

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Precision Measurements of CMB Secondary Anisotropies with the South Pole Telescope

Permalink

<https://escholarship.org/uc/item/77n6v8qt>

Author

Huang, Nicholas

Publication Date

2024

Peer reviewed|Thesis/dissertation

Precision Measurements of CMB Secondary Anisotropies with the South Pole Telescope

by

Nicholas Huang

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Physics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor William Holzappel, Chair

Professor Adrian Lee

Professor Martin White

Fall 2024

Precision Measurements of CMB Secondary Anisotropies with the South Pole Telescope

Copyright 2024
by
Nicholas Huang

Abstract

Precision Measurements of CMB Secondary Anisotropies with the South Pole Telescope

by

Nicholas Huang

Doctor of Philosophy in Physics

University of California, Berkeley

Professor William Holzapfel, Chair

Observations of the sky at millimeter wavelengths are of critical importance to modern cosmology. Measurements of the primary cosmic microwave background (CMB) anisotropies, its secondary anisotropies, and foreground objects all contribute to our understanding the universe. As the oldest light in the universe, the CMB contains imprints from much of cosmic history. Measurements of these secondary anisotropies provides a window into the evolution of the universe, including the growth of the largest structures in the universe, and the process of reionization by early stars and galaxies.

These measurements are only possible with a high-resolution telescope and extremely precise calibration. This dissertation is primarily concerned with the design and calibration of the SPT-3G camera on the South Pole Telescope. In particular, the calibration of the raw detector data, and the measurement of the telescope beams are described in detail. These pieces are essential for making a precise measurement of the mm-wave sky, and properly accounting for the effects of the telescope optics. In the last chapter, we show how all of this comes together to produce a catalog of galaxy clusters discovered through the thermal Sunyaev Zel'dovich effect. We report 89 candidates, with 81 confirmed by optical and infrared follow-up observations. 29 are reported for the first time. This provides a mass-limited catalog of galaxy clusters over 100 square degrees.

Contents

Contents	i
List of Figures	iii
List of Tables	xi
1 The Cosmic Microwave Background	1
1.1 Big Bang Cosmology and the History of the Universe	1
1.2 CMB Observables	6
1.3 CMB Foregrounds and Secondary Anisotropies	8
1.4 Conclusion	18
2 The South Pole Telescope and SPT-3G Camera	19
2.1 SPT-3G Camera Overview	19
2.2 Vacuum Window and IR Filtering	26
2.3 Scan Strategy	32
2.4 Conclusion	33
3 Calibration and Mapmaking	34
3.1 Calibration	34
3.2 Detector Pointing Offsets	45
3.3 Mapmaking	47
3.4 Conclusion	51
4 SPT-3G Beams	52
4.1 Systematic Effects	53
4.2 Data	60
4.3 Map Preparation	62
4.4 Parameter Variations	71
4.5 Beam and Covariance Estimation	72
4.6 Systematic Tests	74
4.7 Results and Discussion	81
4.8 Ongoing Work	83

4.9	Conclusion	85
5	Galaxy Clusters Selected via the Sunyaev-Zel'dovich Effect in the SPTPol 100-Square-Degree Survey	87
5.1	Millimeter-wave Data	88
5.2	Cluster Extraction and Characterization from mm-wave Data	90
5.3	Optical and IR Followup	99
5.4	The Cluster Catalog	101
5.5	Conclusions	106
6	Conclusion	108
	Bibliography	110
A	The Cluster Catalog	124

List of Figures

1.1	A schematic history of the universe. Time increases from the Big Bang at the left to today at the right. Figure from (WMAP Science Team).	2
1.2	The CMB as measured by WMAP. The all-sky map is shown in panel a, and its power spectrum is shown in panel b. The y-axis of panel b is $\ell(\ell + 1)/2\pi$ times the power spectrum. This factor makes a scale-invariant power spectrum constant.	7
1.3	Three maps of the same region of the sky at ~ 150 GHz. They come from WMAP(panel a), <i>Planck</i> (panel b), and SPTPol (panel c) data with approximate resolutions of 13, 7, and 1.5 arcminutes, respectively. In panel c, there are several bright point sources (small white spots) and a few massive galaxy clusters (small black spots) visible. These foreground objects are only visible with higher spatial resolution. Figure adapted from (Benson, 2018).	9
1.4	Power spectra (auto and cross) from SPTPol and SPT-SZ maps. These power spectra show the estimated contribution from a variety of foregrounds and secondary anisotropies. These components of the mm-wave sky only dominate above the CMB on smaller angular scales ($\ell \gtrsim 3000$ or $\theta \lesssim 4'$).	10
1.5	A schematic representation of several possible lines of sight to the CMB. The CMB is emitted on the left, and passes through a variety of galaxies. Black galaxies are mostly neutral, while red galaxies are fully reionized. The change in the color of the line indicates ΔT_{kSZ} , the temperature shift induced by the kSZ effect. The top line of sight interacts with a neutral galaxy, and a galaxy with no peculiar velocity ($v_{\parallel} = 0$), so $\Delta T_{\text{kSZ}} = 0$. In the next line of sight, the kSZ effect from two galaxies cancels. In the third, we have added one additional galaxy that imparts a net kSZ effect. On the final line of sight, all the ionized galaxies have a bulk motion towards the observer, which makes the largest ΔT_{kSZ} . This is a very large scale velocity mode.	11
1.6	Power spectra for the Ostriker-Vishniac effect and primary CMB under different reionization scenarios. As the redshift of reionization (z_{ri}) increases, the optical depth to the CMB increases as well. This decreases the amplitude of the primary CMB, and increases the Ostriker-Vishniac power spectrum. Regardless of the reionization scenario, the Ostriker-Vishniac effect dominates over the primary fluctuations at $\ell \gtrsim 3000$. Figure from Hu and White (1996).	13

1.7	Three simulations of reionization. White regions are neutral, and black regions are fully ionized. $z = 0$ (today) is on the left, and $z = 20$ is on the right. The middle history varies the ionization efficiency of galaxies over cosmic history, leading to longer reionization. The bottom panel decreases the mean free path of ionizing photons, resulting in both longer reionization, and smaller bubbles. Figure adapted from Mesinger et al. (2012)	14
1.8	The spectrum of the tSZ $f(x)$ (top panel) and the apparent change in temperature after passing through a cluster with $y = 10^{-4}$ (bottom). Below $\nu \approx 220$ GHz, the apparent change in temperature is negative. The dashed gray lines show the approximate center of the SPT-3G observing bands (see Chapter 2).	15
2.1	A cross section of the SPT and the SPT-3G camera. Panel a shows the optical path of light coming from the right, and incident on the primary mirror, which is then reflected to the secondary and tertiary mirrors before entering the camera on the bottom right. Panel b shows a detailed view of the camera itself, with 3 alumina cold lenses, the Lyot stop, and the detectors. Both figures are from S22.	20
2.2	The bandpass functions for the three SPT-3G observing bands. The colored lines show the bandpass of each wafer, while the black dashed line shows the average. The gray region shows the typical atmospheric transmission at the South Pole. Figure from S22.	21
2.3	A schematic representation of a bolometer. In the SPT-3G camera, the absorber is a resistor that thermalizes the current from the sinuous antenna, and the thermistor is a TES. Figure from (Benson, 2022).	22
2.4	The resistance as a function of temperature of a typical TES. The transition from normal metal to superconductor is very steep, and nonlinear. Figure from (Benson, 2022).	22
2.5	The SPT-3G focal plane. The detectors are housed in the hexagonal structures in the center. The four temperature stages are visible (4 K, 1 K, IC, and UC). Figure from S22.	25
2.6	A schematic of the vacuum window and IR filtering portions of the camera.	26
2.7	One antireflection surface cut into a HDPE test piece.	27
2.8	The theoretical transmission for one grooved antireflection surface (panel b), and the measured transmission through two surfaces, and 30 mm of HDPE (panel b). Note that the theoretical transmission spectra indicate different transmission depending on the polarization state of the incoming radiation. The grooves on the top and bottom surfaces of our vacuum window are orthogonal, so this differential transmission cancels after passing through both surfaces. The measured spectra were taken at four different locations on the window. The dashed lines in panel a show the centers of the SPT-3G bands.	27
2.9	A schematic representation of our model for the multi-layer IR filter.	29

2.10	The transmission spectra of several different polyethylene foams at two thicknesses in the far (panel a), and mid-IR (panel b). HD foams are made of high density polyethylene, while LD foams are made of low density polyethylene. The numbers after the dash indicate the mean density of the foam, in kg / m^3 , and the final number indicates the thickness of the foam layer, in inches. Measurements and plots by Ki Won Yoon (Yoon, 2016).	30
3.1	Maps of RCW 38 (panel a) and MAT 5a (panel b). The maps are shown with a logarithmic color axis to show the extended emission. These maps are in “source relative” units (see §3.1.1), so each map peaks at approximately 1.0. The dark strips on either sides of bright features are “filtering wings” from the mapmaking process (see §3.3).	36
3.2	TOD from one detector while observing the calibrator. This detector measures the calibrator the signal with $\text{SNR} = 83$ over the full 1-minute observation.	38
3.3	Typical maps of the HII regions RCW 38 (panel a) and MAT 5a (panel b) at 90 GHz. These maps are made using the TOD from a single detector. We show these maps on a linear scale to emphasize the very high SNR we achieve, even with a single detector.	39
3.4	The mean lightcurves, which we use to estimate the calibration error. Panel a shows the observations calibrated using RCW 38, and panel b shows those calibrated using MAT 5a. The subplot titles refer to the mean declination of the subfield in which the bright sources are found. The significant outliers are discussed in more detail in the text; however they are mostly explained by external factors.	42
3.5	Histograms of the mean lightcurves. The orange lines show the same quantity calculated using the original calibration pipeline; the blue histogram shows the calibration pipeline described here. There were some adjustments to the mapmaking process (specifically in masking and glitch finding) that reduce the scatter somewhat (as seen by the higher peak in the central bin). However, the largest difference is the data selection criteria, which simply removes the outlier population centered near 0.2 (see §3.1.2). The gray band indicates the range used to estimate the calibration error.	43
3.6	Histograms of the per-detector signal to noise in long HII region observations over ~ 1 year of observations. The vast majority of detectors have a high SNR on the source over multiple observations. We discard detectors with $\text{SNR} < (60, 60, 20)$ at (90, 150, 220) GHz, which is only a few percent of the total. Note that the x-axis in on the 220 GHz plot is significantly smaller.	44
3.7	The HII region MAT 5a without individual detector pointing offsets applied. The map gets one copy of the source for each pixel with in the focal plane (as long as it has at least one live detector). In this image, each detector is in uncalibrated electrical power, which leads to the variations in source brightness.	45

3.8	A toy model of the coverage from a single observation. The thick gray line shows the true signal, and the dashed gray lines show the boundaries of the map pixels. The colored triangles show two detectors sampling that signal, with a sub-pixel size pointing offset. The top panel shows where the true samples are taken. The bottom panel shows where the samples appear after binning into a map.	46
3.9	Histograms showing the difference between detector offsets under different splits of the data. Panel a shows the difference between pairs of detectors in the same pixel. Both of these detectors share the exact same optical path, and only differ by the polarization state they are sensitive to. Panel b shows the difference between pointing offsets calculated from two different sources, MAT 5a and RCW38. In both cases, the left plot shows differences in the x (R.A.) direction, and the right plot shows the differences in the y (declination) direction.	48
3.10	A map of a very bright source made without masking the source location. The negative strips on either side of the source are “filtering wings,” which can be avoided by masking the source. The minimum of the filtering wing is ~ -0.25 times the peak amplitude of the source.	49
4.1	Calibrator response as a function of frequency (filled circles) and the transfer function for the best-fit time constant for a sampling of detectors. The response for each detector has been normalized to 1.0 at the lowest frequency. Figure by Zhaodi Pan.	53
4.2	The distribution of the mean time constants (τ) for every detector over the 2019 and 2020 observing seasons. The 220 GHz detectors are somewhat faster than the 90 and 150 GHz detectors.	54
4.3	A 1d demonstration of the effect of the time constant on the beam. Panel a shows the true beam (gray dashed), and the beams after the application of a time constant (blue and orange). To first order, each beam is shifted in the direction of the scan. Panel b shows the beam after two different coadding schemes are applied. The orange line shows the standard coaddition of the left- and right-going scans (in which the beam is broadened parallel to the scan direction). The blue line shows the beam when using only the approaching portion of each scan, as we do with the planet maps. In this case, the beam is narrower than expected.	55
4.4	Radial profiles of the (approximate) beam from the brightest source (black dashed), and 60 simulated backgrounds (colored lines). The simulated backgrounds include Gaussian realizations of the CMB, CIB, and thermal and kinematic Sunyaev-Zel’dovich effects, as well as Poisson realizations of radio sources.	56
4.5	The conversion from a beam measured on an AGN (with $I(\nu) \propto \nu^{-0.5}$) to a beam appropriate for the CMB spectrum. This is a negligible effect at 220 GHz, small at 150 GHz, and substantial at 90 GHz.	60
4.6	A comparison of Saturn maps made from left-going (a), and right-going (b) scans, and right-going scans after cleaning (c). The right-going scans show an oscillatory mode which decays from left to right.	64

- 4.7 This shows the Saturn maps coadded without masking, minus the same maps coadded with our masking procedure. The x and y scales are marked in arcminutes from the source center, and the color scale is in units of the peak brightness of the source. The unmasked map is analogous to the orange line in Figure 4.3b, and the masked map is analogous to the blue line. As expected, the unmasked maps are wider in the scan direction, but equal to the masked maps at the center. The 150 GHz map shows the largest effect, as the detectors are operated closer to their saturation power than the other bands. 64
- 4.8 This shows the Saturn maps after accounting for the time constant and offset minus the same maps before accounting for the time constant and offset (these are the maps produced by the masking step). The x and y scales are marked in arcminutes from the source center, and the color scale is in units of the peak brightness of the source. The masked map is analogous to the blue line in Figure 4.3b, while the map after accounting for the time constants and time offset is closest to the mean of the orange and blue lines in Figure 4.3a. The overall effect is to narrow Saturn in the scan direction. 65
- 4.9 The Saturn maps before accounting for the RJ spectrum of Saturn, minus the same maps after accounting for the spectrum. The x and y scales are marked in arcminutes from the source center, and the color scale is in units of the peak brightness of the source. This effect is isotropic, by construction, and makes Saturn broader (since the effective observing frequency on a RJ spectrum is higher than that of an AGN-like spectrum). It is also several times smaller than the other processes applied to the Saturn maps. 66
- 4.10 The map of one AGN before accounting for the time constants, minus the same map after accounting for the time constants. The x and y scales are marked in arcminutes from the source center, and the color scale is in units of the peak brightness of the source. This source is below the mean elevation of the field, so it is more extended in the scan direction (horizontal) than the mean. The very high frequency modes in the scan direction are induced by inverting the transfer function, which has zeros in it. However, these occur at extremely high ℓ ($\approx 40,000$), and are irrelevant to the SPT-3G beam. 68
- 4.11 Radial profiles of the (approximate) beam from the brightest source (black dashed), and 60 simulated backgrounds (colored lines). The simulated backgrounds include Gaussian realizations of the CMB, CIB, and thermal and kinematic Sunyaev-Zel'dovich effects, as well as Poisson realizations of radio sources. Panels a and b show the radial profiles before and after background subtraction, respectively. 69
- 4.12 The fractional error on the beam recovered from a simulation. The dominant effect is a multiplicative scaling, which is degenerate with absolute calibration (and much smaller than our uncertainty on the absolute calibration). The variation with multipole ℓ is $\mathcal{O}(10^{-4})$, and much smaller than the beam uncertainties. 75

- 4.13 The scatter in 100 simulated beams (black dashed line), and the estimated noise from each simulation (blue lines). The black line falls within the expected scatter, indicating the the covariance calculation we use is unbiased. 76
- 4.14 The difference between two AGN maps at different declinations (a source at $\delta = -47.5^\circ$ minus a source at $\delta = -57.6^\circ$). The x and y scales are marked in arcminutes from the source center, and the color scale is in units of the peak brightness of the source. Panel a is before accounting for the time constants and time offset, while b is after. Before processing, the higher declination (lower elevation) source is broadened in the scan direction relative to the lower declination source. This effect does not appear at 220 GHz because the SNR of these sources is insufficient at 220 GHz. After processing, the differences have disappeared. 78
- 4.15 The map of one of the brightest AGN minus a coadd of the Saturn maps. The x and y scales are marked in arcminutes from the source center, and the color scale is in units of the peak brightness of the source. Panel a is before applying the map processing steps (except recentering), while b is after all the processing steps. The dashed circles show the fiducial stitching region. Inside the inner circle, the stitched beam map uses only the AGN; outside the outer circle, the stitched maps use only Saturn; the region between the circles is the transition region. The residuals are significantly reduced at 90 and 220 GHz, but are nearly unaffected at 150 GHz. 79
- 4.16 The SPT-3G beams (B_ℓ) at 90, 150, and 220 GHz (top panel), and their fractional uncertainties (bottom panel). The beams are normalized such that $B_{800} = 1$. The uncertainty shown here is $\sqrt{\mathbf{C}_{bb}}$, i.e. the square root of the diagonal of the covariance matrix. This uncertainty is for qualitative purposes only, and any formal analysis must use the full covariance matrix. 80
- 4.17 Panel a shows the radial profiles of the beams. The high-frequency noise in the radial profiles is due to our corrections for the planet disk and source spectrum. Panel b shows the fraction of the beam area at radius $> r$. This is the reverse cumulative distribution function of the beam area, and is useful as an intuitive reference point for the sidelobe power on a given scale. 81
- 4.18 The most significant contributions to the beam uncertainty. Saturn and AGN refer to variations between the different Saturn and AGN maps, respectively. α_{AGN} refers to uncertainty in the spectral index and detector bandpasses. r_{inner} refers to the inner stitching radius (where we start to use the Saturn maps). BG r_{inner} refers to the inner radius of the region used to fit the background model. The contributions from varying the outer radii of the stitching and background fit regions, and the uncertainty on the time constants and time offset are not shown. 82
- 5.1 Panels a and b each show a 3° -by- 3° cutout at 95 and 150 GHz, respectively, with several galaxy cluster detections circled in orange. These cutouts are taken from the larger map used in this work, which was produced as described in §5.1.2. The black dashed region is shown in Figure 5.2, after the matched filter is applied. 91

- 5.2 Panel a is the region of sky outlined with a black dashed line in Figure 5.1, after the matched filter is applied. This region is 50'0 on each side. The filtering is described in §5.2.1, and this map has been filtered to optimally find objects with $\theta_c = 0'.25$. It contains 2 detected galaxy clusters, SPT-CL J2323-5752 at $z = 1.3$, and SPT-CL J2320-5807 at $z = 0.56$, both of which are detected at a signal-to-noise-ratio greater than 5.0. A composite image of tSZ contours, optical, and infrared images for SPT-CL J2323-5752 is shown in Figure 5.4. Panel b shows the azimuthally-averaged spatial filter used to produce panel a for each observing band. Since the tSZ has a larger magnitude at 95 GHz, the 95 GHz filter remains positive over the entire range. The 150 GHz filter is negative at intermediate l s to subtract CMB fluctuations. The peak of the 95 GHz filter is lower than the peak of the 150 GHz filter due to the relative noise in the 95 and 150 GHz SPTPol maps. 92
- 5.3 Simulated false detection rate, as calculated from approximately 1000 deg² of simulated sky (10 simulations of the sky area used in this work). This shows the density of false detections above a given significance (ξ_{\min}) on the left side, and the expected number of false detections for the entire 100 deg² survey area on the right. We expect 8 ± 2 false detections at our significance cutoff, $\xi_{\min} = 4.6$. . . 96
- 5.4 Four new $z > 1$ clusters identified in the SPTPol 100d catalog. Clockwise from upper left: SPT-CL J2259-5301 at $z = 1.16 \pm 0.09$ detected at $\xi = 5.1$; SPT-CL J2336-5252 ($z = 1.22 \pm 0.09, \xi = 6.7$); SPT-CL J2323-5752 ($z = 1.3 \pm 0.1, \xi = 5.1$), and SPT-CL J2355-5514 ($z = 1.4 \pm 0.1, \xi = 4.9$). The composite *rgb* images are created using Magellan/PISCO *g + r*-band data, Magellan/LDSS3-C *z*-band data, and *Spitzer* IRAC 3.6 μm data from the SSDF; detection contours from the SPTPol $\theta_c=0'.25$ filtered maps are overlaid. 98
- 5.5 The mass and redshift distributions for several recent SZ-selected galaxy cluster catalogs. We plot the estimated mass and redshift for each of the 79 galaxy clusters with measured redshifts from this catalog, 517 from the SPT 2500 deg² survey (B15), 182 from the ACTPol-ED56 survey (Hilton et al., 2018), 91 from the ACT survey (Hasselfield et al., 2013a), and 1653 from the second *Planck* galaxy cluster catalog (Planck Collaboration et al., 2016a). The black line is the forecast detection threshold (50 photon counts) for eROSITA (Pillepich et al., 2012). . . 101
- 5.6 The estimated completeness as a function of M_{500c} for the SPTPol 100d catalog. Completeness is estimated at four redshifts: 0.25, 0.50, 1.00, and 1.50. For fixed mass, the completeness at low redshift decreases, because the angular extent of the clusters becomes similar to the fluctuations in the CMB. We expect our catalog to be complete for $M_{500c} > 2.6 \times 10^{14} M_{\odot} h_{70}^{-1}$ 103

- 5.7 A merging cluster, SPT-CL J2331-5052, identified in the SPTPol 100d sample. This system was detected as two separate clusters in the B15 catalog, but due to the deeper mm-wave data used in this work, it is detected as a single system. The composite image is created using PISCO *gri* images. The contours show tSZ detection significance from the $\theta_c = 0.5$ match-filtered map. The position reported in this catalog is in the upper right, near the center of the $\xi = 16$ contour. In B15, the smaller object (SPT-CL J2332-5053) is at $z = 0.56 \pm 0.04$; SPT-CL J2331-5052 is one of the 13 clusters with spectroscopically measured redshifts at $z = 0.576$. For the more massive object in the upper right of this image, $R_{500} \approx 2.8$ 104
- 5.8 Galaxy cluster masses as measured by SPTPol and XXL. Both the XXL and SPTPol mass uncertainties are calculated using a scaling relation with an intrinsic scatter in the mass-observable relation (see §5.2.4), which leads to an X-ray or SZ mass uncertainty that is highly correlated between clusters in the sample. Only 13 of the 14 clusters found in both surveys are included, because the mass is not provided for one of the XXL clusters. 106

List of Tables

1.1	Parameters of the Λ CDM model. The first six parameters are typically free, while others (including those below the second horizontal line in this table) are derived. Parameter values are from Planck Collaboration et al. (2018).	4
4.1	Model parameters defining the frequency dependency of the beam. ν_{eff} is the effective frequency at which σ_{obs} is measured, assuming a source spectrum $I(\nu) \propto \nu^{-0.5}$. σ_{obs} is the best-fit Gaussian width of our brightest source. $\sigma_{\text{diffraction}}$ is the frequency-dependent portion of the width, and σ_{geom} is the frequency-independent portion. By definition, $\sigma_{\text{geom},90} = 0$, and $\sigma_{\text{obs}}^2 = \sigma_{\text{diffraction}}^2 + \sigma_{\text{geom}}^2$	59
4.2	The AGN used for beam measurement.	61
4.3	Parameters of the input beams used for simulations. For each band, the beam is modeled as $B_\nu(r) \propto \exp(-r^2/2/\sigma_{\text{main}}^2) + a \exp(-r^2/2/\sigma_{\text{sidelobe}}^2)$	62
4.4	Fiducial input parameters for the beam pipeline.	71
A.1	Galaxy cluster candidates with $\xi \geq 4.6$	125
A.2	SPTPol clusters matched to other catalogs.	129

Acknowledgments

Grad school is a long and difficult endeavor, but I have had a great time. That is mainly due to the people I've worked with, both here at UC Berkeley and in the larger SPT collaboration. I can't possibly list everyone who had a positive impact on my experience as a grad student, but I'm going to do my best.

My first and biggest thanks go to my advisor, Bill Holzapfel. He has given me an enormous amount of freedom to find my own path through grad school, and was always available when I got in too deep. Bill has always expected me to think critically about my own work, and learning to meet those expectations has given me the skills to work independently. His advising has given me the opportunity to be the best scientist I can be. And one more "Thanks!" to Bill for showing me all the great bike routes around Berkeley.

The experimental CMB group (members of both Adrian Lee's and Bill Holzapfel's groups) here at UC Berkeley has always been a huge help. In my early days at UC Berkeley, Nick Harrington, Liz George, Bryan Steinbach, and Christian Reichardt all spent time showing me the ropes and teaching me about cosmology, and the SPT. They made me feel welcome, and I was never embarrassed to ask a basic question. Our large shared office turned into a great space for both collaboration and fun, in no small part due to Neil Goeckner-Wald, Charlie Hill, Nathan Whitehorn, Shawn Beckman (and his infinite patience for our stupid jokes), and Ian Shirley. Finally, I want to express my gratitude to my fellow grad students in Bill's group: Anna Coerver, Sam Guns, and Jessica Zebrowski. They all helped to make my experience in grad school a pleasure (and they're now making it very bittersweet to graduate!). As a former theorist, Sam has been my go-to every time I have a strange question about statistics or a formal calculation. Anna has been a great sounding board, and one of the few people I can look forward to seeing in our office post-pandemic. Jessica and I started grad school together, and I can't imagine having a better person with whom to take on grad school. From problem sets, to prelims, to deploying SPT-3G, and untangling data mysteries, we both struggled and helped each other push through.

The UC Berkeley Physics Department provided substantial support for my research. Our R&D machine shop team was instrumental to significant portions of my research projects. As experimental physicists, we often need to design and build custom parts. But with little to no formal training, we can get in over our heads quickly. Having professional (and extremely skilled) machinists to help with designing and making parts bridges that gap. I want to thank Warner Carlisle for consulting on several large parts for the SPT-3G vacuum window. In particular, he helped me get my engineering drawings into a state that would be comprehensible to an outsider, even though we knew the parts couldn't be made in-house. Gordon Long and Tommy Gutierrez worked very closely with us to develop the tooling for the vacuum window. The end product would not have come together without their creativity and problem solving skills. I also deeply appreciate our administrative staff. In particular, Joelle Miles, Anne Takizawa, and Carlos Bustamante made my life noticeably easier with their expertise and consistently upbeat demeanors.

My path to grad school was not exactly normal. I joined the South Pole Telescope collaboration as a winterover, before starting grad school. Without a doubt, that experience set me on the path I am on today. I am eternally grateful to the SPT management (John Carlstrom and Bill Holzapfel, especially) for letting a 22 year-old just out of an undergrad degree operate their \$20-million telescope for an entire year.

None of this would have been possible without the wider SPT collaboration. As a whole, the collaboration has been a very welcoming and safe space to me. I had a lot of support from the collaboration in learning how to be a scientist, and then was given the opportunity to make significant impacts in the collaboration and beyond. I particularly want to thank John Carlstrom for supporting me, and somehow finding time for my concerns in his very busy schedule; Lindsey Bleem, Christian Reichardt, and Tijmen de Haan for guiding me through my first major analysis project; and Tom Crawford for sharing his incredible depth of analysis expertise and near-infinite patience.

Finally, I want to thank people who contributed to this work directly: Bill and Jessica read early drafts of this dissertation, and pointed out many things I'd overlooked; Shawn, Sam, Anna, and Margaret McCarter helped keep me sane and grounded while I put words to paper.

Even though I'm the only author on this dissertation, it really would not have come together without a lot of team work. Thank you!

Chapter 1

The Cosmic Microwave Background

Measurements of the cosmic microwave background (CMB) have been a workhorse of modern cosmology. They have been instrumental in building a simple model of the universe that explains everything from the abundance of helium in the universe, to the distribution of dark matter and galaxies. In this dissertation, we will describe our understanding of the universe, and how the CMB fits into it. From there, we can begin to see how CMB measurements made (and continue to make) significant contributions to that understanding. We will also describe measurements of the CMB using the South Pole Telescope, including the current camera that makes those measurements possible. We will detail how the raw data are calibrated and transformed into an image of the sky, and finally show one of the scientific results that comes from these measurements.

1.1 Big Bang Cosmology and the History of the Universe

Figure 1.1 depicts the entire history of our universe. Three major eras are immediately evident. The first is the very early universe (on the far left), which is extremely hot and dense, and nearly homogeneous. Then come the dark ages, which ended with the ignition of the first stars in the “cosmic dawn.” From there, we get the evolution of galaxies and clusters of galaxies until we get to the highly structured universe we observe today.

In Figure 1.1, the vertical scale indicates the relative size of the universe. We parametrize this by the scale factor a . The scale factor relates the proper distance¹ d to the comoving distance d_c at time t :

$$d = a(t)d_c \tag{1.1}$$

By definition, the $a = 1$ today, and $a \rightarrow 0$ at the Big Bang. The rate of the expansion is the Hubble parameter H :

$$H(t) \equiv \frac{\dot{a}}{a} \tag{1.2}$$

¹Proper distance is the physical distance that can be measured with e.g. a ruler.




Figure 1.1: A schematic history of the universe. Time increases from the Big Bang at the left to today at the right. Figure from (WMAP Science Team).

This expansion leads to the famous Hubble diagram, in which distant galaxies all appear to be receding from the observer, with more distant galaxies receding faster (Hubble, 1929). Since $H(t) > 0$ for all t (i.e. $a(t)$ is monotonically increasing), it is common to use a , or the redshift z , as a proxy for time. The redshift is given by $a = (1 + z)^{-1}$.

Expansion plays a critical role in the very early universe. Immediately after the Big Bang, the universe underwent a period of exponential expansion called inflation. The expansion was superluminal, and enlarged quantum fluctuations of the matter density into macroscopic density variations that source all of the structure in the universe. After inflation, the universe was still extremely hot, and filled with a plasma of fundamental particles. These extreme temperatures exceeded the chemical potential of fundamental particles. Particle-antiparticle pairs were continuously created and annihilated. However, expansion caused the universe to adiabatically cool, and the particle densities to decrease. Once the temperature drops below the chemical potential μ_i of a particle species i , it can no longer be created in any significant number.

As the density decreased, particle interactions became less frequent. At some point, the interaction rate Γ_i for a particle species dropped below H . This means that a given particle was not expected to have an interaction in the time it took the universe to double in size. Once both of these conditions are met ($T < \mu_i$, and $\Gamma_i < H$), the universe can no longer create or destroy a significant number of these particles, so their comoving density is fixed. This is called “freeze out” or decoupling, and we cover the case of photon decoupling in §1.1.2. The equilibrium density of particles at the time of decoupling is responsible for determining

their (comoving) density for the rest of history. For the formation of light atomic nuclei, it is called “Big Bang nucleosynthesis” (BBN), and for the decoupling of photons from matter, it is called “recombination.” BBN occurs at 3 minutes after the Big Bang (Zyla et al., 2020), and recombination happens 380,000 years later (Planck Collaboration et al., 2018).

At the end of recombination, the universe consisted primarily of a photon background, and nearly homogeneous matter. Those photons are still observable at mm wavelengths, and form the CMB. The matter is mostly dark matter with a small fraction of baryons. The baryons themselves were primarily neutral hydrogen, with an even smaller fraction of other light elements. As time passes, the overdense regions of the dark matter gravitationally attract more matter, becoming even more dense. These eventually form distinct high density halos, which are the seeds of galaxies and (much later) galaxy clusters. The baryonic matter is also gathered into these halos. After $\mathcal{O}(10^8)$ years, some of the overdensities of hydrogen have grown massive enough to begin fusion, and the first stars ignite.

The light from these first stars begins the final phase change of the universe. They emit ultraviolet photons that are capable of ionizing the surrounding hydrogen gas. Initially, only small bubbles around the first stars and galaxies are ionized. However, these bubbles slowly grow to encompass the entire universe, in a process called “reionization.” Observations of the Gunn-Peterson trough (Gunn and Peterson, 1965) indicate that reionization is complete ~ 1 billion years after the Big Bang ($z = 5.5-6$) (Becker et al., 2001, 2015; Zhu et al., 2022).

The physics of each of these eras leaves observable tracers. The goal of modern cosmology is to measure those tracers, and reconstruct the history of the universe with great precision.

1.1.1 Λ CDM

Λ CDM is a 6-parameter model that formalizes our understanding of the Big Bang model and the history of the universe (the parameters are defined in Table 1.1). It is built on three major components:

1. Λ , the cosmological constant associated with dark energy
2. Cold dark matter (CDM)
3. Ordinary matter

By modeling these components using general relativity and the Boltzmann equation (see §1.1.2), Λ CDM successfully incorporates the major observations in cosmology. In particular, it is able to predict the photon background that decouples during recombination, the accelerating expansion of the universe, and the large scale distribution of matter into galaxies and galaxy clusters.

The free and derived parameters shown in Table 1.1 represent only one possible configuration. It is possible to refactor the model to exchange derived parameters for fixed parameters, or to fit different combinations of parameters. For example, galaxy surveys are maximally sensitive to $S_8 = \sigma_8 \sqrt{\Omega_m/0.3}$. There are also a variety of extensions to the

Parameter	Value	Description
$\ln(10^{10} A_s)$	3.043 ± 0.014	Amplitude of the primordial density power spectrum
n_s	0.9652 ± 0.0042	Power law index of the primordial density power spectrum
$\Omega_b h^2$	0.02233 ± 0.00015	Baryon energy density times the expansion rate.
$\Omega_m h^2$	0.1198 ± 0.0012	Dark matter energy density times the expansion rate.
τ	0.0540 ± 0.0074	Optical depth to recombination
θ	$0.01041 \pm 3.1 \times 10^{-5}$	Sound horizon at recombination
H_0	67.37 ± 0.54	Expansion rate today
σ_8	0.8101 ± 0.0061	Amplitude of the matter power spectrum
Ω_Λ	0.6847 ± 0.0073	Energy density of dark energy
z_{rec}	1089 ± 0.25	The redshift of recombination

Table 1.1: Parameters of the Λ CDM model. The first six parameters are typically free, while others (including those below the second horizontal line in this table) are derived. Parameter values are from Planck Collaboration et al. (2018).

base Λ CDM model, which account for additional physics. Including inflation in the model (typically) introduces the parameter r , which quantifies the ratio of energy in tensor modes (gravitational waves) to energy in scalar modes (density perturbations) before recombination. Other extensions include the sum of the neutrino masses as a free parameter, or allow dark energy to evolve through its equation of state.

1.1.2 Recombination and the Boltzmann Equation

We now return to the details of particle freeze out, particularly as it applies to photons. Let's assume a reaction of the form $1 + 2 \rightleftharpoons 3 + 4$, where each number labels a particle type, and \rightleftharpoons indicates that the reaction goes in either direction. For example, the reaction for recombination/reionization is $e + p \rightleftharpoons H + \gamma$. For any particular particle species, we can estimate when it freezes out using the Boltzmann equation², which describes the time evolution of its number density. The Boltzmann equation for such a reaction is (Dodelson, 2003)

$$a^{-3} \frac{d(n_1 a^3)}{dt} = n_1^{\text{eq}} n_2^{\text{eq}} \langle \sigma v \rangle \left(\frac{n_3 n_4}{n_3^{\text{eq}} n_4^{\text{eq}}} - \frac{n_1 n_2}{n_1^{\text{eq}} n_2^{\text{eq}}} \right) \quad (1.3)$$

²For a derivation in the context of an expanding universe see Dodelson (2003)

where $\langle\sigma v\rangle$ is the thermally averaged cross section, n_i is the number density of species i , and n_i^{eq} is the equilibrium number density of a species:

$$n_i^{\text{eq}} = \begin{cases} f_i \left(\frac{m_i T}{2\pi}\right)^{3/2} \exp\left(-\frac{m_i}{T}\right) & m_i \gg T \\ g_i \frac{T^3}{\pi^2} & m_i \ll T \end{cases} \quad (1.4)$$

where g_i is the degeneracy of the species. $\Gamma_1 = n_1^{\text{eq}} n_2^{\text{eq}} \langle\sigma v\rangle$ is the rate at which the reaction occurs, and the term in parentheses in equation 1.3 is a measure of how far the current number densities are from equilibrium. In broad strokes, the Boltzmann equation states that the number of a given species moves towards equilibrium at a rate that is proportional to the interaction rate and how far from equilibrium it is. For a species that is still coupled (i.e. Γ_i is large), the only way to keep the number density constant is for the terms in parentheses to cancel. This leads to the Saha equation (Saha, 1920):

$$\frac{n_3 n_4}{n_3^{\text{eq}} n_4^{\text{eq}}} = \frac{n_1 n_2}{n_1^{\text{eq}} n_2^{\text{eq}}} \quad (1.5)$$

Note that the Saha equation only holds when the reaction rate is large. If the reaction rate becomes negligible, the species can fall out of equilibrium (this is the mathematical definition of freeze out). We typically consider a species frozen out if the reaction rate falls below the Hubble rate, i.e.

$$H \approx n_1^{\text{eq}} n_2^{\text{eq}} \langle\sigma v\rangle. \quad (1.6)$$

Let us now consider the case of recombination, when photons decouple from matter. Naively, we would expect recombination to occur at a temperature of $T = 13.6$ eV ($\sim 160,000$ K), which is the ionization energy of a ground state electron in hydrogen. However, any time a free electron is captured by a proton, and goes directly to the ground state, it releases another ionizing photon. As long as the reaction rate is large, that photon will immediately ionize another hydrogen atom, conserving the number of free electrons. To find the temperature at which photons actually decouple, we start with the Saha equation (equation 1.5).

$$\frac{n_e n_p}{n_H} = \left(\frac{m_e T}{2\pi}\right)^{-3/2} \exp\left(\frac{Q}{T}\right) \quad (1.7)$$

where n_e , n_p and n_H are the number densities of electrons, protons and neutral hydrogen, respectively; m_e is the electron mass; T is the temperature; and $Q = 13.6$ eV is the ionization potential of hydrogen³. The Saha equation is only valid near equilibrium, so this will not allow us to derive the full time evolution of n_e through recombination, but it does allow us to estimate when most of the free electrons and protons have combined into hydrogen.

Since the universe has no net electrical charge, $n_e = n_p$. We define the ionization fraction

$$X \equiv \frac{n_e}{n_e + n_H} = \frac{n_e}{n_b} = \frac{n_p}{n_b} \quad (1.8)$$

³We are ignoring the contribution from heavier elements (such as helium) since they exist only in negligible quantities in the early universe.

where n_b is the baryon number density. Substituting this into equation 1.5 gives

$$\frac{1-X}{X^2} = n_b \left(\frac{m_e T}{2\pi} \right)^{-3/2} \exp\left(\frac{Q}{T}\right). \quad (1.9)$$

Finally, we express n_b in terms of the number density of photons, n_γ : $n_b \approx 6 \times 10^{-10} n_\gamma$ (Dodelson, 2003). The ratio of photons to baryons is derived from Ω_b , the energy density of baryons, one of the Λ CDM parameters (see §1.1.1). From equation 1.4, $n_\gamma = 2T^3/\pi^2$, so

$$\frac{1-X}{X^2} \approx 1.2 \times 10^{-9} T^3 \left(\frac{m_e T}{2\pi} \right)^{-3/2} \exp\left(\frac{Q}{T}\right) \quad (1.10)$$

$$= 1.2 \times 10^{-9} \left(\frac{m_e}{2\pi T} \right)^{-3/2} \exp\left(\frac{Q}{T}\right) \quad (1.11)$$

When $T \approx Q$, the right hand side of equation 1.11 is $\sim 10^{-15}$, and $X \approx 1$ (i.e. the universe is completely ionized). If we define T_{rec} as the temperature at which $X = 0.5$, we find $T_{\text{rec}} \approx 0.3 \text{ eV} \approx 3000 \text{ K}$.

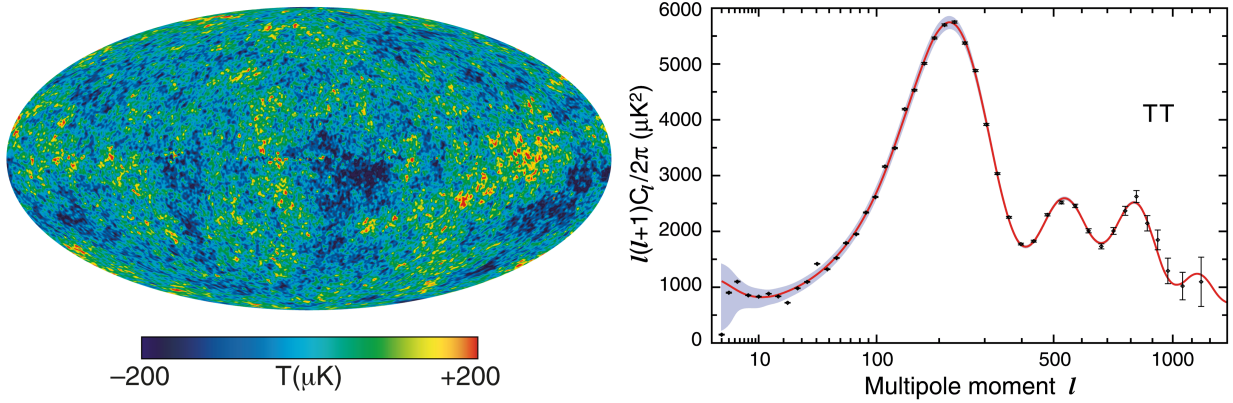
After recombination, this leaves us with a population of photons at $\sim 3000 \text{ K}$. Since they are decoupled, they will free stream through the universe, with a very low probability of interacting with anything else. As expansion proceeds, the temperature of the photons drops, and today they are observable as a 2.726 K background (Fixsen, 2009). This puts their peak intensity in the mm range ($\sim 150 \text{ GHz}$), so the photon background is called the cosmic microwave background. Recombination happens over a (cosmologically) short period of time, so it is often sufficient to approximate recombination as a spherical shell⁴ (which we call the surface of last scattering, since it is the last scattering event for the vast majority of CMB photons).

In this section, we have treated the plasma before recombination as homogeneous. In practice it has small density variations that were sourced by inflation. The density variations also lead to temperature variations in the plasma, which slightly vary the temperature of CMB. So, the CMB is a near-uniform blackbody that gives us an image of the surface of last scattering. It has very small temperature variations which map onto density variations at the surface of last scattering.

1.2 CMB Observables

The Wilkinson Microwave Anisotropy Probe (WMAP) made all-sky observations at mm-wavelengths, and produced a map of the CMB (Figure 1.2a). This shows the very small

⁴One case where we cannot ignore it is when predicting the small-scale variations in the CMB. On scales approximately equal to the duration of recombination times the speed of light, the photons are still able to bring disparate regions into thermal equilibrium. This erases variations at those scales and smaller, in a process called Silk damping (Silk, 1968).



(a) The map of the CMB produced by WMAP. (b) The power spectrum of the map in panel a

Figure 1.2: The CMB as measured by WMAP. The all-sky map is shown in panel a, and its power spectrum is shown in panel b. The y-axis of panel b is $\ell(\ell + 1)/2\pi$ times the power spectrum. This factor makes a scale-invariant power spectrum constant.

deviations from the mean blackbody temperature, which reflect very small temperature differences at the surface of last scattering. The temperature differences are a Gaussian random field, so we are only concerned with the statistics of the map, rather than the map itself.

For a Gaussian random field, the most natural⁵ metric to use is the power spectrum. The power spectrum measures the level of correlation between two points on the map as a function of the distance between those points. Due to the curvature of the sky, we write the power spectrum in terms of the spherical harmonic decomposition of the sky. For a map $M(\theta, \phi)$,

$$M(\theta, \phi) = \sum_{\ell=1}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m} Y_{\ell m}(\theta, \phi) \quad (1.12)$$

where $Y_{\ell m}$ represents the spherical harmonics. Larger ℓ corresponds to smaller angular scales. The power spectrum is

$$C_{\ell} = \left\langle a_{\ell m} a_{\ell' m'}^{\dagger} \right\rangle \quad (1.13)$$

where \dagger indicates the complex conjugate. Due to the orthogonality of the spherical harmonics, $\left\langle a_{\ell m} a_{\ell' m'}^{\dagger} \right\rangle \propto \delta_{\ell\ell'} \delta_{mm'}$, so the average is taken over different values of m . The WMAP power spectrum is shown in Figure 1.2b.

Over a small region, we can approximate the curved sky as a flat plane. In this case, the power spectrum becomes the regular autocorrelation of the map. Specifically, for a map

⁵For a Gaussian random field, the power spectrum contains all of the information. All higher order statistics are zero.

$M(x, y)$, its power spectrum is

$$C_\ell = \frac{1}{2\pi\ell f_{\text{sky}}^2} \int_{\ell_x^2 + \ell_y^2 = \ell^2} \widetilde{M}(\ell_x, \ell_y) \widetilde{M}^\dagger(\ell_x, \ell_y) d\ell_x d\ell_y \quad (1.14)$$

$$= \frac{1}{2\pi\ell f_{\text{sky}}^2} \int_{\ell_x^2 + \ell_y^2 = \ell^2} \left| \widetilde{M}(\ell_x, \ell_y) \right|^2 d\ell_x d\ell_y \quad (1.15)$$

where \widetilde{M} is the Fourier transform of M , and (ℓ_x, ℓ_y) is the Fourier transform pair of (x, y) . f_{sky} is the fraction of the full sky that our small region covers. Its presence in the power spectrum compensates for the reduced variance in a small area, assuming the small patch is statistically identical to the remainder of the sky.

Using the Boltzmann equation and the Λ CDM model, we can make accurate predictions of the CMB power spectrum (Blas et al., 2011; Lewis et al., 2000). We can turn this around, and use measurements of the CMB power spectrum to determine the free parameters of the Λ CDM model (e.g. (Balkenhol et al., 2023; Planck Collaboration et al., 2018; Aiola et al., 2020)).

1.3 CMB Foregrounds and Secondary Anisotropies

When we observe the sky at mm wavelengths, we do not observe only the CMB, nor do we observe a pristine copy of the surface of last scattering. Various sources in the foreground contribute their own mm-wave emission, and a small fraction of CMB photons are perturbed by matter between us and the surface of last scattering. These are referred to as foregrounds and secondary anisotropies, respectively. Figure 1.3 shows the mm-wave sky at three different resolutions. The CMB completely dominates the large scales ($\gtrsim 4'$). However, at small scales, we can see both bright point sources and the thermal Sunyaev-Zel'dovich effect from galaxy clusters (see §1.3.3).

Figure 1.4 shows a series of power spectra that go to much smaller scales (higher ℓ) than the WMAP power spectrum. The red dots are the measured power spectra, and the lines all show different components that contribute to the total. On large angular scales ($\ell \lesssim 3000$ or $\theta \gtrsim 4'$), the CMB dominates the power spectra. However, at smaller scales, foregrounds and secondary anisotropies are brighter than the CMB.

1.3.1 Foregrounds

The most prominent mm-wave foregrounds fall into one of three categories: thermal emission from galactic or extragalactic dust, and synchrotron emission from active galactic nuclei (AGN). The thermal emission from dust mostly comes from star forming regions in our galaxy, or galaxies undergoing significant star formation. The dust is heated to several tens of Kelvin, so dusty sources typically have a rising spectrum (intensity increasing with frequency). The extragalactic dusty sources are also the source of the cosmic infrared background (CIB),

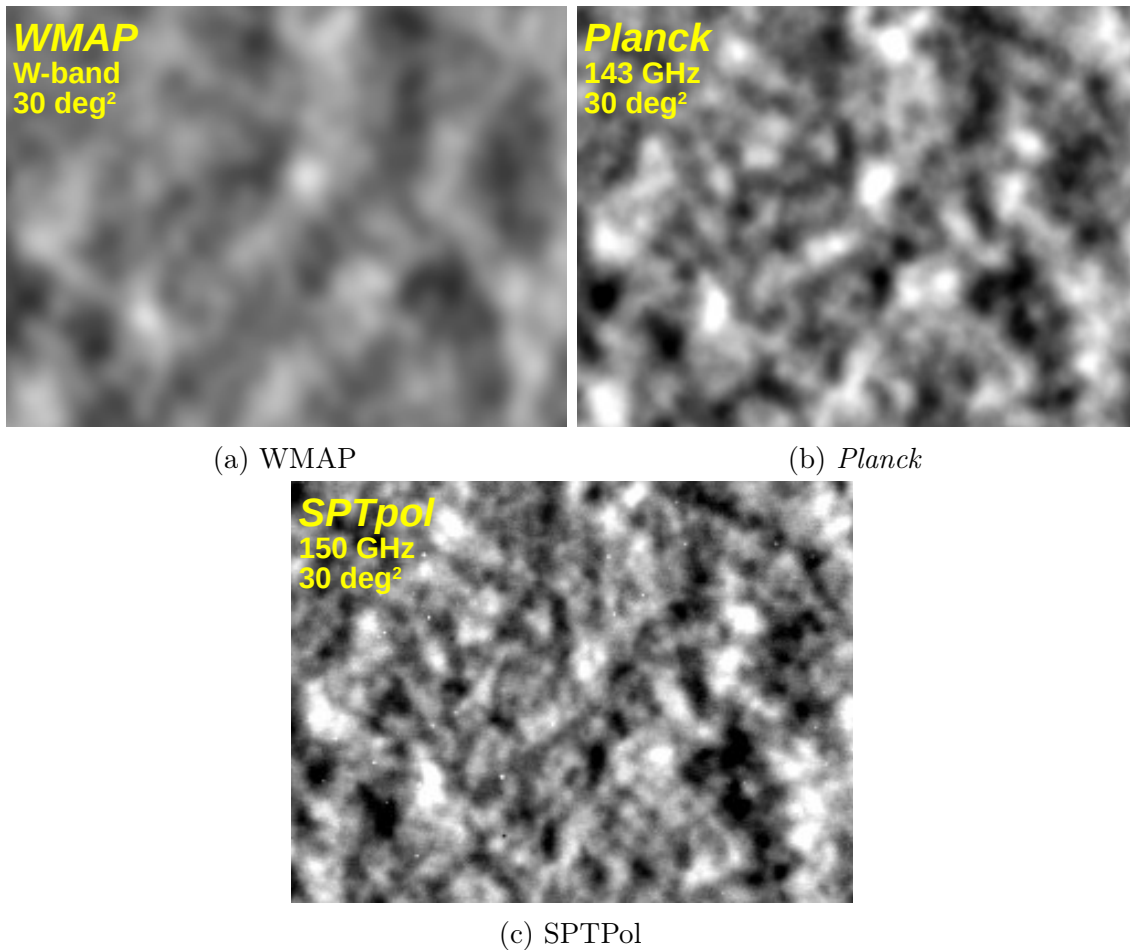


Figure 1.3: Three maps of the same region of the sky at ~ 150 GHz. They come from WMAP (panel a), *Planck* (panel b), and SPTPol (panel c) data with approximate resolutions of 13, 7, and 1.5 arcminutes, respectively. In panel c, there are several bright point sources (small white spots) and a few massive galaxy clusters (small black spots) visible. These foreground objects are only visible with higher spatial resolution. Figure adapted from (Benson, 2018).

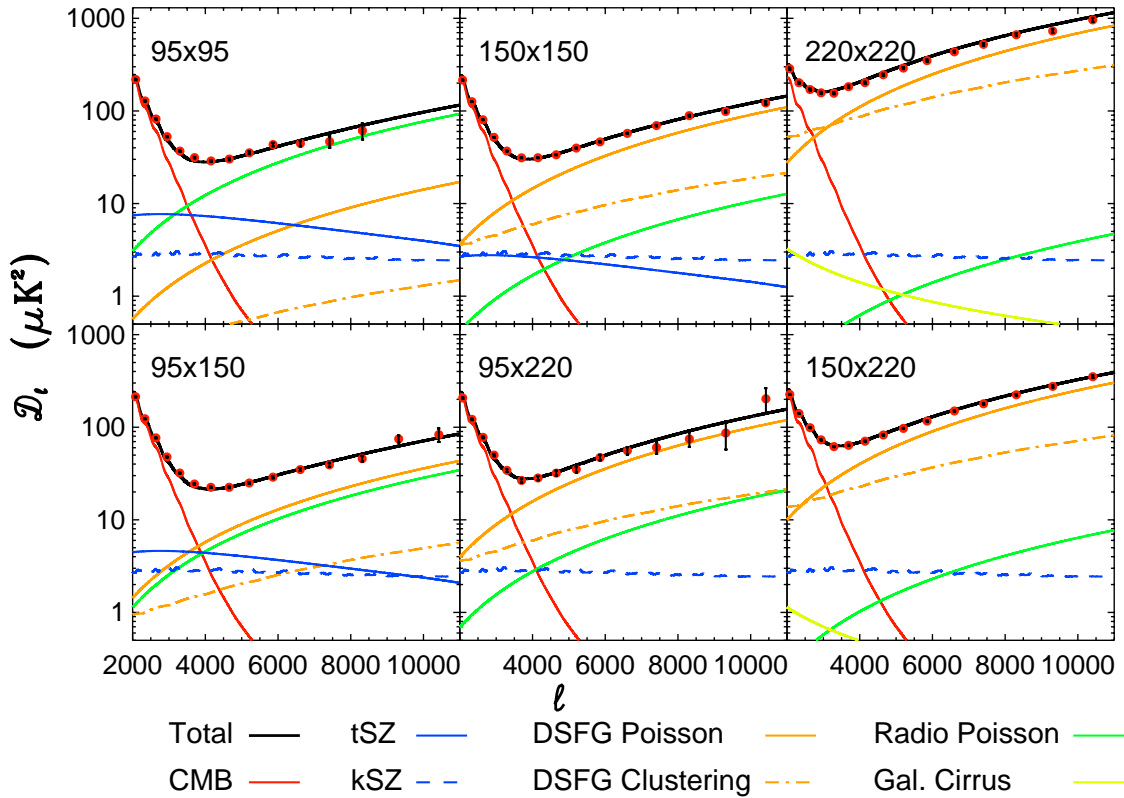


Figure 1.4: Power spectra (auto and cross) from SPTPol and SPT-SZ maps. These power spectra show the estimated contribution from a variety of foregrounds and secondary anisotropies. These components of the mm-wave sky only dominate above the CMB on smaller angular scales ($\ell \gtrsim 3000$ or $\theta \lesssim 4'$).

so we often refer to them as CIB sources. The synchrotron emission from AGN has a falling spectrum. With a sufficiently high resolution observation, the bright individual sources are detectable (see Figure 1.3c). However, the majority of the sources are dim and blend together to form an additional background, which is detectable in the power spectrum (e.g. the orange and green lines in Figure 1.4).

1.3.2 Secondary Anisotropies

Secondary anisotropies are caused by a variety of interactions between CMB photons and matter. There are two gravitational effects: lensing, the bending of light by gravitational potentials; and the integrated Sachs-Wolf effect (ISW), which is caused by the time evolution of gravitational potentials over the time it takes a CMB photon to pass through them. Gravitational lensing operates perpendicular to the line of sight, and tends to smooth out the peaks in the CMB power spectrum. The ISW acts along the line of sight, and changes the temperature of the CMB on very large scales. In this section, we will focus on two different

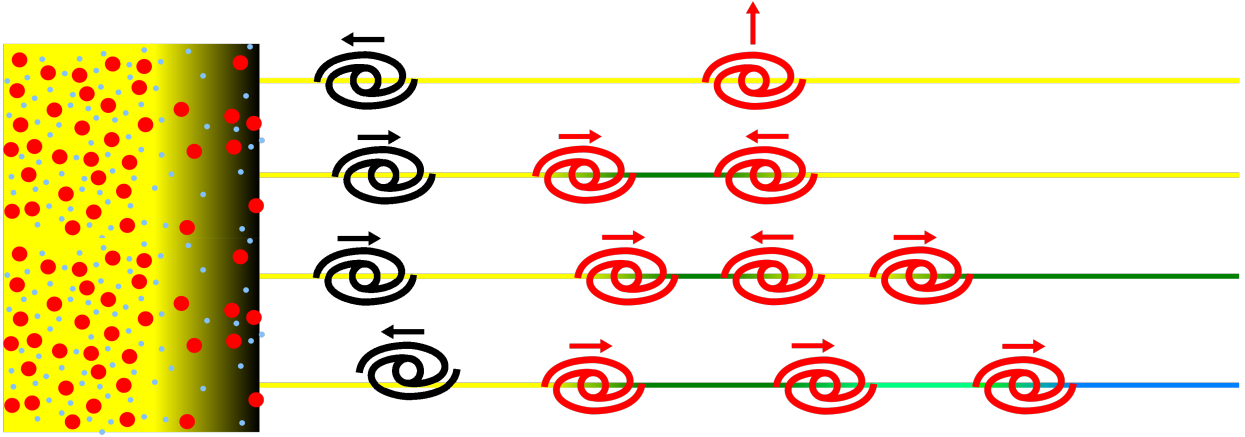


Figure 1.5: A schematic representation of several possible lines of sight to the CMB. The CMB is emitted on the left, and passes through a variety of galaxies. Black galaxies are mostly neutral, while red galaxies are fully reionized. The change in the color of the line indicates ΔT_{kSZ} , the temperature shift induced by the kSZ effect. The top line of sight interacts with a neutral galaxy, and a galaxy with no peculiar velocity ($v_{\parallel} = 0$), so $\Delta T_{\text{kSZ}} = 0$. In the next line of sight, the kSZ effect from two galaxies cancels. In the third, we have added one additional galaxy that imparts a net kSZ effect. On the final line of sight, all the ionized galaxies have a bulk motion towards the observer, which makes the largest ΔT_{kSZ} . This is a very large scale velocity mode.

effects caused by scattering of CMB photons off of free electrons.

The Kinematic Sunyaev-Zel'dovich Effect

In the first, CMB photons scatter off free electrons with a bulk motion relative to the CMB rest frame. This is known as the kinematic Sunyaev-Zel'dovich effect (kSZ) (Sunyaev and Zel'dovich, 1972) This typically happens when CMB photons pass through a galaxy. In the reference frame of that galaxy, the CMB is Doppler shifted, and observed as a dipole. The observed temperature T'_{obs} is

$$T'_{\text{obs}} = T_{\text{CMB}}(1 - v \cos \theta) \quad (1.16)$$

where v is the galaxy's velocity relative to the CMB rest frame (in units where $c = 1$), and θ is the angle of the observation relative to the galaxy's velocity. Since Thomson scattering is isotropic, the scattered radiation has an average temperature of $T'_{\text{scatter}} = \overline{T'_{\text{obs}}}$, where the average is taken over θ . This actually gives the scattered radiation the same temperature as the rest frame CMB (i.e. $T'_{\text{scatter}} = T_{\text{CMB}}$). However, this is in the rest frame of the galaxy.

When we transform back to the CMB rest frame, we find

$$T_{\text{scatter}} = T'_{\text{scatter}}(1 - v \cos \theta) \quad (1.17)$$

$$= T_{\text{CMB}}(1 - v \cos \theta). \quad (1.18)$$

$$(1.19)$$

The fraction of CMB photons that scatter off free electrons in the galaxy is the optical depth τ ,

$$\tau = \int_{\text{LOS}} n_e \sigma_T dl \quad (1.20)$$

where the integral is over an external observer's line of sight, and σ_T is the cross section for Thomson scattering. Now, we can put this all together to express the magnitude of the kSZ effect (due to a single galaxy):

$$\frac{\Delta T_{\text{kSZ}}}{T_{\text{CMB}}} = v \cos \theta \tau \quad (1.21)$$

$$= \sigma_T \int_{\text{LOS}} n_e v_{\parallel} dl \quad (1.22)$$

Two things are immediately notable about this equation. First, the kSZ appears as a change in the CMB temperature, so it is indistinguishable from a true change in the temperature at the surface of last scattering. Second, v_{\parallel} can be either positive or negative. These both have significant impacts on any attempt to measure the total kSZ from many galaxies.

In order to find the total kSZ effect along a line of sight, we simply integrate the effect over all free electrons along a line of sight. Several possible lines of sight are shown in Figure 1.5. In equation 1.22, this simply means allowing n_e and v_{\parallel} to vary with l (the distance along the line of sight). In the late time universe ($z \lesssim 5.5$), we can treat the universe as fully ionized. This is referred to as the homogeneous kSZ, since it is sourced by a homogeneously ionized universe. In that case $n_e \propto 1 + \delta$, where δ is the matter density contrast. This gives an integral with two terms:

$$\frac{\Delta T_{\text{kSZ}}}{T_{\text{CMB}}} \propto \int_{\text{LOS}} v_{\parallel} + \delta v_{\parallel} dl \quad (1.23)$$

If we consider the Fourier transform of v_{\parallel} , it becomes clear that modes shorter than the line of sight cancel in the first term. Those modes have approximately equal positive and negative contributions along the line of sight, and only the extremely slowly varying modes survive. These are firmly in the linear regime, so $\vec{v} \propto \vec{k} \delta(\vec{k})$. The density fluctuations at the scale of the line of sight (which reaches back to $z \approx 5.5$) are negligible, so we can safely ignore the first term. The second term is the correlation of velocity and density modes. This is referred to as the Ostriker-Vishniac effect, and has been explored under various approximations (Ostriker and Vishniac, 1986; Hu and White, 1996; Jaffe and Kamionkowski, 1998). The kSZ effect is undetectable along a single line of sight, with the exception of very extreme objects (Sayers et al., 2013; Adam et al., 2017). Therefore, the primary kSZ observable is the power spectrum. Several possible power spectra for the Ostriker-Vishniac effect are shown in Figure 1.6.




Figure 1.6: Power spectra for the Ostriker-Vishniac effect and primary CMB under different reionization scenarios. As the redshift of reionization (z_{ri}) increases, the optical depth to the CMB increases as well. This decreases the amplitude of the primary CMB, and increases the Ostriker-Vishniac power spectrum. Regardless of the reionization scenario, the Ostriker-Vishniac effect dominates over the primary fluctuations at $\ell \gtrsim 3000$. Figure from Hu and White (1996).

During reionization, the situation becomes even more complicated. We can no longer assume $n_e \propto 1 + \delta$; in principle, we need to know the ionization state along the line of sight. While this can be handled analytically (Hu and White (1996); Jaffe and Kamionkowski (1998) both use analytic visibility functions), it has become common to simulate the dynamics of reionization, and reduce the results to the mm-wave power spectrum (Mesinger et al., 2012; Trac et al., 2022; Gorce et al., 2020). The simulations can account for the various processes that drive reionization in more detail than any completely analytic treatment. For example, Figure 1.7 shows three different reionization scenarios, and their power spectra. The kSZ effect due to reionization is referred to as the patchy kSZ, after the patchy nature of reionization.

1.3.3 Thermal Sunyaev-Zel'dovich Effect

The thermal Sunyaev-Zel'dovich effect (tSZ) is due to inverse Compton scattering of CMB photons off of high-energy electrons. Inverse Compton scattering arises from the interaction of a photon with a relativistic electron. In the rest frame of the electron, the photon is Doppler




Figure 1.7: Three simulations of reionization. White regions are neutral, and black regions are fully ionized. $z = 0$ (today) is on the left, and $z = 20$ is on the right. The middle history varies the ionization efficiency of galaxies over cosmic history, leading to longer reionization. The bottom panel decreases the mean free path of ionizing photons, resulting in both longer reionization, and smaller bubbles. Figure adapted from Mesinger et al. (2012)

shifted to a different energy

$$E'_\gamma = E_\gamma \frac{1 - v \cos \theta}{\sqrt{1 - v^2}} \quad (1.24)$$

where v is the electron velocity, and θ is the angle between the photon's velocity and the electron's velocity in the rest frame of the CMB. If the scattering angle in the electron's rest frame is ϕ' , then the photon's energy after scattering is given by the standard formulation of Compton scattering:

$$E'_{\gamma,f} = E'_\gamma \left(1 + \frac{E'_\gamma}{m_e} (1 - \cos \phi') \right)^{-1} \quad (1.25)$$

Transforming back to the CMB rest frame gives us the final photon energy:

$$E_{\gamma,f} = E'_{\gamma,f} \frac{\sqrt{1 - v^2}}{1 + v \cos \theta_f} \quad (1.26)$$

where $\theta_f \equiv \phi + \theta$ is the angle between the photon's final velocity and the electron's velocity. Let's examine this in the limit where $E_\gamma \ll m_e$. In that case the Compton scattering in the electron's rest frame becomes regular Thomson scattering, and $E'_\gamma \approx E'_{\gamma,f}$. The change in the




Figure 1.8: The spectrum of the tSZ $f(x)$ (top panel) and the apparent change in temperature after passing through a cluster with $y = 10^{-4}$ (bottom). Below $\nu \approx 220$ GHz, the apparent change in temperature is negative. The dashed gray lines show the approximate center of the SPT-3G observing bands (see Chapter 2).

photon's energy only comes from the two Doppler shifts at different angles:

$$E_{\gamma,f} \approx E_{\gamma} \frac{1 - v \cos \theta}{1 + v \cos \theta_f}. \quad (1.27)$$

So, for $\theta = \pi$ (a head on collision), $E_{\gamma,f} \geq E_{\gamma}$. In other words, the photon can gain energy from inverse Compton scattering. Averaging over both θ and θ_f gives

$$\langle E_{\gamma,f} \rangle = \frac{4}{3} \frac{E_{\gamma}}{1 - v^2} \quad (1.28)$$

Since $v \leq 1$, the average inverse Compton scattering event increases the photon energy. In the low-velocity limit ($v^2 \ll 1$), $\langle E_{\gamma,f} \rangle - E_{\gamma} \propto v^2$. So, we can view the kSZ and tSZ as first and second order expansions in v . Usually, we would be less concerned about the v^2 term, however the universe does contain a source of relativistic electrons: galaxy clusters.

Galaxy clusters consist of 10^2 - 10^4 galaxies bound together through mutual gravitational interactions. The space between them (the intracluster medium, or ICM) is filled with a diffuse but extremely hot plasma at $T = 10^6$ - 10^8 K. This gives the electrons typical velocities of $\mathcal{O}(0.01)c$, so these are well into the regime of inverse Compton scattering we have derived. Galaxy clusters are heated by gravitational infall; that is, the gravitational potential their constituents initially have is converted into heat as those constituents fall towards the center

of mass and virialize. In fact, the temperature of the ICM scales with the total mass of a galaxy cluster (Kaiser, 1986; Navarro et al., 1995):

$$T \propto M^{2/3}(1+z). \quad (1.29)$$

The redshift scaling simply accounts for the fact that the mean density scales as a^{-3} , so an otherwise identical galaxy cluster is denser and hotter earlier in the universe.

In order to calculate the magnitude of the tSZ from a cluster, we must average over the electron velocities and directions. This was originally derived in Sunyaev and Zel'dovich (1970, 1972), and has been reviewed by Rephaeli (1995); Birkinshaw (1999); Sunyaev and Zel'dovich (1980). We skip straight to the result:

$$\frac{\Delta T_{\text{tSZ}}}{T_{\text{CMB}}} = f(x) \int n_e \sigma_T \frac{T_e}{m_e} dl \quad (1.30)$$

where $x \equiv \frac{\nu}{T_{\text{CMB}}}$ is the dimensionless frequency, and the $f(x)$ contains all the frequency dependence:

$$f(x) = \left(x \frac{e^x + 1}{e^x - 1} - 4 \right) (1 + \delta_{\text{rc}}) \quad (1.31)$$

where δ_{rc} is a relativistic correction applicable to the most massive clusters, where $T_e \gtrsim 5 \times 10^7$ K (Nozawa et al., 2000). The integral in equation 1.30 is the Comptonization, or Compton y parameter. y depends on sky position through either variations in n_e and T_e within the galaxy cluster, or simply because its depth along the line of sight changes. $Y \equiv \int y d\Omega$ is the integrated Comptonization over an entire galaxy cluster. Y is proportional to $\int T_e n_e d\Omega$, and hence the integrated electron pressure in a galaxy cluster, which is, in turn, proportional to the total thermal energy in the cluster. Given that $T \propto M^{2/3}$,

$$Y \propto M^{5/3}(1+z). \quad (1.32)$$

$f(x)$ is a relatively complicated function, but we know that inverse Compton scattering imparts energy to the CMB photons, so we can sketch out what it should look like. At frequencies above the peak, there are more photons at lower energies. When the photons are shifted to higher energies, the intensity at each frequency is increased (compared to the original blackbody spectrum). At frequencies below the peak of the CMB blackbody, the situation is reversed: there are fewer photons at lower energy. Therefore, the intensity is *lower* after passing through a galaxy cluster. The full calculation of equation 1.31 is shown in Figure 1.8. As expected $f(x) < 0$ at low frequencies, with a null at $\nu \approx 220$ GHz. This is quite an important fact: no other effect increases the intensity of the CMB at some frequencies, and decreases it at others. By combining maps at different frequencies, we can find clusters of galaxies using mm-wave data. Since the tSZ is a spectral distortion, it does not suffer from cosmological dimming. That is to say, ΔT_{tSZ} has no (explicit) dependence on redshift, so two clusters with the same y will have the same tSZ signal, whether they are at $z = 0$ or $z = 2$. See Chapter 5 for an example of a catalog of galaxy clusters generated from mm-wave maps.

Constraints on Cosmology from Galaxy Clusters

However, this begs the question, what use is a catalog of galaxy clusters? This has been covered in multiple reviews (e.g. Carlstrom et al. (2002); Allen et al. (2011)), so we will merely summarize the cosmological utility of galaxy clusters. Galaxy clusters are the most massive gravitationally bound objects in the universe. Therefore, they are highly sensitive to anything that affects the growth of structure on cosmological time scales. Under the Big Bang model, large scale structure is initially sourced by the expansion of quantum fluctuations in the density that are expanded to macroscopic size by inflation. These fluctuations grow over time: the overdensities become more dense as they pull in surrounding matter. Each of these overdensities can be described as a “halo” of matter, with a mass and size. The population of halos can be summarized by the halo mass function $\frac{dN}{dMdz}$. The halo mass function predicts the number of halos in the mass range $(M, M + dm)$ and the redshift range $(z, z + dz)$. Clusters form in the most massive of these halos. The halo mass function is a function of a wide variety of effects, including:

Initial conditions The distribution of the density perturbations from inflation is the starting point for the formation of the halos. The halo mass function depends on the primordial density fluctuations particularly through n_s , the spectral index of the primordial density power spectrum.

Astrophysical processes The baryonic physics occurring in the halos can effect their growth by either injecting heat (through shocks, AGN feedback, star formation, etc.), or remove it (via radiative cooling). This changes the pressure of the baryons, and therefore how the halos grow.

Massive neutrinos When they freeze out of the early universe, neutrinos are initially relativistic. Under these conditions, they cannot be gravitationally captured by the halos, and suppress their growth. However, as they cool, they act like a dark matter component, and enhance the growth of halos.

Galaxy clusters form in the most massive halos, so a catalog of galaxy clusters (with their masses and redshifts) is a measure of the high-mass end of the halo mass function. Therefore, we can use galaxy clusters to help constrain the physics of the growth of structure in the universe.

The constraining power of galaxy clusters is predicated on our ability to measure their mass. The scalings presented here are a significant oversimplification. There are several important assumptions built into $Y \propto M^{5/3}(1+z)$. First, this assumes that the temperature of the ICM is entirely determined by gravitational infall. If astrophysical processes (e.g. AGN feedback, shocks from mergers, etc.) contribute to the ICM temperature, this relationship breaks down. Second, this assumes the only evolution of galaxy clusters with cosmological time is to track the density (through $\rho \propto a^{-3}$). To measure cluster mass from the tSZ, we typically assume a scaling relation of the form

$$T_{\text{tSZ}} \propto M^{B_{\text{SZ}}} E(z)^{C_{\text{SZ}}} \quad (1.33)$$

where $E(z)$ encodes the expansion history of the universe, and B_{SZ} and C_{SZ} are free parameters that determine the scaling between T_{tSZ} and halo mass. There is a significant body of literature devoted to calibrating these scaling relations and extracting cosmological information from cluster catalogs. Various works use simulations (e.g. (da Silva et al., 2004)), gravitational lensing (e.g. (Mantz et al., 2016; Bocquet et al., 2023)), and X-ray observations (e.g. (Benson et al., 2013)).

We can only detect the most massive galaxy clusters through the SZ effect. In fact, such samples are usually mass-limited, since $\Delta T_{SZ} \propto M^{B_{SZ}}$. However, the tSZ signal from all clusters (and any other source of relativistic electrons) is in the maps of the mm-wave sky. We can extract information about these lower-mass halos from the properties of the maps, even if they aren't individually detectable. We can construct a map of the Comptonization (Bleem et al., 2022; Coulton et al., 2024) by combining maps at different frequencies. The same signal can also be extracted from auto- and cross-spectra between maps at different frequencies (Reichardt et al., 2021). Finally, the tSZ injects a significant non-Gaussian component into the maps, which can be measured with higher order statistics (Crawford et al., 2014; Wilson et al., 2012). Measurements of the mm-wave sky give us a window into all of this physics and more.

1.4 Conclusion

The CMB is a powerful probe of cosmology. It provides constraints on the Λ CDM model through the power spectrum. CMB photons have also passed through all of cosmic history since recombination. Even though interactions with matter are rare, we can detect the small changes imprinted on the CMB by intervening matter. However, this requires a highly sensitive camera. The rest of this work describes the South Pole Telescope and the SPT-3G camera, which are used to observe the sky at mm wavelengths. In Chapter 2 we describe the South Pole Telescope and SPT-3G camera. Chapter 3 explains how we calibrate the raw data and turn it into maps of the sky. Chapter 4 presents the measurement of the SPT-3G angular response function (or beam), which is a critical component for understanding the maps. Finally, we present a catalog of galaxy clusters discovered using the SPTPol camera (the predecessor to SPT-3G).

Chapter 2

The South Pole Telescope and SPT-3G Camera

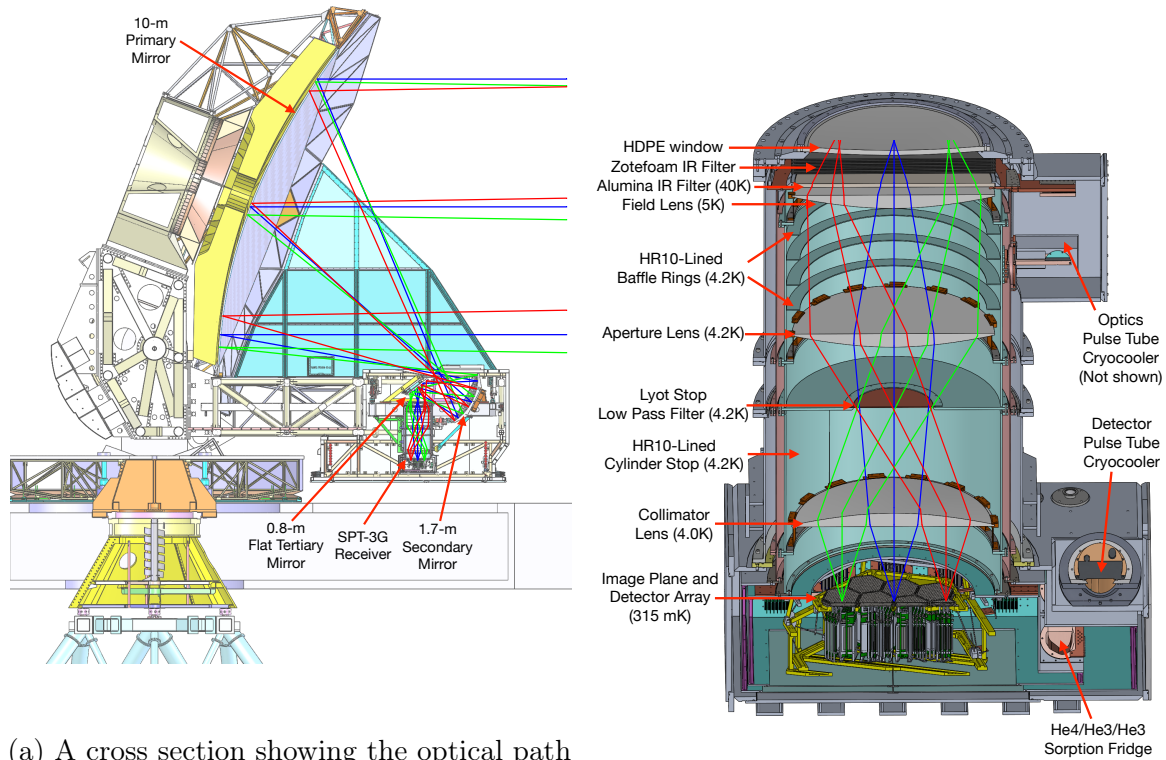
The South Pole Telescope (SPT) is a 10-m off-axis Gregorian telescope located at the Amundsen-Scott South Pole Station (Carlstrom et al., 2011). Since its construction in 2006, the SPT has been home to three mm-wave cameras: SPT-SZ (Chang et al., 2009), SPTPol (Austermann et al., 2012), and, SPT-3G. With a 10-m primary mirror, these cameras have achieved arcminute-scale resolution at 80-280 GHz. The high resolution of the SPT provides access to the small-scale secondary anisotropies and foregrounds that would otherwise be lost to a lower resolution instrument. This chapter describes the SPT-3G camera at a high level. For a more detailed description, see Sobrin et al. (2022) (S22, hereafter).

2.1 SPT-3G Camera Overview

2.1.1 Optics

The SPT has three aluminum mirrors at ambient temperature, as seen in Figure 2.1a. The first optical element is the 10-m diameter primary mirror (shown in light blue). This reflects light into the comoving receiver cabin, which is a heated space that contains the remainder of the optics, the SPT-3G camera, and all the associated electronics. The 1.7-m secondary mirror redirects incident light to a flat tertiary mirror, which sends the light into the camera. The tertiary mirror is only necessary because of space constraints in the receiver cabin.

Figure 2.1b shows a cross section of the camera. The colored lines show rays incident on three different detectors as they pass through the vacuum window, two infrared (IR) filters, the three alumina lenses (beige), the Lyot stop, and the lenslet array directly on top of the detectors. The entire camera is maintained at cryogenic temperatures. This is necessary for the operation of the detectors, and to cool the lenses (which significantly decreases their absorption of mm-wave radiation (Nadolski et al., 2020)).



(a) A cross section showing the optical path through the SPT and into the SPT-3G camera.

(b) A cross section of the SPT-3G camera.

Figure 2.1: A cross section of the SPT and the SPT-3G camera. Panel a shows the optical path of light coming from the right, and incident on the primary mirror, which is then reflected to the secondary and tertiary mirrors before entering the camera on the bottom right. Panel b shows a detailed view of the camera itself, with 3 alumina cold lenses, the Lyot stop, and the detectors. Both figures are from S22.

2.1.2 Detectors

SPT-3G contains ~ 16000 polarization sensitive detectors across three bands, which are centered at 95^1 , 150, and 220 GHz (see Figure 2.2). The detectors are coupled to the sky through broadband sinuous antennae (O’Brien et al., 2010; Suzuki et al., 2014). The sinuous antenna couples mm-wave radiation from the sky (in orthogonal polarizations) onto microstrip transmission lines which carry the signal to a set of three band-defining filters. Each filter allows the frequencies for the relevant band to pass through to a single detector. This yields six detectors per pixel (3 bands in two polarizations). The bandpasses are designed to avoid emission/absorption lines from water vapor and oxygen in the atmosphere (seen in

¹The lowest observing band is centered ~ 95 GHz. However, we refer to it as “90 GHz” throughout this work.




Figure 2.2: The bandpass functions for the three SPT-3G observing bands. The colored lines show the bandpass of each wafer, while the black dashed line shows the average. The gray region shows the typical atmospheric transmission at the South Pole. Figure from S22.

gray in Figure 2.2).

The detectors themselves are transition edge sensors (TESes), a form of bolometer (see Figure 2.3). In general, a bolometer consists of an absorber (the light green rectangle) connected to a thermal bath with a weak thermal link. When the radiation incident on the absorber increases, its temperature increases. In the SPT-3G camera, the sinuous antenna converts incident radiation to a current, which is then thermalized on the absorber through a resistor. The TES is effectively a highly sensitive thermistor, which allows us to measure the temperature of the absorber (and therefore the power coming from the sky). Each TES is made of a superconducting material. The transition from normal metal to superconductor occurs over a very small range in temperature, so a small change in temperature results in a large change in resistance (see Figure 2.4). We apply a voltage bias across the TES, which, when combined with the optical power, heats the absorber enough to keep the TES in the transition between normal metal and superconductor.

When operated under a voltage bias, the TES has a natural negative feedback loop that tends to keep it near the same temperature (Benson, 2022). The electrical power is $P_{\text{elec}} = V^2/R(T)$, where $R(T)$ is the resistance of the TES at temperature T . The change in electrical power with TES temperature is then

$$\frac{dP_{\text{elec}}}{dT} = V^2 \frac{d}{dT} R(T)^{-1} \quad (2.1)$$

$$= \frac{-V^2}{R(T)} \frac{dR}{dT}. \quad (2.2)$$

For a superconductor $\frac{dR}{dT} > 0$ in its transition, so $\frac{dP_{\text{elec}}}{dT} < 0$. In other words, as the temperature increases (due to an increase in the power coming from the sky), the electrical power dissipated in the TES decreases. This works to limit the change in total power on the absorber. We call this effect “electro-thermal feedback”. Its strength is given by the “loop gain” $\mathcal{L} = -\frac{dP_{\text{elec}}}{dP}$,




Figure 2.3: A schematic representation of a bolometer. In the SPT-3G camera, the absorber is a resistor that thermalizes the current from the sinuous antenna, and the thermistor is a TES. Figure from (Benson, 2022).




Figure 2.4: The resistance as a function of temperature of a typical TES. The transition from normal metal to superconductor is very steep, and nonlinear. Figure from (Benson, 2022).

where P is the total power (electrical and optical) on the absorber. We can express P in terms of G , the thermal conductivity of the link connecting the absorber to the thermal bath, and the temperature differential across that link:

$$P = G(T - T_{\text{bath}}). \quad (2.3)$$

Since the thermal bath is assumed to be at a constant temperature, we can calculate the loop gain

$$\mathcal{L} = -\frac{dP_{\text{elec}}}{dP} \quad (2.4)$$

$$= \frac{P_{\text{elec}}}{RG} \frac{dR}{dT} \quad (2.5)$$

$$= \frac{P_{\text{elec}}\alpha}{GT} \quad (2.6)$$

where

$$\alpha \equiv \frac{T}{R} \frac{dR}{dT}. \quad (2.7)$$

Note that $\frac{dR}{dT}$ (and therefore \mathcal{L}) depends on T , since the transition is nonlinear (see Figure 2.4).

We can now calculate the TES sensitivity $S_P = \frac{dP_{\text{elec}}}{dP_{\text{opt}}}$, where P_{opt} is the power from the sky². We want this in terms of loop gain, so we transform variables to R and T .

$$dP_{\text{elec}} = \frac{-P_{\text{elec}}}{R} dR \quad (2.8)$$

and

$$P_{\text{opt}} = G(T - T_{\text{bath}}) - P_{\text{elec}} \quad (2.9)$$

$$dP_{\text{opt}} = GdT + \frac{P_{\text{elec}}}{R} dR. \quad (2.10)$$

So,

$$S_P = \frac{-\frac{P_{\text{elec}}}{R} dR}{GdT + \frac{P_{\text{elec}}}{R} dR} \quad (2.11)$$

$$= \frac{-\frac{P_{\text{elec}}}{R} \frac{dR}{dT}}{G + \frac{P_{\text{elec}}}{R} \frac{dR}{dT}} \quad (2.12)$$

$$= \frac{-\mathcal{L}}{1 + \mathcal{L}}. \quad (2.13)$$

²TES sensitivity is usually defined in terms of the current $S \equiv \frac{dI}{dP_{\text{opt}}}$. We use S_P (in terms of the electrical power) here, because our TESes are biased with a constant voltage, and our calibration pipeline (discussed in Chapter 3) is also based on P_{elec} . In the case of a constant bias voltage V_{bias} , $S = S_P V_{\text{bias}}$.

In the limit of high loop gain ($\mathcal{L} \gg 1$), $S_P = -1$, and the electrical power exactly trades off with optical power, keeping the total power on the absorber constant. Hence, the TES is held at the same temperature.

However, the achieved loop gain of the SPT-3G detectors is $\sim 5 - 10$, and $S_P > -1$. So, as P_{opt} varies, P_{elec} changes by a smaller amount, and the total power varies. This, in turn, varies the temperature of the TES. As long as the temperature stays in a region where $\frac{dR}{dT}$ is approximately constant, the sensitivity (and hence calibration) remains approximately constant. However, if P_{opt} varies enough, the overall calibration of the detector changes. The required change in P_{opt} to make a significant change in S_P is much larger than the variations in the CMB. But, the atmospheric load does vary enough with observing elevation to make this non-linearity a significant effect. We deal with this by re-calibrating the detectors at different elevations, and limiting the elevation range of a single observation (see Chapter 3, and §2.3). Extremely bright astrophysical sources can also cause non-linearity (or even saturation, when the TES temperature rises so much that it leaves the superconducting transition entirely).

2.1.3 Cryogenic Design

Maintaining the operating temperatures required by the detectors demands a significant cryogenic system. The detectors must be held under vacuum, and shielded from blackbody radiation at higher temperatures. This is accomplished with a multi-stage cooling system. The first stage is the vacuum shell (shown in gray in Figure 2.1b). It sits at room temperature, and is only responsible for maintaining a vacuum around the cooled portions of the camera. While most of it is made of aluminum, the vacuum window (at the top of Figure 2.1b) is made out of high density polyethylene (HDPE) to allow light to reach the detectors. The vacuum window is discussed in detail in §2.2. The next two cooling stages are at ~ 30 K and ~ 4 K (shown in brown and turquoise, respectively). The cooling for these stages is provided by two PT415 pulse tube coolers (PTCs). These are closed-cycle helium coolers, and provide ~ 30 W and ~ 1 W of cooling power, respectively. The warmer stage is only used to block radiative power from the ~ 300 K vacuum shell. The 4 K stage both provides the baseline cooling for the detectors, and cools the lenses.

Focal Plane

The detectors are mounted on a four-stage focal plane, shown in Figure 2.5. Starting with the “ultra cold” (UC) stage, each stage is progressively hotter, and has more cooling power. The UC stage operates at 315 mK, with only a few μW of cooling power; the “intercold” (IC) stage operates at 500 mK, and provides ~ 30 μW of cooling power; the “buffer cooler” (1K) stage is at 1 K, and can provide ~ 100 μW of cooling power; the final stage on the focal plane is the 4 K stage, and is connected to the PTC. The stages are mechanically separated by a carbon fiber truss structure. Carbon fiber provides both high mechanical strength (the UC stage weighs 22 kg), and low thermal conductivity (Runyan and Jones, 2008).




Figure 2.5: The SPT-3G focal plane. The detectors are housed in the hexagonal structures in the center. The four temperature stages are visible (4 K, 1 K, IC, and UC). Figure from S22.

The three coldest stages (UC, IC and 1K) are cooled by a helium-filled sorption refrigerator from Chase Research Cryogenics (Bhatia et al., 2000). This refrigerator operates by vacuum pumping (using an activated charcoal adsorber) on a bath of liquid helium (the IC and UC stages use He^3 , while the 1K stage uses regular He^4). Once the helium has completely boiled off, the refrigerator is no longer able to maintain its temperature. We recover the helium by heating the adsorbers. The He^4 in the 1K stage is condensed using the cooling power from the 4 K stage. We then begin pumping on the He^4 to condense the He^3 , which has a boiling point below 4 K. We refer to this process (which takes about 4 hours) as “cycling” or “recycling” the refrigerator. The rate at which the helium boils depends on the heat load on each stage, so decreasing the heat load leads to longer times between cycles. The total load on the UC stage is $4.0 \mu\text{W}$, which allows us ~ 16 hours between cycles.

2.1.4 Readout

To minimize the thermal load on the focal plane from wiring, we use a $68\times$ frequency-multiplexed biasing and readout scheme (Bender et al., 2014, 2016). Under this scheme, each detector is biased with an AC voltage at a different frequency. The detectors are connected in series with a LC filter, that selects only one of the AC voltages. This all occurs on the UC stage, so we are able to reduce the number of conductors between temperature stages by $68\times$. Without a multiplexing system, each detector would require one pair of wires, which would conduct heat between all the different temperature stages.

P_{elec} is typically $\mathcal{O}(1)$ pW, and the TES resistances are $\mathcal{O}(1)$ Ohm, so the typical current in the detectors is $\mathcal{O}(1)$ μA . To measure such a small current, we first amplify it with a SQUID series array (Irwin and Huber, 2001). The SQUIDs are mounted at the 4 K stage of the camera. The amplified current from the SQUIDs is digitized at room temperature with dedicated single-board computers. The digitized signals from each board are sent over a




Figure 2.6: A schematic of the vacuum window and IR filtering portions of the camera.

network to the control computer, which merges the data from all the boards, and writes it to disk. Once an observation is complete, the data from the detectors is merged with the remaining metadata (telescope position, cryogenic temperatures, etc.). This data is both stored locally at the South Pole, and transferred back to North America using a satellite internet link. There is a 24-48 hour delay between initially recording the raw detector data and receiving the full observation in North America.

2.2 Vacuum Window and IR Filtering

The vacuum window and IR filtering are at the nexus of the cryogenic, mechanical, and optical design of the camera. The vacuum window must maintain the integrity of the vacuum shell, and allow mm-wave photons to pass through. The IR filter stack prevents room temperature radiation from heating the colder stages. Without filtering the IR radiation, we would not have enough cooling power to keep the detectors in their operating range. A schematic of this portion of the camera is shown in Figure 2.6. The vacuum window is shown in its approximate deformed shape at the top. The IR filtering occurs in two steps. The first is through a series of thin Zotefoam sheets; the second is via a thick alumina plate, which is connected to the first cold stage.

2.2.1 Vacuum Window

The vacuum window has two primary design constraints. It must have a 680 mm clear aperture to allow the detectors a clear view out to the mirrors. This requires that the window support a load of ~ 8000 pounds of atmospheric pressure. It must also allow as much mm-wave radiation to pass through as possible. There was not a formal specification, but we aimed to allow $\gtrsim 90\%$ transmission. This target ensures that the vacuum window contributes a small fraction of the total loss in the optical path (see Table 2 in S22 for a full accounting of the losses in the optical path). These constraints make HDPE a natural choice: it is mechanically strong, and has low absorption at mm wavelengths (Lamb, 1996).





Figure 2.7: One antireflection surface cut into a HDPE test piece.



(a) The theoretical transmission spectra for one through both surfaces of a 30 mm thick window.

(b) The measured transmission spectrum through two surfaces, and 30 mm of HDPE.

Figure 2.8: The theoretical transmission for one grooved antireflection surface (panel a), and the measured transmission through two surfaces, and 30 mm of HDPE (panel b). Note that the theoretical transmission spectra indicate different transmission depending on the polarization state of the incoming radiation. The grooves on the top and bottom surfaces of our vacuum window are orthogonal, so this differential transmission cancels after passing through both surfaces. The measured spectra were taken at four different locations on the window. The dashed lines in panel a show the centers of the SPT-3G bands.

This leaves us with the problem of reflections at either surface, due to the sudden change in index of refraction. A common solution to this problem is the use of a “metamaterial” antireflection coating. These are features at scales smaller than the minimum observed wavelength that are either added or machined into the the surface of the optical element. Since they are smaller than the wavelength, incoming radiation “sees” only the mean density of the material. If the fill factor of the material is smoothly varied along the optical path, the incoming radiation experiences a smooth transition of the index of refraction. Typically, this manifests as a triangular, conical, or pyramidal shape, depending on the material, and the geometry of the optical element. Metamaterial anti-reflection coatings or layers have seen use in many mm-wave cameras (e.g. Datta et al. (2013); Matsumura et al. (2016); Takaku et al. (2020); Golec et al. (2020, 2022); Beckman (2023)).

We experimented with several ways of machining a metamaterial antireflection layer into a HDPE sheet. Any method involving a traditional spinning cutting tool (e.g. an end mill) resulted in poor definition of the features, and truncated the points of both triangular grooves and pyramids. Eventually, we developed a method for cutting linear triangular grooves with the UC Berkeley Physics Department machine shop. Gordon Long proposed and built a non-spinning cutter made from tool steel. The cutter matches the profile of the desired groove, and is sharpened on the edges so that it acts like a knife when pulled through the HDPE. To cut many parallel grooves, we mounted the HDPE in a CNC machine, and put the cutting tool in the spindle. We then use the CNC machine to pull the cutting tool through the surface of the HDPE, then step over to the next groove, and repeat this process over the entire vacuum window surface. This takes $\sim 12 - 15$ hours per side. Figure 2.7 shows a picture of a test piece with grooves cut into it.

Raguin and Morris (1993) provides a framework for calculating the transmission of a repeated groove. They provide analytic equations that drive the period and height of our grooves. The period is chosen eliminate diffraction spikes at the highest observed frequency (~ 280 GHz); the height is chosen to have low reflectance all the way to the minimum observed frequency (~ 60 GHz). Intuitively, the constraint on the period ensures that we are in the limit that the features are smaller than the minimum observed wavelength, while the constraint on the height ensures that the transition from one index to the other is longer than one wavelength. Using the results in Raguin and Morris (1993), we calculate the transmission spectra for one surface shown in Figure 2.8a, which far exceeds our 90% goal. We negate any birefringence or polarization-dependent transmission by cutting the grooves orthogonally on each side (Schroder et al., 2016). We measured the transmission spectrum of the window with a Fourier transform spectrometer, and found that transmission was $\gtrsim 95\%$ across all three bands. The measured transmission spectrum is shown in Figure 2.8b

The vacuum window is made from a 30 mm (not including the AR grooves) thick sheet of annealed HDPE. The annealing process relieves internal stresses, and results in a stronger sheet at a given thickness. We installed the window in a test fixture and held it under vacuum for several weeks to ensure it, and the supporting structures, were strong enough to withstand the atmospheric pressure load. Over the first three days, the window deformed by 30 mm. After that, there was no measurable increase in deformation. Because of the difference in




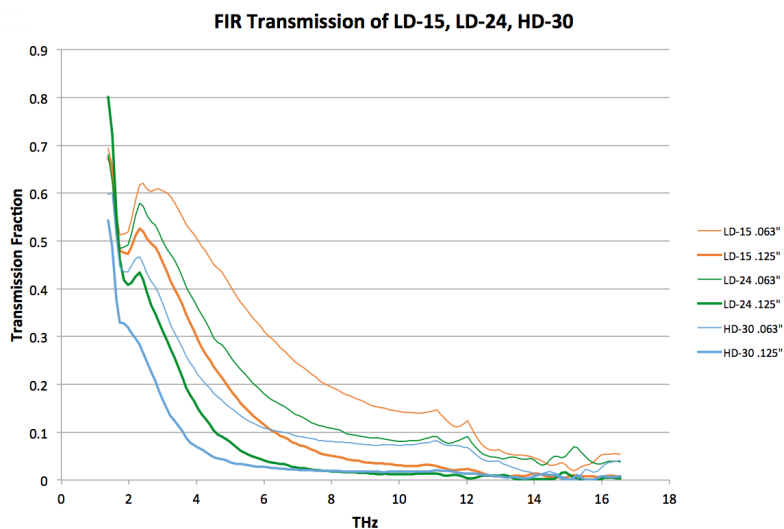
Figure 2.9: A schematic representation of our model for the multi-layer IR filter.

atmospheric pressure between Berkeley and the South Pole, this test exceeds the required strength of the window by $\sim 30\%$.

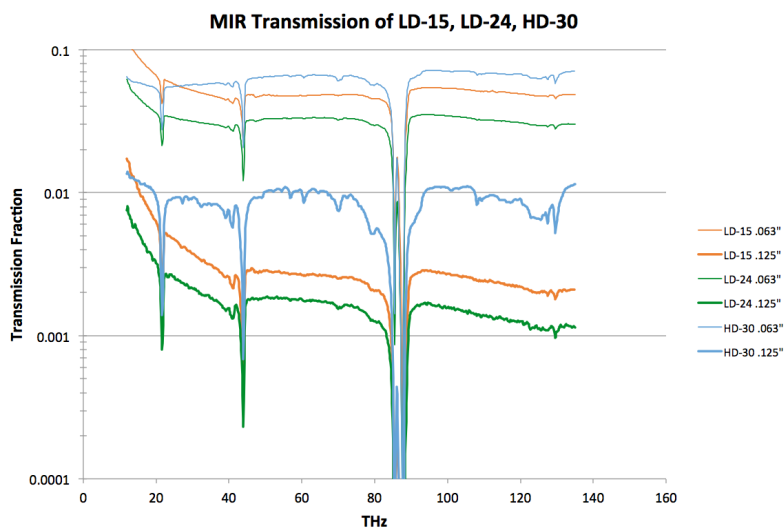
2.2.2 IR Filtering

As discussed above, the PTC coolers provide ~ 30 W of cooling power on the warmest stage, and only 1 W on the 4 K stage. Assuming the window is, or looks out onto, a 273 K blackbody (which is a good approximation above a few hundred GHz), the total radiative power coming through the window is > 100 W. Therefore, we need to block a significant amount of that radiation before it hits the first cold stage. We implemented a very simple filter using 8 thin layers of expanded HDPE foam (specifically, HD30 from Zotefoam). HD30 has a long history in mm-wave observations. It has been used as a vacuum window in the ACBAR, APEX-SZ, SPT-SZ, SPTPol, the BICEP family of telescopes, and POLARBEAR family of telescopes (Runyan et al., 2003; Schwan et al., 2011; Carlstrom et al., 2011; Austermann et al., 2012; Keating et al., 2003; Nguyen et al., 2008; Tomaru et al., 2012; Suzuki et al., 2016).

This type of multi-layer IR filter was initially investigated by Choi et al. (2013) for radio telescopes. A similar filter was implemented in the POLARBEAR-2 experiment in 2015, using polystyrene foam (Inoue et al., 2016). The POLARBEAR-2 implementation has the multi-layer insulation mounted on the first cold stage. The version described in this work (using HD30, and mounted to the vacuum shell) and was simultaneously implemented by BICEP3 (Kang et al., 2018) and SPT-3G in 2017.



(a) Transmission spectra of several Zotefoam products in the far IR range



(b) Transmission spectra of several Zotefoam products in the mid-IR range

Figure 2.10: The transmission spectra of several different polyethylene foams at two thicknesses in the far (panel a), and mid-IR (panel b). HD foams are made of high density polyethylene, while LD foams are made of low density polyethylene. The numbers after the dash indicate the mean density of the foam, in kg / m^3 , and the final number indicates the thickness of the foam layer, in inches. Measurements and plots by Ki Won Yoon (Yoon, 2016).

HD30 is approximately transparent at mm wavelengths, but opaque in the IR range. This is a function of the cell size and density of the foam. HD30 has a density of 30 kg/m^3 , whereas solid HDPE is $\sim 1000 \text{ kg/m}^3$. The cells (empty spaces between solid HDPE) are $\sim 0.4 \text{ mm}$ in HD30. Therefore, at mm wavelengths, the radiation sees only the mean density of HDPE, which is extremely small. However, at wavelengths smaller than the cell size, the radiation sees the edge of each cell as a rapid change in the index of refraction. Therefore, it reflects off of each cell, and has many scatterings inside the foam. This brings the radiation into equilibrium with the foam temperature, which makes the foam an IR absorber. Transmission spectra for several different polyethylene foams are shown in Figure 2.10. $\sim 98\%$ of the power from a 300 K blackbody is at frequencies greater than 5 THz, so any of these foams would block the majority of the radiative power coming through the window.

Modeling this multi-layer insulation is relatively easy. We assume each layer is a perfect absorber, and that it is very thin (so that the surface area is dominated by the top and bottom surfaces). We define P_i as the power emitted by one surface of the i th layer (so that the total power emitted by the i th layer is $2P_i$). The incident power from the vacuum window is P_0 , and the final power emitted by the last layer of stack of n layers is P_n . This is shown graphically in Figure 2.9. If the layers are all in equilibrium, then the total power each layer emits is equal to the power it absorbs from the layer above it and below it:

$$2P_i = P_{i-1} + P_{i+1}. \quad (2.14)$$

This is a recurrence relation with characteristic equation

$$x^2 - 2x + 1 = 0. \quad (2.15)$$

So, the solution is of the form

$$P_i = \alpha + \beta i. \quad (2.16)$$

In order to solve this, we make one final assumption: we assume that $P_{n+1} = 0$. In other words, the power emitted by the alumina backing plate onto the final layer in the stack is negligible. Since the alumina backing plate is actively cooled, and sits at $\sim 40 \text{ K}$, this should be a good assumption. Under this assumption, $\alpha = -\beta(n+1)$. This allows us to find the fraction of the incident power (P_0) that passes through the filter stack:

$$\frac{P_n}{P_0} = \frac{\alpha + \beta n}{\alpha} \quad (2.17)$$

$$= \frac{-\beta(n+1) + \beta n}{-\beta(n+1)} \quad (2.18)$$

$$= \frac{1}{n+1}. \quad (2.19)$$

Assuming the incident power through the window comes from a 273 K blackbody, $P_0 = A\sigma T^4 \approx 115 \text{ W}$, where A is the window area, σ is the Stefan-Boltzmann constant, and T is the ambient temperature outside the camera. We estimate that the total power on the

alumina backing plate is $P_8 = 8.8$ W. This gives $P_8/P_0 \approx 0.077$, which is slightly better than our expectation of 0.11. The total load on the PTC is 33.4 W, with the window loading second to radiative load from the rest of the vacuum shell (at 18.1 W). These loads are estimated from load curves measured in the laboratory at each stage.

We can now quantitatively evaluate the assumption that $P_9 \ll P_8$ (the power emitted by the alumina backing plate is much smaller than the power emitted by the final Zotefoam filter). The window backing plate is 720 mm in diameter and approximately 40 K, so $P_9 \approx 0.06$ W, which is much less than the 8 W emitted by the final Zotefoam layer.

The alumina backing plate operates in a similar mode: it is highly transmissive to mm-wave radiation, but effectively opaque to IR. However, it is actively cooled, so there is no need to add multiple layers (as long as the first cooling stage can absorb the power coming from the Zotefoam filter stack). This also assumes that the thermal conductivity of the alumina filter is sufficient to remove the heat from the center of the filter without allowing its temperature to rise. To ensure this, we use a 7.5 mm thick alumina plate. Alumina is extremely difficult to machine, so instead of cutting grooves into the backing plate, we add layers of PTFE-based dielectrics to smoothly transition the index of refraction (Nadolski et al., 2020).

With these filters in place, we estimate the load on the first lens (at ~ 4 K) is approximately 60 mW (of 800 mW total on the entire 4K stage). The radiative load from the lenses onto the UC stage is 0.16 μ W (compared to a total of 4.0 μ W). The radiative load on the 1K and IC stages is negligible, since their area is very small.

2.3 Scan Strategy

The SPT operates in a continuous scanning mode (as opposed to a “stop-and-stare” mode). Observations are a sequence of scan pairs. In each scan, the telescope moves across the sky at a constant speed in azimuth, while maintaining a constant elevation. Since the SPT is at the geographic South Pole, elevation (in ground coordinates) is almost exactly the opposite of declination (in sky coordinates), so the scans are at nearly constant declination. Each pair of scans is a right-going (increasing in azimuth) and a left-going scan³. After the left-going scan, we increase the elevation by a small amount (0'5 - 15'0). Over the course of an observation (typically $\lesssim 2$ hours), the telescope tracks the rotation of the sky, so that the scans all cover the same range in Right Ascension (R.A.). We only use the portion of the scan where the telescope is moving at a constant velocity in azimuth, and is fixed in elevation. We discard the small sections where the telescope is accelerating (either to turn around, or stepping in elevation).

We primarily observe a ~ 1500 deg² region of the sky given by $20\text{h}40\text{m} < \text{R.A.} < 3\text{h}20\text{m}$ and $-41 > \text{decl.} > -71$. We observe this field in 4 subfields, each of which spans the full R.A.

³Left-going is decreasing in azimuth. Due to the inversion of the local coordinates relative to sky coordinates, left-going scans increase in azimuth. Therefore, left-going scans go from left-to-right in standard map projections.

range, but only 7.5 deg in declination. With the limited declination range, the atmospheric loading does not vary enough between the top and the bottom of the field to cause significant non-linearity in the detectors. The subfields are named by the positions of their centers, so we have ra0hdec-44.75, ra0hdec-52.25, ra0hdec-59.75, and ra0hdec-67.25.

2.4 Conclusion

The SPT-3G camera consists of ~ 16000 polarization sensitive detectors in 3 wide bands, configured around the peak of the CMB blackbody and the maximum transmission of the atmosphere. It is currently in its seventh season of observations, with the a primary science field of 1500 deg^2 . The telescope optics provide arcminute-scale resolution, which allows us to observe the The detectors are TESes, which are read out with a frequency-domain multiplexing system. The raw data consists of the current running through the detectors, as well as the direction the telescope is pointing. We now need to calibrate the data into units referenced to the CMB blackbody, and convert the 1d time-ordered data into a 2d map of the sky. This is the subject of the next chapter.

Chapter 3

Calibration and Mapmaking

In order to extract cosmological information from our data, we must convert it into an image of the sky (a “map”). To do so, we first need to calibrate the raw detector time ordered data (TOD). We reference each detector to a source on the sky of known brightness, to measure the detector’s instantaneous sensitivity. From there, we use an internal calibration source to track any changes in the detector’s sensitivity. Along with the detector TOD, we have TOD for each detector’s pointing (the position on the sky that each detector observes). We use the pointing to determine the location of each TOD sample in the map. This chapter describes the calibration pipeline and mapmaking process. Section §3.1 describes the calibration pipeline in detail. In section §3.2, we describe how we determine where each detector is pointing, and in section §3.3 we briefly describe the mapmaking process.

3.1 Calibration

We make maps in units of the implied change in the CMB temperature at a given location on the sky. This is a first order expansion of the CMB blackbody:

$$\Delta I(\nu, \vec{\theta}) = B(\nu, T_{\text{CMB}}) - B(\nu, T(\vec{\theta})) \approx \Delta T_{\text{CMB}}(\vec{\theta}) \left. \frac{dB}{dT} \right|_{T=T_{\text{CMB}}} \quad (3.1)$$

where $\vec{\theta}$ is the position on the sky, $B(\nu, T)$ is the blackbody spectrum at temperature T , $T_{\text{CMB}} \approx 2.726$ K is the CMB monopole temperature today (Fixsen, 2009), and ΔT_{CMB} is the implied change from the monopole temperature that we store in our maps. By definition, CMB anisotropies have the same magnitude at all frequencies.

For each detector, we record the constant bias voltage and the resultant current (hence we can trivially convert to electrical power). As shown in §2.1.2, the electrical power is proportional to the optical power from the sky. So, the goal of detector calibration is to find the multiplicative factor that converts its TOD from electrical power to ΔT_{CMB} . Under ideal conditions, we could use the sensitivity $S_P = \frac{-\mathcal{L}}{1+\mathcal{L}}$, along with the measured bandpasses and

the blackbody function to directly calculate the conversion factor. However, there are several factors we do not know well enough to directly calculate the conversion to ΔT_{CMB} :

Loop gain The loop gain (see §2.1.2) is a function of the temperature of the TES, which can vary for several reasons. When we recycle the Chase refrigerator, the entire focal plane rises above the superconducting transition temperature for the TESes. When we re-cool the focal plane, we also re-tune the bias voltage on each TES. Day-to-day variations in loading and bath temperature at the time of tuning can lead to different bias voltages and TES temperatures. Due to the finite loop gain of the detectors, their calibration is also sensitive to changes in loading after tuning. We expect changes in the optical loading from two sources: variations in weather conditions may change the column depth of water vapor in the sky above the telescope, which increases the loading (especially at 150 and 220 GHz); in addition, changing the elevation at which we are observing modulates the load by a small amount. The latter effect comes from changing the line-of-sight distance through the atmosphere. It is minimized at zenith, and maximized at the equator.

Atmospheric Opacity In addition to changing the loading on the detectors, variable weather conditions also change the overall transmission of the atmosphere at mm wavelengths.

Optical Efficiency The total optical efficiency of the telescope optics and detectors is not known *a priori*. The optical efficiency is defined as the ratio of the power incident on the aperture to the power deposited on the detectors (P_{opt}). A variety of factors can decrease the optical efficiency, such as absorption in the lenses and vacuum window, the stop efficiency, and the quantum efficiency of the detectors themselves. While many of these quantities can be measured in the lab, we choose to measure them *in situ*.

We do not attempt to measure each of these components separately, and instead determine the calibration using a mix of bright sources on the sky and an internal calibration source. The calibration factors are re-measured for every science observation, to track the time-dependent changes.

3.1.1 Calibration Observations

This starts with an observation of a bright, steady source on the sky. In practice, we primarily use two HII regions, MAT 5a and RCW 38 (see Figure 3.1). These are large star forming regions in our galaxy, so they are very bright. Their emitting region is also very large, which makes it impossible for their brightness to vary significantly over the lifetime of the SPT-3G camera¹. We describe how these maps are made in §3.3.5.

¹Compare this to AGN, which are also very bright (although much farther away), but can change their emission by 10s of percent on week-to-month timescales. Basic causality arguments require that the emission regions be $\lesssim \mathcal{O}(10)$ light-days. For reference, the emitting region of RCW 38 is > 1 lightyear.




Figure 3.1: Maps of RCW 38 (panel a) and MAT 5a (panel b). The maps are shown with a logarithmic color axis to show the extended emission. These maps are in “source relative” units (see §3.1.1), so each map peaks at approximately 1.0. The dark strips on either sides of bright features are “filtering wings” from the mapmaking process (see §3.3).

We do long observations (~ 2 hours) of these HII regions roughly once per week. We also do much shorter observations² (~ 10 minutes) before every observation of the science fields. The final observation in the calibration chain is of an internal calibration source mounted on the secondary mirror, which we refer to as the “calibrator.”

The calibrator is used to track the detector sensitivity S_P (see equation 2.13), which we refer to as the detector “gain.” It chops between a ~ 1000 K blackbody and a blackbody at ambient temperature (~ 280 K). Radiation from either source is coupled to a light pipe that leads to a 0.5-inch aperture in the center of the secondary mirror. This injects a square

²Internally, we refer to the long HII region observations as “fast points” or “pixelrasters.” The latter is intended to indicate that every pixel in the focal plane makes a complete map of the target source. We call the short HII region observations “very fast points.”

wave signal into the detectors, at a commandable frequency. Since the detector beams are not centered on the secondary mirror, we do not expect every detector to receive the same signal amplitude from the calibrator (we call this variation the “illumination pattern” of the calibrator, even though it is a property of the telescope optics, not the calibrator). Therefore it is only used for relative calibration.

The sequence of calibration observations looks like this:

1. Once per week:
 - a) Observe the calibrator
 - b) Long observation of a HII region
2. Once per science observation:
 - a) Observe the calibrator
 - b) Short observation of the same HII region
 - c) Slew the telescope to the science field
 - d) Observe the calibrator

Step 1b allows us to calibrate every detector to the known brightness of the HII region. We use the previous calibrator observation (1a) to measure the illumination pattern of the calibrator on the focal plane. Then, for every observation of the science fields (which may be days later, and after multiple retunings of the detectors), we once again observe the calibrator. Since we now know the illumination pattern, this allows us to remove per-detector calibration variations. Next, we perform a short observation of the same HII region. This allows us to account for changes in the atmospheric opacity since the the long observation, or a change in the temperature of the calibrator thermal source. This correction is applied to all detectors in a band. Finally, we point the telescope to the center of the science field, and do one final calibrator observation. This allows us to compensate for any calibration changes due to the change in atmospheric loading between the elevation of the HII region and the science field.

Calibrator Observations

For each calibrator observation, we point the telescope at an empty part of the sky (i.e. without a detectable source in the field of view). We then chop the thermal source at 4 Hz for 1 minute. This injects a 4 Hz square wave into the detectors (see Figure 3.2). The chopper itself outputs a synchronization signal, indicating when the source is exposed to the detectors, and when it is covered. We analyze the detector TOD to extract the strength of the response each detector has to the calibrator signal.

We first subtract a 4th order polynomial from each detector’s TOD. This gets rid of slow drifts (most likely from atmospheric water vapor fluctuations passing through the field of view in the wind), and removes the mean. The means of the detector TOD are purely a




Figure 3.2: TOD from one detector while observing the calibrator. This detector measures the calibrator the signal with $\text{SNR} = 83$ over the full 1-minute observation.

property of the tuning, and contain no further information. We then apply a software lock-in amplifier, using the sync signal, and record the mean response to each detector in electrical power. To estimate the noise in each detector, we break the 1-minute-long TOD into six chunks, and measure the variation in calibrator response. We store both the mean response and the mean response divided by the variation in the six chunks (i.e. SNR).

The calibrator’s non-uniform illumination pattern and the slow variation³ of the source temperature mean that the detectors’ response to the calibrator in any single observation is meaningless. Fortunately, the illumination pattern is determined by the telescope optics, which are fixed, so we can use the change in response to track relative changes between the detector gains.

We use the change in the calibrator response between short and long observations of the HII regions to track relative changes in the detector gain. Any common-mode change, or variation in the temperature of the thermal source, is accounted for by tracking the apparent change in brightness of the HII region. We also use the change in calibrator response between the short HII region observation and the following observation of the science field. These calibrator observations are separated by only ~ 15 minutes, so we can use them to track relative and absolute changes in gain.

Long HII Region Observations

The long observations are designed such that each detector makes a complete map of the source. Figure 3.3 shows a typical single-detector map of both MAT 5a and RCW 38 at 90 GHz. These maps will have a variety of peak amplitudes and sub-pixel offsets (due to the limited coverage in the vertical direction, see §3.2). The HII regions are a few times larger

³The temperature of the thermal source varies by $\lesssim 1\%$ on hour timescales. This means it is only useful for tracking a focal plane wide variation in calibration factors on scales of a few hours, but less than half a day.

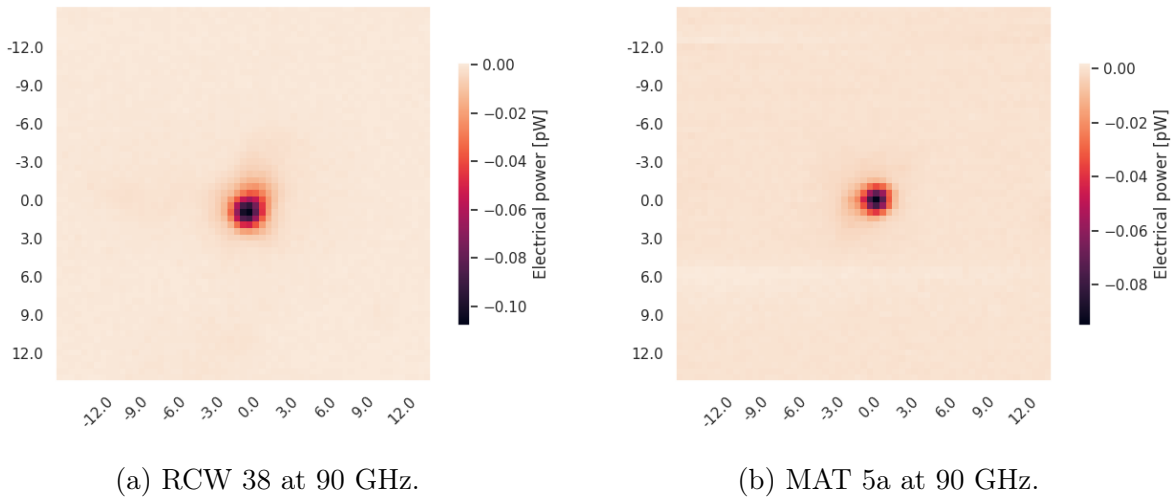


Figure 3.3: Typical maps of the HII regions RCW 38 (panel a) and MAT 5a (panel b) at 90 GHz. These maps are made using the TOD from a single detector. We show these maps on a linear scale to emphasize the very high SNR we achieve, even with a single detector.

than the beam, and have a complicated internal structure. Therefore, we can't fit them to a simple analytic function, and adopt a template-based iterative fitter.

We start by making an initial guess at the calibration factor for each detector, and its position. This only needs to be somewhat close, so we take the position of the source as the location of the brightest pixel. For the calibration, we use the inverse of the maximum pixel value in a $4' \times 4'$ box centered on the nominal source position. We then rank the detectors by their SNR in the previous calibration observation, and use the top 60 detectors to form a template.

To make the template, we first multiply each map by our initial guess at the calibration. This will set the brightest pixel in each map to 1 (in arbitrary power units). We then interpolate each map so that the brightest pixel is at the center of the map, and coadd the 60 we have selected. Next, we fit the original single-detector maps to the template, with three free parameters: the amplitude, and two position offsets. These fit values then begin the next iteration: we rebuild the template with the updated values, and re-fit them. We run a total of three iterations (after three iterations, the changes are no longer meaningful).

This produces a template with arbitrary normalization (although the brightest pixel in the template is expected to be near 1.0), and an amplitude for each detector that makes its map match the template. The best-fit amplitudes are the calibration factors for each detector to an arbitrary unit, defined by the amplitude of the template. We also calculate the integrated flux in a $4' \times 4'$ box at the center of the final template map. These values are stored for later use.

Short HII Region Observations

In the short observations, no single detector makes a complete map of the source. In fact, not every detector even sees the source. We observe the source with three narrow strips (in azimuth) across the focal plane. This is sufficient to make a high SNR measurement of the source when all the detectors in a band are coadded into a single map.

We make these maps in “source relative” units. This is a unit system that matches the arbitrary units of the template in the previous long HII region observation. To construct the source relative calibration, we first multiply each detector TOD by the ratio of its calibrator response before this observation, and before the previous long HII region observation. This accounts for changes in the relative calibration between detectors. Then we multiply each TOD by the best-fit amplitude calculated from the long HII region observation. This applies the original calibration derived from the long HII region.

We use the maps to measure the integrated flux in the same $4' \times 4'$ box centered on the source. If the calibration has not changed since the last long HII region observation, the integrated flux in this map will be the same as the integrated flux measured in the long HII region observation. Any difference is indicative of a change to the calibration.

3.1.2 Final Calibration

The calibrator observations, long HII region observations, and short HII region observations give us enough information to assemble the final calibration for the science observations. Let us first define some terms:

$$F_{\text{long},\nu} \quad F_{\text{short},\nu}$$

The integrated flux of a HII region from a long or short observation, per band. This has dimensions of electrical power \times solid angle.

$$A_{\text{HII},i,\nu}$$

The best-fit amplitude for the i th bolometer in band ν from the long HII region observation, per detector. This quantity is dimensionless.

$$R_{\text{cal,long},i,\nu} \quad R_{\text{cal,short},i,\nu} \quad R_{\text{cal,field},i,\nu}$$

The calibrator response of the i th bolometer in band ν , before a long, or short HII region observation, or before an observation of the science field. $R_{\text{cal},i,\nu}$ has dimensions of electrical power.

$$F_{\text{HII},\nu,\text{true}}$$

The known integrated flux of the HII region in band ν . These are measured from SPT-SZ maps that were calibrated to *Planck*. $F_{\text{HII},\nu,\text{true}}$ has dimensions of $\Delta T_{\text{CMB}} \times$ solid angle.

Finally, we introduce one derived quantity: the illumination pattern of the calibrator over the focal plane.

$$I_{\text{cal},i,\nu} \equiv \frac{R_{\text{cal, long}, i, \nu}}{A_{\text{HII}, i, \nu}} \quad (3.2)$$

Under these definitions, the conversion factor applied to get the source relative units for the short HII region observation is $I_{\text{cal},i,\nu}/R_{\text{cal, short}, i, \nu}$. The estimated integrated flux in each single detector map of the HII region is $A_{\text{HII}, i, \nu} F_{\text{long}, \nu}$. Combining this with $F_{\text{HII}, \nu, \text{true}}$ gives the conversion factor for that detector from electrical power to ΔT_{CMB} . However, that is the instantaneous conversion, which is only valid for the time and elevation of the long HII region observation. We need to use the remaining values to build up the more robust calibration chain.

We account for changes in detector gains by dividing $I_{\text{cal},i,\nu}$ by the calibrator response (i.e. $I_{\text{cal},i,\nu}/R_{\text{cal, field}, i, \nu}$). We also need to account for the variation in atmospheric opacity. For this, we simply use the ratio of HII region integrated fluxes between the short and long observations. And finally, we need to convert this quantity to ΔT_{CMB} . Everything so far is referenced back to $F_{\text{long}, \nu}$, so we use the ratio of $F_{\text{HII}, \nu, \text{true}}$ and $F_{\text{long}, \nu}$. This gives us the final calibration factor in dimensions of ΔT_{CMB} per electrical power:

$$C_{i,\nu} \equiv \frac{I_{\text{cal},i,\nu}}{R_{\text{cal, field}, i, \nu}} \frac{F_{\text{long}, \nu}}{F_{\text{short}, \nu}} \frac{F_{\text{HII}, \nu, \text{true}}}{F_{\text{long}, \nu}} \quad (3.3)$$

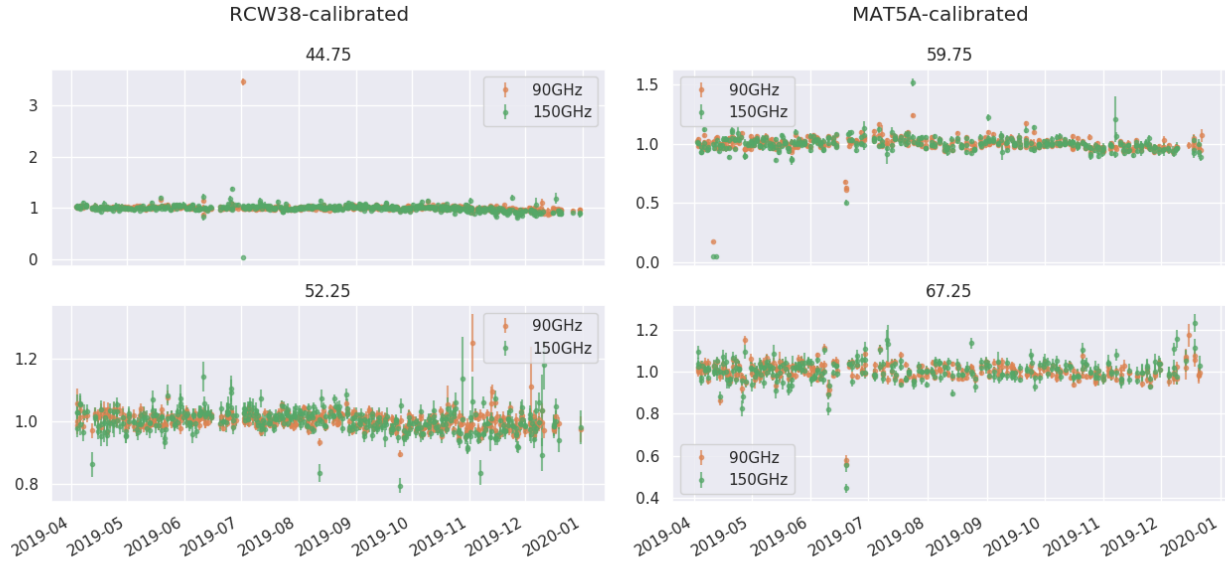
$$= \frac{I_{\text{cal},i,\nu}}{R_{\text{cal, field}, i, \nu}} \frac{F_{\text{HII}, \nu, \text{true}}}{F_{\text{short}, \nu}}. \quad (3.4)$$

The first term in equation 3.3 tracks the change in detector gains, although it may also incorporate a change in the temperature of the calibrator thermal source; the second term tracks changes in atmospheric opacity, and removes any variation in the temperature of the thermal source; the final term applies the absolute calibration to ΔT_{CMB} . On the second line, we have canceled the two $F_{\text{long}, \nu}$ terms.

Upon closer examination, it becomes evident that the weekly long HII region observations are unnecessary. As a reminder, we expect the $I_{\text{cal},i,\nu}$ is fixed, since it is purely a property of the relative locations of the secondary mirror and the focal plane. $F_{\text{long}, \nu}$ completely disappears from the final equation (which is not surprising— we can use the integrated flux in the short HII region observation to convert to ΔT_{CMB} and to track atmospheric opacity simultaneously). The calibration only needs to track the change in opacity from some arbitrary point. Therefore, we could use any one long HII region observation, or a coadd of several and get equivalent results. However, the amount of time we spend on long HII region observations is already a small fraction of the total observation time, so the value of getting rid of them is minimal.

Comparison with Bright Sources in the Field

Our observations do not include any other sources that are bright enough to act as a consistency check on the per-detector calibration. However, we can check the overall



(a) Mean source variability in RCW-38-calibrated subfields. ra0hdec-44.75 on top, and ra0hdec-52.25 below. (b) Mean source variability in MAT-5a-calibrated subfields. ra0hdec-59.75 on top, and ra0hdec-67.25 below.

Figure 3.4: The mean lightcurves, which we use to estimate the calibration error. Panel a shows the observations calibrated using RCW 38, and panel b shows those calibrated using MAT 5a. The subplot titles refer to the mean declination of the subfield in which the bright sources are found. The significant outliers are discussed in more detail in the text; however they are mostly explained by external factors.

calibration using bright point sources in the science fields. All of the brightest sources we observe are AGN, which have a spectrum that falls with increasing frequency. This, unfortunately, means that they do not have sufficient SNR at 220 GHz to provide a useful consistency check on the calibration. AGN are also somewhat variable, but their variability is uncorrelated. Therefore, we use the average brightness of many sources.

We first measure the brightness of 30 sources in each field. To measure their brightness, we fit a beam-sized Gaussian to each source, and take the best-fit amplitude as the source brightness. We do this for many observations to build a lightcurve for each source. Given the distribution of source fluxes (with just a few very bright sources and many dimmer ones) taking the mean of the lightcurves at this point would result in a mean dominated by just a few sources. This would introduce spurious calibration variation from the variability of those particular sources. Therefore, we normalize all the light curves such that their mean is 1.0. Finally, we apply a weighting scheme to decrease the influence of more variable sources. To calculate the weights, we subtract the mean lightcurve from each of the individual lightcurves. We assign the inverse variance of each of the mean-subtracted lightcurves as its weight. We then recalculate the mean using these weights. This downweights sources that vary independently from the rest of the sources, but preserves a coherent variability. If the




Figure 3.5: Histograms of the mean lightcurves. The orange lines show the same quantity calculated using the original calibration pipeline; the blue histogram shows the calibration pipeline described here. There were some adjustments to the mapmaking process (specifically in masking and glitch finding) that reduce the scatter somewhat (as seen by the higher peak in the central bin). However, the largest difference is the data selection criteria, which simply removes the outlier population centered near 0.2 (see §3.1.2). The gray band indicates the range used to estimate the calibration error.

calibration is varying between observations, it should appear as a coherent variability in all of the sources.

Figure 3.4 shows the mean lightcurves after this processing. As expected, they are roughly consistent with 1 (indicating that the calibration pipeline is not injecting variability between observations). However, there are some very significant outliers, which turn out to be caused by external factors. The two biggest outliers in the ra0hdec-44.75 field around September 2019 were caused by a problem with the thermal source in the calibrator. It failed, and was running at a much lower temperature. This caused most of the detectors to be cut when making maps of the short RCW 38 observation, to the point that the source was not imaged at all. Similarly, the cluster of three observations around mid-April of 2019 in the ra0hdec-59.75 subfield were caused by a failure of the chopper itself. The low points around mid-July 2019 in both ra0hdec-59.75 and ra0hdec-67.25 were caused by a problem with the elevation drive that injected a huge noise signal into the short MAT 5a observations. Finally, the two high points in mid-August 2019 in the ra0hdec-59.75 field were caused by particularly bad weather. The summer months (November through March) exhibit higher noise due to the expected increase in atmospheric noise. We used these failures to develop data selection procedures in the calibration pipeline (discussed in §3.1.2).

Finally, Figure 3.5 shows histograms of the mean lightcurves. This figure shows a comparison to an older calibration pipeline (“v1”, in orange), as well. The pipeline discussed here is shown in blue. Since the distribution is highly consistent with 1, we use the gray




Figure 3.6: Histograms of the per-detector signal to noise in long HII region observations over ~ 1 year of observations. The vast majority of detectors have a high SNR on the source over multiple observations. We discard detectors with $\text{SNR} < (60, 60, 20)$ at (90, 150, 220) GHz, which is only a few percent of the total. Note that the x-axis in on the 220 GHz plot is significantly smaller.

region to estimate the statistical error of the calibration pipeline. This will be strictly an overestimate, since some amount of intrinsic source variability will leak into the final light curve. We find the RMS scatter of the lightcurves is 3.7% at 90 GHz and 6.0% at 150 GHz. Since we handle absolute calibration later (see §3.1.3), these bad calibrations only apply the wrong weight to an observation when we make a coadd. Small misweightings inflate the final noise by approximately the square root of the error on the weight. Therefore, the observed scatter only results in a $\sim 2\text{-}3\%$ increase in the final noise from miscalibration between observations.

While we cannot directly test the per-detector relative calibration, this analysis shows that the mean detector calibration per observation is correct. Based on the required SNR of the observations that make up the calibration, we expect (very roughly) expect $\sim 5\text{-}10\%$ errors in the per-detector calibration, and therefore only a $\sim 2\text{-}5\%$ contribution to the total noise from misweighting between detectors.

Calibration Quality Control

As seen in Figure 3.4, there are a handful of observations with extremely bad calibrations. The extreme outliers that have excessively high calibrations will contribute up to several times the weight of the other observations. Therefore, it is important to remove extremely bad calibrations, and there are marginal gains from removing the less extreme outliers.

We remove detectors from both the HII region maps if their calibrator response SNR is < 20 . This removes detectors with either low sensitivity, or high noise. We also choose not to process long HII observations that have fewer than 6000 good detectors. This is about half the number of detectors that typically pass all data selection cuts. If we do not process a long HII region observation, we can simply fall back to the previous good observation, so there is no loss of science data. This cut is responsible for removing the largest outliers seen in Figures 3.4 and 3.5.




Figure 3.7: The HII region MAT 5a without individual detector pointing offsets applied. The map gets one copy of the source for each pixel within the focal plane (as long as it has at least one live detector). In this image, each detector is in uncalibrated electrical power, which leads to the variations in source brightness.

In order to assess the quality of the calibration, we calculate the SNR of the HII region for each single detector map. Here, the SNR is defined as the flux in a $4' \times 4'$ box divided by the RMS of the map outside the signal-dominated region (this masking is discussed in §3.3.5). Note that the “signal” and “noise” here are not defined in the same way (and don’t even have the same units). However, the cut thresholds are arbitrary, so we do not need the formally correct SNR here. The distribution of this SNR metric is shown in Figure 3.6. By default, we cut detectors that have $\text{SNR} < (60, 60, 20)$ at $(90, 150, 220)$ GHz.

3.1.3 Absolute Calibration

As mentioned in §3.1.2, we use SPT-SZ maps (which are calibrated to *Planck*) to measure the “true” HII region flux. This is not perfect, as SPT-SZ and SPT-3G have different bandpass functions and beams, so we expect the flux (and hence final calibration) to be off by 5-10%. When highly precise absolute calibration is required, we use cross spectra between SPT-3G and *Planck* maps (which are calibrated to the CMB dipole (Planck Collaboration et al., 2016b)). This still has a $\sim 1\%$ uncertainty, so we usually leave the absolute calibration of our power spectra as a free parameter in our cosmological fits. For more details on this process, see Dutcher et al. (2021).

3.2 Detector Pointing Offsets

In order to reconstruct an image of the sky with the full resolution of the SPT, we must account for the pointing offset of each detector. Without doing so, a single source would

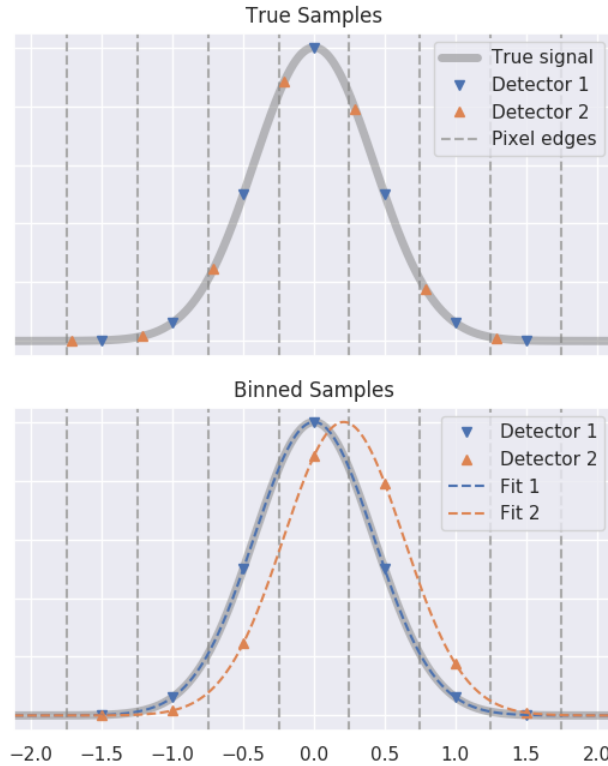


Figure 3.8: A toy model of the coverage from a single observation. The thick gray line shows the true signal, and the dashed gray lines show the boundaries of the map pixels. The colored triangles show two detectors sampling that signal, with a sub-pixel size pointing offset. The top panel shows where the true samples are taken. The bottom panel shows where the samples appear after binning into a map.

appear once for each pixel (for example, see Figure 3.7). We fit the individual detector pointing offsets using maps of HII regions.

There are some subtleties in the way this fit interacts with our scan strategy. By scanning continuously in azimuth, but stepping in elevation, we get very different coverage in the two dimensions. Horizontally, our maps are very well sampled, while they are not well sampled in the vertical direction. The maps are not sparse, however (i.e. every pixel on the source has several samples). Figure 3.8 shows a toy model of our vertical coverage. The thick gray line shows the true signal, and the dashed gray lines show the boundaries of the map pixels. The colored triangles show two detectors sampling that signal, with a sub-pixel size pointing offset. In the top panel, we show where each detector samples the signal in the sky. The blue, downward triangle detector samples exactly at the pixel centers, while the orange, upward triangle detector samples near the pixel edges. In the lower panel, we show the effect of binning these two detector TOD into a map. Since we do not keep any information about subpixel pointing, it now appears that the detectors were sampled at the same location, and the source has a spurious offset. In the map (the lower plot), the apparent offset of the

detector is ~ 0.4 pixels. From the top panel, we can see that its true offset is ~ -0.4 pixels. This example shows the worst case scenario, where the sample spacing (i.e. elevation step in the scan) is equal to the pixel spacing. If we do not account for this effect, we would inject a random error into the vertical offset of each detector. In a map made with many detectors, it would manifest as a convolution of the beam with a $1'$ -tall tophat function in the vertical direction (and no effect in the horizontal direction), which would significantly degrade the SPT's sensitivity on small scales.

There are two obvious solutions for poor coverage in a map. We can increase the map pixel size, although if the pixel size becomes comparable to the source size, we begin to increase our localization error. We can also increase the sampling density. This is impossible in a single observation (increasing the density in a single observation would make the observation too long). However, we have many (~ 75 , at the time this analysis was done) observations with small elevation offsets. We can coadd these to make a map that is well-sampled both horizontally and vertically.

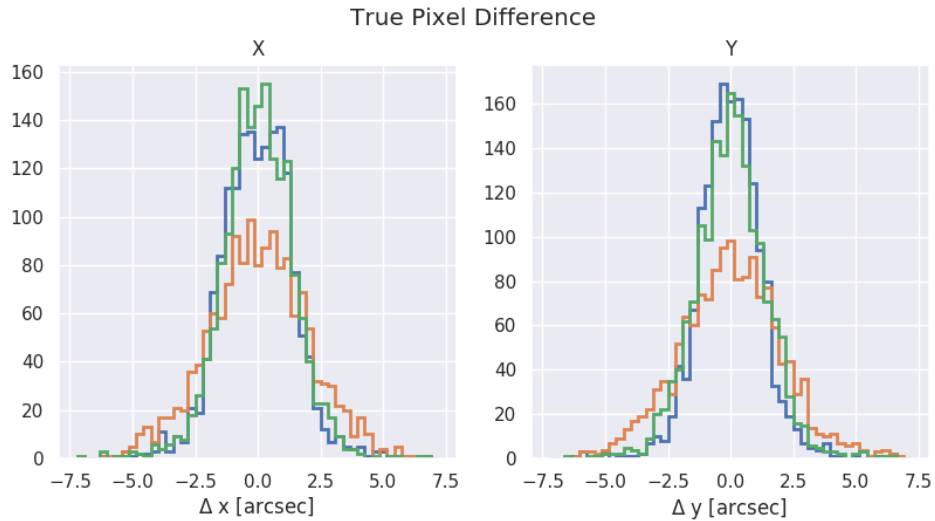
In order to do this, we must first take out an overall pointing offset between observations. So, we first make coadds of all the detectors in one band for each observation. We measure the flux-weighted centroid in each observation, and apply these offsets to a new set of single-detector maps. We then coadd all the observations for each detector to get a densely-sampled map. In order to avoid assigning bad offsets to some detectors, we cut any detector that has fewer than 10 usable maps. To find the pointing offset for each detector, we use the iterative fitting method described in §3.1.1.

Finally, we performed some systematic tests on the calculated offsets. We expect that different detectors in the same pixel and band (i.e. co-pointed detectors observing orthogonal polarizations) should have the same pointing offsets, and that we should get the same offsets whether we use MAT 5a as or RCW 38 maps to find the offsets. In both cases, we find that the difference in the offsets is consistent with zero, as shown in Figure 3.9. Furthermore, these differences are small compared to the arcminute-scale beams, and errors at this level will have a negligible impact on our science results.

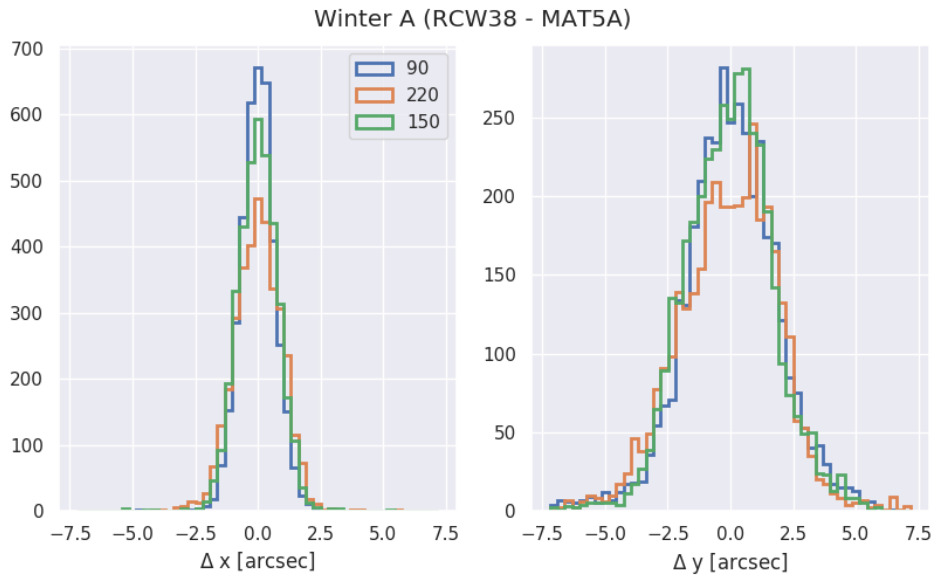
3.3 Mapmaking

With calibrated TOD, and measured detector offsets, we turn our attention to the actual process of making maps. Mapmaking consists of four major steps:

1. Data selection
2. Filtering
3. Weighting
4. Binning



(a) Differences between detector offsets in polarization pairs.



(b) Differences between detector offsets calculated using RCW 38 maps and MAT 5a maps.

Figure 3.9: Histograms showing the difference between detector offsets under different splits of the data. Panel a shows the difference between pairs of detectors in the same pixel. Both of these detectors share the exact same optical path, and only differ by the polarization state they are sensitive to. Panel b shows the difference between pointing offsets calculated from two different sources, MAT 5a and RCW38. In both cases, the left plot shows differences in the x (R.A.) direction, and the right plot shows the differences in the y (declination) direction.




Figure 3.10: A map of a very bright source made without masking the source location. The negative strips on either side of the source are “filtering wings,” which can be avoided by masking the source. The minimum of the filtering wing is ~ -0.25 times the peak amplitude of the source.

3.3.1 Data Selection

In the first step, we discard bad data. The data are discarded at the scan level: we evaluate our selection criteria for every detector in every scan. We remove detectors for a variety of reasons, including

- Signal-to-noise ratio of < 20 in the most recent calibrator observation (see §3.1.1)
- Missing or low-quality calibration data (see §3.1.2)
- Missing identifying information (e.g. some detectors are missing pointing offsets due to producing low quality maps of the HII regions)
- Excess glitches (glitches are short transient spikes in the detector TOD, which lie above the detector noise).

We typically cut ~ 800 of ~ 12000 ($\sim 7\%$) live detectors.

3.3.2 Filtering

We then filter the TOD. The filters are applied to one detector TOD over one scan at a time. These filters have two primary purposes. The first is to reduce low frequency noise, which comes from both the instrument and from brightness variations in the atmosphere. The second is to apply an anti-aliasing low pass filter, to remove signals that on the scale

of the map pixel size or smaller. We implement three different time domain filters. The anti-aliasing low-pass filter is implemented in the Fourier domain. We apply a high-pass filter in the time domain, by fitting and subtracting sinusoidal modes. For each scan, we re-calculate the cutoff frequencies of the high- and low-pass filters to keep the spatial scale of the filters the same across an observation. Finally, we implement a polynomial filter, which fits and remove Legendre polynomials. The polynomial filter is useful for removing very slowly varying modes that the high-pass filter struggles to remove. These modes contain power on scales that are similar to or larger than the scan, and therefore do not map cleanly onto low-frequency Fourier modes (a slope across the entire scan, for example).

The use of the polynomial filter and high-pass filters assumes that the individual detector TOD are noise dominated. This is generally true, as the CMB variations are much smaller than the combined atmospheric and instrument noise in a single detector. However, sufficiently bright point sources will exceed the individual detector noise floor, and bias those filters. This manifests as “filtering wings,” where the flux from the source is spread along the scan direction (see Figure 3.10 for an example). Therefore, we allow some TOD samples to be masked in the polynomial and high-pass filters. This excludes any TOD samples within a given radius of bright sources when fitting the sinusoids or polynomials. However, we do subtract the filtered modes from the masked regions. We typically mask sources that are detected at $\text{SNR} \gtrsim 10$ in a coadd of all the observations, although this threshold may vary depending on the use case for the maps. For most analyses using the 2019 and 2020 data of the 1500 deg² field, we have chosen to mask sources that are brighter than 6 mJy at 150 GHz. This results in ~ 2000 masked sources over 1500 deg², and a $\sim 1\%$ loss in total area.

3.3.3 Weighting

To account for the different detector noise levels, we apply an inverse variance weighting to each detector in each scan. We have two methods of calculating the variance. We either use the band-limited variance (i.e. the sum over each detectors power spectrum in some frequency range, selected to ensure the variance only contains white noise), or we use the total variance, with optional masking in the time domain. The former is used for most maps, but the latter is useful for small maps where a single source might dominate the variance.

3.3.4 Binning

Finally, we bin the filtered and weighted TOD into a map. The value of each pixel in the map is the weighted mean of all the TOD samples that fall within the boundary of that pixel. In order to determine which TOD samples go into each pixel, we use the telescope pointing, combined with the detector offsets described in §3.2.

3.3.5 HII Region Mapmaking

We make maps of the HII regions in the same way as described above with one key difference: we obviously cannot require that the detectors have valid calibrations, since we are in the process of calculating them. We also set the data selection criteria to the most permissive setting possible, so as to produce as many valid calibrations as possible (and leave the decision to cut potentially noisy detectors for a later time). The maps are made in dimensions of electrical power or “source relative” units (as detailed in §3.1.1). The maps are $30' \times 30'$ with 0.5 pixels.

Example maps of RCW 38 and MAT 5a are shown in Figure 3.1. Since the emission of both sources is quite extended, we use large masks (a minimum of $7'$ in radius). RCW 38 requires an irregularly shaped mask. There are some sources in the observed region around RCW 38 that are not masked, and the results of filtering an unmasked source are visible as extended negative wings on either side. However, these are at $\sim 1 - 2\%$ the peak amplitude of RCW 38, and do not impinge on the main source. We do mask MAT 5b, the dimmer, but more extended, source visible at the left edge of the MAT 5a map. This is necessary because MAT 5b is similarly bright, and is at the same declination as MAT 5a, so the filtering wings would intersect the primary source. The additional unmasked sources near both HII regions fall outside of the default $30' \times 30'$ map.

3.4 Conclusion

We now have maps of the sky, calibrated relative to the CMB (i.e. in ΔT_{CMB}). However, these maps are not unbiased. The time domain filtering removes information at certain scales (and preferentially in the scan direction). We can measure this effect using simulations (see Dutcher et al. (2021), for example). The telescope optics also remove information at small scales. We show how the effect of the optics is measured in the next chapter. Both of these effects must be either removed or accounted for in a science analysis, as we will see in the final chapter.

Chapter 4

SPT-3G Beams

The sky contains information at all scales. However, the resolution of the SPT is limited by diffraction at the scale of the primary mirror. This imposes a low-pass filter on the scale of 1 arcminute. In order to recover the true sky signal, we need to characterize the telescope response as a function of angle. This is known as the “beam” (in other fields, it is more commonly called the “point spread function”). The signal that the telescope measures is the true sky signal convolved with the beam. This has several important impacts on the data:

1. A power spectrum measured directly from the maps will be biased low on smaller scales due to the beam. By measuring the beam, we can recover that power, and accurately report the true power spectrum.
2. We can optimally extract objects of known shape from our maps using a matched filter. However, their shape has been convolved with the beam, so we must include its effects to get the optimal recovery. See Chapter 5 for an application of the matched filter to finding clusters of galaxies.
3. The flux from point sources is spread out over the beam, and an accurate flux measurement requires knowledge of the beam to recover all of the flux.

This chapter describes the measurement of the SPT-3G beams. We use observations of both planets and bright sources (all of which are AGN) in our science field. The beam falls off quickly with radius, so the AGN can only be used out to a radius of $\sim 5'$. Planets are much brighter, and have $\text{SNR} > 1$ to several tens of arcminutes. However, the planets are bright enough to saturate our detectors near the center of the source. Therefore, we “stitch” the core of the AGN maps to the planet maps. This gives us a low noise measurement of the beam at all scales, although it does require a careful accounting of systematics. We make many of these composite maps, and use their cross spectra to estimate the beam and its covariance.

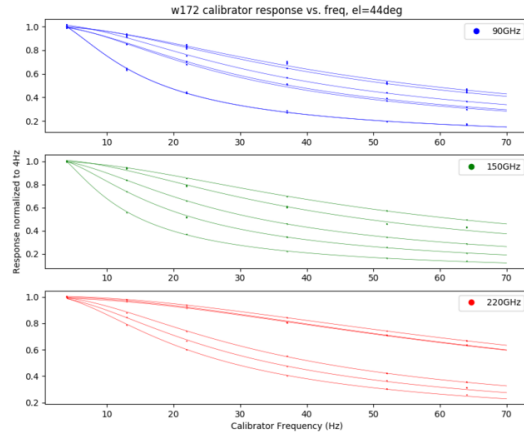


Figure 4.1: Calibrator response as a function of frequency (filled circles) and the transfer function for the best-fit time constant for a sampling of detectors. The response for each detector has been normalized to 1.0 at the lowest frequency. Figure by Zhaodi Pan.

4.1 Systematic Effects

This section describes several systematics we must account for in the process of measuring the beam. Here, we primarily describe the effects; the mitigations are described in detail in §4.3.

4.1.1 Saturation

We use planet observations to get a high SNR on the outskirts of the beam, where the signal may be reduced by orders of magnitude from the center. However, the extreme brightness of the planets deposits enough power on the detectors to make them nonlinear, and eventually saturate. This makes the data near the center of planet unusable. Furthermore, once the detectors saturate, it can take them $\mathcal{O}(0.1)$ seconds to recover. This makes the portion of the scan after passing over the planet unusable.

4.1.2 Detector Time Constants

Our detectors have a non-zero response time, which we characterize as the time constant (τ) in a single-pole filter. In the time domain, this acts as a simple delay between the sky signal and the signal recorded from the detector (to first order). In the Fourier domain, the time constant decreases the detector response as frequency increases. We use a sequence of calibrator (§3.1.1) observations over a wide range of frequencies to measure the detector time constants. We then fit those to a single-pole filter model to measure the detector time





Figure 4.2: The distribution of the mean time constants (τ) for every detector over the 2019 and 2020 observing seasons. The 220 GHz detectors are somewhat faster than the 90 and 150 GHz detectors.

constants. The response and best-fit filter models are shown in Figure 4.1 for a handful of detectors.

In principle, the time constants can vary day-to-day, as they depend on the optical load on each detector. However, we find that they are fairly stable, and we can use their mean, rather than accounting for any time variability. The mean detector time constants are shown in Figure 4.2

Finally, we also have a temporal offset of $\delta t = -4.6$ ms between the detector TOD and the pointing TOD. To first order, this effect is identical to the time constants in the time domain (the time constants have higher orders, but the time offset does not). Purely by chance, the time offset is of approximately the same magnitude as the time constants, but has the opposite sign. Their combined effect is therefore smaller than one would anticipate, but still significant. In particular, it causes discrepancies between beam measurements with AGN at different elevations. To avoid including this as an uncertainty in the final beam, we must remove the time constant and offset effects from the maps.

In the map domain, time constants also cause a shift, to first order. But the map-space shift is determined by the product of the time constant and the on-sky scan speed. Since our scans are at a fixed speed in R.A., the magnitude of the shift is proportional the cosine of the declination. Therefore, in a map made from left-going scans only, everything is shifted slightly towards the left (towards positive R.A.), and the shift is larger closer to the equator. In a map made from right-going scans only, everything is shifted to the right (with the same declination dependence). This is shown schematically in Figure 4.3a. The gray dashed line shows the true sky signal, and the blue and orange lines show the recorded signal in right- and left-going scans, respectively. We use maps that combine both the left- and right-going scans for both our science analyses and beam measurements. The net effect of the time constant and time offset is to broaden the beam parallel to the scan direction, and the magnitude of



(a) True beam (gray), and the beam as seen in (b) The beam after coadding the full left and the left-going (orange) and right-going (blue) right scans (orange), and when using only the scans. approaching portion of the scan (blue).

Figure 4.3: A 1d demonstration of the effect of the time constant on the beam. Panel a shows the true beam (gray dashed), and the beams after the application of a time constant (blue and orange). To first order, each beam is shifted in the direction of the scan. Panel b shows the beam after two different coadding schemes are applied. The orange line shows the standard coaddition of the left- and right-going scans (in which the beam is broadened parallel to the scan direction). The blue line shows the beam when using only the approaching portion of each scan, as we do with the planet maps. In this case, the beam is narrower than expected.

the broadening depends on declination (the orange line in Figure 4.3b).

Due to the effects of the saturation, we also make maps that use only the portion of a scan that approaches the planet (the receding portion is discarded). In this format, the time constants have the opposite effect (it makes the beam narrower parallel to the scan direction). The blue line in Figure 4.3b shows the result of coadding only the approaching portion of each scan.

Assuming the time constants act as a single pole filter, both their effect and the effect of the time offset on the map can be described by a Fourier domain filter.

$$F(\theta, \ell_x) = [1 - i\ell_x \tau v \cos \theta]^{-1} + \exp(i\ell_x \delta t v \cos \theta) \quad (4.1)$$

The first term is the effect of the time constant, and the second term is the time offset. τ is the time constant (which depends on the observing band), ℓ_x is Fourier partner of the horizontal coordinate of the map, v is the scan speed, and θ is the declination of the source. v is positive in right-going scans, and negative in left-going scans. In principle, we should allow θ to vary over the map. However the maps used for the beam only span $30'$ in declination, so the variation of $\cos \theta$ is small, and we ignore it.

To remove the effects of the time constant and offset, we deconvolve this filter from the maps. However, the maps used for our science analyses have not had these effects removed.




Figure 4.4: Radial profiles of the (approximate) beam from the brightest source (black dashed), and 60 simulated backgrounds (colored lines). The simulated backgrounds include Gaussian realizations of the CMB, CIB, and thermal and kinematic Sunyaev-Zel’dovich effects, as well as Poisson realizations of radio sources.

Therefore, we need to apply the mean effect of the time constants and offset over the entire field. The mean effect is given by the square root of the average of $F(\theta, \ell_x)^2$, where the average is taken over the entire field. Since these effects depend only on declination, the mean can be taken over declination, with a weight for the area observed at each declination. Specifically,

$$F_{\text{mean}}(\ell_x)^2 = \Omega^{-1} \int_{\text{field}} d\Omega F(\theta, \ell_x)^2 \quad (4.2)$$

$$= \frac{\Delta\phi}{\Omega} \int_{\theta_{\min}}^{\theta_{\max}} d\theta \cos\theta F(\theta, \ell_x)^2 \quad (4.3)$$

where $(\theta_{\min}, \theta_{\max})$ is the declination extent of the field, $\Delta\phi$ is the width of the field in R.A., and Ω is the area of the entire field.

4.1.3 Background Contamination

Both types of sources (planets and AGN) are much brighter than astrophysical backgrounds (by multiple orders of magnitude). However, on scales of 2’-4’ (where we stitch the AGN to the planet, see §4.3.3), the astrophysical backgrounds become a significant contaminant. This potentially biases the stitching by changing the apparent amplitude of the AGN. Figure 4.4 shows the relative amplitude of the AGN and astrophysical backgrounds. On very large scales ($\gtrsim 10'$), astrophysical backgrounds become significant compared to the planets. However, we are able to directly remove these backgrounds (see §4.3.1).

4.1.4 Intra-band Frequency Dependence

The presence of the Lyot stop in our optical path injects some dependence on the observing frequency into the beam. Let's consider the limit where the illumination across the opening of the Lyot stop is flat (i.e. the beam coming from the detectors is much larger than the Lyot stop). In this case, the illumination of the primary mirror is constant with frequency (and defined by the edge of the Lyot stop). Then, the beam coming from the primary mirror is diffraction-limited, and therefore its width scales as the inverse of the observing frequency (ν^{-1}). Now, let's consider the opposite limit, where the beam from the detectors is entirely contained within the opening of the Lyot stop. In that case, the diffraction-limited detector beam determines the illumination of the primary. The illumination of the primary scales as ν^{-1} , and therefore the beam width on the sky is constant with frequency.

Our 90 GHz band is near the first limit (flat illumination of the Lyot stop), and the higher frequencies (150 and 220 GHz) are somewhere between the two limits. When we measure the beam using AGN (which have a falling spectrum), the result is not appropriate for a different spectrum (such as the CMB, or the CIB which both rise with frequency). To model this, we approximate each beam as a Gaussian with

$$\sigma_{\text{obs}}(\nu_{\text{eff}})^2 = \sigma_{\text{geom}}^2 + \sigma_{\text{diffraction}}(\nu_{\text{eff}})^2 \quad (4.4)$$

where σ_{obs} is the observed beam width, σ_{geom} is the frequency-independent part of the beam width (from geometric optics), and $\sigma_{\text{diffraction}}$ is the frequency-dependent part (from diffraction). ν_{eff} is the effective frequency at which the observation is made. Under our assumption that the 90 GHz band evenly illuminates the Lyot stop, $\sigma_{\text{obs},90} = \sigma_{\text{diffraction}}$. In other words, the 90 GHz band behaves like a standard diffraction-limited beam, with its width inversely proportional to observing frequency. This allows us to assign

$$\sigma_{\text{diffraction}}(\nu_{\text{eff}}) = \sigma_{\text{obs},90} \frac{\nu_{\text{eff},90}}{\nu_{\text{eff}}} \quad (4.5)$$

We now need to determine ν_{eff} .

Following an internal memo by Tom Crawford (Crawford, 2022), we define ν_{eff} as the frequency at which a beam with a δ -function bandpass has the same solid angle as our observed beam. We start by defining the observed beam as

$$B^{\text{obs}}(\hat{\theta}) \equiv \frac{P(\hat{\theta})}{P(0)}. \quad (4.6)$$

This states that the measured beam (in map space) is the power received by the SPT from a point source at an angle $\hat{\theta}$ from the source, divided by the power received when pointing directly at the source. We now expand $P(\hat{\theta})$ and $P(0)$ in terms of the source spectrum $I(\nu)$, the bandpass function $t(\nu)$, the solid angle of the point source Ω_s , the effective aperture area

$A(\nu)$, and $B(\nu, \hat{\theta})$, the beam for a δ -function bandpass at frequency ν :

$$B^{\text{obs}}(\hat{\theta}) = \frac{\int \Omega_s A(\nu) I(\nu) t(\nu) B(\nu, \hat{\theta}) d\nu}{\int \Omega_s A(\nu) I(\nu) t(\nu) B(\nu, 0) d\nu} \quad (4.7)$$

$$= \frac{\int A(\nu) I(\nu) t(\nu) B(\nu, \hat{\theta}) d\nu}{\int A(\nu) I(\nu) t(\nu) d\nu} \quad (4.8)$$

In the second line, we have assigned $B(\nu, 0) = 1$, which just sets the overall normalization of the beam. This just states that the observed beam is a weighted mean of the δ -function bandpass beams, weighted by the power received at each frequency.

The solid angle of the beam is

$$\Omega_b^{\text{obs}} = \int B^{\text{obs}}(\hat{\theta}) d\Omega \quad (4.9)$$

$$= \int \frac{\int A(\nu) I(\nu) t(\nu) B(\nu, \hat{\theta}) d\nu}{\int A(\nu) I(\nu) t(\nu) d\nu} d\Omega \quad (4.10)$$

$$= \frac{\int A(\nu) I(\nu) t(\nu) \int B(\nu, \hat{\theta}) d\Omega d\nu}{\int A(\nu) I(\nu) t(\nu) d\nu} \quad (4.11)$$

We now need to evaluate $\int B(\nu, \hat{\theta}) d\Omega$. First, we assume $A(\nu) = \text{const.}$, so $\int B(\nu, \hat{\theta}) d\Omega \propto \nu^{-2}$. We use ν_{eff} as the pivot frequency, such that

$$\int B(\nu, \hat{\theta}) d\Omega = \Omega_b(\nu_{\text{eff}}) \left(\frac{\nu_{\text{eff}}}{\nu} \right)^2 \quad (4.12)$$

Inserting this assumption into equation 4.11 gives

$$\Omega_b^{\text{obs}} = \frac{\int A(\nu) I(\nu) t(\nu) \int B(\nu, \hat{\theta}) d\Omega d\nu}{\int A(\nu) I(\nu) t(\nu) d\nu} \quad (4.13)$$

$$= \frac{\int I(\nu) t(\nu) \Omega_b(\nu_{\text{eff}}) \left(\frac{\nu_{\text{eff}}}{\nu} \right)^2 d\nu}{\int I(\nu) t(\nu) d\nu} \quad (4.14)$$

$$= \Omega_b(\nu_{\text{eff}}) \nu_{\text{eff}}^2 \frac{\int I(\nu) t(\nu) \nu^{-2} d\nu}{\int I(\nu) t(\nu) d\nu} \quad (4.15)$$

To determine the value of ν_{eff} , we set $\Omega_b^{\text{obs}} = \Omega_b(\nu_{\text{eff}})$, and solve for ν_{eff} :

$$\nu_{\text{eff}} = \left[\frac{\int I(\nu) t(\nu) \nu^{-2} d\nu}{\int I(\nu) t(\nu) d\nu} \right]^{-\frac{1}{2}}. \quad (4.16)$$

We now fit the free parameters of this model using observations of the brightest source in our science field. Since we assume that the 90 GHz evenly illuminates the Lyot stop, the only

Band	ν_{eff} GHz	σ_{obs} arcmin	$\sigma_{\text{diffraction}}$ arcmin	σ_{geom} arcmin
90 GHz	93.7	0'66	0'66	0'
150 GHz	147.1	0'50	0'42	0'26
220 GHz	219.9	0'42	0'28	0'31

Table 4.1: Model parameters defining the frequency dependency of the beam. ν_{eff} is the effective frequency at which σ_{obs} is measured, assuming a source spectrum $I(\nu) \propto \nu^{-0.5}$. σ_{obs} is the best-fit Gaussian width of our brightest source. $\sigma_{\text{diffraction}}$ is the frequency-dependent portion of the width, and σ_{geom} is the frequency-independent portion. By definition, $\sigma_{\text{geom},90} = 0$, and $\sigma_{\text{obs}}^2 = \sigma_{\text{diffraction}}^2 + \sigma_{\text{geom}}^2$.

free parameters are $\sigma_{\text{geom},150}$, and $\sigma_{\text{geom},220}$. We use the best-fit circular Gaussian to define σ_{obs} for each band. The spectrum of AGN are generally falling power laws. For this source, we find $I(\nu) \propto \nu^{-0.5}$ is a good fit. Finally, we use the measured bandpasses for the SPT-3G receiver. The results are shown in Table 4.1. This shows that the 150 GHz band is close to the diffraction limit, and the 220 GHz band is evenly split between geometric effects and diffraction.

This model allows us to approximate the change in the beam for different source spectra. Under the Gaussian beam approximation, we simply (de)convolve another Gaussian to convert beams between sources with different spectra. In Fourier space, this is a multiplication by a Gaussian. For example, to convert a Gaussian beam measured on AGN to a beam appropriate for the CMB, we multiply by

$$\exp(-\ell^2(\sigma_{\text{CMB}}^2 - \sigma_{\text{AGN}}^2)/2) \quad (4.17)$$

where σ_{CMB} , and σ_{AGN} are the Gaussian widths for the CMB and AGN respectively. Using our model, this can restated in terms of effective frequencies:

$$\sigma_{\text{AGN} \rightarrow \text{CMB}}^2 \equiv \sigma_{\text{CMB}}^2 - \sigma_{\text{AGN}}^2 \quad (4.18)$$

$$= (\sigma_{\text{geom}}^2 + \sigma_{\text{diffraction}}(\nu_{\text{eff,CMB}})^2) - (\sigma_{\text{geom}}^2 + \sigma_{\text{diffraction}}(\nu_{\text{eff,AGN}})^2) \quad (4.19)$$

$$= \sigma_{\text{obs},90}^2 \nu_{\text{eff,AGN},90}^2 (\nu_{\text{eff,CMB}}^{-2} - \nu_{\text{eff,AGN}}^{-2}) \quad (4.20)$$

Since the CMB spectrum rises with frequency, and the AGN spectrum falls with frequency, and the SPT-3G bandpasses are approximately tophats, $\nu_{\text{eff,CMB}} > \nu_{\text{eff,AGN}}$. Therefore, $\sigma_{\text{AGN} \rightarrow \text{CMB}}^2 < 0$, and the multiplicative conversion in Fourier space rises with ℓ . In other words, the CMB beam is narrower than the AGN beam in map space. This conversion is shown in Figure 4.5.

This model has a few problems. First, there is an inconsistency in how ν_{eff} is defined. We have calculated ν_{eff} assuming that the beam area Ω is proportional to ν^2 , and the aperture area A is fixed. Our model uses a more complex formulation of the beam area, which includes




Figure 4.5: The conversion from a beam measured on an AGN (with $I(\nu) \propto \nu^{-0.5}$) to a beam appropriate for the CMB spectrum. This is a negligible effect at 220 GHz, small at 150 GHz, and substantial at 90 GHz.

terms that are both linear in ν and constant. Our model does assume A is fixed (at least within a band), however that cannot be the case at 150 and 220 GHz, where σ_{geom} is of the same order as $\sigma_{\text{diffraction}}$. We have also fixed $\sigma_{\text{geom},90} = 0$, under the assumption that the 90 GHz band over-illuminates the Lyot stop significantly. Finally, we assume that frequency dependence is scale independent, which we also know is not true. For example, we expect some large scale response in the beam from imperfections in the alignment of the panels that make up the primary mirror (Ruze scattering (Ruze, 1952, 1966)). This effect increases with frequency, rather than decreasing, as our model for the main lobe does. There is work in progress to address some of these concerns (see §4.8.1). Nonetheless, we show that the data are not inconsistent with this model (§4.6.4).

4.2 Data

We use two types of sources to measure the beams: planets, and AGN. The planets provide an extremely high signal-to-noise ratio (SNR), but are so bright that they cause the detectors to saturate at radii of $r \lesssim 2'$. We use AGN in our science field to fill in the core of the beam, where the planet maps are unreliable.

4.2.1 AGN Maps

We selected the 10 brightest sources in the 1500 d² winter field to measure the beam. However, two of them have dimmer sources within a few arcmin, so we reject those two sources, leaving us with 8. Of these sources, one is clearly brighter than the others, at ~ 3

Source Name	R.A.	Declination	Flux [mJy]		
	deg	deg	90 GHz	150 GHz	220 GHz
0210-51	32.693	-51.017	2788	2245	1827
2329-47	352.324	-47.505	1223	865	600
2056-47	314.068	-47.246	1096	882	685
0309-60	47.484	-60.978	1050	806	617
0133-52	23.274	-52.001	1028	820	651
2357-53	359.472	-53.187	874	638	461
0253-54	43.372	-54.697	833	652	502
0049-57	12.497	-57.641	826	603	442

Table 4.2: The AGN used for beam measurement.

mJy, while the remainder are around ~ 1 mJy at 90 GHz. See Table 4.2 for their fluxes in all three bands, and positions.

Since these sources are in our main science field, we have ~ 1000 observations of each source. We make small maps centered on each source from the 2019 and 2020 observations. We use observations of entire subfields, but only create small output maps. The TOD filtering follows the process outlined in §3.3, with a $15'0$ mask around each source, the high pass filter is set at $\ell = 500$, and a 9th order polynomial subtraction. At the $0'1$ scale, the SPT has effectively no response, so we do not apply a time-domain low pass filter. The beam itself provides sufficient anti-aliasing at these scales. The masking prevents the TOD filtering from impacting the beam, since we only use the AGN maps at scales smaller than $5'$ in radius. The maps are $30' \times 30'$, and are pixelized at $0'1$.

We coadd the maps of each source, where each map is weighted by its inverse variance. To measure the variance we first mask out the source with a $5'0$ mask, then apply a $3'0 \times 3'0$ mean filter, and calculate the total variance. This weighting scheme captures noise from atmospheric fluctuations that are common between detectors, and therefore not correctly accounted for in the map weights. It also allows us to weight on exactly the most important spatial scales.

4.2.2 Planet Maps

We use 9 dedicated observations of Saturn, performed in 2019 and 2020. In principle, we could use other planets (previous cameras on the SPT have used Mars, Venus and Jupiter). However, we are limited by which planets are far enough above the horizon at the South Pole. We make the Saturn maps with a much larger mask and lighter TOD filtering to recover the large scales. We use a $30'0$ mask, 2nd order polynomial, and set the high pass filter¹ at

¹Since the high pass filter is masked, we do not remove the large scale modes present in the planet.

Band	σ_{main} arcmin	σ_{sidelobe} arcmin	a
90 GHz	0.66	2.8	0.011
150 GHz	0.48	1.2	0.033
220 GHz	0.39	0.75	0.086

Table 4.3: Parameters of the input beams used for simulations. For each band, the beam is modeled as $B_\nu(r) \propto \exp(-r^2/2/\sigma_{\text{main}}^2) + a \exp(-r^2/2/\sigma_{\text{sidelobe}}^2)$.

$\ell = 200$. The planet maps are $6^\circ \times 6^\circ$, and pixelized at $0'.1$, to match the AGN maps. We also apply a small perturbation to the telescope pointing to track the movement of Saturn over the course of a 6-hour observation. Instead of producing one map per observation, we split each observation into left- and right-going scans, and make maps from these separately.

4.2.3 Simulations

We use simulations to test that the beam analysis pipeline is properly constructed. To construct the inputs to the simulation, we use a very approximate model of the beams. We assume each band can be modeled as the sum of two Gaussians. One (which is narrower and has a higher amplitude) represents the main lobe of the beam. The second is much wider and lower amplitude, and represents the sidelobes. The exact details of the input beam are unimportant, but they do need to have approximately the right signal levels in the stitching region, so the sidelobe contribution is important. In summary each band is modeled as $B_\nu(r) \propto \exp(-r^2/2/\sigma_{\text{main}}^2) + a \exp(-r^2/2/\sigma_{\text{sidelobe}}^2)$, with individual band parameters as shown in Table 4.3. These radial profiles are inserted into maps with perfect circular symmetry (i.e. no azimuthal dependence).

Depending on the simulation, we may also use two different types of noise: white noise and simulations of the mm-wave sky. The simulations include Gaussian realizations of the CMB, CIB, tSZ, kSZ, and Poisson realizations of radio sources with 150 GHz flux less than 6 mJy. The power spectra used to generate the Gaussian realizations are the best-fit models from (Reichardt et al., 2021).

4.3 Map Preparation

This section describes how we prepare the planet and AGN maps, and then stitch them together to form the composite beam map.

4.3.1 Saturn Preparation

We apply 4 processes to the planet maps.

1. CMB cleaning
2. Remove a scan-synchronous signal from right-going scans
3. Recentering
4. Masking and coadding left- and right-going maps

CMB Cleaning

Even though the signal-to-noise ratio on the planet is very high (> 1000), the large scales still contain a significant CMB contribution. Fortunately, we are able to clean the CMB from the planet maps. Since planets move in (R.A., Dec.), we can use previous observations of the current location of the planet to remove the CMB. In this case, we use maps from the *Planck* satellite. For each band and planet observation, we mock observe map of the closest *Planck* band. In a mock observation, we simulate an observation of a given input sky with the SPT. We use the real pointing information from an observation to determine where in the input map each detector is pointing, then assign those values to the simulated detector TOD. From there, we proceed with the usual mapmaking procedure. This ensures that the same modes are removed by our time domain filtering. We then subtract the mock-observed map from our planet maps to remove the CMB.

Scan-synchronous cleaning

In the 2020 observations of Saturn (7 of 9), the telescope elevation has not settled before starting the right-going scans. This modulates the atmospheric loading on the detectors, and leaves a noticeable oscillatory mode in the right-going maps. The mode is purely in the x /R.A. direction, because the scans are at fixed elevation. Figure 4.6 shows left- and right-going maps for comparison. The right-going map has a large oscillatory mode that starts on the right side, and decays as it moves towards the left.

Since the injected mode is so well synchronized between scans, we can effectively remove it directly from the maps. First, we mask Saturn with a $40''$ circle, and the the two visible sources closest to the center of the map with a $3''$ radius. This prevents subtracting real signal (the Saturn mask) or spreading a contaminating signal in the y direction (the additional source masks). We then average the map in the y direction, to measure the average scan synchronous mode. Finally, we subtract that mode from every row of pixels. Figure 4.6c shows the right-going map after the scan-synchronous mode is removed.

Recentering

We are unable to use our usual pointing corrections for the planet maps, since those are based on bright sources scattered throughout a large field. However, we can simply center our maps on the planet itself. For each map, we fit a Gaussian to the planet, and find the offset that would bring the peak of that Gaussian to the center of the map. We then




Figure 4.6: A comparison of Saturn maps made from left-going (a), and right-going (b) scans, and right-going scans after cleaning (c). The right-going scans show an oscillatory mode which decays from left to right.




Figure 4.7: This shows the Saturn maps coadded without masking, minus the same maps coadded with our masking procedure. The x and y scales are marked in arcminutes from the source center, and the color scale is in units of the peak brightness of the source. The unmasked map is analogous to the orange line in Figure 4.3b, and the masked map is analogous to the blue line. As expected, the unmasked maps are wider in the scan direction, but equal to the masked maps at the center. The 150 GHz map shows the largest effect, as the detectors are operated closer to their saturation power than the other bands.

apply that offset in the Fourier domain, to avoid additional smoothing caused by map-space interpolation.

Masking and Coaddition

Our detectors saturate as they pass over Saturn. This causes the time constants to increase significantly, and the detector to become highly non-linear. It may take several time constants for the detector to recover. Therefore, we mask the receding portion of scans that come within $2\prime 0$ at 90 and 150 GHz, or $1\prime 0$ at 220 GHz. The effect of the masking is shown




Figure 4.8: This shows the Saturn maps after accounting for the time constant and offset minus the same maps before accounting for the time constant and offset (these are the maps produced by the masking step). The x and y scales are marked in arcminutes from the source center, and the color scale is in units of the peak brightness of the source. The masked map is analogous to the blue line in Figure 4.3b, while the map after accounting for the time constants and time offset is closest to the mean of the orange and blue lines in Figure 4.3a. The overall effect is to narrow Saturn in the scan direction.

in Figure 4.7. These masks cover the opposite sides of the map in the left- and right-going maps, so when we coadd them together, we get a complete map of the planet. However, this requires an additional step to account for the time constants.

Time Constants and Offset

As described in §4.1.2, this masking narrows the beam in the scan direction. Since the receding portion of the scan is contaminated by saturation and non-linear effects, we cannot simply deconvolve the time constant and time offset from the maps. Therefore, we apply an iterative forward modeling approach. We begin by generating separate transfer functions for the left- and right-going maps from the measured time constants and time offset. We then apply these transfer functions to a biased estimate of the true planet map (in this case, the coadd of the left- and right-going maps with the receding portions of the map masked). This produces biased left- and right-going maps, which we coadd with the same mask to produce a doubly-biased coadd. We then take the difference between the doubly-biased map, and the estimate of the true map, which produces an estimate of the effect we are trying to remove. We subtract that from the estimate of the true map, and repeat this process until we have converged on the true map. With a satisfactory estimate of the true map, we apply the mean filter for the field (see equation 4.2). We test this method on the brightest AGN in our science field, and find that only two iterations are required to converge, with residuals of $< 1\%$ of the peak amplitude of the source.

For the unmasked portions of the map, we first deconvolve the transfer function of the




Figure 4.9: The Saturn maps before accounting for the RJ spectrum of Saturn, minus the same maps after accounting for the spectrum. The x and y scales are marked in arcminutes from the source center, and the color scale is in units of the peak brightness of the source. This effect is isotropic, by construction, and makes Saturn broader (since the effective observing frequency on a RJ spectrum is higher than that of an AGN-like spectrum). It is also several times smaller than the other processes applied to the Saturn maps.

time constant and time offset from each map (equation 4.1). We use the mean of the time constants measured immediately before and after each Saturn observation. Then, we convolve with the mean transfer function for the 1500 deg² field (equation 4.2). For the science field, we use the mean time constants measured over 2019 and 2020 (Figure 4.2). For each band, we take the weighted mean over all detectors in that band, with weights taken from the average weight each detector receives in the mapmaking process. This results in time constants of (6.9, 6.5, 4.3) ms at (90, 150, 220) GHz. This makes the beam in the Saturn maps match the mean beam in the 1500 deg² science maps. The effect of correcting for the time constant and offset is shown in Figure 4.8.

Spectrum Conversion

Saturn is much hotter than the CMB, with brightness temperature of ~ 100 K across all three of our bands (Planck Collaboration et al., 2017; Hasselfield et al., 2013b; Dunn et al., 2005). So, we assume Saturn follows a Rayleigh-Jeans spectrum ($I \propto \nu^2$). We apply the spectrum correction described in §4.1.4. Instead of converting from an AGN-like spectrum to a CMB spectrum, we instead convert from the RJ spectrum to an AGN-like spectrum. Since we are working with maps, we must also construct the transfer function in two dimensions. Finally, we also add a small broadening for the pointing error in the field observations. The maps of the science field have a small random error in pointing reconstruction, which we treat as a circular Gaussian with standard deviation $\sigma_{\text{jitter}} = 3''$. Since we work with individually recentered Saturn observations, we need to broaden the maps by σ_{jitter} to match the AGN

maps. With these adaptations, the transfer function is

$$\exp -(\ell_x^2 + \ell_y^2)(\sigma_{\text{RJ}}^2 - \sigma_{\text{AGN}}^2 + \sigma_{\text{jitter}}^2)/2 \quad (4.21)$$

(compare this to equation 4.17). This effect is shown in Figure 4.9.

4.3.2 AGN Preparation

We apply 4 processes to the AGN maps before stitching them:

1. Convolution with the planet disk
2. Correct for the time constants and offset
3. Recentering
4. Background subtraction

The first three operations are applied in Fourier space. Finally, we combine all of the AGN maps into a single stack, with each AGN weighted by its signal-to-noise ratio.

Planet Disk

The finite extent of Saturn cannot be completely neglected. To account for it, we convolve the AGN (which *are* point-like) with a disk that has a radius of $8''.55$. This may seem like an odd choice at first, since Saturn has bright rings that are several times larger than the planet itself. However, the rings are simply very dim at mm wavelengths². The BIMA collaboration made high resolution images of Saturn (Dunn et al., 2005) which allow separation of the rings from the planet. The rings are only $\sim 15 - 20\text{K}$, compared to the 100 K planetary disk. Therefore, we ignore emission from the rings.

Time Constants and Offset

Each AGN is subject to a slightly different time constant and time offset transfer function because of their different elevations. To make them all match the mean over the 1500 deg^2 field, we first remove the transfer function for the the declination of each AGN. We then apply the mean transfer function over the 1500 deg^2 . Figure 4.10 shows the resulting change to one of the AGN maps.

²The brightness of the rings in optical images comes from reflected sun light, which is a negligible contribution to the mm-wave emission from Saturn or its rings (Esposito et al., 1984; Cuzzi, 1984).





Figure 4.10: The map of one AGN before accounting for the time constants, minus the same map after accounting for the time constants. The x and y scales are marked in arcminutes from the source center, and the color scale is in units of the peak brightness of the source. This source is below the mean elevation of the field, so it is more extended in the scan direction (horizontal) than the mean. The very high frequency modes in the scan direction are induced by inverting the transfer function, which has zeros in it. However, these occur at extremely high ℓ ($\approx 40,000$), and are irrelevant to the SPT-3G beam.

Recentering


The map centers used for mapmaking were based on preliminary source positions, and so the sources lie slightly off the center of the map. We correct for this by fitting a Gaussian to the full-depth coadd, and find the shift that puts the peak of the Gaussian at the center of the map. While the individual observations have low enough noise to perform this recentering, there is some scatter in the pointing between observations. The same scatter is in the science observations, and effectively broadens the beam, so we need to include it here. This operation is performed in Fourier space to avoid the smoothing caused by interpolation.

Background Subtraction

In order to stitch the planet and AGN maps, we are forced to use the AGN maps at 2-3 times the FWHM (see §4.3.3). The signal from the source itself is many times lower than the peak in this range, and therefore we need to carefully consider sources of noise. We measure the CMB at high signal to noise ratio, so the instrumental noise will never be relevant. The CMB itself (and other astrophysical backgrounds) make up the majority of the noise. The CMB power spectrum peaks at much larger scales (\sim degree scales, as opposed to the several-arcminute scales considered here), so the noise can be approximated very well by a low-order expansion. In this section, we show that a second order background model can effectively subtract the majority of the noise, and provides a higher signal-to-noise ratio for the stitching.



(a) Before background subtraction.



(b) After background subtraction

Figure 4.11: Radial profiles of the (approximate) beam from the brightest source (black dashed), and 60 simulated backgrounds (colored lines). The simulated backgrounds include Gaussian realizations of the CMB, CIB, and thermal and kinematic Sunyaev-Zel'dovich effects, as well as Poisson realizations of radio sources. Panels a and b show the radial profiles before and after background subtraction, respectively.

Assuming the beam is symmetric about its center of mass (which is a good approximation), the contribution of the background modes are effectively azimuthally averaged about the center of the source. Any mode that is odd about the center of the source (such as a slope) contributes no noise. Therefore, we chose to model the background as constant offset and a slope in radius. The slope in radius is a cone shape in the 2d map, which maps onto an infinite number of smoothly varying modes in x and y .

To fit the model parameters (the slope and offset), we use the coadd of planet maps described above. The fit uses the 90 and 150 GHz maps, since we expect the CMB to be completely dominant in those bands (the 220 GHz maps contain a significant CIB contribution). We fit the parameters over an annulus centered on the source, with $2'.25 < r < 6'.0$. The lower limit ensures that we don't include the portion of planet that is contaminated by saturation, and the upper limit constrains the fit to regions where the planet has a high signal-to-noise ratio. We run the fit on the coadd of each AGN, and then subtract the background (over the entire map, not just the fitting region).

To test the efficacy of the background subtraction, we use simulated maps with astrophysical backgrounds included. Figure 4.11 shows the simulated radial profiles before and after background subtraction. These used an annulus at slightly larger radius than the final beam maps, but the qualitative results are unaffected: this process effectively removes background noise from the AGN maps.

4.3.3 Stitching

With the maps prepared, we can now combine the planet and AGN maps into a single beam map. The planet and AGN have dramatic differences in brightness, so we scale the planet map to match the AGN map. We cannot use the core of the beam (where the SNR is highest) in the fit, because the planet maps are contaminated by saturation in that region. This adds noise to the fit, since we are forced into a region of lower SNR on the AGN. Therefore, we do the fit in two stages. First, we fit the AGN map to the stack of all AGN maps, with only an amplitude scaling as a free parameter. This fit uses the full AGN map ($30'.0 \times 30'.0$), and includes the region of maximum SNR. We then fit the stack of all the AGN to the planet map, using only a small region at a radius of $r_{\text{inner}} \leq r \leq r_{\text{outer}}$, again with only an amplitude scaling as a free parameter. Using the stack gives us maximum SNR on stack-to-planet amplitude. The product of the two best-fit amplitudes gives the final amplitude scaling to make the planet map match the AGN.

To stitch the maps together, we first apodize the AGN map with a mask that linearly transitions from 1 to 0 over $r_{\text{inner}} \leq r \leq r_{\text{outer}}$. We apodize the planet map with the opposite mask (i.e. transition from 0 to 1 from r_{inner} to r_{outer}). Finally, we add the two apodized maps together. We stitch every AGN map to every Saturn map, creating a total of $n_{\text{AGN}} \times n_{\text{planet}} = 72$ beam maps.

	r_{inner}	r_{outer}	BG r_{inner}	BG r_{outer}	$\bar{\alpha}_{\text{AGN}}$	δt	τ	r_{disk}	σ_{jitter}
	arcmin	arcmin	arcmin	arcmin		ms	ms	arcsec	arcsec
90 GHz	2'25	3'5					6.88		
150 GHz	1'5	3'0	2'25	6'0	-0.622	-4.6	6.53	8''55	3''0
220 GHz	1'0	2'5	-	-			4.27		

Table 4.4: Fiducial input parameters for the beam pipeline.

4.4 Parameter Variations

Through all of the steps outlined above, we have accrued 9 important input parameters for the beam pipeline. They are

r_{inner} The inner radius of the stitching region, per band.

r_{outer} The outer radius of the stitching region, per band.

BG r_{inner} The inner radius of background fitting region.

BG r_{outer} The outer radius of background fitting region.

$\bar{\alpha}_{\text{AGN}}$ Mean spectral index of the AGN.

δt Time offset between detector and pointing TOD.

τ Mean of detector time constants, per band.

r_{disk} Radius of the disk of Saturn.

σ_{jitter} Pointing reconstruction uncertainty.

Their fiducial values are listed in Table 4.4

For the purposes of measuring the beam uncertainty, we repeat this process with some of the input parameters varied. This allows us to approximately marginalize over systematic variations from our analysis choices in the mean beam, and include those variations in the final uncertainty. We vary the stitching region, moving the inner and outer radii by $\pm 0'5$ separately. Under ideal circumstances, this would not cause any change to the beam. However, the Saturn and AGN maps do not match exactly, and varying the stitching region accounts for this

systematic uncertainty. We also vary the region used to fit the background that is subtracted from the AGN maps, for the same reason. Finally, we vary the time offset between pointing and detector TOD to account for uncertainty in the time constants and time offset. In the observations used to determine the time offset, we find that the observation-to-observation scatter on the time constants is ~ 0.7 ms. It is also highly correlated between bands, so simply apply a ± 0.7 ms variation to the time offset (rather than varying the time constants for each band individually). The mean time constants themselves are extremely well-determined over two years of observations, so any additional uncertainty is negligible.

4.5 Beam and Covariance Estimation

We use cross spectra of the different stitched maps to estimate the beam window function. The beam window function is the power spectrum of the beam map, and labeled W_ℓ . We cannot use every possible pair of stitched maps, as we would generate a noise bias from taking the cross spectrum of one of the input maps with itself. For example, we cannot cross the stitched map made from $\text{AGN}_1, \text{planet}_1$ with the stitched map made from $\text{AGN}_1, \text{planet}_2$. The planet maps are different, but the AGN map is the same, and has the same noise near the center of the source (i.e. below the stitching radius). Since the beam normalization is arbitrary, we fix the beam to unity at an ℓ that is appropriate for the analysis. The beam normalization is completely degenerate with absolute calibration, which we determine from cross spectra with *Planck* maps. The ℓ range used to do this varies depending on the analysis goals (e.g. the high- ℓ TT power spectrum uses $1200 < \ell < 1500$, whereas the mid- ℓ analyses use $600 < \ell < 1200$). For the purposes of this chapter, we have fixed the beam to 1 at $\ell = 800$.

The beam maps are sufficiently small that the flat-sky approximation works very well, so we are able to use the standard Fourier transform (rather than the spherical harmonic transform). Before taking the Fourier transform, we apodize the beam maps to remove noise-dominated regions. The apodization mask is unity at $r < 0.75^\circ$, and has a cosine edge taper smoothly goes to 0 at $r = 1.25^\circ$. We tested several variations of the apodization mask (including a maximal version that only apodizes the noisy map edges), and found no significant differences.

Including all the parameter variations, we have 792 beam maps, and can create $\mathcal{O}(10^5)$ unique cross spectra. Actually calculating all these spectra is infeasible, so we actually calculate a semi-random subset of these spectra to measure the mean beam window function. The generation of random subsets is motivated by the expected correlation structure of the spectra. We randomly select a subset of the stitched maps such that each planet and AGN map is used at most once (this requires that we leave out one of the planet maps, since we have 8 AGN maps and 9 planet maps). For each map, we randomly select a parameter to vary. For most parameters (everything except the inner and outer stitching radii), we then chose the sign of the variation, and use it for all three bands. For the inner and outer stitching radii, we choose the sign of the variation at random for each band. This is because

the choice of stitching region is independent for the three bands, and therefore we do not want to capture any covariance from varying those parameters together. This forms a set of maps with independent noise. We repeat this selection process 1000 times to get a reasonable sampling of different cross spectra and parameter variations.

To calculate the mean window function, we simply take the mean of all these spectra. Finally, we account for two semi-analytic effects: the planetary disk, and the AGN spectrum. As described above, we convolve the AGN maps with a disk of $8''.55$ in radius to match the size of Saturn. We now deconvolve the same disk from B_ℓ (the square root of the beam window function). We also use the frequency dependence model described in §4.1.4 to convert B_ℓ from the AGN spectrum to the CMB spectrum. To do so, we use the mean of the measured AGN spectral indices (which is $\bar{\alpha}_{\text{AGN}} = -0.622$).

To estimate the uncertainty, we extend the procedure described in Appendix A of Lueker et al. (2010) (hereafter L10). L10 describes a method of calculating the covariance of all unique spectra created from a set of independent maps. For a set of n_{maps} maps, the covariance of the mean of their cross spectra is given by

$$\mathbf{C}_{\ell\ell'}^{\text{window}} = \frac{2f(n_{\text{maps}})}{n_{\text{maps}}^4} \sum_{\lambda} \sum_{\alpha \neq \lambda} \left(\Delta D_{\ell}^{\lambda\alpha} \Delta D_{\ell'}^{\lambda\alpha} + 2 \left[\sum_{\beta \neq \lambda\alpha} \Delta D_{\ell}^{\lambda\alpha} \Delta D_{\ell'}^{\lambda\beta} \right] \right) \quad (4.22)$$

where $\mathbf{C}_{\ell\ell'}^{\text{data}}$ is covariance matrix of the cross spectra, $\Delta D_{\ell}^{\lambda\alpha}$ is the mean-subtracted cross spectrum formed from maps λ and α , and $f(n_{\text{maps}})$ is a factor that accounts for the finite number of maps. Since we have non-independent beam maps, this method is not directly applicable. However, the previous paragraph describes an algorithm to select subsets of the beam maps that are independent. Therefore, we apply the L10 covariance calculation to each set of independent maps, and take the mean over all of the resulting covariance matrices. L10 uses 300 different maps, and assumes $f(n_{\text{maps}}) \approx 1$. With only 8 AGN maps, we must use account for this factor properly.

$$f(n_{\text{maps}}) = \frac{n_{\text{maps}}^4}{n_{\text{maps}}(n_{\text{maps}} - 1)^2(n_{\text{maps}} - 2)} \quad (4.23)$$

and $f(8) = 1.74$. We calculate $\mathbf{C}_{\ell\ell'}^{\text{window}}$ for 1000 different sets of independent beam maps, and take their mean.

$\mathbf{C}_{\ell\ell'}^{\text{window}}$ is the covariance of the beam window function (B_ℓ^2), however we want the covariance of B_ℓ . Applying standard error propagation formulae, we have

$$\mathbf{C}_{\ell\ell'}^{\text{data}} = \frac{\langle \mathbf{C}_{\ell\ell'}^{\text{window}} \rangle}{4B_\ell B_{\ell'}}. \quad (4.24)$$

At this point, we need to account for uncertainty in our AGN spectral index and bandpasses. We assign the spectra index an uncertainty of $\sigma(\alpha_{\text{AGN}}) = 0.16$. This is two times the largest uncertainty in the spectral index of any AGN. To estimate the uncertainty this adds to the beam, we calculate the change in B_ℓ assuming the average AGN spectral index is

$\alpha = \bar{\alpha}_{\text{AGN}} \pm \sigma(\alpha_{\text{AGN}})$. We then create the covariance matrix from such an uncertainty, and add it to the covariance matrix calculated from the data. That is,

$$\Delta B_\ell = B_\ell(\alpha = -0.622 - 0.16) - B_\ell(\alpha = -0.622 + 0.16) \quad (4.25)$$

$$\mathbf{C}_{\ell\ell'}^{\text{AGN spectrum}} = \Delta B_\ell \Delta B'_\ell. \quad (4.26)$$

$$\mathbf{C}_{\ell\ell'}^{\text{beam}} = \mathbf{C}_{\ell\ell'}^{\text{data}} + \mathbf{C}_{\ell\ell'}^{\text{AGN spectrum}} \quad (4.27)$$

$\mathbf{C}_{\ell\ell'}^{\text{beam}}$ is the final covariance of the beam.

We also produce a binned radial profile of the beam. This is not currently used in any science analysis, but can be useful for gaining an intuitive understanding of the beam. We first azimuthally average every stitched map (including those created from varying input parameters), with bins that are 0'.1 wide, then take their mean. Finally, we apply corrections for the planet disk, and the AGN spectrum (in the future, these should be done at the map level).

4.6 Systematic Tests

Since the beam measurement is such a broadly used calibration product, and directly affects our power spectra, we have put this process through a battery of systematics tests. These tests are either based on the simulations described in §4.2.3, or on differences between different data products.

4.6.1 Pipeline Tests

The pipeline tests us our simulations to make sure we recover the simulation inputs.

Beam Pipeline

To test for bugs and inconsistencies in our pipeline, we run a simulation without any noise. However, we also need to inject several systematics into the input maps. First, we use the our model of the beam dependence on observing frequency to expand or contract the input beam, depending on the source. For the planet maps, we assume a RJ spectrum ($\alpha = 2$). For the AGN maps, we use the mean spectral index of the 8 AGN used to measure the real beams ($\alpha = -0.622$) The planet maps are then convolved with a disk of radius 8''.55. Finally, we apply the time constants and time offset. This is done in map space (rather than the time domain, where the effect natively lives) to save time. For each map, we construct the effective map-space filter (equation 4.1), using the measured time constants and source declinations. For the AGN maps, the time constants are the same between maps, but the declination varies. For the planet maps, we make separate left- and right-going maps, using the appropriate transfer function. Each planet map is assigned the time constants measured from one real observation of Saturn. The time offset is fixed for all maps. Finally, each map



Figure 4.12: The fractional error on the beam recovered from a simulation. The dominant effect is a multiplicative scaling, which is degenerate with absolute calibration (and much smaller than our uncertainty on the absolute calibration). The variation with multipole ℓ is $\mathcal{O}(10^{-4})$, and much smaller than the beam uncertainties.

is scaled to match the observed amplitude of the corresponding source (the planet maps are all given different amplitudes, based on the amplitudes of the different Saturn observations). To determine the true beam, we simply take the power spectrum of one of the input maps, and remove the pixel window function.

While this is a relatively thorough application of the systematics, it does have a few deficiencies. First, there are some systematics we have not simulated. In particular, we have not included the saturation effects in the planet maps. However, the methodology used here is primarily intended to remove the effects of the saturation, so ignoring it in the simulations should have little effect. Second, we use the same models (and the same code) to inject and remove the systematics. Therefore, any inaccuracy in the model is not reflected in these simulations. For example, modeling the individual detector time constants as a single mean time constant for all the detectors in one band is somewhat inaccurate. It would be more accurate to model the time constant transfer function as the mean of many single-pole filters with the individual detector time constants. However, that still assumes that the detectors behave as single-pole filters in the time domain, which may not be exactly true. Nonetheless, we believe that these models are close enough (and we account for their inaccuracies in the uncertainty).

We run these maps through the beam pipeline with as few modifications as possible. In fact, we only have to remove the treatment of pointing jitter in the pipeline. Nearly all of the processing to account for systematics is invertible (the one exception is the masking applied to the receding portion of the planet scans). Therefore, we should recover the input beams extremely accurately. In fact, we do. Figure 4.12 shows the fractional change between the input and recovered beams. These are dominated by a global scaling, which is degenerate with absolute calibration (and much less than the uncertainty on our absolute calibration).




Figure 4.13: The scatter in 100 simulated beams (black dashed line), and the estimated noise from each simulation (blue lines). The black line falls within the expected scatter, indicating the the covariance calculation we use is unbiased.

This likely comes from floating point arithmetic discrepancies. The beam normalization is calculated by interpolating the complex beam window function at the pivot ℓ . The recovered B_ℓ is the square root of the real part of the beam window function, so some discrepancy is not surprising. The remaining variation is at the level of a few parts in 10^4 , which is much smaller than the final beam uncertainties. We do not fully understand the difference between 150 GHz and the other bands. However, this only appears when the time constants and offset are included in the simulations, which allows us to partake in some motivated speculation. The 150 GHz detectors are most sensitive to additional loading, and have the longest time constants (this comes about from operating near their saturation power). Therefore, the difference between the time constants in the Saturn observations and the AGN observations is largest at 150 GHz. Due to the information loss from masking the receding portions of the Saturn scans, we are forced to recover information in an iterative fashion, which is not exact. These facts most likely conspire to create a small mismatch between the Saturn maps and the AGN maps at 150 GHz, that couples into the beam as the long wavelength oscillation we see here.

Covariance

To test the covariance calculation, we run a set of simulations with white noise added to the Saturn and AGN maps. The white noise amplitude set to roughly match the SNR of the individual AGN and Saturn maps, although this is not particularly important for this test. For each simulation, we create 9 Saturn maps and 8 AGN maps, with their amplitude scaled to match the observed amplitudes in the data maps. We do not apply any of the systematics described in the previous section, since we want the only noise contribution to be from the white noise. We then run a minimal version of the beam pipeline, that consists of the background subtraction, stitching, beam estimation, and covariance calculation only. We perform 100 of these simulations, and compare the scatter in the recovered beams with the

covariance estimated from each simulation. This is shown in Figure 4.13, where we can see that the covariance estimates generally agree with the scatter in the recovered beam.

4.6.2 Data-based Tests

While the simulation-based tests are effective for catching implementation errors in the pipeline, they do not test whether the models we are applying are a good representation of the data. In this sections, we show two further tests based on differences between subsets of the data that directly test the models. These are similar in spirit the the null tests commonly used in a power spectrum analysis, but we do not have sufficient information to construct rigorous passing criteria.

4.6.3 Elevation Dependence of AGN

As described in §4.3.2, we account for the elevation-dependent effects of the time constant and time offsets. Without the corrections, we would expect to see features along the scan direction when comparing AGN maps at different declinations. After the map processing, we expect those differences to be consistent with zero (or at least significantly reduced).

To test this, we take the difference between two sources, one at declination $\delta = -57.6^\circ$, the other at $\delta = -47.5^\circ$. Since they have different intrinsic fluxes, we must normalize their amplitudes before taking the difference. We do this by fitting for a single multiplicative amplitude. Otherwise, there are no free parameters (the time constants and offset are fixed by external measurements). The differences (both before and after processing) are shown in Figure 4.14. After processing, there is no detectable difference between these sources.

4.6.4 Agreement Between Planet and AGN Maps

Our AGN and planet observations have several significant differences:

- The planet observations occur with much higher loading on the detectors, which makes their time constants longer.
- We use a different on-sky scan speed
- Planets have rising spectra (Saturn follows a Rayleigh-Jeans spectrum in our bands), compared to the falling spectra of AGN.
- Saturn is not point-like.

However, we expect to find the the same underlying beam in both the AGN and planet maps. We can use the difference between the AGN and planet maps (with an amplitude scaling as a free parameter) to see if our processing is making the AGN and planet maps more similar.

The expectations in this case are somewhat complicated. Some of they systematic effects we account for partially cancel each other. For example, accounting for the elevation

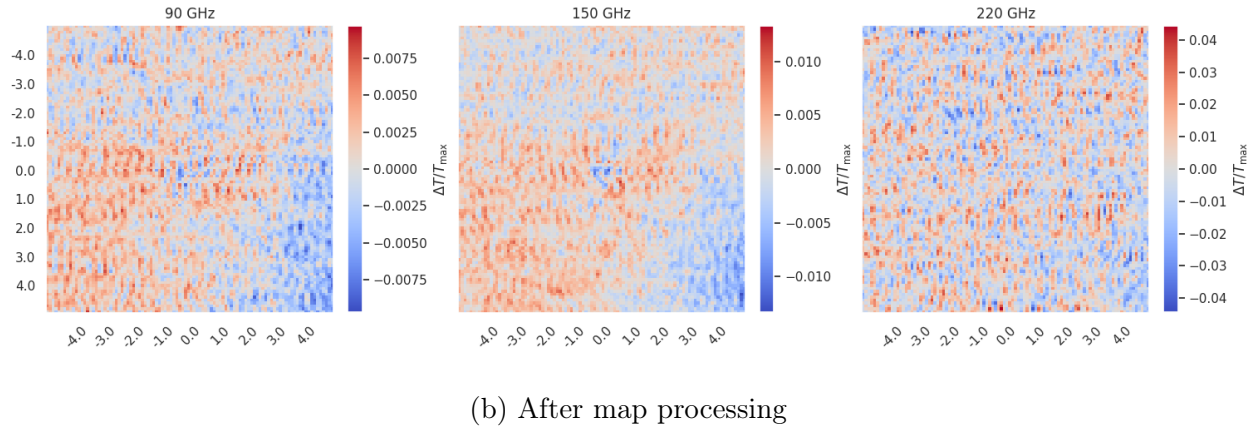
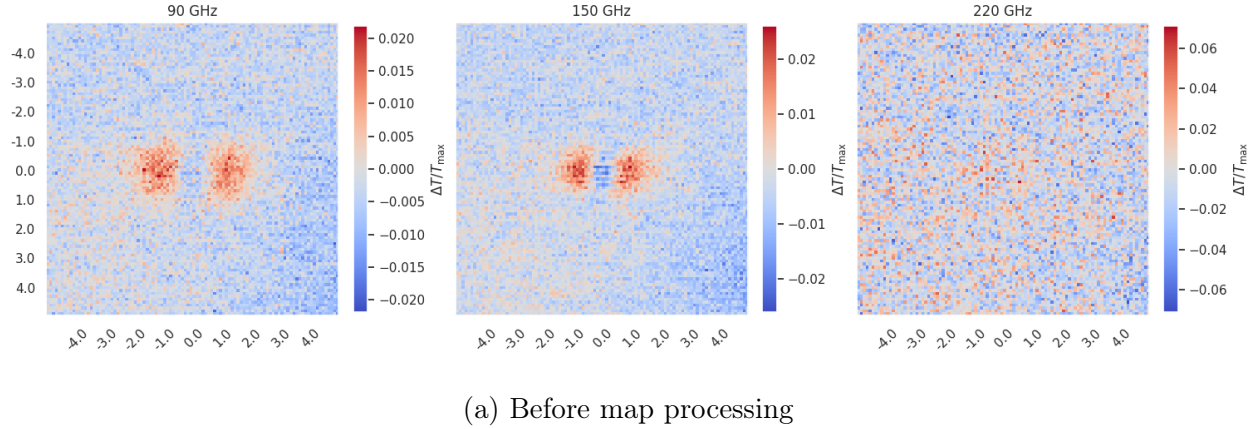




Figure 4.14: The difference between two AGN maps at different declinations (a source at $\delta = -47.5^\circ$ minus a source at $\delta = -57.6^\circ$). The x and y scales are marked in arcminutes from the source center, and the color scale is in units of the peak brightness of the source. Panel a is before accounting for the time constants and time offset, while b is after. Before processing, the higher declination (lower elevation) source is broadened in the scan direction relative to the lower declination source. This effect does not appear at 220 GHz because the SNR of these sources is insufficient at 220 GHz. After processing, the differences have disappeared.



(a) Before map processing



(b) After map processing

Figure 4.15: The map of one of the brightest AGN minus a coadd of the Saturn maps. The x and y scales are marked in arcminutes from the source center, and the color scale is in units of the peak brightness of the source. Panel a is before applying the map processing steps (except recentering), while b is after all the processing steps. The dashed circles show the fiducial stitching region. Inside the inner circle, the stitched beam map uses only the AGN; outside the outer circle, the stitched maps use only Saturn; the region between the circles is the transition region. The residuals are significantly reduced at 90 and 220 GHz, but are nearly unaffected at 150 GHz.





Figure 4.16: The SPT-3G beams (B_ℓ) at 90, 150, and 220 GHz (top panel), and their fractional uncertainties (bottom panel). The beams are normalized such that $B_{800} = 1$. The uncertainty shown here is $\sqrt{C_{bb}}$, i.e. the square root of the diagonal of the covariance matrix. This uncertainty is for qualitative purposes only, and any formal analysis must use the full covariance matrix.

dependence of the time constants and time offset makes Saturn narrower in the scan direction, while accounting for using only the approaching scans broadens Saturn in the scan direction. Therefore, we cannot separate each systematic, and expect the agreement between Saturn and the AGN to improve. However, the final maps should agree much better than the unprocessed maps. We show the difference between the coadd of Saturn maps and the brightest AGN (at $\delta = -51.0^\circ$) in Figure 4.15. As in the previous section, we must normalize each source for its intrinsic brightness. For this set of three maps, there are a total of five free parameters: 3 amplitudes to account for the difference in brightness, and the 2 background parameters (a slope and offset, which are applied across all three bands). All the other parameters of the map processing are fixed by external measurements. Therefore, it is somewhat unsurprising that we do not get perfect agreement between Saturn and this AGN. We expect some disagreement at the core ($\lesssim 0'.5$), due to the detectors saturating on Saturn. However, the residuals extend far beyond that core. Therefore, we set the inner stitching radius outside the majority of the residual, and include the change in the beam with variation of the stitching radii in the uncertainty.



(a) Radial profiles of the beams.

(b) Fractional area of the beam at radii larger than r .

Figure 4.17: Panel a shows the radial profiles of the beams. The high-frequency noise in the radial profiles is due to our corrections for the planet disk and source spectrum. Panel b shows the fraction of the beam area at radius $> r$. This is the reverse cumulative distribution function of the beam area, and is useful as an intuitive reference point for the sidelobe power on a given scale.

4.7 Results and Discussion

We present the final SPT-3G beams and their fractional uncertainties in Figure 4.16. The beams and uncertainty are calculated using all Saturn maps, AGN maps, and parameter variations.

4.7.1 Radial Profiles

Figure 4.17 shows the radial profiles of the beams. The 90 GHz radial profile shows a distinct non-Gaussian feature at $\sim 2.5'$. This is due to the greater effect of the Lyot stop at 90 GHz. 220 GHz has a larger response than 150 GHz at large radii ($r \gtrsim 6'$), which is most likely due to Ruze scattering from imperfections in the primary mirror. This is also visible in the fractional beam area as a function of radius (Figure 4.17b), which shows several times more sidelobe response at 220 GHz than 150 GHz. We fit a Gaussian to each radial profile, and find the full width at half maximum (FWHM) is ($1'.54$, $1'.16$, $0'.99$) at (90, 150, 220) GHz. A more thorough analysis of these radial profiles is ongoing (see §4.8).

4.7.2 Uncertainty Contributions

The covariance calculation described in §4.5 cannot be used to isolate the contribution of a single systematic, unfortunately. Since it operates on cross spectra between different Saturn




Figure 4.18: The most significant contributions to the beam uncertainty. Saturn and AGN refer to variations between the different Saturn and AGN maps, respectively. α_{AGN} refers to uncertainty in the spectral index and detector bandpasses. r_{inner} refers to the inner stitching radius (where we start to use the Saturn maps). BG r_{inner} refers to the inner radius of the region used to fit the background model. The contributions from varying the outer radii of the stitching and background fit regions, and the uncertainty on the time constants and time offset are not shown.

and AGN maps, it is impossible to remove the uncertainty due to noise in the maps, and variations between Saturn observations and different AGN. However, we can approximate the contribution of each uncertainty term for the purposes of understanding their relative contributions to the total uncertainty. To estimate the uncertainty from varying each input parameter, we first calculate the mean beam for every parameter value (this is only two beams per parameter). We then approximate the uncertainty from that parameter taking the standard deviation of the beams over the different parameter values. For the uncertainty caused by differences between the AGN and Saturn maps, we have to adopt a slightly more complicated procedure. As noted above, it is impossible to isolate a single AGN or Saturn map. Therefore, we select all the cross spectra that are made from a single AGN. For each AGN, we then construct the mean beam from those cross spectra, and take their standard deviation. We repeat the same process for each Saturn observation, as well.

The most significant contributions to the uncertainty are shown in Figure 4.18. On the largest scales, variations in the Saturn maps dominate the uncertainty. This is from atmospheric noise in the Saturn maps. The contribution from the uncertainty in the spectral index and detector bandpasses is primarily significant at 90 GHz. This is due in part to the shape of the 90 GHz bandpass (it is the least uniform), and the higher fractional bandwidth. Finally, the contribution from varying the inner stitching radius shows an interesting band dependence. This uncertainty comes from discrepancies between the observed beam profile in the Saturn and AGN maps. As a reminder, Figure 4.15b shows the difference between the brightest AGN and a coadd of the Saturn maps, with the fiducial stitching region. Based just on the magnitude of the residuals, we would expect 90 GHz to receive the smallest

contribution from varying r_{inner} (as it does), and 150 GHz to receive the largest (it does not). However, we are more sensitive to the change in residual at r_{inner} (the inner dashed circle in Figure 4.15b). This may indicate that we are choosing r_{inner} sub-optimally at 150 and 220 GHz. If we increased its fiducial value, we would decrease the uncertainty contribution from varying r_{inner} . However, we would increase the contribution from map noise in the AGN maps (since the AGN maps have lower SNR than the Saturn maps, particularly at the higher frequencies).

4.8 Ongoing Work

In this section, we discuss some ongoing work to potentially improve the beam measurement.

4.8.1 Towards a Physical Beam Model

As mentioned in §4.1.4, there are a few deficiencies in our model of frequency dependence in the beam. We are working on developing a physically motivated model for the beam that more appropriately captures the frequency dependence. For now, this model only includes the main lobe of the beam (not the sidelobes). We consider four components in the model:

1. Gaussian illumination pattern of the primary mirror by the detector lenslets.
2. Truncation of that illumination pattern by the Lyot stop.
3. Gaussian broadening of the diffracted beam due to a frequency independent geometric aberration.
4. Averaging the frequency dependent beam over the detector bandpass weighted by a CMB spectrum

We model the illumination pattern for a δ -function bandpass on the primary mirror as a truncated Gaussian, then use the Fraunhofer diffraction equation to generate the beam on the sky. We then integrate over the real bandpass function, and apply the geometric aberration.

The illumination pattern on the primary is

$$P(r, \nu) = \exp\left(\frac{-r^2}{2(\sigma_0 \frac{\nu_0}{\nu})^2}\right) \Pi(R) \quad (4.28)$$

where σ_0 is the width of the Gaussian beam from the lenslets at frequency ν_0 , ν_0 is a fiducial frequency and $\Pi(R)$ is a disk of radius R :

$$\Pi(R) = \begin{cases} 1 & r \leq R \\ 0 & r > R \end{cases} \quad (4.29)$$

For SPT-3G $R \approx 3.75$ m. Note that $P(r)$ is the illumination pattern for the electric field. In the terminology of §4.1.4, $\sigma_{\text{diffraction}}$ is roughly equivalent to $\sigma_0\nu_0/\nu$ (although the model is not identical, and we do expect somewhat different values).

To get the beam on the sky, we apply the Fraunhofer diffraction equation.

$$U(\rho, z, \nu) = 2\pi \int_0^\infty P(r, \nu) J_0\left(\frac{2\pi r \rho}{c/\nu z}\right) r dr \quad (4.30)$$

$$U(\theta, \nu) = 2\pi \int_0^\infty P(r, \nu) J_0\left(\frac{2\pi r \theta}{c/\nu}\right) r dr \quad (4.31)$$

where z is the distance to the image plane, ρ is the distance from the origin on the image plane, and J_0 is the Bessel function. In the second line, we have taken the small angle approximation ($\rho/z = \theta$) to put this in terms of angles on the sky. The beam on the sky is the square of the electric field,

$$B_{\text{ideal}}(\theta, \nu) = U(\theta, \nu)^2. \quad (4.32)$$

The ideal subscript indicates that this is without any geometric aberration.

To get the beam for our wide bandpass function, we need to integrate $B_{\text{ideal}}(\theta, \nu)$ over the bandpass and source spectrum. One might think we should integrate U over the bandpass and source spectrum, since U represents the electric field. However, at any significant separation between frequencies, the electric field phases will decohere very quickly. Our bandpasses and observing spectra are slowly varying, so any adjacent frequencies where phase coherence matters have very nearly the same illumination pattern. Therefore, the beam over a wide bandpass function $t(\nu)$ when observing a source with a spectrum given by $I(\nu)$ is

$$B_{\text{ideal}}(\theta) = \frac{\int A(\nu) \Omega_s I(\nu) t(\nu) B_{\text{ideal}}(\theta, \nu) d\nu}{\int A(\nu) \Omega_s I(\nu) t(\nu) d\nu} \quad (4.33)$$

where $A(\nu) = 2\pi \int P(r, \nu) r dr$ is effective collecting area of the SPT at frequency ν , and Ω_s is the solid angle of the source. While these integrands are quite complicated, this is just the mean of $B_{\text{ideal}}(\theta, \nu)$ weighted by the total power received at each frequency. Finally, we need to account for geometric aberration, which we model as a convolution with a Gaussian:

$$B(\vec{\theta}) = B_{\text{ideal}}(\vec{\theta}) \otimes \exp\left(-\frac{\vec{\theta}^2}{2\sigma_{\text{geom}}^2}\right) \quad (4.34)$$

where \otimes indicates convolution, and we have replaced θ with $\vec{\theta}$ to emphasize that this convolution must be done in the 2D plane.

This model has two free parameters: σ_0 , and σ_{geom} . We can simultaneously fit these using the radial profiles of the beams at low radii. We restrict the fit to low radii to ensure that the beams are dominated by the main lobe.

4.8.2 Polarized Beam Measurements

Thus far, all SPT analyses have assumed that the response to polarized and unpolarized sources is identical. However, there is some evidence that this is not the case (Ge et al. *in prep*, Chou et al. *in prep*). In particular, it seems that the side lobes may have a lower polarization efficiency. Physically, we might expect this from either scattering off the gaps between the panels that make up the primary mirror, or any optical path that involves multiple scatterings inside the camera. In both of our upcoming papers, we explore the possibility of allowing the side lobes to have a lower polarization efficiency than the main lobe. We do so by first modeling the main lobe (Ge et al. uses the model presented here). We then assume that the side lobes are whatever remains in the beam after subtracting the main lobe, and allow them to have a free amplitude in the analysis of polarization spectra.

We are also exploring the possibility of measuring the polarized beams directly. This is quite challenging, as there are no point-like sources with polarized flux similar to the AGN we use. We have found two sources with $\mathcal{O}(100)$ mJy of polarized flux. While these are much dimmer than the few Jy AGN we use in this work, the polarization noise is also lower (both the atmosphere and the CMB have low polarization fractions). We are in the process of planning dedicated observations of these two sources. In addition, we are beginning to explore the possibility of flying our own bright polarized source at the South Pole. This is challenging, because we need to put the source in the far field of the optics (which is at least a few km from the SPT). We also need the source to be at a reasonable elevation ($\gtrsim 25^\circ$) above the horizon to avoid excess atmospheric loading. Finally we need the source's position to be stable at a scale much smaller than the beam scales³ we need to measure ($\lesssim 30''$). At 3 km, this requires position stability of ~ 0.5 m.

4.8.3 Time Constants

Until very recently, our science maps have ignored the effects of the time constants and time offset. Since they very nearly cancel, this is acceptable. However, future data products will remove both effects from each detector in the time domain. To analyze these data products, we will need a new beam measurement, with the time constants and time offset removed at the TOD level, as well. In principle, this makes the beam pipeline much simpler: there are several processing steps that only exist because we have to account for the different effects of the time constants and time offset with observing elevation.

4.9 Conclusion

The telescope beam is convolved with the sky, so it affects our observations at all scales. In this chapter, we described the measurement of the SPT-3G beam using bright sources, and confirmed that our beams fall in the 1-2 arcminute range. We also make a careful accounting

³In principle, we only need to measure the side lobes, on scales of $> 2'$.

of the uncertainty in the beam, and the different contributions to that uncertainty. These are critical elements to a measurement of the CMB power spectrum, where they are used to both correct for the effects of the telescope optics, and properly account for uncertainties in the final spectra (e.g. Dutcher et al. (2021) and Ge et al. *in prep.*). In the final chapter, we also show how knowledge of the beam is important for finding resolved objects.

Chapter 5

Galaxy Clusters Selected via the Sunyaev-Zel'dovich Effect in the SPTPol 100-Square-Degree Survey

This chapter describes the construction of a catalog of galaxy clusters using the tSZ signal. The catalog was previously published as Huang et al. (2020). This work uses data from the SPTPol camera, which is the predecessor the SPT-3G (which has been the primary focus of this work). The cameras are technically very similar: the detectors are TESes, read out with a frequency-domain multiplexing system; however, SPTPol had only two observing bands, centered at 95 and 150 GHz. Since they are mounted on the same telescope, we achieve approximately the same angular resolution. The calibration, mapmaking, and beam measurement procedures are equivalent in spirit, but vary in their details. For instance, the HII region calibration is done in the time domain for the SPTPol data.

In §1.3.3, we have described the thermal Sunyaev-Zel'dovich effect as a spectral change to the CMB induced by clouds of relativistic electrons. Galaxy clusters provide a plentiful supply of relativistic electrons in the ICM. As mentioned, we can use the unique spectrum of the tSZ to isolate the galaxy clusters with maps across multiple observing frequencies.

Compared to traditional methods of finding galaxy clusters (either using the X-ray emission from the ICM, or overdensities in optical/infrared galaxy catalogs), the tSZ is less effective at finding low-mass, low-redshift clusters. On the other hand, the traditional probes rely on intrinsic emission from the galaxy cluster, which is subject to cosmological dimming. At infrared wavelengths, the K correction compensates for most of the cosmological dimming (see e.g. Gonzalez et al., 2019, and references therein). However, infrared surveys are affected by projection effects and have higher mass-observable scatter than tSZ-selected clusters. tSZ surveys are able to provide approximately mass-limited catalogs across the full redshift range, with the maximum redshift set by the increasing rarity of high-mass clusters at high redshift. At the highest redshifts, more massive clusters have not had time to form.

As a reminder from Chapter 1, the magnitude of the tSZ is given by

$$\Delta T_{\text{tSZ}} = T_{\text{CMB}} f(x) \int n_e \frac{k_B T_e}{m_e c^2} \sigma_T dl \quad (5.1)$$

$$\equiv T_{\text{CMB}} f(x) y \quad (5.2)$$

(Sunyaev and Zel'dovich, 1972), where $x \equiv h\nu/k_B T_{\text{CMB}}$ is the dimensionless frequency, n_e is the electron number density, T_e is the electron temperature, σ_T is the Thomson cross section, k_B is the Boltzmann constant, c is the speed of light, and the integral is along the line of sight. f describes the frequency dependence of the tSZ:

$$f(x) = \left(x \frac{e^x + 1}{e^x - 1} - 4 \right) (1 + \delta_{\text{rc}}) \quad (5.3)$$

where δ_{rc} is a relativistic correction which is several percent for massive clusters with $T_e > 5$ keV (Nozawa et al., 2000). Below 217 GHz, $f < 0$, which leads to the CMB decrement noted above. Since the tSZ is a spectral effect, it is independent of the distance to the galaxy cluster. Furthermore, y is proportional to the thermal energy integrated along the line of sight. The total tSZ signal of the cluster (Y_{sz}), defined as the integral of y over the transverse extent of the cluster, is mathematically equivalent to the total thermal pressure of the cluster. It is expected to be tightly correlated with cluster mass (Motl et al., 2005). This makes the tSZ an effective tool for building galaxy cluster catalogs for cosmological analyses (see Carlstrom et al. 2002 for a review).

We present a catalog of galaxy clusters found in one of the deepest high-resolution CMB maps currently available. This pushes the cluster detection threshold to lower mass, and represents the first of several catalogs that will use data of similar or greater depth. By decreasing the mass threshold, we have also increased the effective redshift limit. The catalog consists of 89 galaxy cluster candidates, of which 81 have been confirmed using optical and near-infrared data. Of the confirmed clusters, 29 are presented for the first time. This chapter is organized as follows: in §5.1, we describe the mm-wave data and processing; §5.2 discusses the cluster search methodology and characterization; §5.3 describes the optical and infrared follow up used to confirm cluster candidates and estimate their redshifts; the catalog is presented and compared with other cluster catalogs in §5.4; finally, we discuss conclusions and upcoming work in §5.5. Selected data reported in this work (and any updates to the clusters in this catalog) will be available at <http://pole.uchicago.edu/public/data/spt-sz-clusters>.

5.1 Millimeter-wave Data

The cluster sample presented in this work was derived from two-band millimeter-wave (mm-wave) data taken with the SPTPol camera on the South Pole Telescope (SPT, Carlstrom et al. 2011). In this section, we describe the observations and the data processing used to produce mm-wave maps.

5.1.1 Telescope and Observations

The SPT is a 10-meter-diameter telescope located within 1 km of the geographic South Pole, at the National Science Foundation’s Amundsen-Scott South Pole Station, one of the premier sites on Earth for mm-wave observations. The 10-meter aperture results in diffraction-limited angular resolution of ~ 1 arcmin at 150 GHz, which is well-matched to the angular size of high-redshift galaxy clusters. Combined with the redshift-independent surface brightness of the tSZ, this makes the SPT a nearly ideal instrument for discovering high-redshift clusters through the tSZ. The first camera on the SPT was used to conduct the 2500-square-degree SPT-SZ survey. The resultant cluster catalog contains 42 clusters above $z = 1$ with typical masses of $3 \times 10^{14} M_{\odot}$, the most extensive sample of massive, high-redshift systems selected via ICM observables in the literature (Bleem et al., 2015b, hereafter B15).

The SPTPol camera (Austermann et al., 2012), installed on the telescope in 2012, consists of 1536 detectors, 1176 configured to observe at 150 GHz, and 360 configured to observe at 95 GHz. The first SPTPol observing season and part of the second season were spent observing a roughly 100-square-degree field centered at right ascension (R.A) $23^{\text{h}}30^{\text{m}}$, declination -55° (bounded by $23^{\text{h}} < \text{R.A.} < 24^{\text{h}}$, $-60^{\circ} < \text{decl.} < -50^{\circ}$, the same definition as the RA23H30DEC–55 field in B15). This is referred to as the SPTPol 100d field. It was actually observed as two separate sub-fields: a “lead” field and a “trail” field. This strategy was adopted to mitigate the effects of ground-based signals. No ground signals were ever detected, so we analyze both fields as one. Observations occurred between March and November of 2012, and in March and April of 2013. Of approximately 6600 observations of the field, 6150 observations at 95 GHz and 6040 observations at 150 GHz are used in this work. Maps made from the weighted sum of all SPTPol observations of this field have rough noise levels (in CMB temperature fluctuation units) of $6.5 \mu\text{K-arcmin}$ at 150 GHz and $11.2 \mu\text{K-arcmin}$ at 95 GHz, a factor of 3-4 lower than the typical noise levels for the SPT-SZ maps used in B15, but over a much smaller area.¹ We thus expect a lower total number of clusters compared to B15 but a higher density and a larger fraction of systems at high redshift.

5.1.2 Mapmaking

The mapmaking for this work is performed as described in §3.3 The data from every individual observation of the field at each observing frequency is calibrated, filtered, and binned into 0.25 arcmin map pixels using the Sanson-Flamsteed projection (Calabretta and Greisen, 2002). The equivalent Fourier-domain filtering from each of these steps is a high-pass in the scan direction with a cutoff of angular multipole $\ell \sim 400$, a low-pass in the scan direction with a cutoff of $\ell \sim 20,000$, and an isotropic high-pass with a cutoff of $\ell \sim 200$, respectively. The low-pass filter is set well above the resolution of the SPT, and is only used to prevent aliasing of high-frequency time-domain noise. In the two high-pass filtering steps, compact

¹The SPTPol camera is sensitive to both the total intensity and the polarization of incoming radiation, but we only use the total intensity information in this work.

emissive sources with flux density greater than 6.4 mJy at 150 GHz are masked to avoid filtering artifacts (the same sources were masked in B15).

5.1.3 Beams and Calibration

For SPTPol, we measure the instrument response as a function of angle (i.e., instrument beams) on planets, particularly Mars. The main lobe of the beam at both frequencies is well approximated by an azimuthally symmetric Gaussian, with full width at half maximum (FWHM) $\sim 1'.6$ and $\sim 1'.1$ at 95 and 150 GHz respectively. Uncorrected shifts in absolute pointing between individual-observation maps leads to a further smearing of the beam in coadded maps, resulting in final beam FWHM of $\sim 1'.7$ and $\sim 1'.2$ at 95 and 150 GHz. In the matched filter described in the next section, we use a Gaussian approximation to the beam at each frequency.

The relative calibration among detectors in the SPTPol focal plane, and the absolute calibration used in this work, are derived using a combination of detector response to an internal thermal source and to the Galactic HII region RCW38. We use a calibration procedure identical to that described in Schaffer et al. (2011), and we refer the reader to that work for details. We apply an additional calibration factor to our simulations to account for any discrepancies between the RCW38-based calibration and absolute calibration against SPT-SZ maps (see §5.2.2).


5.2 Cluster Extraction and Characterization from mm-wave Data

In this section, we summarize the procedure used to extract the cluster signal from maps of the microwave sky. The method used in this work is very similar to that of previous SPT publications. For a more detailed treatment, see Williamson et al. (2011), Reichardt et al. (2013), Vanderlinde et al. (2010, hereafter V10), and B15.

5.2.1 Cluster Extraction

The SPTPol maps contain signals from several classes of sources, each of which has its own spatial and spectral characteristics. We first describe the spatial behavior of these sources. At large angular scales, our maps are dominated by the CMB. On the smallest scales, power comes mostly from point sources, such as dusty galaxies and radio-bright sources. Galaxy clusters populate the intermediate regime, between the large scale CMB fluctuations and the point sources.

The intensity of these signals also varies by observing frequency. The amplitude of the CMB and kSZ is preserved across observing frequencies, because the maps are calibrated in CMB temperature units. Radio-loud active galactic nuclei (AGN) appear with a falling spectrum, while dusty galaxies have a rising spectrum. Finally, the tSZ spectrum is given in



(a) 95 GHz map

(b) 150 GHz map

Figure 5.1: Panels a and b each show a 3° -by- 3° cutout at 95 and 150 GHz, respectively, with several galaxy cluster detections circled in orange. These cutouts are taken from the larger map used in this work, which was produced as described in §5.1.2. The black dashed region is shown in Figure 5.2, after the matched filter is applied.

equation 5.3. For a given cluster, the magnitude of the tSZ is greater at 95 GHz than 150 GHz. We model the map as follows:

$$T(\vec{\theta}, \nu_i) = B(\theta, \nu_i) * \left[\Delta T_{\text{sz}}(\vec{\theta}, \nu_i) + N_{\text{astro}}(\vec{\theta}, \nu_i) \right] + N_{\text{noise}}(\vec{\theta}, \nu_i) \quad (5.4)$$

where $\vec{\theta}$ is the position on the sky, ν_i is the observing band, ΔT_{sz} is the tSZ, as defined in equation 5.1, N_{astro} is all other signals fixed on the sky (such as emissive point sources), N_{noise} is any noise term not fixed on the sky (such as instrumental noise), $B(\theta, \nu_i)$ represents the effects of the SPTPol beam, as well as filtering operations applied during mapmaking, and $*$ is the convolution operator.

Using the known spatial and spectral forms of the astrophysical noise terms described above, and the measured non-astrophysical noise (see §5.2.2), we construct a simultaneous spatial-spectral filter to optimally extract the cluster signal. The process is similar to that described in Melin et al. (2006). In the Fourier domain, the filter takes the form

$$\psi(\mathbf{l}, \nu_i) = \sigma_\psi^2 \sum_j \mathbf{N}_{ij}^{-1}(\mathbf{l}) f(\nu_j) S_{\text{filt}}(\mathbf{l}, \nu_j) \quad (5.5)$$

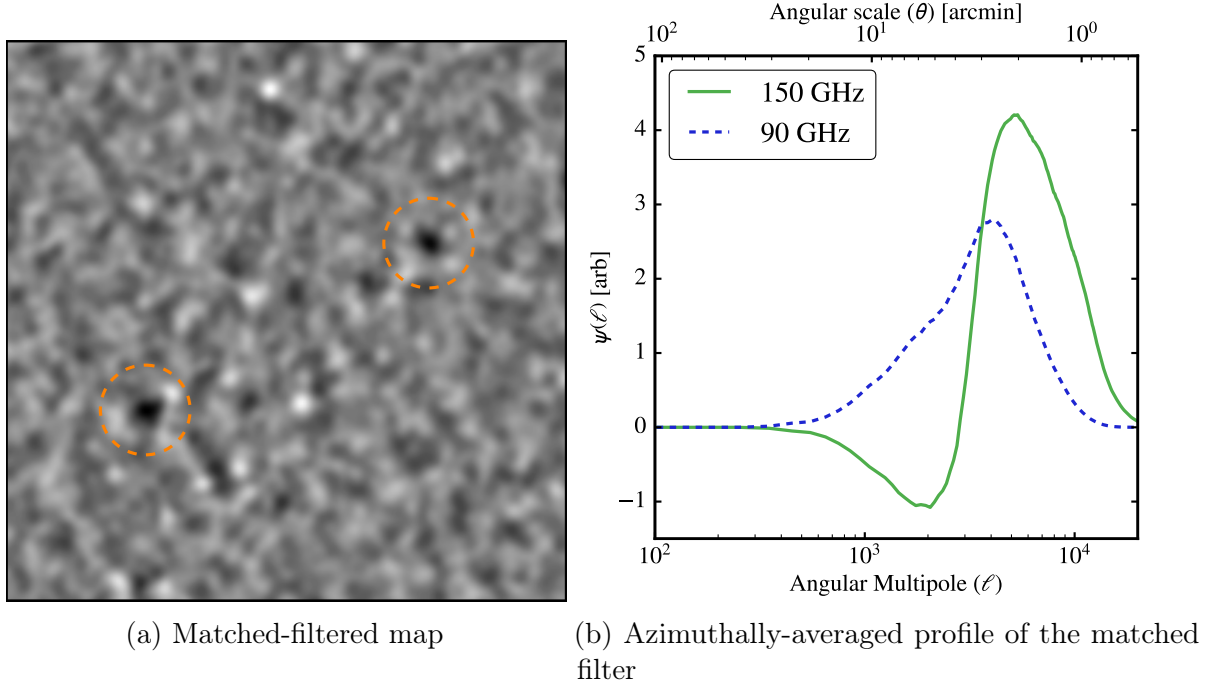


Figure 5.2: Panel a is the region of sky outlined with a black dashed line in Figure 5.1, after the matched filter is applied. This region is $50'0$ on each side. The filtering is described in §5.2.1, and this map has been filtered to optimally find objects with $\theta_c = 0'25$. It contains 2 detected galaxy clusters, SPT-CL J2323-5752 at $z = 1.3$, and SPT-CL J2320-5807 at $z = 0.56$, both of which are detected at a signal-to-noise-ratio greater than 5.0. A composite image of tSZ contours, optical, and infrared images for SPT-CL J2323-5752 is shown in Figure 5.4. Panel b shows the azimuthally-averaged spatial filter used to produce panel a for each observing band. Since the tSZ has a larger magnitude at 95 GHz, the 95 GHz filter remains positive over the entire range. The 150 GHz filter is negative at intermediate ℓ s to subtract CMB fluctuations. The peak of the 95 GHz filter is lower than the peak of the 150 GHz filter due to the relative noise in the 95 and 150 GHz SPTPol maps.

where σ_ψ^2 is the predicted variance of the filtered map, given by

$$\sigma_\psi^{-2} = \int d^2l \sum_{i,j} f(\nu_i) S_{\text{filt}}(\mathbf{l}, \nu_i) \mathbf{N}_{ij}^{-1}(\mathbf{l}) f(\nu_j) S_{\text{filt}}(\mathbf{l}, \nu_j) \quad (5.6)$$

The band-band, pixel-pixel covariance matrix is represented by \mathbf{N} , and S_{filt} is the Fourier transform of the search template convolved with $B(\theta, \nu_i)$. We use an isothermal projected β -model (Cavaliere and Fusco-Femiano, 1976) with β fixed at 1 for the SZ surface brightness template:

$$S = \Delta T_0 (1 + \theta^2/\theta_c^2)^{-\frac{3}{2}\beta + \frac{1}{2}}. \quad (5.7)$$

The normalization (ΔT_0) is a free parameter, and we search over a range of core radii (θ_c). V10 explored the use of more sophisticated models, but found no significant improvement.

The covariance matrix \mathbf{N} is constructed using model power spectra for the astrophysical terms, and measured noise properties for the remaining terms. The astrophysical portion is made up of the lensed CMB, point sources, and kSZ and tSZ from unresolved sources. The CMB power spectrum is calculated using the best-fit WMAP7 + SPT Λ CDM parameters (Komatsu et al., 2011; Keisler et al., 2011). The tSZ background is taken to be flat in $l(l+1)$ space, with the level taken from Lueker et al. (2010). The kSZ spectrum is taken from Shirokoff et al. (2011). While newer results are available for these spectra (e.g., George et al. 2015; Planck Collaboration et al. 2018), the matched filter is insensitive to the details of the spectra at this level. The combination of instrumental and atmospheric noise is measured from the data (see §5.2.2).

Before applying the matched filter, we mask point sources above 6.4 mJy (at 150 GHz) with a radius of 4 arcminutes from the source center. This reduces spurious detections caused by ringing around the brightest sources. We also apply an additional 8 arcminute veto around these sources, in which we reject any candidate objects. A small number of spurious detections are still included, which we remove via visual inspection.² After masking, the total search area is 94.1 deg².

We use 12 different filters, with θ_c evenly spaced between 0'.25 and 3'.0. After applying the filter in the Fourier domain, we transform back to map space to locate cluster candidates. Candidates are selected by their signal-to-noise ratio. Our observation strategy causes small variations in noise in both R.A. and declination. The overlap region between the lead and trail fields is somewhat deeper than the non-overlapping regions, and instrumental and atmospheric noise are larger closer to the equator. Our simulations (see §5.2.2) naturally include both noise variations. We account for the declination-dependence of the noise by splitting the field into 90' strips of constant declination and ignore the R.A.-dependence when calculating the signal-to-noise ratio of cluster candidates. Both variations are small (typically $\leq 5\%$), so this will not bias our mass estimation (see §5.2.4).

In each strip, we fit a Gaussian to the distribution of all unmasked pixel values within 5 σ of the mean. We use the standard deviation of the fitted Gaussian as our measure of noise. The pixel distribution is highly Gaussian within 4 σ of the mean, and the fitted standard deviation varies only a small amount when fitting all pixels within 2, 3, 4 or 5 σ of the mean. Therefore, we expect that using the Gaussian approximation of our noise distribution to be sufficiently accurate. This was changed from previous SPT analyses, which used the root mean square (RMS) of the strip. For each strip, we divide the filtered map by the noise, resulting in a signal-to-noise map. The updated method of measuring the noise typically results in a 6-10% increase in the signal-to-noise ratio at the locations of clusters. Galaxy cluster candidates are found using a **SExtractor**-like algorithm (Bertin and Arnouts, 1996) on each filtered map. For each candidate, we maximize the signal-to-noise ratio (ξ) over R.A., declination, and filter scale. A small region of the mm-wave maps is shown in Figure 5.1. The matched filter, and the results of applying it to this region are shown in Figure 5.2.

²Our visual inspection is statistically consistent with masking point sources above 3 mJy in simulations.

5.2.2 Simulations

In order to characterize our detection algorithm, we use simulated maps of the sky. Each map has several components:

CMB The Code for Anisotropies in the Microwave Background (CAMB; Lewis et al. 2000) is used to generate CMB spectra from the WMAP7+SPT best-fit Λ CDM parameters (Komatsu et al., 2011; Keisler et al., 2011). These are used to generate Gaussian random fields for the CMB realizations.

Radio Sources We generate point sources at random locations, with fluxes at 150 GHz based on the model from De Zotti et al. (2005). We assume 100% spatial correlation between 95 and 150 GHz, and a spectral index $\alpha_{\text{radio}} = -0.9$ (George et al., 2015) with Gaussian scatter of 0.47 (Mocanu et al., 2013).³ Sources have fluxes between 1.0 and 6.4 mJy. The upper cutoff was chosen to match the masking threshold used on the real data.

Dusty Sources Dusty sources are modeled as two Gaussian random fields: one for the Poisson contribution, and another for the clustered term. We again assume 100% spatial correlation between bands. We adopt the model from George et al. (2015) and use the best-fit values for spatial and spectral behavior from that work: the Poisson term is assumed to be of the form $C_\ell = \text{const.}$, and is normalized such that $D_\ell = \ell(\ell + 1) / (2\pi) C_\ell = 9.16 \mu\text{K}^2$ at $\ell = 3000$ in the 150 GHz band; the clustered term is normalized such that $D_{3000} = 3.46 \mu\text{K}^2$ at 150 GHz, and $D_\ell \propto \ell^{0.8}$. George et al. (2015) finds the best fit spectral indices to be $\alpha_{\text{clustered}} = 3.27$ and $\alpha_{\text{poisson}} = 4.27$.

tSZ N-body simulations are used to project tSZ halos onto simulated skies. Halo locations and masses are taken from the Outer Rim N-body simulation (Heitmann et al., 2019). We use the method described in Flender et al. (2016) to map halos to a location on the sky. The pressure profile described in Battaglia et al. (2012) is added at the location of each halo to create a map of the tSZ. The simulated maps contain contributions from halos with $M_{500c} > 6.25 \times 10^{12} M_\odot h_{70}^{-1}$, spanning a redshift range from $z = 0.01$ to $z = 3.0$.⁴ These maps contain a sufficient number of low-mass clusters to reproduce a power spectrum consistent with George et al. (2015). While the amplitude of the kSZ is similar to that of the tSZ, we neglect the kSZ because it is much smaller than other noise terms.

Noise The instrumental and atmospheric noise terms are not simulated. Instead, we use noise realizations calculated from the data. Noise realizations are generated by splitting the maps into equal-weight halves, and then differencing them. We split the data into

³We define the spectral index such that the specific intensity of a species i is $I_{\nu,i} \propto \nu^{\alpha_i}$.

⁴ M_{500c} is the mass within r_{500c} of the cluster center, where r_{500c} is the radius at which the average density of the galaxy cluster is 500 times the critical density at the cluster's redshift.

halves containing different maps for each noise realization, and normalize based on the number of input maps.

We do not account for any spatial correlation between simulated components. In particular, this leads to an underestimate of the number of tSZ decrements that are partially filled by radio sources. Bleem et al. (2020) performs an analysis of the radio intensity at the locations of tSZ-selected galaxy clusters, and finds that this is not a large effect. However, this work uses deeper mm-wave data, so the Bleem et al. (2020) result should be considered a lower limit to the radio contamination. The effects of lensing dusty sources and CMB on cluster catalogs was studied in Hezaveh et al. (2013). They conclude that lensing has little to no effect on the source density in tSZ-selected cluster surveys. Finally, we have not included the kSZ at all. However, we have tested its impact using the simulations from Flender et al. (2016), and find that it adds variance, but no bias. We find that the kSZ is responsible for $\sim 5\%$ of the intrinsic scatter in the scaling relation between signal-to-noise and mass (see §5.2.4).

These simulations are used to characterize both the expected number of false detections, and the normalization when calculating cluster mass. In order to properly account for the SPTPol beams and calibration, we first filter a map containing all the simulated components with the SPTPol filter transfer function and the beam. For this step, we use the measured beam instead of the Gaussian approximation that is used when constructing the matched filter. Then, we apply an overall calibration factor to the map for each band. This calibration factor is calculated via comparisons to calibrated SPT-SZ maps of the same field (see Keisler et al. 2015). Finally, the instrumental and atmospheric noise is added.

5.2.3 Sample Purity

We expect some fraction of our cluster candidates to be false detections caused by astrophysical and instrumental noise. In order to estimate the rate of false detections, we create a set of simulated sky maps containing all noise terms, but no tSZ signal. While, in principle, this means we are missing a noise term from sub-threshold galaxy clusters, the power in the tSZ is much smaller than other noise terms on the spatial scales that the matched filter is sensitive to. Furthermore, at current noise levels, chance associations of below-threshold clusters that overlap to form an apparently higher significance cluster are very rare. We search for clusters in these maps, using the procedure described above (including by-eye cleaning of spurious detections associated with bright point sources). The remaining detections are assumed to be representative of the false detections we find in the real data. At our cutoff of $\xi > 4.6$ we expect 91% purity (i.e. 8 false detections). See §5.4 for more details on the cutoff. The false detection rate is shown as a function of cutoff significance in Figure 5.3.




Figure 5.3: Simulated false detection rate, as calculated from approximately 1000 deg² of simulated sky (10 simulations of the sky area used in this work). This shows the density of false detections above a given significance (ξ_{\min}) on the left side, and the expected number of false detections for the entire 100 deg² survey area on the right. We expect 8 ± 2 false detections at our significance cutoff, $\xi_{\min} = 4.6$.

5.2.4 Galaxy Cluster Mass

We estimate masses for each cluster based on its detection significance (ξ) in SPT maps. Our method for estimating masses is described in detail in Reichardt et al. (2013), Benson et al. (2013) and Bocquet et al. (2019). In this section, we briefly describe the method.

For each galaxy cluster, we calculate the posterior probability for mass

$$P(M|\xi) = \left. \frac{dN}{dMdz} \right|_z P(\xi|M, z) \quad (5.8)$$

where ξ is the measured signal-to-noise ratio of the cluster, $\frac{dN}{dMdz}$ is the assumed mass function (Tinker et al., 2008), and $P(\xi|M)$ is the ξ -mass scaling relation. As discussed in V10, ξ itself is a biased mass estimator, so we introduce the unbiased SPT detection significance ζ :

$$\zeta \equiv \sqrt{\langle \xi \rangle^2 - 3}. \quad (5.9)$$

For $\langle \xi \rangle > 2$, this removes bias caused by the maximization over three degrees of freedom (R.A., declination, and filter scale) performed when producing the catalog. $\langle \xi \rangle$ is the mean significance with which a cluster would be detected in an ensemble of tSZ surveys. For each cluster, we assume that $P(\xi)$ is a Gaussian distribution with unit width, centered on the

measured signal-to-noise ratio, and marginalize over the scatter. Finally, we parameterize the relationship between the unbiased detection significance and cluster mass as

$$\zeta = A_{\text{sz}} \left(\frac{M_{500c}}{3 \times 10^{14} M_{\odot} h_{70}^{-1}} \right)^{B_{\text{sz}}} \left(\frac{E(z)}{E(0.6)} \right)^{C_{\text{sz}}} \quad (5.10)$$

with normalization A_{sz} , slope B_{sz} , redshift evolution C_{sz} . We assume log-normal scatter $\sigma_{\ln \zeta}$ on the $\zeta - M$ scaling relation.

As defined, A_{sz} is dependent on the noise levels in the maps. In order to use the same scaling relation for a variety of fields, we redefine $A_{\text{sz}} \rightarrow \gamma_{\text{field}} A_{\text{sz}}$. To find γ_{field} , we run a modified version of the cluster extraction algorithm. We filter 10 simulated maps, using all the components (noise and tSZ) described above. We also filter the maps containing only the tSZ (signal-only maps). Instead of running the source extraction on the signal-to-noise ratio maps, we divide the signal-only maps by the noise (calculated as described in §5.2.1). This gives us a direct measure of $\langle \xi \rangle$ instead of ξ , by removing the scatter associated with the noise terms, and reduces the number of simulations required. To construct a catalog, we maximize the signal-to-noise over the filter scale at the central pixel of each halo. We fit the scaling relation using all halos with $M_{500c} > 2 \times 10^{14} M_{\odot} h_{70}^{-1}$ and $z > 0.25$, to match the SPTPol cluster sample. Note that the size of the cluster on the sky becomes comparable to CMB fluctuations for galaxy clusters at $z \leq 0.25$, and our mass estimates may be systematically low at these redshifts.

In previous SPT publications, we have calculated γ_{field} from a common set of simulations. We use a different set of cluster simulations to calculate γ_{field} in this work than have been used in previous SPT publications. To account for this, we recalculate γ_{field} for each of the 19 sub-fields in the SPT-SZ observing region using our new simulations, and we compare these results to the SPT-SZ values of γ_{field} published in de Haan et al. (2016). The ratios of the SPT-SZ field scalings in de Haan et al. (2016) to those calculated using our new simulations agree at the percent level among the 19 fields, and we take the median of this ratio and apply it to the γ_{field} we calculate for the SPTPol 100d field. This results in a final scaling for the SPTPol 100d field of $\gamma_{\text{field}} = 2.66$. Because γ_{field} is defined relative to V10, this implies that for the same cluster detected in this data and in V10, we should find a value of ξ 2.66 times larger here. A quick check using the most significant object in both catalogs, SPT-CL J2337-5942, bears this out: that cluster is detected in this work with $\zeta = 38.9$ and in V10 with $\zeta = 14.8$ (a factor of 2.63 smaller).

In order to provide consistent, comparable mass estimates between cluster catalogs produced by the SPT collaboration, we adopt the mass estimation method used in B15. In this method, the cosmological parameters are held fixed and only the scaling relation parameters are varied (such that the fitting procedure becomes equivalent to abundance-matching to the fixed cosmology). The cluster catalog we use for the scaling-relation fit is the updated B15 catalog presented in Bocquet et al. (2019).⁵ The resulting best-fit parameters

⁵The updated catalog provides better redshift measurements on a handful of clusters and redshifts for a few previously unconfirmed clusters.




Figure 5.4: Four new $z > 1$ clusters identified in the SPTPol 100d catalog. Clockwise from upper left: SPT-CL J2259-5301 at $z = 1.16 \pm 0.09$ detected at $\xi = 5.1$; SPT-CL J2336-5252 ($z = 1.22 \pm 0.09, \xi = 6.7$); SPT-CL J2323-5752 ($z = 1.3 \pm 0.1, \xi = 5.1$), and SPT-CL J2355-5514 ($z = 1.4 \pm 0.1, \xi = 4.9$). The composite *rgb* images are created using Magellan/PISCO *g* + *r*-band data, Magellan/LDSS3-C *z*-band data, and *Spitzer* IRAC 3.6 μm data from the SSDF; detection contours from the SPTPol $\theta_c=0'.25$ filtered maps are overlaid.

are $A_{\text{SZ}} = 4.07$, $B_{\text{SZ}} = 1.65$, $C_{\text{SZ}} = 0.63$, and $\sigma_{\ln \zeta} = 0.18$. As a check, we run the same fit using only the cluster sample presented in this work. As expected, the constraints are not as tight as those provided by the fit using the B15 sample, but they are consistent.

5.3 Optical and IR Followup

As in previous SPT cluster publications, we use the existence of overdensities of red-sequence galaxies at the locations of SZ cluster candidates to both confirm candidates as clusters as well as to obtain redshift estimates for confirmed systems. Galaxy overdensities are identified in optical and near-infrared imaging observations; such imaging was obtained for all 66 cluster candidates at $\xi \geq 5$ and 94% (84/89) of the candidates at $\xi \geq 4.6$. Optical and infrared imaging for four high-redshift clusters is shown in Figure 5.4.

5.3.1 Optical Data

The optical imaging for this work was primarily conducted using the Parallel Imager for Southern Cosmology Observations (PISCO; Stalder et al. 2014), a simultaneous 4-band (*griz*) imager with a $9'$ field-of-view mounted on the 6.5-meter Magellan/Clay telescope at Las Campanas Observatory. The simultaneous imaging in multiple bands allows PISCO to provide efficient follow-up of faint targeted sources such as galaxy clusters.

PISCO data were reduced using a custom-built pipeline that incorporates standard image processing corrections (overscan, debiasing, flat-fielding, illumination) as well as additional PISCO-specific corrections that account for non-linearities introduced by bright sources. Reduced images are further processed through the PHOTPIPE pipeline (Rest et al., 2005; Garg et al., 2007; Miknaitis et al., 2007) for astrometric calibration and to prepare for coaddition which is performed with the SWarp algorithm (Bertin et al., 2002). Astrometry is tied to bright stars from the Dark Energy Survey public release (Dark Energy Survey Collaboration et al., 2018) and sources from the second Gaia data release (Gaia Collaboration et al., 2018). Sources are identified in the coadded imaging data utilizing the SExtractor algorithm (Bertin and Arnouts, 1996) (v 2.8.6) in dual-image mode, star-galaxy separation is accomplished using the *SG* statistic introduced in Bleem et al. (2015a), and photometry is calibrated using the Stellar Locus Regression (SLR; High et al. 2009).

Additional optical imaging was drawn from targeted observations with the LDSS3-camera on Magellan/Clay (5 systems) as well as from a reprocessing of public data from the Blanco Cosmology Survey (Bleem et al. 2015a; 5 systems). These data were processed following the methods described in B15 and Bleem et al. (2015a), respectively.

5.3.2 Near-Infrared Data

Spitzer/IRAC imaging (Fazio et al., 2004) at $3.6 \mu\text{m}$ and $4.5 \mu\text{m}$ is available for the majority (91%) of the SPTPol 100d sample. These data come from both targeted imaging obtained over the course of the SPT-SZ survey (8%, see B15 for details) as well as from data in the *Spitzer*-South Pole Telescope Deep Field (SSDF; Ashby et al. 2013). The 94-square-degree SSDF in particular was designed to overlap with deep observations by the South Pole Telescope in this 100d field, although it does not cover the entire extent of the mm-wave

maps used in this work. Our near-infrared imaging with Spitzer is of sufficient depth to identify galaxy clusters at $z \lesssim 1.5$.

Near-Infrared Completeness

We use SSDF galaxy catalogs to estimate the probability that otherwise unconfirmed candidates are false detections. The method we use was developed in Song et al. (2012); and we only give a brief overview here. For each cluster candidate, we fit a β -model with $\beta = 1$ and a background density to the observed number density of galaxies brighter than 18.5 magnitude (Vega) in the $3.6 \mu\text{m}$ channel. We also perform this fit for 10,000 random locations in the SSDF footprint. The probability of a candidate being a false detection is then the fraction of the random locations that have a larger β -model amplitude than the fit at the location of the cluster candidate. This statistic is reported as P_{blank} for all candidates without redshifts and galaxy clusters with $z > 0.7$ in Table A. We consider candidates with $P_{\text{blank}} < 0.05$ confirmed galaxy clusters.

5.3.3 Redshift Estimation

Cluster candidates are characterized using tools that identify excesses of red-sequence galaxies at cluster locations. These tools, as well as the models used for the colors of red-sequence galaxies as a function of redshift, are described in detail in B15; here we simply note details specific to the analysis of this SPTPol sample.

As this is one of the first analyses to make use of PISCO data for confirming and characterizing SPT cluster candidates, it was necessary to calibrate the cluster red-sequence models to match the PISCO photometric system. Further details are given in Bleem et al. (2020), which performs a detailed comparison of redshifts derived from PISCO and a subsample of the B15 catalog with spectroscopic redshifts. Similar to other *griz*-band photometry derived from previous targeted observations of SPT clusters, we have transformed the PISCO photometric system to that of the Sloan Digital Sky Survey (SDSS; Abazajian et al. 2009) using instrumental color terms in the SLR calibration process.⁶ The PISCO imaging presented here was obtained as part of a broader effort to obtain high-quality data with good seeing on SPT-SZ and SPTPol clusters for a number of additional follow-up studies. In the course of this effort, imaging was obtained for a large fraction of galaxy clusters identified in the 2500d SPT-SZ survey with spectroscopic confirmations (Ruel et al., 2014; Bayliss et al., 2016); 51 of these systems were used to calibrate the red-sequence models for redshift estimation with the PISCO data; from this calibration we determine cluster redshifts to be accurate to $\sigma_z/(1+z) \sim 0.015 - 0.02$.

⁶This transformation works well in the *gri*-bands, but results in a poorer match in the *z*-band, in part because PISCO has a narrower *z*-filter (similar to that to be employed on LSST Ivezić et al. 2007). This photometric calibration scheme is sufficient for our work here, where having a well-calibrated color-redshift model is the driving requirement. However, in the future PISCO data will be tied to the LSST photometric system.




Figure 5.5: The mass and redshift distributions for several recent SZ-selected galaxy cluster catalogs. We plot the estimated mass and redshift for each of the 79 galaxy clusters with measured redshifts from this catalog, 517 from the SPT 2500 deg² survey (B15), 182 from the ACTPol-ED56 survey (Hilton et al., 2018), 91 from the ACT survey (Hasselfield et al., 2013a), and 1653 from the second *Planck* galaxy cluster catalog (Planck Collaboration et al., 2016a). The black line is the forecast detection threshold (50 photon counts) for eROSITA (Pillepich et al., 2012).

Finally, we have also taken advantage of additional spectroscopy of high-redshift clusters obtained following the publication of B15 (e.g., Bayliss et al. 2016; Khullar et al. 2019) to improve our calibration of near-infrared color-redshift models.

5.4 The Cluster Catalog

The SPTPol 100d cluster catalog contains 89 galaxy cluster candidates, 81 of which are optically confirmed and 29 are reported for the first time. In Table A, we present the galaxy cluster candidates found with $\xi > 4.6$. We originally chose a cutoff of $\xi = 4.5$, so that the

marginal false detection rate⁷ is approximately 50%. This choice is largely motivated by the need to follow up our candidates with dedicated optical imaging, and was used to select targets for followup with PISCO. In the interim, we discovered that it was necessary to change our noise calculation (as noted in §5.2.1), which caused ξ for each candidate to increase by approximately 0.1. This left us with nearly complete followup above $\xi = 4.6$, and a similar marginal false detection rate.

For each candidate, we provide its location, mass, detection significance (ξ), the β -model core radius (θ_c) of the filter scale that maximizes detection significance, redshift (if one is available) and P_{blank} (see §5.3.2, for candidates without redshifts and clusters with $z > 0.7$). Masses are estimated according to §5.2.4, and redshift estimation is discussed in §5.3. We have 2 candidates with $P_{\text{blank}} \leq 0.05$, but no optical data. We consider these candidates confirmed clusters, but are unable to provide redshifts or masses at this time.

The median redshift of the catalog is 0.60, with a maximum redshift of 1.38. There are 23 clusters (29% of the sample) with $z > 0.8$, 15 of which (18% of the sample) are above $z = 1.0$. The median mass of the catalog is $M_{500c} = 2.7 \times 10^{14} M_{\odot} h_{70}^{-1}$, and spans a range from $1.8 \times 10^{14} M_{\odot} h_{70}^{-1}$ to $8.3 \times 10^{14} M_{\odot} h_{70}^{-1}$. The full redshift and mass distribution in comparison to other catalogs is shown in Figure 5.5.

Based on our simulations, we expect to find 2 ± 1 false detections in this catalog (which equates to 98% purity) above $\xi = 5.0$. Our optical follow up is consistent with this estimate: it indicates that two of our candidates (SPT-CL J2349-5140 and SPT-CL J0002-5214) are likely false detections. For the full catalog ($\xi > 4.6$) we expect 8 ± 2 false detections; our followup indicates that there are 7 false detections. In addition, 1 candidate cannot be confirmed due to bright stars in the foreground. Ignoring this unconfirmed candidate, we obtain a catalog purity of 92%. The expected number of false detections is described in detail in §5.2.3 and plotted as a function of minimum significance in Figure 5.3. Our optical and near-infrared data on these candidates are of sufficient depth to confirm clusters out to $z \lesssim 1.5$.

Figure 5.6 shows the estimated completeness for this catalog. Our selection function is a Heaviside function in significance ($\Theta(\xi - 4.6)$). We convert this to a completeness in mass and redshift using the $\zeta - M$ scaling relation discussed in §5.2.4. We find that this catalog is expected to include all galaxy clusters in the survey region more massive than $M_{500c} = 2.6 \times 10^{14} M_{\odot} h_{70}^{-1}$ and at a redshift greater than $z = 0.25$.

Figure 5.5 shows the mass and redshift distribution for several galaxy cluster surveys. The ACT, ACTPol, SPT-SZ, and SPTPol selection functions are nearly independent of redshift. There is a slight slope, caused by two effects. At low redshift, the size of the cluster becomes similar that that of CMB fluctuations and residual atmospheric noise (see e.g., Schaffer et al. 2011). At moderate to high redshift, the slope is predicted by the self-similar evolution of galaxy clusters. Self-similarity implies that clusters of the same mass will have higher temperatures at higher redshift, which causes the tSZ to be brighter for more distant clusters.

⁷The marginal false detection rate is defined as the fraction of candidates we expect to be false detections in the range $\xi_0 - d\xi < \xi < \xi_0$, where ξ_0 is the cutoff significance.




Figure 5.6: The estimated completeness as a function of M_{500c} for the SPTPol 100d catalog. Completeness is estimated at four redshifts: 0.25, 0.50, 1.00, and 1.50. For fixed mass, the completeness at low redshift decreases, because the angular extent of the clusters becomes similar to the fluctuations in the CMB. We expect our catalog to be complete for $M_{500c} > 2.6 \times 10^{14} M_{\odot} h_{70}^{-1}$.

Self-similar evolution further predicts that $C_{sz} = \frac{2}{3}$ Kravtsov et al. (2006). tSZ-selected cluster catalogs generally do not provide strong constraints on self-similarity (Bocquet et al. 2019 provides only a $\sim 50\%$ constraint on C_{sz}). McDonald et al. (2017) used X-ray data to study clusters below $z = 1.9$ and found the bulk of the ICM to be remarkably self-similar. The *Planck* catalog is limited to lower redshift because the SZ signal from distant clusters is diluted by the relatively large *Planck* beams ($7'$ at 143 GHz).

5.4.1 Comparison to Other Cluster Catalogs

In this section, we compare the catalog presented in this work with other cluster catalogs. We compare against galaxy clusters in the Simbad⁸ database, as well as the second *Planck* cluster catalog (Planck Collaboration et al., 2016a), B15, and the XXL 365 catalog (Adami et al., 2018). Two objects are considered a match if they fall within $5'0$ of each other below $z = 0.3$. For higher redshift objects, we apply a threshold of $2'0$ (except for comparisons to the *Planck* catalog, where we use a $4'0$ radius to compensate for the larger *Planck* instrument beam). The full results can be found in Table A. Here, we highlight some recent catalogs.

The most directly comparable catalog is the SPT-SZ catalog presented in B15. The full SPT-SZ 2500 deg² catalog contains 0.22 confirmed clusters per square degree, while our catalog has nearly four times the density, at 0.86 clusters per square degree. The median

⁸<http://simbad.u-strasbg.fr/simbad/>




Figure 5.7: A merging cluster, SPT-CL J2331-5052, identified in the SPTPol 100d sample. This system was detected as two separate clusters in the B15 catalog, but due to the deeper mm-wave data used in this work, it is detected as a single system. The composite image is created using PISCO *gri* images. The contours show tSZ detection significance from the $\theta_c = 0'.5$ match-filtered map. The position reported in this catalog is in the upper right, near the center of the $\xi = 16$ contour. In B15, the smaller object (SPT-CL J2332-5053) is at $z = 0.56 \pm 0.04$; SPT-CL J2331-5052 is one of the 13 clusters with spectroscopically measured redshifts at $z = 0.576$. For the more massive object in the upper right of this image, $R_{500} \approx 2'.8$.

redshifts of the two catalogs are similar ($z_{\text{med}} = 0.60$ for this work, and $z_{\text{med}} = 0.55$ in B15), but our catalog contains a much larger fraction at high redshift. 18% of the clusters with measured redshifts in this work are more distant than $z = 1$, while only 8% of the clusters in B15 have $z > 1$.

The B15 catalog contains 28 candidates in the RA23H30DEC-55 field, and we find 22 of them in this work. One cluster (SPT-CL J2332-5053) is optically confirmed in B15, but not included in this catalog. It is part of an interacting system with SPT-CL J2331-5052, which was also noted in previous X-ray observations (Andersson et al., 2011). Our source finding algorithm groups all connected pixels above a fixed threshold. Due to the deeper data used in this work, there is a high significance “bridge” that connects the two clusters detected in B15. By increasing the detection threshold above the minimum value in the bridge in our source finding algorithm, we can force this detection to become separate objects. We find two clusters, one at (R.A., decl.) of (352.96122, -50.864841) with $\xi = 19.08$, and the other at (353.02512, -50.892534) with $\xi = 9.59$, both with $\theta_c = 0'.5$. Both objects are at the

same redshift (the smaller object, reported only in B15 is at $z = 0.56 \pm 0.04$, while SPT-CL J2331-5052 is at $z = 0.576$). For the more massive object, $R_{500} \approx 2'.8$, which is more than half the projected distance between the two objects. In Table A we report the single cluster SPT-CL J2331-5052. This choice is largely arbitrary for this catalog, and we have chosen to be consistent with our stated selection function. However, future catalogs will find many objects like this, and the choice to report them as single clusters or more than one must be accounted for in any cosmological analyses. Figure 5.7 shows a composite image with combined optical imaging and tSZ contours of this system. The remaining 5 candidates found in B15 are not optically confirmed, which is consistent with the expected false detection rate. Finally, one candidate (SPT-CL J2321-5418) is not optically confirmed in this catalog or B15, due to several bright stars in the foreground of the optical imaging.

The sky coverage in this work extends slightly beyond the boundaries of the ra23h30dec-55 field in B15, and as a result we find an additional 4 clusters which are reported in different fields in the B15 catalog; this brings the total number of matching clusters in both our catalog and B15 to 26. The estimated masses for these two catalogs are consistent with the estimated errors.

The *Planck* galaxy cluster catalog (Planck Collaboration et al., 2016a) contains 4 clusters in this field, all of which are included in both this catalog and the SPT 2500 deg² catalog. The cluster redshifts agree and the scatter in mass between these clusters is consistent with the scatter between the clusters in B15 and the *Planck* cluster catalog. A more detailed comparison of *Planck* and SPTPol masses is included in Bleem et al. (2020).

The XXL is an X-ray survey performed using *XMM-Newton* (Pierre et al., 2017), and is one of the deepest X-ray surveys to date. It covers 50 deg², 25 of which overlap with the sky area in this work. In the overlapping area, XXL finds 154 galaxy clusters, the majority of which are at sufficiently low mass or redshift as to be undetectable in this analysis. 14 of these clusters are also found in the SPTPol data. The masses of clusters in the XXL catalog are reported using two different methods. Since we report masses in M_{500c} , we compare to the XXL masses calculated using an X-ray temperature scaling relation (since these are also in M_{500c}). Of the 14 matching clusters, only 13 have scaling relation masses reported in Adami et al. (2018). We compare the masses of the 13 clusters, and find that they are statistically consistent, with a median ratio of $1.08_{+0.06}^{-0.04}$ and RMS scatter of 41%. Dividing out the median difference yields a scatter of 38%. The scatter is dominated by the highest mass cluster, however the SPTPol mass estimate is only 1.2σ discrepant from the estimated XXL mass. The masses of the matching clusters are shown in Figure 5.8. We also compare redshifts between the two catalogs. Most of the redshifts are consistent, but one cluster is significantly discrepant. Our optical imaging suggests that there is a very low richness cluster at the reported redshift ($z = 0.81$) of XLSSC 549. There is also a much larger cluster at higher redshift ($z = 1.2$), which we report as SPT-CL J2334-5308. Given that the mass reported for XLSSC 549 is very similar to that of SPT-CL J2334-5308 ($M_{500c} = 2.58 \pm 0.97 \times 10^{14} M_{\odot} h_{70}^{-1}$ and $M_{500c} = 2.25 \pm 0.38 \times 10^{14} M_{\odot} h_{70}^{-1}$, respectively), we believe that the reported redshift for XLSSC 549 is incorrect. Assuming that the X-ray temperature remains fixed, the change in redshift would bring the mass of XLSSC 549 closer to the mass measured in this work.




Figure 5.8: Galaxy cluster masses as measured by SPTPol and XXL. Both the XXL and SPTPol mass uncertainties are calculated using a scaling relation with an intrinsic scatter in the mass-observable relation (see §5.2.4), which leads to an X-ray or SZ mass uncertainty that is highly correlated between clusters in the sample. Only 13 of the 14 clusters found in both surveys are included, because the mass is not provided for one of the XXL clusters.

5.5 Conclusions

This work documents the assembly and properties of the first catalog of galaxy clusters using data from the SPTPol camera. Using a matched filter on intensity maps from both 95 and 150 GHz bands, we construct a catalog of 89 candidates at signal-to-noise ratio $\xi > 4.6$. Of these, 29 candidates are reported for the first time, and 81 are confirmed using optical and infrared observations. Based on simulations, we expect the sample to be 91% pure (8 false detections), and we find that the number of unconfirmed candidates (7) is consistent with this expectation.

We use a program of optical and infrared follow up to both confirm candidates and estimate their redshift. The median redshift of the catalog is $z \sim 0.6$, with 23 (29%) clusters at $z > 0.8$, and 15 (18%) at redshift greater than 1.0. While this catalog does not contain clusters at the highest redshifts seen in B15 (which contains 6 confirmed clusters at $z \geq 1.4$), the fraction at $z > 1$ is substantially larger. We estimate galaxy cluster masses based on the detection significance and redshift of each cluster. Using the fixed set of cosmological parameters and scaling relation from B15, we find masses ranging from $M_{500c} = 1.8 \times 10^{14} M_{\odot} h_{70}^{-1}$ to $8.3 \times 10^{14} M_{\odot} h_{70}^{-1}$, with a median mass of $2.7 \times 10^{14} M_{\odot} h_{70}^{-1}$.

This work was followed by two more cluster catalogs using data from the SPTPol camera. The next catalog covers 27 times the area at higher noise levels (the 2700 square degree

SPTPol-ECS field, Bleem et al. 2020), while the other covers 5 times the area at similar noise levels (the SPTPol 500 square degree field, Bleem et al. 2024). This program is continuing with the SPT-3G camera, which began collecting data in early 2018. The SPT-3G survey will cover 15 times the area included in this work. Data from this survey will achieve noise levels of (3.0, 2.2, 8.8) $\mu\text{K-arcmin}$ at (95, 150, 220) GHz, approximately ($3 \times$, $2 \times$) lower at (95, 150) GHz than the map noise in this work. These noise levels will allow us to detect over 3000 galaxy clusters with redshifts extending beyond $z = 2.0$, and masses as low as $9 \times 10^{13} M_{\odot} h_{70}^{-1}$. On a similar timescale, a new X-ray satellite, eROSITA, will survey the entire sky, and produce an extremely large cluster catalog (Merloni et al., 2012). Although eROSITA will detect lower mass clusters, our catalog already includes clusters below the eROSITA detection threshold at $z > 1$. This work is still in progress, with early results shown in Sobrin and SPT-3G Team 2022.

The sky area used in this work has also been surveyed at several other wavelengths. In fact, we have used some of those datasets in this work. In the infrared, nearly the entire area has been observed by *Spitzer*/IRAC; a similar area was observed in the far infrared with the SPIRE instrument on the *Herschel Space Observatory* (Pilbratt et al., 2010); the Dark Energy Survey (Dark Energy Survey Collaboration et al., 2018) has covered the region in 5 optical and near-infrared bands (*grizY*); the XXL survey has observed 25 deg^2 of the SPTPol 100d field; and the Australian Telescope Compact Array (ATCA) was used to survey 86 deg^2 of the field at 1.1-3.1 GHz (O’Brien et al., 2018), as well as the entire Southern sky at 20 GHz (Ekers et al., 2007). Since these observations are sensitive to different physics (galactic emission in the optical and infrared, ICM temperature in X-rays, and electron pressure in the microwave), they provide semi-independent measures of intrinsic cluster properties. Combined, this multi-wavelength data set can give us a much clearer picture of the physics at work in the galaxy clusters than any single probe.

Chapter 6

Conclusion

This dissertation has discussed observations of the Cosmic Microwave Background with the South Pole Telescope. The CMB is the oldest light in the universe, originating from the surface of last scattering, when the universe suddenly became neutral. Its primary anisotropies in the CMB provide state-of-the-art constraints on the Λ CDM parameters (such as H_0 the expansion rate of the universe). CMB photons traverse the majority of the universe before we can observe them, and a small fraction of them are affected by structures that evolved over cosmic history. These secondary anisotropies allow us to constrain significant processes in cosmic history, such as reionization and the growth of large scale structure. To measure the small secondary anisotropies, we need a highly sensitive camera with high resolution at mm wavelengths. The SPT is capable of exactly these observations.

The 10-m primary mirror of the SPT provides arcminute-scale resolution at mm wavelengths. SPT-3G (the third-generation camera on the SPT) houses 16000 polarization sensitive detectors in three bands centered at 95, 150, and 220 GHz. The detectors are TESes, which are operated below temperatures below 1 K. While many systems are important to maintaining the temperatures, we describe the vacuum window and infrared filtering in detail.

The TESes have a time-variable sensitivity that must be calibrated out to reproduce images of the sky. The SPT-3G calibration pipeline uses a combination of bright galactic sources (HII regions) and an internal calibration source to account for changes in the sensitivity. To achieve the maximum resolution of the SPT-3G camera, we must determine where (on the sky) each detector points relative to the others. With these components in place, we can convert the raw time-ordered data into images of the sky. In our mapmaking pipeline, we perform a basic data selection to remove detectors that have either too much noise or too little sensitivity. We then filter the remaining detector TOD to remove low-frequency noise (which comes from the atmosphere and the readout electronics). Finally, we bin the filtered TOD into an image of the sky.

Before we can interpret these images, we need to account for the effects of the telescope optics. We use observations of bright extragalactic sources (AGN) and Saturn to construct the instrument beam. Saturn is bright enough to saturate our detectors when they pass close to the center of the planet, so we “stitch” AGN near the center of the maps to Saturn at the

outskirts. The beam is a critical part of any analysis of the maps, so we perform a variety of systematic tests to ensure our measurement is unbiased. We also develop a method to measure the uncertainty in the beam despite the complicate noise correlations induced by the stitching process. Finally, we show how all of these parts come together by constructing a catalog of galaxy clusters using data from the SPTPol camera.

Bibliography

- K. N. Abazajian, J. K. Adelman-McCarthy, M. A. Agüeros, et al. The Seventh Data Release of the Sloan Digital Sky Survey. *ApJS*, 182:543–558, June 2009. doi: 10.1088/0067-0049/182/2/543.
- G. O. Abell, H. G. Corwin, Jr., and R. P. Olowin. A catalog of rich clusters of galaxies. *ApJS*, 70:1–138, May 1989. doi: 10.1086/191333.
- R. Adam, I. Bartalucci, G. W. Pratt, et al. Mapping the kinetic Sunyaev-Zel’dovich effect toward MACS J0717.5+3745 with NIKA. *A&A*, 598:A115, February 2017. doi: 10.1051/0004-6361/201629182.
- C. Adami, P. Giles, E. Koulouridis, et al. The XXL Survey. XX. The 365 cluster catalogue. *A&A*, 620:A5, Nov 2018. doi: 10.1051/0004-6361/201731606.
- Simone Aiola, Erminia Calabrese, Loïc Maurin, et al. The Atacama Cosmology Telescope: DR4 maps and cosmological parameters. *J. Cosmology Astropart. Phys.*, 2020(12):047, December 2020. doi: 10.1088/1475-7516/2020/12/047.
- S. W. Allen, A. E. Evrard, and A. B. Mantz. Cosmological Parameters from Observations of Galaxy Clusters. *ARA&A*, 49:409–470, September 2011. doi: 10.1146/annurev-astro-081710-102514.
- K. Andersson, B. A. Benson, P. A. R. Ade, et al. X-Ray Properties of the First Sunyaev-Zel’dovich Effect Selected Galaxy Cluster Sample from the South Pole Telescope. *ApJ*, 738:48–+, September 2011. doi: 10.1088/0004-637X/738/1/48.
- M. L. N. Ashby, S. A. Stanford, M. Brodwin, et al. The Spitzer South Pole Telescope Deep Field: Survey Design and Infrared Array Camera Catalogs. *ApJS*, 209:22, December 2013. doi: 10.1088/0067-0049/209/2/22.
- J. E. Austermann, K. A. Aird, J. A. Beall, et al. SPTpol: an instrument for CMB polarization measurements with the South Pole Telescope. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 8452 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, September 2012. doi: 10.1117/12.927286.

- L. Balkenhol, D. Dutcher, A. Spurio Mancini, et al. Measurement of the CMB temperature power spectrum and constraints on cosmology from the SPT-3G 2018 TT, TE, and EE dataset. *Phys. Rev. D*, 108(2):023510, July 2023. doi: 10.1103/PhysRevD.108.023510.
- N. Battaglia, J. R. Bond, C. Pfrommer, and J. L. Sievers. On the Cluster Physics of Sunyaev-Zel'dovich and X-Ray Surveys. II. Deconstructing the Thermal SZ Power Spectrum. *ApJ*, 758:75, October 2012. doi: 10.1088/0004-637X/758/2/75.
- M. B. Bayliss, J. Ruel, C. W. Stubbs, et al. SPT-GMOS: A Gemini/GMOS-South Spectroscopic Survey of Galaxy Clusters in the SPT-SZ Survey. *ApJS*, 227:3, November 2016. doi: 10.3847/0067-0049/227/1/3.
- George D. Becker, James S. Bolton, Piero Madau, et al. Evidence of patchy hydrogen reionization from an extreme Ly α trough below redshift six. *MNRAS*, 447(4):3402–3419, March 2015. doi: 10.1093/mnras/stu2646.
- Robert H. Becker, Xiaohui Fan, Richard L. White, et al. Evidence for Reionization at $z \sim 6$: Detection of a Gunn-Peterson Trough in a $z=6.28$ Quasar. *AJ*, 122(6):2850–2857, December 2001. doi: 10.1086/324231.
- Shawn M. Beckman. *Advancements in Microwave Optics for Measuring Polarization Anisotropies in the Cosmic Microwave Background*. PhD thesis, University of California, Berkeley, Department of Physics, January 2023.
- A. N. Bender, P. A. R. Ade, A. J. Anderson, et al. Integrated performance of a frequency domain multiplexing readout in the SPT-3G receiver. In Wayne S. Holland and Jonas Zmuidzinas, editors, *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VIII*, volume 9914 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 99141D, July 2016. doi: 10.1117/12.2232146.
- Amy N. Bender, Jean-François Cliche, Tijmen de Haan, et al. Digital frequency domain multiplexing readout electronics for the next generation of millimeter telescopes. In Wayne S. Holland and Jonas Zmuidzinas, editors, *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII*, volume 9153 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 91531A, July 2014. doi: 10.1117/12.2054949.
- B. A. Benson. SPT Talk Archive. From a set of images made by Brad Benson for general use by the SPT collaboration, 2018.
- B. A. Benson. TES Bolometers. 2022 Summer School on CMB Detectors and Instrumentation, 2022. URL https://kicp-workshops.uchicago.edu/2022-CMB-SCHOOL/depot/talk-benson-bradford__2.pdf.

- B. A. Benson, T. de Haan, J. P. Dudley, et al. Cosmological Constraints from Sunyaev-Zel'dovich-selected Clusters with X-Ray Observations in the First 178 deg² of the South Pole Telescope Survey. *ApJ*, 763(2):147, February 2013. doi: 10.1088/0004-637X/763/2/147.
- E. Bertin and S. Arnouts. SExtractor: Software for source extraction. *A&AS*, 117:393–404, June 1996.
- E. Bertin, Y. Mellier, M. Radovich, et al. The TERAPIX Pipeline. In D. A. Bohlender, D. Durand, & T. H. Handley, editor, *Astronomical Data Analysis Software and Systems XI*, volume 281 of *Astronomical Society of the Pacific Conference Series*, pages 228–+, 2002.
- R. S. Bhatia, S. T. Chase, S. F. Edgington, et al. A three-stage helium sorption refrigerator for cooling of infrared detectors to 280 mK. *Cryogenics*, 40(11):685–691, January 2000. doi: 10.1016/S0011-2275(00)00072-2.
- M. Birkinshaw. The sunyaev zel'dovich effect. *Physics Reports*, 310:97, 1999.
- Diego Blas, Julien Lesgourgues, and Thomas Tram. The Cosmic Linear Anisotropy Solving System (CLASS). Part II: Approximation schemes. *J. Cosmology Astropart. Phys.*, 2011(7):034, July 2011. doi: 10.1088/1475-7516/2011/07/034.
- L. E. Bleem, B. Stalder, M. Brodwin, et al. A New Reduction of the Blanco Cosmology Survey: An Optically Selected Galaxy Cluster Catalog and a Public Release of Optical Data Products. *ApJS*, 216:20, January 2015a. doi: 10.1088/0067-0049/216/1/20.
- L. E. Bleem, B. Stalder, T. de Haan, et al. Galaxy Clusters Discovered via the Sunyaev-Zel'dovich Effect in the 2500-Square-Degree SPT-SZ Survey. *ApJS*, 216:27, February 2015b. doi: 10.1088/0067-0049/216/2/27.
- L. E. Bleem, S. Bocquet, B. Stalder, et al. The SPTpol Extended Cluster Survey. *ApJS*, 247(1):25, March 2020. doi: 10.3847/1538-4365/ab6993.
- L. E. Bleem, T. M. Crawford, B. Ansarinejad, et al. CMB/kSZ and Compton- γ Maps from 2500 deg² of SPT-SZ and Planck Survey Data. *ApJS*, 258(2):36, February 2022. doi: 10.3847/1538-4365/ac35e9.
- L. E. Bleem, M. Klein, T. M. C. Abbot, et al. Galaxy Clusters Discovered via the Thermal Sunyaev-Zel'dovich Effect in the 500-square-degree SPTpol Survey. *The Open Journal of Astrophysics*, 7:13, February 2024. doi: 10.21105/astro.2311.07512.
- S. Bocquet, J. P. Dietrich, T. Schrabback, et al. Cluster Cosmology Constraints from the 2500 deg² SPT-SZ Survey: Inclusion of Weak Gravitational Lensing Data from Magellan and the Hubble Space Telescope. *ApJ*, 878(1):55, Jun 2019. doi: 10.3847/1538-4357/ab1f10.
- S. Bocquet, S. Grandis, L. E. Bleem, et al. SPT Clusters with DES and HST Weak Lensing. I. Cluster Lensing and Bayesian Population Modeling of Multi-Wavelength Cluster Datasets. *arXiv e-prints*, art. arXiv:2310.12213, October 2023. doi: 10.48550/arXiv.2310.12213.

- M. K. Braid and H. T. MacGillivray. A finding list of southern clusters of galaxies. I. *MNRAS*, 182:241–248, January 1978.
- M. R. Calabretta and E. W. Greisen. Representations of celestial coordinates in FITS. *A&A*, 395:1077–1122, December 2002. doi: 10.1051/0004-6361:20021327.
- J. E. Carlstrom, G. P. Holder, and E. D. Reese. Cosmology with the Sunyaev-Zel’dovich Effect. *ARA&A*, 40:643–680, 2002.
- J. E. Carlstrom, P. A. R. Ade, K. A. Aird, et al. The 10 Meter South Pole Telescope. *PASP*, 123:568–581, May 2011. doi: 10.1086/659879.
- A. Cavaliere and R. Fusco-Femiano. X-rays from hot plasma in clusters of galaxies. *A&A*, 49:137, May 1976.
- C. L. Chang, P. A. R. Ade, K. A. Aird, et al. SPT-SZ: a Sunyaev-Zel’dovich survey for galaxy clusters. In Betty Young, Blas Cabrera, and Aaron Miller, editors, *The Thirteenth International Workshop on Low Temperature Detectors - LTD13*, volume 1185 of *American Institute of Physics Conference Series*, pages 475–477. AIP, December 2009. doi: 10.1063/1.3292381.
- J. Choi, H. Ishitsuka, S. Mima, et al. Radio-transparent multi-layer insulation for radiowave receivers. *Review of Scientific Instruments*, 84(11):114502-114502-6, November 2013. doi: 10.1063/1.4827081.
- William Coulton, Mathew S. Madhavacheril, Adriaan J. Duivenvoorden, et al. Atacama Cosmology Telescope: High-resolution component-separated maps across one third of the sky. *Phys. Rev. D*, 109(6):063530, March 2024. doi: 10.1103/PhysRevD.109.063530.
- T. M. Crawford. Bandpasses and conversion factors redux. Internal memo, 2022.
- T. M. Crawford, K. K. Schaffer, S. Bhattacharya, et al. A Measurement of the Secondary-CMB and Millimeter-wave-foreground Bispectrum using 800 deg² of South Pole Telescope Data. *ApJ*, 784(2):143, April 2014. doi: 10.1088/0004-637X/784/2/143.
- J. N. Cuzzi. Physical Properties of Saturn’s Rings. In A. Brahic, editor, *Planetary Rings*, page 211, January 1984.
- Antonio C. da Silva, Scott T. Kay, Andrew R. Liddle, and Peter A. Thomas. Hydrodynamical simulations of the Sunyaev-Zel’dovich effect: cluster scaling relations and X-ray properties. *MNRAS*, 348(4):1401–1408, March 2004. doi: 10.1111/j.1365-2966.2004.07463.x.
- Dark Energy Survey Collaboration, T. M. C. Abbott, F. B. Abdalla, et al. The Dark Energy Survey Data Release 1. *ArXiv e-prints*, January 2018.

- R. Datta, C. D. Munson, M. D. Niemack, et al. Large-aperture wide-bandwidth antireflection-coated silicon lenses for millimeter wavelengths. *Appl. Opt.*, 52(36):8747, December 2013. doi: 10.1364/AO.52.008747.
- T. de Haan, B. A. Benson, L. E. Bleem, et al. Cosmological Constraints from Galaxy Clusters in the 2500 Square-degree SPT-SZ Survey. *ApJ*, 832:95, November 2016. doi: 10.3847/0004-637X/832/1/95.
- G. De Zotti, R. Ricci, D. Mesa, et al. Predictions for high-frequency radio surveys of extragalactic sources. *A&A*, 431:893–903, March 2005. doi: 10.1051/0004-6361:20042108.
- Scott Dodelson. *Modern Cosmology*. Academic Press, 2003.
- David E. Dunn, Imke de Pater, Melvyn Wright, et al. High-Quality BIMA-OVRO Images of Saturn and its Rings at 1.3 and 3 Millimeters. *AJ*, 129(2):1109–1116, February 2005. doi: 10.1086/424536.
- D. Dutcher, L. Balkenhol, P. A. R. Ade, et al. Measurements of the E -mode polarization and temperature-E -mode correlation of the CMB from SPT-3G 2018 data. *Phys. Rev. D*, 104(2):022003, July 2021. doi: 10.1103/PhysRevD.104.022003.
- R. D. Ekers, M. Massardi, and E. M. Sadler. AT20G: an all sky blind survey at 20GHz. In *From Planets to Dark Energy: the Modern Radio Universe*, page 47, 2007.
- L. W. Esposito, J. N. Cuzzi, J. B. Holberg, et al. Saturn’s rings - Structure, dynamics, and particle properties. In T. Gehrels and M. S. Matthews, editors, *Saturn*, pages 463–545. 1984.
- G. G. Fazio, J. L. Hora, L. E. Allen, et al. The Infrared Array Camera (IRAC) for the Spitzer Space Telescope. *ApJS*, 154:10–17, September 2004. doi: 10.1086/422843.
- D. J. Fixsen. The Temperature of the Cosmic Microwave Background. *ApJ*, 707:916–920, December 2009. doi: 10.1088/0004-637X/707/2/916.
- Samuel Flender, Lindsey Bleem, Hal Finkel, et al. Simulations of the Pairwise Kinematic Sunyaev-Zel’dovich Signal. *ApJ*, 823:98, Jun 2016. doi: 10.3847/0004-637X/823/2/98.
- Gaia Collaboration, A. G. A. Brown, A. Vallenari, et al. Gaia Data Release 2. Summary of the contents and survey properties. *ArXiv e-prints*, April 2018.
- A. Garg, C. W. Stubbs, P. Challis, et al. Light Curves of Type Ia Supernovae from Near the Time of Explosion. *AJ*, 133:403–419, February 2007. doi: 10.1086/510118.
- E. M. George, C. L. Reichardt, K. A. Aird, et al. A Measurement of Secondary Cosmic Microwave Background Anisotropies from the 2500-Square-degree SPT-SZ Survey. *ApJ*, 799:177, February 2015. doi: 10.1088/0004-637X/799/2/177.

- Joseph E. Golec, Jeffrey J. McMahon, Aamire Ali, et al. Design and fabrication of metamaterial anti-reflection coatings for the Simons Observatory. In Ramón Navarro and Roland Gejl, editors, *Advances in Optical and Mechanical Technologies for Telescopes and Instrumentation IV*, volume 11451 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 114515T, December 2020. doi: 10.1117/12.2561720.
- Joseph E. Golec, Shreya Sutariya, Rebecca Jackson, et al. Simons Observatory: broadband metamaterial antireflection cuttings for large-aperture alumina optics. *Appl. Opt.*, 61(30): 8904, October 2022. doi: 10.1364/AO.472459.
- Anthony H. Gonzalez, Daniel P. Gettings, Mark Brodwin, et al. The Massive and Distant Clusters of WISE Survey. I. Survey Overview and a Catalog of > 2000 Galaxy Clusters at $z \approx 1$. *The Astrophysical Journal Supplement Series*, 240:33, Feb 2019. doi: 10.3847/1538-4365/aafad2.
- A. Gorce, S. Ilić, M. Douspis, et al. Improved constraints on reionisation from CMB observations: A parameterisation of the kSZ effect. *A&A*, 640:A90, August 2020. doi: 10.1051/0004-6361/202038170.
- James E. Gunn and Bruce A. Peterson. On the Density of Neutral Hydrogen in Intergalactic Space. *ApJ*, 142:1633–1636, November 1965. doi: 10.1086/148444.
- M. Hasselfield, M. Hilton, T. A. Marriage, et al. The Atacama Cosmology Telescope: Sunyaev-Zel’dovich selected galaxy clusters at 148 GHz from three seasons of data. *J. Cosmology Astropart. Phys.*, 7:008, July 2013a. doi: 10.1088/1475-7516/2013/07/008.
- Matthew Hasselfield, Kavilan Moodley, J. Richard Bond, et al. The Atacama Cosmology Telescope: Beam Measurements and the Microwave Brightness Temperatures of Uranus and Saturn. *ApJS*, 209(1):17, November 2013b. doi: 10.1088/0067-0049/209/1/17.
- Katrin Heitmann, Hal Finkel, Adrian Pope, et al. The Outer Rim Simulation: A Path to Many-Core Supercomputers. *arXiv e-prints*, art. arXiv:1904.11970, Apr 2019.
- Y. Hezaveh, K. Vanderlinde, G. Holder, and T. de Haan. Lensing Noise in Millimeter-wave Galaxy Cluster Surveys. *ApJ*, 772:121, August 2013. doi: 10.1088/0004-637X/772/2/121.
- F. W. High, C. W. Stubbs, A. Rest, et al. Stellar Locus Regression: Accurate Color Calibration and the Real-Time Determination of Galaxy Cluster Photometric Redshifts. *AJ*, 138: 110–129, July 2009. doi: 10.1088/0004-6256/138/1/110.
- M. Hilton, M. Hasselfield, C. Sifón, et al. The Atacama Cosmology Telescope: The Two-season ACTPol Sunyaev-Zel’dovich Effect Selected Cluster Catalog. *ApJS*, 235:20, March 2018. doi: 10.3847/1538-4365/aaa6cb.
- W. Hu and M. White. CMB anisotropies in the weak coupling limit. *A&A*, 315:33–39, November 1996. doi: 10.48550/arXiv.astro-ph/9507060.

- N. Huang, L. E. Bleem, B. Stalder, et al. Galaxy Clusters Selected via the Sunyaev-Zel'dovich Effect in the SPTpol 100-square-degree Survey. *AJ*, 159(3):110, March 2020. doi: 10.3847/1538-3881/ab6a96.
- Edwin Hubble. A Relation between Distance and Radial Velocity among Extra-Galactic Nebulae. *Proceedings of the National Academy of Science*, 15(3):168–173, March 1929. doi: 10.1073/pnas.15.3.168.
- Y. Inoue, P. Ade, Y. Akiba, et al. POLARBEAR-2: an instrument for CMB polarization measurements. In *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VIII*, volume 9914 of *Proc. SPIE*, page 99141I, July 2016. doi: 10.1117/12.2231961.
- K. D. Irwin and M. E. Huber. SQUID operational amplifier. *IEEE Transactions on Applied Superconductivity*, 11(1):1265–1270, March 2001. doi: 10.1109/77.919580.
- Ž. Ivezić, J. A. Smith, G. Miknaitis, et al. Sloan Digital Sky Survey Standard Star Catalog for Stripe 82: The Dawn of Industrial 1% Optical Photometry. *AJ*, 134:973–998, September 2007. doi: 10.1086/519976.
- Andrew H. Jaffe and Marc Kamionkowski. Calculation of the Ostriker-Vishniac effect in cold dark matter models. *Phys. Rev. D*, 58(4):043001, August 1998. doi: 10.1103/PhysRevD.58.043001.
- D. H. Jones, M. A. Read, W. Saunders, et al. The 6dF Galaxy Survey: final redshift release (DR3) and southern large-scale structures. *MNRAS*, 399:683–698, October 2009. doi: 10.1111/j.1365-2966.2009.15338.x.
- N. Kaiser. Evolution and clustering of rich clusters. *MNRAS*, 222:323–345, September 1986. doi: 10.1093/mnras/222.2.323.
- Jae Hwan Kang, P. A. R. Ade, Z. Ahmed, et al. 2017 upgrade and performance of BICEP3: a 95GHz refracting telescope for degree-scale CMB polarization. In Jonas Zmuidzinas and Jian-Rong Gao, editors, *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*, volume 10708 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 107082N, July 2018. doi: 10.1117/12.2313854.
- Brian G. Keating, Peter A. R. Ade, James J. Bock, et al. BICEP: a large angular scale CMB polarimeter. In Silvano Fineschi, editor, *Polarimetry in Astronomy*, volume 4843 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 284–295, February 2003. doi: 10.1117/12.459274.
- R. Keisler, C. L. Reichardt, K. A. Aird, et al. A Measurement of the Damping Tail of the Cosmic Microwave Background Power Spectrum with the South Pole Telescope. *ApJ*, 743: 28, December 2011. doi: 10.1088/0004-637X/743/1/28.

- R. Keisler, S. Hoover, N. Harrington, et al. Measurements of sub-degree b-mode polarization in the cosmic microwave background from 100 square degrees of sptpol data. *ApJ*, 807(2): 151, 2015. URL <http://stacks.iop.org/0004-637X/807/i=2/a=151>.
- G. Khullar, L. E. Bleem, M. B. Bayliss, et al. Spectroscopic Confirmation of Five Galaxy Clusters at $z > 1.25$ in the 2500 deg² SPT-SZ Survey. *ApJ*, 870(1):7, Jan 2019. doi: 10.3847/1538-4357/aaeed0.
- E. Komatsu, K. M. Smith, J. Dunkley, et al. Seven-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Interpretation. *ApJS*, 192:18–+, February 2011. doi: 10.1088/0067-0049/192/2/18.
- A. V. Kravtsov, A. Vikhlinin, and D. Nagai. A New Robust Low-Scatter X-Ray Mass Indicator for Clusters of Galaxies. *ApJ*, 650:128–136, October 2006. doi: 10.1086/506319.
- James W. Lamb. Miscellaneous data on materials for millimetre and submillimetre optics. *International Journal of Infrared and Millimeter Waves*, 17(12):1997–2034, December 1996. doi: 10.1007/BF02069487.
- A. Lewis, A. Challinor, and A. Lasenby. Efficient Computation of Cosmic Microwave Background Anisotropies in Closed Friedmann-Robertson-Walker Models. *ApJ*, 538:473–476, August 2000. doi: 10.1086/309179.
- M. Lueker, C. L. Reichardt, K. K. Schaffer, et al. Measurements of Secondary Cosmic Microwave Background Anisotropies with the South Pole Telescope. *ApJ*, 719:1045–1066, August 2010. doi: 10.1088/0004-637X/719/2/1045.
- Adam B. Mantz, Steven W. Allen, R. Glenn Morris, et al. Weighing the giants- V. Galaxy cluster scaling relations. *MNRAS*, 463(4):3582–3603, December 2016. doi: 10.1093/mnras/stw2250.
- Tomotake Matsumura, Karl Young, Qi Wen, et al. Millimeter-wave broadband antireflection coatings using laser ablation of subwavelength structures. *Appl. Opt.*, 55(13):3502, May 2016. doi: 10.1364/AO.55.003502.
- T. Mauch, T. Murphy, H. J. Buttery, et al. Sumss: A wide-field radio imaging survey of the southern sky - ii. the source catalogue. *MNRAS*, 342:1117–1130, March 2003. URL <https://ui.adsabs.harvard.edu/abs/2003MNRAS.342.1117M>.
- M. McDonald, S. W. Allen, M. Bayliss, et al. The Remarkable Similarity of Massive Galaxy Clusters from $z \sim 0$ to $z \sim 1.9$. *ApJ*, 843(1):28, Jul 2017. doi: 10.3847/1538-4357/aa7740.
- J.-B. Melin, J. G. Bartlett, and J. Delabrouille. Catalog extraction in SZ cluster surveys: a matched filter approach. *A&A*, 459:341–352, November 2006. doi: 10.1051/0004-6361:20065034.

- F. Menanteau, J. P. Hughes, L. F. Barrientos, et al. Southern Cosmology Survey. II. Massive Optically Selected Clusters from 70 Square Degrees of the Sunyaev-Zel'dovich Effect Common Survey Area. *ApJS*, 191:340–351, December 2010. doi: 10.1088/0067-0049/191/2/340.
- A. Merloni, P. Predehl, W. Becker, et al. eROSITA Science Book: Mapping the Structure of the Energetic Universe. *arXiv e-prints*, art. arXiv:1209.3114, Sep 2012.
- Andrei Mesinger, Matthew McQuinn, and David N. Spergel. The kinetic Sunyaev-Zel'dovich signal from inhomogeneous reionization: a parameter space study. *MNRAS*, 422(2): 1403–1417, May 2012. doi: 10.1111/j.1365-2966.2012.20713.x.
- G. Miknaitis, G. Pignata, A. Rest, et al. The ESSENCE Supernova Survey: Survey Optimization, Observations, and Supernova Photometry. *ApJ*, 666:674–693, September 2007. doi: 10.1086/519986.
- L. M. Mocanu, T. M. Crawford, J. D. Vieira, et al. Extragalactic Millimeter-wave Point-source Catalog, Number Counts and Statistics from 771 deg² of the SPT-SZ Survey. *ApJ*, 779:61, December 2013. doi: 10.1088/0004-637X/779/1/61.
- P. M. Motl, E. J. Hallman, J. O. Burns, and M. L. Norman. The Integrated Sunyaev-Zeldovich Effect as a Superior Method for Measuring the Mass of Clusters of Galaxies. *ApJ*, 623: L63–L66, April 2005. doi: 10.1086/430144.
- A. Nadolski, J. D. Vieira, J. A. Sobrin, et al. Broadband, millimeter-wave antireflection coatings for large-format, cryogenic aluminum oxide optics. *Appl. Opt.*, 59(10):3285, April 2020. doi: 10.1364/AO.383921.
- Julio F. Navarro, Carlos S. Frenk, and Simon D. M. White. Simulations of X-ray clusters. *MNRAS*, 275(3):720–740, August 1995. doi: 10.1093/mnras/275.3.720.
- Hien Trong Nguyen, John Kovac, Peter Ade, et al. BICEP2/SPUD: searching for inflation with degree scale polarimetry from the South Pole. In William D. Duncan, Wayne S. Holland, Stafford Withington, and Jonas Zmuidzinas, editors, *Millimeter and Submillimeter Detectors and Instrumentation for Astronomy IV*, volume 7020 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 70201F, July 2008. doi: 10.1117/12.787997.
- S. Nozawa, N. Itoh, Y. Kawana, and Y. Kohyama. Relativistic Corrections to the Sunyaev-Zeldovich Effect for Clusters of Galaxies. IV. Analytic Fitting Formula for the Numerical Results. *ApJ*, 536:31–35, June 2000. doi: 10.1086/308938.
- A. N. O'Brien, R. P. Norris, N. F. H. Tothill, and M. D. Filipović. The spatial correlation of bent-tail galaxies and galaxy clusters. *MNRAS*, 481:5247–5262, December 2018. doi: 10.1093/mnras/sty2642.

- Roger O'Brient, Peter Ade, Kam Arnold, et al. A dual-polarized multichroic antenna-coupled TES bolometer for terrestrial CMB Polarimetry. In Wayne S. Holland and Jonas Zmuidzinas, editors, *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy V*, volume 7741 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 77410J, July 2010. doi: 10.1117/12.857801.
- J. P. Ostriker and E. T. Vishniac. Generation of microwave background fluctuations from nonlinear perturbations at the era of galaxy formation. *ApJ*, 306:L51, July 1986. URL http://adsabs.harvard.edu/cgi-bin/nph-bib_query?bibcode=1986ApJ...306L..51O&db_key=AST.
- A. Pascut and T. J. Ponman. The Chandra Deep Group Survey - cool core evolution in groups and clusters of galaxies. *MNRAS*, 447(4):3723–3744, Mar 2015. doi: 10.1093/mnras/stu2688.
- M. Pierre, C. Adami, M. Birkinshaw, et al. The XXL survey: First results and future. *Astronomische Nachrichten*, 338(334):334–341, Mar 2017. doi: 10.1002/asna.201713352.
- G. L. Pilbratt, J. R. Riedinger, T. Passvogel, et al. Herschel Space Observatory. An ESA facility for far-infrared and submillimetre astronomy. *A&A*, 518:L1, July 2010. doi: 10.1051/0004-6361/201014759.
- A. Pillepich, C. Porciani, and T. H. Reiprich. The X-ray cluster survey with eRosita: forecasts for cosmology, cluster physics and primordial non-Gaussianity. *MNRAS*, 422:44–69, May 2012. doi: 10.1111/j.1365-2966.2012.20443.x.
- Planck Collaboration, P. A. R. Ade, N. Aghanim, et al. Planck 2015 results. XXVII. The second Planck catalogue of Sunyaev-Zeldovich sources. *A&A*, 594:A27, Sep 2016a. doi: 10.1051/0004-6361/201525823.
- Planck Collaboration, P. A. R. Ade, N. Aghanim, et al. Planck 2015 results. V. LFI calibration. *A&A*, 594:A5, September 2016b. doi: 10.1051/0004-6361/201526632.
- Planck Collaboration, Y. Akrami, M. Ashdown, et al. Planck intermediate results. LII. Planet flux densities. *A&A*, 607:A122, November 2017. doi: 10.1051/0004-6361/201630311.
- Planck Collaboration, N. Aghanim, Y. Akrami, et al. Planck 2018 results. VI. Cosmological parameters. *ArXiv:1807.06209*, July 2018.
- Daniel H. Raguin and George Michael Morris. Analysis of antireflection-structured surfaces with continuous one-dimensional surface profiles. *Appl. Opt.*, 32(14):2582–2598, May 1993. doi: 10.1364/AO.32.002582.
- C. L. Reichardt, B. Stalder, L. E. Bleem, et al. Galaxy Clusters Discovered via the Sunyaev-Zel'dovich Effect in the First 720 Square Degrees of the South Pole Telescope Survey. *ApJ*, 763:127, February 2013. doi: 10.1088/0004-637X/763/2/127.

- C. L. Reichardt, S. Patil, P. A. R. Ade, et al. An Improved Measurement of the Secondary Cosmic Microwave Background Anisotropies from the SPT-SZ + SPTpol Surveys. *ApJ*, 908(2):199, February 2021. doi: 10.3847/1538-4357/abd407.
- Y. Rephaeli. Comptonization of the cosmic microwave background: The sunyaev-zel'dovich effect. *ARA&A*, 33:541, 1995.
- A. Rest, C. Stubbs, A. C. Becker, et al. Testing LMC Microlensing Scenarios: The Discrimination Power of the SuperMACHO Microlensing Survey. *ApJ*, 634:1103–1115, December 2005. doi: 10.1086/497060.
- J. Ruel, G. Bazin, M. Bayliss, et al. Optical Spectroscopy and Velocity Dispersions of Galaxy Clusters from the SPT-SZ Survey. *ApJ*, 792:45, September 2014. doi: 10.1088/0004-637X/792/1/45.
- M. C. Runyan and W. C. Jones. Thermal conductivity of thermally-isolating polymeric and composite structural support materials between 0.3 and 4 K. *Cryogenics*, 48:448–454, September 2008. doi: 10.1016/j.cryogenics.2008.06.002.
- M. C. Runyan, P. A. R. Ade, R. S. Bhatia, et al. ACBAR: The Arcminute Cosmology Bolometer Array Receiver. *ApJS*, 149(2):265–287, December 2003. doi: 10.1086/379099.
- J. Ruze. The effect of aperture errors on the antenna radiation pattern. *Il Nuovo Cimento*, 9(3):364–380, March 1952. doi: 10.1007/BF02903409.
- J. Ruze. Antenna Tolerance Theory – A Review. *IEEE Proceedings*, 54:633–642, April 1966.
- Megh Nad Saha. Liii. ionization in the solar chromosphere. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 40(238):472–488, 1920. doi: 10.1080/14786441008636148. URL <https://doi.org/10.1080/14786441008636148>.
- J. Sayers, T. Mroczkowski, M. Zemcov, et al. A Measurement of the Kinetic Sunyaev-Zel'dovich Signal Toward MACS J0717.5+3745. *ApJ*, 778(1):52, November 2013. doi: 10.1088/0004-637X/778/1/52.
- K. K. Schaffer, T. M. Crawford, K. A. Aird, et al. The First Public Release of South Pole Telescope Data: Maps of a 95 deg² Field from 2008 Observations. *ApJ*, 743:90, December 2011. doi: 10.1088/0004-637X/743/1/90.
- Arne Schroder, Axel Murk, Pavel Yagoubov, and Ferdinand Patt. Design and Characterization of the ALMA Band 5 Vacuum Window. *IEEE Transactions on Terahertz Science and Technology*, 6(1):156–162, January 2016. doi: 10.1109/TTHZ.2015.2497139.
- D. Schwan, P. A. R. Ade, K. Basu, et al. Invited Article: Millimeter-wave bolometer array receiver for the Atacama pathfinder experiment Sunyaev-Zel'dovich (APEX-SZ) instrument. *Review of Scientific Instruments*, 82(9):091301–091301, September 2011. doi: 10.1063/1.3637460.

- E. Shirokoff, C. L. Reichardt, L. Shaw, et al. Improved Constraints on Cosmic Microwave Background Secondary Anisotropies from the Complete 2008 South Pole Telescope Data. *ApJ*, 736:61–+, July 2011. doi: 10.1088/0004-637X/736/1/61.
- Joseph Silk. Cosmic Black-Body Radiation and Galaxy Formation. *ApJ*, 151:459, February 1968. doi: 10.1086/149449.
- M. F. Skrutskie, R. M. Cutri, R. Stiening, et al. The Two Micron All Sky Survey (2MASS). *AJ*, 131:1163–1183, February 2006. doi: 10.1086/498708.
- J. A. Sobrin, A. J. Anderson, A. N. Bender, et al. The Design and Integrated Performance of SPT-3G. *ApJS*, 258(2):42, February 2022. doi: 10.3847/1538-4365/ac374f.
- Joshua Sobrin and SPT-3G Team. The First Galaxy Cluster Catalog from the SPT-3G 1500 deg² Survey. In *APS April Meeting Abstracts*, volume 2022 of *APS Meeting Abstracts*, page G14.002, April 2022.
- J. Song, A. Zenteno, B. Stalder, et al. Redshifts, Sample Purity, and BCG Positions for the Galaxy Cluster Catalog from the First 720 Square Degrees of the South Pole Telescope Survey. *ApJ*, 761:22, December 2012. doi: 10.1088/0004-637X/761/1/22.
- B. Stalder, A. A. Stark, S. M. Amato, et al. PISCO: the Parallel Imager for Southern Cosmology Observations. In *Ground-based and Airborne Instrumentation for Astronomy V*, volume 9147 of *Proc. SPIE*, page 91473Y, July 2014. doi: 10.1117/12.2054933.
- R. A. Sunyaev and Y. B. Zel'dovich. The Spectrum of Primordial Radiation, its Distortions and their Significance. *Comments on Astrophysics and Space Physics*, 2:66–+, March 1970.
- R. A. Sunyaev and Y. B. Zel'dovich. The Observations of Relic Radiation as a Test of the Nature of X-Ray Radiation from the Clusters of Galaxies. *Comments on Astrophysics and Space Physics*, 4:173–+, November 1972.
- R.A. Sunyaev and Y.B. Zel'dovich. Microwave background radiation as a probe of the contemporary structure and history of the universe. *ARAA*, 18:537, 1980. doi: 10.1146/annurev.aa.18.090180.002541.
- A. Suzuki, K. Arnold, J. Edwards, et al. Multi-Chroic Dual-Polarization Bolometric Detectors for Studies of the Cosmic Microwave Background. *Journal of Low Temperature Physics*, 176(5-6):650–656, September 2014. doi: 10.1007/s10909-013-1049-5.
- A. Suzuki, P. Ade, Y. Akiba, et al. The Polarbear-2 and the Simons Array Experiments. *Journal of Low Temperature Physics*, 184(3-4):805–810, August 2016. doi: 10.1007/s10909-015-1425-4.

- Ryota Takaku, Shaul Hanany, Yurika Hoshino, et al. Demonstration of anti-reflective structures over a large area for CMB polarization experiments. In Jonas Zmuidzinas and Jian-Rong Gao, editors, *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy X*, volume 11453 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 114531A, December 2020. doi: 10.1117/12.2562028.
- Jeremy Tinker, Andrey V. Kravtsov, Anatoly Klypin, et al. Toward a Halo Mass Function for Precision Cosmology: The Limits of Universality. *ApJ*, 688(2):709–728, December 2008. doi: 10.1086/591439.
- Takayuki Tomaru, Masashi Hazumi, Adrian T. Lee, et al. The POLARBEAR-2 experiment. In Wayne S. Holland and Jonas Zmuidzinas, editors, *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VI*, volume 8452 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 84521H, September 2012. doi: 10.1117/12.926158.
- Hy Trac, Nianyi Chen, Ian Holst, et al. AMBER: A Semi-numerical Abundance Matching Box for the Epoch of Reionization. *ApJ*, 927(2):186, March 2022. doi: 10.3847/1538-4357/ac5116.
- R. Šuhada, J. Song, H. Böhringer, et al. The XMM-BCS galaxy cluster survey. I. The X-ray selected cluster catalog from the initial 6 deg². *A&A*, 537:A39, January 2012. doi: 10.1051/0004-6361/201117214.
- K. Vanderlinde, T. M. Crawford, T. de Haan, et al. Galaxy Clusters Selected with the Sunyaev-Zel’dovich Effect from 2008 South Pole Telescope Observations. *ApJ*, 722:1180–1196, October 2010. doi: 10.1088/0004-637X/722/2/1180.
- W. Voges, B. Aschenbach, T. Boller, et al. The ROSAT all-sky survey bright source catalogue. *A&A*, 349:389–405, September 1999.
- R. Williamson, B. A. Benson, F. W. High, et al. A Sunyaev-Zel’dovich-selected Sample of the Most Massive Galaxy Clusters in the 2500 deg² South Pole Telescope Survey. *ApJ*, 738:139–+, September 2011. doi: 10.1088/0004-637X/738/2/139.
- Michael J. Wilson, Blake D. Sherwin, J. Colin Hill, et al. Atacama Cosmology Telescope: A measurement of the thermal Sunyaev-Zel’dovich effect using the skewness of the CMB temperature distribution. *Phys. Rev. D*, 86(12):122005, December 2012. doi: 10.1103/PhysRevD.86.122005.
- WMAP Science Team. URL <https://map.gsfc.nasa.gov/media/060915/index.html>.
- Ki Won Yoon. Private communication, 2016.

Yongda Zhu, George D. Becker, Sarah E. I. Bosman, et al. Long Dark Gaps in the Ly β Forest at $z < 6$: Evidence of Ultra-late Reionization from XQR-30 Spectra. *ApJ*, 932(2): 76, June 2022. doi: 10.3847/1538-4357/ac6e60.

P.A. Zyla et al. Review of Particle Physics. *PTEP*, 2020(8):083C01, 2020. doi: 10.1093/ptep/ptaa104. and 2021 update.

Appendix A

The Cluster Catalog

In this appendix, we present the full catalog of galaxy cluster candidates.

Table A.1: Galaxy cluster candidates with $\xi \geq 4.6$

SPT ID	R.A.	Decl.	Best		Redshift	M_{500c} ($10^{14}h_{70}^{-1}M_{\odot}$)	Imaging	P_{blank}	Notes
	(J2000)	(J2000)	ξ	θ_c					
SPT-CL J0000-5748	0.2479	-57.8081	14.64	0.25	0.702	$4.72^{+0.50}_{-0.59}$	3	0.002	1
SPT-CL J0000-6020	0.0323	-60.3405	7.18	0.50	0.762 ± 0.049	$2.90^{+0.39}_{-0.46}$	1	...	
SPT-CL J0001-5440	0.4132	-54.6695	9.11	0.50	0.820 ± 0.082	$3.37^{+0.40}_{-0.48}$	3	0.009	
SPT-CL J0001-5614	0.4862	-56.2410	5.41	0.25	0.428 ± 0.036	$2.55^{+0.45}_{-0.46}$	3	...	
SPT-CL J0002-5017	0.6515	-50.2889	5.45	0.25	0.901 ± 0.033	$2.25^{+0.40}_{-0.41}$	1	...	
SPT-CL J0002-5214	0.5985	-52.2388	5.88	0.25	3	0.445	2
SPT-CL J0002-5557	0.5048	-55.9624	7.22	0.75	1.150 ± 0.097	$2.62^{+0.36}_{-0.41}$	3	0.016	
SPT-CL J2259-5301	344.8284	-53.0308	5.08	0.25	1.160 ± 0.094	$1.98^{+0.34}_{-0.41}$	3	0.014	
SPT-CL J2259-5349	344.7941	-53.8236	7.17	0.25	0.258 ± 0.023	$3.28^{+0.44}_{-0.52}$	3	...	
SPT-CL J2259-5431	344.9783	-54.5260	7.79	0.75	0.390 ± 0.043	$3.38^{+0.42}_{-0.51}$	3	...	3
SPT-CL J2300-5331	345.1749	-53.5208	10.69	1.00	0.262	$4.31^{+0.49}_{-0.58}$	3	...	
SPT-CL J2300-5617	345.0003	-56.2849	9.53	0.25	0.153	$4.11^{+0.47}_{-0.58}$	3	...	4
SPT-CL J2301-5317	345.3371	-53.2843	5.26	1.50	0.348 ± 0.025	$2.54^{+0.44}_{-0.49}$	3	...	
SPT-CL J2301-5546	345.4486	-55.7759	5.45	0.75	0.748	$2.35^{+0.40}_{-0.44}$	3	0.005	
SPT-CL J2303-5114	345.8057	-51.2406	5.72	0.50	0.288 ± 0.023	$2.76^{+0.45}_{-0.50}$	3	...	
SPT-CL J2304-5007	346.0036	-50.1167	4.79	0.50	0.590 ± 0.030	$2.22^{+0.39}_{-0.46}$	3	...	
SPT-CL J2304-5718	346.1080	-57.3099	6.25	0.25	0.897 ± 0.033	$2.53^{+0.38}_{-0.42}$	3	0.064	
SPT-CL J2305-5719	346.2706	-57.3261	5.67	0.25	0.654 ± 0.043	$2.48^{+0.43}_{-0.44}$	3	...	
SPT-CL J2306-5120	346.6121	-51.3465	8.11	0.50	1.260 ± 0.102	$2.78^{+0.36}_{-0.40}$	3	0.001	
SPT-CL J2307-5440	346.7843	-54.6681	5.50	0.75	0.688 ± 0.044	$2.40^{+0.43}_{-0.43}$	3	...	
SPT-CL J2309-5710	347.2520	-57.1777	6.17	0.25	0.364 ± 0.035	$2.87^{+0.44}_{-0.49}$	3	...	
SPT-CL J2310-5239	347.7022	-52.6609	4.91	0.50	*	...	2	0.047	
SPT-CL J2310-5919	347.5696	-59.3203	6.12	0.25	0.768 ± 0.049	$2.56^{+0.40}_{-0.44}$	3	0.020	
SPT-CL J2311-4944	347.8888	-49.7397	4.95	1.50	0.584 ± 0.041	$2.29^{+0.41}_{-0.47}$	1	...	
SPT-CL J2311-5522	347.8989	-55.3724	7.47	0.75	0.217 ± 0.022	$3.42^{+0.45}_{-0.52}$	3	...	
SPT-CL J2311-5820	347.9955	-58.3363	5.47	0.25	0.930 ± 0.087	$2.25^{+0.40}_{-0.41}$	3	0.020	
SPT-CL J2312-5101	348.2421	-51.0292	4.81	1.00	0.350 ± 0.024	$2.38^{+0.42}_{-0.49}$	3	...	
SPT-CL J2314-5554	348.5355	-55.9016	4.88	0.25	0.712 ± 0.044	$2.18^{+0.38}_{-0.46}$	3	0.074	
SPT-CL J2316-5027	349.1854	-50.4550	4.73	0.50	1.120 ± 0.092	$1.90^{+0.33}_{-0.40}$	3	0.025	
SPT-CL J2316-5454	349.2113	-54.9020	9.18	1.00	0.371 ± 0.035	$3.80^{+0.44}_{-0.53}$	3	...	
SPT-CL J2317-5000	349.3259	-50.0018	4.87	0.25	1.110 ± 0.090	$1.95^{+0.34}_{-0.41}$	3	0.002	

SPT ID	R.A. (J2000)	Decl. (J2000)	Best		Redshift	M_{500c} ($10^{14}h_{70}^{-1}M_{\odot}$)	Imaging	P_{blank}	Notes
			ξ	θ_c					
SPT-CL J2317-5357	349.3365	-53.9649	6.17	0.25	0.395 ± 0.035	$2.85^{+0.44}_{-0.48}$	3	...	
SPT-CL J2317-5852	349.4379	-58.8807	5.08	0.50	0.594 ± 0.042	$2.32^{+0.40}_{-0.47}$	3	...	
SPT-CL J2318-5059	349.5380	-50.9834	4.90	0.75	0.350 ± 0.024	$2.41^{+0.42}_{-0.49}$	3	...	
SPT-CL J2318-5617	349.7028	-56.2885	5.55	0.25	0.545 ± 0.039	$2.53^{+0.42}_{-0.46}$	3	...	
SPT-CL J2319-5842	349.8642	-58.7125	6.94	0.50	0.298 ± 0.023	$3.20^{+0.42}_{-0.50}$	3	...	
SPT-CL J2320-5233	350.1251	-52.5641	6.45	0.25	0.755 ± 0.046	$2.68^{+0.38}_{-0.44}$	3	0.001	
SPT-CL J2320-5807	350.0694	-58.1311	5.80	1.00	0.562 ± 0.040	$2.59^{+0.44}_{-0.45}$	3	...	
SPT-CL J2321-5418	350.3925	-54.3109	4.68	0.75	3	0.445	5
SPT-CL J2323-5752	350.8787	-57.8831	5.14	0.25	1.300 ± 0.097	$1.92^{+0.35}_{-0.40}$	3	0.006	
SPT-CL J2325-5116	351.3778	-51.2824	5.01	0.25	0.940 ± 0.082	$2.08^{+0.36}_{-0.43}$	3	0.002	
SPT-CL J2325-5316	351.4188	-53.2766	4.60	0.25	0.350 ± 0.026	$2.31^{+0.40}_{-0.49}$	3	...	
SPT-CL J2325-5815	351.3448	-58.2545	5.72	0.25	0.556 ± 0.039	$2.57^{+0.42}_{-0.47}$	3	...	
SPT-CL J2327-5137	351.7801	-51.6231	6.22	1.25	0.338 ± 0.024	$2.90^{+0.44}_{-0.49}$	3	...	
SPT-CL J2328-5533	352.1780	-55.5605	7.88	0.25	0.773 ± 0.031	$3.08^{+0.39}_{-0.46}$	3	0.016	
SPT-CL J2328-5550	352.0383	-55.8466	4.65	0.75	0.770 ± 0.033	$2.07^{+0.35}_{-0.45}$	3	0.003	
SPT-CL J2329-5831	352.4712	-58.5241	10.81	0.50	0.719 ± 0.045	$3.87^{+0.44}_{-0.52}$	3	0.001	
SPT-CL J2330-5955	352.5054	-59.9275	4.90	2.50	3	0.445	
SPT-CL J2331-5052	352.9767	-50.8720	19.08	0.50	0.576	$5.75^{+0.58}_{-0.70}$	3	...	1,6
SPT-CL J2331-5736	352.8991	-57.6143	8.40	0.50	1.380 ± 0.102	$2.77^{+0.34}_{-0.39}$	3	0.055	7
SPT-CL J2332-5220	353.1338	-52.3482	4.95	1.00	0.460 ± 0.026	$2.36^{+0.41}_{-0.47}$	3	...	
SPT-CL J2332-5358	353.1076	-53.9745	18.25	1.50	0.402	$5.83^{+0.59}_{-0.71}$	3	...	1,8
SPT-CL J2334-5308	353.5188	-53.1397	6.02	0.25	1.200 ± 0.094	$2.25^{+0.38}_{-0.38}$	3	0.001	
SPT-CL J2334-5938	353.6940	-59.6479	7.09	0.25	0.400 ± 0.027	$3.15^{+0.42}_{-0.49}$	3	...	
SPT-CL J2335-5434	353.8866	-54.5795	5.26	0.75	1.030 ± 0.088	$2.11^{+0.37}_{-0.42}$	3	0.059	
SPT-CL J2336-5252	354.0877	-52.8725	6.74	0.25	1.220 ± 0.094	$2.45^{+0.35}_{-0.41}$	3	0.009	
SPT-CL J2336-5352	354.0112	-53.8683	5.86	0.75	0.549 ± 0.039	$2.63^{+0.42}_{-0.47}$	3	...	
SPT-CL J2337-5912	354.3994	-59.2048	6.25	1.00	0.599 ± 0.041	$2.73^{+0.40}_{-0.46}$	3	...	
SPT-CL J2337-5942	354.3532	-59.7074	38.90	0.25	0.775	$8.32^{+0.82}_{-0.96}$	3	0.000	1
SPT-CL J2339-5008	354.9618	-50.1427	4.81	0.25	*	...	2	0.005	
SPT-CL J2339-5550	354.8686	-55.8416	5.08	0.50	0.403 ± 0.036	$2.44^{+0.43}_{-0.49}$	3	...	
SPT-CL J2340-5958	355.0851	-59.9711	4.96	0.25	3	0.758	
SPT-CL J2341-5119	355.3003	-51.3299	22.05	0.25	1.003	$5.61^{+0.58}_{-0.67}$	3	0.002	

SPT ID	R.A. (J2000)	Decl. (J2000)	Best		Redshift	M_{500c} ($10^{14}h_{70}^{-1}M_{\odot}$)	Imaging	P_{blank}	Notes
			ξ	θ_c					
SPT-CL J2341-5138	355.4472	-51.6395	5.42	0.75	0.567 ± 0.040	$2.45^{+0.45}_{-0.45}$	3	...	
SPT-CL J2341-5640	355.4535	-56.6678	5.93	0.25	0.480 ± 0.038	$2.70^{+0.42}_{-0.48}$	3	...	
SPT-CL J2341-5724	355.3527	-57.4161	13.31	0.50	1.259	$3.85^{+0.41}_{-0.49}$	3	0.004	
SPT-CL J2342-5411	355.6940	-54.1869	10.42	0.25	1.075	$3.44^{+0.40}_{-0.47}$	3	0.022	1
SPT-CL J2343-5024	355.8396	-50.4016	9.03	0.25	0.879 ± 0.033	$3.29^{+0.39}_{-0.47}$	3	0.009	
SPT-CL J2344-5655	356.0958	-56.9292	4.85	0.50	0.500 ± 0.038	$2.31^{+0.41}_{-0.47}$	3	...	
SPT-CL J2349-5113	357.3841	-51.2259	6.05	0.75	0.416 ± 0.036	$2.79^{+0.45}_{-0.46}$	3	...	
SPT-CL J2349-5138	357.4703	-51.6417	4.77	0.25	3	0.857	
SPT-CL J2349-5140	357.2968	-51.6744	5.05	0.25	3	0.887	
SPT-CL J2350-5107	357.6060	-51.1253	4.63	0.75	3	0.960	
SPT-CL J2350-5301	357.7262	-53.0218	10.11	0.75	0.541 ± 0.039	$3.88^{+0.45}_{-0.53}$	3	...	
SPT-CL J2351-5005	357.8298	-50.0936	4.94	0.75	0.580 ± 0.041	$2.28^{+0.40}_{-0.47}$	3	...	
SPT-CL J2351-5452	357.9044	-54.8825	11.40	0.75	0.384	$4.37^{+0.47}_{-0.58}$	3	...	1
SPT-CL J2352-5251	358.2098	-52.8631	5.06	1.00	0.470 ± 0.037	$2.39^{+0.42}_{-0.48}$	3	...	
SPT-CL J2354-5106	358.5402	-51.1013	6.02	0.50	0.308 ± 0.024	$2.85^{+0.47}_{-0.48}$	3	...	
SPT-CL J2354-5632	358.7191	-56.5499	8.21	0.75	0.563 ± 0.040	$3.37^{+0.40}_{-0.49}$	3	...	
SPT-CL J2355-5055	358.9464	-50.9325	8.39	1.00	0.320	$3.64^{+0.44}_{-0.52}$	3	...	
SPT-CL J2355-5156	358.8482	-51.9474	8.24	0.25	0.704 ± 0.044	$3.24^{+0.39}_{-0.47}$	3	0.007	
SPT-CL J2355-5258	358.9372	-52.9833	6.97	0.25	0.712 ± 0.044	$2.87^{+0.40}_{-0.46}$	3	0.007	
SPT-CL J2355-5514	358.8668	-55.2497	4.91	0.25	1.320 ± 0.098	$1.84^{+0.32}_{-0.40}$	3	0.005	
SPT-CL J2355-5850	358.9619	-58.8468	4.94	0.25	0.970 ± 0.084	$2.04^{+0.37}_{-0.43}$	3	0.001	
SPT-CL J2355-6002	358.7942	-60.0428	4.78	0.25	1.220 ± 0.095	$1.85^{+0.32}_{-0.40}$	3	0.049	
SPT-CL J2357-5421	359.2691	-54.3594	7.06	0.25	0.920 ± 0.081	$2.75^{+0.37}_{-0.43}$	3	0.014	
SPT-CL J2357-5953	359.2865	-59.8988	4.67	0.75	3	0.768	
SPT-CL J2358-5229	359.5318	-52.4840	7.69	0.50	0.638 ± 0.042	$3.15^{+0.41}_{-0.47}$	3	...	
SPT-CL J2359-5010	359.9302	-50.1708	9.69	0.50	0.775	$3.55^{+0.41}_{-0.50}$	3	0.003	1

SPT ID	R.A.	Decl.	Best		Redshift	M_{500c}	Imaging	P_{blank}	Notes
	(J2000)	(J2000)	ξ	θ_c		($10^{14}h_{70}^{-1}M_{\odot}$)			

The first column is the ID of the candidate. The second and third column give the position. The fourth column gives the signal-to-noise ratio at which the cluster is found, and the fifth column gives the scale of the β profile which maximizes the signal-to-noise ratio. The sixth column lists the redshift, and the seventh column gives the calculated mass for clusters with measured redshift. The eighth column indicates our available followup for each cluster. 1 means optical imaging only, 2 means near-infrared only, and 3 indicates both optical and infrared imaging. The ninth column has the statistic P_{blank} (the probability that the near-infrared data are consistent with a blank field, see §5.3.2) for clusters in the SSDF without measured redshifts, and those with $z > 0.7$. Redshifts derived from spectroscopic data are reported without uncertainties, but they are typically accurate to $\sim 0.1\%$ (see Ruel et al., 2014; Bayliss et al., 2016). The data in this catalog will be hosted at <http://pole.uchicago.edu/public/data/sptsz-clusters>.

¹ Strong-lensing cluster.

² There is a group of galaxies at $z = 0.44 \pm 0.04$.

³ There may be a background cluster at $z \approx 1.1$.

⁴ Cluster masses at low redshift ($z < 0.25$) are only approximate. See §5.2.4.

⁵ Bright star impedes confirmation.

⁶ This cluster is currently undergoing a merger; see §5.4.1.

⁷ There is also a foreground cluster at $z = 0.29 \pm 0.02$.

⁸ The mass is biased low by a factor of ~ 1.5 owing to contamination from a magnified high-redshift dusty star forming galaxy; see B15 and Andersson et al. (2011) for details.

* Confirmed cluster based on infrared imaging only. See §5.3.2.

Table A.2: SPTPol clusters matched to other catalogs.

SPT ID	Catalogs with Matches	SPTPol Redshift	Literature Redshift	Redshift Reference
SPT-CL J0000-5748	1	0.70	0.70	[1]
SPT-CL J0001-5440	1	0.82	0.82	[1]
SPT-CL J0002-5557	1	1.15	1.15	[1]
SPT-CL J2259-5431	1,2	0.39	0.39	[1]
SPT-CL J2300-5331	1,2,3,4	0.26	0.26	[1]
SPT-CL J2300-5617	1,2,3,6,7	0.15	0.15	[1]
SPT-CL J2301-5546	1	0.75	0.75	[1]
SPT-CL J2306-5120	1	1.26	1.26	[1]
SPT-CL J2311-5522	3	0.22
SPT-CL J2311-5820	1	0.93	0.93	[1]
SPT-CL J2312-5101	8	0.35	0.33	[8]
SPT-CL J2314-5554	8	0.71	0.75	[8]
SPT-CL J2316-5454	1,8,9,10,11,12	0.37	0.37	[1]
SPT-CL J2317-5357	12	0.40	0.38	[12]
SPT-CL J2318-5059	8	0.35	0.33	[8]
SPT-CL J2318-5617	8,9,12	0.55	0.55	[12]
SPT-CL J2321-5418	1
SPT-CL J2325-5316	8,12	0.35	0.37	[12]
SPT-CL J2327-5137	1	0.34	0.34	[1]
SPT-CL J2328-5533	10,12	0.77	0.81	[12]
SPT-CL J2328-5550	10,12	0.77	0.80	[12]
SPT-CL J2329-5831	1	0.72	0.72	[1]
SPT-CL J2331-5052	1	0.58	0.58	[1]
SPT-CL J2332-5220	8,9,12	0.46	0.46	[12]
SPT-CL J2332-5358	1,2,7,8,9,12	0.40	0.40	[1]
SPT-CL J2334-5308	12	1.20	0.81	[12]
SPT-CL J2335-5434	10,12	1.03	0.67	[12]
SPT-CL J2336-5352	8,9,12	0.55	0.52	[12]
SPT-CL J2337-5912	1	0.60	0.60	[1]
SPT-CL J2337-5942	1,2	0.78	0.77	[1]
SPT-CL J2339-5550	8,12	0.40	0.38	[12]
SPT-CL J2341-5119	1,13	1.00	1.00	[1]
SPT-CL J2341-5640	12	0.48	0.47	[12]
SPT-CL J2341-5724	1	1.26	1.26	[14]
SPT-CL J2342-5411	1,12	1.07	1.08	[1]
SPT-CL J2350-5301	1,9	0.54	0.54	[1]
SPT-CL J2351-5452	1,8,9	0.38	0.38	[1]
SPT-CL J2354-5632	1,7,9	0.56	0.56	[1]
SPT-CL J2355-5055	1,7	0.32	0.32	[1]
SPT-CL J2358-5229	1	0.64	0.64	[1]
SPT-CL J2359-5010	1	0.78	0.77	[1]

SPT ID	Catalogs with Matches	SPTPol Redshift	Literature Redshift	Redshift Reference
--------	--------------------------	--------------------	------------------------	-----------------------

Cluster candidates with matches in other published catalogs. We take two objects to be a match if they are within $2'0$ ($5'0$) if their redshift is higher (lower) than $z = 0.3$, except *Planck* clusters, which are considered to be a match within $4'0$ at $z > 0.3$. When available, we provide the redshift of the matching cluster.

[1] Bleem et al. (2015b); [2] Voges et al. (1999); [3] Abell et al. (1989); [4] Skrutskie et al. (2006); [5] Jones et al. (2009); [6] Braid and MacGillivray (1978); [7] Planck Collaboration et al. (2016a); [8] Bleem et al. (2015a); [9] Menanteau et al. (2010); [10] Šuhada et al. (2012); [11] Mauch et al. (2003); [12] Adami et al. (2018); [13] Pascut and Ponman (2015); [14] Khullar et al. (2019);