Spectrometer, a module for Prime-Cam on FYST
- **ES** - electronics space; one of the DAQ electronics rack areas on FYST
- **F/#** - the f-number or focal ratio of an optical system
- **FDM** - frequency-domain multiplexing in the context of chapters two and four; in the context of chapter five, fuzzy dark matter
- **FPGA** - field programmable gate array
- **FWHM** - full-width half-maximum
- **FYST** - The Fred Young Submillimeter Telescope, the name of the telescope for the CCAT Observatory
- **GHS** - gas handling system, a component of the Bluefors DR setup
- **GPS** - global positioning system
- **GUI** - graphic user interface
- **HK** - housekeeping data, the data from all non-detector monitoring systems

- **IP** - internet protocol; a set of rules for how packets of data communicate across networks
- **IS1** - instrument space 1; where Prime-Cam lives on FYST
- **LAT** - the SO Large Aperture Telescope
- **LATR** - the SO Large Aperture Telescope Receiver
- **LC** - a type of fiber connector (likely Lucent Connector after the company that developed it)
- **LF** - low frequency observing band for SO, comprising the 30 and 40 GHz windows
- **LSS** - large-scale structure
- **MF** - mid frequency observing band for SO, comprising the 90 and 150 GHz windows
- **MKID** - microwave kinetic inductance detector
- **μ mux** - microwave multiplexing, a type of FDM
- **ML** - maximum likelihood, referring to a type of mapmaking commonly used in our field
- **MTP** - multi-fiber termination push-on, a type of connector for optical fiber cables
- **NFS** - network file system
- **NIST** - the National Institute of Standards and Technology; the SO and CCAT detector arrays are made and tested at their laboratory in Boulder, CO
- **NTP** - network time protocol

- **NUC** - next unit of computing; name for a series of small form-factor computers
- **OCS** - the Observatory Control System developed for FYST by the CCAT software team
- **OMT** - orthomode transducer, a common antenna technology for mm wavelength TESs
- **OpsDB** - CCAT operations database
- **PCS** - the Prime-Cam Control System, the specific SO-OCS agents developed for Prime-Cam
- **PSF** - point spread function
- **PT** - pulse tube
- **PTC** - pulse tube compressor; some people may use PTC to refer to pulse tube cryocooler, but here we will distinguish the pulse tube itself (PT) from its compressor
- **PTE** - probability to exceed
- **PTP** - precision time protocol
- **QCD** - quantum chromodynamics, the theory of the strong nuclear force in the Standard Model of particle physics
- **RAID** - Redundant Array of Independent Disks
- **RF** - radio frequency
- **RFSoc** - Radio Frequency System on a Chip
- **RJ45** - registered jack 45; standard Ethernet cable connector type
- **ROC** - remote observing (or observation) coordinator
- **SAT** - the SO Small Aperture Telescopes

- **SFP** - small form-factor pluggable interface; on our devices, refers to the 1 Gbit fiber connections
- **SFP+** - small form-factor pluggable interface plus; on our devices, refers to the 10 Gbit fiber connections
- **SMuRF** - SLAC (Stanford Linear Accelerator Center) Microresonator Radio Frequency electronics
- **SNMP** - simple network management protocol, a common network communication protocol
- **SO** - Simons Observatory
- **SOCS** - the Simons Observatory Control System package containing agents for SO-OCS
- **SO-OCS** - the Observatory Control System from Simons Observatory (so designated to avoid confusing it with CCAT's OCS)
- **SPT-3G** - South Pole Telescope, Third Generation
- **SQUIDS** - superconducting quantum interference devices
- **SSD** - solid state drive
- **TCP** - transmission control protocol, a common internet communication protocol
- **TCS** - the Telescope Control System, a subsystem of the CCAT OCS for interacting with the ACU
- **TDM** - time-domain multiplexing
- **TES** - transition edge sensor
- **TiN** - titanium nitride

- **TOD** - time-ordered data; often written as TODs to refer collectively to the different sets of time-ordered data for each observation with a telescope
- **UDP** - user datagram protocol, a common method for sending information in packets over an internet protocol network
- **UFM** - universal focal plane module, the SO detector package
- **UHF** - ultrahigh frequency observing band for SO, comprising the 220 and 280 GHz windows
- **UMM** - universal multiplexing module, the readout-only component of an SO UFM
- **USB** - universal serial bus
- **VNA** - vector network analyzer
- **WIMP** - weakly interacting massive particle, a type of dark matter model

APPENDIX B

ZEMAX TIPS AND TRICKS

This appendix seeks to collect information about the details of the optimization and analysis procedures used by the author and collaborators for optimizing the designs found in Ch. 3. Zemax OpticStudio has many settings and tools, and it can be difficult to identify best practices for using them. The material presented here, however, should be taken as just that – a set of best practices, some of which are highly specific to Prime-Cam optics designs. Some of these practices have caveats about their usefulness, but it is my hope that the information collected here might aid future CCAT collaborators in understanding how previous designs were optimized and in producing future designs in Zemax. It is entirely possible that there are better methods than those presented here; the Zemax Help Center is a useful place to learn about recommended settings and procedures when trying to decide how to run the best optimization or analyses. I also make no pretension to a comprehensive overview of the optimization process; it is better to think of this appendix as a collection of important information to consider rather than a how-to guide for optimizing CCAT optics designs.

The initial Prime-Cam designs were made in older versions of Zemax, but the final ones presented in this thesis were all made in Ansys Zemax OpticStudio 2023 R1.00.

A note for opening Prime-Cam Zemax files: the Prime-Cam files contain custom aperture files to describe the shape of the mirrors and the window. When you open them for the first time (and occasionally even after that), it may say that it cannot find the custom aperture files and will ask you to save new ones. This is okay – it will just save separate UDA files for the information contained in the Lens

Data editor. Just click through the boxes and save the files in the recommended locations.

Likewise, it could complain about missing a glass index (MM-WAVE), which can be downloaded from the CCAT wiki (site:opticsdesign). If you are loading the ZAR archive files from the CCAT wiki, they should contain all of these extra files automatically, so hopefully you won't need to do these things. Whenever possible, use the ZAR feature for sharing files instead of sending raw ZOS or ZMX files. Older versions of Zemax may not be able to open ZAR files.

B.1 Optical Parameters to Check

Here is a list of the various optical parameters that should be checked during the optimization process for Prime-Cam modules:

1. The Strehl ratio at the relevant band center wavelengths should generally be above 0.8.
2. The chief ray angles at the image plane should generally be less than 2 deg.
3. Footprints at M1 and stop – make sure the Lyot stop is both properly filled and is the limiting aperture for the system. The footprint for all fields should be a circle of the size of the Lyot stop (the diameter of which may also need to be optimized to control the illumination on the primary) without any fields showing chords that indicate vignetting on other surfaces. The primary mirror should be illuminated appropriately, which is a minor axis diameter of $\sim 5.5 - 5.8$ m for Prime-Cam.
4. The focal plane size from the footprint at the image should match the diameter you plan on filling with detectors.

5. The overall length from the front of the window to the focal surface should be 1180 mm or shorter to fit in Prime-Cam. This is typically the first entry in our Merit Function editor so we can easily check the length.
6. The lens physical parameters should match the physical size of your silicon and interface correctly with mechanical design. Some of the important mechanical parameters are shown in the machining drawing for lens two of EoR-Spec in Fig. B.1.
7. When preparing a design for machining, check the lens thickness solves to make sure they are being calculated for the appropriate lens radius and desired thickness at the edge. To check these, click on the “E” next to the thickness for the front side of each lens. Within the Edge Thickness solve window that comes up, I usually use the full mechanical semi-diameter as the radial height (though the SO ID12 design sometimes used a radius a little less than this). For calculating the target thickness, we first made the drawings for the lenses that included a mounting flange chosen so that a 78.4 mm diameter blade running tangent to the surface of the Zemax-designed lens (i.e prior to adding 0.32 mm on each side of the silicon for the AR coating, which we do in the drawings but not in Zemax) would have 200 microns of clearance above the flange for the entire length of the flange. Fig. B.2 shows the results of this calculation in Solidworks for lens two of EoR-Spec. The flange itself sticks out 8 mm from the edge of the lens. I then adjusted the thickness in the edge thickness solve to match the thickness of the lens required by the drawings to accommodate this central mounting flange. This entire process should be done *after* optimization; it does not affect the optical performance of the lenses to change the central cylindrical thickness as long as you do not alter the curved surfaces. You may want to think about

it during the optimization process, however, because you need to ensure you do not produce a design with lenses too thick for the available silicon blanks.

8. The $F/\#$ calculated at image plane (and at stop for EoR-Spec as a measure of collimation) should be roughly similar to the other Prime-Cam modules.
9. You may wish to use the Huygens PSF method to extract the FWHM and ellipticity of beams. We are looking to limit the median ellipticity and the number of fields that are outliers in ellipticity and FWHM.
10. The sags should be below 14 mm on all lens surfaces if required for AR coating the lenses. The initial SO designs did not require this, but our Chicago collaborators requested it for future designs because it is much easier for them. This requirement led to us making both sides of the lenses curved instead of having all the sag on one side.
11. Check the optical clearances on the lenses – make sure your footprint is sufficiently far in to prevent beam being clipped by poor silicon near edges and/or clamp; generally 10 mm clearance on the radius is sufficient.

For analyzing these things, I either looked at these metrics directly in Zemax (mostly Lens Data, the Multi-Configuration Editor, and the Merit Function Editor) or used the macros and Python scripts written by Patricio Gallardo and available at https://github.com/Patriciogallardo/zemax_tools to save the outputs of various Zemax analyses. Some of the scripts were locally modified or developed further by me (especially for calculating the $F/\#$ at the Lyot stop and for Zemax PSF analyses); these are available locally to our group members on lnx201 under my user folder, but they are not on GitHub. With the newer versions of Zemax that use the Python API, Patricio has also developed many tools and scripts that are available to do this work in a much more efficient way than




Figure B.1: The machining drawing for lens two of the EoR-Spec optical design. Some of the relevant mechanical constraints shown on the drawing are the optically-active diameter (336 mm, smaller than the total diameter of 352 mm which includes the clamping region), the central thickness of the lens including the sag of both sides and the thickness of the clamp (section A-A), and the shape of the clamp (detail B). Figure credit: Rodrigo Freundt.

the old way of using Zemax macros to save text files that are then analyzed separately by Python scripts. I never used these extensively, but I encourage any future collaborators to talk with Patricio and explore them in greater detail.

B.2 Zemax Settings

In this section, I want to focus on a few settings in the System Explorer that are relevant for understanding our designs. The SO designs had a wide range of wavelengths in the wavelengths settings. I would recommend shrinking these




Figure B.2: The Solidworks CAD drawing of the flange for lens two of the EoR-Spec optical design. The flange width of 4.5 mm (D4) is fixed by ensuring that the black dashed-dotted line, the tangent line that the metamaterial antireflection cutting blade makes to the lens surface, has 0.2 mm clearance on each side of the flange (D6 and D9). The 8 mm length of the flange (D3) is fixed by the clamp size in the mechanical design. Figure credit: Rodrigo Freundt.

to only the relevant wavelengths for the design, in part so that the optimization focuses on them, but also so that you do not confuse yourself and use the wrong wavelength on accident. Several Zemax analysis tools default to using wavelength one, so set that to the wavelength you care most about.

The Field Data Editor sets the fields that are used in many analysis tools. You should think of these as the locations being sampled in the focal plane to assess optical performance. The SO ID12 design only uses twelve fields that are spread out in an unpredictable pattern to sample different aspects of the focal plane; twelve is the maximum number that can be viewed simultaneously on a footprint diagram. For most of my designs, I changed this to a uniform diamond of twenty-five fields to ensure that I was covering the focal plane fully during optimization and analysis. Fig. B.3 shows those fields for all the Prime-Cam designs. Note that the maximum x-angle is larger than the maximum y-angle – this is a result of the orientation of the hexagonal window, and it may need to be adjusted if the window is rotated. The LATR window and the Prime-Cam windows do not have the same orientation, so the extreme fields must differ for CCAT and SO designs.

Another option for the fields often used by some of our collaborators is a radial layout that samples the center of the focal plane more densely than the outer annulus.

Setting the aperture correctly is important for PSF analyses and for ensuring that you are setting the illumination of the primary mirror appropriately during optimization and analysis. One option is to set the aperture type to “Entrance Pupil Diameter” and the aperture value to your target value (usually around 5500 mm, as in ID12). Over time, I began to use a different method. For optimization, the default method that Patricio and I used is to clear vignetting and set the aperture to Float by Stop Size for each optimization so that the optimization ensures that the Lyot stop is fully filled. Using this approach ensures that you do not need to worry about distortions to the Lyot stop or constraining the Lyot stop in the merit function. It does mean, however, that you need to reset the aperture




Figure B.3: The twenty-five fields used for optimizing and analyzing CCAT designs

to get a good look at what your design is doing under the conditions you want, thought you can still see the Strehl well enough to know whether the optimization is working.

After each optimization, I would usually check the footprint at the primary mirror. If the illumination looks bad, that might be an indication that you need to tweak your Lyot stop size and reoptimize, though you should be careful not to shrink the Lyot stop needlessly because it reduces your mapping speed. If the illumination was very problematic, it often meant that a different optimization approach was required.

Generally, the illumination was fairly good. In that case, I would set the

aperture size to an oversized value like 7000 mm and then would set the vignetting factors for each configuration before doing the analysis. As long as the Lyot stop is the only aperture-defining surface for the design, the vignetting factors in the Field Data Editor ensure that you are only tracing the rays that pass through the Lyot stop, which should make your illumination and analysis tools the most accurate reflection of what our real system would be like. Using an undersized aperture (e.g. 5500 mm) may be sufficient and is certainly easier for many applications.

To set the illumination for each configuration separately, I used a Python API script that Patricio developed¹ to set them in the Multi-Configuration Editor. For each field, it added a FVCX, FVCY, FVDX, FVDY, and FTAN field, then cycled through each configuration and set the vignetting factors. It did not seem that clicking the button in the Field Data Editor did this separately for each configuration, which is why we used this more complex method.

Finally, a word about real ray aiming: in the System Explorer under Ray Aiming, there is an option to turn on Real Ray Aiming. Generally, the Zemax documentation argues that this setting is not necessary. This is likely true for the broadband designs. For the EoR-Spec design, however, I did find that using real ray aiming seemed to give more accurate calculations of the $F/\#$ around the stop, enabling the optimization process to find a better design with aspheric lenses than my early biconic lens designs.

The major caveat with real ray aiming is that it seems to make the designs

¹https://github.com/Patriciogallardo/zemax_tools/blob/master/ZOS_API_scripts/S4camScripts/time_reverse/TMP/set1_lyot_stop_scripts/10_set_vignetting.py - to run, start by opening the Interactive Extension window in Zemax under the Programming tab. Copy the script linked above into the \sim /Documents/Zemax/ZOS-API Projects/PythonZOSConnection (or the equivalent location of the Zemax folder) on your computer, then right-click on it, click Edit with IDLE, edit the number of configs and row start as needed, and run it by going to Run→Run Module.

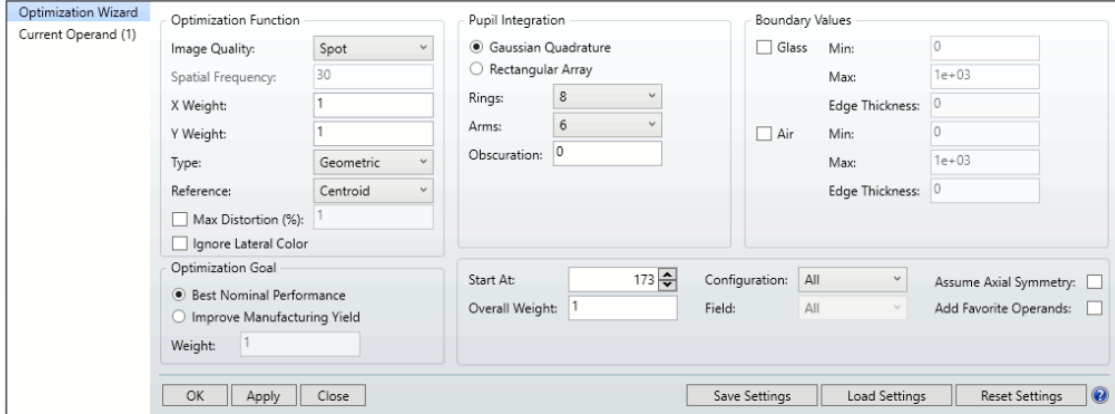


Figure B.4: Settings that I commonly used for the default sequential merit function.

somewhat fragile. For example, if I optimize the design with real ray aiming and then disable it and try to tweak the values of the lens parameters, I occasionally saw nonsensical Strehl plots and other indications that analysis tools were not always working as expected. This affects the “as-machined” (v10) design for EoR-Spec, which has some jagged features in what had previously been smooth Strehl plots. I do not fully understand what real ray aiming is doing here, but I believe that this is a numerical artifact and not an issue with the design itself. For future designs, however, it is not clear that real ray aiming helps considerably and it may make your analyses look a bit strange.

B.3 Merit Function and Optimization Settings Considerations

For the base merit function, I most commonly used PTV spot (or Geometric spot) as the main optical quality metric. RMS wavefront also seemed to yield good results at times. For the Gaussian Quadrature, I usually used 8 rings and

6 arms because the higher number of rays seemed to yield better results early on in developing the EoR-Spec design. Later, I found information in the Zemax documentation about how to set these by starting at a low value and incrementing them until the merit function does not change by more than 1%, so I probably could have used a smaller number of rays (which would have sped things up). Fig. B.4 shows the settings I usually used.

For the rest of the merit function, the full merit functions can be found in the relevant ZAR files. The original SO files had constraints to maintain the effective focal plane size, limit the variation of the $F/\#$ across the focal plane, prevent vignetting, keep rays perpendicular to the 40 K filters, and maintain a good Lyot stop and flat focal plane. I left all of these in place, though I did play around at times with their relative weights. If a design consistently tended towards failing to maintain any of these, I often altered the lens parameters to a fairly generic lens that did not fail any of these, then increased the weight on the constraints that were failing and tried again.

For all designs, I added sag constraints on all surfaces as in Fig. B.5. Some of these never had issues or were not being optimized and could be disabled by setting their weight to zero. Keeping these constraints under 14 mm was often challenging in the optimization when the designs were having difficulty finding a good solution for the other constraints.

For EoR-Spec, I also added constraints on the $F/\#$ at the Lyot stop that used a custom Zemax macro (ZPL00.ZPL; can be found in my lnx201 folder or the ZAR files) with the ZPLM constraint type that calculated the $F/\#$ using the extreme rays as presented in section 3.1. I found that if I set all of these to the target of 100, it often failed to reach any of them, but the results were still good if I set the

	Type	Cfg#											
1	CONF ▾	1											
2	BLNK ▾	length of optics											
3	TTHI ▾	26	67						0.000	0.000	1175.448	0.000	
4	OPLT ▾	3							1178.0...	20.000	1178.000	0.000	
5	BLNK ▾	window to first lens (not inc ~17mm window sag)											
6	TTHI ▾	26	44						0.000	0.000	213.873	0.000	
7	BLNK ▾	sag of L3											
8	SSAG ▾	58	0	168.0...	0.000	0	0	0	0.000	0.000	11.338	0.000	
9	CONS ▾								13.800	0.000	13.800	0.000	
10	MAXX ▾	8	9						13.800	10.000	13.800	0.000	
11	BLNK ▾	sag of L3 back											
12	SSAG ▾	59	1	168.0...	0.000	0	0	0	0.000	0.000	-7.057	0.000	
13	OPGT ▾	12							-13.800	2.000	-13.800	0.000	
14	BLNK ▾	sag of L1 back											
15	SSAG ▾	45	1	189.0...	0.000	0	0	0	0.000	0.000	-12.078	0.000	
16	CONS ▾								-14.000	0.000	-14.000	0.000	
17	MINN ▾	15	16						-14.000	0.000	-14.000	0.000	
18	BLNK ▾	sag of L2 back											
19	SSAG ▾	50	1	168.0...	0.000	0	0	0	0.000	0.000	-9.909	0.000	
20	CONS ▾								-13.800	0.000	-13.800	0.000	
21	MINN ▾	19	20						-13.800	0.000	-13.800	0.000	
22	BLNK ▾	sag of L2 front											
23	SSAG ▾	49	1	168.0...	0.000	0	0	0	0.000	0.000	5.500	0.000	
24	OPLT ▾	23							13.800	0.000	13.800	0.000	
25	BLNK ▾	sag of L4											
26	SSAG ▾	63	1	155.0...	0.000	0	0	0	0.000	0.000	7.111	0.000	
27	CONS ▾								14.000	0.000	14.000	0.000	
28	MAXX ▾	26	27						14.000	2.000	14.000	0.000	

Figure B.5: Sag constraints in the merit function

central target to 100 and the outer targets to something slightly lower (e.g. 70).

For all designs, I optimized all possible locations in Prime-Cam simultaneously and usually used configuration one as the default configuration in the merit function. If a module is going to go into one location for certain, something that could be explored is setting the configuration in the merit function to that location. Very limited tests using other configurations in the merit function for the stop F/# optimization with EoR-Spec did not seem to reveal much of a difference, especially if all configurations were used in the default merit function.

While optimizing, I also found that it was best to run the optimization for a

short period of time and check the results. If you let it run for hours, the merit function might tick down a little bit further, but it also has a greater tendency to fall into local minima that do not have desirable properties. The best results were obtained by tweaking the values of the variables being used, running the optimization for a couple of minutes until the merit function made a major drop, checking the results, tweaking which variables are used and the weights of merit function operands to guide the design in the direction you want, and iterating this process several times.

B.4 Tolerancing Procedure

Once a design was finalized, we did a rough tolerancing procedure using the Zemax tolerancing data editor to ensure that the design was not overly sensitive to the precise mechanical locations and orientations of the lenses and focal plane.

For this procedure, I largely followed steps suggested by Patricio:

1. Save the existing merit function as a Zemax MF file so that you can recover it later.
2. Make a new merit function that only uses the STRH constraint to sample the Strehl at various locations in the relevant configuration(s). For EoR-Spec_rra_asph_v9.zos, I used a regular set of field points that cover about half the radius of the focal plane. We sample out to this radius because most of the detectors are in that region. I sampled the STRH constraints at 128x128 (option 3; this made it take many hours to run, but should be more accurate), used the lowest wavelength (since that gives the tightest

tolerance), and set the target and weight to one so that the merit function will decrease for better Strehl. We generally looked for the change in the tolerancing parameters that would cause a 2% degradation in the collective merit function of these Strehl constraints.

3. Open the tolerancing data editor under the Tolerancing tab in the Zemax toolbar. Set the test wavelength to the shortest relevant central wavelength, which will have the tightest tolerances. Zemax puts in the radii of lenses, the z location of various components (mostly the lenses), tilts and decenters for those components, and indices of refraction by default, but we generally only care about the locations, tilts, and decenters for all lenses. The others can be deleted or ignored (by checking the “ignore this operand during tolerancing” box in the operand properties). I usually just leave the default starting range and the default way of calculating the distances for each lens by grouping together a range of surfaces. This does not separately call out the Lyot stop location, instead lumping it in with L3. Fig. B.6 shows the operands used for EoR-Spec (along with the final results after running the next step).
4. Run the tolerancing function in Zemax. I had to delete the (empty) parameters for surfaces 43, 48, and 57 (the coordinate breaks before lenses) from the multi-configuration editor before it would let me do the tolerancing for EoR-Spec. Within the tolerancing dialog box, I used inverse limit tolerancing with the Merit Function as the criterion and a limit of 0.176, which was 0.02 greater than the original merit function. I did not use any compensators or Monte Carlo runs. Running this feature then shows the results in the tolerancing data editor.

We stopped with the results at this point since they were generally larger than 1 mm and could easily be met. The exception was the distance between lens four

and the image plane, but these are held at the same temperature and the gap between them is easiest to maintain in our mechanical designs.

Patricio would often do an additional step in his tolerancing. Starting from the output of this first step, he would shrink the tolerances by a factor of $\sqrt{\text{dof}}$ where dof is the number of parameters varied during tolerancing (the “degrees of freedom”). He would then run a series of Monte Carlo simulations to vary the parameters within that tighter tolerance and make sure that the large majority of simulations stay within the overall Strehl tolerance. The reason for doing this is that the first step of tolerancing varies each element one by one through the desired range to see the limits at which the target merit function degradation is met. In a real system, you will have offsets from the design in every element, not just one of them. The Monte Carlo simulations account for this by drawing an offset for each variable at random many times to build up a distribution of expected Strehl degradation in a real system. We decided it was not necessary to do this for our rough EoR-Spec tolerancing.

	Type	Surf		Nominal	Min	Max	Comment
1	TWAV ▾				860.000		Default test wavelength.
2	TFRN ▾	31		0.000	-1.000	1.000	Default radius tolerances.
3	TFRN ▾	32		0.000	-1.000	1.000	
4	TFRN ▾	34		0.000	-1.000	1.000	
5	TFRN ▾	35		0.000	-1.000	1.000	
6	TFRN ▾	38		0.000	-1.000	1.000	
7	TFRN ▾	39		0.000	-1.000	1.000	
8	TFRN ▾	41		0.000	-1.000	1.000	
9	TFRN ▾	42		0.000	-1.000	1.000	
10	TTHI ▾	31	31	1.000	-2.292	20.000	z location of tube
11	TTHI ▾	43	48	0.000	-7.018	20.000	lens1
12	TTHI ▾	48	53	0.000	-1.327	10.000	lens2
13	TTHI ▾	55	62	0.000	-4.178	20.000	lens3
14	TTHI ▾	62	67	0.000	-7.141	0.564	lens4
15	TTHI ▾	67	67	0.000	-0.703	9.166	array
16	TPAR ▾	43	1	0.000	-5.000	5.000	Default element dec/tilt tolerances 44...
17	TPAR ▾	43	2	0.000	-5.000	5.000	
18	TPAR ▾	43	3	0.000	-2.000	2.000	
19	TPAR ▾	43	4	0.000	-2.000	2.000	
20	TPAR ▾	48	1	0.000	-5.000	5.000	Default element dec/tilt tolerances 44...
21	TPAR ▾	48	2	0.000	-5.000	5.000	
22	TPAR ▾	48	3	0.000	-2.000	2.000	
23	TPAR ▾	48	4	0.000	-2.000	2.000	
24	TPAR ▾	57	1	0.000	-5.000	5.000	Default element dec/tilt tolerances 44...
25	TPAR ▾	57	2	0.000	-5.000	5.000	
26	TPAR ▾	57	3	0.000	-2.000	2.000	
27	TPAR ▾	57	4	0.000	-2.000	2.000	
28	TPAR ▾	62	1	0.000	-5.000	5.000	Default element dec/tilt tolerances 44...
29	TPAR ▾	62	2	0.000	-5.000	5.000	
30	TPAR ▾	62	3	0.000	-2.000	2.000	
31	TPAR ▾	62	4	0.000	-2.000	2.000	
32	TIND ▾	31		1.560	-1.000E-03	1.000E-03	Default index tolerances.

Figure B.6: The tolerance data editor for the EoR-Spec tolerancing. The parameters in gray were not used in tolerancing, so the distances between lenses and relative tilts and decenters were the only things toleranced. Units are in mm. The tightest tolerance is the distance between L4 and the array (0.564 mm, moving L4 closer to the array), but this is also the most easily constrained parameter by our carbon fiber strut, so we felt comfortable proceeding.

BIBLIOGRAPHY

- [1] N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo, and et al. Planck 2018 results - VI. Cosmological parameters. *Astronomy & Astrophysics*, 641:A6, September 2020. ISSN 1432-0746. doi: 10.1051/0004-6361/201833910. URL <http://dx.doi.org/10.1051/0004-6361/201833910>.
- [2] J. C. Mather, E. S. Cheng, R. E. Eplee, Jr., R. B. Isaacman, S. S. Meyer, R. A. Shafer, R. Weiss, E. L. Wright, et al. A Preliminary Measurement of the Cosmic Microwave Background Spectrum by the Cosmic Background Explorer (COBE) Satellite. *Astrophysical Journal Letters*, 354:L37, May 1990. doi: 10.1086/185717. URL <https://doi.org/10.1086/185717>.
- [3] Planck Collaboration. Planck 2018 results - I. Overview and the cosmological legacy of Planck. *Astronomy & Astrophysics*, 641:A1, September 2020. doi: 10.1051/0004-6361/201833880. URL <http://www.aanda.org/articles/aa/abs/2020/09/aa33880-18/aa33880-18.html>.
- [4] S. K. Choi, M. Hasselfield, S. P. Ho, B. Koopman, M. Lungu, et al. The Atacama Cosmology Telescope: a measurement of the Cosmic Microwave Background power spectra at 98 and 150 GHz. *Journal of Cosmology and Astroparticle Physics*, 2020(12):045–045, Dec 2020. ISSN 1475-7516. doi: 10.1088/1475-7516/2020/12/045. URL <http://dx.doi.org/10.1088/1475-7516/2020/12/045>.
- [5] S. R. Dicker, P. A. Gallardo, J. E. Gudmundsson, P. D. Mauskopf, et al. Cold optical design for the Large Aperture Simons Observatory telescope. In *Ground-based and Airborne Telescopes VII*, volume 10700 of *Proc. SPIE*,

- pages 1064–1076, 2018. doi: 10.1117/12.2313444. URL <https://doi.org/10.1117/12.2313444>.
- [6] E. M. Vavagiakis, C. J. Duell, et al. CCAT-prime: design of the Mod-Cam receiver and 280 GHz MKID instrument module. In Jonas Zmuidzinas and Jian-Rong Gao, editors, *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XI*, volume 12190, page 1219004. International Society for Optics and Photonics, SPIE, 2022. doi: 10.1117/12.2630115. URL <https://doi.org/10.1117/12.2630115>.
- [7] S. K. Choi, C. J. Duell, et al. CCAT-prime: Characterization of the First 280 GHz MKID Array for Prime-Cam. *Journal of Low Temperature Physics*, 209:849–856, Aug 2022. doi: 10.1007/s10909-022-02787-9. URL <https://doi.org/10.1007/s10909-022-02787-9>.
- [8] C. J. Duell et al. CCAT: nonlinear effects in 280 GHz aluminum kinetic inductance detectors. In *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XII*, volume 13102, pages 638–647. SPIE, August 2024. doi: 10.1117/12.3020507. URL <https://doi.org/10.1117/12.3020507>.
- [9] Z. B. Huber, Y. Li, E. M. Vavagiakis, et al. The Simons Observatory: Magnetic Shielding Measurements for the Universal Multiplexing Module. *Journal of Low Temperature Physics*, 209:667–676, Nov 2022. ISSN 1573-7357. doi: 10.1007/s10909-022-02875-w. URL <https://doi.org/10.1007/s10909-022-02875-w>.
- [10] B. J. Koopman, J. Lashner, L. J. Saunders, M. Hasselfield, et al. The Simons Observatory: Overview of data acquisition, control, monitoring,

- and computer infrastructure. In *Software and Cyberinfrastructure for Astronomy VI*, page 6, December 2020. doi: 10.1117/12.2561771. URL <https://doi.org/10.1117/12.2561771>.
- [11] S. K. Choi et al. Sensitivity of the Prime-Cam Instrument on the CCAT-Prime Telescope. *Journal of Low Temperature Physics*, 199(3-4):1089–1097, Mar 2020. ISSN 1573-7357. doi: 10.1007/s10909-020-02428-z. URL <http://dx.doi.org/10.1007/s10909-020-02428-z>.
- [12] Z. B. Huber, L. T. Lin, E. M. Vavagiakis, R. G. Freundt, et al. CCAT: Prime-Cam optics overview and status update. In *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XII*, volume 13102, pages 784–793. SPIE, August 2024. doi: 10.1117/12.3020373. URL <https://doi.org/10.1117/12.3020373>.
- [13] M. A. Fedderke, P. W. Graham, and S. Rajendran. Axion dark matter detection with CMB polarization. *Phys. Rev. D*, 100(1):015040, July 2019. doi: 10.1103/PhysRevD.100.015040. URL <https://link.aps.org/doi/10.1103/PhysRevD.100.015040>. Publisher: American Physical Society.
- [14] DES Collaboration, E. O. Nadler, A. Drlica-Wagner, K. Bechtol, S. Mau, R. H. Wechsler, V. Gluscevic, K. Boddy, A. B. Pace, T. S. Li, et al. Constraints on Dark Matter Properties from Observations of Milky Way Satellite Galaxies. *Phys. Rev. Lett.*, 126(9):091101, March 2021. doi: 10.1103/PhysRevLett.126.091101. URL <https://link.aps.org/doi/10.1103/PhysRevLett.126.091101>. Publisher: American Physical Society.
- [15] V. Iršič, M. Viel, M. G. Haehnelt, J. S. Bolton, and G. D. Becker. First Constraints on Fuzzy Dark Matter from Lyman- α Forest Data and Hydrodynamical Simulations. *Phys. Rev. Lett.*, 119(3):031302, July 2017. doi:

- 10.1103/PhysRevLett.119.031302. URL <https://link.aps.org/doi/10.1103/PhysRevLett.119.031302>. Publisher: American Physical Society.
- [16] CAST Collaboration. New CAST limit on the axion-photon interaction. *Nature Phys*, 13(6):584–590, June 2017. ISSN 1745-2481. doi: 10.1038/nphys4109. URL <https://www.nature.com/articles/nphys4109>. Publisher: Nature Publishing Group.
- [17] P. J. E. Peebles. *Cosmology’s Century*. Princeton University Press, 2020. ISBN 9780691196022.
- [18] S. Staggs, J. Dunkley, and L. Page. Recent discoveries from the cosmic microwave background: a review of recent progress. *Reports on Progress in Physics*, 81, Feb 2018. doi: 10.1088/1361-6633/aa94d5. URL <https://doi.org/10.1088/1361-6633/aa94d5>.
- [19] D. Samtleben, S. Staggs, and B. Winstein. The Cosmic Microwave Background for Pedestrians: A Review for Particle and Nuclear Physicists. *Annual Review of Nuclear and Particle Science*, 57(1):245–283, Nov 2007. ISSN 1545-4134. doi: 10.1146/annurev.nucl.54.070103.181232. URL <http://dx.doi.org/10.1146/annurev.nucl.54.070103.181232>.
- [20] S. Dodelson and F. Schmidt. *Modern Cosmology*. Elsevier Science, 2020. ISBN 9780128159484.
- [21] D. Scott and G.F. Smoot. The Review of Particle Physics: Cosmic Microwave Background. *Prog. Theor. Exp. Phys.*, 2024. URL <https://pdg.lbl.gov/2024/reviews/rpp2024-rev-cosmic-microwave-background.pdf>.
- [22] O. Lahav and A.R. Liddle. The Review of Particle Physics: Cosmological

- Parameters. *Prog. Theor. Exp. Phys.*, 2024. URL <https://pdg.lbl.gov/2024/reviews/rpp2024-rev-cosmological-parameters.pdf>.
- [23] J. D. Jackson. *Classical Electrodynamics*. John Wiley and Sons, 3rd edition, 1999. ISBN 9780471309321.
- [24] S. Weinberg. *Cosmology*. Oxford University Press, 2008. ISBN 9780198526827.
- [25] M. Zaldarriaga. Nature of the $E - B$ decomposition of CMB polarization. *Phys. Rev. D*, 64(10):103001, October 2001. doi: 10.1103/PhysRevD.64.103001. URL <https://link.aps.org/doi/10.1103/PhysRevD.64.103001>. Publisher: American Physical Society.
- [26] K.A. Olive and J.A. Peacock. The Review of Particle Physics: Big-Bang Cosmology. *Prog. Theor. Exp. Phys.*, 2024. URL <https://pdg.lbl.gov/2024/reviews/rpp2024-rev-bbang-cosmology.pdf>.
- [27] S. Aiola, E. Calabrese, L. Maurin, S. Naess, B. L. Schmitt, et al. The Atacama Cosmology Telescope: DR4 maps and cosmological parameters. *Journal of Cosmology and Astroparticle Physics*, 2020(12):047–047, Dec 2020. ISSN 1475-7516. doi: 10.1088/1475-7516/2020/12/047. URL <http://dx.doi.org/10.1088/1475-7516/2020/12/047>.
- [28] The Simons Observatory Collaboration. The Simons Observatory: Science goals and forecasts. *J. Cosmol. Astropart. Phys.*, 2019(02):056–056, February 2019. ISSN 1475-7516. doi: 10.1088/1475-7516/2019/02/056. URL <https://iopscience.iop.org/article/10.1088/1475-7516/2019/02/056>. arXiv:1808.07445 [astro-ph].

- [29] A. Kosowsky, M. Milosavljevic, and R. Jimenez. Efficient cosmological parameter estimation from microwave background anisotropies. *Physical Review D*, 66(6), Sep 2002. ISSN 1089-4918. doi: 10.1103/physrevd.66.063007. URL <http://dx.doi.org/10.1103/PhysRevD.66.063007>.
- [30] D. Dutcher, L. Balkenhol, P. A. R. Ade, Z. Ahmed, E. Anderes, A. J. Anderson, M. Archipley, J. S. Avva, K. Aylor, P. S. Barry, and et al. Measurements of the E-mode polarization and temperature-E-mode correlation of the CMB from SPT-3G 2018 data. *Physical Review D*, 104(2), Jul 2021. ISSN 2470-0029. doi: 10.1103/physrevd.104.022003. URL <http://dx.doi.org/10.1103/PhysRevD.104.022003>.
- [31] A. Lewis and S. Bridle. Cosmological parameters from CMB and other data: A Monte Carlo approach. *Physical Review D*, 66(10), Nov 2002. ISSN 1089-4918. doi: 10.1103/physrevd.66.103511. URL <http://dx.doi.org/10.1103/PhysRevD.66.103511>.
- [32] A. Lewis. CosmoMC README, 2021. URL <https://cosmologist.info/cosmomc/readme.html>.
- [33] J. Ellis and D. Wands. The Review of Particle Physics: Inflation. *Prog. Theor. Exp. Phys.*, 2024. URL <https://pdg.lbl.gov/2024/reviews/rpp2024-rev-inflation.pdf>.
- [34] A. G. Riess, A. V. Filippenko, et al. Observational evidence from supernovae for an accelerating universe and a cosmological constant. *The Astronomical Journal*, 116(3):1009–1038, Sep 1998. ISSN 0004-6256. doi: 10.1086/300499. URL <http://dx.doi.org/10.1086/300499>.
- [35] S. Perlmutter, G. Aldering, G. Goldhaber, R. A. Knop, P. Nugent, The Su-

- pernova Cosmology Project, et al. Measurements of Omega and Lambda from 42 High-Redshift Supernovae. *The Astrophysical Journal*, 517(2): 565–586, Jun 1999. ISSN 1538-4357. doi: 10.1086/307221. URL <http://dx.doi.org/10.1086/307221>.
- [36] D.H. Weinberg and M. White. The Review of Particle Physics: Dark Energy. *Prog. Theor. Exp. Phys.*, 2024. URL <https://pdg.lbl.gov/2024/reviews/rpp2024-rev-dark-energy.pdf>.
- [37] G. Bertone, D. Hooper, and J. Silk. Particle dark matter: evidence, candidates and constraints. *Physics Reports*, 405(5):279–390, 2005. ISSN 0370-1573. doi: <https://doi.org/10.1016/j.physrep.2004.08.031>. URL <https://www.sciencedirect.com/science/article/pii/S0370157304003515>.
- [38] L. Baudis and S. Profumo. The Review of Particle Physics: Dark Matter. *Prog. Theor. Exp. Phys.*, 2024. URL <https://pdg.lbl.gov/2024/reviews/rpp2024-rev-dark-matter.pdf>.
- [39] T. Damour. The Review of Particle Physics: Experimental Tests of Gravitational Theory. *Prog. Theor. Exp. Phys.*, 2024. URL <https://pdg.lbl.gov/2024/reviews/rpp2024-rev-gravity-tests.pdf>.
- [40] J. W. Fowler, M. D. Niemack, S. R. Dicker, et al. Optical design of the Atacama Cosmology Telescope and the Millimeter Bolometric Array Camera. *Appl. Opt., AO*, 46(17):3444–3454, June 2007. ISSN 2155-3165. doi: 10.1364/AO.46.003444. URL <https://opg.optica.org/ao/abstract.cfm?uri=ao-46-17-3444>. Publisher: Optica Publishing Group.
- [41] D. S. Swetz et al. OVERVIEW OF THE ATACAMA COSMOLOGY TELESCOPE: RECEIVER, INSTRUMENTATION, AND TELESCOPE SYS-

- TEMS. *The Astrophysical Journal Supplement Series*, 194(2):41, jun 2011. doi: 10.1088/0067-0049/194/2/41. URL <https://dx.doi.org/10.1088/0067-0049/194/2/41>.
- [42] R. J. Thornton et al. THE ATACAMA COSMOLOGY TELESCOPE: THE POLARIZATION-SENSITIVE ACTPol INSTRUMENT. *ApJS*, 227(2):21, December 2016. ISSN 0067-0049. doi: 10.3847/1538-4365/227/2/21. URL <https://dx.doi.org/10.3847/1538-4365/227/2/21>. Publisher: The American Astronomical Society.
- [43] S. W. Henderson et al. Advanced ACTPol Cryogenic Detector Arrays and Readout. *J Low Temp Phys*, 184(3):772–779, August 2016. ISSN 1573-7357. doi: 10.1007/s10909-016-1575-z. URL <https://doi.org/10.1007/s10909-016-1575-z>.
- [44] T. Louis, G. E. Addison, M. Hasselfield, et al. The Atacama Cosmology Telescope: cross correlation with Planck maps. *J. Cosmol. Astropart. Phys.*, 2014(07):016, July 2014. ISSN 1475-7516. doi: 10.1088/1475-7516/2014/07/016. URL <https://dx.doi.org/10.1088/1475-7516/2014/07/016>.
- [45] F. J. Qu, B. D. Sherwin, M. S. Madhavacheril, D. Han, K. T. Crowley, et al. The Atacama Cosmology Telescope: A Measurement of the DR6 CMB Lensing Power Spectrum and Its Implications for Structure Growth. *The Astrophysical Journal*, 962(2):112, feb 2024. doi: 10.3847/1538-4357/acfe06. URL <https://dx.doi.org/10.3847/1538-4357/acfe06>.
- [46] E. M. Vavagiakis, P. A. Gallardo, V. Calafut, et al. The Atacama Cosmology Telescope: Probing the baryon content of SDSS DR15 galaxies with the thermal and kinematic Sunyaev-Zel’dovich effects. *Phys. Rev. D*, 104(4):043503, August 2021. doi: 10.1103/PhysRevD.104.043503. URL <https://>

link.aps.org/doi/10.1103/PhysRevD.104.043503. Publisher: American Physical Society.

- [47] V. Calafut, P. A. Gallardo, E. M. Vavagiakis, S. Amodeo, S. Aiola, J. E. Austermann, N. Battaglia, E. S. Battistelli, J. A. Beall, R. Bean, and et al. The Atacama Cosmology Telescope: Detection of the pairwise kinematic Sunyaev-Zel'dovich effect with SDSS DR15 galaxies. *Physical Review D*, 104 (4), Aug 2021. ISSN 2470-0029. doi: 10.1103/physrevd.104.043502. URL <http://dx.doi.org/10.1103/PhysRevD.104.043502>.
- [48] Y. Li, E. Biermann, S. Naess, et al. The Atacama Cosmology Telescope: Systematic Transient Search of 3 Day Maps. *The Astrophysical Journal*, 956 (1):36, oct 2023. doi: 10.3847/1538-4357/ace599. URL <https://dx.doi.org/10.3847/1538-4357/ace599>.
- [49] E. Calabrese, J. C. Hill, H. T. Jense, A. La Posta, et al. The Atacama Cosmology Telescope: DR6 Constraints on Extended Cosmological Models, March 2025. URL <http://arxiv.org/abs/2503.14454>. arXiv:2503.14454 [astro-ph].
- [50] S. Naess et al. The Atacama Cosmology Telescope: arcminute-resolution maps of 18 000 square degrees of the microwave sky from ACT 2008-2018 data combined with Planck. *J. Cosmol. Astropart. Phys.*, 2020(12):046, December 2020. ISSN 1475-7516. doi: 10.1088/1475-7516/2020/12/046. URL <https://dx.doi.org/10.1088/1475-7516/2020/12/046>.
- [51] S. Naess, Y. Guan, A. J. Duivenvoorden, M. Hasselfield, Y. Wang, et al. The Atacama Cosmology Telescope: DR6 Maps, March 2025. URL <http://arxiv.org/abs/2503.14451>. arXiv:2503.14451 [astro-ph].

- [52] Z. Xu and The Simons Observatory Collaboration. The Simons Observatory: The Large Aperture Telescope (LAT). *Res. Notes AAS*, 5(4):100, April 2021. ISSN 2515-5172. doi: 10.3847/2515-5172/abf9ab. URL <https://dx.doi.org/10.3847/2515-5172/abf9ab>. Publisher: The American Astronomical Society.
- [53] N. Zhu, T. Bhandarkar, G. Coppi, A. M. Kofman, J. L. Orlowski-Scherer, Z. Xu, et al. The Simons Observatory Large Aperture Telescope Receiver. *ApJS*, 256(1):23, September 2021. ISSN 0067-0049, 1538-4365. doi: 10.3847/1538-4365/ac0db7. URL <https://iopscience.iop.org/article/10.3847/1538-4365/ac0db7>.
- [54] Z. Xu, T. Bhandarkar, G. Coppi, A. M. Kofman, J. L. Orlowski-Scherer, N. Zhu, et al. The Simons Observatory: the Large Aperture Telescope Receiver (LATR) Integration and Validation Results. In *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy X*, page 183, December 2020. doi: 10.1117/12.2576151. URL <https://dx.doi.org/10.1117/12.2576151>.
- [55] A. M. Ali et al. Small Aperture Telescopes for the Simons Observatory. *J Low Temp Phys*, 200(5-6):461–471, September 2020. ISSN 0022-2291, 1573-7357. doi: 10.1007/s10909-020-02430-5. URL <https://link.springer.com/10.1007/s10909-020-02430-5>.
- [56] K. Kiuchi et al. Simons Observatory Small Aperture Telescope overview. In *Ground-based and Airborne Telescopes VIII*, page 162, December 2020. doi: 10.1117/12.2562016. URL <https://dx.doi.org/10.1117/12.2562016>.
- [57] K. T. Crowley et al. The Simons Observatory: calibration and characterization of the first small-aperture telescope. In *Millimeter, Submillimeter,*

- and Far-Infrared Detectors and Instrumentation for Astronomy XII*, volume 13102, pages 502–515. SPIE, August 2024. doi: 10.1117/12.3020564. URL <https://doi.org/10.1117/12.3020564>.
- [58] The Simons Observatory Collaboration. The Simons Observatory: Science Goals and Forecasts for the Enhanced Large Aperture Telescope, March 2025. URL <http://arxiv.org/abs/2503.00636>. arXiv:2503.00636 [astro-ph].
- [59] C. J. Duell, E. M. Vavagiakis, et al. CCAT-prime: Designs and status of the first light 280 GHz MKID array and mod-cam receiver. *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy X*, Dec 2020. doi: 10.1117/12.2562757. URL <http://dx.doi.org/10.1117/12.2562757>.
- [60] S. C. Chapman, A. I. Huber, A. K. Sinclair, J. D. Wheeler, et al. CCAT-prime: the 850 GHz camera for prime-cam on FYST. In Jonas Zmuidzinas and Jian-Rong Gao, editors, *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XI*, volume 12190, page 1219005. International Society for Optics and Photonics, SPIE, 2022. doi: 10.1117/12.2630628. URL <https://doi.org/10.1117/12.2630628>.
- [61] T. Nikola et al. CCAT-prime: the epoch reionization spectrometer for prime-cam on FYST. In Jonas Zmuidzinas and Jian-Rong Gao, editors, *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XI*, volume 12190, page 121900G. International Society for Optics and Photonics, SPIE, 2022. doi: 10.1117/12.2629338. URL <https://doi.org/10.1117/12.2629338>.
- [62] U. U. Graf, N. Honingh, I. Barrueto, K. Jacobs, B. Klein, R. Molina,

- A. Reyes, P. A. Tapia, and J. Stutzki. CHAI, the CCAT-prime Heterodyne Array Instrument. In *Proceedings of the 30th International Symposium on Space Terahertz Technology*, volume 30, page 77, 2019.
- [63] G. J. Stacey et al. CCAT-prime/FYST: a status report on the ultra-widefield submillimeter observatory on Cerro Chajnantor. In *Ground-based and Airborne Telescopes IX*, volume 12182, pages 365–384. SPIE, August 2022. doi: 10.1117/12.2630380. URL <https://dx.doi.org/10.1117/12.2630380>.
- [64] CCAT-Prime Collaboration. CCAT-prime Collaboration: Science Goals and Forecasts with Prime-Cam on the Fred Young Submillimeter Telescope. *The Astrophysical Journal Supplement Series*, 264(1):7, dec 2022. doi: 10.3847/1538-4365/ac9838. URL <https://dx.doi.org/10.3847/1538-4365/ac9838>.
- [65] C. K. Walker. *Terahertz Astronomy*. CRC Press, 2016.
- [66] S. C. Parshley, M. D. Niemack, R. Hills, S. R. Dicker, et al. The optical design of the six-meter CCAT-prime and Simons Observatory telescopes. In Roberto Gilmozzi, Heather K. Marshall, and Jason Spyromilio, editors, *Ground-based and Airborne Telescopes VII*, page 145, Austin, United States, July 2018. SPIE. ISBN 978-1-5106-1953-1 978-1-5106-1954-8. doi: 10.1117/12.2314073. URL <https://dx.doi.org/10.1117/12.2314073>.
- [67] M. D. Niemack. Designs for a large-aperture telescope to map the CMB $10\times$ faster. *Appl. Opt.*, 55(7):1688, March 2016. ISSN 0003-6935, 1539-4522. doi: 10.1364/AO.55.001688. URL <https://opg.optica.org/abstract.cfm?URI=ao-55-7-1688>.
- [68] S. Hanany, M. D. Niemack, and L. Page. CMB Telescopes and Optical

- Systems. In *Planets, Stars and Stellar Systems*, pages 431–480. Springer, Dordrecht, 2013. ISBN 978-94-007-5621-2. doi: 10.1007/978-94-007-5621-2_10. URL https://link.springer.com/referenceworkentry/10.1007/978-94-007-5621-2_10.
- [69] C. Dragone. Offset multireflector antennas with perfect pattern symmetry and polarization discrimination. *The Bell System Technical Journal*, 57(7): 2663–2684, September 1978. ISSN 0005-8580. doi: 10.1002/j.1538-7305.1978.tb02171.x. URL <https://ieeexplore.ieee.org/document/6773968>. Conference Name: The Bell System Technical Journal.
- [70] E. Hecht. *Optics*. Pearson, 5th edition, 2017.
- [71] Z. Xu, G. E. Chesmore, et al. The Simons Observatory: Metamaterial Microwave Absorber (MMA) and its Cryogenic Applications. *Appl. Opt.*, 60(4): 864, February 2021. ISSN 1559-128X, 2155-3165. doi: 10.1364/AO.411711. URL <https://dx.doi.org/10.1364/AO.411711>.
- [72] R. Datta, C. D. Munson, M. D. Niemack, J. J. McMahon, J. Britton, E. J. Wollack, et al. Large-aperture wide-bandwidth antireflection-coated silicon lenses for millimeter wavelengths. *Appl. Opt.*, 52(36):8747–8758, Dec 2013. doi: 10.1364/AO.52.008747. URL <https://opg.optica.org/ao/abstract.cfm?URI=ao-52-36-8747>.
- [73] S. M. Simon, J. E. Golec, et al. Feedhorn development and scalability for Simons Observatory and beyond. In Jonas Zmuidzinas and Jian-Rong Gao, editors, *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*, page 147, Austin, United States, July 2018. SPIE. ISBN 978-1-5106-1969-2 978-1-5106-1970-8. doi: 10.1117/12.2313405. URL <https://dx.doi.org/10.1117/12.2313405>.

- [74] J. E. Austermann, J. A. Beall, S. A. Bryan, B. Dober, J. Gao, G. Hilton, J. Hubmayr, P. Mauskopf, C. M. McKenney, S. M. Simon, J. N. Ullom, M. R. Vissers, and G. W. Wilson. Millimeter-Wave Polarimeters Using Kinetic Inductance Detectors for TolTEC and Beyond. *Journal of Low Temperature Physics*, 193(3-4):120–127, May 2018. ISSN 1573-7357. doi: 10.1007/s10909-018-1949-5. URL <http://dx.doi.org/10.1007/s10909-018-1949-5>.
- [75] M. J. Griffin, J. J. Bock, and W. K. Gear. Relative performance of filled and feedhorn-coupled focal-plane architectures. *Appl. Opt.*, 41(31):6543, November 2002. ISSN 0003-6935, 1539-4522. doi: 10.1364/AO.41.006543. URL <https://opg.optica.org/abstract.cfm?URI=ao-41-31-6543>.
- [76] E. M. Vavagiakis et al. Prime-Cam: a first-light instrument for the CCAT-prime telescope. In Jonas Zmuidzinas and Jian-Rong Gao, editors, *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*, volume 10708, page 107081U. International Society for Optics and Photonics, SPIE, 2018. doi: 10.1117/12.2313868. URL <https://doi.org/10.1117/12.2313868>.
- [77] Eve M. Vavagiakis. *Measuring the Sunyaev-Zel’dovich Effects with Current and Future Observatories*. PhD thesis, Cornell University, August 2021. URL https://www.classe.cornell.edu/~mdn49/theses/Vavagiakis_thesis_final_2021.pdf.
- [78] N. Galitzki et al. The Simons Observatory: instrument overview. In Jonas Zmuidzinas and Jian-Rong Gao, editors, *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*, page 3, Austin, United States, July 2018. SPIE. ISBN 978-1-5106-1969-2 978-1-5106-

- 1970-8. doi: 10.1117/12.2312985. URL <https://dx.doi.org/10.1117/12.2312985>.
- [79] H. McCarrick, E. Healy, et al. The Simons Observatory microwave SQUID multiplexing detector module design. *ApJ*, 922(1):38, November 2021. ISSN 0004-637X, 1538-4357. doi: 10.3847/1538-4357/ac2232. URL <https://dx.doi.org/10.3847/1538-4357/ac2232>. arXiv:2106.14797 [astro-ph].
- [80] H. McCarrick et al. The 90 and 150 GHz universal focal-plane modules for the Simons Observatory, December 2021. URL <http://arxiv.org/abs/2112.01458>. arXiv:2112.01458 [astro-ph].
- [81] E. Healy, D. Dutcher, et al. The Simons Observatory 220 and 280 GHz Focal-Plane Module: Design and Initial Characterization. *J Low Temp Phys*, July 2022. ISSN 0022-2291, 1573-7357. doi: 10.1007/s10909-022-02788-8. URL <https://dx.doi.org/10.1007/s10909-022-02788-8>.
- [82] B. Prasad, B. Westbrook, S. Beckman, A. Mangu, C. Raum, A. Lee, A. Suzuki, D. Dutcher, K. Zheng, S. Staggs, B. R. Johnson, and Y. Wang. The Simons Observatory: design, fabrication and characterization of low-frequency detectors. In *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XII*, volume 13102, pages 370–386. SPIE, August 2024. doi: 10.1117/12.3020813. URL <https://dx.doi.org/10.1117/12.3020813>.
- [83] K.D. Irwin and G.C. Hilton. Transition-Edge Sensors. In Claus E. Ascheron, Hans J. Kölsch, Werner Skolaut, and Christian Enss, editors, *Cryogenic Particle Detection*, volume 99, pages 63–150. Springer Berlin Heidelberg, Berlin, Heidelberg, July 2005. ISBN 978-3-540-20113-7 978-3-540-31478-3. doi: 10.

1007/10933596_3. URL http://link.springer.com/10.1007/10933596_3.
Series Title: Topics in Applied Physics.

- [84] P. Day, H. LeDuc, B. Mazin, et al. A broadband superconducting detector suitable for use in large arrays. *Nature*, 425:817–821, Oct 2003. doi: 10.1038/nature02037. URL <https://doi.org/10.1038/nature02037>.
- [85] J. Zmuidzinas. Superconducting microresonators: Physics and applications. *Annual Review of Condensed Matter Physics*, 3(1):169–214, 2012. doi: 10.1146/annurev-conmatphys-020911-125022. URL <https://doi.org/10.1146/annurev-conmatphys-020911-125022>.
- [86] B. Dober, J. A. Austermann, J. A. Beall, D. Becker, G. Che, H. M. Cho, M. Devlin, S. M. Duff, N. Galitzki, J. Gao, C. Groppi, G. C. Hilton, J. Hubmayr, K. D. Irwin, C. M. McKenney, D. Li, N. Lourie, P. Mauskopf, M. R. Vissers, and Y. Wang. Optical demonstration of thz, dual-polarization sensitive microwave kinetic inductance detectors. *Journal of Low Temperature Physics*, 184(1-2):173–179, Dec 2015. ISSN 1573-7357. doi: 10.1007/s10909-015-1434-3. URL <http://dx.doi.org/10.1007/s10909-015-1434-3>.
- [87] A. Middleton, S. K. Choi, S. Walker, et al. Ccat: Led mapping and characterization of the 280 ghz tin kid array. *IEEE Transactions on Applied Superconductivity*, 35(5):1–4, 2025. doi: 10.1109/TASC.2024.3517564. URL <https://doi.org/10.1109/TASC.2024.3517564>.
- [88] A. K. Vaskuri, J. D. Wheeler, J. E. Austermann, M. R. Vissers, et al. 280-GHz aluminum MKID arrays for the Fred Young Submillimeter Telescope. *Journal of Astronomical Telescopes, Instruments, and Systems*, 11

- (2):026005, 2025. doi: 10.1117/1.JATIS.11.2.026005. URL <https://doi.org/10.1117/1.JATIS.11.2.026005>.
- [89] X. Liu, W. Guo, Y. Wang, M. Dai, L. F. Wei, B. Dober, C. M. McKenney, G. C. Hilton, J. Hubmayr, J. E. Austermann, J. N. Ullom, J. Gao, and M. R. Vissers. Superconducting micro-resonator arrays with ideal frequency spacing. *Applied Physics Letters*, 111(25):252601, December 2017. ISSN 0003-6951. doi: 10.1063/1.5016190. URL <https://doi.org/10.1063/1.5016190>.
- [90] M. R. Vissers, J. Wheeler, J. Austermann, A. Vaskuri, J. Hubmayr, J. Gao, Z. Huber, J. Imrek, and J. Ullom. Improving the yield of CCAT MKID arrays with post-measurement lithographic corrections. In Jonas Zmuidzinas and Jian-Rong Gao, editors, *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XII*, volume PC13102, page PC131020Z. International Society for Optics and Photonics, SPIE, 2024. doi: 10.1117/12.3020661. URL <https://doi.org/10.1117/12.3020661>.
- [91] E. M. Vavagiakis et al. The Simons Observatory: Magnetic Sensitivity Measurements of Microwave SQUID Multiplexers. *IEEE Transactions on Applied Superconductivity*, 31(5):1–5, August 2021. ISSN 1558-2515. doi: 10.1109/TASC.2021.3069294. URL <https://ieeexplore.ieee.org/document/9388904>. Conference Name: IEEE Transactions on Applied Superconductivity.
- [92] J. A. Connors, Z. Ahmed, J. Austermann, E. V. Denison, B. Dober, J. M. D’Ewart, J. C. Frisch, S. W. Henderson, R. Herbst, G. C. Hilton, Z. B. Huber, J. Hubmayr, Y. Li, J. A. B. Mates, M. Niemack, L. R. Vale, D. Van Winkle, E. Vavagiakis, Z. Whipps, E. Young, C. Yu, and J. Ullom. Magnetic Field

- Sensitivity of Microwave SQUID Multiplexers. *J Low Temp Phys*, 209(3): 710–717, November 2022. ISSN 1573-7357. doi: 10.1007/s10909-022-02806-9. URL <https://doi.org/10.1007/s10909-022-02806-9>.
- [93] D. Flanigan, B. R. Johnson, M. H. Abitbol, S. Bryan, R. Cantor, P. Day, G. Jones, P. Mauskopf, H. McCarrick, A. Miller, and J. Zmuidzinas. Magnetic field dependence of the internal quality factor and noise performance of lumped-element kinetic inductance detectors. *Applied Physics Letters*, 109(14):143503, October 2016. ISSN 0003-6951. doi: 10.1063/1.4964119. URL <https://doi.org/10.1063/1.4964119>.
- [94] A. K. Sinclair, J. Burgoyne, A. I. Huber, C. Murphy, S. K. Choi, et al. CCAT: detector noise limited performance of the RFSoc-based readout electronics for mm/sub-mm/far-IR KIDs. In *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XII*, volume 13102, pages 883–891. SPIE, August 2024. doi: 10.1117/12.3020557. URL <https://doi.org/10.1117/12.3020557>.
- [95] Cody J. Duell. *Superconducting Microwave Resonators for Cosmology and Astrophysics with CCAT*. PhD thesis, Cornell University, 2024. URL https://www.classe.cornell.edu/~mdn49/theses/Duell_Thesis_final_2024.pdf.
- [96] R. (Matt) Xie, A. K. Sinclair, J. Burgoyne, S. Chapman, and A. Huber. CCAT: multirate DSP for sub-mm astronomy: polyphase synthesis filter bank on FPGA for enhanced MKID readout. In *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XII*, volume 13102, pages 438–453. SPIE, August 2024. doi: 10.1117/12.3020586. URL <https://doi.org/10.1117/12.3020586>.

- [97] A. I. Huber et al. CCAT: design and performance of densely packed, high-frequency, dual-polarization kinetic inductance detectors for the Prime-Cam 850 GHz module. In *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XII*, volume 13102, pages 43–55. SPIE, August 2024. doi: 10.1117/12.3020419. URL <https://doi.org/10.1117/12.3020419>.
- [98] J. R. Burgoyne, A. K. Sinclair, et al. CCAT: FYST prime-cam readout software: a framework for massively scalable KID arrays. In *Software and Cyberinfrastructure for Astronomy VIII*, volume 13101, pages 1131–1147. SPIE, July 2024. doi: 10.1117/12.3019028. URL <https://doi.org/10.1117/12.3019028>.
- [99] 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. In *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*, volume 10708, pages 284–299. SPIE, July 2018. doi: 10.1117/12.2314435. URL <https://doi.org/10.1117/12.2314435>.
- [100] C. Yu, Z. Ahmed, J. C. Frisch, S. W. Henderson, M. Silva-Feaver, et al. SLAC microresonator RF (SMuRF) electronics: A tone-tracking readout system for superconducting microwave resonator arrays. *Review of Scientific Instruments*, 94(1):014712, January 2023. ISSN 0034-6748. doi: 10.1063/5.0125084. URL <https://doi.org/10.1063/5.0125084>.
- [101] D. V. Nguyen, S. Bhimani, N. Galitzki, B. J. Koopman, J. Lashner, L. Newburgh, M. Silva-Feaver, and K. Yamada. The Simons Observatory: alarms and detector quality monitoring. In *Millimeter, Submillimeter, and Far-*

- Infrared Detectors and Instrumentation for Astronomy XII*, volume 13102, pages 1010–1023. SPIE, August 2024. doi: 10.1117/12.3018481. URL <https://doi.org/10.1117/12.3018481>.
- [102] S. Bhimani, J. Lashner, et al. The Simons Observatory: deployment and current configuration of the observatory control system for SAT-MF1 and data access software systems. In *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XII*, volume 13102, pages 525–539. SPIE, August 2024. doi: 10.1117/12.3020237. URL <https://doi.org/10.1117/12.3020237>.
- [103] BICEP2 Collaboration. Detection of B -Mode Polarization at Degree Angular Scales by BICEP2. *Phys. Rev. Lett.*, 112(24):241101, June 2014. doi: 10.1103/PhysRevLett.112.241101. URL <https://link.aps.org/doi/10.1103/PhysRevLett.112.241101>. Publisher: American Physical Society.
- [104] M. Tegmark. How to Make Maps from Cosmic Microwave Background Data without Losing Information. *ApJ*, 480(2):L87, May 1997. ISSN 0004-637X. doi: 10.1086/310631. URL <https://iopscience.iop.org/article/10.1086/310631/meta>. Publisher: IOP Publishing.
- [105] P. A. Gallardo, J. Gudmundsson, B. J. Koopman, F. T. Matsuda, S. M. Simon, et al. Systematic uncertainties in the Simons Observatory: optical effects and sensitivity considerations. In Jonas Zmuidzinas and Jian-Rong Gao, editors, *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*, volume 10708, page 107083Y. International Society for Optics and Photonics, SPIE, 2018. doi: 10.1117/12.2312971. URL <https://doi.org/10.1117/12.2312971>.

- [106] M. Born and E. Wolf. *Principles of Optics*. Cambridge University Press, 7th edition, 2019.
- [107] N. F. Cothard et al. The Design of the CCAT-prime Epoch of Reionization Spectrometer Instrument. *Journal of Low Temperature Physics*, 199(3-4): 898–907, Jan 2020. ISSN 1573-7357. doi: 10.1007/s10909-019-02297-1. URL <http://dx.doi.org/10.1007/s10909-019-02297-1>.
- [108] R. Freundt, Y. Li, D. Henke, et al. CCAT: a status update on the EoR-Spec instrument module for Prime-Cam. In *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XII*, volume 13102, pages 343–352. SPIE, August 2024. doi: 10.1117/12.3018477. URL <https://doi.org/10.1117/12.3018477>.
- [109] Z. B. Huber et al. CCAT-prime: the optical design for the epoch of reionization spectrometer. In Jonas Zmuidzinas and Jian-Rong Gao, editors, *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XI*, volume 12190, page 121902H. International Society for Optics and Photonics, SPIE, 2022. doi: 10.1117/12.2629744. URL <https://doi.org/10.1117/12.2629744>.
- [110] A. I. Huber, S. C. Chapman, A. K. Sinclair, L. D. Spencer, et al. CCAT-prime: optical and cryogenic design of the 850 GHz module for Prime-Cam. In Jonas Zmuidzinas and Jian-Rong Gao, editors, *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XI*, volume 12190, page 121901D. International Society for Optics and Photonics, SPIE, 2022. doi: 10.1117/12.2630370. URL <https://doi.org/10.1117/12.2630370>.
- [111] P. A. Gallardo, R. Puddu, K. Harrington, B. Benson, J. E. Carlstrom, S. R.

- Dicker, N. Emerson, J. E. Gudmundsson, M. Limon, J. McMahon, J. M. Nagy, T. Natoli, M. D. Niemack, S. Padin, J. Ruhl, S. M. Simon, and The CMB-S4 Collaboration. Freeform three-mirror anastigmatic large-aperture telescope and receiver optics for CMB-S4. *Appl. Opt., AO*, 63(2):310–321, January 2024. ISSN 2155-3165. doi: 10.1364/AO.501744. URL <https://opg.optica.org/ao/abstract.cfm?uri=ao-63-2-310>. Publisher: Optica Publishing Group.
- [112] R. Schaaf, M. Nolta, R. Higgins, T. Nikola, and R. A. Timmermann. CCAT-prime: observatory control software for FYST. In *Software and Cyberinfrastructure for Astronomy VII*, volume 12189, pages 922–927. SPIE, August 2022. doi: 10.1117/12.2629999. URL <https://dx.doi.org/10.1117/12.2629999>.
- [113] B. J. Koopman, S. Bhimani, N. Galitzki, M. Hasselfield, J. Lashner, H. Nakata, L. Newburgh, D. V. Nguyen, T. Sakuma, and K. Yamada. The Simons Observatory: deployment of the observatory control system and supporting infrastructure. In *Software and Cyberinfrastructure for Astronomy VIII*, volume 13101, pages 1463–1481. SPIE, July 2024. doi: 10.1117/12.3020710. URL <https://doi.org/10.1117/12.3020710>.
- [114] T. Sakuma, Y. Guan, M. Hasselfield, B. Koopman, L. Newburgh, and D. Nguyen. Nextline, February 2025. URL <https://doi.org/10.5281/zenodo.14889734>.
- [115] Jack Lashner. *Data acquisition systems and detector characterization for Simons Observatory*. PhD thesis, University of Southern California, January 2023.
- [116] Jason Shing-Yan Leung. *From the Ends of the Earth to the Edge of the*

Universe: Building CMB Experiments in Extreme Environments. Ph.D., University of Toronto (Canada), 2024. URL <https://www.proquest.com/pqdtglobal/docview/3034132279/abstract/A087B32F682C4522PQ/1>. ISBN: 9798382192758.

- [117] S. Weinberg. A New Light Boson? *Phys. Rev. Lett.*, 40:223–226, Jan 1978. doi: 10.1103/PhysRevLett.40.223. URL <https://link.aps.org/doi/10.1103/PhysRevLett.40.223>.
- [118] F. Wilczek. Problem of Strong P and T Invariance in the Presence of Instantons. *Phys. Rev. Lett.*, 40:279–282, Jan 1978. doi: 10.1103/PhysRevLett.40.279. URL <https://link.aps.org/doi/10.1103/PhysRevLett.40.279>.
- [119] P. Svrcek and E. Witten. Axions in string theory. *Journal of High Energy Physics*, 2006(06):051, jun 2006. doi: 10.1088/1126-6708/2006/06/051. URL <https://dx.doi.org/10.1088/1126-6708/2006/06/051>.
- [120] A. Arvanitaki, S. Dimopoulos, S. Dubovsky, N. Kaloper, and J. March-Russell. String axiverse. *Phys. Rev. D*, 81:123530, Jun 2010. doi: 10.1103/PhysRevD.81.123530. URL <https://link.aps.org/doi/10.1103/PhysRevD.81.123530>.
- [121] D. J. E. Marsh. Axion cosmology. *Physics Reports*, 643:1–79, July 2016. ISSN 0370-1573. doi: 10.1016/j.physrep.2016.06.005. URL <https://www.sciencedirect.com/science/article/pii/S0370157316301557>.
- [122] K. Ehret, M. Frede, S. Ghazaryan, M. Hildebrandt, E. Knabbe, D. Kracht, A. Lindner, J. List, T. Meier, N. Meyer, D. Notz, J. Redondo, A. Ringwald, G. Wiedemann, and B. Willke. New ALPS results on hidden-sector lightweights. *Physics Letters B*, 689(4):149–155, 2010. ISSN 0370-

2693. doi: <https://doi.org/10.1016/j.physletb.2010.04.066>. URL <https://www.sciencedirect.com/science/article/pii/S0370269310005526>.
- [123] R. Ballou, G. Deferne, M. Finger, M. Finger, L. Flekova, J. Hosek, S. Kunc, K. Macuchova, K. A. Meissner, P. Pagnat, M. Schott, A. Siemko, M. Slunicka, M. Sulc, C. Weinsheimer, and J. Zicha. New exclusion limits on scalar and pseudoscalar axionlike particles from light shining through a wall. *Phys. Rev. D*, 92:092002, Nov 2015. doi: 10.1103/PhysRevD.92.092002. URL <https://link.aps.org/doi/10.1103/PhysRevD.92.092002>.
- [124] ADMX Collaboration. Axion Dark Matter eXperiment: Run 1A analysis details. *Phys. Rev. D*, 109:012009, Jan 2024. doi: 10.1103/PhysRevD.109.012009. URL <https://link.aps.org/doi/10.1103/PhysRevD.109.012009>.
- [125] A. Ejlli, F. Della Valle, U. Gastaldi, G. Messineo, R. Pengo, G. Ruoso, and G. Zavattini. The PVLAS experiment: A 25 year effort to measure vacuum magnetic birefringence. *Physics Reports*, 871:1–74, August 2020. ISSN 0370-1573. doi: 10.1016/j.physrep.2020.06.001. URL <https://www.sciencedirect.com/science/article/pii/S0370157320302428>.
- [126] R. Hložek, D. J E Marsh, and D. Grin. Using the full power of the cosmic microwave background to probe axion dark matter. *Monthly Notices of the Royal Astronomical Society*, 476(3):3063–3085, May 2018. ISSN 0035-8711. doi: 10.1093/mnras/sty271. URL <https://doi.org/10.1093/mnras/sty271>.
- [127] R. Hlozek, D. Grin, D. J. E. Marsh, and P. G. Ferreira. A search for ultralight axions using precision cosmological data. *Phys. Rev. D*, 91(10):103512, May

2015. doi: 10.1103/PhysRevD.91.103512. URL <https://link.aps.org/doi/10.1103/PhysRevD.91.103512>. Publisher: American Physical Society.
- [128] W. Hu, R. Barkana, and A. Gruzinov. Fuzzy Cold Dark Matter: The Wave Properties of Ultralight Particles. *Phys. Rev. Lett.*, 85:1158–1161, Aug 2000. doi: 10.1103/PhysRevLett.85.1158. URL <https://link.aps.org/doi/10.1103/PhysRevLett.85.1158>.
- [129] L. Hui, J. P. Ostriker, S. Tremaine, and E. Witten. Ultralight scalars as cosmological dark matter. *Phys. Rev. D*, 95(4):043541, February 2017. doi: 10.1103/PhysRevD.95.043541. URL <https://link.aps.org/doi/10.1103/PhysRevD.95.043541>. Publisher: American Physical Society.
- [130] N. Bar, K. Blum, and C. Sun. Galactic rotation curves versus ultralight dark matter: A systematic comparison with SPARC data. *Phys. Rev. D*, 105:083015, Apr 2022. doi: 10.1103/PhysRevD.105.083015. URL <https://link.aps.org/doi/10.1103/PhysRevD.105.083015>.
- [131] CAST Collaboration. New Upper Limit on the Axion-Photon Coupling with an Extended CAST Run with a Xe-Based Micromegas Detector. *Phys. Rev. Lett.*, 133(22):221005, November 2024. doi: 10.1103/PhysRevLett.133.221005. URL <https://link.aps.org/doi/10.1103/PhysRevLett.133.221005>. Publisher: American Physical Society.
- [132] A. Payez, C. Evoli, T. Fischer, M. Giannotti, A. Mirizzi, and A. Ringwald. Revisiting the SN1987A gamma-ray limit on ultralight axion-like particles. *J. Cosmol. Astropart. Phys.*, 2015(02):006, February 2015. ISSN 1475-7516. doi: 10.1088/1475-7516/2015/02/006. URL <https://dx.doi.org/10.1088/1475-7516/2015/02/006>.

- [133] D.M.C. Marsh, H. R. Russell, A. C. Fabian, B. R. McNamara, P. Nulsen, and C. S. Reynolds. A new bound on axion-like particles. *J. Cosmol. Astropart. Phys.*, 2017(12):036, December 2017. ISSN 1475-7516. doi: 10.1088/1475-7516/2017/12/036. URL <https://dx.doi.org/10.1088/1475-7516/2017/12/036>.
- [134] J. Sisk-Reynés, J. H. Matthews, C. S. Reynolds, H. R. Russell, R. N. Smith, and D. M. C. Marsh. New constraints on light axion-like particles using Chandra transmission grating spectroscopy of the powerful cluster-hosted quasar H1821+643. *Monthly Notices of the Royal Astronomical Society*, 510(1):1264–1277, February 2022. ISSN 0035-8711. doi: 10.1093/mnras/stab3464. URL <https://doi.org/10.1093/mnras/stab3464>.
- [135] S. M. Carroll, G. B. Field, and R. Jackiw. Limits on a Lorentz- and parity-violating modification of electrodynamics. *Phys. Rev. D*, 41(4):1231–1240, February 1990. doi: 10.1103/PhysRevD.41.1231. URL <https://link.aps.org/doi/10.1103/PhysRevD.41.1231>. Publisher: American Physical Society.
- [136] A. Caputo, L. Sberna, M. Frías, D. Blas, P. Pani, L. Shao, and W. Yan. Constraints on millicharged dark matter and axionlike particles from timing of radio waves. *Phys. Rev. D*, 100(6):063515, September 2019. doi: 10.1103/PhysRevD.100.063515. URL <https://link.aps.org/doi/10.1103/PhysRevD.100.063515>. Publisher: American Physical Society.
- [137] A. Castillo, J. Martin-Camalich, J. Terol-Calvo, et al. Searching for dark-matter waves with PPTA and QUIJOTE pulsar polarimetry. *J. Cosmol. Astropart. Phys.*, 2022(06):014, June 2022. ISSN 1475-7516. doi: 10.1088/

- 1475-7516/2022/06/014. URL <https://dx.doi.org/10.1088/1475-7516/2022/06/014>. Publisher: IOP Publishing.
- [138] X. Xue, S. Dai, H. N. Luu, T. Liu, J. Ren, J. Shu, Y. Zhao, A. Zic, et al. First Pulsar Polarization Array Limits on Ultralight Axion-like Dark Matter, December 2024. URL <http://arxiv.org/abs/2412.02229>. arXiv:2412.02229 [astro-ph].
- [139] M.M. Ivanov, Y.Y. Kovalev, M.L. Lister, A.G. Panin, A.B. Pushkarev, T. Savolainen, and S.V. Troitsky. Constraining the photon coupling of ultralight dark-matter axion-like particles by polarization variations of parsec-scale jets in active galaxies. *J. Cosmol. Astropart. Phys.*, 2019(02):059, February 2019. ISSN 1475-7516. doi: 10.1088/1475-7516/2019/02/059. URL <https://dx.doi.org/10.1088/1475-7516/2019/02/059>.
- [140] POLARBEAR Collaboration. Exploration of the polarization angle variability of the Crab Nebula with POLARBEAR and its application to the search for axionlike particles. *Phys. Rev. D*, 110(6):063013, September 2024. ISSN 2470-0010, 2470-0029. doi: 10.1103/PhysRevD.110.063013. URL <https://link.aps.org/doi/10.1103/PhysRevD.110.063013>.
- [141] T. Louis, A. La Posta, Z. Atkins, H. T. Jense, et al. The Atacama Cosmology Telescope: DR6 Power Spectra, Likelihoods and Λ CDM Parameters, March 2025. URL <http://arxiv.org/abs/2503.14452>. arXiv:2503.14452 [astro-ph].
- [142] T. Namikawa, Y. Guan, O. Darwish, B. D. Sherwin, et al. Atacama Cosmology Telescope: Constraints on cosmic birefringence. *Phys. Rev. D*, 101(8):083527, April 2020. doi: 10.1103/PhysRevD.101.083527. URL

- <https://link.aps.org/doi/10.1103/PhysRevD.101.083527>. Publisher: American Physical Society.
- [143] BICEP/Keck Collaboration. BICEP/Keck XII: Constraints on axion-like polarization oscillations in the cosmic microwave background. *Phys. Rev. D*, 103(4):042002, Feb 2021. ISSN 2470-0010, 2470-0029. doi: 10.1103/PhysRevD.103.042002. URL <https://link.aps.org/doi/10.1103/PhysRevD.103.042002>.
- [144] BICEP/Keck Collaboration. BICEP/Keck XIV: Improved constraints on axionlike polarization oscillations in the cosmic microwave background. *Phys. Rev. D*, 105(2):022006, Jan 2022. doi: 10.1103/PhysRevD.105.022006. URL <https://link.aps.org/doi/10.1103/PhysRevD.105.022006>. Publisher: American Physical Society.
- [145] SPT-3G Collaboration, K. R. Ferguson, et al. Searching for axionlike time-dependent cosmic birefringence with data from SPT-3G. *Phys. Rev. D*, 106(4):042011, August 2022. doi: 10.1103/PhysRevD.106.042011. URL <https://link.aps.org/doi/10.1103/PhysRevD.106.042011>. Publisher: American Physical Society.
- [146] POLARBEAR Collaboration. Constraints on axionlike polarization oscillations in the cosmic microwave background with POLARBEAR. *Phys. Rev. D*, 108(4):043017, August 2023. doi: 10.1103/PhysRevD.108.043017. URL <https://link.aps.org/doi/10.1103/PhysRevD.108.043017>. Publisher: American Physical Society.
- [147] E. Hivon, K. M. Górski, C. B. Netterfield, B. P. Crill, S. Prunet, and F. Hansen. MASTER of the Cosmic Microwave Background Anisotropy

- Power Spectrum: A Fast Method for Statistical Analysis of Large and Complex Cosmic Microwave Background Data Sets. *ApJ*, 567(1):2, March 2002. ISSN 0004-637X. doi: 10.1086/338126. URL <https://iopscience.iop.org/article/10.1086/338126/meta>. Publisher: IOP Publishing.
- [148] J. T. VanderPlas. Understanding the Lomb-Scargle Periodogram. *ApJS*, 236(1):16, May 2018. ISSN 0067-0049. doi: 10.3847/1538-4365/aab766. URL <https://dx.doi.org/10.3847/1538-4365/aab766>. Publisher: The American Astronomical Society.
- [149] G. P. Centers, J. W. Blanchard, J. Conrad, N. L. Figueroa, A. Garcon, A. V. Gramolin, D. F. Jackson Kimball, M. Lawson, B. Pelssers, J. A. Smiga, A. O. Sushkov, A. Wickenbrock, D. Budker, and A. Derevianko. Stochastic fluctuations of bosonic dark matter. *Nature Communications*, 12(1):7321, December 2021. ISSN 2041-1723. doi: 10.1038/s41467-021-27632-7. URL <https://www.nature.com/articles/s41467-021-27632-7>. Publisher: Nature Publishing Group.
- [150] J. R. Eskilt and E. Komatsu. Improved constraints on cosmic birefringence from the WMAP and Planck cosmic microwave background polarization data. *Phys. Rev. D*, 106:063503, Sep 2022. doi: 10.1103/PhysRevD.106.063503. URL <https://link.aps.org/doi/10.1103/PhysRevD.106.063503>.
- [151] S. Das, T. A. Marriage, et al. THE ATACAMA COSMOLOGY TELESCOPE: A MEASUREMENT OF THE COSMIC MICROWAVE BACKGROUND POWER SPECTRUM AT 148 AND 218 GHz FROM THE 2008 SOUTHERN SURVEY. *ApJ*, 729(1):62, February 2011. ISSN 0004-637X.

doi: 10.1088/0004-637X/729/1/62. URL <https://dx.doi.org/10.1088/0004-637X/729/1/62>. Publisher: The American Astronomical Society.