Instrumentation and Analysis Miscellanea Regarding the Cosmology Large Angular Scale Surveyor

by

Matthew Aidan Petroff

A dissertation submitted to Johns Hopkins University in conformity with the requirements for the degree of Doctor of Philosophy

Baltimore, Maryland

July 2021

Abstract

The Cosmology Large Angular Scale Surveyor (CLASS) is an array of polarization-sensitive millimeter-wave telescopes that observes ~70% of the sky in frequency bands centered near 40 GHz, 90 GHz, 150 GHz, and 220 GHz from a high-altitude site in the Atacama desert of northern Chile. It seeks to measure polarization anisotropy in the cosmic microwave background (CMB), with a particular emphasis on measuring the optical depth due to reionization via large-angular-scale polarization *E*-modes, as well as searching for primordial polarization *B*-modes, a detection of which would provide strong evidence for cosmological inflation. This dissertation starts by providing an overview of physical cosmology, before describing the science goals and instrument design of CLASS. It then describes various instrument components that were developed, describes a novel 3D-printed millimeter-wave absorber, and describes the control and systems software used to operate the telescopes. Analysis efforts are then covered, specifically the modeling and detection of atmospheric circular polarization due to Zeeman-splitting of molecular oxygen emission lines in the geomagnetic field and a method of cleaning CMB foregrounds from full-sky maps that utilizes machine learning techniques.

Primary reader: Tobias A. Marriage

Secondary reader: Charles L. Bennett

Acknowledgments

I would like to thank my advisor, Prof. Toby Marriage, and my unofficial co-advisor, Prof. Chuck Bennett, for their support. I would not have been able to complete graduate school without them, and their helpful advice and thoughtful comments have helped me grow as a scientist. I would also like to thank the remaining member of my advisory committee, Prof. Stephan McCandliss, and the remaining members of my defense committee, Prof. David Chuss, Prof. Danielle Speller, and Dr. Joshua Peek.

It has been a privilege to work with the CLASS collaboration, and I extend thanks to my office-mates, Zhilei Xu and Aamir Ali, and to the other CLASS collaborators whom I have worked closely with, John Appel, Joseph Cleary, Jullianna Denes Couto, Joseph Eimer, Tom Essinger-Hileman, Katie Harrington, John Karakla, and Deniz Valle, in particular. I would also like to thank key collaborators on the published papers that are incorporated into this dissertation, specifically Graeme Addison, Karwan Rostem, Janet Weiland, and Ed Wollack.

I am grateful to my family and friends for their support and encouragement, especially to my parents; my twin brother, Chris; and my friend, Mai Nguyen.

The research described in this dissertation was supported in part by National Science Foundation Division of Astronomical Sciences grant numbers 0959349, 1429236, 1636634,

1654494, and 2034400; NASA grant numbers 80NSSC19K0526 and 80NSSC20K0445; a Space@Hopkins seed grant; and a Rowland fellowship. The work is in part based on observations obtained with Planck, an ESA science mission with instruments and contributions directly funded by ESA Member States, NASA, and Canada. The research made use of NASA's Astrophysics Data System and the Legacy Archive for Microwave Background Data Analysis (LAMBDA), part of the High Energy Astrophysics Science Archive Center (HEASARC); HEASARC/LAMBDA is a service of the Astrophysics Science Division at the NASA Goddard Space Flight Center. Part of the research utilized computational resources at the Maryland Advanced Research Computing Center (MARCC). CLASS is located in the Parque Astronómico Atacama in northern Chile under the auspices of the Agencia Nacional de Investigación y Desarrollo (ANID).

The research described in this dissertation was made possible thanks to the Astropy (Astropy Collaboration et al. 2013; Astropy Collaboration et al. 2018), Asymptote (Bowman and Hammerlindl 2008), CAMB (Lewis et al. 2000), CLASS (Blas et al. 2011), GetData (Wiebe et al. 2015), Healpix (Gorski et al. 2005), Healpy (Zonca et al. 2019; Gorski et al. 2005), Matplotlib (Hunter 2007), MSISE-00 (Hirsch 2019), Numba (Lam et al. 2015), NumPy (Walt et al. 2011; Harris et al. 2020), Pixell (Næss et al. 2021), Planck Sky Model (Delabrouille et al. 2013), PolSpice (Chon et al. 2004), PySM (Thorne et al. 2017), SciPy (Virtanen et al. 2020), TensorFlow (Abadi et al. 2016), and TensorFlow Large Model Support (Le et al. 2019) software packages.

Contents

Abstract ii

Acknowledgments iii

List of Tables ix

List of Figures x

1 Cosmology 1

1.1 Cosmic dynamics . . . 2

1.2 Inflation . . . 7

1.3 Cosmic microwave background . . . 9

1.3.1 ΛCDM model . . . 13

1.3.2 Polarization . . . 15

1.3.3 Foregrounds . . . 23

2 Cosmology Large Angular Scale Surveyor 29

2.1 Science goals . . . 30

2.2 Instrument design . . . 33

3 Bespoke Instrumentation Electronics 38

3.1 Vibration measurement . . . 38

3.1.1 Hardware design . . . 38
3.1.2 Analysis . . . 44
3.2 VPM control hardware design . . . 48
3.3 Diode thermometry readout electronics . . . 53
3.4 Optics cage warm thermometry and heaters . . . 54

4 A 3D-printed Broadband Millimeter-wave Absorber 59
4.1 Space-filling curves . . . 62
4.2 Material selection . . . 63
4.3 Electromagnetic modeling . . . 65
4.4 Fabrication . . . 69
4.5 Measurement . . . 69
4.6 Conclusions . . . 71

5 Design of Assorted Hardware 74
5.1 XPS foam infrared absorbing filters . . . 74
5.2 Photogrammetry . . . 76
5.3 3D-printed anti-reflective layer . . . 78
5.4 Silicone-based broadband absorber . . . 83
5.5 3D-printed dog-leg absorber geometry . . . 84
5.6 Smooth-walled rectangular-to-circular waveguide transition . . . 86

6 Single-track Absolute Rotary Color Optical Encoder 90
6.1 Color graph . . . 91
6.2 Generating an encoder track . . . 94
6.3 Prototype design and construction . . . 96
6.4 Discussion . . . 98

7 Control and Systems Software **100**

7.1 Network and software architecture . . . 101

7.1.1 Network layout . . . 101

7.1.2 Software structure . . . 103

7.1.3 Scheduling . . . 105

7.1.4 Databases . . . 105

7.2 Hardware interfaces . . . 108

7.2.1 Detectors . . . 109

7.2.2 Mount . . . 110

7.2.3 Variable-delay polarization modulator . . . 111

7.3 Data pipeline . . . 112

7.3.1 Acquisition . . . 112

7.3.2 Packaging . . . 113

7.3.3 Transfer . . . 118

7.4 Web interface . . . 121

7.5 Lessons learned . . . 124

7.6 Conclusions . . . 127

8 A First Detection of Atmospheric Circular Polarization at Q Band **128**

8.1 Atmospheric emission theory . . . 132

8.1.1 Layer attenuation . . . 132

8.1.2 Radiative transfer . . . 138

8.1.3 Atmosphere and magnetic field properties . . . 139

8.1.4 Primary source of polarized emission . . . 139

8.2 Simulation results . . . 141

8.3 Comparison with observations . . . 144

8.4 Conclusion 151

9 Full-sky Cosmic Microwave Background Foreground Cleaning Using Machine Learning 152

9.1 Neural network architecture 154

9.2 Training data 157

9.3 Training procedure 159

9.4 Application to Planck observations 162

9.4.1 Comparison of dust models 166

9.5 Conclusions 167

References 170

Vita 194

List of Tables

4.1 Comparison between types of low-profile graded-index absorbers 61

4.2 Dielectric function measurement results for bulk plastic samples 65

6.1 Colors that comprise color graph 94

8.1 Relative intensity factor $\xi(N, M)$ 135

8.2 Atmosphere simulation fit parameters at the CLASS observing site for the CLASS Q-band telescope 144

List of Figures

1.1 Evolution of Universe 10

1.2 Planck temperature map 11

1.3 CMB angular power spectra 14

1.4 Neutrino mass hierarchy 16

1.5 Thomson scattering 17

1.6 E- and B-mode polarization 20

1.7 CMB polarization from acoustic oscillations 21

1.8 Gravitational wave anisotropy 22

1.9 CMB polarization from gravitational waves 22

1.10 Gravitational lensing 24

1.11 Planck 353 GHz temperature map 25

1.12 CMB temperature foregrounds 26

1.13 WMAP K-band polarization map 27

1.14 CMB polarized foregrounds 27

2.1 CLASS observing site location 30

2.2 Comparison of experiment frequency and angular scale coverage 31

2.3 CMB angular power spectra for varying A_s and τ 32

2.4 CLASS optics diagram 35

2.5 VPM polarization modulation 37

3.1 Overview of accelerometer system 41
3.2 Accelerometer system readout and sensor hardware 42
3.3 Positions of accelerometers 45
3.4 Servo vibration–focal plane temperature correlation 46
3.5 VPM control system hardware requirements 49
3.6 Overview of VPM control hardware design 51
3.7 Assembled VPM control electronics 52
3.8 Diode thermometry readout electronics 55
3.9 Temperature sensor modules 57
3.10 Cage heater control box 58

4.1 Unit cell used in simulations 66
4.2 Labeled cross section of the absorber wedge 66
4.3 Results of FEA reflection simulations of absorber 67
4.4 Prototype carbon-loaded HIPS Hilbert curve absorber 70
4.5 Detailed view of the prototype absorber 70
4.6 Alternative realization of absorber geometry 71
4.7 Reflectance measurement results 72

5.1 XPS foam filters 76
5.2 Prototype with 3D-printed anti-reflective layer 80
5.3 Filter chop test setup 83
5.4 Silicone absorber 85
5.5 Dog-leg absorber geometry 86
5.6 Rectangular-to-circular waveguide transition 89

6.1 Chromaticity diagram 93
6.2 Encoder track visualization 94

6.3 Generated 0.1° encoder track . . . 95
6.4 Assembled color encoder prototype with printed test track . . . 97
6.5 Installed color encoder prototype . . . 98

7.1 Network architecture overview . . . 102
7.2 Software structure overview . . . 104
7.3 Example of SCHEDULER syntax . . . 106
7.4 Overview of synchronous data packaging process . . . 116
7.5 Overview of the CLASS data product structure . . . 119
7.6 Web status interface . . . 123
7.7 Web-scheduler-interface main screen . . . 125
7.8 Web-scheduler new-schedule interface . . . 125

8.1 Altitude dependence of Zeeman emission . . . 140
8.2 Frequency dependence (ν) of polarized atmospheric signal at zenith . . . 142
8.3 Simulated azimuth and zenith angle dependence of the atmospheric Stokes V signal . . . 143
8.4 Simulated azimuth profiles of the atmospheric V signal . . . 143
8.5 Example binned azimuth profiles shown along with sinusoidal best fit lines for three combinations of detector pairs and boresight rotation angles . . . 146
8.6 Zenith angle dependence of peak-to-peak observed ΔV signal compared to simulation . . . 147
8.7 Stokes V signal mapped onto the projected celestial sky using pointings from CLASS Q-band telescope observations . . . 150

9.1 Artificial neural network architecture overview . . . 156
9.2 Simulated maps from one example in the test set . . . 160

9.3 Temperature angular power spectra of one example of simulated maps from the test set 161
9.4 Planck maps 164
9.5 Planck temperature angular power spectra 165
9.6 Dust model comparison 168

Chapter 1

Cosmology

Cosmology is the study of the Universe as a gestalt, including its origin and past, present, and future evolution. On large scales, the Universe has been observed to be homogeneous and isotropic (e.g., Partridge and Wilkinson 1967; Smoot et al. 1992; Yadav et al. 2005; Bennett et al. 2013), so small-scale structures such as stars and galaxies can be averaged over. This homogeneity means that there are no preferred locations, so observations on large scales should not depend on the observer's location. The isotropy means that there are no preferred directions—and, by extension, no center—to the Universe, so observations on large scales should not depend on the observer's viewing direction. The *Copernican principle* states that Earth is not in a special place in the Universe (Bondi 1952). When this principle is combined with observations of the Universe on scales $\gtrsim$100 Mpc, it leads to an axiom that states that the Universe is homogeneous and isotropic on large scales; this axiom is known as the *cosmological principle* and is foundational to the study of physical cosmology.

Additional observations that are foundational to cosmology are the fact that the night sky is dark and the fact that the farther away a galaxy is, the faster it is receding. Based on Newtonian gravity, an infinitely-old universe would need to be infinitely large and contain

an infinite number of stars to avoid collapse (Halley 1720). In such a universe, every point on the sky would terminate on the surface of a star and would thus be as bright as the Sun (Olbers 1826). As is self-evident, this is not the case, and neither Halley's nor Olbers's solutions to this paradox hold up to scrutiny (Jaki 1967). To explain the absence of both gravitational collapse and a bright night sky, the Universe must thus have a finite age. The fact that the farther away a galaxy is the faster it is receding (Hubble 1929) means that the Universe is expanding. In such an expanding universe of a finite age, our own, it is only sensible to attempt to look back in time to model the universe's evolution and origin. This leads to the hot big-bang paradigm, which postulates that the early Universe was extremely hot and dense and that it has been expanding, cooling, and decreasing in density through to the present.

1.1 Cosmic dynamics

The proportionality between galactic redshift and distance was discovered by Edwin Hubble in 1929 (Hubble 1929). Redshift, z, the fractional Doppler shift of emitted light due to relative motion, is defined as

$$z \equiv \frac{\nu_e}{\nu_o} - 1 = \frac{\lambda_o}{\lambda_e} - 1, \tag{1.1}$$

where ν_o, ν_e, λ_o, and λ_e are the observed and emitted frequency and observed and emitted wavelength, respectively (Hogg 1999). By comparing the redshifts and distances of 20 galaxies, Hubble found a linear relationship between them known as Hubble's law,

$$z = H_0 \cdot r, \tag{1.2}$$

where H_0 is the Hubble constant, with dimensions of inverse time, and r is distance (Hubble 1929). The naught in H_0 refers to the present value of the proportionality between z and r; in general, this proportionality factor, H, is time dependent. This relationship, specifically the Hubble constant, quantifies the Universe's rate of expansion, and this expansion is responsible for *cosmological redshift*, the redshift that remains once the proper motions of the source and observer are accounted for. The current best estimates for H_0 are $(67.4 \pm 0.5)\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$ as measured from the cosmic microwave background by the Planck satellite (Planck Collaboration VI 2020) and $(73.2 \pm 1.3)\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$ as measured using a cosmic distance ladder based on geometric parallaxes, the period–luminosity relationship of Cepheids, and supernovae light curves (Riess et al. 2021). These measurements—and early- and late-Universe measurements in general—are in tension (Addison et al. 2018), which means that one or more of the measurements has unaccounted for systematic errors or that there is new physics to be discovered.

Besides the rate of expansion, another property inherent to our Universe is its spatial curvature. Since the Universe is homogeneous and isotropic on large scales, there are three possibilities for its curvature: flat (zero curvature), uniform positive curvature, or uniform negative curvature. The relationship that connects coordinates to the physical distance between two points is known as the metric, and the class of metrics that describe these curvatures, in spherical polar coordinates, are

$$(ds)^2 = \frac{(dr)^2}{1 - \kappa r^2} + r^2 (d\Omega)^2, \tag{1.3}$$

where

$$(d\Omega)^2 \equiv (d\theta)^2 + \sin^2\theta (d\phi)^2 \tag{1.4}$$

and

$$\kappa = \begin{cases} +1 & \text{positive curvature, spherical} \\ 0 & \text{flat, Euclidean} \\ -1 & \text{negative curvature, hyperspherical;} \end{cases} \tag{1.5}$$

κ is the curvature constant and R is the radius of curvature (Weinberg 2008). However, this still excludes time and thus does not compensate for the Universe's expansion. The Friedmann–Lemaître–Robertson–Walker metric takes into account this expansion (or contraction) and is expressed as

$$-(d\tau)^2 = -(dt)^2 + a(t)^2 \left(\frac{(dr)^2}{1-\kappa r^2} + r^2 (d\Omega)^2 \right), \tag{1.6}$$

where $a(t)$ is a scale factor that dictates how distances change over time (Friedmann 1922; Friedmann 1924; Lemaître 1927; Robertson 1935; Walker 1937). The time, t, is *cosmological proper time*, that is time as measured by an observer for whom the Universe expands uniformly; the spatial coordinates are co-moving coordinates, which are time-invariant in a perfectly homogeneous, isotropic universe.

In order to model our observed Universe, another parameter is needed, the *cosmological constant*, Λ, having dimensions (time)$^{-2}$. The source of the cosmological constant—so called dark energy—is not well understood, but it is required for models of the Universe to match observations (Perlmutter et al. 1997; Riess et al. 1998). With the cosmological constant, the *Friedmann equation*, which, using general relativity, links $a(t)$, κ, and R_0, the starting radius of curvature, is

$$\left(\frac{\dot{a}}{a} \right)^2 = \frac{8\pi G}{3} \rho(t) - \frac{\kappa}{R_0^2 a(t)^2} + \frac{\Lambda}{3}, \tag{1.7}$$

where $\rho(t)$ is the energy density of the Universe (Friedmann 1922; Friedmann 1924). This relationship describes the expansion (or contraction) of the Universe in terms of spatial

curvature, energy density, and the cosmological constant. Although very important in cosmology, the Friedmann equation does not contain enough information to fully describe how the scale factor, $a(t)$, evolves with time, since the energy density, $\rho(t)$, is also an unknown. A second important equation, the *fluid equation*, can be derived from the first law of thermodynamics and is expressed as

$$\dot{\rho} + 3\frac{\dot{a}}{a}[\rho(t) + P(t)] = 0, \tag{1.8}$$

where $P(t)$ is pressure. This equation uses the adiabatic nature of the Universe's expansion to derive a relationship between its scale factor and its energy density (Robertson 1933). By combining the Friedmann equation and the fluid equation, an equation describing the acceleration of the Universe's expansion over time can be derived:

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3}[\rho(t) + 3P(t)] + \frac{\Lambda}{3}. \tag{1.9}$$

Although this acceleration equation fully describes the Universe's expansion over time if values of ρ and P are known, the Friedmann equation and the fluid equation are a set of two independent equations with three unknowns, $a(t)$, $\rho(t)$, and $P(t)$, so another equation is needed to solve for these unknowns. This remaining equation is an equation of state, a relationship between energy density and pressure. While equations of state are complicated in general, an energy component of the Universe can be modeled as a dilute gas on cosmological scales, so its equation of state is

$$P = w\rho, \tag{1.10}$$

where w is a dimensionless number, with $w \leq 1$. However, the exact value of w is dependent on the contents of the Universe (Wainwright and Goode 1980). Using the

Friedmann equation, the fluid equation, and the equations of state of a universe's various components along with appropriate boundary conditions, the evolution of a universe's energy density, scale factor, and pressure is fully described.

The energy density, ρ, of the Universe and the cosmological constant, Λ, can be represented as dimensionless density parameters Ω_ρ and Ω_Λ by

$$\Omega_\rho \equiv \frac{8\pi G\rho_0}{3H_0^2} \tag{1.11}$$

$$\Omega_\Lambda \equiv \frac{\Lambda}{3H_0^2}, \tag{1.12}$$

where ρ_0 is the energy density evaluated at the present epoch. A third density parameter for spatial curvature, Ω_κ, is then defined by

$$\Omega_\rho + \Omega_\Lambda + \Omega_\kappa = 1. \tag{1.13}$$

Except for the addition needed to cancel dimensions, H_0^{-2}, these parameters are of forms similar to that of parts of the Friedmann Equation. Due to the Universe's expansion, distances are dependent on time. The *proper distance*, or ruler distance, between two points can be found using the Friedmann–Lemaître–Robertson–Walker metric using a fixed time, t. However, this cannot be measured directly on cosmological scales. Therefore, a more common distance measure is the *co-moving distance*, which is the proper distance multiplied by $1 + z$ for an infinitesimal co-moving distance δD_C; this redshift correction accounts for expansion. To calculate the total line-of-sight co-moving distance, these infinitesimal contributions are integrated over such that

$$D_C = D_H \int_0^z \left[\Omega_\rho(1+z)^3 + \Omega_\kappa(1+z)^2 + \Omega_\Lambda\right]^{-1/2} dz', \tag{1.14}$$

where D_H is the Hubble distance. Since co-moving distance accounts for Hubble flow, it is a fundamental distance measure in cosmology (Hogg 1999).

Currently, the most widely accepted cosmological model is the ΛCDM model. This model describes a flat universe dominated by cold dark matter, baryons, and a cosmological constant. The energy densities of these components are denoted Ω_c, Ω_b, and Ω_Λ. These are related to the previously discussed energy density, Ω_ρ, by $\Omega_\rho = \Omega_c + \Omega_b$ (Spergel et al. 2003). This model will be further discussed later in this chapter.

1.2 Inflation

In its most basic form, the widely accepted hot big-bang paradigm of the Universe is not without its shortcomings and fails to explain the flatness problem, the horizon problem, and the origin of the energy density fluctuations that produced the Universe's large-scale structure. The flatness problem is that while the Universe is nearly flat today, extrapolating back in time leads to a universe that was much, much flatter in the past, with

$$|1 - \Omega| \lesssim 1 \times 10^{-55} \tag{1.15}$$

at the Planck time, where $\Omega = \Omega_\rho + \Omega_\Lambda$ is the overall energy density. This would make our Universe an extremely unlikely occurrence, since the tiniest deviation would have resulted in a closed universe that collapsed on itself in the distant past or an open universe with a density far lower than what is observed at the present (Dicke and Peebles 1979; Guth 1981). The horizon problem stems from the homogeneity and isotropy of the Universe on large scales. The cosmic microwave background (CMB), the oldest light visible in the Universe (further discussed in the next section), is very nearly homogeneous and isotropic (Smoot et al. 1992), yet under the standard hot big-bang paradigm, opposite points on the surface of

last scattering would not be causally connected, much less in thermal equilibrium (Rindler 1956; Guth 1981). This would make the CMB's homogeneity and isotropy an extreme coincidence. Finally, the hot big-bang paradigm takes energy density fluctuations as an input parameter to predict the large-scale structure of the Universe, but it does not offer an explanation as to their origin; thus, another theory is required.

A solution to these three perceived shortcomings of the standard big-bang paradigm is *inflation*. Inflation hypothesizes that there was a period in the early Universe when its expansion was accelerating outwards, $\ddot{a} > 0$. For inflation, the acceleration equation, Equation (1.9), results in $P < -\rho/3$ when $\ddot{a} > 0$ and an equation of state parameter $w < -1/3$. Most inflation theories attribute this negative pressure to a positive cosmological constant, Λ_i, with $w = -1$. During this inflationary epoch, the energy density was dominated by a cosmological constant, so the Hubble parameter H_i remained fixed during this epoch, and the scale factor grew exponentially over time (Starobinskiĭ 1979; Guth 1981; Linde 1982; Albrecht et al. 1982).

With negative pressure, $|1 - \Omega|$ decreases over time. Since the Universe expanded exponentially during inflation,

$$|1 - \Omega(t)| \propto e^{-2H_i t}, \tag{1.16}$$

so the difference between Ω and unity decreased exponentially over time. After dozens of e-foldings of inflation, even a universe with fairly-strong initial curvature could reach the observed flatness of our current Universe. During inflation, the Universe's horizon size grows exponentially. Extrapolating back to before the massive growth in horizon size from dozens of e-foldings of inflation results in a very small horizon size before inflation began, which is in casual contact, solving the horizon problem (Guth 1981). Finally, inflation offers an explanation for the origin of the energy density fluctuations that provided the seed for

the formation of large-scale structure in the Universe. Due to inflation's rapid expansion of space, the pre-inflation universe was small enough for quantum fluctuations to affect its energy density. Thus, inflation predicts that these random quantum fluctuations seeded the Universe's large-scale structure. As the Universe expanded and gravity took over as the dominant force on large scales, regions with higher density coalesced into filaments, and regions of lower density became voids (Bond et al. 1996).

1.3 Cosmic microwave background

Approximately 13.8 billion years ago (Planck Collaboration VI 2020), the expansion history of the Universe began with inflation, and the primordial quark–gluon fluid quickly condensed into a baryon-containing plasma. For the next ∼380 000 years, the Universe expanded and cooled, with electrons binding to protons to form the first atoms—neutral hydrogen and some helium—toward the end of this period, a process known as *recombination* (Peebles 1968).[1] Prior to recombination, the primeval plasma was opaque to electromagnetic radiation due to Thomson scattering off of the abundant free electrons resulting in a short mean-free-path length for photons. However, the free-electron density rapidly decreased during recombination, increasing the mean-free-path length of photons to be greater than the Hubble length, H^{-1}, allowing photons to free stream. This decoupling created a surface of last scattering. The photons free streaming from this surface are known as the CMB, which contains an imprint of the Universe prior to decoupling, with density perturbations being the primary source of CMB anisotropy. At larger angular scales, these density perturbations produced CMB anisotropy through the Sachs–Wolfe and integrated Sachs–Wolfe effects, which result from the redshifting and blueshifting of photons due to gravitational potentials (Sachs and Wolfe 1967). At smaller angular

[1] The "re" portion of this name is a misnomer, which is kept for historical reasons.

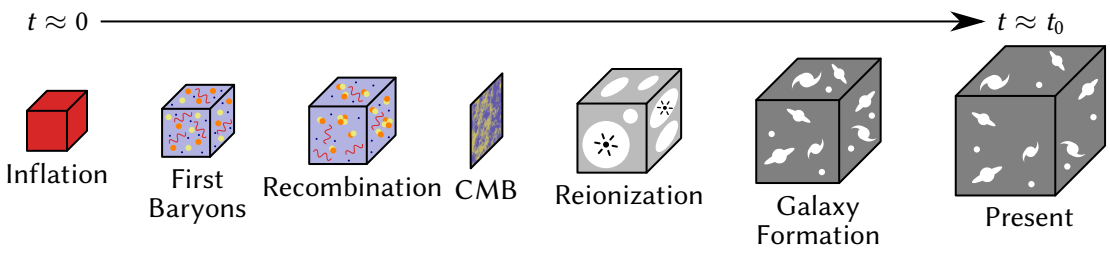

Figure 1.1: A schematic overview of the evolution of the Universe is shown. Shortly after inflation, the first baryons and atomic nuclei form. This is followed by the first atoms during recombination as the Universe expands and cools. Decoupling happens shortly thereafter, which creates a surface of last scattering, the CMB. Next, matter coalesces into the first stars, which reionize the Universe. Galaxies then form, and the Universe continues to expand through to the present. Based in part on a figure in Barkana and Loeb (2001) and a figure by Wayne Hu.

scales, the anisotropy results from acoustic oscillations of the photon–baryon fluid prior to decoupling. The CMB marks the surface of the sphere visible to a particular observer; it has a co-moving radius equal to the co-moving distance traveled by light since the Universe became transparent. As a product of the finite speed of light, it does not imply that the Universe has a center. Over the next several-hundred-million years, this neutral hydrogen—as well as dark matter—coalesced under the influence of gravity into the first stars, which reionized the Universe (Barkana and Loeb 2001). This reionization has caused additional Thomson scattering in the subsequent millennia, which has left an additional imprint in the CMB. As the Universe has continued to expand and cool through to the present, CMB photons have been redshifted such that the blackbody temperature of the CMB has decreased to its present value of 2.73 K (Fixsen 2009), which results in peak emission at millimeter wavelengths. A schematic overview of this history is shown in Figure 1.1.

The CMB was first detected as an antenna temperature excess by Penzias and Wilson (1965) and determined to be consistent with predictions for a hot big-bang cosmology (Alpher and Herman 1948; Dicke et al. 1965). An important advance was then made by the Cosmic Background Explorer (COBE) satellite mission, launched in 1989, which

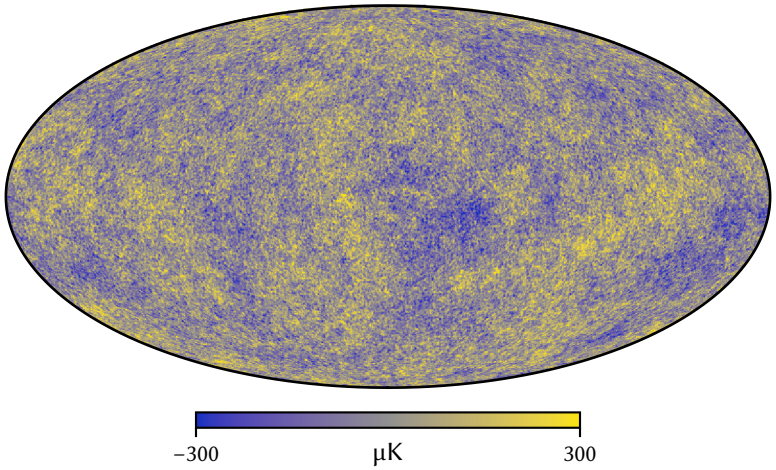

Figure 1.2: CMB temperature anisotropy as measured by the Planck satellite (Planck Collaboration I 2020) are shown in a Mollweide projection in Galactic coordinates. The monopole and dipole components have been subtracted, and masked regions have been inpainted.

made a definitive detection of CMB anisotropy (Smoot et al. 1992) and showed that the CMB is—or is nearly—a perfect blackbody (Mather et al. 1990; Fixsen et al. 1996). As COBE had low angular resolution, ground-based and balloon-borne experiments such as MAT/Toco (Torbet et al. 1999), BOOMERanG (Melchiorri et al. 2000), and MAXIMA (Hanany et al. 2000) made improved measurements at smaller angular scales on patches of the sky throughout the late 1990s. The Wilkinson Microwave Anisotropy Probe (WMAP) satellite mission (Bennett et al. 2013), launched in 2001, made the first detailed full-sky maps of the CMB, and the Planck satellite mission (Planck Collaboration I 2020), launched in 2009, improved upon these maps in angular resolution and sensitivity and increased coverage to higher frequencies. A map of CMB temperature anisotropy, as measured by Planck, is shown in Figure 1.2. In the past two decades, ground-based and balloon-borne experiments (e.g., Filippini et al. 2010; Carlstrom et al. 2011; Swetz et al. 2011; Bischoff et al. 2013; Harrington et al. 2016; BICEP2 Collaboration et al. 2018) have continued to make improved measurements of the CMB and millimeter sky in both sensitivity and angular resolution.

In order to derive useful information from the CMB anisotropy, the information contained in a sky map must be reduced to a more concise form. In particular, the statistics of the anisotropy are what is of interest to cosmology. To start, a map can be converted from pixel space to frequency space by expressing its anisotropy in terms of spherical harmonics, $Y_\ell^m(\hat{n})$,

$$\Delta T(\hat{n}) \equiv T(\hat{n}) - T_0 = \sum_{\ell m} a_{\ell m} Y_\ell^m(\hat{n}), \tag{1.17}$$

where $T(\hat{n})$ is the temperature in direction $\hat{n}$, T_0 is the CMB monopole temperature, $a_{\ell m}$ are coefficients, and the summation is over $\ell \in \mathbb{Z}^+$ and $m \in [-\ell, \ell]$. Furthermore, since temperature contains no imaginary component, $a^*_{\ell m} = a_{\ell(-m)}$. Based on the cosmological principle, rotational invariance is assumed, so the imaginary phase component of the $a_{\ell m}$ coefficients can be dropped and

$$\langle T(\hat{n}) T(\hat{n}') \rangle = \sum_{\ell m} C_\ell Y_\ell^m(\hat{n}) Y_\ell^{-m}(\hat{n}') = \sum_\ell C_\ell \left(\frac{2\ell + 1}{4\pi} \right) P_\ell(\hat{n} \cdot \hat{n}') = \sum_\ell D_\ell P_\ell(\hat{n} \cdot \hat{n}'), \tag{1.18}$$

where P_ℓ are Legendre polynomials and the C_ℓ—or, alternatively, the D_ℓ—coefficients constitute an angular power spectrum (Weinberg 2008). This spectrum contains a wealth of information to which cosmological models can be fit, generally using Markov Chain Monte Carlo techniques (Lewis and Bridle 2002). However, there is a fundamental limit to the precision of such measurements, since we only have one sky to observe.[2] There is a finite number of multipole moments at each multipole, ℓ, with the lowest ℓ, corresponding to the largest angular scales, having the fewest moments. This restriction is known as the *cosmic variance limit.*

[2]For this discussion, methods for measuring the CMB at other locations, such as that proposed by Deutsch et al. (2018), will not be considered.

1.3.1 ΛCDM model

As previously mentioned, the ΛCDM model, which assumes a flat universe dominated by cold dark matter, baryons, and a cosmological constant, is presently the most widely accepted cosmological model. This model contains six independent parameters, which can be used to derive a variety of other parameters of cosmological interest. Although various choices for the independent parameters exist, one choice is to use the physical baryon density, $\Omega_b h^2$, the physical cold dark matter density, $\Omega_c h^2$, the Hubble constant, H_0, the curvature perturbations for the pivot scale $k_0 = 0.05\,\mathrm{Mpc}^{-1}$, A_s, the scalar spectral index, n_s, and the optical depth due to reionization, τ (Spergel et al. 2003; Planck Collaboration VI 2020).[3] The effects of these parameters on the CMB's temperature angular power spectrum and E-mode polarization angular power spectrum, which will be discussed in the next section, are shown in Figure 1.3. As can be seen, different parameters have different effects on the angular power spectra, although there are degeneracies, so best-fit values for ΛCDM parameters can be determined from the CMB angular power spectra. From the independent parameters, other parameters of cosmological significance such as the age of the Universe, t_0, the dark energy density, Ω_Λ, and the density of matter fluctuations averaged over a sphere with radius $8h^{-1}\,\mathrm{Mpc}^{-1}$, σ_8, can be derived. In the standard six-parameter ΛCDM model, some parameters such as the tensor-to-scalar, r, and the sum of the three neutrino masses, $\sum m_\nu$, are held fixed, although these parameters can be allowed to vary in extended models.

The sum of the neutrino masses is extremely small but is known to be non-zero from results of neutrino oscillation experiments (Fukuda et al. 1998; Ahmad et al. 2001). Neutrinos come in three flavors, ν_e, ν_μ, and ν_τ, which correspond to the leptons that result from their interactions. Additionally, neutrinos have three possible mass states, denoted ν_1, ν_2, and ν_3, and neutrinos in each mass state interact as the three neutrino flavors

[3] Note that H_0 and h both refer to the Hubble constant, just in different units.

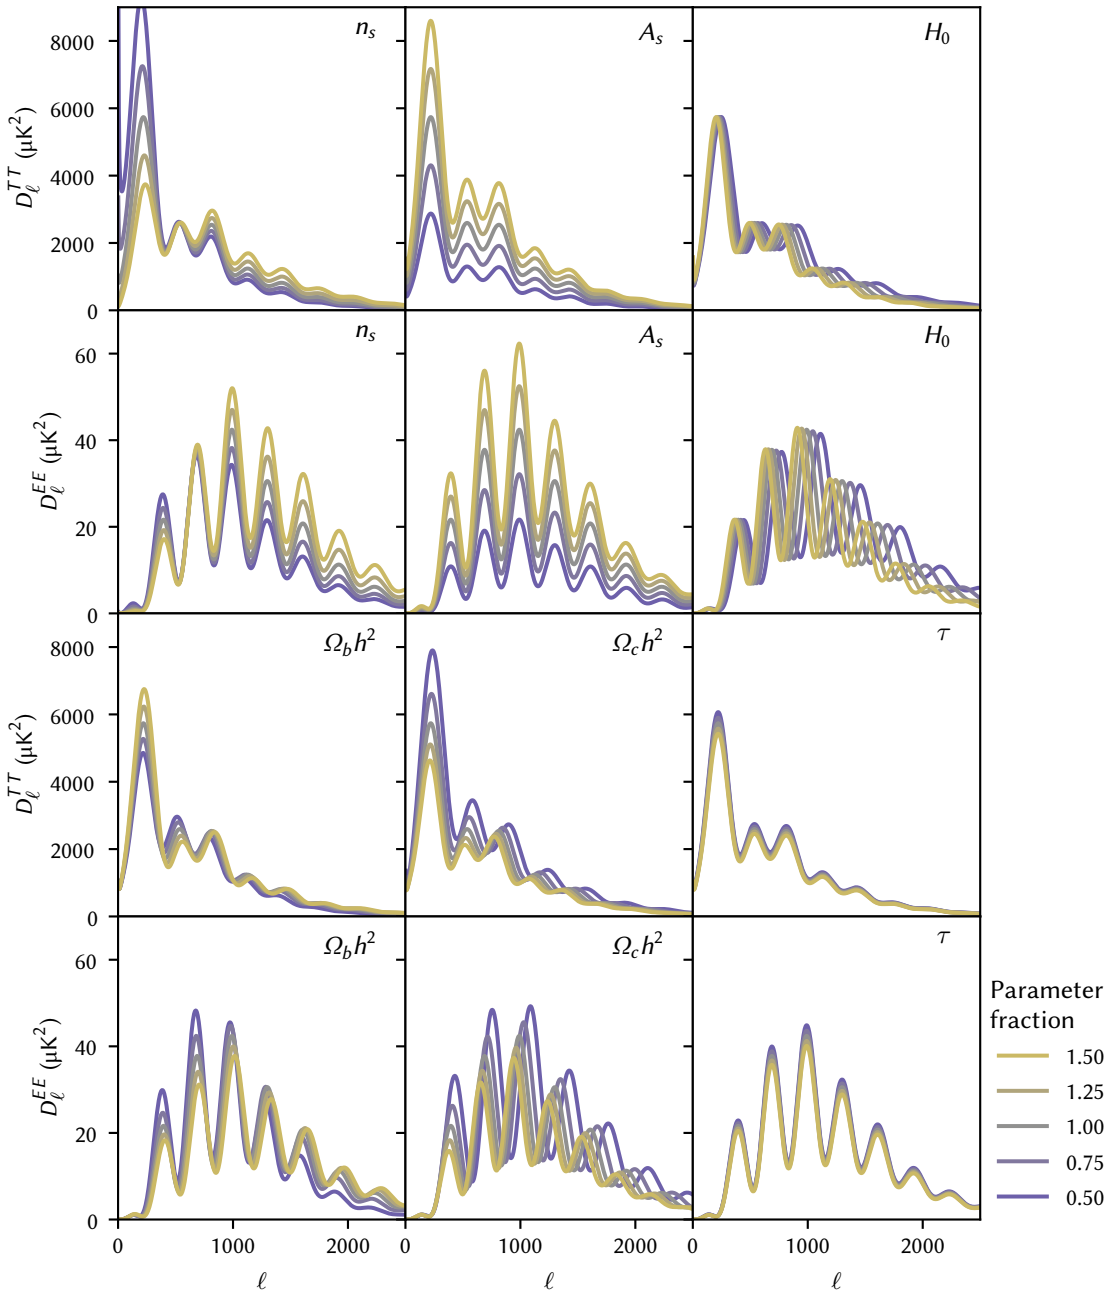

Figure 1.3: The effects of the n_s, A_s, H_0, $\Omega_b h^2$, $\Omega_c h^2$, and τ ΛCDM parameters on the CMB temperature (TT) and E-mode polarization (EE) angular power spectra are shown. For each parameter, the TT and EE spectra are plotted in separate panels with the parameter set to 0.50, 0.75, 1.00, 1.25, and 1.50 times its fiducial value with the other ΛCDM parameters held constant. The fiducial values are from the Planck 2018 results (Planck Collaboration VI 2020), and the spectra were simulated with `CAMB` (Lewis et al. 2000).

with different probabilities. In the Sun, electron neutrinos, ν_e, are produced, which are primarily in the ν_1 mass state. These neutrinos gain a tiny amount of additional mass when interacting with electrons, transitioning to mass state ν_2. Thus, measurements of solar neutrinos can be used to produce a signed mass difference, Δm^2_{sol}, which constrains the mass difference between ν_1 and ν_2. A second mass difference constraint, Δm^2_{atm}, results from measurements of neutrinos produced in Earth's upper atmosphere from its interaction with cosmic rays. These neutrinos are primarily muon neutrinos, ν_μ, which are mostly in mass state ν_3. This second mass difference, between ν_3 and ν_2, is more than an order of magnitude larger than that between ν_1 and ν_2. However, its sign is unknown, which allows for two different neutrino mass state orderings, which are known as the *normal hierarchy* and the *inverted hierarchy* (Maltoni et al. 2004); an overview of these hierarchies is shown in Figure 1.4. In the normal hierarchy, ν_3 is the heaviest mass state, while in the inverted hierarchy, it is the lightest mass state. Thus, the neutrino mass states in the normal hierarchy are ordered, from lightest to heaviest, ν_1, ν_2, ν_3, while in the inverted hierarchy they are ordered ν_3, ν_1, ν_2. This, combined with the Δm^2_{sol} and Δm^2_{atm} mass difference measurements, produces a different lower-limit prediction for $\sum m_\nu$ for each hierarchy. Therefore, a precise measurement of $\sum m_\nu$ may help to determine which hierarchy is correct, although this depends on the mass of the lowest-mass state.

1.3.2 Polarization

The initial source of the electromagnetic radiation in the CMB is thermal emission, which is unpolarized. However, the CMB is actually polarized, albeit weakly. This polarization is due to Thomson scattering (Thomson 1903; Jackson 1998) of photons off of free electrons, and a schematic overview of Thomson scattering is shown in Figure 1.5. Polarized radiation incident on an electron induces motion along the radiation's polarization axis, which, in turn, causes some of the radiation to be scattered, primarily in the direction perpendicular

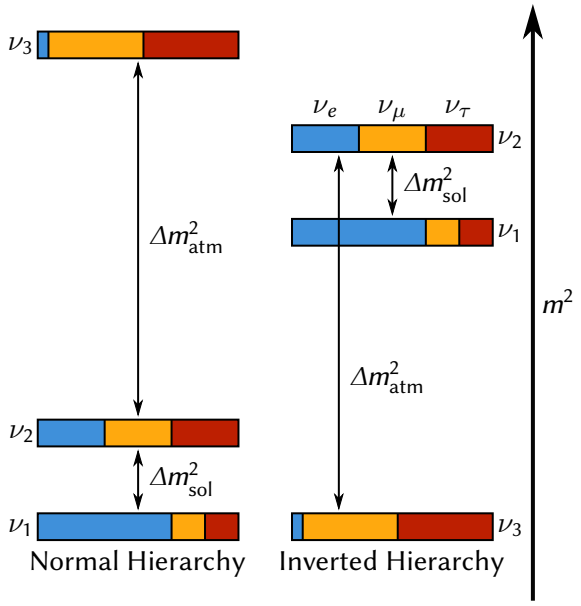

Figure 1.4: A schematic overview of the normal and inverted neutrino mass hierarchies is shown. The relative orderings of the three neutrino mass states in the two hierarchies are shown, along with the neutrino flavors that each mass state interacts as. The two hierarchies predict different neutrino mass sums, so a precise measurement of the sum of the neutrino masses may help determine which hierarchy is correct.

to both the direction of incidence and the polarization axis. While this scattered radiation is polarized, the source radiation was also polarized, so it alone does not explain the source of CMB polarization. If, instead, the incident radiation is unpolarized but still only from a single direction, this radiation induces motion in the plane perpendicular to the direction of incidence; this scatters radiation primarily in the same plane, which is polarized. While this converts a fraction of the incident unpolarized radiation into polarized radiation, it is not a realistic case, since incident radiation would be expected from all directions. However, since the scattering does not depend on the direction in which the radiation is traveling along the incident axis, incident radiation from just one direction is simply an extreme case of quadrupolar anisotropy. If unpolarized radiation is incident on a free electron in two perpendicular axes in equal amounts, the scattered radiation is unpolarized, but if the radiation intensity on these axes differs, polarization is induced.

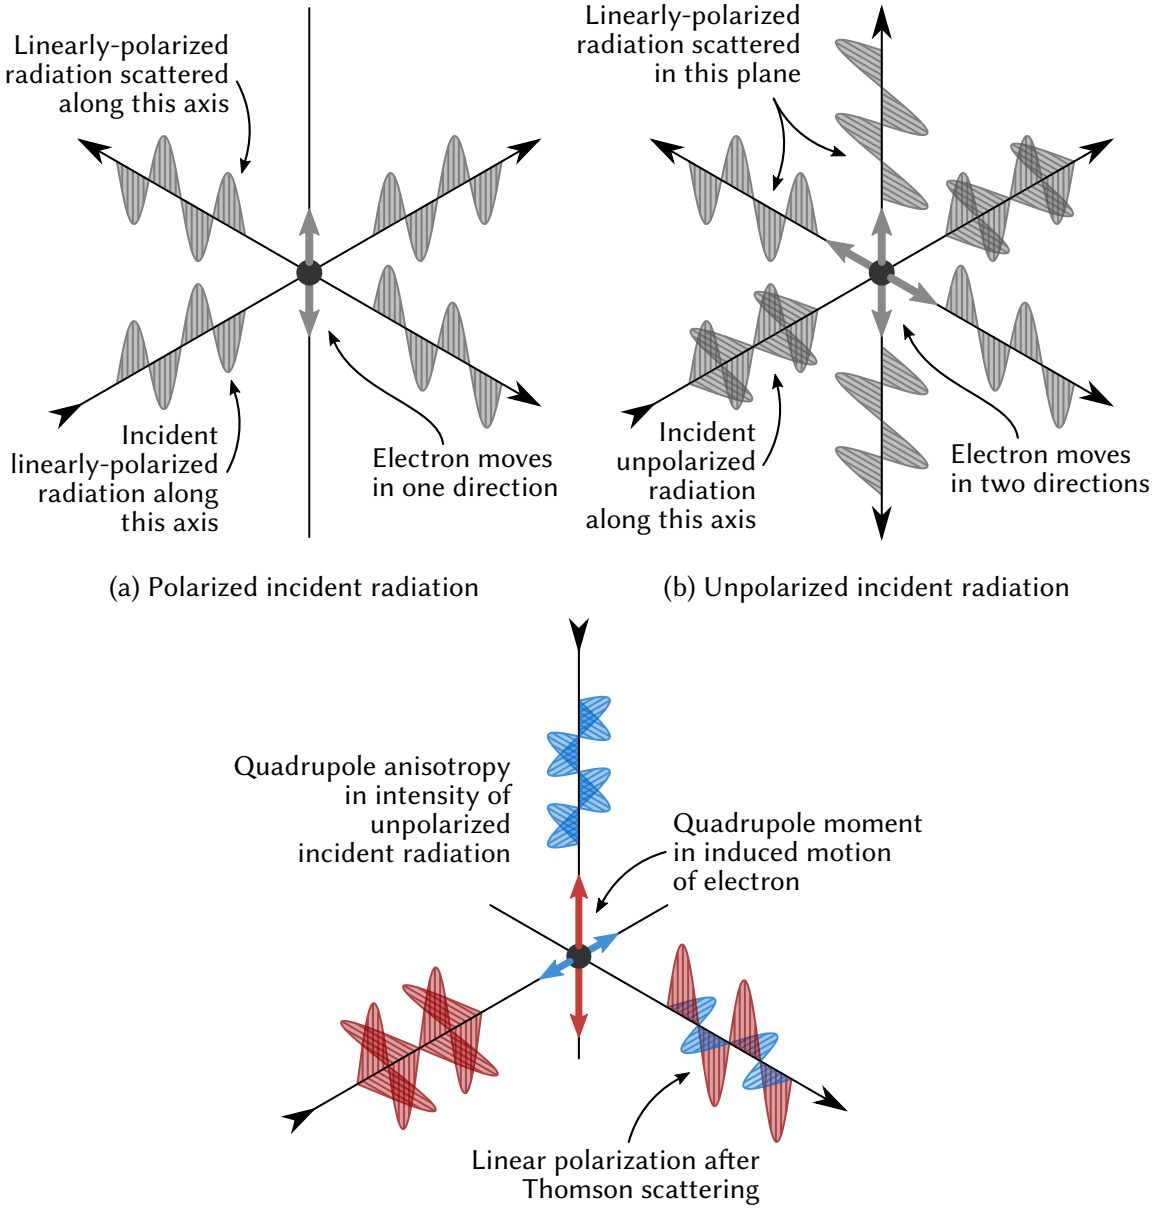

(a) Polarized incident radiation

(b) Unpolarized incident radiation

(c) Anisotropic unpolarized incident radiation

Figure 1.5: A schematic overview of Thomson scattering is shown. When polarized radiation is incident on an electron, it induces movement of the electron along the polarization axis, which scatters a fraction of the radiation primarily perpendicular to both the direction of incidence and the polarization axis; for unpolarized radiation, the induced movement and scattering is along both axes of the plane perpendicular to the direction of incidence. When the intensity of unpolarized incident radiation from multiple directions forms a quadrupole anisotropy, this produces a quadrupole moment in the induced movement of the electron, which causes the scattered radiation to be linearly polarized. Based on figures in Hedman (2002).

Since temperature anisotropy are a spin-0 field, the spherical harmonics discussed previously can fully describe it. Polarization anisotropy, however, are a spin-2 field, so the non-trivial choice of which basis functions to use must first be made before it can be described using spherical harmonics. The basis functions generally used to describe polarization are the Stokes parameters, where Q is used to describe vertical and horizontal linear polarization, U is used to describe linear polarization along the diagonals, and V is used to describe circular polarization (Born and Wolf 1959). Since these basis functions describe linear polarization in terms of a global direction reference, they depend on the coordinate system orientation, but there is no obvious choice for a coordinate system for describing polarization of cosmological significance, since the Universe is isotropic. Thus, alternative basis functions are necessary.

The standard linear polarization decomposition used for CMB analysis is the E–B decomposition, which characterizes linear polarization in terms of a scalar, curl-free E-mode component and a pseudoscalar, divergence-free B-mode component (Seljak and Zaldarriaga 1997; Kamionkowski et al. 1997). A schematic representation of this decomposition is shown in Figure 1.6. The B-mode component changes sign under a parity transformation, but the E-mode component does not. The naming of the components is derived from an analogy to electric (E) and magnetic (B) fields. Furthermore, these components are defined to be non-local, so a spot on the sky with zero polarization can still have non-zero E-mode or B-mode values, and a uniformly-polarized region can have E-mode and B-mode values of zero. The directions are defined in terms of the direction of the polarization magnitude gradient. More specifically,

$$E(\boldsymbol{\theta}) = \int d^2\widetilde{\boldsymbol{\theta}}\, \omega\left(\widetilde{\theta}\right) Q_r\left(\boldsymbol{\theta} + \widetilde{\boldsymbol{\theta}}\right) \tag{1.19}$$

$$B(\boldsymbol{\theta}) = \int d^2\widetilde{\boldsymbol{\theta}}\, \omega\left(\widetilde{\theta}\right) U_r\left(\boldsymbol{\theta} + \widetilde{\boldsymbol{\theta}}\right), \tag{1.20}$$

where

$$Q_r\left(\boldsymbol{\theta}+\widetilde{\boldsymbol{\theta}}\right)=Q\left(\boldsymbol{\theta}+\widetilde{\boldsymbol{\theta}}\right)\cos\left(2\widetilde{\phi}\right)-U\left(\boldsymbol{\theta}+\widetilde{\boldsymbol{\theta}}\right)\sin\left(2\widetilde{\phi}\right) \tag{1.21}$$

$$U_r\left(\boldsymbol{\theta}+\widetilde{\boldsymbol{\theta}}\right)=Q\left(\boldsymbol{\theta}+\widetilde{\boldsymbol{\theta}}\right)\sin\left(2\widetilde{\phi}\right)+U\left(\boldsymbol{\theta}+\widetilde{\boldsymbol{\theta}}\right)\cos\left(2\widetilde{\phi}\right) \tag{1.22}$$

are Stokes parameters redefined in terms of radial and tangential directions, $\omega(\widetilde{\theta})$ is a weight defined as

$$\omega(\widetilde{\theta})=\begin{cases}-1/\widetilde{\theta}^2 & \widetilde{\theta}>0\\ 0 & \widetilde{\theta}=0,\end{cases} \tag{1.23}$$

$\boldsymbol{\theta}$ is the direction of the point on the sky in question, $\widetilde{\boldsymbol{\theta}}$ is the direction of a nearby point on the sky, $\widetilde{\theta}$ is the angular distance between $\boldsymbol{\theta}$ and $\widetilde{\boldsymbol{\theta}}$, and $\widetilde{\phi}$ is the direction from $\boldsymbol{\theta}$ to $\widetilde{\boldsymbol{\theta}}$ (Zaldarriaga 2001). Circularly-polarized emission, which is expected to be many orders of magnitude smaller than linearly-polarized emission from theoretical predictions for the CMB (Inomata and Kamionkowski 2019), is still described by the Stokes V parameter in the E–B decomposition.

With a mechanism for producing linear polarization, Thomson scattering of photons off of free electrons with quadrupolar anisotropy in the incident radiation, and a basis with which to describe this polarization, the E–B decomposition, what remains is the source of these quadrupolar anisotropy. The acoustic oscillations that created the density perturbations responsible for the CMB temperature anisotropy also created velocity perturbations, which are out of phase from the density perturbations. The velocity gradients resulting from these velocity perturbations create quadrupolar anisotropy, which, due to Thomson scattering, produce linear polarization. Since these polarization anisotropy are sourced from the same scalar perturbations as the temperature anisotropy, the quadrupoles are aligned with the direction of the polarization magnitude gradient, resulting in pure E-mode

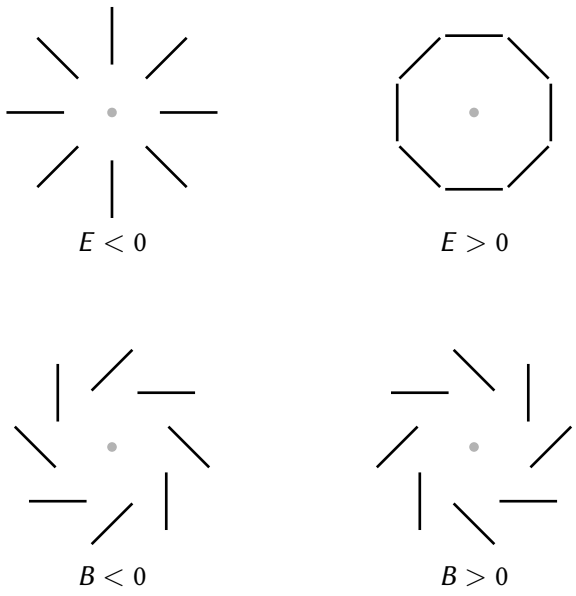

Figure 1.6: A schematic representation of E- and B-mode polarization is shown. Linear polarization forms E-modes when it is curl-free, while it forms B-modes when it is divergence-free. The designations are derived from an analogy to electric (E) and magnetic (B) fields.

polarization. An overview of this process is shown in Figure 1.7.

An alternative source of quadrupolar anisotropy is from tensor perturbations, gravitational waves in particular. As a gravitational wave propagates, it squeezes and stretches space-time, which redshifts and blueshifts the photons contained in the affected space. This directly creates quadrupolar anisotropy. An overview of this effect is shown in Figure 1.8. Crucially, these gravitational waves do not need to be aligned with the scalar density or velocity perturbations. Thus, if a gravitational wave were to intersect the surface of last scattering it could create either B-mode polarization or E-mode polarization in the CMB, depending on its alignment, which is shown in Figure 1.9.

Scalar perturbations produce only E-mode polarization, while both vector and tensor perturbations can produce B-mode polarization. Since vector perturbations—vortices—are not expected to be present (Seljak and Zaldarriaga 1997), a detection of primordial B-mode polarization would be convincing evidence of tensor perturbations. As the only expected

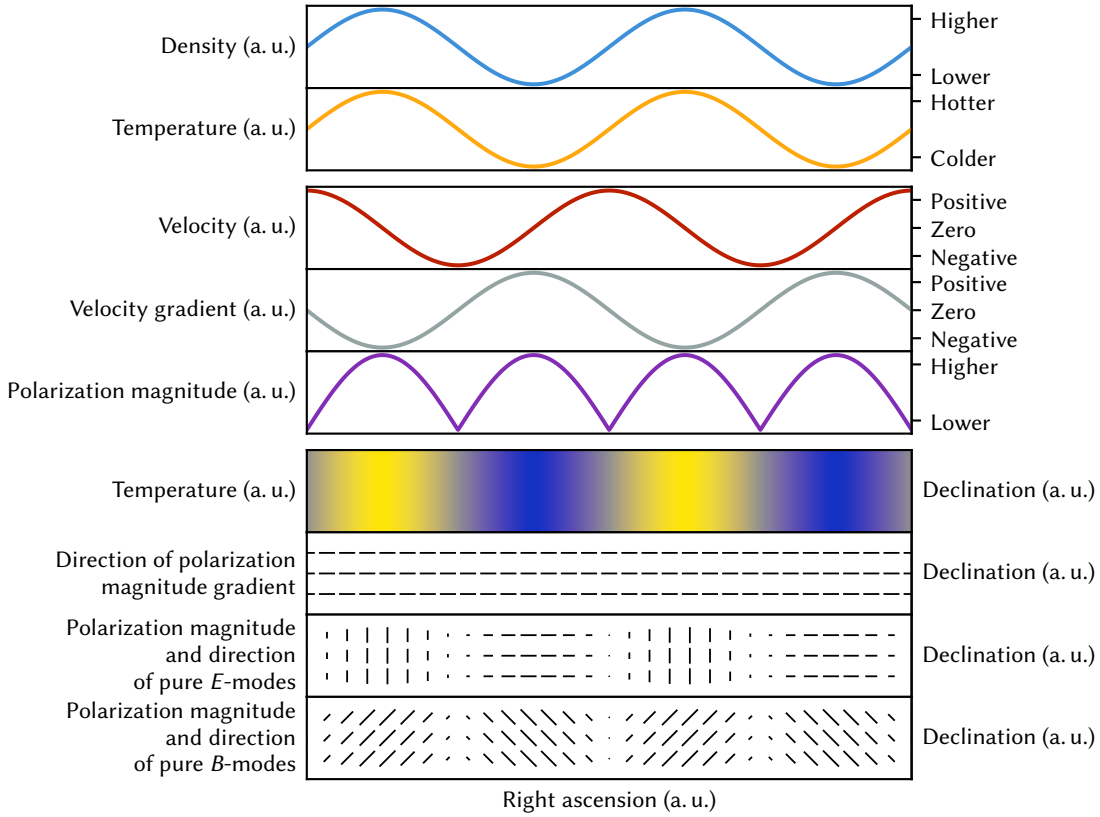

Figure 1.7: An overview of how acoustic oscillations result in CMB polarization is shown. Shown in the top panel group, acoustic oscillations produce density perturbations in the primordial plasma, which result in temperature perturbations at the surface of last scattering. These oscillations also produce velocity perturbations that are out of phase from the density perturbations, which result in velocity gradients (that are in phase with the density perturbations). The primary source of CMB polarization is from Thomson scattering resulting from quadrupolar anisotropy sourced by these velocity gradients. This is shown in the middle panel group. The bottom panel group shows the result if these perturbations are parallel to the surface of last scattering in right ascension and constant in declination. The polarization magnitude gradient is parallel to the perturbation direction, and the resulting linear polarization is pure *E*-modes, which are either parallel to or perpendicular to the polarization magnitude gradient. Also shown is the polarization direction necessary for pure *B*-modes, although *B*-modes are not sourced by the discussed scalar perturbations.

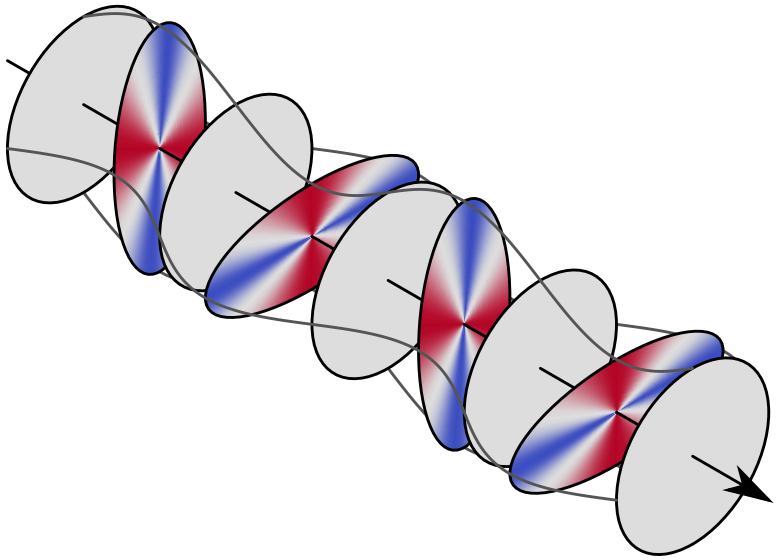

Figure 1.8: Quadrupolar anisotropy caused by gravitational waves is shown. As a gravitational wave squeezes and stretches space-time, it redshifts and blueshifts photons, creating quadrupolar anisotropy. The arrow denotes the direction of travel of the gravitational wave.

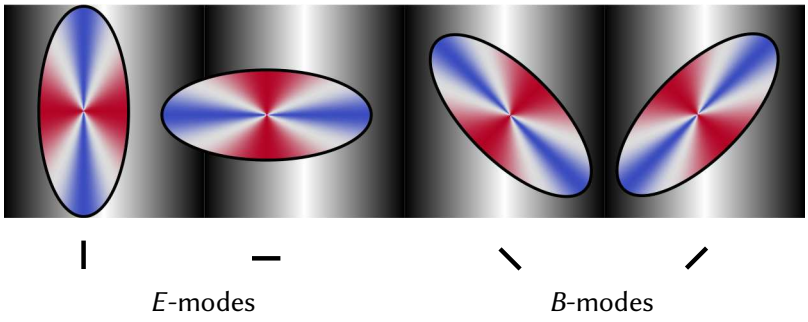

Figure 1.9: Gravitational waves with propagation directions parallel to the line of sight are shown above a grayscale gradient representing CMB temperature fluctuations on the surface of last scattering. The polarization magnitude gradient is aligned with the temperature gradient but since the orientations of the gravitational waves are independent, the gravitational waves can produce either *E*-modes or *B*-modes, depending on their alignment with the temperature fluctuations and the surface of last scattering.

source of tensor perturbations is gravitational waves and the only expected source of gravitational waves is inflation, a detection of primordial B-mode polarization would be convincing evidence of inflation. However, not all B-mode polarization is primordial. As shown in Figure 1.10, gravitational lensing can convert E-mode polarization into B-mode polarization; the reverse is also true, but this is not a concern, since primordial B-mode polarization, if it exists, is much weaker than its E-mode counterpart and would consequently have no significant effect on the scientific interpretation of the E-mode power spectrum. Due to the prevalence of galaxies and large-scale structure between our telescopes and the CMB and due to the weak primordial B-mode signal, gravitational lensing of E-mode polarization into B-mode polarization must be corrected for, particularly at the smaller angular scales where it is most prevalent. However, extragalactic gravitational potentials are not the only foregrounds obscuring the CMB, as will be discussed in the next section.

1.3.3 Foregrounds

Although of immense astrophysical interest, our Milky Way galaxy sits between us and the CMB, obscuring our view and manifesting itself as a significant foreground at the millimeter and submillimeter wavelengths used for CMB observations. As an example, Figure 1.11 shows the sky as measured by the Planck satellite's 353 GHz frequency band in temperature, where the sky is dominated by Galactic thermal dust emission (Planck Collaboration III 2020). This emission is concentrated along the Galactic plane, and toward the Galactic center in particular, but there are no regions of the sky that are completely free of contamination. Synchrotron, free–free, and spinning dust emissions also contribute to the Galactic temperature foreground, which exceeds the CMB's brightness at many frequencies. Synchrotron is emitted by high-energy electrons following Galactic magnetic fields, free–free is emitted by optically-thin interstellar plasma, and spinning dust emission,

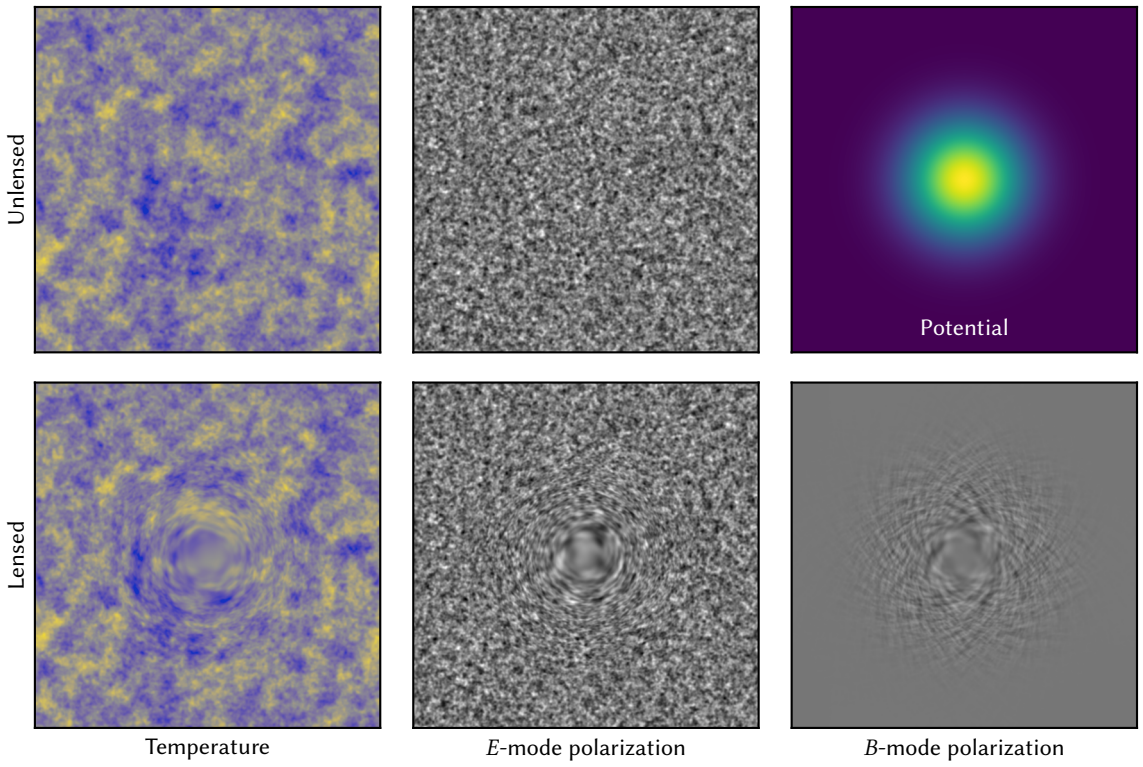

Figure 1.10: An overview of CMB gravitational lensing is shown. The top row shows temperature and *E*-mode polarization anisotropy for a flat sky patch, as well as a (greatly-exaggerated) gravitational potential; this input sky patch has no *B*-mode polarization. The bottom row shows the temperature, *E*-mode polarization, and *B*-mode polarization anisotropy when the gravitational potential lenses the input sky patch. Due to the lensing, a fraction of the *E*-mode polarization is converted into *B*-mode polarization. Based on a figure in Hu and Okamoto (2002); simulations were performed using `Pixell` (Næss et al. 2021).

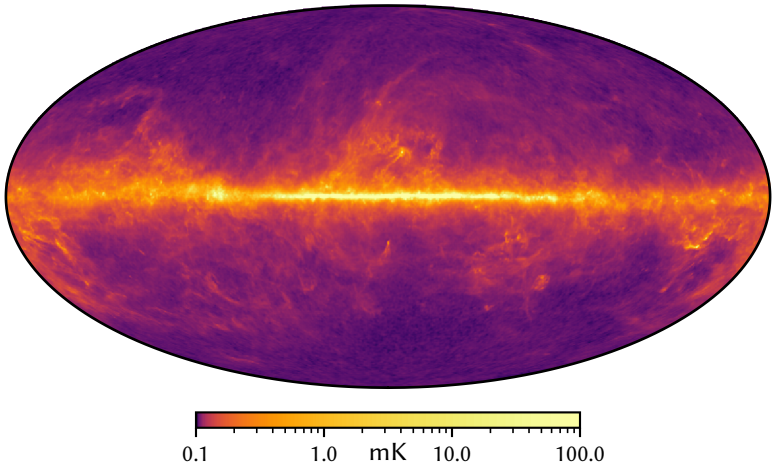

Figure 1.11: The submillimeter sky as measured by the Planck satellite's 353 GHz frequency band (Planck Collaboration III 2020) is shown in a Mollweide projection in Galactic coordinates with a logarithmic colormap. The map is smoothed with a 0.5° FWHM Gaussian window function.

sometimes referred to as (part of) anomalous microwave emission, is likely due to electric dipole radiation from rapidly spinning dust grains (Planck Collaboration XXV 2016). Fortunately, the spectral emission profiles of these foreground components differ from that of the CMB, allowing for them to be separated, similar to how a color photograph allows isoluminant objects of different colors to be differentiated between, a task that would be impossible in grayscale. The spectral emission profiles of the various Galactic temperature foregrounds are shown in Figure 1.12. The foreground emission minimum is near 100 GHz, with synchrotron, free–free, and spinning dust emission dominating at lower frequencies and thermal dust emission dominating at higher frequencies.

Galactic emission is also observable in linear polarization, as can be seen in the map of polarization magnitude and linear polarization direction made using the WMAP satellite's K-band channel (Bennett et al. 2013) and shown in Figure 1.13. This view is dominated by synchrotron emission, since the emissions due to free–free and spinning dust have negligible polarization (Planck Collaboration XXV 2016). Thermal dust emission is also polarized, due to the tendency of dust grains to align with magnetic fields. The spectral

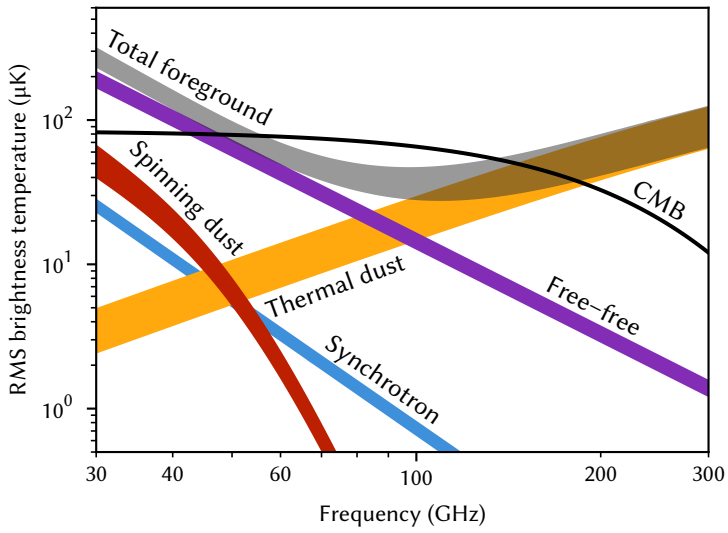

Figure 1.12: An overview of the spectral dependence of CMB temperature foregrounds is shown. Galactic synchrotron, free–free, and spinning dust emissions are significant at lower frequencies, while Galactic thermal dust emission is significant at higher frequencies. Data are from a `PySM3` (Thorne et al. 2017) fiducial simulation at `NSIDE` = 64; the plotted foreground bands display the difference between 10% and 20% sky cuts using the Planck 2015 Galactic plane masks (Planck Collaboration I 2016).

emission profiles of polarized Galactic foregrounds are shown in Figure 1.14. With the lack of polarized free–free and spinning dust emission, the foreground minimum is shifted to somewhat lower frequencies. However, the ratio of Galactic foreground emission to CMB emission is higher in linear polarization than it is in temperature.

In addition to these diffuse Galactic foregrounds, there are also extragalactic foregrounds to consider. A diffuse component that must be considered at higher frequencies is the cosmic infrared background (CIB), which contains the dust emissions of unresolved high-redshift galaxies (Planck Collaboration XXX 2014). Various point sources, both Galactic and extragalactic, must also be considered, with the strongest sources generally masked out of maps prior to higher-level analysis. Galaxy clusters also contribute through the Sunyaev–Zel′dovich effect (Zeldovich and Sunyaev 1969). As the CIB and Sunyaev–Zel′dovich effect are only very weakly polarized, they do not need to be considered for CMB

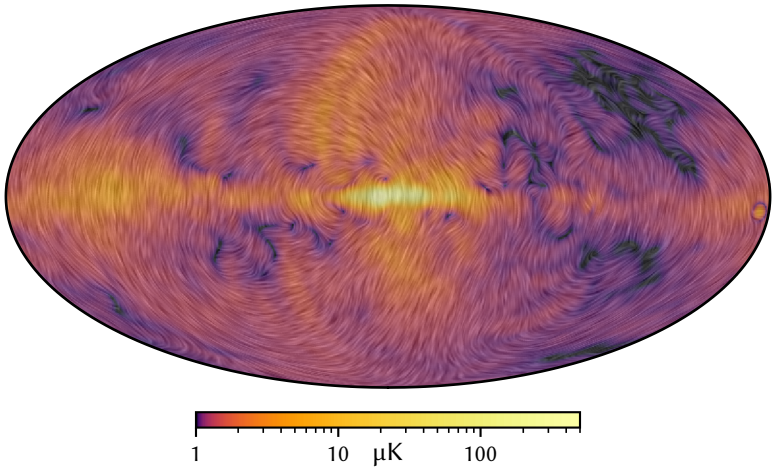

Figure 1.13: The K-band sky as measured by the WMAP satellite (Bennett et al. 2013) is shown in a Mollweide projection in Galactic coordinates. Polarization magnitude is shown with a logarithmic colormap, while a line-integral convolution is used to show polarization direction. The map is smoothed with a 5° FWHM Gaussian window function.

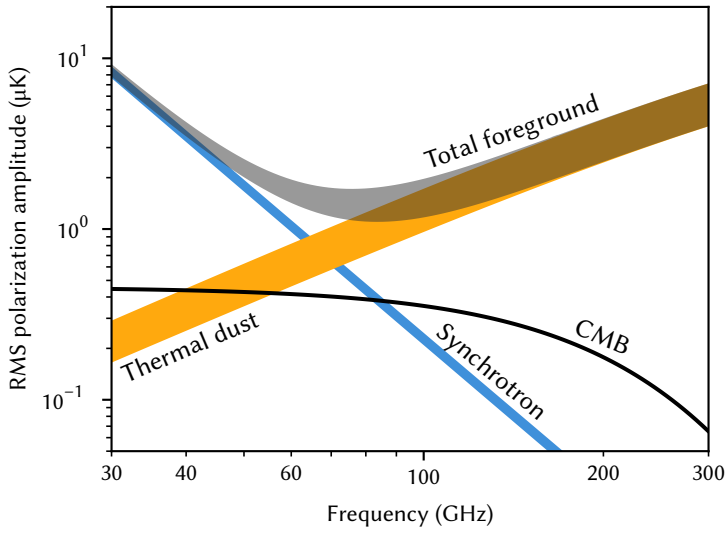

Figure 1.14: An overview of the spectral dependence of CMB polarized foregrounds is shown. Emission due to Galactic synchrotron dominates at lower frequencies, while emission from Galactic thermal dust dominates at higher frequencies. Polarized emissions from free–free and spinning dust are negligible. Data are from a `PySM3` (Thorne et al. 2017) fiducial simulation at `NSIDE` = 64; the plotted foreground bands display the difference between 10% and 20% sky cuts using the Planck 2015 Galactic plane masks (Planck Collaboration I 2016).

polarization measurements. Point sources can be polarized, but they only contribute at smaller angular scales and only do so weakly (Datta et al. 2019). Finally, for ground-based observations, the contributions of Earth's atmosphere must be considered. Fortunately, the atmosphere does not have significant linear polarization at millimeter wavelengths (Keating et al. 1998), so it is more of a concern for temperature and circular polarization observations. Atmospheric circular polarization will be further discussed in Chapter 8.

Chapter 2

Cosmology Large Angular Scale Surveyor

The Cosmology Large Angular Scale Surveyor (CLASS) is an array of millimeter-wave telescopes that observe the polarization of the CMB on large angular scales ($\ell \gtrsim 2$) from a high-altitude site in the Atacama desert of northern Chile. This location near the equator allows for ~70% of the sky to be observed on a daily basis. The sky is observed in frequency bands centered near 40 GHz, 90 GHz, 150 GHz, and 220 GHz, which straddle the polarized foreground minimum (Eimer et al. 2012; Essinger-Hileman et al. 2014; Harrington et al. 2016). The high altitude of the site, approximately 5200 m above sea level, places the telescopes above a significant fraction of the Earth's atmosphere, which reduces the detrimental effects of atmospheric emission on the observations. Furthermore, the dryness of the site reduces the effects of water-vapor emission. The location of the site and a photo of it are shown in Figure 2.1.

(a) Site location on map

(b) Site as of February 2021

Figure 2.1: The CLASS observing site is located at an approximately 5200 m elevation in the Atacama desert of northern Chile. This location near the equator allows for approximately 70% of the sky to be observed on a daily basis.

2.1 Science goals

The primary science goals of CLASS are two-fold, to detect or place an upper limit on the B-mode signal of primordial gravitational waves and to measure the optical depth due to reionization, τ, to near the cosmic variance limit (Watts et al. 2015; Watts et al. 2018). CLASS aims to achieve these goals by targeting the largest angular scales on the sky, $\ell \gtrsim 2$, while also observing at near the spectral minimum of polarized CMB foreground emission. An overview of the frequency and angular scale coverage of various current ground-based and balloon-borne CMB experiments, including CLASS, is shown in Figure 2.2. As can be seen, CLASS is one of the only suborbital experiments to target the largest angular scales and is the only one to do so at near the foreground minimum. Targeting the largest angular scales allows CLASS to constrain the effects of reionization on the CMB, since these are the angular scales most affected by reionization. The gravitational waves caused by inflation would leave a B-mode imprint on the CMB polarization during both recombination and reionization. However, the imprints left during recombination and reionization occur

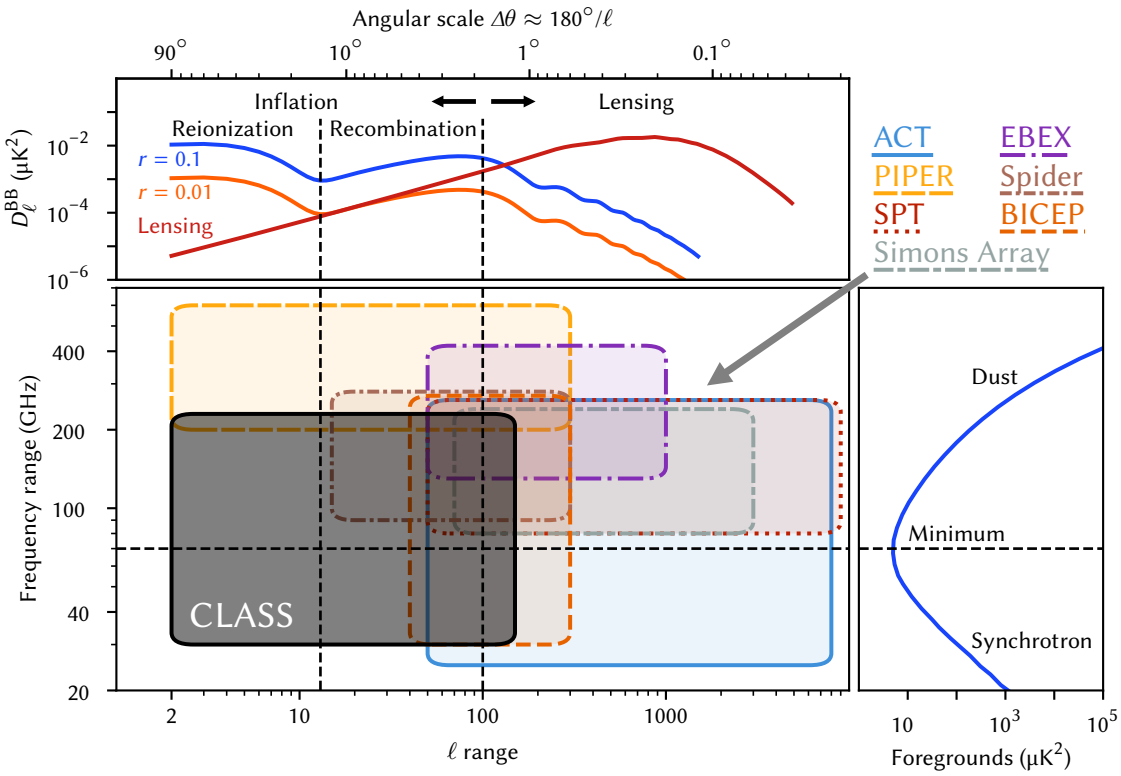

Figure 2.2: A comparison of the frequency and angular scale coverage of various current ground-based and balloon-borne CMB experiments is shown. The bottom-left panel shows the frequency and angular scale coverage of the experiments, while the right panel shows an overview of the frequency dependence of polarized CMB foregrounds, and the top panel shows the angular scale dependence of the inflationary B-mode signal for different values of the tensor-to-scalar ratio, r, as well as that of gravitational lensing. As can be seen, CLASS is unique in its position covering the largest angular scales while also straddling the CMB foreground minimum. Based on a figure by Duncan Watts.

at different angular scales, producing two different peaks in the B-mode angular power spectrum. As the reionization peak only affects $\ell \lesssim 10$, targeting the largest angular scales allows CLASS to constrain this peak.

The other crucial measurement that can only be made at the largest angular scales is that of the optical depth due to reionization, τ. As is shown in Figure 2.3, the ΛCDM parameters A_s and τ have significant degeneracy in temperature and in E-mode polarization for $\ell \gtrsim 10$. Thus, breaking this degeneracy involves measuring E-mode polarization on the

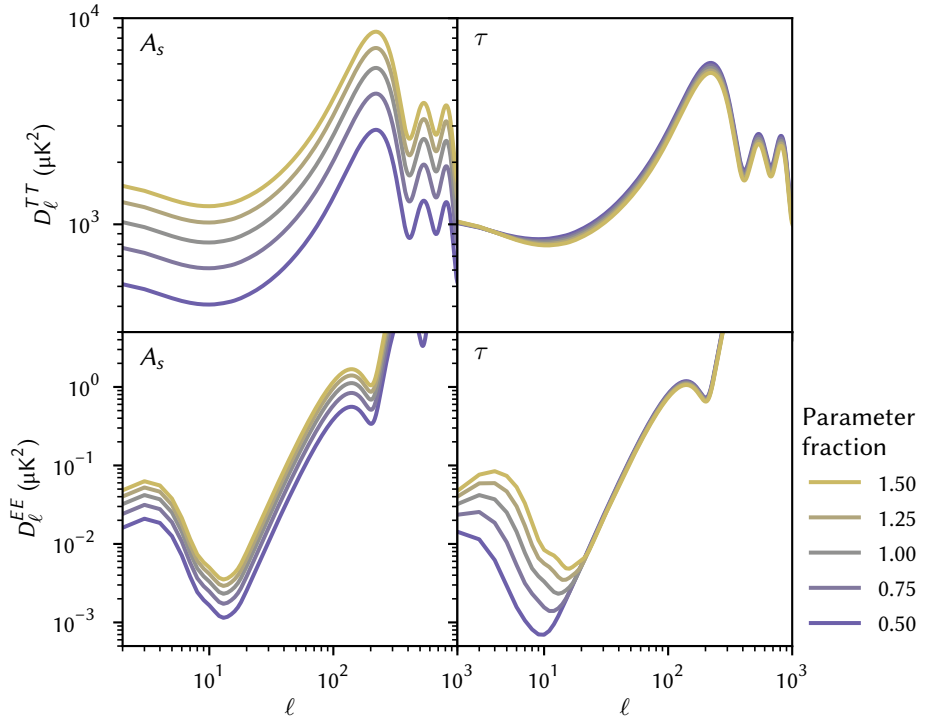

Figure 2.3: The effects of the A_s and τ ΛCDM parameters on the CMB temperature (TT) and E-mode polarization (EE) angular power spectra are shown, with a focus on the largest angular scales. For each parameter, the TT and EE spectra are plotted in the top and bottom rows, respectively, with the parameter set to 0.50, 0.75, 1.00, 1.25, and 1.50 times its fiducial value with the other ΛCDM parameters held constant. As can be seen, these two parameters are mostly degenerate in the TT spectrum but can be disentangled using the low-ℓ EE spectrum. The fiducial values are from the Planck 2018 results (Planck Collaboration VI 2020), and the spectra were simulated with `CAMB` (Lewis et al. 2000).

largest angular scales, $\ell \lesssim 10$. As τ is currently the least-constrained ΛCDM parameter, CLASS can make an important contribution by making an improved measurement.

By making an improved measurement of τ, CLASS can also contribute beyond cosmology by helping to constrain the sum of the neutrino masses, $\sum m_\nu$. When ΛCDM is extended to include $\sum m_\nu$, there are degeneracies between $\sum m_\nu$, the clustering amplitude at $8h^{-1}\,\mathrm{Mpc}^{-1}$ (σ_8), and the physical cold dark matter density ($\Omega_c h^2$). These degeneracies can be reduced with an improved measurement of the primordial co-moving curvature power spectrum amplitude, A_s. Since A_s and τ are also degenerate, particularly in temperature,

an improved measurement of τ from the polarization data helps to break this degeneracy, which results in an improved measurement of A_s. This improved measurement of A_s in turn helps to break the degeneracies involving $\sum m_\nu$, leading to an improved measurement of it (Allison et al. 2015; Watts et al. 2018).

2.2 Instrument design

The CLASS telescope array currently consists of a Q-band receiver with a center frequency near 40 GHz, a W-band receiver with a center frequency near 90 GHz, and a dichroic G-band receiver with center frequencies near 150 GHz and 220 GHz. These receivers were deployed in 2016, 2018, and 2019, respectively. A second W-band receiver is also being assembled. The receivers and associated optics are placed on alt-azimuth mounts, which additionally each contain a third axis for rotating around the boresight pointing direction. This boresight axis allows for the relative polarization angle between the telescopes and the sky to be changed. Two co-pointed receivers are placed on each mount, with the first mount containing the Q-band receiver and the existing W-band receiver. The second mount contains the G-band receiver and will contain the second W-band receiver, once it is deployed (Harrington et al. 2016).

The CLASS telescopes use an optics design based on a combination of warm reflective optics and cryogenic refractive optics. Starting from the sky, light passes through a forebaffle before reflecting off of a variable-delay polarization modulator (VPM) as the first primary optical element. This VPM, which will be described in further detail below, serves as a fast front-end polarization modulator, which switches the polarization sensitivity of the system at 10 Hz. The light then reflects off of two monolithic aluminum mirrors and is directed through an ultra-high-molecular-weight polyethylene (UHMWPE) vacuum window and into the cryogenic receiver. Once in the receiver, the light passes through a

cold stop and two refractive lenses, which focus the light into smooth-walled feedhorns (Eimer et al. 2012). In the case of the Q- and W-band receivers, the lenses are made from high-density polyethylene (HDPE), while the G-band receiver features silicon lenses. The feedhorns direct the light onto orthomode transducers, which direct the two orthogonal linear polarization states through on-chip band-defining filters and into separate transition-edge-sensor (TES) bolometers (Dahal et al. 2018). The cryogenic receivers also each feature a set of filters in the optical path, primarily for reducing thermal loading due to infrared radiation (Iuliano et al. 2018). A majority of these filters serve as radio-transparent multilayer insulation (RT-MLI), which is described in further detail in Section 5.1. An overview of this optics design is shown in Figure 2.4.

The CLASS cryostats utilize a Cryomech PT415 two-stage pulse tube cryocooler[1] in combination with a Bluefors ^{3}He/^{4}He dilution refrigerator.[2] The first pulse-tube stage is cooled to below 60 K, while the second is cooled to below 4 K. The dilution refrigerator cools the cryostat's two coldest stages, one to approximately 1 K and the second to below 100 mK. The focal plane is mounted to the 100 mK stage, while one lens each is mounted to the 1 K and 4 K stages. The stages are all horizontally cantilevered (Iuliano et al. 2018). The focal plane measures the bolometric power of the incoming radiation using feedhorn-coupled TES bolometers, which are constructed using a molybdenum–gold bilayer that can be tuned to a specific superconducting transition temperature. As these bolometers are biased to be at a temperature near the middle of their superconducting transition, minute changes in incoming power have a significant effect on the electrical resistance of the detector. Superconducting quantum interference devices (SQUIDs) are used to read out the detectors via time-division multiplexing and to amplify their signals (Dahal et al. 2018). This readout is controlled by and the signal is digitized by Multi-Channel Electronics

[1] Cryomech, Inc., Syracuse, NY; `https://www.cryomech.com/`
[2] Bluefors Oy, Helsinki, Finland; `https://bluefors.com/`

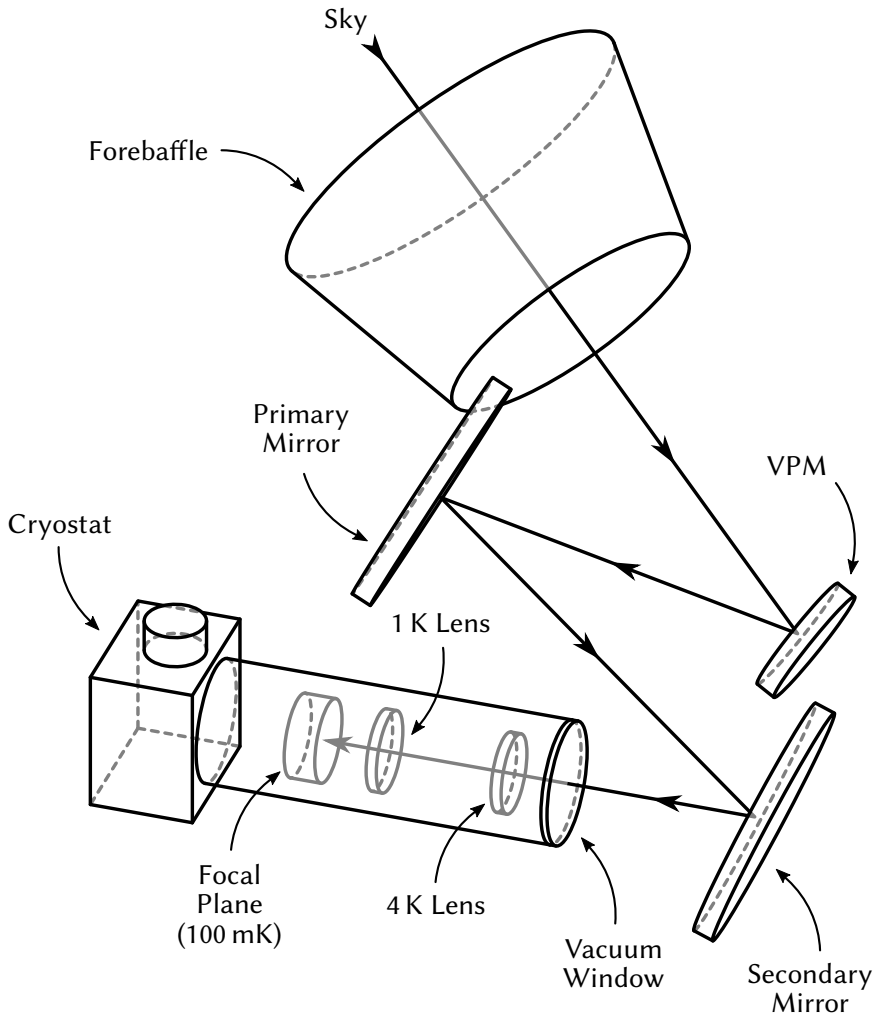

Figure 2.4: An overview of the optics design of the CLASS telescopes is shown. Starting from the sky, light passes through a forebaffle before reflecting off the VPM, which serves as the first optical element. Light then reflects off the primary mirror and the secondary mirror before passing through the vacuum window and entering the cryostat. It then passes through a stop and two cold lenses, as well as various filters, before being directed via feedhorns onto the OMTs connected to the TES detectors.

(MCE) units (Battistelli et al. 2008).

VPMs utilize a flat movable mirror placed behind a linearly polarizing array of parallel wires to induce a varying phase delay between polarization states perpendicular to and parallel to the direction of the array wires. When combined with detectors sensitive to linear polarization, modulating the mirror position, and thus the distance between the mirror and the wire array, results in the modulation of one linear polarization state, instrument Stokes U in the case of CLASS, into circular polarization, Stokes V, and vice versa (Chuss et al. 2012). Instrument Stokes parameters are defined relative to the coordinate system of the telescope's detectors. An overview of this mechanism is shown in Figure 2.5. This modulation increases polarization measurement stability by utilizing phase-sensitive detection and allows for the measurement of large angular scale modes on the sky (Miller et al. 2016). The CLASS VPMs utilize a 60 cm aluminum-honeycomb flat mirror mounted onto a mirror transport mechanism (MTM) utilizing cross-leaf rotary flexures made from tempered spring steel, which is driven by voice-coil actuators (Harrington et al. 2018). The MTM keeps the mirror parallel to the wire array throughout its throw, and the voice-coil actuators provide a non-contact force. The VPMs are designed to operate for in excess of one billion cycles, and the voice-coil actuators and spring steel flexures provide the necessary durability. An additional voice coil–flexure system is used with a reaction mass to cancel out vibrations. However, the array–mirror distance needs to be precisely controlled, which is challenging given the control input delay inherent to voice-coil actuators. This led to the development of custom control electronics, which are described in Section 3.2.

The CLASS observing strategy involves scanning in azimuth while keeping the altitude angle fixed at 45°. The telescopes scan in 720° sweeps in azimuth before switching direction. Originally, these scans were performed at $1\,^{\circ}\,\mathrm{s}^{-1}$, but the speed has since been increased to $2\,^{\circ}\,\mathrm{s}^{-1}$, which was enabled in part by a reduction in vibrations. The process used to find and fix these vibrations is described in Section 3.1. The boresight rotation angle is stepped

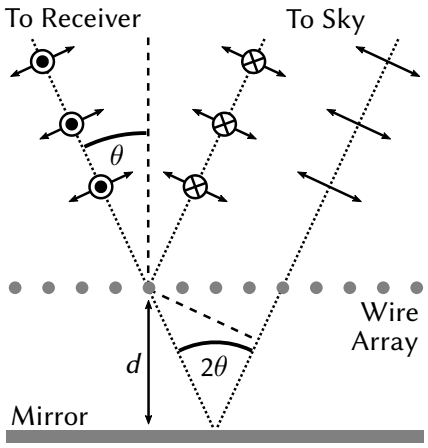

Figure 2.5: The mechanism by which the VPM modulates polarization is shown. Incident light that is polarized parallel to the wire array reflects off of it, while light that is polarized perpendicular to the wire array passes through it and is reflected off of the flat mirror that sits behind the wire array. Changing the array–mirror distance, d, affects the phase delay between the orthogonal polarization axes. Depending on the distance, the polarization can be unaffected, linear polarization can be converted into circular polarization and vice versa, or a partial conversion can take place. Modulating the array–mirror distance therefore also modulates the polarization of the incident light. Based on a figure in Harrington et al. (2018).

on a daily basis in 15° increments from −45° to +45°. Since a VPM only measures one linear Stokes parameter at a time, this boresight rotation is necessary to fully measure linear polarization on the sky, but it also serves as a crucial cross-check against systematic errors. Since the daytime sky is transparent in the millimeter-wave regime, CLASS observes day and night throughout the year, aiming for a high overall observing efficiency (Appel et al. 2019).

Chapter 3

Bespoke Instrumentation Electronics

3.1 Vibration measurement

3.1.1 Hardware design

Shortly after first light, in May 2016, it became apparent that vibrations were negatively affecting telescope performance. Excessive vibrations were causing the VPM to turn off, particularly with high wind, and were causing focal plane heating. These end effects provided little illumination as to the sources of the vibrations and were dependent on too many factors to provide an immediate indication as to whether a change helped or hurt the vibrations. The failure of early attempts to fix the vibrations demonstrated that comprehensive measurements of the vibrations were needed in order to understand them, such that they could be reduced, or ideally, eliminated. At the time, the vibrations could only be measured indirectly, via the VPM and mount encoders. These measurements are problematic, since they only provide insight into a few fixed locations, and it was unclear whether or not they provided accurate representations of the vibrations in question. Therefore, a dedicated system to directly measure the vibrations was needed.

The spectrum produced by performing a discrete Fourier transform on a time series of

acceleration measurements shows the frequencies and intensities of vibrations. Thus, a set of continuously-recording accelerometers, with appropriate bandwidth, were needed to directly measure the telescope's vibrations. Ideally, these sensors would be both small and low cost, such that it would be feasible to deploy many sensors at once to simultaneously measure vibrations at various points on the telescope. Precision calibration of the sensors, both in intensity and direction, was not necessary. Unfortunately, available accelerometer modules mostly fell into two categories: precision sensors that offered continuous readout but were prohibitively expensive or battery-powered data loggers that did not provide continuous readout. The battery-powered data loggers contained the ideal sensor, accelerometers built using microelectromechanical systems (MEMS); similar MEMS-based accelerometers are also found in smartphones, which is why they are widely available and very inexpensive. A new packaging and readout method for these sensors was therefore needed.

MEMS-based accelerometers are manufactured in small, surface-mount packages and are designed to be read out using circuitry collocated on the same printed circuit board. Most modern accelerometers provide a digital readout interface using either the Serial Peripheral Interface (SPI) bus or the Inter-Integrated Circuit (I^2C) bus. However, both buses were designed for only short-distance communications, which is troublesome for the desired accelerometer readout system. Fortunately, various integrated circuits are available to extend the operating distance of the I^2C bus, so it was chosen for this reason as well as its need for fewer signaling lines than SPI. There are existing buses that were designed for long-distance communication, such as RS-485, but these would require a microcontroller at each sensor to provide protocol conversion, which would increase complexity. The STMicroelectronics LIS3DSH three-axis digital MEMS accelerometer was chosen;[1] it provides 16-bit data output at up to 1.6 kHz. A later revision of the sensor

[1] LIS3DSH datasheet: `https://www.st.com/resource/en/datasheet/lis3dsh.pdf`

module design uses an Analog Devices ADXL355 three-axis digital MEMS accelerometer instead,[2] which is a sensor with 20-bit data output and one-sixth the readout noise density.

I^{2}C was originally developed by Philips Semiconductor in 1982 and was first standardized a decade later. For its physical layer, it uses two bidirectional open-drain lines, one as a clock line and one as a data line; both lines have pull-up resistors connected. The bus speed is up to 100 kbit s^{-1} for *standard mode* or up to 400 kbit s^{-1} for *fast mode*; arbitrary lower speeds are allowed. Nodes are either controllers or peripherals and have either a 7-bit or 10-bit address.[3] The maximum allowed bus capacitance is 400 pF, which limits maximum cable length; this is set by a combination of the 3 mA maximum drive current and the rise time requirements determined by the bus speed.

In order to extend the range of I^{2}C, the maximum allowable bus capacitance needs to be increased, which can be achieved by using higher drive currents. The P82B715 I^{2}C bus extender was chosen for this purpose.[4] By placing one device on each end of the cable, the maximum allowable bus capacitance is increased to 3000 pF. The device works using unidirectional analog current amplification to increase the drive current by a factor of ten. As Ethernet or telephone cable usually has a capacitance of 50–70 pF m^{-1}, this allows for cable lengths in excess of 30 m. However, the capacitances of different cables add, so the maximum cable length would depend on the number of sensors connected. To avoid this issue, PCA9510A I^{2}C bus buffers are used on the central readout module,[5] between each bus extender and the main I^{2}C bus. These buffers isolate each cable capacitance from the rest of the system, so the number of devices no longer affects the maximum cable length. In the original revision, telephone cable with 6P4C modular connectors was used to connect the sensor modules to the readout module, but this proved to be problematic

[2] ADXL355 datasheet: `https://www.analog.com/media/en/technical-documentation/data-sheets/adxl354_adxl355.pdf`

[3] I^{2}C-bus specification: `https://www.nxp.com/docs/en/user-guide/UM10204.pdf`

[4] P82B715 datasheet: `https://www.ti.com/lit/ds/symlink/p82b715.pdf`

[5] PCA9510A datasheet: `https://www.nxp.com/docs/en/data-sheet/PCA9510A.pdf`

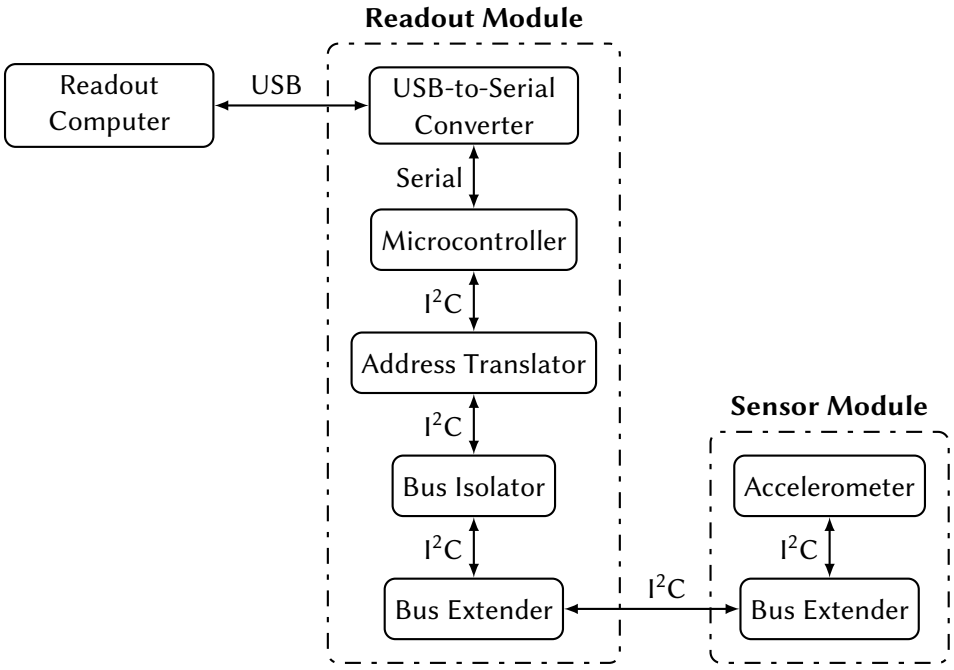

Figure 3.1: An overview of the accelerometer system is shown. The sensor modules contain an accelerometer and an I^2C bus extender and connect to the readout module via shielded Category 5e cable. The readout module contains sets of I^2C bus extender, bus isolator, and address translator chips for interfacing with the sensor modules. Finally, a microcontroller and a USB-to-serial converter are used to read out the sensors and communicate with the readout computer over USB.

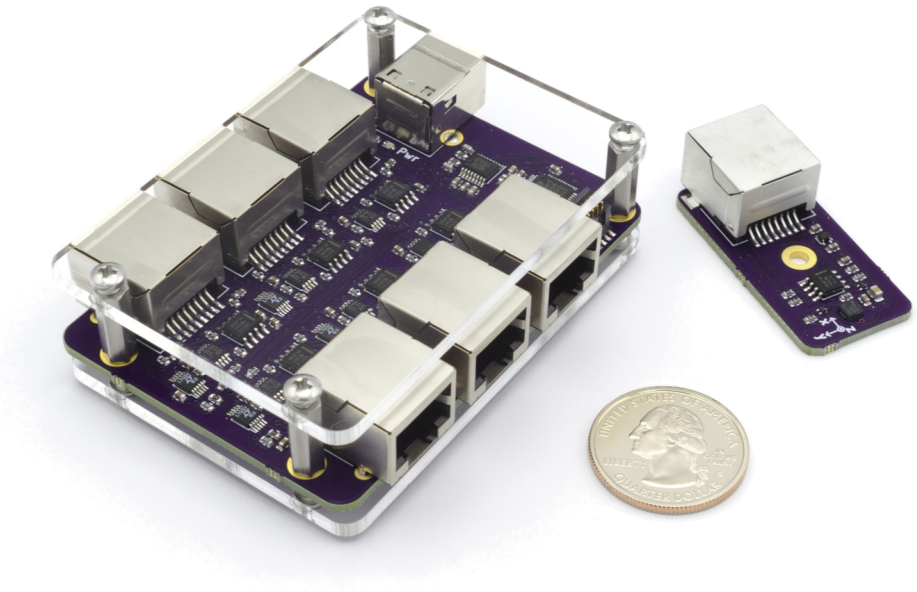

Figure 3.2: The accelerometer system readout and sensor hardware are shown, with a U.S. quarter for scale. The sensors are connected to the readout module using Category 5e cable and 8P8C modular connectors.

for longer cable runs. Thus, the cabling was switched to shielded Category 5e cables with 8P8C connectors, a cable type that is most commonly used for Ethernet.

As mentioned previously, each node on an I^2C bus has an address. While different devices generally have different addresses, identical devices almost always have the same address or the choice between a small subset of addresses. The LIS3DSH accelerometer that was chosen only supports two addresses, something that is common among I^2C sensors; furthermore, the address is set by tying one of the sensor's pins either high or low, so setting two sensor modules to have different addresses would require assembling them differently or using some sort of jumper. To avoid this issue, LTC4316 I^2C address translators were used, which allow I^2C addresses to be transparently altered in real time.[6] These address translators are placed between the main I^2C bus and each bus buffer such

[6]LTC4316 datasheet: `https://www.analog.com/media/en/technical-documentation/data-sheets/4316fa.pdf`

that when identical sensors, with identical addresses, are connected to each cable, they can be accessed using different addresses.

The final consideration that needs to be taken into account in choosing how many sensors to allow is their data rate. A combination of the sensor data rate and the I^2C bus bandwidth limits the maximum number of sensors. In the case of the telescope, frequencies up to 100 Hz are of the most interest. Taking the Nyquist frequency into account, this requires a minimum sampling rate of 200 Hz. As the two closest sampling rates supported by the LIS3DSH accelerometer are 100 Hz and 400 Hz, 400 Hz was chosen. Sampling three 16-bit numbers, one for each of the three accelerometer axes, at 400 Hz results in a raw data rate of 19.2 kbit s^{-1}. For the LIS3DSH, each data read requires an additional three bytes of overhead for addressing and register selection, bringing the data rate up to 28.8 kbit s^{-1}. As using I^2C's 100 kbit s^{-1} *standard mode* is preferable due to its looser rise time requirements, this limits us to three accelerometers per bus. This results in a bus bandwidth utilization of 86.4%, not counting pauses between transmissions. In practice, fully saturating the bus does not tend to be a good idea, and the pauses are significant enough that the bus speed was raised to 130 kbit s^{-1} to comfortably support three accelerometers.

The remaining portion of the vibration measurement system is the computer interface and readout. For this, an FTDI FT230X USB-to-serial converter chip was used in conjunction with an STMicroelectronics STM32F030 32-bit ARM microcontroller. The microcontroller's two I^2C buses are used to read three accelerometers each, allowing the readout module to read six accelerometers at once. Four hundred times a second, the microcontroller reads all six sensors, computes a checksum of the data, packetizes it using *consistent overhead byte stuffing* (Cheshire and Baker 1999), and sends it to the host computer via the FT230X USB-to-serial interface. The firmware also includes support for hot swapping in different accelerometers, allows for remote reset of the readout system, and implements a watchdog timer, which resets the readout system if it hangs, to ensure reliability. Finally, a script

running on the host computer timestamps the data and logs it to disk.

Once design and assembly of the prototype accelerometer system were completed, it was deployed to the field for testing. This testing turned up issues with longer cable runs, which motivated the switch from telephone cable to shielded Ethernet cable and an increase in the signaling voltage from 3.3 V to 5 V.[7] Once the revised design was produced, it was deployed to replace the initial prototype. A dozen accelerometer modules were affixed to various locations on the telescope using double-sided adhesive tape and were connected to a pair of readout modules using pre-terminated shielded Category 5e cables. The locations where accelerometer modules were installed are shown in Figure 3.3. With the accelerometer modules installed, data were collected both during normal observations and during targeted tests.

3.1.2 Analysis

Initial ideas to reduce vibrations were implemented and then evaluated using data from the accelerometer system and by looking to see if the issues that prompted the improvement attempts, focal-plane heating and VPM shutoffs, improved. Stiffing portions of the elevation axis and the support plate underneath the boresight-axis bearing by welding additional bracing to the telescope mount had little effect on either the vibrations or the symptoms, and bracing the bottom of the column that attached the azimuth axis to its encoder reduced the vibrations at the accelerometer on the column but did not improve the symptoms. Re-tuning the control loop for the azimuth servo motors led to a significant reduction in vibrations throughout the telescope, but it did not eliminate all the symptoms. Additional analysis of the accelerometer data and focal-plane temperature data determined that the focal plane heating correlated with the azimuth scanning direction, and the source of the vibrations was localized to the vicinity of the secondary azimuth servo motor. At

[7] This involved adding a voltage regulator to each sensor module, since the accelerometer requires 3.3 V.

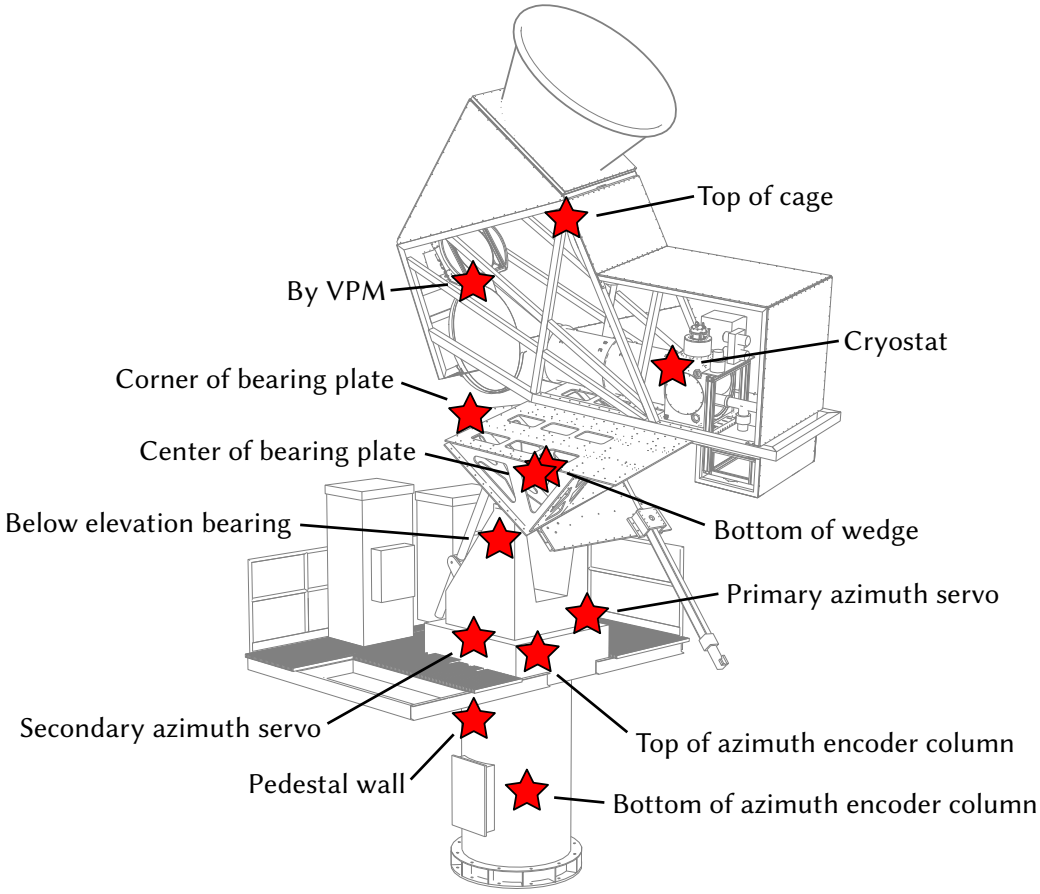

Figure 3.3: The locations where accelerometers modules were installed on the telescope are shown. The modules were attached using double-sided adhesive tape, and two readout modules were used to support the dozen sensors.

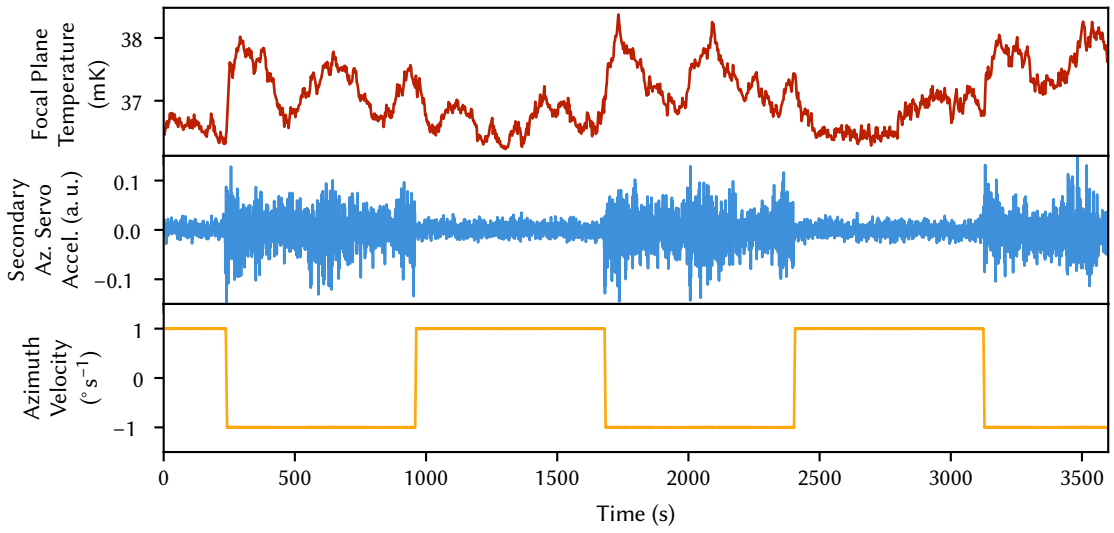

Figure 3.4: The correlation between the secondary-azimuth-servo vibration and the focal-plane temperature for a one-hour period is shown. The top panel shows the focal-plane temperature, the middle panel shows the vibrations as measured by the accelerometer installed on the secondary azimuth servo, and the bottom panel shows the azimuth velocity. Note the clear correlation between negative azimuth velocity and increased vibrations and the less obvious—but still noticeable—correlation between the increased vibrations and an increase in focal-plane temperature.

one point, the accelerometer module attached to this motor came loose in a way that amplified the vibration signal to the extent that it was clearly visible in the accelerometer time streams; this is shown in Figure 3.4. Once the accelerometer module was reattached properly, additional data were taken. By placing the accelerometer time stream data in azimuth-angle and scan-direction bins and taking Fourier transforms of the binned data, it was determined that the vibrations were dependent on azimuth angle in addition to azimuth scan direction.

The three components that could have contributed to the vibrations seen at the secondary azimuth servo motor were the motor itself, the gearbox that increased its torque and drove the primary azimuth-axis ring gear, and the shaft coupling that connected the motor to the gearbox. Since the coupling was the easiest and cheapest item to replace, replacing it was the first remedy attempted. The original coupling contained a torque-limiting slip

clutch, and this was replaced with a simpler flexible shaft coupling.[8] Additional data were collected while the telescope was scanning after the replacement, which showed that the fix was successful. Both the azimuth-angle and azimuth-scan-direction correlations in the accelerometer data were eliminated, as was the focal-plane-heating effect. Although this fix, as well as the previously discussed mitigations, significantly reduced the observed vibrations and reduced focal-plane heating, the VPM would still turn off at times, particularly when it was windy at the observing site. Additional hardware changes to improve the VPM performance are discussed in the next section.

The strongest vibration lines while the telescope is stationary are at 29.3 Hz and 57 Hz. Laboratory measurements using the accelerometer system confirmed that these vibrations originate with the pulse-tube cryocooler system, with the 57 Hz vibration line coming from the helium compressor and the 29.3 Hz vibration line coming from the cold-head-remote-motor-and-rotary-valve assembly. While scanning, there are additional vibration lines that result from interactions with the main azimuth-axis ring gear, which contains 270 teeth, or one every 4/3°; this produces vibrations at the scan frequency divided by the tooth spacing, e.g., at 0.75 Hz for a $1\,^{\circ}\,s^{-1}$ constant-velocity azimuth scan, and at higher multiples of this frequency. A combination of the vibration reductions enabled by the accelerometer system, the VPM control system improvement discussed in the next section, and the more symmetric mass distribution on the telescope resulting from a second receiver being installed on the mount eventually allowed azimuth scan speeds to be increased from $1\,^{\circ}\,s^{-1}$ to $2\,^{\circ}\,s^{-1}$, which potentially helps with recovery of large angular scales.

[8] There is additional over-torque protection in the form of limits on the servo motor drive current.

3.2 VPM control hardware design

While the original Q-band VPM worked, it did not do well with external disturbances, partly since it was operating above the resonance frequency of its spring system and partly because its controller did not have enough computing resources to run sophisticated control algorithms. To improve this, the flexures were stiffened to raise the resonance frequency above the 10 Hz operating frequency for subsequent VPMs, and new control electronics were designed, the latter of which is described here. Commercially-available motion controllers are normally used to control systems that are much easier to predict than the coupled-spring system used in the VPM and thus do not have enough computing power to run the comparatively complex algorithms needed to operate the VPM with disturbance rejection. Thus, custom control hardware was needed.

In addition to providing copious amounts of processing power, the control hardware needs to be able to read the VPM's four quadrature encoders, control its two voice-coil actuators, and provide an Ethernet interface for control input and data output. Thus, the electronics can be divided into two main components, a logic board and a voice-coil-driver board. An initial prototype for the logic board was built around a BeagleBone Black[9] single-board computer, which contains a Texas Instruments AM3358 processor and 512 MB of DDR3 RAM. The BeagleBone Black was chosen over alternatives such as the Raspberry Pi[10] due to the AM3358 processor's extensive suite of hardware peripheral interfaces, including three dedicated quadrature encoder inputs, and its two programmable real-time unit (PRU) coprocessors, one of which is used to implement a fourth quadrature encoder input. Other peripherals used include Ethernet, I^2C, SPI, UART, and GPIO pins. After the initial prototype, the logic board was transitioned to the Octavo Systems OSD3358 system-in-package, which integrates the critical components of the BeagleBone Black into

[9] `https://beagleboard.org/black`
[10] `https://www.raspberrypi.org/`

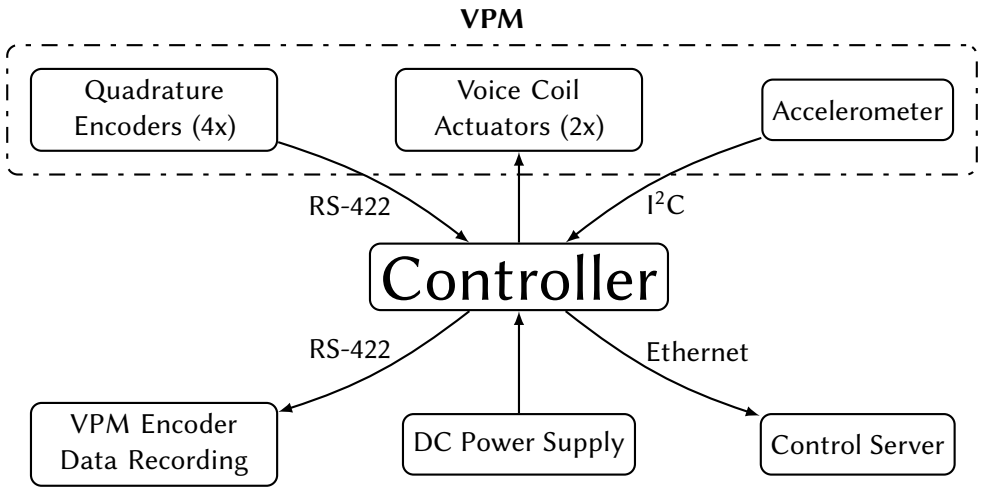

Figure 3.5: An overview of the VPM control system hardware requirements is shown. The controller needs to be able to drive the voice-coil actuators, read the position encoders, check for vibrations, enable synchronous recording of the encoder positions, and communicate with a control server.

a single package. Switching from the BeagleBone Black to completely-custom hardware was necessary because the BeagleBone Black does not expose all of the AM3358 processor's pins needed for the quadrature encoder inputs, while the OSD3358 does. Furthermore, using the OSD3358 system-in-package instead of copying the equivalent discrete layout from the BeagleBone Black greatly reduces the design complexity. In addition to the OSD3358, the logic board contains four pairs of four-channel RS-422 differential receivers and drivers, one pair for each of the four quadrature encoders; an interface to the voice coil driver board; JTAG and UART (via USB) debug interfaces; a microSD card slot for firmware; an I^2C bus (with bus extender) for an external accelerometer;[11] and an Ethernet PHY, which is Precision Time Protocol capable (for potential future use). The logic board uses a four-layer PCB.

The voice-coil-driver board contains two H-bridge drivers, one for each coil. Each driver consists of a Texas Instruments DRV8701P gate driver,[12] four Infineon BSC010N04LSI

[11]This interface is identical to that used for CLASS's vibration measurement system.
[12]Datasheet: `https://www.ti.com/lit/ds/symlink/drv8701.pdf`

MOSFETs,[13] and an Infineon TLI4970-D025T4 Hall-effect current sensor.[14] Each coil driver is designed for roughly 400 W of output power, limited by the amount of bulk capacitance, although the VPM draws considerably less power than this. The driver board also contains a 5 V step-down voltage regulator for powering the logic board and encoders and includes full reverse-polarity protection for its input power connection. Due to the large current requirements, the coil driver PCB uses 70 µm thick copper layers,[15] which are twice the standard thickness.

The logic board and voice-coil-driver board are connected using a flexible flat cable (FFC) for control signals and terminal blocks and wire for power. The entire system is packed into a compact extruded-aluminum enclosure with the help of five additional PCBs, which route electrical signals to and from connectors and serve as faceplates and mechanical supports. The logic board is supported by these additional PCBs, while the driver board is screwed to the top of the enclosure. One faceplate contains a terminal block for the 12 V input power, an Ethernet port for networking, and four additional 8P8C modular connectors for routing encoder signals from the RS-422 drivers to the mount control computer via shielded Category 5e cable; the terminal block is connected to the driver board via wire, while the other connectors are soldered directly to the logic board. An additional PCB connects to the other end of the encoder output cables and contains termination resistors for the RS-422 signals and a 78 pin D-subminiature connector for attaching to the appropriate input card on the mount control computer. The opposite faceplate contains four 15-pin D-subminiature connectors, one for each encoder; two terminal blocks, one for each voice coil actuator; and an 8P8C modular connector for attaching an external accelerometer over I^2C; the two terminal blocks are connected to the driver board via wires, while the other signals are connected to the logic board via PCBs.

[13] Datasheet: `https://www.infineon.com/dgdl/?fileId=db3a3043353fdc16013552d84dc147dc`
[14] Datasheet: `https://www.infineon.com/dgdl/?fileId=5546d4625607bd1301562bdf09d8339f`
[15] More commonly expressed in archaic units as "2 oz copper."

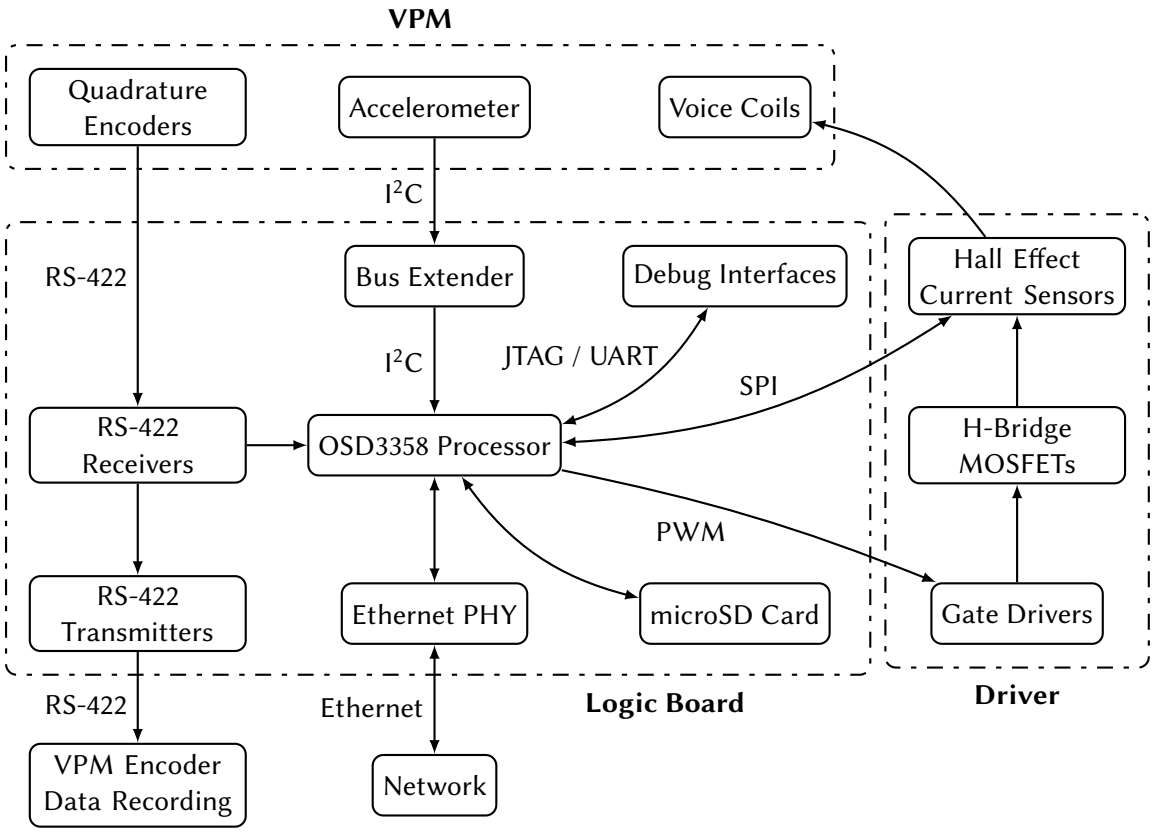

Figure 3.6: An overview of the VPM control hardware design is shown. It is based around two major PCBs, one for the control logic and one for the voice-coil drivers. Both of these subsystems interface with the VPM itself.

All of the components in the system are rated for industrial-temperature-range operation, rated to work between at least −40 °C and 85 °C.

On the firmware side of the design, a version of the Das U-Boot boot loader is stored on a microSD card and is loaded by the processor on boot. U-Boot is commonly used for embedded systems and has an abundance of options for loading the main device firmware. In this case, it is configured to load a firmware image off of the telescope site's central server via a Network File System (NFS) mount. Loading the main firmware over the network instead of using local storage allows the firmware to be quickly changed during development and to be much more easily upgraded in the field. The device's firmware is built using Texas Instrument's TI-RTOS real-time operating system and written in

Figure 3.7: The assembled VPM control electronics are shown. Besides the logic board and the voice-coil-driver board, additional PCBs are used to hold connectors, route electrical signals, provide mechanical support, and serve as faceplates. This allows for the electronics to be assembled into a compact extruded-aluminum enclosure.

C. After initializing the hardware, the firmware listens over TCP for commands and sends telemetry via UDP. The main control loop is implemented using a timer-driven hardware interrupt, while the networking stack is managed using the RTOS's scheduler. As mentioned previously, one of the processor's two PRUs is used to implement the fourth quadrature encoder input; the runtime of this program, written in assembly, is completely deterministic, taking the exact same number of clock cycles to run each iteration.[16] The second PRU loads the encoder count into memory for the main processor to access, since memory access is not deterministic. Additionally, it monitors the driver board's error and overcurrent outputs and is used to implement limits on the encoders such that if either driver has an error or any encoder limit is exceeded, both voice coil outputs are immediately

[16]The PRUs use a pipeline-free architecture, so all instructions that do not require bus access are executed in exactly one clock cycle. The firmware is then built without bus access and structured such that each branch of conditionals has the same number of instructions by adding no-ops to the shorter branch to ensure each iteration takes the same number of clock cycles.

disabled and an error is thrown. Disabling both the mirror and reaction axes at the same time if one of the axes encounters an error is necessary to avoid potentially-damaging vibrations.

The new control electronics provide vastly more processing power than the previously used commercial motion controller, such that computational power is no longer a constraint. Since the firmware is also completely custom, a nicer interface can be used between the control hardware and the external telescope control software. Furthermore, the use of custom electronics allows for a more compact all-in-one control solution that is easier to install and interface with. Should the MCE sync box be eliminated, the Precision-Time-Protocol-capable hardware will allow the controller to directly record the VPM encoder positions, eliminating the need for external logging.

3.3 Diode thermometry readout electronics

The forward voltage drop of a silicon diode is temperature dependent. Thus, such diodes can be used as thermometers by biasing them with a constant current and measuring the forward voltage drop. Such diode-based sensors are used for most thermometry within the CLASS cryostats. These diodes are each biased with a 10 μA current supplied using an LM234 constant-current source[17] operating in a temperature-compensated configuration. The forward voltages are digitized using Linear Technology LTC2449 24-bit analog-to-digital converters (ADCs).[18] As each ADC has 16 input channels, two are used to support the 24 input channels of each of CLASS's diode readout boxes. Readout is facilitated by an STMicroelectronics STM32F030F4 microcontroller,[19] which reads data from the ADCs and sends it out over Ethernet to a readout computer via a WIZnet W5500 TCP/IP

[17] Datasheet: `https://www.st.com/resource/en/datasheet/lm234.pdf`

[18] Datasheet: `https://www.analog.com/media/en/technical-documentation/data-sheets/2444589fc.pdf`

[19] Datasheet: `https://www.st.com/resource/en/datasheet/stm32f030f4.pdf`

embedded Ethernet controller.[20] The data are sent via UDP packets; this allows for seamless resumption of data logging if either the diode readout box or the readout script running on the readout computer is restarted, since the readout box's microcontroller stores the destination for the UDP packets in non-volatile memory. A previous version of this instrument is described in detail in Gothe (2015). Unfortunately, the previous version injected noise into the detectors, likely due to its use of a USB interface. The revised version, shown assembled in Figure 3.8, kept the same measurement circuit but switched the microcontroller and communications interface. It also utilizes a completely new circuit board layout that eliminated the previous design's excessive wiring and switches from primarily using through-hole components to using surface-mount components for everything except for the connectors.

3.4 Optics cage warm thermometry and heaters

It can be useful to monitor and record the temperatures of the various telescope components located in the telescope's optics cage, such as the VPM, primary and secondary mirror, and the exterior of the pulse-tube cold head. To this end, a digital thermometry system was designed and installed. This system was designed around the Texas Instruments TMP107 digital temperature sensor,[21] which has a resolution of 0.016 °C and an accuracy of ±0.4 °C. It communicates using a bidirectional one-wire interface that allows for up to 32 sensors to be daisy-chained together and addressed based on their position in the chain. The daisy-chaining greatly simplifies installation over alternatives such as thermocouples, which require each sensor to be wired individually. To further simplify installation, the CLASS thermometry modules were designed to use pre-terminated telephone cable with 6P4C modular connectors. Each of CLASS's thermometry modules contains a temperature

[20] Datasheet: `https://wizwiki.net/wiki/doku.php?id=products:w5500:datasheet`
[21] Datasheet: `https://www.ti.com/lit/ds/symlink/tmp107.pdf`

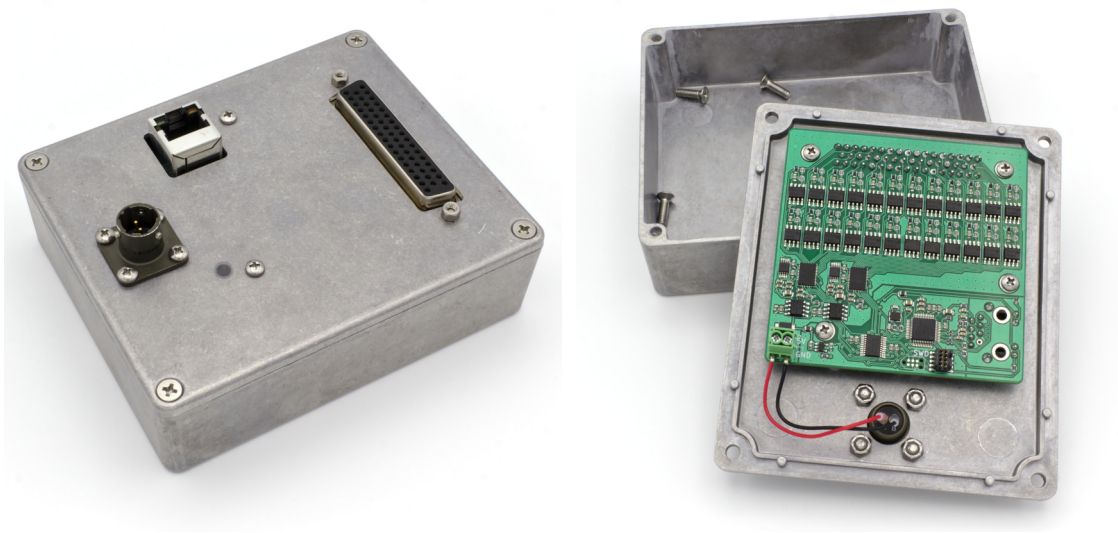

(a) Outside (b) Inside

Figure 3.8: The cryogenic diode thermometry readout electronics are shown. The left panel shows the outside of the readout box, which contains a power connector, a power status LED, an Ethernet connector, and a 50-pin D-subminiature connector, for connecting the diode sensors. The right panel shows the inside of the box, with the readout circuit board mounted to the inside of the box's lid and wires running from the power connector to a terminal block on the circuit board; the Ethernet connector, D-subminiature connector, and power status LED are all soldered directly to the circuit board.

sensor, a bypass capacitor, and two 6P6C modular connector jacks, which are all surface-mount components installed on the top side of the printed circuit board (PCB), leaving the bottom of the PCB flat. The bottom of the PCB contains a copper pour, which isn't covered with solder mask. This pour is connected to the top of the PCB using a large via, which also serves as a mounting hole. On the top of the PCB, this via is connected to the ground lead of the temperature sensor using a high-thermal-conductivity low-capacitance surface-mount capacitor, which thermally connects the copper pour on the bottom of the PCB to the temperature sensor but keeps it effectively isolated electrically. This connection—combined with the exposed copper on the flat bottom of the PCB—allows the temperature sensor to be thermally sunk to the surface it is connected to. These sensors have been installed in the optics cages primarily using thermally-conductive double-sided adhesive tape, although a mounting screw and thermal compound can also be used. Assembled sensor modules are shown in Figure 3.9. The sensors communicate via a simple serial protocol. Originally, a readout module with a USB-to-serial interface chip was used, but this has since been replaced with a readout module utilizing a microcontroller and Ethernet interface with UDP data transmission, similar to that used by the diode readout box described in the previous section. This readout module is powered via an 802.3af power-over-Ethernet (PoE) module.

With an accurate, reliable method of measuring the temperature inside the optics cage, it makes sense to heat the inside of the cage to regulate its temperature, since improved thermal stability could potentially help increase the stability of the CMB observations. However, the detectors are sensitive to radio-frequency interference (RFI), so a heater control system must be designed with this in mind. To this end, a heater control system was designed using variable zero-crossing solid-state relays to power 800 W positive temperature coefficient (PTC) fan heaters. A pair of DBK Cirrus 80 (FGC2001) heaters[22] were

[22]Datasheet: `https://datasheets.dbk-worldwide.com/EN/Datasheets/PTCFanHeaters/DBKITM_`

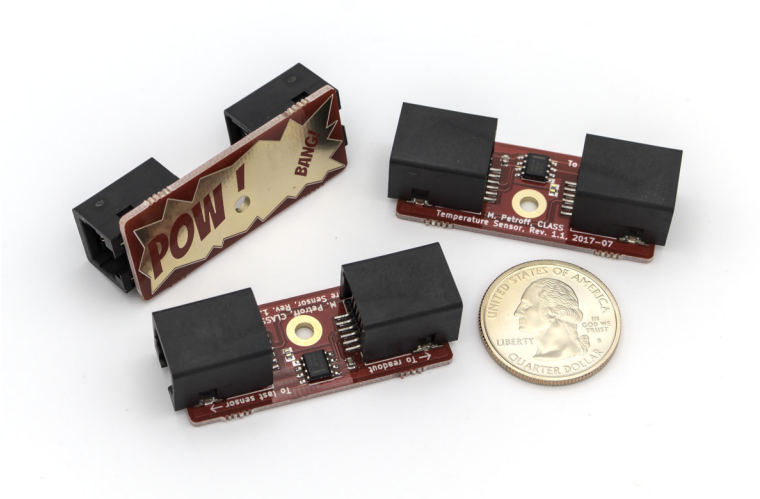

Figure 3.9: Assembled temperature sensor modules are shown with a U.S. quarter for scale. Note the exposed copper pour on the bottom of the sensor module, which allows the sensor to be thermally sunk to what it is attached to.

installed in each side of the optics cage on the first telescope mount, using a combination of machined aluminum parts and 3D-printed plastic parts, with each heater pair controlled by a Power-io HMA-3V50E variable relay[23] and with the heater fans controlled by a single Power-io IO-OAC-280 non-variable zero-crossing solid-state relay.[24] A third variable relay was included to allow for a forebaffle heater to be installed in the future. Both the heaters and the fans use 120 V single-phase AC power, and fuses are included for overcurrent protection. A Moxa ioLogik E1241 input/output module[25] is used to control the relays via an Ethernet interface. The assembled heater control system box is shown in Figure 3.10. As solid-state relays are used, there are no mechanical contacts, which could produce RFI when switched. The relays also only switch on or off when the AC power waveform is at zero potential, further reducing the risk of RFI. The relays and control electronics are also all installed in a grounded metal enclosure. Despite the RFI mitigations and despite

Cirrus%2080_eng.pdf

[23] Product page: https://www.power-io.com/products/hma.htm

[24] Product page: https://www.power-io.com/products/oac.htm

[25] Datasheet: https://www.moxa.com/getmedia/d8e9f25e-8831-4541-b6f4-092f4527a17a/moxa-iologik-e1200-series-datasheet-v2.2.pdf

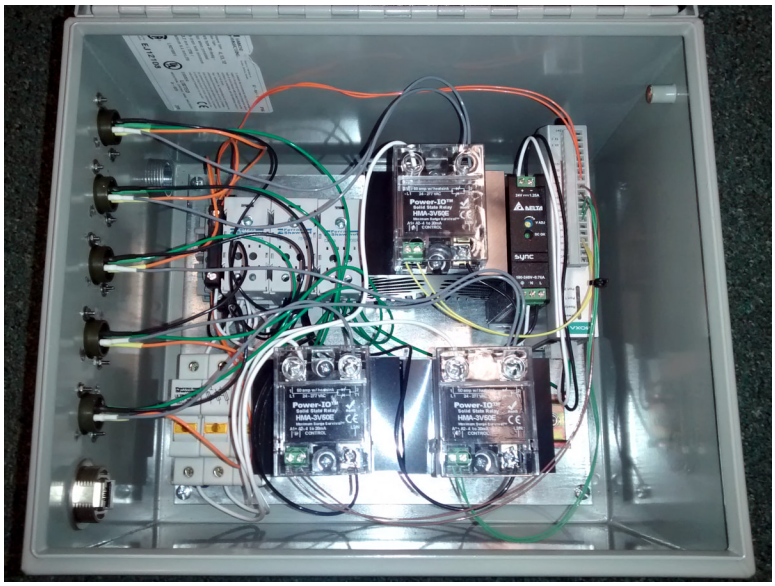

Figure 3.10: The assembled optics cage heater control box is shown. The components are mounted on a pair of DIN rails in a metal enclosure.

the lack of RFI detected, this heating system was never used due to RFI concerns. Thus, a heater control loop was not developed, and no heating system was installed on the second telescope mount once it was deployed.

Chapter 4

A 3D-printed Broadband Millimeter-wave Absorber

The ongoing development of increasingly sensitive cosmic microwave background (CMB) telescopes requires commensurate improvements in the control of systematic errors. One such source of error is from stray light, which needs to be controlled and terminated through the use of millimeter-wave absorbers, both under ambient conditions and within cryogenic receivers (Eimer et al. 2012; Harrington et al. 2016). Since the field is moving towards the use of multichroic detectors with wide frequency bands in a shared optical path (Abitbol et al. 2017), broadband absorbers are required. Per this application, the frequency range from 30 GHz to 230 GHz is of interest, since it covers the CMB emission peak as well as synchrotron and thermal dust foregrounds. The maximum allowable reflectivity is dependent on where in the instrument the absorbers are used and is often difficult to quantify precisely; in general, lower reflectivity is better, but this must be considered in the context of cost, volume, and thermal design constraints. Broadband millimeter-wave absorbers should also prove useful in W-band radar applications or, more generically, as

Except for some minor changes, this chapter was previous published as Petroff et al. (2019).

terminations or glint reduction media in an optical bench.

In recent years, additive manufacturing in the form of 3D printing has become increasingly common, in particular fused filament fabrication (FFF).[1] FFF-based printing works by extruding a plastic filament through a heated nozzle mounted on a CNC stage such that an object is built up layer-by-layer (Jones et al. 2011). This allows for rapid prototyping and allows for the manufacturing of easily-customized designs and one-off parts. In this work, this ease of customization is applied to the fabrication of millimeter-wave absorbers. 3D printing has been previously used for manufacturing free-space electromagnetic absorbers (Sanz-Izquierdo and Parker 2014; Kronberger and Soboll 2016; Ren and Yin 2018). To the best of our knowledge, however, only comparatively narrowband, resonator-based absorbers have been previously produced with such techniques.

Graded-index absorbers generally take the form of an array of pyramidal structures (Emerson 1973). In the limit where the wavelength is similar to or greater than the feature size, the pyramids form a smooth gradient in effective permittivity, greatly reducing reflections. In the limit where the wavelength is smaller than the feature size, the pyramids cause the incident light to reflect off the absorber structure multiple times, with some radiation absorbed at each interaction. Since these two limits are governed by the feature pitch and the wavelength, it can be helpful to use the parameterization p/λ_g, where p is the pitch and $\lambda_g = \lambda_0/\sqrt{\epsilon_r'}$ is the wavelength in the absorber material as a function of the free-space wavelength, λ_0, and the real component of the absorber material's dielectric function, ϵ_r'; this parameterization can be derived from the antenna theory analogue of a graded-index absorber (Chuss et al. 2017).

While a periodic pyramidal structure makes for an effective absorber, it is not well suited for FFF-based printing. When sliced into layers for FFF-based printing, each pyramid slice is disconnected from the others in a given layer. Thus, the filament extrusion process

[1] This is also referred to by the phrase "fused deposition modeling" and trademarked acronym FDM.

Table 4.1: Comparison between types of low-profile graded-index absorbers, specifically cast (or injection molded) pyramid arrays, foams with conductivity gradients, and the 3D-printed absorber presented in this work.

	Absorber type		
Property	Cast	Foam	3D-printed
Cryogenic compatibility	yes	no	yes
Easily customized	no	no	yes
Reflectivity	very low	low	low

must be stopped and the filament retracted for each and every pyramid slice; this procedure prevents the creation of sharp points, since small amounts of plastic can be drawn back into the nozzle, and possibly causes stringing—thin strands of plastic stretched over gaps—that must be manually removed after the print is finished. Furthermore, prints are generally weakest along their layer lines, making points liable to break off. To avoid these issues, a geometric approximation of a space-filling curve is used, which fills the plane with a continuous wedge.

In this chapter, we demonstrate that 3D printing allows for the rapid production of broadband millimeter-wave absorbers. These easily customizable absorbers can achieve adequate performance for stray light termination at low cost and as a thermal plastic—not a foam—they can be used cryogenically. The fabricated absorbers are not intended to serve as extremely low reflectivity calibration targets, such as pyramidal array (Bai et al. 2017) and cone-shaped (Houtz et al. 2017) targets that have been detailed in the literature. Furthermore, due to the intended use case, only a low power regime has been evaluated. A comparison between the fabricated absorbers and other common types of low-profile graded-index absorbers is presented in Table 4.1.

4.1 Space-filling curves

A space-filling curve is a curve that passes through every point of a two-dimensional region with positive Jordan measure, in this case area (Sagan 1994).[2] This type of mathematical monster was first described by Peano (1890). Shortly thereafter, Hilbert described another such curve but also described an iterative sequence of geometric approximations to his curve (Hilbert 1891). Importantly, this type of geometric approximation can be physically realized.

Space-filling curves are self-similar (Sagan 1994). Thus, absorbers created from different order geometric approximations of space-fillings curves will exhibit similar behavior. In this work, a Hilbert curve, which fills a square, is primarily used. However, other shapes are possible. For example, an absorber designed around a generalized Gosper curve (Gardner 1976; Fukuda et al. 2001), which roughly fills a regular hexagon, was also prototyped, since it allows one to more easily cover a circle with absorber tiles. For $n > 2$, all electromagnetic absorbers with n-fold rotational symmetry have no net polarization response (Mackay 1989). Although a Hilbert curve has only 1-fold symmetry—or 2-fold symmetry in the case of the loop variant, the Moore curve—geometric approximations of the Hilbert curve have asymptotically equal numbers of uniformly distributed horizontal and vertical line segments (Moon et al. 2001); thus, they should not show a net polarization response in reflectivity, unlike 2-fold symmetric absorbers created from sets of parallel wedges.

To create an absorber, a wedge was modeled such that the peak of the wedge follows the centerline of a geometric approximation of the Hilbert curve. The wedge then extends down such that the halfway points between segments of the curve form troughs. Furthermore, only a shell is printed, leaving the inside hollow, to reduce printing time and save on

[2] For such a curve, the Lebesgue covering (topological) dimension is equal to the Hausdorff (fractal) dimension. Thus, a space-filling curve is *not* a fractal in the purest sense, since a fractal requires the Hausdorff dimension to be strictly greater than the geometric dimension (Sagan 1994; Mandelbrot 1977).

material use. A rendering of a second-order Hilbert curve absorber model can be seen in Figure 4.1.

A similar procedure can also be followed to create absorbers from other space-filling curves with appropriate symmetry, although this still only allows for a few different shapes. To create custom-shaped absorbers, a solid geometry intersection operation can be performed between the existing, square Hilbert curve absorber model and an extrusion of the desired absorber shape, e.g., a circular absorber can be created by computing the intersection between a cylinder and a Hilbert curve absorber model. By scaling the model's dimensions, it is easy to adjust the size and pitch of the absorber's features.

4.2 Material selection

Two different material candidates were tested, a commercially available carbon-loaded conductive polylactic acid (PLA) filament[3] and a high impact polystyrene (HIPS) filament custom extruded[4] from a commercially available carbon-loaded conductive HIPS pellet.[5] The PLA has the advantage of being an off-the-shelf filament, and PLA is generally easier to print than HIPS. However, since the carbon-loaded HIPS comes in pellet form, it can easily be mixed with normal HIPS pellets during filament extrusion to control the conductivity and dielectric properties of the resulting filament. It may thus be desirable to find a source of carbon-loaded PLA pellets—or pulverized PLA to mix with a powdered additive—to maintain this advantage and combine it with the ease of use of PLA.

The complex relative dielectric functions of the carbon-loaded PLA filament, the carbon-loaded HIPS, plain PLA,[6] and plain HIPS[7] plastics were characterized using a pair of

[3]Proto-pasta Conductive PLA; ProtoPlant, Inc., Vancouver, WA; `https://www.proto-pasta.com/`
[4]Filabot EX2 filament extruder; Triex LLC, Barre, VT; `https://www.filabot.com/`
[5]PS-715; Modern Dispersions, Inc., Leominster, MA; `http://www.moderndispersions.com/`
[6]Ingeo Biopolymer 4043D; NatureWorks LLC, Minnetonka, MN; `https://www.natureworksllc.com/`
[7]Avantra 8130; PMC Polymer Products, Inc., Stockertown, PA; `http://pmcpolymerproducts.com/`

filled rectangular waveguide sections[8] for each material, which each form a Fabry-Pérot resonator; the waveguide section thicknesses were nominally chosen such that the thinner and thicker sections would constrain the real, ϵ_r', and imaginary, ϵ_r'', components of the dielectric function, respectively (Wollack et al. 2008). To fill the waveguide sections, the plastic samples were heated until soft and then compressed, eliminating air bubbles; excess material was then removed, and the sections were lapped flat. Using a modeling approach described in existing literature (Wollack et al. 2008), the dielectric functions were extracted from scattering parameters measured from 22 GHz to 40 GHz using a vector network analyzer (VNA)[9] calibrated with Thru-Reflect-Line (TRL) standards, sampled at 801 points. The results of these measurements are summarized in Table 4.2. These measurements target the lower range of the absorber's response, which, in viewing the absorber as a lossy adiabatic structure, has the highest sensitivity to the magnitude of the dielectric function and to the geometry. In practice, the dielectric function for a dilute non-magnetic conductively-loaded thermal plastic is anticipated to be a weak function of frequency (Sihvola 1999) as has been observed in dielectric mixture systems (Wollack et al. 2008) analogous to those employed here.

Of the materials tested, the carbon-loaded HIPS is the best candidate for an absorber, since it is an absorptive dielectric that is not too reflective. As ϵ_r'' increases, increasing the loss, ϵ_r' also generally increases, increasing the reflectance; thus, a compromise must be made between these two parameters. The carbon-loaded PLA is much more conductive and behaves like a poor metal, making it too reflective for this application. Without carbon loading, the HIPS is a low-loss dielectric, which makes it potentially useful in millimeter-wave optics applications, while the PLA is a relatively lossy dielectric, with properties similar to that of nylon.

[8] WR28; broad-wall 0.280", guide height 0.140"

[9] Agilent PNA-X; Keysight Technologies, Inc., Santa Rosa, CA; `https://www.keysight.com/`

Table 4.2: Dielectric function measurement results for bulk plastic samples in waveguide section fixtures are presented. The observed relative statistical error of the dielectric function for the average of each set of samples is < 0.01 for the real component and < 0.12 for the imaginary component. The systematic error due to calibration errors and sample preparation variation is estimated to be $< 2\,\%$ for the real component and $< 10\,\%$ for the imaginary component.

Material	ϵ_r'	ϵ_r''
Avantra 8130 HIPS	2.49	0.003
Ingeo 4043D PLA	2.80	0.03
PS-715 Conductive HIPS	7.7	2.1
Proto-pasta Conductive PLA	15.	15.

4.3 Electromagnetic modeling

A small unit cell of the absorber was modeled in the ANSYS HFSS software package,[10] which performs an electromagnetic finite element analysis (FEA). Periodic boundary conditions were implemented using a pair of perfect electric conductor (PEC) planes and a pair of perfect magnetic conductor (PMC) planes around the unit cell. A second-order Hilbert curve was used as a unit cell for the absorber, since it has greater symmetry than a first-order curve—and thus should have lower residual polarization response—while having reduced simulation complexity when compared to higher-order curves. The simulation setup is shown in Figure 4.1. Reflectivity at normal incidence was simulated.

Different wall thicknesses, including a solid cross section, were simulated for the conductive HIPS, as were orthogonal polarizations. Additionally, the conductive HIPS was simulated with varying degrees of conductivity. Possible wall thicknesses are quantized by the diameter of the 3D printer's extrusion nozzle, 0.5 mm in this case. Unless otherwise noted, a double nozzle diameter wall thickness, with a thickness of 1 mm and manufactured with two parallel passes of the printer's extrusion nozzle, and a 3:1 wedge height to wedge width aspect ratio was used (before truncation due to nozzle diameter). A 2:1 aspect ratio was also tried; since this affects p/λ_g, it should result in a somewhat different frequency

[10] ANSYS HFSS 19.1; ANSYS, Inc., Canonsburg, PA; `https://www.ansys.com/`

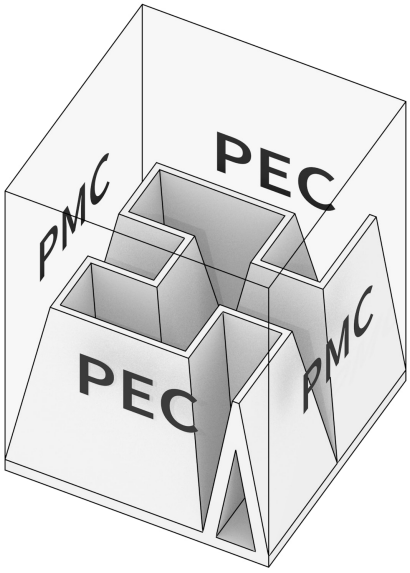

Figure 4.1: The unit cell used in simulations, a second-order Hilbert curve, is shown. Periodic boundary conditions were established by using perfect electric conductors (PECs) for the front and back faces of the bounding prism and perfect magnetic conductors (PMCs) for the left and right faces of the bounding prism. The top face of the prism was used as a wave port, and the bottom face was set as a PEC.

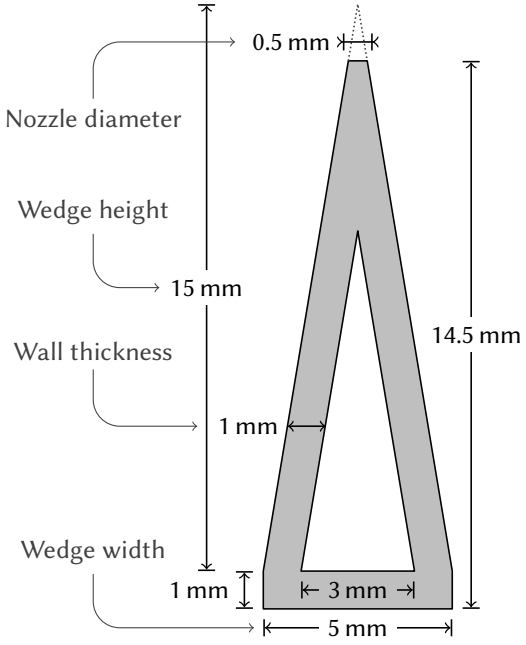

Figure 4.2: A cross section of the absorber wedge is shown, labeled with the dimensions used for the measured prototype, which utilizes a double wall printed with a 0.5 mm nozzle. Due to the nozzle diameter, the wedge tip is truncated.

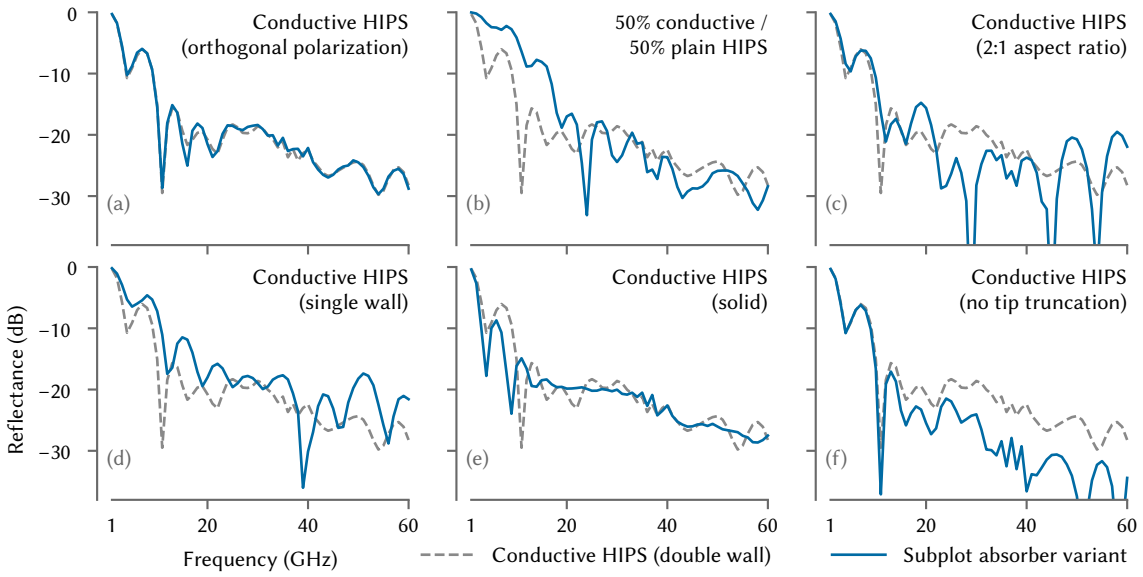

Figure 4.3: The results of FEA reflection simulations of absorber variants (solid) are shown against the baseline simulation of the geometry and material used for the fabricated prototype (dashed). Each simulation was performed from 1 GHz to 60 GHz at an interval of 1 GHz. The reflectance is shown as a function of frequency for six absorber variants.

response. Finally, a simulation was performed without wedge tip truncation to evaluate how the geometry performs when it is not limited by manufacturing constraints. To constrain the design space, a fixed height for the absorber was used.

A cross section of the simulated model, including dimension labels, is shown in Figure 4.2; the aspect ratio refers to this wedge cross section, not the unit cell shown in Figure 4.1. The single wall width simulation used a 0.5 mm thick wall, while the triple wall width simulation used a 1.5 mm thick wall. To simulate orthogonal polarization, the PEC and PMC boundaries were switched to the opposite walls from what is shown in Figure 4.1. Simulations were performed from 1 GHz to 60 GHz at an interval of 1 GHz. Higher frequencies were omitted from simulations to reduce the computation time required; since the reflectivity levels off for higher frequencies, the frequency range simulated is enough to evaluate the relative performance of the simulated absorber variants. The results of these simulations are shown in Figure 4.3.

The simulation data show good polarization symmetry, which supports the choice of unit cell. The equal mix of conductive HIPS and plain HIPS shows slightly better performance than the conductive HIPS alone, but it was decided that this marginal improvement was not worth the extra effort in filament manufacturing, so a mix was not used for prototyping. The solid-cross-section simulation results were quite similar to the double-wall simulation results, so the double wall was used, as this reduces printing time and material usage without negatively affecting the reflectivity. The single wall and 2:1 aspect ratio simulations show better performance at some frequencies but are worse at others and are more inconsistent across a range of frequencies, so these design variations were not used for the prototype absorber. The similarity of the reflectivities of the geometry and material variants is expected, since adiabatic graded-index absorbers are relatively insensitive to their exact geometry and dielectric function (Janaswamy 1992; Holloway et al. 1997).

When the wedge tip was not truncated, reflectivity was significantly decreased, showing that the flatness of the wedge tip that results from manufacturing limitations is a factor that limits performance. If the physical dimensions of the absorber are scaled up, the relative size of the tip truncation will decrease, since the diameter of the nozzle tip—and thus the absolute size of the tip truncation—remains constant. Due to this decrease in the relative size of the tip truncation, scaling up the absorber size should improve the reflectivity beyond what would be expected from the shift due to the change in p/λ_g. However, this potential increase in performance must be weighed against the increase in absorber thickness.

4.4 Fabrication

To prepare a solid model for an absorber, a script was written to procedurally generate the model to meet specific design parameters using a solid modeling scripting library.[11] With the solid model in hand, G-code instructions for the 3D printer were generated using slicer software.[12] The solid model generation code, the resulting solid models, and the final G-code have been made available (Petroff 2018).

This G-code instruction file was then used to print the absorber prototype on a LulzBot TAZ 6 FFF 3D printer[13] with a 0.5 mm diameter extrusion nozzle. The tested absorber prototype was manufactured using carbon-loaded HIPS and is a fifth-order Hilbert curve with a square footprint of 160 mm by 160 mm and a total height of 14.5 mm; a double wall and 3:1 aspect ratio were used. A cross section of the tested absorber is shown in Figure 4.2. The prototype absorber is shown in Figure 4.4, and a detailed view of the prototype is shown in Figure 4.5. An alternative realization of the absorber geometry is shown in Figure 4.6.

4.5 Measurement

The fabricated absorber prototype was measured using a VNA coupled to a free-space quasi-optical setup. This setup is described in detail in previous literature (Chuss et al. 2017). Measurements were taken in two waveguide bands, from 63 GHz to 115 GHz and from 140 GHz to 215 GHz; in terms of p/λ_g, this is 2.9 to 5.3 and 6.5 to 10.0. The results are shown in Figure 4.7. While the simulations were performed with a PEC backing behind the absorber, no such conductive backing was used when performing the measurements; however, the power transmitted through the absorber was less than −50 dB, so the effect

[11] CadQuery v1.2.0, `https://github.com/dcowden/cadquery`
[12] Cura v2.6.1, `https://github.com/Ultimaker/Cura`
[13] Aleph Objects, Inc., Loveland, CO; `https://www.lulzbot.com/`

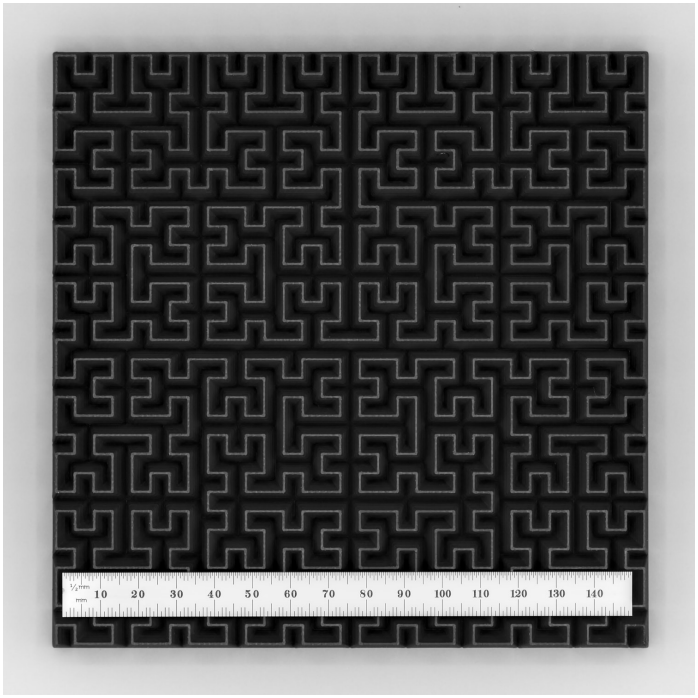

Figure 4.4: The prototype carbon-loaded HIPS Hilbert curve absorber is shown. The footprint of the absorber is a 160 mm square.

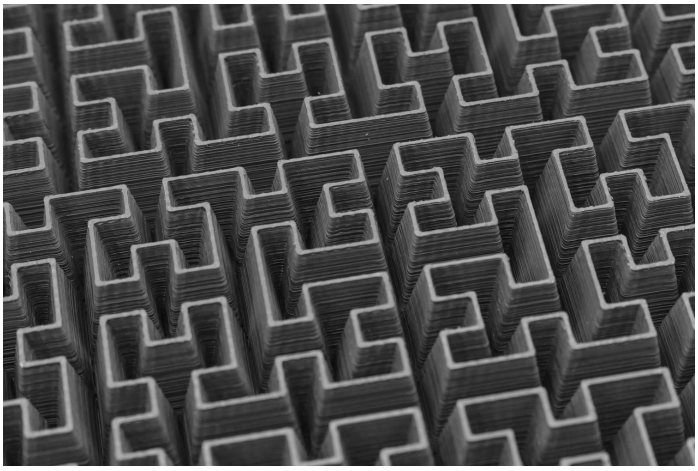

Figure 4.5: A detailed view of the prototype absorber is shown. Layer lines from the 3D printing process are visible.

Figure 4.6: An alternative realization of the absorber geometry is shown. Due to the geometry's size, it will likely not perform well at millimeter wavelengths. While the material's dielectric properties have not been measured, the imaginary component is likely too low for good electromagnetic performance; it is edible, however.

of the backing or lack thereof is insignificant. There is good agreement between the measurements and electromagnetic simulations, providing additional validation of the choice of unit cell used for the simulations. Furthermore, prototypes have been successfully used cryogenically at 60 K for stray light absorption, surviving dozens of thermal cycles, and have been successfully cooled using liquid nitrogen for use as cold loads.

4.6 Conclusions

A broadband graded-index absorber was designed around a geometric approximation of the space-filling Hilbert curve and was 3D-printed using FFF. Its reflectance was then measured in the frequency range 63 GHz to 215 GHz and found to be better than −20 dB at normal incidence, which is suitable for the stray light absorption application being considered; evaluation at oblique incidence is left for future work. The use of a space-filling curve overcomes the limitations imposed by FFF while also providing additional mechanical robustness when compared to a traditional tiled pyramid design. With a single 3D printer,

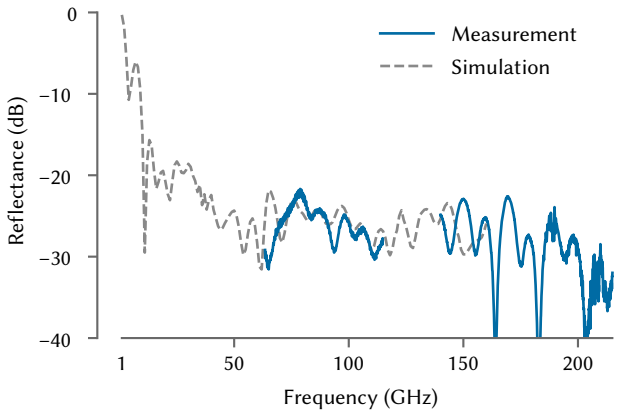

Figure 4.7: Reflectance measurement results (solid) are shown along with reflectance values predicted by electromagnetic simulation (dashed).

one absorber can be fabricated per day, but this production rate can be easily scaled up by parallelizing fabrication across multiple 3D printers. To extend the concept from stray light absorbers to cryogenic calibration targets, the reflectivity needs to be further reduced, and the thermal gradient caused by plastic's poor thermal conductivity needs to be addressed, such as through the use of a thermalizing core, which could be made by adding a metal insert into holes left in the bottom of the 3D print or by filling the currently hollow interior with an epoxy designed to have higher thermal conductivity.

The absorber could be further refined by using a smaller extrusion nozzle on the 3D printer, allowing for a sharper peak to the wedge, which should improve performance. Since additive manufacturing allows for rapid prototyping, multiple wedges profiles could be easily tried and their performances compared. The use of a multi-material FFF printer could allow for different amounts of carbon loading to be used for the outer and inner walls of the absorber to decrease reflection or to potentially create a gradient in the carbon loading, something that would be difficult to accomplish with more traditional manufacturing techniques. Moving away from FFF, a selective laser sintering (SLS) [stereolithography (SLA)] 3D printer could be used with carbon powder or stainless steel powder mixed in with the nylon powder [resin] to further extend the ease of customization of the absorbers,

including for covering curved surfaces, by utilizing a more traditional pyramidal absorber structure. The fine resolution of SLA 3D printers is also well suited for the creation of extremely low reflectivity calibration targets.

Chapter 5

Design of Assorted Hardware

5.1 XPS foam infrared absorbing filters

Polystyrene is almost completely transparent at millimeter wavelengths while being almost completely opaque in the infrared regime, making it an ideal candidate for an absorptive infrared filter; the main issue with the material on its own is reflection. Extruded polystyrene (XPS) foam[1] is a very-low-density form of polystyrene; due to its low density, its refractive index is very close to one, resulting in almost no reflection. In vacuum, each filter absorbs almost all of the incoming infrared power, heats up, and then reradiates the power, half on one side of the filter and half on the other. Thus, stacking a set of these filters creates a radio-transparent multilayer insulation (RT-MLI) (Choi et al. 2013), which allows the incoming infrared power to be greatly attenuated with very little in-band loss.

Unfortunately, much of the commercially-available XPS foam contains additives that are emissive in the millimeter. Three different brands of XPS foam insulation were tested,[2] but all were found to be emissive. Fortunately, a source of non-emissive XPS foam, intended

[1] This is normally sold at home improvement stores as insulation.

[2] Owens Corning FOAMULAR 150, Lowe's / Kingspan GreenGuard XPS, DiversiFoam CertiFoam 15

for use with the construction of model airplanes, was located.[3] As all three emissive samples were different colors and since the foam insulation sourced in Japan by Choi et al. (2013) was not emissive, a fire-retardant additive is a likely source of the emissivity. The non-emissive model airplane foam was found to burn more easily than the emissive foam insulations, giving further credence to this theory.

The model airplane foam is delivered as an approximately 6 mm thick sheet, with a skin layer on each side, and has a density of approximately 36 $\mathrm{kg\,m^{-3}}$ when the skin layers are included; its cell size was measured to be 200 µm to 300 µm. Using an end mill designed for cutting foam,[4] the skin layer and approximately 2 mm of foam were removed from each side of the foam sheet by milling in a spiral pattern using a CNC milling machine.[5] The milling machine was also used to cut the filters into a circle of the necessary diameter. As polystyrene is a rigid plastic, XPS foam machines more easily than foams made from more malleable plastics such as polyethylene. A sheet of XPS foam insulation was used as a machining bed, which was milled flat prior to beginning the filter machining process; filters were affixed to the bed using double-sided adhesive tape. The milling machine was used since it was available, having been previously used to machine plastic lenses and absorptive filters of the Q-band and W-band receivers. However, a CNC router and a vacuum-table bed would have made the machining process easier.[6] Each machined filter is 2 mm thick. With the skin layers removed, the foam has a density of approximately 31 $\mathrm{kg\,m^{-3}}$. These filters are then stacked on nylon threaded rods, with laser-cut acetal plastic spacer rings between filters, for use in the CLASS cryostats. The machining process and a completed filter stack are shown in Figure 5.1.

[3] Model Plane Foam, https://www.modelplanefoam.com/; manufactured by R. L. Adams Plastics, Inc., https://www.goadams.com/

[4] Amana Tool HSS1614, https://www.amanatool.com/hss1614-high-speed-steel-hss-foam-cutting-straight-v-flute-1-2-dia-x-1-1-4-x-1-2-shank-router-bit.html

[5] Tormach PCNC 1100 Personal CNC Mill, https://tormach.com/

[6] The rotary stage used on the milling machine prevented the use of a vacuum table.

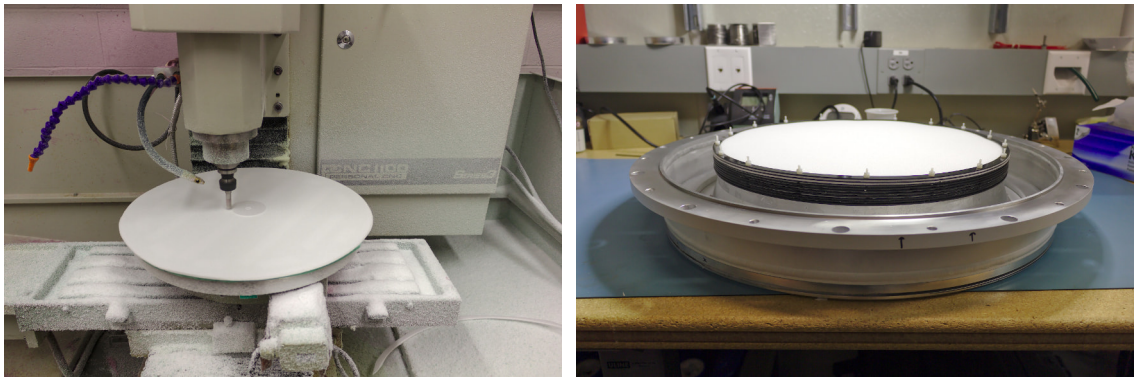

(a) Machining filter (b) Assembled filter stack

Figure 5.1: XPS foam filters. On the left, the machining process of the XPS foam filters is shown. An assembled filter stack is shown on the right.

5.2 Photogrammetry

In order to align the CLASS warm optics to the cold optics within the cryostat, the positions of the VPM and mirrors with regard to the cryostat position need to be accurately measured. For this, CLASS uses photogrammetry. Coded optical alignment targets are placed throughout the optics cage, and precision uncoded optical targets are installed in known positions on the cryostat and mirrors. A scale bar is also installed in the optics cage. Using a calibrated camera, a set of several hundred photos is taken, and the coded and uncoded targets are identified using computer vision algorithms. The identification number of a coded target can be directly determined from a single image, which helps with the next step. A bundle adjustment is then run, which estimates camera lens parameters, the position of the camera for all of the photos, and the positions of all of the optical targets using trigonometry. Finally, the arbitrary scale of the photogrammetry reconstruction is set to a physical scale via measurements of targets on the scale bar.

For final measurements in the field, a Geodetic Systems V-STARS N photogrammetry system is used.[7] However, the optics also needed to be aligned prior to shipping the

[7] Geodetic Systems, Inc., Melbourne, FL; `https://www.geodetic.com/`

optics cage, to ensure that the optics could be adjusted into the proper locations. As the V-STARS system belongs to a collaborator in Chile, it could not be used for this purpose, and acquiring a second copy of the system was prohibitively expensive. Thus, a new system was needed, so an AICON Scanreference system,[8] which is an order of magnitude less expensive than the V-STARS system, was purchased. Unfortunately, the camera used with this system has a fairly narrow field of view, necessitating a larger number of optical targets and a larger number of photos to complete the reconstruction. Furthermore, the software provided with the system is slow, and the reconstruction will completely fail if it encounters any sort of error.

Although the system required a large number of coded targets to complete a reconstruction of the optics cage, it only included 149 magnetic coded targets and digital files from which an additional 192 coded targets could be printed. Unfortunately, the magnetic coded targets were too small to work reliably at the distances required to photograph the optics cage, leaving only the digital files, which could be used to print appropriately-sized targets. As the 192 printable targets were not enough, the manufacturer was asked to provide files for the remaining targets. Instead of providing the files, they attempted to sell a multi-thousand-dollar software package for creating them. Fortunately, some research turned up an expired patent that described the ring coding system used for the optical targets (Thielbeer 1997). The patent contains all of the details needed to generate the ring codes for the coded targets, except for the exact parameters and numbering scheme used. This missing information was fairly straightforward to figure out—the targets are 14-bit with no restrictions on the number of transitions from black segments to white segments and are ordered by increasing binary value. Scripts to generate the ring codes and to generate a set of printable targets were then written using this information. With the output of these scripts, appropriately-sized stickers with all 516 possible coded targets

[8]Hexagon DEU01 GmbH, Braunschweig, Germany; `https://www.scanreference.com/`

were produced. This enabled the pre-shipment optics alignment to be completed.

5.3 3D-printed anti-reflective layer

When light encounters a boundary between materials with differing refractive indices, a fraction of the light is reflected, as dictated by the Fresnel equations. When transitioning between vacuum or air, with a refractive index of ~1 and a material with a higher index, such as the ~1.58 of high-impact polystyrene (HIPS) plastic in millimeter wavelengths, a layer with an index equal to the square root of the index of the higher-index material can be used to reduce reflections. When the thickness of the lower-index layer is equal to a quarter wavelength of the incident light, reflections are eliminated due to destructive interference. As the thickness and index diverge from a quarter wavelength of the incident light's wavelength and the square root of the index of the substrate, respectively, the reflection increases. If portions of the substrate material are removed from a layer with feature sizes that are sub-wavelength, a metamaterial layer is created with an effective refractive index that is between that of the substrate material and air, which is sometimes referred to as a simulated dielectric. This allows for the creation of anti-reflective layers without the need for a separate material, and such layers are used for the anti-reflective layers on the CLASS lenses, created using subtractive machining.

While such metamaterial anti-reflective layers are commonly constructed using subtractive machining, it should also be possible to create such layers using additive manufacturing. To evaluate the feasibility of such a technique, a prototype was developed to be manufactured using extrusion-based fused filament fabrication 3D printing for use at Q band. This prototype was modeled via a transfer matrix method as a solid HIPS-plastic layer of a given thickness sandwiched between two layers of lower index of a given thickness, with the entire stack surrounded by vacuum. An optimization to minimize reflection was

then run over the frequency range of the CLASS Q-band receiver, with the thickness of the solid plastic core and the thicknesses and indices of the metamaterial layers as free parameters. This optimization determined the optimal thicknesses of the metamaterial layers and solid layer to be 1.7 mm and 2.5 mm, respectively. The target index for the metamaterial layers was 1.17, which is less than the square root of the index of the solid layer due to interactions between the three layers.

To produce the prototype, a flat, solid substrate was first printed in the shape of a square out of HIPS plastic. On top of this, a metamaterial anti-reflective layer was printed using a 0.25 mm diameter extrusion nozzle; this layer consisted of a single-nozzle-width meander in the form of a geometric approximation of a Hilbert curve. Once an initial prototype was printed, it was imaged with a flatbed scanner with a black material placed behind the white, semi-translucent prototype. This allowed the fill fraction of the metamaterial layer to be estimated using an image lightness thresholding algorithm, which was necessary to account for manufacturing tolerances. This fill fraction measurement was used to iteratively adjust the printed pattern in order to reach the desired fill fraction. The thicknesses of the samples were also verified using calipers and adjusted as needed. Once the final manufacturing parameters were determined, two larger samples were produced. These were then solvent-welded together using methyl ethyl ketone in order to create a sample with anti-reflective layers on each side, with care being taken to mask the metamaterial layers to avoid damaging them with the solvent. Originally, attempts were made to directly print a prototype with anti-reflective layers on each side, but these attempts were unsuccessful since the material layer would not stick well enough to the 3D printer's print bed, causing the print to fail. The completed prototype is shown in Figure 5.2.

The performance of the prototype was tested using the fielded CLASS Q-band receiver during the period of time when the optics cage was removed to facilitate upgrades to the

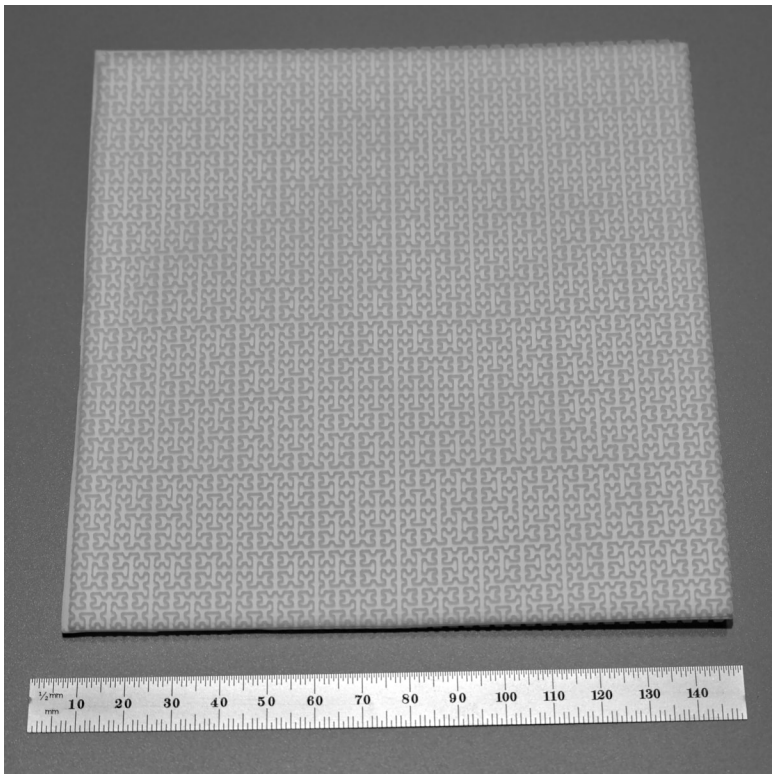

Figure 5.2: A completed prototype with a 3D-printed anti-reflective layer is shown. Note that the basic unit cell of the Hilbert curve meander is smaller than the ∼7 mm wavelength of the high-frequency band edge of the CLASS Q-band receiver.

receiver and to facilitate deployment of the first W-band receiver on the same mount. The lack of the optics cage allowed a flat mirror to be installed in front of the receiver window to direct the beam toward zenith, as shown in Figure 5.3. The receiver window was necked down using an aluminum plate with a 5 cm diameter circular aperture in the center, to allow for measurement of the small prototype.

Data collection started with observations toward zenith, and an Eccosorb HR-25 absorber sheet[9] backed by aluminum-foil-covered cardboard was then installed in front of the window aperture. The 3D-printed filter was then chopped twice between the window aperture and the Eccosorb sheet, and the same was done for a 1.02 mm thick sheet of plain HIPS, as a control. The Eccosorb sheet was then removed, and its aluminum-foil-covered back was chopped twice in front of the window, with it pushed flush against the aperture plate each time, to ensure it was parallel to the plate. By analyzing the time streams recorded during this process for detectors located near the center of the focal plane, the reflectivity of the filter can be determined. The Eccosorb sheet approximates zero reflectance, while the aluminum sheet approximates 100% reflection back into the cryostat, so the reflectance of the filter can be calculated as

$$R_{\text{filter}} = \frac{M_{\text{filter}} - M_{\text{aluminum}}}{M_{\text{Eccosorb}} - M_{\text{aluminum}}}, \tag{5.1}$$

where the M values are the detector data values in arbitrary units corresponding to the subscript labels. This measurement showed the filter reflectance to be ~0.5%, which is considerably larger than the predicted ~0.001%, although this can easily be attributed to manufacturing uncertainties and experimental error. The reflectance of the non-anti-reflective-coated reference sheet was measured to be 17.5%, which is close to the 17.4% expected reflectance. A formal error analysis was not performed.

[9]Datasheet: `https://www.laird.com/sites/default/files/2021-01/RFP-DS-HR%20092815.pdf`

To measure the emissivity, a different procedure was used. Data collection again started with observations toward zenith, and the filter was chopped twice in front of the window aperture. The Eccosorb sheet was then also chopped twice in front of the window aperture. The Eccosorb sheet approximates a blackbody with 100% emissivity and temperature equal to the ambient temperature, T_{ambient}, approximately 270 K, while the sky temperature, T_{sky}, is approximately 12 K in band. Assuming the filter is also at ambient temperature, the filter's emissivity can be calculated as

$$\varepsilon_{\text{filter}} = \left(\frac{M_{\text{filter}} - M_{\text{sky}}}{M_{\text{Eccosorb}} - M_{\text{sky}}} \right) \left(\frac{T_{\text{ambient}} - T_{\text{sky}}}{T_{\text{ambient}}} \right). \tag{5.2}$$

The emissivity was measured to be ∼13%. This is a bit higher than would be expected given the bulk properties of HIPS, suggesting that the additional emissivity might be due to the fact that the basic unit cell of the Hilbert curve pattern is only a factor of a few smaller than the wavelength. As the fabrication method did not allow for a smaller pattern, this suggests that the developed technique may be more useful for longer wavelengths, such as at K band.

Additional measurements of the prototype were made using a VNA coupled to a free-space quasi-optical setup, similar to that described in Chuss et al. (2017). This measurement placed the average in-band reflectance for the CLASS Q-band telescope at ∼0.05%, showing that the anti-reflective layers performed well. The measurement also allowed for simulation parameters to be fit to the data, resulting in a dielectric function of 1.25+0.04i and a thickness of 1.71 mm for the metamaterial layers and a dielectric function of 2.49+0.06i and a thickness of 2.49 mm for the solid layer. The thicknesses matched the designed values, but the refractive index of the metamaterial layers, 1.12, was lower than expected. Additionally, the loss was significantly higher than expected, suggesting that the HIPS used to produce this prototype had higher loss than the HIPS sample measured in Table 4.2,

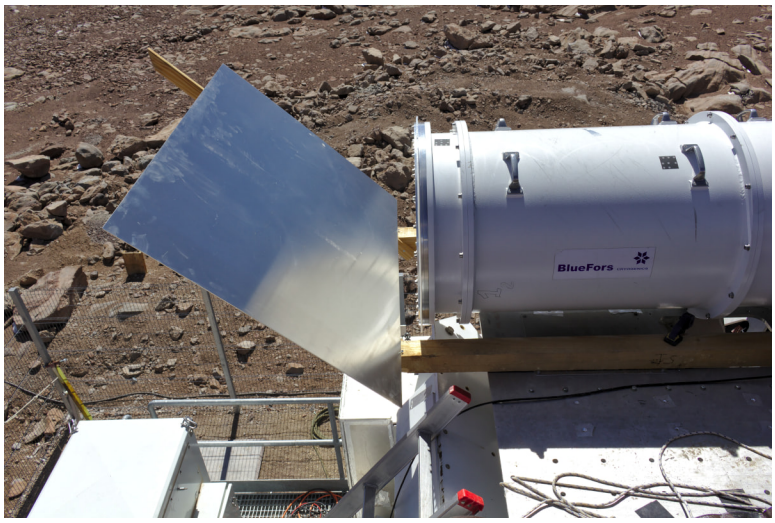

Figure 5.3: The CLASS Q-band receiver is shown installed on the telescope mount in the Atacama desert, without the optics cage installed. A flat mirror has been installed in front of the receiver window to direct the beam toward zenith.

which was from a different material supplier, or that the methyl ethyl ketone solvent used to assemble the prototype had a detrimental effect.

5.4 Silicone-based broadband absorber

Polyurethane-foam absorbers, such as Eccosorb HR, quickly degrade when installed at the CLASS observing site due to ultraviolet (UV) exposure and must additionally be protected from the occasional snowstorm, motivating the development of an alternative broadband millimeter-wave absorber. Such an alternative should ideally be flexible, for ease of installation on the curved forebaffles of the CLASS telescopes, and needs to be able to be produced more quickly than the 3D-printed absorber described in Chapter 4, in order to produce enough to cover the forebaffles in a reasonable amount of time. To this end, development of a carbon-loaded cast silicone absorber was explored. Silicone—particularly with carbon loading—is UV resistant. Additionally, it is flexible, and there is no concern that the carbon loading will wash out of the material.

To produce such an absorber, a pyramidal design with a 5 mm base length and a 10 mm height was chosen. A casting mold was then manufactured using a Formlabs Form 2 stereolithography (SLA) 3D printer out of Formlabs *standard clear* photopolymer resin.[10] Smooth-On Mold Star 15 silicone[11] was mixed with approximately 5% carbon black by volume, poured in the mold, vacuum degassed, allowed to cure, and then removed from the mold. This initial prototype was measured at Q band using a VNA coupled to a free-space quasi-optical setup, similar to that described in Chuss et al. (2017). The reflectance started at −6 dB at 34 GHz and decreased to −14 dB by 50 GHz. As the reflectance of the initial prototype was a bit high, a second prototype was produced with 10% carbon black by volume. The reflectance for the second prototype started at −8 dB at 34 GHz and decreased to −17 dB by 50 GHz. Based on these measurements, the absorber is expected to perform well at W band but would likely need a larger pyramid size for better performance at Q band. To test the durability of the absorber, a prototype was secured to the roof of a shipping container at the CLASS observing site, in direct sun. After 29 months, it showed no signs of degradation. The mold used to produce the absorbers and a prototype absorber are shown in Figure 5.4. Despite quicker production times than the 3D-printed absorber design and reasonable performance, development of this absorber was not pursued further, since producing enough of the absorber to cover the entirety of the CLASS forebaffles would have still taken an unreasonable amount of time and effort.

5.5 3D-printed dog-leg absorber geometry

Based on a concept presented in Wollack et al. (2014), an absorber geometry with a dog-leg feature was developed for use with SLA-based 3D printing. This geometry is based on a pyramidal geometry, but the region along the base of the pyramids is replaced with a dog-

[10] Formlabs, Inc., Somerville, MA; `https://formlabs.com/`

[11] Datasheet: `https://www.smooth-on.com/tb/files/MOLD_STAR_15_16_30_TB.pdf`

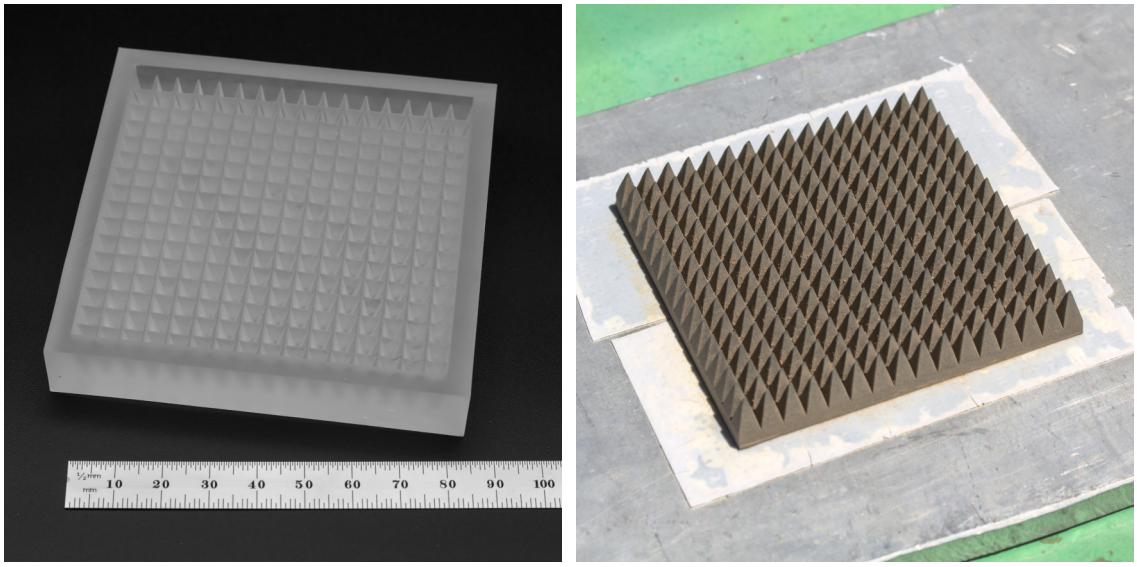

(a) Mold (b) Absorber

Figure 5.4: On the left, the mold used to cast the silicone absorber prototypes is shown. On the right, a prototype absorber is shown after 29 months installed in direct sun at the CLASS observing site; although it is a bit dusty, it shows no signs of degradation.

leg feature, which hides the base of the slot from view. When the absorber is operated in the geometric limit, this feature should reduce the absorber's reflectance, since light reaching the base of the pyramids must bounce and be absorbed multiple times. A prototype was produced using a Formlabs Form 2 SLA 3D printer out of Formlabs *durable* photopolymer resin, proving the feasibility of the concept; this prototype is shown in Figure 5.5. The prototype survived testing with liquid nitrogen, suggesting that it can handle cryogenic use. For such a geometry to be used as a broadband absorber, an SLA-compatible photopolymer resin that is also a lossy dielectric needs to be developed, but the development of such a resin has been left for future work.

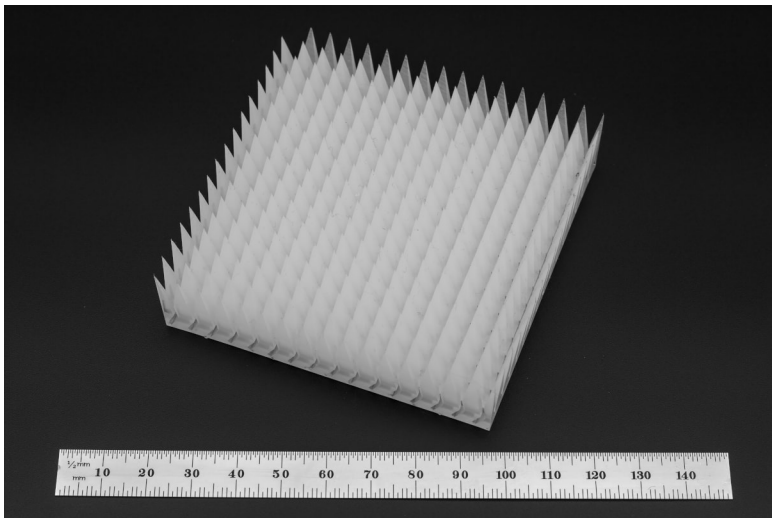

Figure 5.5: A 3D-printed prototype of a dog-leg geometry for a broadband absorber is shown. The dog-leg feature prevents single-bounce reflection off of the region at the base of the pyramidal structures.

5.6 Smooth-walled rectangular-to-circular waveguide transition

Abruptly transitioning from one waveguide cross-section to another, such as from a rectangular waveguide to a circular waveguide, produces reflections. Thus, a transition geometry is needed to minimize such reflections. The simplest transition is to linearly interpolate from one cross-section to another, but this is not optimal, as it leads to either unnecessarily-high reflectance for a fixed transition length or an unnecessarily-long transition for a fixed maximum reflectance. One common approach is to use a stepped Chebyshev transformer, which uses multiple transition steps, each of which is a quarter wavelength long, to minimize the magnitude and ripple of the frequency-dependent reflectance in the frequency band of interest (Wollack 1995). While this works well for transitions manufactured using subtractive machining of aluminum or copper, a smooth, non-stepped transition is preferable for use with SLA-based 3D printing as it avoids abrupt overhangs. As photopolymer resins are poor conductors, the waveguide parts are copper plated after printing.

In order to produce a smooth transition, the generalized-superellipse method of Cogollos et al. (2018) was extended to work with rectangular-to-circular waveguide transitions. A generalized superellipse is defined parametrically using the equation

$$\left(\frac{x}{k_x}\right)^m + \left(\frac{y}{k_y}\right)^n = 1, \tag{5.3}$$

where k_x and k_y are the semi-radii of the superellipse, $|x| \leq k_x$ and $|y| \leq k_y$ are dimensions along the 2D plane, and $m, n \in \mathbb{R}^+$ control the curvature. A planar transition between two microstrips of differing widths is constructed by taking one quadrant of the superellipse, where k_x is equal to the length of the transition and k_y is equal to the difference between the widths of the two microstrips, and forming a taper profile where the taper edge follows the superellipse. This can be extended to transition from one rectangular waveguide cross-section to another by using two different superellipses, one for each waveguide dimension, as was done in Cogollos et al. (2018). It can be further extended to a rectangular-to-circular waveguide transition by first also parameterizing the waveguide cross-section as a generalized superellipse, with $m = n = \infty$ and $m = n = 2$ for the rectangular and circular ends, respectively. To transition from a circular waveguide to a rectangular waveguide, the two superellipses for the rectangular-to-rectangular waveguide transition are supplemented with a third superellipse to set the transition profile of the squareness parameter, with $k_x = k_y = 1$. This gives final parametric equations for the transition surface of

$$x(u,v) = \left[\frac{w}{2} + \left(\frac{d-w}{2}\right)\left(1-u^{m_x}\right)^{1/n_x}\right] |\cos v|^{\alpha} \frac{\cos v}{|\cos v| + \epsilon} \tag{5.4}$$

$$y(u,v) = \left[\frac{h}{2} + \left(\frac{d-h}{2}\right)\left(1-u^{m_y}\right)^{1/n_y}\right] |\sin v|^{\alpha} \frac{\sin v}{|\sin v| + \epsilon} \tag{5.5}$$

$$z(u,v) = l \cdot u, \tag{5.6}$$

where

$$\alpha = \left(1 - u^{m_s}\right)^{1/n_s} + \epsilon, \tag{5.7}$$

$u \in [0, 1]$, $v \in [-\pi, \pi)$, w is the width of the rectangular waveguide, h is the height of the rectangular waveguide, d is the diameter of the circular waveguide, l is the length of the transition, and $\epsilon = 0.001$ is a small number to avoid numerical issues; n_x, m_x, n_y, m_y, n_s, and m_s are parameters to optimize over.

A transition with length $1.234\lambda_c$ from a rectangular waveguide with dimensions $0.392\lambda_c$ and $0.784\lambda_c$, e.g., a WR10 waveguide, to a circular waveguide with diameter $0.919\lambda_c$ was designed with 40% fractional bandwidth; λ_c denotes the center wavelength. The diameter of the circular waveguide was chosen to have an identical cut-off frequency of the TE_{10} mode as the rectangular waveguide. The transition was parametrically modeled in the ANSYS HFSS software package,[12] which performs an electromagnetic finite element analysis (FEA), and a genetic-algorithm optimization was performed to find parameters that minimized the maximum in-band reflectance. This resulted in generalized-superellipse curvature parameters of $m_x = 1.378$ and $n_x = 1.033$ for the broad-wall transition, $m_y = 1.393$ and $n_y = 0.5659$ for the narrow-wall transition, and $m_s = 1.171$ and $n_s = 0.6149$ for the squareness-parameter transition. The in-band reflectance was below −35 dB across the entire band; this compares favorably to a linear transition of the same length, which only has in-band reflectance below −22 dB across the entire band. A wire-frame rendering of the optimized transition is shown in Figure 5.6. This design has yet to be fabricated, although smooth-walled copper-plated SLA-printed waveguide components should work at frequencies up through at least W band.

[12] ANSYS HFSS 2019 R1; ANSYS, Inc., Canonsburg, PA; `https://www.ansys.com/`

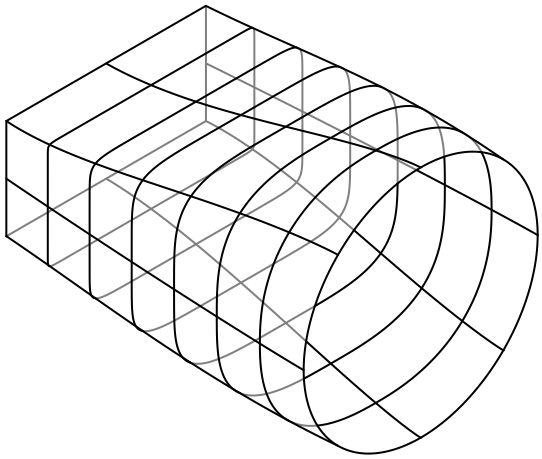

Figure 5.6: A wire-frame rendering of the optimized rectangular-to-circular waveguide transition is shown. Note the smooth walls and the non-linear nature of the transition.

Chapter 6

Single-track Absolute Rotary Color Optical Encoder

Measuring absolute angular position is traditionally done using either a Gray code or a quadrature encoder with an additional index track, both of which are commonly read using either optical or magnetic sensors. Gray codes allow for instantaneous measurement but require a large number of tracks for high resolution, e.g., twelve tracks to achieve 0.1° resolution for a rotary encoder. Since the value of only one track changes at a time with a Gray code, Gray codes are resistant to sensor read errors; read errors caused by track segment transitions only result in a misread of the adjacent position instead of a position elsewhere on the encoder track (Gray 1947). Indexed quadrature encoders require only two tracks but need to first be rotated to find an index mark before outputting an absolute angular position.

However, particularly-space-constrained applications only have room for a single track. Single-track Gray codes are a particular class of Gray codes that allow multiple sensors to be placed at different positions on a single track instead of using multiple tracks with one sensor per track, as is done with a traditional Gray code, but single-track Gray codes still

require at least as many sensors as a traditional Gray code (Hiltgen et al. 1996). Naively using a colored track allows this binary coding scheme to be extended to base-n, where n is the number of colors used, thereby reducing the number of read heads required. However, in order to maintain Gray code's resistance to position errors resulting from track segment transitions, only a subset of color transitions can be used, necessitating the use of a color graph. While rare, color has been previously used for optical encoders (Iafolla et al. 2020), but the representation of color as a graph is a novel concept.

6.1 Color graph

When a color sensor crosses the boundary between two adjacent track segments that are different colors, the sensor measures a mixture of the two colors. For example, if one segment is blue and the other is yellow, the sensor will measure green during the transition; if green is also a color used on the encoder track, this would introduce a read error. Thus, only transitions between blue and green and then between green and yellow can be allowed if Gray code's resistance to read errors resulting from transitions is to be maintained.

These color transitions can be quantified by using a Euclidean chromaticity plane, e.g., the CIE 1931 xy plane. At its 1931 meeting, the Colorimetry Committee of the Commission Internationale de l'Éclairage (CIE) established a set of three reference primary colors, X, Y, and Z, which form a linear color space that establishes both chromaticity and luminance. This volume is then projected via

$$x = \frac{X}{X + Y + Z} \tag{6.1}$$

$$y = \frac{Y}{X + Y + Z} \tag{6.2}$$

to form the xy chromaticity plane (Fairchild 2013). When two isoluminant colors are mixed on this plane, the resulting color is on the line that connects the two source colors.

On this plane, the entire color gamut visible to a typical human observer can be plotted, as is shown in Figure 6.1, with the boundaries formed by monochromatic light and the line of purples. Thus, a three-color method such as the red, green, and blue (RGB) primaries, used for additive color applications such as emissive digital displays, or the cyan, magenta, yellow, and black (CMYK) primaries,[1] used for standard process color printing, cannot be used to represent the entire visible gamut. RGB additive primaries form a triangle in the gamut with one primary forming each corner, while CMYK subtractive color, of particular interest since a color encoder track is printed via this method, forms a hexagon with three corners formed by the primaries and the other three corners formed by equal mixtures of two different primaries. This printable gamut is plotted in Figure 6.1.

Since the result of mixing two colors is on the line between the two colors, encoder track segment transitions fall on this line. In order to prevent read errors from being caused by transitions, the midpoint of said line needs to be closer to its two endpoints than any other color used on the encoder track. Thus, an equilateral triangular grid is overlayed on the printable gamut, and adjacent nodes of the grid are used to select colors for the encoder track. To maximize robustness, the grid's offset and orientation and the node selection should be done such that the grid's scale can be maximized, thus maximizing the distance between adjacent nodes.

A graph consisting of the corners and center of a regular hexagon with adjacent nodes connected was chosen for creating a 0.1° resolution encoder. This seven-color graph is plotted on the xy plane in Figure 6.1, and the values of the seven colors are shown in Table 6.1.

[1] Black ink is a special case.

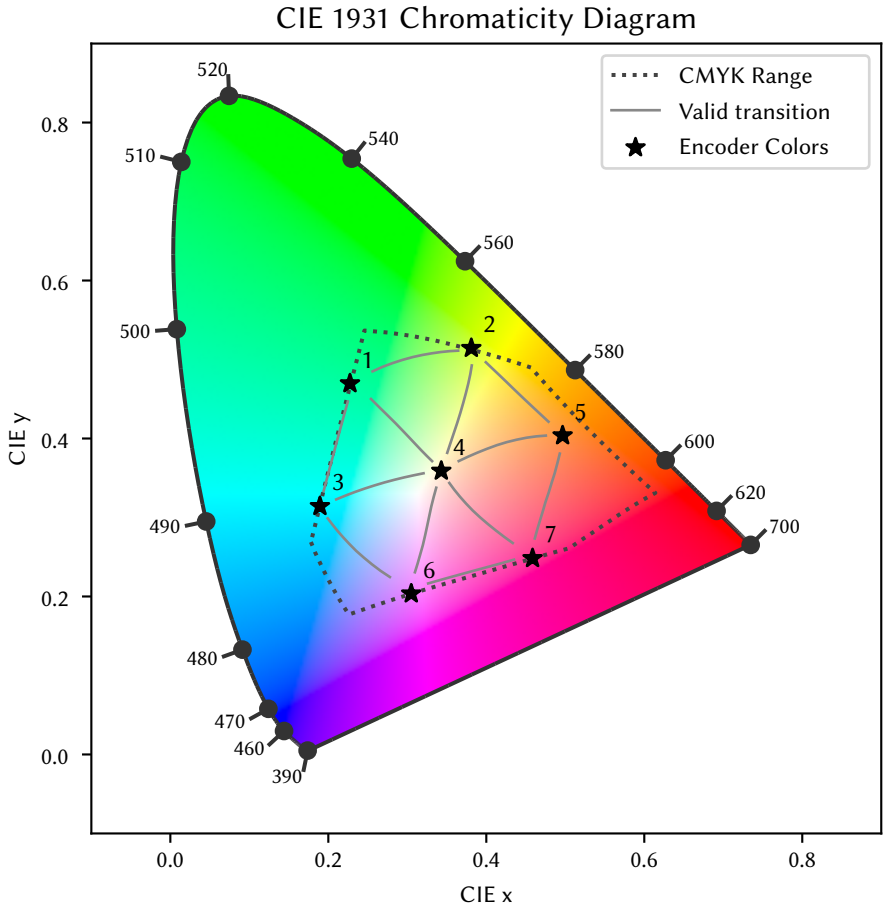

Figure 6.1: A chromaticity diagram showing the visible gamut, the CMYK printable gamut, the encoder colors graph, and the valid color transitions is presented. The Specification for Web Offset Publications (SWOP) color profile is used for the CMYK range boundary as it describes a fairly generic CMYK process. The numbers along the edge of the horseshoe shape refer to the wavelength of monochromatic light in nanometers. As the full visible gamut is much larger than what can be printed on paper or displayed on a computer monitor, the colors used here are for visualization purposes only.

Table 6.1: The colors that comprise the color graph are presented. The number refers to the encoder color ID, x and y are CIE 1931 chromaticity coordinates, and the CMYK color values are expressed as percentages, derived using $Y = 38$ (CIE 1931 Y, not CMYK Y).

No.	x	y	C	M	Y	K
1	0.227	0.470	99	0	76	0
2	0.381	0.514	49	0	100	0
3	0.189	0.314	100	0	21	0
4	0.343	0.359	31	22	24	0
5	0.497	0.404	12	54	83	0
6	0.305	0.204	76	100	0	30
7	0.459	0.248	22	100	4	20

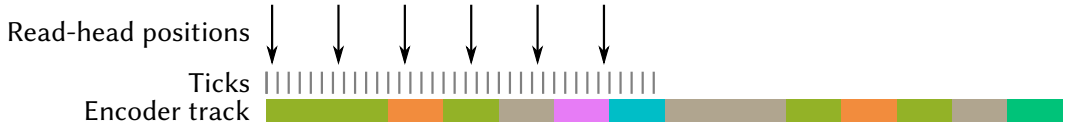

Figure 6.2: A visualization of the encoder track is shown. The arrows denote the relative positions of the six read heads, the ticks below them denote the track divisions that correspond to the encoder resolution, and the colored boxes at the bottom show the encoder track. Note that the first colored box is 11 ticks wide, and the next five colored boxes are each 5 ticks wide. This segment pattern then repeats but with different colors.

6.2 Generating an encoder track

In order to minimize the encoder's size, the read heads are placed as close together as allowed by the restriction that only one sensor changes at a time; this results in sensors spaced by the number of track divisions that is equal to the number of sensors. For a 0.1° resolution encoder, six read heads were used with the aforementioned color graph. With this information, a repeating track segment pattern needs to first be constructed; the number of divisions in this pattern must be equal to the number of read heads times the read head spacing, in this case 36 divisions. Numerous such patterns can be trivially constructed; the pattern used was 11–5–5–5–5–5, since this maximizes the width of the smallest segment, which in turn minimizes the chance of a read head measuring an adjacent segment. A visualization of this track pattern is shown in Figure 6.2.

With a segment pattern and color graph in hand, the encoder pattern can then be

Pattern: [11, 5, 5, 5, 5, 5]
Color sequence: [2, 5, 2, 4, 6, 3, 4, 2, 5, 2, 4, 1, 4, 7, 4, 1, 2, 5, 4, 1, 4, 3, 4, 6, 3, 1, 4, 5, 7, 6, 7, 4, 1, 2, 4, 7, 6, 3, 1, 2, 1, 4, 3, 1, 3, 1, 2, 1, 2, 4, 2, 5, 2, 5, 2, 1, 2, 5, 7, 5, 4, 6, 3, 6, 3, 4, 3, 6, 7, 5, 7, 4, 5, 4, 5, 2, 5, 2, 1, 3, 4, 7, 4, 2, 1, 3, 4, 6, 3, 6, 4, 7, 5, 4, 3, 1, 4, 2, 1, 3, 6, 3, 1, 2, 4, 1, 3, 4, 5, 4, 3, 4, 6, 4, 6, 4, 2, 1, 4, 7, 5, 4, 7, 4, 3, 4, 3, 6, 3, 6, 7, 4, 5, 7, 6, 7, 5, 7, 6, 7, 4, 6, 4, 6, 4, 1, 2, 4, 3, 1, 4, 3, 6, 4, 1, 3, 4, 2, 4, 2, 1, 2, 1, 2, 4, 1, 2, 5, 7, 6, 7, 5, 4, 7, 4, 7, 5, 7, 6, 4, 1, 4, 6, 3, 4, 5, 7, 5, 2, 1, 2, 1, 4, 3, 6, 7, 6, 3, 1, 2, 5, 2, 1, 4, 3, 6, 4, 6, 3, 6, 7, 6, 4, 2, 1, 2, 4, 2, 1, 2, 4, 6, 3, 4, 6, 4, 7, 4, 1, 2, 1, 3, 1, 3, 6, 7, 6, 4, 7, 5, 7, 6, 3, 1, 3, 4, 1, 3, 4, 3, 1, 4, 1, 4, 2, 4, 2, 4, 3, 6, 7, 6, 4, 1, 4, 7, 6, 7, 5, 4, 1, 4, 1, 4, 5, 4, 1, 2, 5, 2, 5, 7, 5, 4, 1, 2, 1, 4, 7, 6, 7, 6, 4, 2, 5, 2, 5, 4, 3, 6, 7, 4, 1, 3, 1, 3, 1, 4, 7, 5, 2, 1, 2, 5, 4, 1, 3, 6, 4, 1, 2, 5, 7, 5, 2, 5, 2, 1, 4, 5, 2, 5, 7, 5, 2, 1, 3, 4, 2, 4, 6, 4, 7, 4, 6, 3, 1, 2, 4, 6, 7, 4, 2, 1, 2, 1, 3, 4, 6, 4, 1, 4, 7, 4, 3, 4, 5, 2, 5, 7, 6, 4, 6, 7, 6, 3, 6, 4, 1, 3, 1, 2, 5, 7, 5, 7, 6, 3, 6, 4, 2, 4, 5, 2, 4, 2, 4, 1, 3, 1, 4, 1, 3, 4, 3, 4, 2, 4, 5, 7, 4, 2, 4, 3, 4, 6, 4, 2, 4, 7, 4, 1, 4, 6, 4, 5, 4, 6, 7, 5, 2, 5, 2, 4, 5, 4, 7, 4, 6, 4, 1, 3, 6, 4, 7, 6, 7, 5, 2, 1, 4, 5, 4, 5, 7, 6, 3, 6, 7, 6, 7, 5, 2, 5, 4, 3, 1, 2, 1, 3, 1, 4, 1, 2, 4, 7, 5, 2, 4, 2, 1, 4, 5, 7, 5, 4, 1, 3, 6, 7, 6, 3, 6, 3, 4, 7, 6, 4, 5, 7, 4, 1, 3, 6, 3, 6, 3, 1, 2, 5, 7, 4, 5, 7, 5, 2, 5, 4, 2, 5, 7, 5, 7, 4, 5, 2, 4, 1, 3, 1, 2, 4, 2, 4, 1, 2, 4, 3, 6, 7, 5, 4, 1, 3, 4, 3, 6, 3, 4, 5, 4, 1, 4, 3, 1, 3, 4, 7, 5, 2, 4, 3, 1, 4, 6, 7, 5, 7, 5, 2, 4, 2, 4, 5, 4, 3, 4, 1, 4, 5, 7, 4, 2, 5, 4, 6, 4, 1, 2, 5, 4, 5, 4, 6, 7, 6, 3, 4, 7, 6]

Figure 6.3: The generated 0.1° encoder track values are shown. Color sequence numbers refer to the colors shown in Figure 6.1 and Table 6.1. As the track is cyclic, the start position of the track is arbitrary.

generated. This was done by performing a random walk on the color graph. After each step, the entries in the list of colors corresponding to the read heads at each division were compared to ensure that there were no repeated codes. If a repeated code was found, an exhaustive search was performed with the other possible colors to try to continue. If continuing was not possible, the last color step was removed, and a second exhaustive search was performed. If progress was still not possible, the process was restarted. After $O(10^8)$ iterations, a valid cyclic encoder track was found; this track is presented in Figure 6.3. Searches were performed with either one fewer sensor or with a color graph consisting of one fewer color, but no valid tracks were found. Although no tracks were found for such configurations, their existence has not been ruled out.

6.3 Prototype design and construction

A prototype sensor was designed and constructed using six identical AMS TCS34725 RGB color sensor integrated circuits[2] connected via a Texas Instruments PCA9548A I^2C multiplexer[3] to an STMicroelectronics STM32L031F6 microcontroller.[4] To reduce the circuit board size, the color sensors were placed on the backside of the PCB, while the rest of the components were placed on the front.

Since the color sensors have a wide field of view, ~120° FWHM, some sort of optics are needed to reduce this. An illumination source also needs to be installed near the sensors. For the prototype, both of these problems were solved by two additional PCBs. First, a spacer PCB with a circular cutout for each color sensor was used to clear the vertical height of the color-sensor chips. Then, a second PCB with a small drill hole centered above each color sensor was created to serve as an aperture stop to restrict the field of view; LEDs were also mounted to this PCB close to each aperture to illuminate the encoder track. Both PCBs were designed with castellated edges to allow them to be soldered to the main PCB after the color sensors were installed, both for mechanical attachment and to carry power to the LEDs. Matte-black solder mask was also used on all three PCBs to reduce problematic reflections. The assembled prototype can be seen in Figure 6.4.

Firmware was written for the microcontroller to read out the six color sensors and output the current absolute position via a UART serial interface. First, the RGB values are read from the sensors. Then, these RGB values are converted to XYZ values by applying calibration matrices, and then the XYZ values are converted to xy chromaticity values. An individual calibration matrix was created for each sensor by measuring printed patches of the seven encoder-track colors and then fitting the matrix using these RGB data and the

[2]Datasheet: `https://ams.com/documents/20143/36005/TCS3472_DS000390_3-00.pdf`
[3]Datasheet: `https://www.ti.com/lit/ds/symlink/pca9548a.pdf`
[4]Datasheet: `https://www.st.com/resource/en/datasheet/stm32l031f6.pdf`

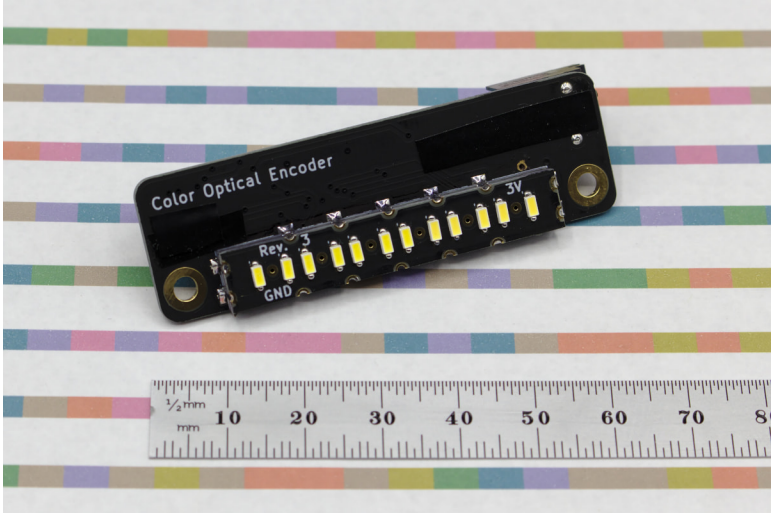

Figure 6.4: The assembled prototype sensor is shown with a printed test track. The yellow rectangular components on the PCB are LEDs, and the untented vias in between them serve as apertures for the color sensors.

fiducial xy values. Next, Euclidean distance is used to match each measured xy value to the closest encoder track color; these distances are also summed and recorded as an error estimate. Since the microcontroller has limited flash memory, only the track pattern and color sequence are directly stored. To determine the angular position, all 3600 positions are looped through and their color lists constructed on the fly until a match is found; if no match is found, the read is reported as an error. Finally, the measured position and error estimate are transmitted to a host system via a serial interface.

An encoder track was manufactured by printing the encoder track pattern onto matte-white paper using an inkjet photo printer; toner-based printing processes were found to introduce problematic reflections, necessitating the use of an inkjet printer. The track was split into 20 segments in order to fit in the available page area. These strips were then affixed to the outside edge of a ~1.5 m diameter aluminum ring that is free to rotate, and the assembled sensor was mounted adjacent to the ring. The mounted track and sensor are shown in Figure 6.5.

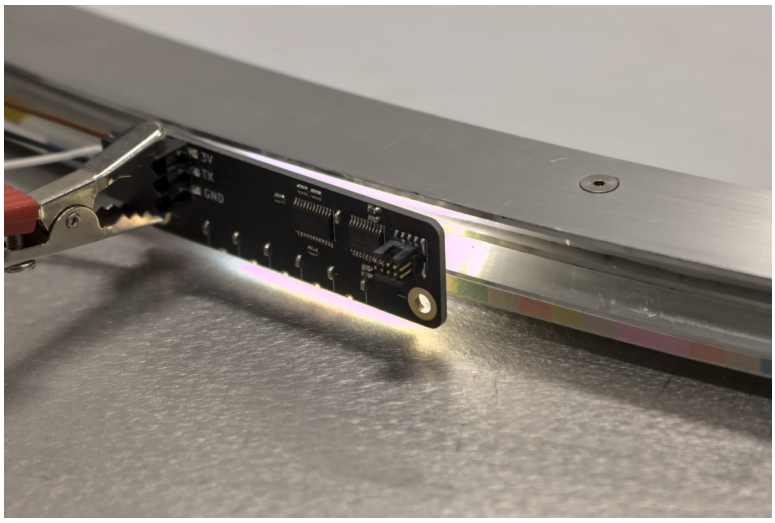

Figure 6.5: The fully assembled encoder prototype is shown installed in a position to read an encoder track installed on the edge of a ~1.5 m diameter ring.

6.4 Discussion

The rotating ring used for testing is part of a near-field calibration device for CLASS. The calibrator needs to be able to slide into a thin slot on the telescopes, at the base of the forebaffle, for certain measurements to be performed, which introduced significant thickness constraints. Due to these constraints, the rotating ring only has 3.2 mm of thickness available for an encoder track. Commercially-available absolute linear encoders were considered for measuring the angular position, but almost all of the encoders that were considered required tapes that were too tall to fit on the rotating ring, motivating the development of this single-track encoder. Furthermore, these encoders were much higher resolution than necessary, and the additional resolution resulted in tight alignment tolerances, which were impractical to maintain due to how the ring is mounted.

The developed color optical encoder concept is far more forgiving to misalignment than the previously-mentioned encoders, to the extent that it will even function when carefully held by hand. The encoder concept can also be used for a linear encoder, simply by dropping the cyclic requirement when generating the encoder track or by using just

a portion of the existing encoder track. However, while the concept is sound, further development is needed for the physical implementation to make it more practical. In particular, proper optics should improve performance, and the seven colors used should be optimized with the color sensors and printing process in mind. In the end, the plans for motorizing and automating the calibration device were dropped, so further development of the encoder concept was not pursued.

Chapter 7

Control and Systems Software

In order for CLASS to perform CMB measurements, software and computing and networking equipment are necessary to perform observations, collect time-ordered data, and package and transfer data for further analysis. To this end, this chapter describes the architecture of the software used to control the telescopes, acquire data from the various instruments, schedule observations, monitor the status of the instruments and observations, create archival data packages, and transfer data packages to North America for analysis. The telescopes operate at a high-altitude site in northern Chile located on Cerro Toco and are connected to the internet via a facility in the nearby town of San Pedro de Atacama. Data are then transferred to a data server located on the Johns Hopkins University campus in Baltimore, Maryland for further analysis. The remainder of this chapter is structured as follows. Section 7.1 gives an overview of the network and software architecture, Section 7.2 describes how this software is interfaced to some key subsystems, Section 7.3 details the data packaging and transfer pipeline, and Section 7.4 outlines the web interface used to monitor status and schedule observations. Finally, we present some lessons learned in Section 7.5 and conclude in Section 7.6.

Except for some minor changes, this chapter was previous published as Petroff et al. (2020b).

7.1 Network and software architecture

With multiple receivers and numerous housekeeping systems, more than a dozen computers are used to operate CLASS. Networking and a cohesive software control suite are therefore required to coordinate the operation of these systems such that science data can be collected, transferred off-site, and later analyzed.

7.1.1 Network layout

The CLASS network is divided into two main components, the site network on Cerro Toco and the supporting systems down the mountain in San Pedro de Atacama. These components are connected using a wireless network link operating at 5 GHz. This approximately 43 km line-of-sight radio link is provided by two Ubiquiti AF-5X transceivers[1] and is able to sustain data rates in excess of 100 Mbps. The site end of the wireless link is located on a tower ~170 m from the telescopes, at the edge of a cliff and with a clear line-of-sight to the San Pedro de Atacama end of the link. The site end of the link is powered via a solar panel and batteries and is connected to the rest of the site network via a fiber optic cable, providing complete electrical isolation, to protect against lightning damage. All computers are powered via uninterruptible power supplies (UPSs) so that they remain online during the brief power losses sustained while switching between the site's two generators. Similarly, the equipment in San Pedro de Atacama is powered with a UPS; this UPS is capable of providing several hours of backup power, which is necessary due to the frequency and duration of power outages at this location. An overview of the CLASS network is given in Figure 7.1.

The San Pedro de Atacama portion of the network consists of a router, an analysis computer, and a remote access computer. The ANALYSIS MACHINE is used for doing prelim-

[1] Ubiquiti Inc.; `https://www.ui.com/`

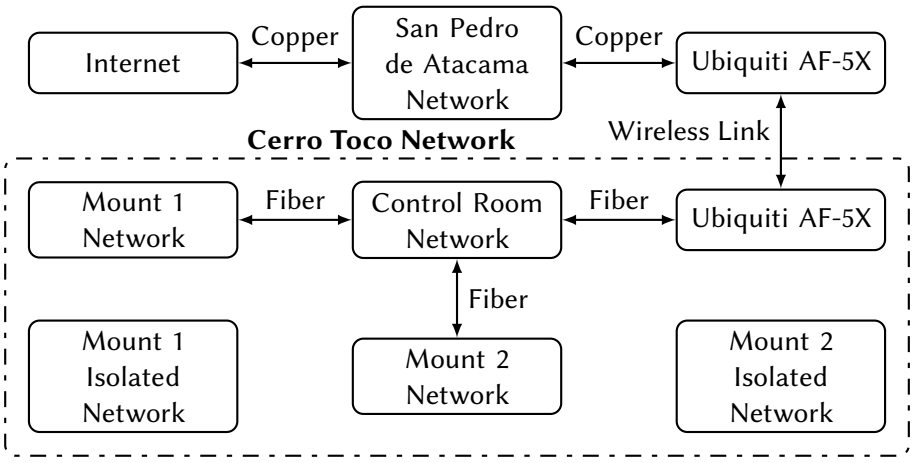

Figure 7.1: Network architecture overview. Network segments in San Pedro de Atacama and on Cerro Toco are connected via a wireless link; other segments are connected via either copper or fiber optic Ethernet cables.

inary analysis of data collected by the telescope. Additionally, it is used to store telescope data once it is transferred down the mountain, until the data are copied to North America. The Cerro Toco portion of the network consists of three major segments: the control room, the first mount, and the second mount. The CONTROL ROOM and MOUNTS all use managed Ethernet switches, which provide gigabit connectivity over both copper and fiber optic cables, and are interconnected via fiber optic cables. Each MOUNT has an additional switch for communication between the mount computers and the mount servo motors, which are isolated from the primary network. The radio link to San Pedro de Atacama is connected to the CONTROL ROOM switch. The MOUNTS each contain two computers for controlling the mount; two housekeeping computers, one for each cryostat; a star camera; four servo motors, which are on an isolated network; and two VPM CONTROLLERS, one for each telescope. The CONTROL ROOM contains a server rack with a CENTRAL SERVER, a display and status computer, a command terminal, a star camera host computer, and detector readout computers, one for each detector readout unit.[2] Additionally, there are

[2]The detector readout units are on the mount and interface with the detectors; they are connected to the readout computers via fiber optic cables.

laptop computers that can be moved around as needed, to fulfill miscellaneous needs. A Wi-Fi access point is also installed, although this is turned off during observations.

7.1.2 Software structure

CLASS's software infrastructure is currently built around the Ubuntu 20.04 Linux distribution[3] and the Python programming language,[4] specifically version 3.8. These releases are both supported through at least 2025. Software is developed using the Git[5] distributed version control system. All machines run the Ubuntu 20.04 operating system, except for the cryostat virtual machines that run Windows 7 for compatibility with vendor-provided proprietary software and the mount computers that run the VxWorks[6] real-time operating system (RTOS). With the exception of some C code required to interface with some hardware and a few shell scripts, all control and data acquisition software running on the Linux servers is written in Python. The control scripts are run as `systemd`[7] services so that they start on boot.

The basic software structure consists of Python scripts running each of the telescopes' subsystems, with a central COMMAND SCRIPT to coordinate operations. Since control and data acquisition software is running on multiple computers on the network simultaneously, a control and communications system for these distributed systems is required. The PYRO[8] package, version 4, was chosen for this purpose. Each subsystem is operated by a separate Python script, which includes a PYRO interface. These PYRO interfaces expose the scripts' control function to the network. The COMMAND SCRIPT is then able to call these exposed functions as if they were local functions, with PYRO seamlessly taking care of all of the

[3] Canonical Ltd.; `https://www.ubuntu.com/`
[4] Python Software Foundation; `https://www.python.org/`
[5] `https://git-scm.com/`
[6] Wind River Systems Inc.; `https://www.windriver.com/`
[7] `https://www.freedesktop.org/wiki/Software/systemd/`
[8] `https://pyro4.readthedocs.io/`

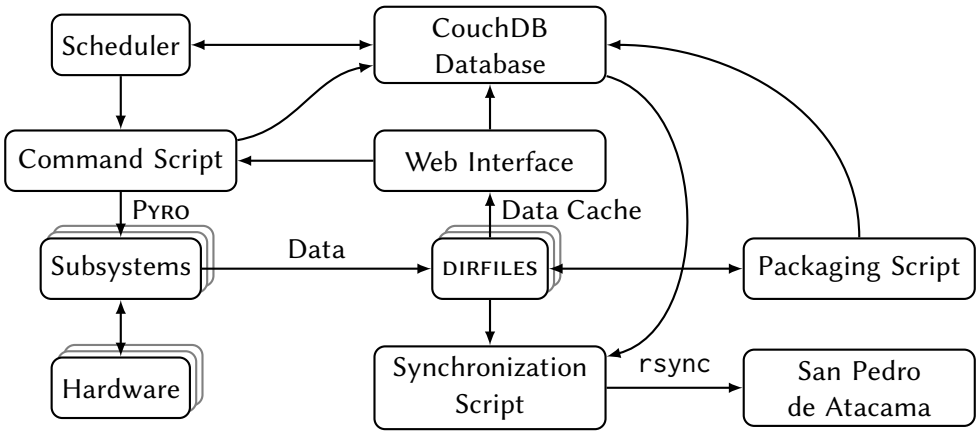

Figure 7.2: Software structure overview. Boxes represent various components or locations, and arrows represent data or command flow. The software is structured around Python scripts, with a web interface used for user interaction and DIRFILES used for data recording.

network communications. Thus, subsystems are controlled by calling the COMMAND SCRIPT, which in turn uses PYRO to call a function on the corresponding subsystem.

Figure 7.2 contains an overview of the software structure. All subsystems write data to the CENTRAL SERVER over the network, and a centralized database is used to store metadata. A SCHEDULER is used to execute preplanned observing routines, and a web interface is used to monitor system statuses and control the SCHEDULER. The standard data format is the DIRFILE standard for time-ordered data.[9] This filesystem-based, column-oriented format has been used for previous cosmology experiments, including the Atacama Cosmology Telescope (Switzer et al. 2008) and the EBEX (Milligan et al. 2010) balloon-borne experiment. It is supported by its reference GETDATA implementation (Wiebe et al. 2015), which includes the `pygetdata` Python bindings, and can be easily visualized using the Kst plotting software.[10]

[9] `http://getdata.sourceforge.net/dirfile.html`
[10] KDE e.V.; `https://kst-plot.kde.org/`

7.1.3 Scheduling

While a central COMMAND SCRIPT is sufficient, and preferable, for testing, its use for regular observations is highly inefficient and error-prone. Instead, a scheduling system is needed so an observation can be planned in advance. For a survey experiment like CLASS, once a reasonable observation strategy is found and tested, it will be used again and again. Therefore, a scheduling system that facilitates running such preplanned observations needs to be devised. One way to implement such a system would be to replace the COMMAND SCRIPT with a scheduler, calling the same PYRO functions as the command script. Although this method has the appeal of being simple and straightforward, it has the disadvantage of possibly producing different behavior than the COMMAND SCRIPT, which would lead to much confusion once a schedule is created using commands that were first tested by hand using the COMMAND SCRIPT. Therefore, the CLASS SCHEDULER calls the COMMAND SCRIPT, to ensure the exact same behavior.[11] The results of the commands are logged to a database. A schedule entry is created using the arguments to the COMMAND SCRIPT, exactly what would be typed into the command terminal. These entries are combined with lines containing a simple time delta syntax, relative to either the schedule's start time or the last time directive, as well as optional comment lines. An example of the SCHEDULER syntax can be seen in Figure 7.3.

7.1.4 Databases

While the science data are written to disk, it is helpful to keep various sorts of metadata in a database for easy access and manipulation. For CLASS, metadata about chunks of science data are stored in a database as well as a log of all commands executed and SCHEDULER data.

[11] More specifically, it calls a symbolic link to the COMMAND SCRIPT. The COMMAND SCRIPT uses this difference in the name of the script called to log the fact that the call was made by the SCHEDULER. There is also an additional mode that evaluates the command arguments but does not execute anything, which allows command syntax to be verified.

```
# This is a comment
MCEQ -autobias

# The following sets the commands after it to be run 2 minutes
# after the last time directive
/+2m
MCEQ -acqdata 50000000

# The following sets the commands after it to be run
# 2 days, 20 hours, 5 minutes, 30 seconds after the start time of the schedule
/2d20h5m30s
MCEQ -stop
```

Figure 7.3: Example of SCHEDULER syntax. Examples are shown of individual commands, comments, and relative and absolute time deltas.

Metadata for the science data collected are entered into a database to facilitate both data transfer and analysis. Therefore, a distributed database system is needed such that it can be accessed locally at the telescope site, in San Pedro de Atacama, and in North America. Three desirable properties for distributed computer systems, in this case a distributed database, are consistency, availability, and partition tolerance. However, Brewer's theorem states that only two of the three can be achieved by any one system (Gilbert and Lynch 2002). Consistency requires all copies of the database to return identical results when queried, availability requires the database to always be accessible, and partition tolerance requires the system to still function even if the network link between the different copies is severed. For example, if a database system enforces both partition tolerance and consistency, the database will be read-only, and thus only partially accessible, if the network link goes down, since the link is needed to keep the different copies consistent. For CLASS, availability and partition tolerance are the most important properties, since both the radio link to the telescope site and the internet connection to North America might be unreliable, but this should not stop data collection, which also involves adding metadata to the database. Although this could in general lead to merging problems when the database system tries

to restore consistency once a severed network link is reestablished, the software running at CLASS's three locations modify different database fields, so no conflicts will arise. Most data are entered at the telescope site, with the software at the two auxiliary sites only modifying fields pertaining to their local data storage locations.

To this end, the CouchDB database[12] was chosen for handling metadata, with a server at each location running a database server instance that hosts a full copy of the databases to allow for fast, low-latency database queries. CouchDB is a NoSQL database that stores data as JSON documents and is highly available and partition tolerant, with eventual consistency. NoSQL databases differ from traditional SQL databases by not requiring a fixed, predefined table structure; instead, each document can have any data keys it needs, functioning more like a dictionary than a spreadsheet. Instead of using structured queries to retrieve data, mapping and reduction functions are written and used; in the CouchDB parlance, these are known as "views" and are written in JavaScript.

The CouchDB instance hosts databases to store commands executed, schedules, the results of commands executed for schedules, and metadata about the science data products. The COMMANDS DATABASE records commands executed by the COMMAND SCRIPT along with a timestamp and Git revision hashes for both the COMMAND SCRIPT and the target script; this logging happens regardless of how the COMMAND SCRIPT was called—directly or by the SCHEDULER. Each commit in a Git repository has a unique 160-bit SHA-1 hash (National Institute of Standards and Technology 2015) associated with it; since all of the site software is contained in a Git repository, recording these revision hashes uniquely identifies the exact version of the scripts that were executed, for future reference. The SCHEDULES DATABASE contains schedules for use by the SCHEDULER. Each entry contains the actual schedule entered, a parsed copy of the schedule for use by the SCHEDULER, an ID number, descriptive tags, the schedule's author, a start time, and a status field.

[12] Apache Software Foundation; `https://couchdb.apache.org/`

Once the schedule finishes running, or fails, an end time is added. Failed schedules also record the line that failed. These entries are created by the SCHEDULER's web interface and are modified by it and the actual SCHEDULER. The SCHEDULER also records the results of commands executed in a JOBS DATABASE; each entry contains the command executed, a timestamp, the command's output, and if it failed. The final database contains metadata for each packaged chunk of CLASS science data. In addition to these standard metadata, which are also kept with the data packages, the database stores the location of the data packages at each site: Cerro Toco, San Pedro de Atacama, and North America. Each site only modifies the location field for itself, so synchronization issues after a network outage are avoided. These location fields are used to facilitate data transfer and analysis.

7.2 Hardware interfaces

To capture data, a physical instrument is required. To record these data and to control these instruments and other hardware, computer interfaces to the hardware are required, as is software to use these hardware interfaces. Although the hardware of different CLASS subsystems differs considerably, their software interfaces were designed to be similar, using a Python script that interfaces either directly with the hardware or through vendor-provided software to expose a PYRO control interface. Where custom electronics are used, such as for cryogenic diode-based thermometry, warm thermometry in the telescopes' optics cages, and for control of the telescopes' VPMs, these are designed around an Ethernet interface. While TCP is used for sending commands, UDP is used to stream data in a manner that is robust to connectivity interruptions and equipment restarts. For equipment that uses an RS-232 serial interface, such as AC resistance bridges (used for reading out ruthenium oxide cryogenic thermometers), magnetometers, the diesel tank level sensor, and the site weather station, Ethernet-to-serial converters are used to allow access over the site

network; this simplifies operations, since the control scripts for these equipment can be run from the CENTRAL SERVER.[13] Along with scientific instruments, other site infrastructure, such as the generators, can also be controlled remotely.

The only exceptions to this general control architecture are the cryostats. They are controlled using a proprietary GUI software provided by their manufacturer, Bluefors,[14] which only runs under the Windows operating system and does not provide much in the way of a scriptable interface. Therefore, the cryostats are operated manually, with the REMOTE DESKTOP PROTOCOL used to access the GUI. However, as the cryostats only need to be controlled when they are being cooled down or warmed up for maintenance, not during normal operation, this lack of integration with the rest of the site operations software does not impede observations.

7.2.1 Detectors

The telescopes' transition edge sensor detectors (Denis et al. 2009; Dahal et al. 2020) are read out via time-division multiplexing using Multi-Channel Electronics (MCE) units (Battistelli et al. 2008) developed at the University of British Columbia.[15] There is one MCE per receiver and currently one or two receivers per telescope mount, although the two mounts will eventually have two receivers each. The detectors are connected to the MCE using a set of SQUID multiplexers at the focal plane and a SQUID amplifier series array at the 4 K stage of the cryostat. The MCE is triggered using a FRAME PULSE, delivered with a clock signal over fiber optics from a SYNC BOX, one per telescope mount, which is a device used to synchronize data collection by multiple MCEs and the telescope mount. The FRAME PULSE assigns a 32-bit ID number to the readout frame, which consists of a full focal

[13] A USB-to-serial interface is used instead for the weather station, since it is mounted on the outside of the control room, a short distance from the CENTRAL SERVER.

[14] Bluefors Oy; `https://bluefors.com/`

[15] `https://e-mode.phas.ubc.ca/mcewiki/index.php/Main_Page`

plane readout, and delivers pulses to trigger reading out each detector. The FRAME PULSE is delivered at slightly over 200 Hz.[16] The MCE, mounted on the cryostat, communicates with a HOST COMPUTER in the CONTROL ROOM via fiber optics and a PCI readout card. This HOST COMPUTER uses MCE ACQUISITION SOFTWARE (MAS) SCRIPTS, provided by the University of British Columbia, to control the MCE and record data. A PYRO interface allows use of these scripts over the network and also provides a routine for measuring I–V response and applying appropriate detector biasing. Data are written to an NFS mount on the CENTRAL SERVER.

7.2.2 Mount

Each telescope mount is controlled by an industrial MOUNT COMPUTER running the VxWorks RTOS, which loads its software over the network at boot time. The MOUNT COMPUTER, also known as the antenna control unit (ACU), controls the mount's four servo motors, two for azimuth and one each for elevation and boresight angle, over an isolated Ethernet network and also reads in various encoders and tiltmeters. The MOUNT COMPUTER also receives FRAME PULSE information from the SYNC BOX over RS-485 (instead of the fiber optics used for the MCEs) for synchronization with the MCE units and has a GPS receiver to obtain precise timing information; as the MCEs only record FRAME PULSE information and not time, the mount's timing information is also used to assign times to detector data when the data are packaged. The mount's position is determined using encoders that feed a pointing model, and these data, as well as timing and status information, are written to an NFS mount on the CENTRAL SERVER. The pointing model is derived from Moon and planet observations (Xu et al. 2020), although the initial pointing model was constructed using star camera observations processed using the Astrometry.net blind astrometric solver (Lang

[16] Running slightly fast is preferable to running slightly slow, since it ensures the recorded data chunks are less than ten minutes in length. This simplifies the data packaging process, since it guarantees that a data chunk only needs to be split once.

et al. 2010).[17] The MOUNT COMPUTER exposes a Telnet interface and a TCP interface to the site network. The Telnet interface allows for interactive control via a terminal and status monitoring, while the TCP interface is used for scripted control of the mount. A PYRO interface running on the CENTRAL SERVER translates this TCP interface into a standard PYRO interface for the COMMAND SCRIPT.

The mount monitor and control system, which is written in C++, implements the programs used for scanning in azimuth, with sun avoidance to preclude the boresight from coming within 20° of the Sun; sky dips, for atmospheric analysis; and drift scans of planets and the Moon, which are used for pointing and beam analysis. The JPL DE430 ephemerides (Folkner et al. 2014) are used by the monitor and control system for tracking the Moon, Sun, and planets. Parameters used by these programs are passed to the monitor and control system via the PYRO interface.

7.2.3 Variable-delay polarization modulator

The VPM consists of a fixed wire array and a movable mirror, to modulate the incoming polarization signal. The mirror is driven by a closed-loop motion control system. This system consists of voice coils to actuate the mirror and encoders to measure its position. The encoder signals are duplicated, with one copy going to the VPM CONTROLLER and another copy going to the ACU. The VPM CONTROLLER, built around an industrial applications processor and running a network-booted RTOS, uses this information to maintain proper mirror position, while the ACU logs the encoder data. This duplication is necessary because the VPM CONTROLLER is not connected to the SYNC BOX, and the VPM encoder positions need to be recorded synchronously with the detector data. A control script running on the CENTRAL SERVER communicates with and configures this CONTROLLER via TCP over

[17] The star camera consists of a networked machine vision camera in a waterproof enclosure that is rigidly attached to the telescope mount.

an Ethernet connection. Finally, a Pyro interface exposes a standard network interface for VPM control.

7.3 Data pipeline

In order to measure the polarization of the CMB, data need to be acquired from the telescopes' detectors and various housekeeping systems, combined, packaged, and transferred to North America for analysis. This packaging needs to happen in real time, to prevent a backlog and maximize observing time; furthermore, the process should be lossless and verifiable to prevent data loss and identify potential data corruption.

7.3.1 Acquisition

Data acquisition starts with scheduling an observation and acquiring data from hardware instruments. As data are acquired, they are written as dirfiles on the central server, either directly or using NFS mounts. These are structured with a directory for each computer acquiring data, with subdirectories for each acquisition system. Individual dirfiles are named using a timestamp formatted as `%Y-%m-%d-%H-%M-%S`, e.g., `2020-03-22-20-00-00`. Dirfiles consist of a directory that contains a plain-text format file and separate binary files for each stream of time-ordered data. Additional derived fields can be defined as linear combinations of existing fields, e.g., to scale raw data to SI units; aliases, e.g., to label the location of thermometers without having to edit the acquisition code; and extracted bit-fields, e.g., to combine multiple status flags in one raw data field.

Data are divided into ten-minute chunks as they are acquired. For asynchronous housekeeping data, which are not acquired using the sync box, a new chunk is started at clock-aligned intervals, e.g., at 00:00, 00:10, 00:20, etc., so the first data chunk in the series will be less than ten minutes in length, e.g., six minutes if data collection started at

twenty-four minutes past the hour. These divisions are done when the time in seconds since the Unix Epoch[18] modulo 600 seconds rolls over. Synchronous data are acquired in arbitrarily-aligned chunks of approximately ten minutes in length.[19] All timestamps are in Universal Time, with asynchronous data using the computer's NTP-synchronized system clock in UTC and the mount using UT1, derived from GPS time; the detector data are not acquired with timestamps, but they are synchronized with the mount data using the SYNC BOX. The time is stored in the DIRFILES as seconds since the Epoch.

7.3.2 Packaging

Once data are acquired, they must be packaged into a standard data product. The standard CLASS data product is a clock-aligned ten-minute DATA CHUNK consisting of time-ordered data, with one DIRFILE per mount for synchronous data and individual DIRFILES for each asynchronous data acquisition system. Synchronous data are combined in this data product, since this can be done without altering any data or losing any information. Since the DIRFILE format keeps each data field in a separate file and the GETDATA library does not load the data until a specific field is read, this combining of data from multiple receivers does not add any overhead when only one receiver's data need to be read. Asynchronous data are kept separate due to the losses inherent to resampling and interpolation. This is a trade-off between making things simpler during analysis and preventing data loss, since interpolation eventually needs to be done. A packaging script is run every ten minutes, at 00:02, 00:12, 00:22, etc. It processes data from two data periods prior and before, e.g., the 00:32 run will only process data from the 00:10 period and before; this ensures that the non-ten-minute-boundary-aligned synchronous data collection for the data period has finished. For example, synchronous data collection might start at 00:19, which would

[18] `1970-01-01T00:00:00Z`

[19] The chunks are 120 000 frame pulses long, or just under ten minutes, "redefining" the second as 200 FRAME PULSES.

include data in the 00:10 period; this collection would not end until 00:29, unlike the asynchronous data collection, which would start a new data chunk at 00:20. Therefore, the maximum delay, plus a few minutes of buffer, is used.

Synchronous

The detector and ACU data sampling on a given mount is synchronized, but the two mounts are not synchronized with each other. An overview of the synchronous data packaging process can be seen in Figure 7.4. The first step in packaging synchronous data is to divide them into CLASS's standard ten-minute data intervals. This process starts with dividing each data chunk collected by the ACU into two parts, `A` and `B`, where seconds since the Epoch modulo 600 seconds rolls over; the ACU data must be the first processed, since they are the only synchronous data that contain timestamps, instead of just FRAME PULSE numbers. These divided chunks are then combined to form ten-minute chunks aligned with the Epoch, with zero padding added before, after, and between chunks when needed to form a full ten-minute chunk. The field containing the FRAME PULSE number is filled in with the appropriate values instead of zero padding like the rest of the fields; this allows for easier synchronization with the detector data. A status field is also added to the resulting DIRFILE that shows where the ACU data are valid, as opposed to where they are zero-padded, stored in an 8-bit integer field.

Next, the detector data from each MCE unit on the mount are divided using the FRAME PULSE numbers recorded in both the MCE data and the now aligned ACU data. As with the ACU data, these chunks are then combined to form aligned ten-minute chunks, again with zero padding added before, after, and between chunks as needed; again, a status field is added, but different bits in the 8-bit integer field are used. Care must be taken, since the FRAME PULSE counter occasionally rolls over.[20] Finally, the raw data files in each of

[20] As the FRAME PULSE is a 32 bit unsigned integer, this happens roughly every eight months. This counter

the MCE DIRFILES are moved into subdirectories in the ACU DIRFILE. The ACU DIRFILE FORMAT FILE is altered to include the MCE DIRFILE FORMAT FILES as child fragments, which prefix the MCE DIRFILE field names with the appropriate MCE name, and the data valid status fields are combined using a bitwise OR. After this process is complete, the resulting DIRFILE can be treated as an asynchronous data chunk in the next step of the packaging process.

The merged data are then compressed, with `gzip` compression used for the mostly floating-point ACU data and FLAC compression[21] used for the fixed-precision MCE data. As FLAC was designed as an audio compression codec, it does not support the 32-bit data recorded by the MCE, but GETDATA avoids this limitation by splitting each 32-bit integer into two 16-bit integers and compressing the data as two separate channels. However, this is still not optimal, since FLAC will treat these integers as smaller than 16 bits if the most significant bits are never used. CLASS operates the MCEs in DATA MODE 10, which uses the lower 25 bits of the 32-bit integer to store low-pass-filtered detector data as a signed integer and the upper 7 bits to store a flux jump counter as another signed integer. Thus, when the flux jump counter is zero or positive, the upper half of the 32-bit integer is treated as having fewer than 16 bits, often as having as few as 9–10 bits, but is treated as having the full 16 bits when the flux jump counter is negative. Therefore, the data packaging script removes the 7-bit flux jump counter from the combined data stream and converts it to a separate 8-bit signed integer field, which is FLAC-compressed. The filtered detector data that remain in the original 32-bit field is then treated as a 25-bit integer when FLAC compression is applied. This scheme performs well on the CLASS data, reducing disk space usage while still providing rapid decompression and data access.

also resets any time the SYNC BOX is power-cycled.

[21] Xiph.Org Foundation; `https://xiph.org/flac/`

The process starts with misaligned synchronous data.

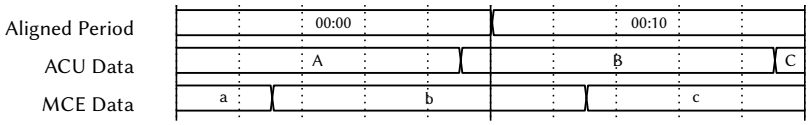

ACU data are then split at the proper alignment based on timestamps...

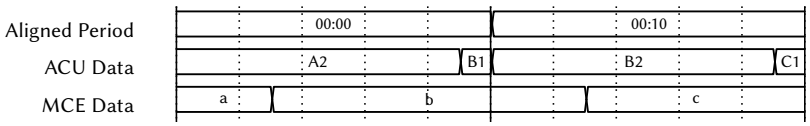

...and then joined to form aligned chunks.

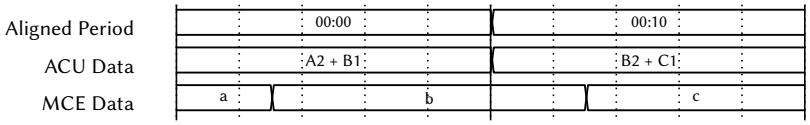

MCE data are then split at the proper alignment based on FRAME PULSE numbers in the ACU data...

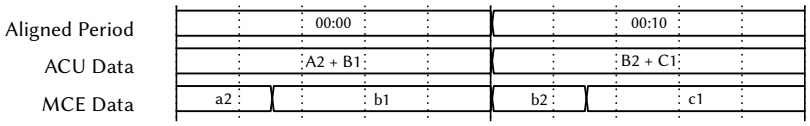

...and then joined to form aligned chunks.

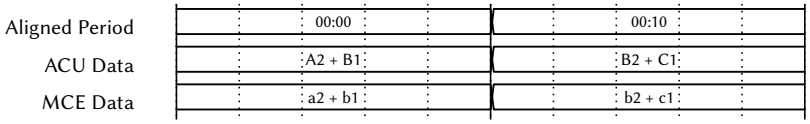

Figure 7.4: Overview of synchronous data packaging process. The procedure by which ACU (mount) and MCE (detector) data are aligned and joined is shown.

Asynchronous

Asynchronous data are recorded to timestamped DIRFILEs in a directory structure based on the computer recording the data. However, the final data product structure starts with a timestamped directory, with subdirectories for different types of data, a hierarchy divorced from the computer systems used to collect the data. Figure 7.5 contains an overview of this structure. Therefore, the first step of the asynchronous data packaging process, which includes the merged synchronous data, rearranges the asynchronous DIRFILEs from their original structure to the final data product structure, in a new location. As the vendor-provided software for operating the cryostats only records CSV files, these data are converted to DIRFILEs at this step. While the DIRFILEs are originally recorded uncompressed, at this stage they are compressed using the `gzip` compression format to save space.[22] Images from each of the telescope site's monitoring cameras are also packaged, since this provides a method of checking for bad weather or unusual site activity should artifacts be found in the telescope data during analysis.

Next, the DIRFILEs are converted to uncompressed ZIP file[23] archives. An uncompressed ZIP file concatenates all of the files it contains together and then includes a file offset table that allows for its contents to be located and read in a manner that allows for random reads. Encapsulating the DIRFILE into a ZIP archive reduces the number of files in the data package by more than an order of magnitude, which provides significant speed improvements for data transfer and backup operations, due to the small size of many of the files contained in the unencapsulated DIRFILEs.[24] Finally, a JSON metadata file is created that contains basic information about the DATA CHUNK, including the files included, their sizes, and SHA-1 checksums, for later verification; this file is saved in the root of the DATA CHUNK. The

[22] FLAC compression is not used for the asynchronous data, since many of the data fields involve floating point numbers, which FLAC compression does not support.

[23] ISO/IEC 21320-1:2015

[24] Unfortunately, this extended DIRFILE functionality is only available as a patch, since the GETDATA library has not received updates in some time.

location of the DATA CHUNK on disk is then appended to the metadata information, and this information is saved to the database.

7.3.3 Transfer

Once the data are packaged, they need to be transferred off the mountain to San Pedro de Atacama and then to North America. Data integrity needs to be verified after each transfer, and old data eventually need to be deleted from the CENTRAL SERVER to free up disk space. Data are transferred from the telescope site's CENTRAL SERVER to the ANALYSIS MACHINE in San Pedro de Atacama via the CLASS network's radio link. A synchronization script runs on the site's CENTRAL SERVER every ten minutes, which transfers the data files using `rsync`.[25] This script uses a lock file to ensure only one copy of itself is running at a time, and a certificate is used to authenticate with the ANALYSIS MACHINE. A verification script is also run every ten minutes on the ANALYSIS MACHINE in San Pedro de Atacama. This script checks each unverified DATA CHUNK on the machine against the local copy of the database and sees if the DATA CHUNK has a local path assigned to it. If it does not have a local path, the DATA CHUNK is verified against the database, and the CHUNK's local path is added to the database. If the verification fails, an error is thrown, and the local path is not added to the database, so it is treated as if the data were never transferred. The corrupted data might be fixed by `rsync` with the next synchronization run—else, manual intervention is required to delete the corrupted data. If the DATA CHUNK already has a local path, the data have already been verified, so the CHUNK is ignored. Since the synchronization script uses `rsync`'s delayed updates feature, files do not appear until the transfer is complete, so issues with incomplete transfers are mitigated. This process is repeated to transfer data from the ANALYSIS MACHINE in San Pedro de Atacama to a data server located in Baltimore.

[25] `https://rsync.samba.org/`

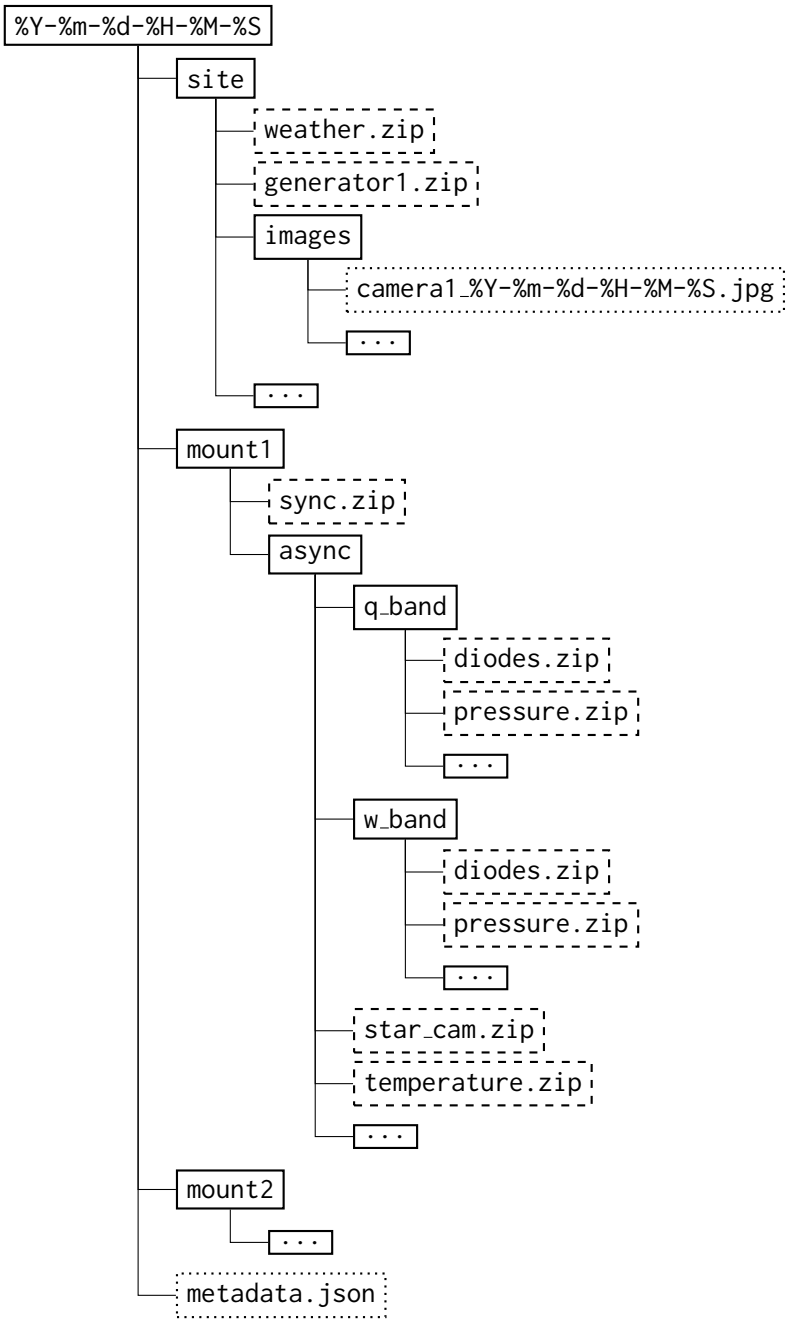

Figure 7.5: Overview of the CLASS data product structure. Solid boxes represent directories, dashed boxes represent DIRFILES, and dotted boxes represent other files. Only a subset of the DIRFILES and images are shown. The subdirectories in the async directory contain asynchronous data associated with a specific receiver, while the sync.zip DIRFILE combines ACU data for a given mount with detector data from both of its receivers.

Archiving of data and deletion of old data

While a permanent copy of the data will be kept on the file server in Baltimore, the data will not be permanently kept on disk on-site or in San Pedro de Atacama, as disk space will need to be freed for new observations. On both the CENTRAL SERVER on Cerro Toco and the ANALYSIS MACHINE in San Pedro de Atacama, a script is run on a daily basis to free up space by removing old DATA CHUNKS that have already been transferred and verified, to ensure a minimum amount of free disk space; as long as enough space is free, data will not be deleted. When disk space needs to be freed, the database will be checked for DATA CHUNKS that have already been transferred to San Pedro de Atacama in the case of the telescope site's CENTRAL SERVER and transferred to North America in the case of the ANALYSIS MACHINE. Of these DATA CHUNKS, data will be deleted in chronological order, with the oldest data first, until enough free disk space is available; currently, this threshold is set at 75% of the disk capacity. If this process is not able to free enough space, an error will be thrown, and manual intervention will be required. The only existing copy of a DATA CHUNK will never be automatically deleted.

Ensuring data integrity

A number of steps are taken throughout the data packaging and transfer process to ensure data integrity. At each step in the process, data are not deleted until after the next step in the process has finished. Updating the database with a DATA CHUNK's location at a new location is always the last step in the process, to avoid the possibility of the database showing that the data are there when they are not actually there; if something goes wrong in the transfer process, some or all of the data might be there, but the database will be conservative and list the data as not being present.

As the last step in the packaging procedure, SHA-1 checksums are created for each file and stored in the database. These checksums are unique, but repeatable, 160-bit

cryptographic hashes based on the data; if even one bit in the source file changes, a completely different hash will be generated, so transfer errors of either the data or the checksums are easily detected. After each transfer step, new checksums are generated for each file, and these checksums are compared to the copies stored in the database. If the checksums match, the transfer was successful. If they don't match, the transfer failed, and data corruption occurred; in this case, the data are treated as if they were not transferred at all, and an error is thrown. Along with being replicated across multiple locations, regular offline backups are made of the CLASS CouchDB database that stores the science data product metadata.

The primary CLASS data server located in Baltimore utilizes a ZFS file system.[26] ZFS provides strong data integrity guarantees, including against silent data corruption, by making extensive use of checksums and through regular integrity checking. A `RAIDZ2` configuration is used to protect against disk failures, and regular file system snapshots are taken to protect against inadvertent data deletion, caused either by human error or malware. Nightly backups are made to a replica data server, located in a different location, to provide additional protection. These backups rely on ZFS's ability to efficiently send and receive file system snapshots and are configured in such a way that even if one of the two servers is compromised by a malicious actor, it would be difficult to erase both copies of the data. Finally, the data packages are additionally backed up to a cloud storage provider. As of June 2021, CLASS has 22 TB of compressed data packages recorded to disk.

7.4 Web interface

A web interface is run on the CLASS site's CENTRAL SERVER to provide status information, facilitate execution of commands and scheduling of observations, and provide a wiki to

[26] `https://openzfs.org/`

store site operations information and documentation. Except for the wiki, which uses the MediaWiki software package,[27] the interface runs using Python and the Django web framework[28] and uses the Bootstrap[29] front-end framework.

The status display provides pertinent status information about the telescope systems, including cryostat temperatures, pressures, and flows; available disk space; and site environmental conditions. A screenshot of the status display overview section can be seen in Figure 7.6, which provides a color-coded dashboard for quickly evaluating the current instrument health as well as sparklines (Tufte 2006) for evaluating how parameters have changed over the past day. In addition to the overview, detailed pages are available for each cryostat and mount. Here, data are graphed client-side using a JavaScript charting library; multiple time ranges can be selected, and real-time updates are pushed using a WebSockets connection for the shortest time range. A script continuously reads in housekeeping data from the currently recording DIRFILES and stores it in a day-long circular cache; this cache is flushed to disk every ten minutes, so it will still contain data if the caching script needs to be restarted. This caching script exposes a PYRO interface, which the web interface back-end uses to access the data. The script additionally triggers the pushing of the real-time updates over the WebSockets connection. Due to the packaging process, only the currently recording DIRFILE is accessible, necessitating the cache, which is also useful to improve performance. Although the web interface back-end could directly cache the data, this is not done, since the web server runs multiple, simultaneous copies of the back-end to maintain responsiveness; this would cause duplicate disk access and, more importantly, cause duplicate real-time updates to be pushed. In addition to providing information for the status display, the caching script also sends alerts[30] for certain failures

[27] Wikimedia Foundation; `https://www.mediawiki.org/`

[28] Django Software Foundation; `https://www.djangoproject.com/`

[29] `https://getbootstrap.com/`

[30] These are currently sent to a channel in the *Slack* chat tool used by the CLASS collaboration, although the destination of the alerts can be easily changed.

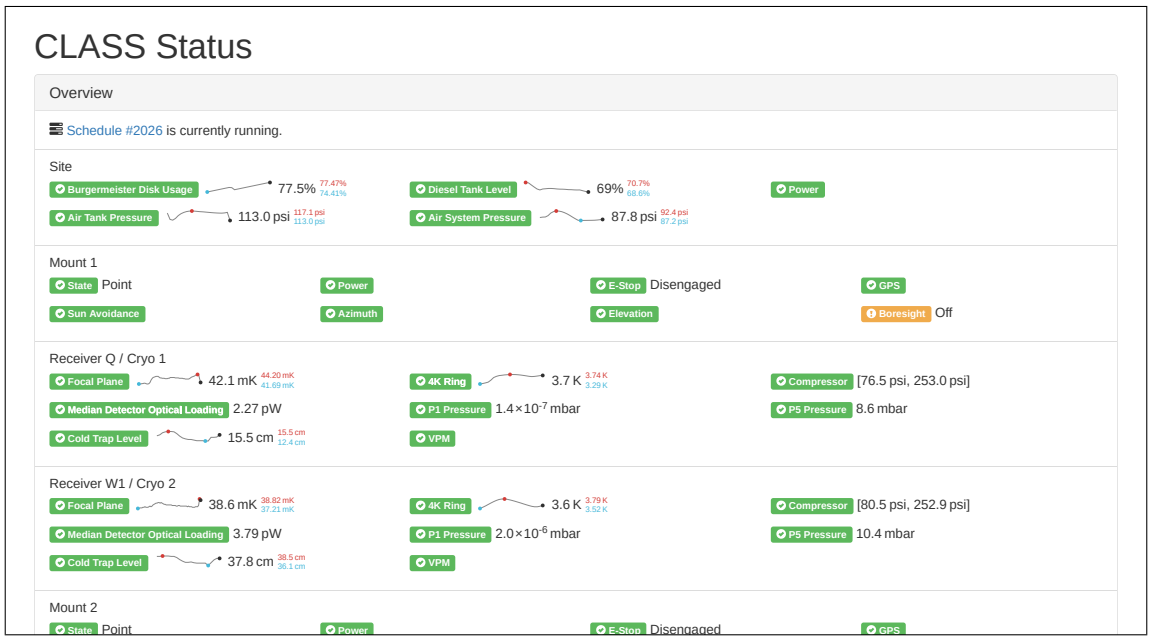

Figure 7.6: Web status interface. This screenshot shows an example view from the web status interface overview section, which allows one to quickly evaluate instrument health.

that require immediate attention, such as a cryostat warming up, although in some cases alerts are directly triggered by the corresponding hardware interface script instead. The status display also provides access to results from automated preliminary analysis scripts, which are run every morning on the ANALYSIS MACHINE in San Pedro de Atacama, and to current views from the various site observation cameras, either as static images updated once per minute or as live video.

The SCHEDULER interface provides an easy-to-use, graphical method for scheduling observations on the telescopes. Login credentials are required to access and use the SCHEDULER interface. The main screen shows the currently running schedule if one is running, the next pending schedule if no schedule is currently running but one is pending, or the last schedule run if none are currently running or pending. Figure 7.7 shows this interface. A sidebar provides a list of recent schedules and their status information, with a link to a complete history of schedules. When viewed, each schedule can be displayed either in the raw form it was entered in, with basic syntax highlighting, or in a parsed

form. Status, start time, title, author, schedule number, and tagging information are also displayed. If a schedule completed, the end time is shown. If a schedule failed to complete, the failure time and failed command are displayed. Pending schedules can be edited, and pending or running schedules can be canceled. An interface for adding new schedules is also accessible from the main screen. This prompts a user to enter metadata including a title, start time, and, optionally, tags; a date and time picker is provided to assist the user with start time entry. The author field is automatically filled in with the user's name, but can also be edited.[31] Then comes a box for entering the schedule; basic syntax highlighting is provided. Any time in the editing process, the user can switch between the edit box and a verification tab, which verifies the syntax of the entered schedule; lines containing incorrect syntax are highlighted. Once finished, the schedule can be submitted and will be run by the SCHEDULER. Previously-run schedules also provide a button for opening the new schedule interface with its fields pre-filled with data from the previous schedule. This interface for adding schedules can be seen in Figure 7.8. A separate interface is also provided for executing, and displaying the results of, commands outside the context of a schedule. This interface is normally locked out while a schedule is running to prevent accidental execution of commands, but it can be unlocked if necessary. Additionally, it displays a list of recently run commands and shows in real-time which other users have the command interface open, to encourage communication.

7.5 Lessons learned

Some parts of the architecture described above were different when the first CLASS telescope was deployed in 2016, and other parts should have been done differently in hindsight or would be done differently without certain hardware interface restrictions.

[31] The username of the user that entered the schedule is also recorded to the database as a separate field but is not displayed.

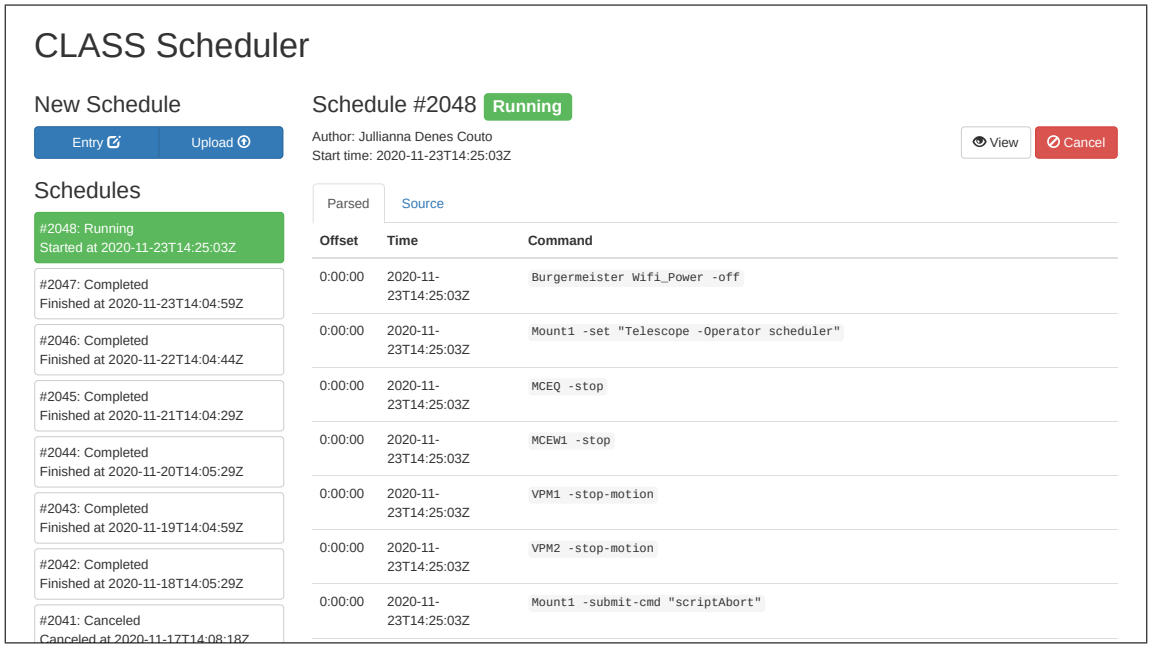

Figure 7.7: Web-scheduler-interface main screen. The main screen of the web scheduler interface is shown in this screenshot, displaying the currently running schedule.

Figure 7.8: Web-scheduler new-schedule interface. This screenshot shows a new schedule being written; metadata can be added, the start time can be adjusted, and the schedule can be validated before it is submitted to be run.

Other changes have been incremental updates, such as transitioning from Python 2 to Python 3.[32] One lesson that was learned is that interfacing with instruments over Ethernet with a control script running on the CENTRAL SERVER is easier and more reliable than running a control script on another computer and writing data over NFS, since the former is a "pull" configuration, while the latter is a "push" configuration; "pull" configurations are less likely to run into issues if network connectivity is temporarily lost or if the CENTRAL SERVER needs to be rebooted. With this insight, thermometry data recording was transitioned from using USB interfaces to using Ethernet interfaces, which improved reliability. Ideally, this transition would be extended to its logical conclusion, using a single server to operate the telescopes and all the instruments. Unfortunately, limitations with the MCE interfaces and cryostat interfaces prevent this. MCE interface limitations additionally hamper improvements to data acquisition synchronization. Without these limitations, the detector, pointing, and VPM encoder data acquisition hardware could be configured to receive time information over the network via the precision time protocol standard, which allows for sub-microsecond time distribution. Data frames could then be synchronized with the start of GPS time seconds, as long as the data acquisition rate is limited to an integer number of samples per second, allowing for the elimination of the dedicated synchronization hardware that is currently used. Another lesson learned is that a large number of small files makes data transfer and backup operations much slower than they would otherwise be. This led to a transition from standard DIRFILES to DIRFILES encapsulated in ZIP files, which reduced the number of files stored on disk by more than an order of magnitude and significant performance improvements for data transfer and backup operations. In hindsight, it would have also likely been better if the data packages were based on hour-long chunks instead of ten-minute chunks for the same reason, but

[32] Although Python 3 was fairly well established when software development started in 2014, Python 2 was originally used, since, at the time, the GETDATA library did not yet support Python 3.

the current data package chunking scheme is too ingrained to change now for what would be a more incremental improvement.

7.6 Conclusions

The current CLASS control and systems software architecture has been described. This architecture is based around Python scripts, with DIRFILES used to record data and the PYRO library used for network-based control interfaces. Data are packaged into ten-minute long chunks, which are aligned with the start of the hour. Scheduling and status monitoring are performed using a web interface. Throughout, there is a focus on ensuring data integrity. It is our hope that our documenting of the CLASS software architecture here can help inform software development for future experiments. The CLASS software architecture has been successful at fulfilling its intended purpose and has proven to be sufficiently reliable.

Chapter 8

A First Detection of Atmospheric Circular Polarization at Q Band

In the presence of Earth's magnetic field, molecular oxygen in the atmosphere experiences Zeeman splitting of its magnetic dipole transitions. This produces polarized emission in the millimeter-wave regime, primarily circular polarization, which manifests as a foreground for polarization-sensitive ground-based cosmic microwave background (CMB) experiments such as the Cosmology Large Angular Scale Surveyor (CLASS) (Eimer et al. 2012; Essinger-Hileman et al. 2014; Harrington et al. 2016). The effect of this foreground has previously been discussed qualitatively by Keating et al. (1998) and quantified by Hanany and Rosenkranz (2003) and Spinelli et al. (2011). However, previous attempts to observe this foreground, such as by MIPOL (Mainini et al. 2013), have not been successful in detecting it.

Molecular oxygen has strong emission lines in the 50 GHz to 70 GHz range, as well as a line at 118.8 GHz, which are in the frequency range of interest for CMB observations. There are also water vapor and ozone emission lines near this frequency range, but unlike

Except for some minor changes, this chapter was previous published as Petroff et al. (2020c).

diatomic oxygen, these molecules do not experience Zeeman splitting and thus do not produce polarized emission (Liebe 1981). Zeeman splitting of molecular oxygen is generally considered in the context of remote sensing of the temperature of the mesosphere, where individual Zeeman-split emission lines can be resolved; pressure broadening obscures individual lines at lower altitudes (Meeks and Lilley 1963).

Modeling of the polarized emission of Zeeman-split oxygen began with the seminal works of Lenoir (1967) and Lenoir (1968), with further development of atmospheric emission models by Liebe (1981) and Liebe (1989), Liebe et al. (1992) and Liebe et al. (1993), Rosenkranz and Staelin (1988), and others. Near the emission lines, both linearly and circularly polarized emission have been detected from orbit (Schwartz et al. 2006; Kunkee et al. 2008). Ground-based detections of linear polarization have also been made at the emission lines, at 234 GHz (Pardo et al. 1995) and 53 GHz (Navas-Guzmán et al. 2015). However, atmospheric remote sensing instruments generally do not observe polarization at frequencies far from the emission lines.

As the circular polarization predicted by standard cosmological models is many orders of magnitude smaller than that of the linearly polarized signal (Inomata and Kamionkowski 2019), CMB experiments are designed to be primarily sensitive to linear polarization, which is itself an extremely faint signal (Hu and White 1997). For ground-based experiments, the linearly polarized signal from Zeeman-split molecular oxygen is expected to be on the nK level, roughly four orders of magnitude weaker than the corresponding circularly polarized signal (Hanany and Rosenkranz 2003). As this is much fainter than can be detected with current CMB instruments, and as the circularly polarized component is much stronger, atmospheric Zeeman emission is primarily considered a CMB foreground at the largest angular scales ($\ell \gtrsim 2$), fixed in topocentric coordinates, due to potential circular-to-linear polarization leakage as a result of instrument non-idealities (O'Dea et al. 2007). For linear polarization, emission and scattering by ice crystal clouds in the upper troposphere is a

larger atmospheric contaminant (Pietranera et al. 2007; Takakura et al. 2019).

For large angular scales at Q band (~30 GHz to ~50 GHz), diffuse Galactic synchrotron emission is expected to be the largest contributor to extraterrestrial circular polarization, with circular polarization due to Faraday conversion of linear polarization induced by Population III stars exceeding this contribution at smaller angular scales (King and Lubin 2016). Another possible source of Faraday conversion is via galaxy cluster magnetic fields, at primarily small angular scales (Cooray et al. 2003). Beyond synchrotron emission and Faraday conversion, additional potential sources of circular polarization include scattering by the cosmic neutrino background (Mohammadi 2014), primordial magnetic fields (Giovannini 2009), photon–photon interactions in neutral hydrogen (Sawyer 2015), and cosmic birefringence via coupling of the Chern–Simons term (Carroll et al. 1990), as well as postulated new physics (Zarei et al. 2010; Tizchang et al. 2016). As these predicted signals are at most on the nK level, they are well below current detection thresholds.

The most stringent previously published upper limit on CMB circular polarization was set by the balloon-borne Spider instrument, utilizing non-idealities in its half-wave plate polarization modulators at frequencies near 95 GHz and 150 GHz (Nagy et al. 2017). As Spider observed from the stratosphere, above much of the atmosphere and thus much of the atmospheric emission, it was not sensitive to atmospheric circularly polarized emission. However, for ground-based experiments to significantly improve on this limit, circularly polarized atmospheric emission must first be detected and subtracted. As with Mipol (Mainini et al. 2013), previous ground-based measurements did not have the requisite sensitivity to detect circularly polarized atmospheric emission (Lubin et al. 1983; Partridge et al. 1988). An improved upper limit on extraterrestrial circular polarization utilizing CLASS observations is presented in Padilla et al. (2020).

Through its use of Variable-delay Polarization Modulators (VPMs) (Chuss et al. 2012; Harrington et al. 2018), CLASS is uniquely capable of measuring large angular scale circular

polarization. CLASS currently observes the microwave sky in frequency bands centered near 40 GHz, 90 GHz, 150 GHz, and 220 GHz, using three telescope receivers; a fourth receiver will be deployed in the future. The present analysis focuses on the first two years of observations from the Q-band telescope, which is centered near 40 GHz. CLASS is designed to map the polarization of the CMB at large angular scales over 70% of the sky to detect or place an upper limit on the B-mode signal of primordial gravitational waves and to measure the optical depth due to reionization, τ, to near the cosmic variance limit (Watts et al. 2015; Watts et al. 2018).

VPMs utilize a movable mirror placed behind a linearly polarizing array of parallel wires to induce a varying phase delay between polarization states perpendicular to and parallel to the direction of the array wires. When combined with detectors sensitive to linear polarization, modulating the mirror position, and thus the distance between the mirror and the wire array, results in the modulation of one linear polarization state, instrument Stokes U in the case of CLASS, into circular polarization, Stokes V, and vice versa. This modulation increases polarization measurement stability by utilizing phase-sensitive detection and allows for the measurement of large angular scale modes on the sky. As Stokes U and V are modulated instead of Stokes Q and U as is the case of half-wave plate modulators more commonly used in CMB instruments (Kusaka et al. 2018), CLASS has significant capability to measure circular polarization (Harrington et al. 2021). Furthermore, its V detection capability is a more direct measurement than sensitivity obtained via half-wave plate non-idealities, which are highly frequency dependent in a poorly constrained manner (Nagy et al. 2017).

The remainder of this chapter is organized as follows. In Section 8.1, we present the theory behind polarized Zeeman emission of molecular oxygen in the atmosphere that is used to simulate the expected signal. Next, in Section 8.2, we present the results of these simulations for the CLASS observing site in the Atacama Desert of northern Chile. Then,

we compare the simulation results to data from the first era of observations of the CLASS Q-band receiver in Section 8.3. Finally, we conclude in Section 8.4.

8.1 Atmospheric emission theory

In its electronic ground state, molecular oxygen has a spin quantum number $S = 1$ due to the unpaired spins of two electrons, resulting in a magnetic dipole moment. This magnetic dipole moment results in transitions between rotational states of the molecule's electronic and vibrational ground states with millimeter-wave emission. The spin quantum number couples with the total rotational angular momentum quantum number N, which must be odd due to the exclusion principle, to yield the rotational quantum number J. This results in three possible values for J per N, $J = N, N \pm 1$. Selection rules allow for an N^+ transition, $(J = N) \rightarrow (J = N + 1)$, and an N^- transition, $(J = N) \rightarrow (J = N - 1)$. In the absence of an external magnetic field, this emission is unpolarized, but a non-zero external magnetic field induces Zeeman splitting, which produces polarized emission (Berestetskii et al. 1982). An external magnetic field splits the transition lines due to a given J into $2J + 1$ lines corresponding to the magnetic quantum number M, where $-J \leq M \leq J$; M expresses the projection of the molecular magnetic moment on the external magnetic field vector.

8.1.1 Layer attenuation

Molecular oxygen

In a coherency matrix formalism (Lenoir 1967), the attenuation of a given atmosphere layer due to resonances in molecular oxygen can be defined as

$$\boldsymbol{G}_Z(\nu) = \frac{1}{2} \sum_i S_i \sum_{\Delta M=-1}^{1} \boldsymbol{\rho} \sum_{M=-N}^{N} \xi(N, M) F(\nu, \nu_k), \tag{8.1}$$

where S_i ($\mathrm{GHz\,km^{-1}}$) is the intensity of the unsplit line, $\boldsymbol{\rho}$ is the transition matrix, ξ is the intensity of the Zeeman component relative to the unsplit line, F (GHz) is the line profile, ν (GHz) is the frequency of evaluation, and ν_k (GHz) is the frequency of the Zeeman line. Note that this is field attenuation, which is a factor of two smaller than power attenuation, in units of $\mathrm{Np\,km^{-1}}$. The outer summation is performed over the first thirty-eight resonance lines, starting at $N^{\pm} = 1$ and ending at $N^{\pm} = 37$. Higher-order resonance lines are excluded due to lack of available line broadening and mixing data; as these lines are much weaker, the effect of excluding them is minimal.

The intensity of the unsplit line ($\mathrm{GHz\,km^{-1}}$) is

$$S_i = S_{296} \frac{n p_{\mathrm{air}} P}{RT} \left(\frac{T_{296}}{T} \right)^{2.5} \exp \left[\frac{E''}{T_{296} k_B} \left(1 - \frac{T_{296}}{T} \right) \right], \tag{8.2}$$

where S_{296} ($\mathrm{MHz\,m^2\,mol^{-1}}$) is the fiducial intensity at 296 K, $n \approx 0.2095$ is the volume fraction of molecular oxygen in the atmosphere (Machta and Hughes 1970),[1] $p_{\mathrm{air}} = 1 - p_w$ is the fractional partial pressure of dry air, P (hPa) is the pressure of the atmosphere layer, R ($\mathrm{J\,mol^{-1}\,K^{-1}}$) is the molar gas constant, T (K) is the physical temperature of the atmosphere layer, $T_{296} = 296\,\mathrm{K}$ is a reference temperature, E'' (J) is the lower-state energy of the transition, and k_B is the Boltzmann constant ($\mathrm{J\,K^{-1}}$). S_{296} and E'' values are from HITRAN (Gordon et al. 2017). As per Liebe et al. (1993), the fractional partial pressure of water vapor is

$$p_w = 2.408 \cdot 10^{11} \frac{u}{P} \left(\frac{T_{300}}{T} \right)^{5} \exp \left[-22.644 \left(\frac{T_{300}}{T} \right) \right], \tag{8.3}$$

where u is the fractional relative humidity and $T_{300} = 300\,\mathrm{K}$ is a reference temperature.

[1] This decreases at very high altitudes due to photodissociation (Penndorf 1949), but, as will be shown later, the relevant signal is primarily from the lower atmosphere.

For $\Delta M = 0$,

$$\boldsymbol{\rho}_{\pi} = \begin{pmatrix} 0 & 0 \\ 0 & \sin^2\theta \end{pmatrix} = \frac{\sin^2\theta}{2}\boldsymbol{\sigma}_I - \frac{\sin^2\theta}{2}\boldsymbol{\sigma}_Q, \tag{8.4}$$

and for $\Delta M = \pm 1$,

$$\begin{aligned} \boldsymbol{\rho}_{\sigma\pm} &= \begin{pmatrix} 1 & \mp i\cos\theta \\ \pm i\cos\theta & \cos^2\theta \end{pmatrix} \\ &= \frac{1+\cos^2\theta}{2}\boldsymbol{\sigma}_I + \frac{1-\cos^2\theta}{2}\boldsymbol{\sigma}_Q \pm \cos\theta\boldsymbol{\sigma}_V, \end{aligned} \tag{8.5}$$

where θ is the angle between the line of sight and the geomagnetic field vector and

$$\boldsymbol{\sigma}_I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad \boldsymbol{\sigma}_Q = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \quad \boldsymbol{\sigma}_V = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}.$$

Thus, the amplitude of circular polarization is maximized when the line of sight is aligned with the direction of the magnetic field, and the linear polarization is maximized when the two vectors are perpendicular.

While the Zeeman effect in molecular oxygen can be approximated to reasonable accuracy by Hund's case (b) (Townes and Schawlow 1955), we use the more exact calculations laid out in Larsson et al. (2019). In this formalism, the frequencies of the Zeeman lines are given by $\nu_k = \nu_i + \delta_{\Delta M}$, where ν_i (GHz) [from HITRAN (Gordon et al. 2017)] is the frequency of the unsplit line and

$$\delta_{\Delta M} = -\frac{\mu_B B}{h}\left[g_{N=J}M + g_{N=J\pm1}(M+\Delta M)\right] \cdot 10^{-9} \tag{8.6}$$

is the Zeeman frequency shift in GHz. Here, B (T) is the magnetic field strength, μ_B (J T^{-1}) is the Bohr magneton, h (J s) is the Planck constant, and $g_{N=J}$ and $g_{N=J\pm1}$ are the numerical

Table 8.1: Relative intensity factor $\xi(N, M)$

	N^+ line	N^- line
$\Delta M = 0$	$\frac{3[(N+1)^2 - M^2]}{(N+1)(2N+1)(2N+3)}$	$\frac{3(N+1)(N^2 - M^2)}{N(2N+1)(2N^2+N-1)}$
$\Delta M = \pm 1$	$\frac{3(N \pm M + 1)(N \pm M + 2)}{4(N+1)(2N+1)(2N+3)}$	$\frac{3(N+1)(N \pm M)(N \pm M - 1)}{4N(2N+1)(2N^2+N-1)}$

Note—For Zeeman components of O_2 lines, where $\pm M \leq N$ (Liebe 1981).

Zeeman coefficients corresponding to the given $N^{\pm}$ transition from Table 2 of Larsson et al. (2019). The values of the relative intensity factor $\xi(N, M)$ are shown in Table 8.1; note that these normalize to two in the absence of a magnetic field, but when combined with $\boldsymbol{\rho}$, the combination yields the identity matrix.

Following Larsson et al. (2014) and Melsheimer et al. (2005) the line profile is defined as

$$F(\nu, \nu_k) = \left(\frac{\nu}{\nu_k}\right)^2 \frac{1}{\Delta\nu_D\sqrt{\pi}} \left[(1 + g_l P^2 - i y_l P) w(z_-) + (1 + g_l P^2 + i y_l P) w(z_+)\right], \tag{8.7}$$

where the Faddeeva function (for $\mathrm{Im}(z) > 0$) (Faddeyeva and Terent'ev 1961) is,

$$w(z) = \frac{i}{\pi} \int_{-\infty}^{\infty} \frac{e^{-t^2}}{z - t} dt, \tag{8.8}$$

with

$$z_{\pm} = \frac{\nu \pm \nu_k \pm \delta\nu_l P^2 + i\Delta\nu_p}{\Delta\nu_D}. \tag{8.9}$$

Here, g_l is the second-order line shape correction, y_l is the first-order phase correction, and $\delta\nu_l$ is the second-order frequency correction. These are defined with temperature dependence by

$$Z_l(T) = \left[Z_l^0 + Z_l^1 \left(\frac{T_{300}}{T} - 1\right)\right] \left(\frac{T_{300}}{T}\right)^{x_Z}, \tag{8.10}$$

where Z_l is the first- or second-order coefficient y_l, g_l, or $\delta\nu_l$ and $y_l^{0,1}$ (hPa^{-1}), $g_l^{0,1}$ (hPa^{-2}),

and $\delta\nu_l^{0,1}$ (GHz hPa^{-2}) are all from Table 1 of Makarov et al. (2011); the exponent x_Z is 0.8 for y_l and 1.6 for g_l and $\delta\nu_l$. The pressure broadening half width is

$$\Delta\nu_p = \gamma_{\text{air}} p_{\text{air}} P \left(\frac{T_{296}}{T}\right)^{x_\nu} + \gamma_w p_w P \left(\frac{T_{296}}{T}\right)^{x_\nu} \tag{8.11}$$

for

$$\gamma_{\text{air},w}(N) = A_\gamma + \frac{B_\gamma}{1 + c_1 N + c_2 N^2 + c_3 N^4}, \tag{8.12}$$

where A_γ (GHz hPa^{-1}), B_γ (GHz hPa^{-1}), c_1, c_2, and c_3 are parameters from column four of Table 3 of Koshelev et al. (2016) for γ_{air} and from column two of Table 2 of Koshelev et al. (2015) for γ_w; the exponent $x_\nu = 0.75412$ is from Table 1 of Koshelev et al. (2016). Also used is the $1/e$ Doppler half width (Herbert 1974; Varghese and Hanson 1984),

$$\Delta\nu_D = \frac{\nu_k}{c}\sqrt{\frac{2k_B T}{M_{O_2}}} \tag{8.13}$$

in GHz, where c (m s^{-1}) is the speed of light and M_{O_2} (kg mol^{-1}) is the molar mass of molecular oxygen [from HITRAN (Gordon et al. 2017)]. The $1/(\Delta\nu_D\sqrt{\pi})$ factor is a normalization (Armstrong 1967). At high pressures, the Faddeeva function simplifies to a Lorentz shape function, so the combination of the two Faddeeva functions reduces to a Van Vleck–Weisskopf (Van Vleck and Weisskopf 1945) line shape. At low pressures, when $P \to 0$, the line mixing and pressure correction effects are eliminated, so the function behaves similar to a Voigt line shape function, although with slightly reduced amplitude on the line wings. As will later be shown, the high-pressure case is more important, so the line shape function was chosen such that it is most accurate in that regime.

Dry air and water vapor

While the Zeeman effect is the sole significant non-transient source of polarized atmospheric emission in the millimeter spectrum, there are other unpolarized sources, quantified by a dry air continuum and water vapor contributions. Following Tretyakov (2016), the dry air and water vapor non-resonant continua term is

$$\alpha_c(\nu, T) = \left[C_w^0 \left(\frac{T_{300}}{T} \right)^{x_w} p_w^2 + C_{\text{air}}^0 \left(\frac{T_{300}}{T} \right)^{x_{\text{air}}} p_{\text{air}} p_w + C_{\text{dry}}^0 \left(\frac{T_{300}}{T} \right)^{x_{\text{dry}}} p_{\text{air}}^2 \right] P^2 \nu^2, \tag{8.14}$$

where C_w^0, C_{air}^0, C_{dry}^0 ($\text{Np}\,\text{km}^{-1}\,\text{hPa}^{-2}\,\text{GHz}^{-2}$) and x_w, x_{air}, x_{dry} are numerical coefficients from Table 5 of Tretyakov (2016). Following Rosenkranz (1998), the water vapor contribution is

$$\alpha_w(\nu, T) = \sum_j S_j(T) \left[f_j(\nu) + f_j(-\nu) \right] \tag{8.15}$$

summed over the water vapor lines shown in Table 1 of Rosenkranz (1998), where the line profile is defined as

$$f_j(\nu) = \begin{cases} \dfrac{\nu^2 \gamma_j}{\pi \nu_j^2} \left[\dfrac{1}{\left(\nu - \nu_j\right)^2 + \gamma_j^2} - \dfrac{1}{\nu_c^2 + \gamma_j^2} \right] & |\upsilon - \upsilon_j| < \upsilon_c \\ 0 & |\upsilon - \upsilon_j| \geq \upsilon_c \end{cases} \tag{8.16}$$

with

$$\gamma_i = w_s p_w P \left(\frac{T_{300}}{T} \right)^{x_s} + w_f p_{\text{air}} P \left(\frac{T_{300}}{T} \right)^{x_f}, \tag{8.17}$$

where $\nu_c = 750\,\text{GHz}$ is a cutoff frequency and w_s, w_f ($\text{GHz}\,\text{km}^{-1}$) and x_s, x_f are numerical coefficients from Table 1 of Rosenkranz (1998). This formulation resembles the Van Vleck–Weisskopf line shape but includes a high-frequency cutoff, as proposed by Clough et al. (1989). The line intensities, S_j ($\text{GHz}\,\text{km}^{-1}$), follow Equation (8.2) but with np_{air} replaced with p_w; S_{296} ($\text{MHz}\,\text{m}^2\,\text{mol}^{-1}$) and E'' (J) values are again from HITRAN (Gordon et al. 2017),

along with the line frequencies, ν_j (GHz). The full coherency matrix (Np km^{-1}) is then

$$\boldsymbol{G}(\nu) = \boldsymbol{G}_Z(\nu) + \frac{\alpha_c}{2}\boldsymbol{I} + \frac{\alpha_w}{2}\boldsymbol{I}, \tag{8.18}$$

where $\boldsymbol{I}$ is the identity matrix.

8.1.2 Radiative transfer

As we are interested in polarization, a tensor radiative transfer approach is required to model how radiation propagates through the atmosphere. Using a plane-parallel approximation, the atmosphere is divided into layers of thickness $\Delta z = \sec\phi \cdot 0.2\,\text{km}$ starting from ground level, 5.2 km in the case of CLASS, and ending at 100 km; ϕ denotes the zenith angle. Using the approach described in Lenoir (1967) and Lenoir (1968) and including phase as per Rosenkranz and Staelin (1988), the brightness temperature coherency matrix (Randa et al. 2008) of a given atmosphere layer is defined as

$$\boldsymbol{T}_B(z) = e^{-\boldsymbol{G}\Delta z}\boldsymbol{T}_B(z_0)e^{-\boldsymbol{G}^\dagger\Delta z} + T(z)\left(\boldsymbol{I} - e^{-\boldsymbol{G}\Delta z}e^{-\boldsymbol{G}^\dagger\Delta z}\right) \tag{8.19}$$

using matrix exponentiation, where $\boldsymbol{T}_B(z_0)$ is the brightness temperature coherency matrix of the atmosphere layer before it, $T(z)$ (K) is the physical temperature of the atmosphere layer, and † represents the conjugate transpose operation. Since we are observing from the ground, unlike the satellite observations described in Lenoir (1968), we start at 100 km with $\boldsymbol{T}_B(z_0) = T^{\text{RJ}}_{\text{CMB}}\boldsymbol{I}$ and calculate the propagation downward layer-by-layer until the ground, where $T^{\text{RJ}}_{\text{CMB}}$ is the CMB monopole (Fixsen 2009) brightness temperature at the observing frequency. Note that as the brightness temperature coherency matrix contains matrix elements defined in terms of the brightness temperature of a single polarization

of radiation,[2] there is an extra factor of one half in the conversion to Stokes parameters when compared to a standard coherency matrix (Born and Wolf 1959), e.g., $I = \frac{1}{2}\,\mathrm{Tr}(\boldsymbol{T}_B)$.

8.1.3 Atmosphere and magnetic field properties

In order to perform the aforementioned calculations, data are needed for the atmospheric temperature and pressure profiles, as well as the geomagnetic field vector direction and magnitude at the location in question. The NRLMSISE-00 atmosphere model (Picone et al. 2002), averaged over the full year, was used to calculate the temperature and pressure profiles used; for the lower atmosphere, this model was compared to data acquired via radiosonde launches from the Chajnantor Plateau, which is adjacent to the CLASS observing site, during the ALMA site characterization campaign[3] and was found to be in good agreement. The magnetic field was calculated at ground level using the National Oceanic and Atmospheric Administration's Enhanced Magnetic Model 2017 (EMM2017) (Chulliat et al. 2015). While the magnetic field strength does decrease with altitude, this effect is minor, so it was omitted from the calculations.

8.1.4 Primary source of polarized emission

Zeeman emission from atmospheric molecular oxygen is generally considered to be a mesospheric effect, since individual emission lines can be discerned at these altitudes due to the lack of pressure broadening. However, we are interested in polarized emission at frequencies far from the resonance lines, where the discernibility of individual emission lines is not relevant. To determine which altitude region of the atmosphere is primarily responsible for the polarized emission, we simulated a series of observations to the north

[2] See the footnote of Section 1.2.2 of Janssen (1993) for a discussion of differing brightness temperature definitions.

[3] Data retrieved from `http://legacy.nrao.edu/alma/site/Chajnantor/instruments/radiosonde/`.

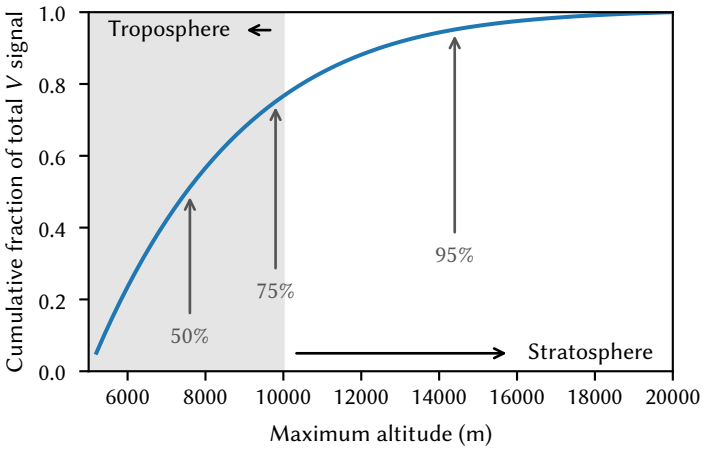

Figure 8.1: Altitude dependence of Zeeman emission, comparing the total Stokes V signal observed for various simulation altitude cutoffs to the fiducial case of 100 km maximum altitude. The simulated observations are to the north with a 45° zenith angle, starting from 5200 m at the CLASS observing site for the CLASS Q-band telescope. The light gray background denotes the troposphere, while the white background denotes the stratosphere.

with a 45° zenith angle from the CLASS observing site in the frequency band of the CLASS Q-band telescope; the properties of both the site and telescope are further detailed in the next section. Instead of starting at 100 km altitude to calculate the layer-by-layer propagation downward, the starting altitude is reduced incrementally and compared to the 100 km fiducial case, the results of which are shown in Figure 8.1. This series of calculations shows that the lower atmosphere is the primary contributor to the polarized Zeeman emission far from the resonance lines, with three-quarters of the signal contributed by the troposphere. At lower altitudes, the atmospheric pressure is higher, so there are more oxygen molecules per unit volume, which leads to more emission. Thus, while the mesospheric emission is most significant when one wishes to resolve individual Zeeman emission lines, tropospheric emission is most significant far from the resonance lines.

8.2 Simulation results

The CLASS observing site is located at 5.2 km elevation in the Atacama Desert of northern Chile, with coordinates 22.959 75° S, 67.787 26° W. For this location, at ground level, the EMM2017 magnetic field model was evaluated for 2017 January 1, a date near the middle of the observing period that will be described in Section 8.3; this resulted in a magnetic field of 22 738 nT oriented with an azimuth angle of −5.9° and a zenith angle of 68.8°. As previously mentioned, the NRLMSISE-00 atmosphere model was used to calculate temperature and pressure profiles. To include water vapor effects, 10% relative humidity is assumed, which is a typical value for the CLASS observing site during good weather; this corresponds to ∼0.6 mm of precipitable water vapor (PWV). As the observations are made away from the water vapor resonance lines, the relative humidity only affects the polarized signal via effecting small changes to the partial pressure of molecular oxygen. Thus, water vapor effects are small for the frequencies of interest, so the exact value is not critical.

Using the atmosphere temperature and pressure profiles described above, a full radiative transfer simulation can be performed, as described by Equation (8.19), using coherency matrices described by Equation (8.18). It is informative to first consider the frequency dependence of the polarized emission, which is shown in Figure 8.2. The primary features are the cluster of resonance lines in the 50 GHz to 70 GHz range and the 1^- line at 118.8 GHz; since the sign of the circular polarization is reversed when transitioning from frequencies below the resonance frequency of a given line and above it, there are nulls due to the interactions between the 1^- line and the other resonance lines. Furthermore, the polarized signal is attenuated due to water vapor at 183 GHz and to a lesser extent at 22 GHz.

Next, effects from 32.3 GHz to 43.7 GHz, the frequency band of the CLASS Q-band telescope (Appel et al. 2019), are considered. A sky plot is shown in Figure 8.3, which shows both the azimuth and zenith angle dependence of the circularly polarized atmospheric

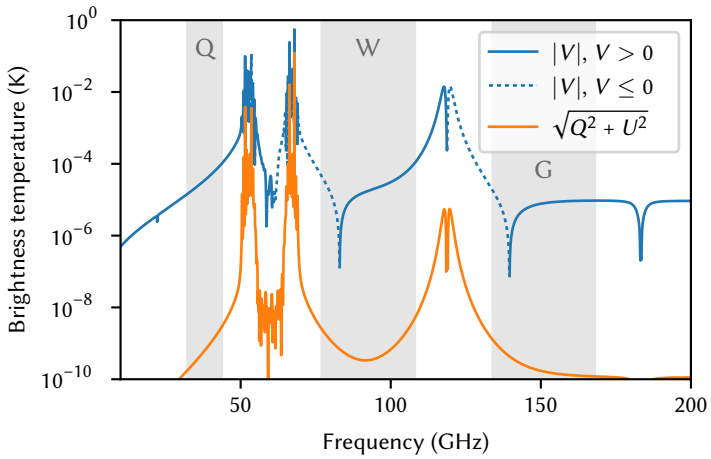

Figure 8.2: Frequency dependence (ν) of polarized atmospheric signal at zenith for the CLASS observing site, both for circular polarization ($|V|$), shown in blue) and linear polarization ($\sqrt{Q^2 + U^2}$), shown in orange). The light gray bands indicate CLASS observing frequencies, with the lowest frequency band corresponding to the Q-band telescope.

signal. A detailed view of the azimuth dependence is shown in Figure 8.4. The magnitude of the circularly polarized signal is strongest when the azimuth is aligned with the magnetic declination angle, as expected, and is also stronger in the north than the south, since the line-of-sight is better aligned with the magnetic field in that direction, as the magnetic field vector points above the horizon to the north. Since a larger air mass is observed closer to the horizon, the signal is also stronger further from zenith. The effect of the 1.5° full width at half maximum beams of the CLASS Q-band telescope on the zenith angle and azimuth dependence was evaluated and found to be negligible, so the effect was not considered in the remainder of the analysis.

For rapid evaluation of the atmosphere model, a function of the form

$$V = a \cdot \tan(b \cdot \phi) \cdot \cos(\psi - c) + d \tag{8.20}$$

can be precisely fit to the simulations, where ϕ is the zenith angle and ψ is the azimuth angle, with fit parameters shown in Table 8.2; zenith angles from 30° to 60° were used in

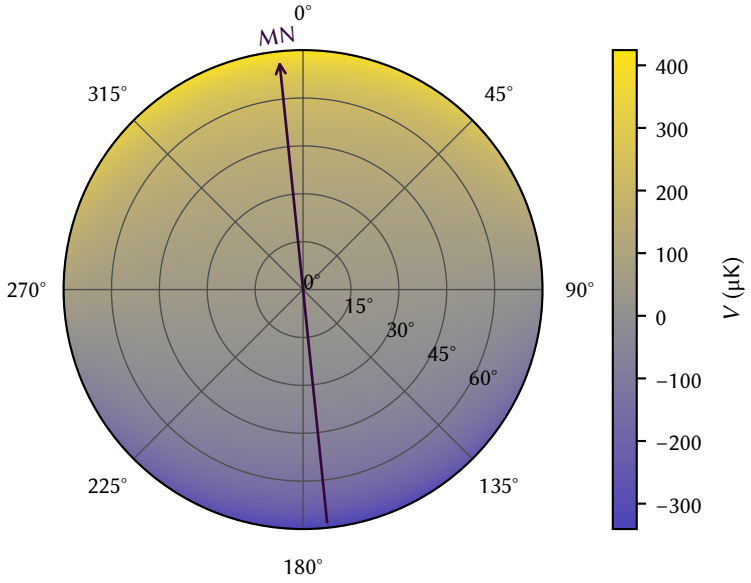

Figure 8.3: Simulated azimuth and zenith angle dependence of the atmospheric Stokes V signal at the CLASS observing site for the CLASS Q-band telescope. Azimuth is shown for a full 360°, and zenith angle is shown for 0° to 75°. The dark purple arrow indicates magnetic north.

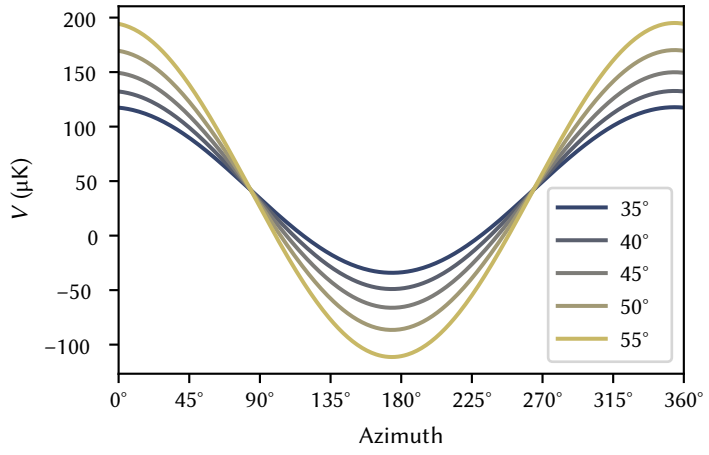

Figure 8.4: Simulated azimuth profiles of the atmospheric V signal at the CLASS observing site for CLASS Q-band telescope are shown for different zenith angles.

Table 8.2: Atmosphere simulation fit parameters at the CLASS observing site for the CLASS Q-band telescope

Parameter	Value
a	1.106×10^{-4} K
b	9.848×10^{-1}
c	$-5.9°$
d	4.185×10^{-5} K

NOTE—When using b with Equation (8.20), ϕ should be in radians.

the fit. With a and d specified in K, V is also in K; b is a dimensionless scale factor, and ψ and c are angles. The average fit residual is ∼200 nK. Parameter c, the azimuth offset, corresponds exactly to the magnetic declination. The simulation code used to produce these results has been published (Petroff 2019).

8.3 Comparison with observations

From the CLASS Era 1 survey, nighttime data recorded during the period from 2016 September through 2018 February are used for the present analysis; some nights are excluded due to particularly poor weather conditions or due to operational difficulties, such as interruptions to power or cryogenic systems. For each nightly observation, the telescope was scanned with the boresight center pointing at a 45° zenith angle through an azimuthal range of ±360° at a rate of $1\,^{\circ}\,\mathrm{s}^{-1}$. For the entirety of each night, the boresight rotation angle of the telescope relative to the horizon remained fixed at −45°, −30°, −15°, 0°, +15°, +30°, or +45°; the boresight rotation angle was changed daily such that each angle was observed on a weekly basis. Boresight rotation combined with individual feedhorn pointing offsets provides access to a range of zenith angles from 35° to 55° with this scanning strategy.

Of the 28 detector pairs in the Q-band telescope that were operational during the Era 1

survey (Appel et al. 2019), 25 are used for the present analysis, with the remaining three conservatively rejected due to atypical noise properties. Anomalous artifacts found in detector timestreams are excised, along with windows surrounding them chosen such that any filtering operations in the timestream processing pipeline do not convolve the artifacts with surrounding data. The Stokes V signal is extracted from pair-differenced detector timestreams via demodulation with the VPM polarization transfer function (Chuss et al. 2012; Harrington et al. 2021), where each pair corresponds to the two detectors with orthogonal linear polarization sensitivity in each feedhorn. These resulting V timestreams are further checked for stable noise properties, with variance cuts used to eliminate data with abnormally high noise. This data processing pipeline will be described in detail by J. Eimer et al. (2021, in preparation). After all data cuts, 47927.6 detector pair-hours of data remain. Daytime data were excluded from the present analysis, as additional pipeline developments are required to properly handle their reduced stability and artifacts; however, this does not preclude future use of these data. Furthermore, sun avoidance maneuvers alter the telescope's scan strategy during parts of the day and prevent coverage of the full azimuth range, making these data less suitable for the present analysis.

To confirm that an observed circularly polarized signal is due to atmospheric Zeeman emission, three properties should be satisfied: the azimuth angles of maximum and minimum signal should align with the magnetic declination, the signal should show an appropriate zenith angle dependence, and the signal should have approximately the correct amplitude. To evaluate the observed signal, detector data were divided by detector pair and boresight angle, processed as previously described, binned by azimuth with inverse variance weighting, and fit with a sinusoidal profile; variance was evaluated for timestream segments after splitting the demodulated timestreams into sweeps of constant direction scanning in azimuth. As ground pickup contamination, likely due to $T \rightarrow V$ leakage, was visible in binned azimuth profiles for some combinations of detector pairs and boresight

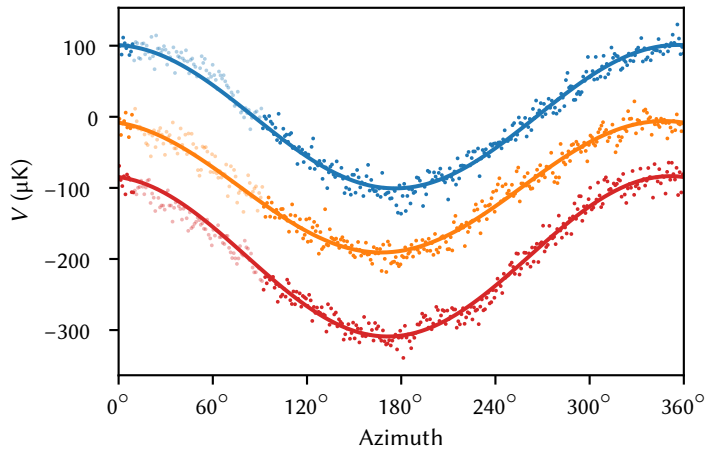

Figure 8.5: Example binned azimuth profiles are shown along with sinusoidal best fit lines for three combinations of detector pairs and boresight rotation angles. The data are plotted with arbitrary amplitude offsets; error bars are not shown. The lighter data points indicate the region excluded from the sinusoidal fits due to the ground elevation angle cut. The profile in blue is from a zenith angle of 43.9° and a boresight rotation angle of −45°, the profile in orange is from a zenith angle of 46.7° and a boresight rotation angle of 0°, and the profile in red is from a zenith angle of 52.8° and a boresight rotation angle of +45°.

rotation angles, the azimuth range where the ground elevation angle is above 6°, i.e., where the peak of Cerro Toco is, was masked before fitting; this corresponds to 10° to 92° azimuth. A covariance threshold was also used to remove poor fits. As each azimuth sweep is mean-subtracted, only the peak-to-peak amplitude of the azimuth-dependent signal is measured, not its absolute offset. Example binned azimuth profiles and sinusoidal fits are shown in Figure 8.5.

The fits result in a signal maximum aligned with an azimuth angle of (−5.5 ± 0.6)°, which is consistent with the expected magnetic declination, the angle between magnetic north and true north, of −5.9°; this error was calculated using bootstrapping after performing an additional data cut based on the residual after subtracting the sinusoidal best fits from the binned azimuth profiles. The other mechanism by which the Earth's magnetic field is expected to affect the detector data is through magnetic pickup in the detectors and Superconducting Quantum Interference Device (SQUID) multiplexers and amplifiers.

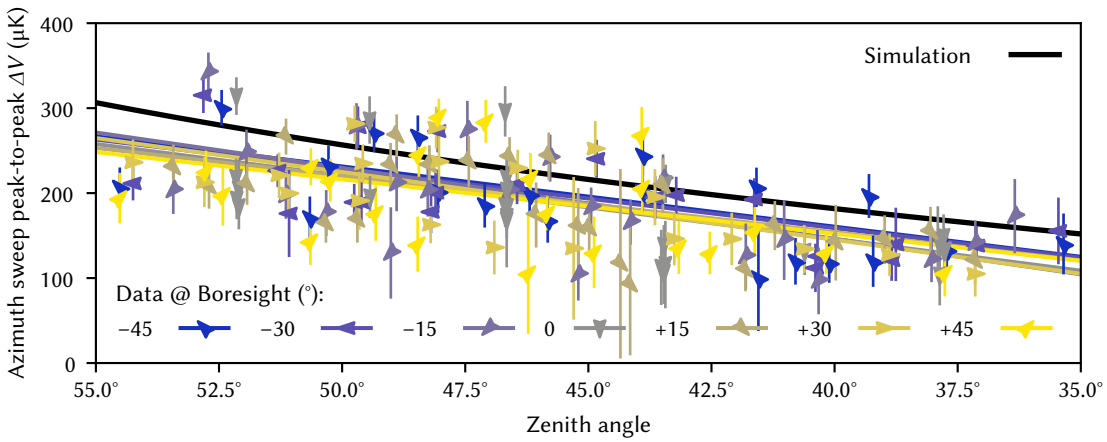

Figure 8.6: Zenith angle dependence of peak-to-peak observed ΔV signal (shown in color) compared to simulation (shown in black). The data are split by both feedhorn and boresight rotation angle, with an individual point plotted for each combination. Error bars show one standard deviation calculated on the residual after subtracting the sinusoidal best fits from the binned azimuth profiles. Linear best fit lines are also shown for each boresight rotation angle; the reduced χ^2 for these fits range from 1.8 to 3.4, with an average reduced χ^2 of 2.5.

As this pickup is not modulated by the VPM and should therefore not be visible in the demodulated timestreams, this alone is strong evidence of a detection of a signal of atmospheric origin. Next, the zenith angle dependence of the peak-to-peak signal in the data was evaluated and compared to the simulation, the results of which are shown in Figure 8.6. A data point is shown for each combination of detector pairs and boresight rotation angles; zenith angle differences are due to individual feedhorn pointing offsets, which change relative to the boresight center pointing at different boresight rotation angles. An expected zenith angle dependence is seen, giving further evidence that the observed signal is of atmospheric origin. The excess scatter is thought to be due to uncertainties in the preliminary detector relative efficiency calibrations used in the analysis combined with potential bandpass mismatches between detectors.

As can be seen in Figure 8.6, the measured amplitude is consistent between different boresight rotation angles and between different detector pairs, providing a check against systematic errors. Additionally, data were split by date of observation to check for changes

over time; these splits were also found to be consistent. While there are expected to be slight changes over time to the atmospheric signal due to evolution of the geomagnetic field, these changes are much smaller than the error of the measurements; EMM2017 predicts the yearly change in field strength to be on the sub-percent level and the yearly change in magnetic declination to be approximately -0.2°.

A ground loop between the VPM control electronics and the detector readout electronics would in theory be able to introduce modulated signal into the SQUID magnetic pickup. However, detector data collected while the VPM was running and the receiver window was covered with a metal plate to block optical signals was inspected and found to contain no modulated signal, ruling out this possibility.

To further exclude the possibility of the detected signal being due to magnetic pickup, the physical geometry of the telescope's receiver can be considered. The detectors and SQUIDs are all in the same plane relative to one another despite having different optical pointing offsets, i.e., the detectors all have the same physical orientation relative to the magnetic field. Thus, the peak-to-peak amplitude of their magnetic pickup—which should not be in the demodulated timestreams to start with—would only be dependent on the zenith angle of the telescope's boresight pointing. Since all observations were taken with a boresight pointing zenith angle of 45°, SQUID magnetic pickup would not result in a zenith angle dependence, contrary to the detected signal. Additionally, changing the boresight rotation angle alters the physical orientation of the SQUIDs relative to the magnetic field and thus yields a boresight rotation angle dependence to the magnetic pickup, again contrary to the detected signal.

Using Equation (8.20) and the fit parameters in Table 8.2, a simulated map was created using pointings from the aforementioned observations, which is shown in Figure 8.7 along with a map of CLASS V data and a map showing residuals. To create the data map, the V signal was accumulated into pixels using inverse variance weighting of data, using the

same data processing and variance calculation procedures used for fitting the sinusoidal profiles. For the simulated map, azimuth sweeps were mean subtracted and weighted to match the processing of the data. Note that the overall dipole pattern is visible in both maps; the dipole amplitude is 162 μK in the simulated map and 130 μK in the data map, with a direction difference of 0.7°. The dipole amplitude ratio between the simulated map and the data map was used to scale the simulated map before subtracting it from the data map to produce the residual map. If the azimuth sweeps are not mean subtracted when creating the simulated map, the dipole changes in amplitude by +6 μK and in direction by 0.4°, demonstrating that the mean subtraction only has a minor effect on the signal. Areas of higher noise are due to uneven sky coverage.

The measured amplitude in both the sinusoidal profiles and map are in reasonable agreement with that predicted by the simulations. The remaining ∼20% discrepancy is likely due to some combination of calibration errors in the detector data and shortcomings of the simulation. Although the detector calibration was done using Stokes I, the demodulated linear polarization of the Crab Nebula matches previous observations (Xu et al. 2020), so it is unlikely that a discrepancy of the magnitude observed between observation and simulation is due to an error in either the VPM polarization transfer function calibration or detector calibration. The detector bandpass uncertainty is also low enough that a bandpass error cannot fully explain the discrepancy (Appel et al. 2019). This leaves a shortcoming of the simulation as the most likely source of the discrepancy. The Van Vleck–Weisskopf line shape is known to be inaccurate far from the resonance lines (Hill 1986), so it may be reasonable to attribute the difference to this inaccuracy, since we are observing at frequencies far from the resonance lines, and the majority of the observed signal is from the lower atmosphere, where the line profile used in the simulations reduces to a Van Vleck–Weisskopf line shape. Furthermore, the line mixing and pressure broadening parameters used in the simulation were also all measured near the resonance lines. As

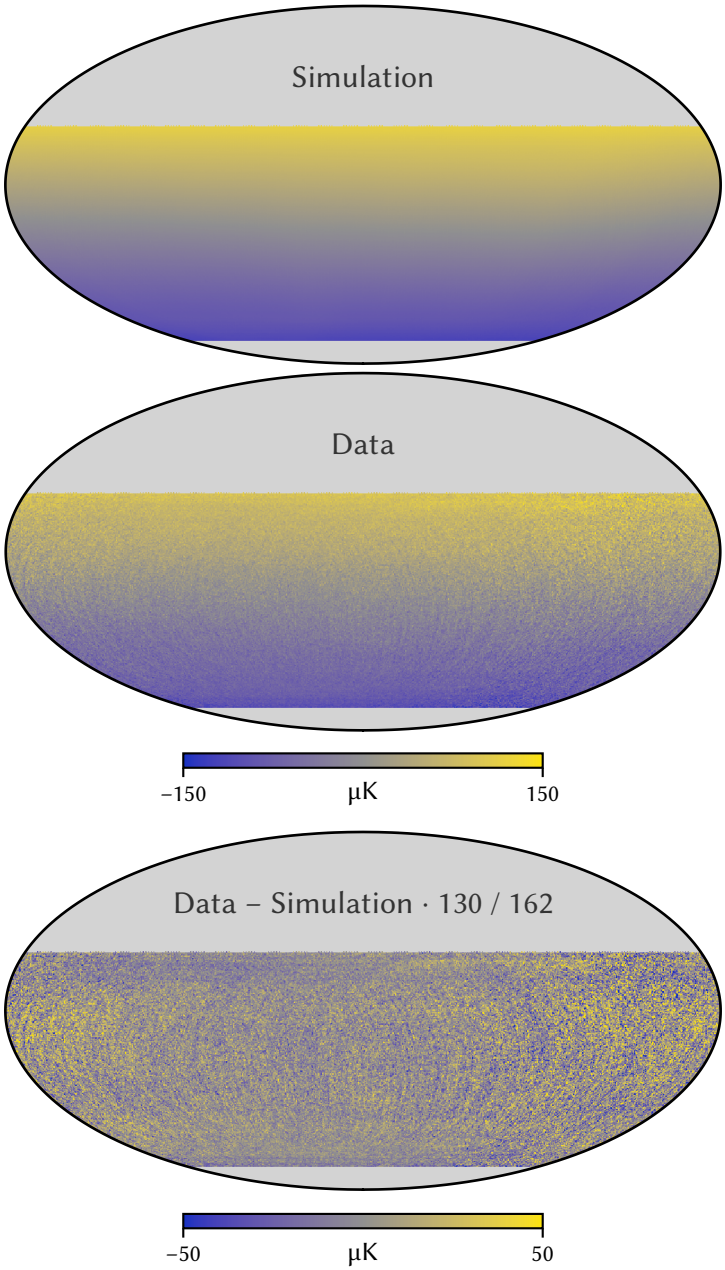

Figure 8.7: Stokes *V* signal mapped onto the projected celestial sky using pointings from CLASS Q-band telescope observations. The top plot uses a simulated atmospheric signal, the middle plot uses data, and the bottom plot shows the difference between the two maps, with the simulation data scaled by the amplitude difference of the dipoles fitted to the two previous maps. The color scale is identical for the top two maps and different for the bottom map. Note that this data map was produced with a different analysis pipeline than the maps presented in Padilla et al. (2020).

Zeeman splitting of molecular oxygen resonance lines in the presence of Earth's magnetic field is the only theorized source of circularly polarized atmospheric emission, we conclude that this is the source of the detected signal.

8.4 Conclusion

Expanding on prior models and utilizing recent spectroscopic data, a model for circularly polarized atmospheric emission from Zeeman splitting of molecular oxygen was presented. This model was then used to simulate the atmosphere of the CLASS observing site at the frequencies of the CLASS Q-band telescope. An analysis of circular polarization timestreams observed by the CLASS Q-band telescope, utilizing the Stokes V measurement capability of a VPM, was then compared to the simulations and shown to be a strong detection of this atmospheric emission; this is believed to be the first such detection of circular polarization in a frequency band used for ground-based CMB observations. The amplitude of the signal differed by ~20% between the simulations and observations but is still in good agreement. As the atmospheric signal is orders of magnitude larger than any theoretical cosmological or astrophysical signal, its subtraction is necessary for setting improved upper limits on said signals and providing more rigorous observational tests. Although this signal can be adequately removed empirically, such as by fitting a dipole sky pattern or by creating templates that describe the signal as a function of azimuth, modeling the signal helps us verify our understanding of it. After the subtraction of the atmospheric signal, Stokes V serves primarily as a null channel for CLASS; as no residual signal is expected, a VV angular power spectrum that is not consistent with zero would suggest the presence of either an unmitigated systematic error or a non-standard cosmological signal.

Chapter 9

Full-sky Cosmic Microwave Background Foreground Cleaning Using Machine Learning

The cosmic microwave background (CMB), a relic of the early Universe, encodes information about the physics of the first moments of the Universe. This wealth of information has proven crucial to modern precision cosmology. However, we can only observe the CMB through the local contamination of the Milky Way and other galaxies in the intervening Universe, which has fluctuations that are many times brighter than the anisotropy of the primordial signal (Planck Collaboration I 2020). Fortunately, the spectral indices of the foregrounds, both Galactic and extragalactic, differ from that of the CMB, allowing for the signals to be discriminated between through observations made at multiple frequencies. The most straightforward method to do so is to perform a simple internal linear combination of sky maps observed at different frequencies, fitting templates to the various components based on their spectral indices (Bennett et al. 2003). Numerous other more

Except for some minor changes, this chapter was previous published as Petroff et al. (2020a).

advanced foreground cleaning techniques have also been developed in an attempt to more completely remove foregrounds and reduce systematic errors (e.g., Martínez-González et al. 2003; Hansen et al. 2006; Bobin et al. 2007; Cardoso et al. 2008; Eriksen et al. 2008; Katayama and Komatsu 2011; Remazeilles et al. 2011; Basak and Delabrouille 2012; Ichiki et al. 2019). Reduction of systematic errors is crucial as such errors bias the cosmological parameters extracted from the CMB map.

Machine learning, specifically deep learning using artificial neural networks, has been increasingly applied to astrophysical problems (see reviews in Ball and Brunner 2010; Carleo et al. 2019; Fluke and Jacobs 2019), including to the CMB (e.g., Auld et al. 2008; Ciuca and Hernández 2017; He et al. 2018; Caldeira et al. 2019; Münchmeyer and Smith 2019; Puglisi and Bai 2020; Vafaei Sadr and Farsian 2020; Yi et al. 2020). Of particular note are previous applications of neural networks to CMB foreground cleaning, but these approaches have either not taken into account morphological features (Nørgaard-Nielsen 2010) or have not been applied to the full sky (Aylor et al. 2021). An artificial neural network is a neuroscience-inspired construct that consists of a set of units connected with trainable weights and nonlinear transformations, which can be used to approximate arbitrary functions (Russell and Norvig 2010). Convolutional neural networks (CNNs) are a type of artificial neural network particularly well suited for cleaning foregrounds from maps, as they work on images and process structural information instead of only individual pixels (Krizhevsky et al. 2012). However, traditional CNNs are limited to rectangular images and thus are not well suited to full-sky maps of the CMB, although they can be applied to flat sky projections of limited parts of the sky. Recently, multiple techniques have been developed to extend CNNs to the sphere, allowing for the full sky to be considered at once (Cohen et al. 2018; Esteves et al. 2018; Kondor et al. 2018; Jiang et al. 2019; Krachmalnicoff and Tomasi 2019; Perraudin et al. 2019). We apply one such technique to cleaning foregrounds from full-sky temperature maps of the millimeter and submillimeter sky, leaving the primordial CMB

signal, as well as possible residuals. We also assess its ability to evaluate how well different simulations match observations.

The remainder of this chapter is organized as follows. First, in Section 9.1, we discuss the artificial neural network architecture used. Next, in Section 9.2, we discuss the simulations used to train the model. We then describe the training procedure in Section 9.3 and apply the trained model to the simulations. The trained model is then applied to Planck sky maps in Section 9.4. Finally, we conclude in Section 9.5.

9.1 Neural network architecture

In concordance with the cosmological principle, a neural network architecture for cosmological applications should ideally encode neither position nor direction and should encompass the entire sky. The *DeepSphere* architecture of Perraudin et al. (2019) fulfills these properties; it is a convolutional architecture, so it does not encode position, and it learns radially symmetric kernels, so it is approximately rotationally equivariant,[1] all while covering the full sky. This architecture represents the Healpix equal-area sphere pixelization (Gorski et al. 2005) as a graph. Convolutions are performed as operations on the graph using Chebyshev polynomials, while downscaling takes advantage of Healpix's hierarchical nature, combining four pixels into their single parent pixel. However, this is not a complete solution, since the technique, as presented by Perraudin et al. (2019), only produces outputs suitable for regression or classification problems.

CMB foreground cleaning falls under the category of image-to-image problems in the machine-learning parlance. Thus, we extend the widely used *U-Net* architecture of Ronneberger et al. (2015) to the sphere. This technique consists of both contracting and expanding paths, which decrease and increase the image resolution, respectively, in order

[1]The improved weighting scheme presented in Defferrard et al. (2020) was used.

to encode features on multiple scales; *skip connections* are used to bypass smaller levels of the contracting–expanding path to transfer details and bypass the information bottleneck of the smallest level. On the contracting path, after a convolution is applied, a pooling operation is used to combine a block of four pixels into a single pixel; in this case, we use *max pooling*, which takes the maximum value of the four input pixels. On the expanding path, each pixel is subdivided into four identical child pixels, and a convolution is then performed on the higher resolution image. Each convolution layer uses a set of trainable filters and is generally followed by a nonlinear activation layer.

To extend this architecture to the Healpix sphere, we use Healpix's hierarchical nature to either combine four pixels into a single pixel or to subdivide a pixel into four child pixels, depending on whether the resolution is being decreased or increased. For the contracting path, we apply a DeepSphere graph-based convolution before pooling, while for the expanding path, we apply a DeepSphere graph-based convolution after pixels are duplicated to produce a higher resolution map.

Finally, a method of combining features from separate frequency maps is needed. We accomplish this by using a separate contracting path for each of the input frequency maps, with trainable multiplicative scale and additive bias weights for each frequency at each detail level. While a single contracting path that takes all of the frequency channels as input could have been used with a greater number of convolution filters, this would have resulted in a larger model with a greater number of trainable parameters. Instead, we exploit a priori knowledge of astrophysical foregrounds and construct an architecture that allows the model to learn resolution-dependent internal linear combinations in addition to morphological features. On the single expanding path, the scaled and biased frequency maps are summed with the upscaled, convolved map from the smaller level in the path. After this step, a nonlinear activation is applied to the summed map, and it is upscaled and convolved. The process is repeated for the more detailed result. A schematic overview of

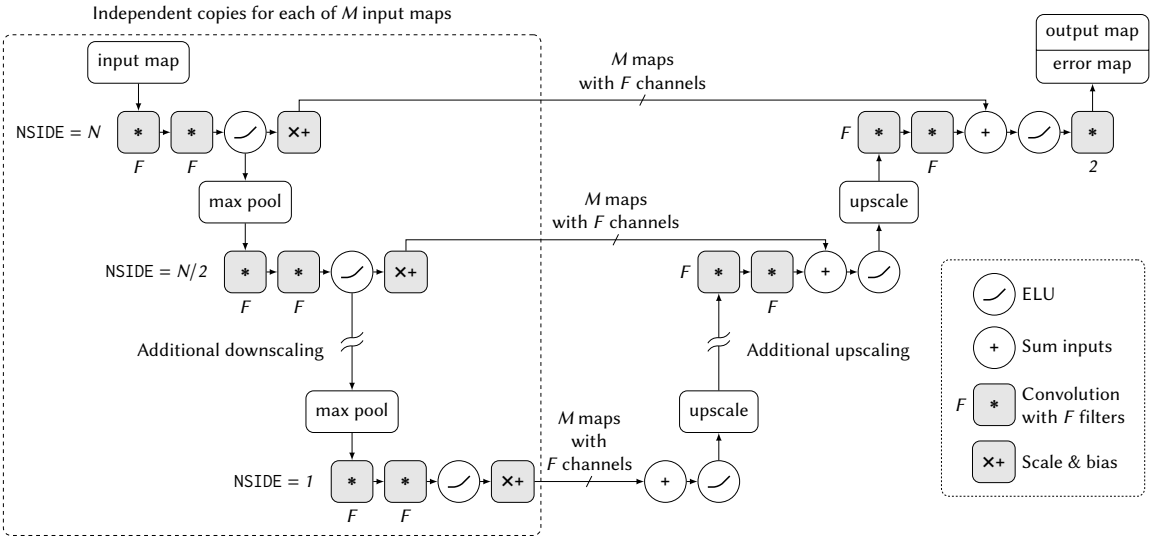

Figure 9.1: Artificial neural network architecture overview. Shading denotes nodes with trainable parameters.

this architecture is shown in Figure 9.1. All nonlinear activations are performed using the exponential linear unit (ELU) operation (Clevert et al. 2016). The number and sizes of the convolution filters will be discussed in Section 9.3. Each location where a convolution is performed consists of a pair of convolution layers, since this was found to reduce the loss function (defined below) more than a single convolution layer with twice the number of filters (but the same total number of filters).

For machine learning to be maximally useful in scientific contexts, accurate error estimation is crucial. This estimated error can be divided into aleatoric uncertainty, uncertainty due to noise in the observations, i.e., statistical uncertainty, and epistemic uncertainty, uncertainty due to error in the model, i.e., model-based systematic uncertainty (Kendall and Gal 2017). To this end, we used a Bayesian neural network with *Concrete dropout* (Gal et al. 2017). *Dropout* is a process that randomly drops units, and their connections, from an artificial neural network during training, which was developed to prevent overfitting during the training process (Srivastava et al. 2014). Dropout was later shown to be a Bayesian approximation that can be used to estimate a network's epistemic uncertainty by

also randomly dropping units over multiple evaluations, but the dropout probability had to be tuned to calibrate the uncertainty estimate (Gal and Ghahramani 2016); Concrete dropout allows for automatic tuning of the dropout probability. Following Kendall and Gal (2017), we have the network estimate aleatoric uncertainty by producing an error map in addition to a CMB map. This is accomplished by training the network to minimize the loss function

$$\mathcal{L} = \frac{1}{D}\sum_{i=0}^{D}\left[\frac{1}{2}\exp(-\log\hat{\sigma}_i^2)||\mathbf{y}_i - \hat{\mathbf{y}}_i||^2 + \frac{1}{2}\log\hat{\sigma}_i^2\right], \tag{9.1}$$

where $\mathbf{y}$ is the true (simulated) CMB map, $\hat{\mathbf{y}}$ is the predicted CMB map, and $\hat{\sigma}^2$ is the predicted variance, all averaged over D pixels. As per Kendall and Gal (2017), the log variance is used as the model output for numerical stability reasons. The two types of uncertainty are combined by evaluating the network many times with dropout. Each iteration, a random noise realization is drawn from a Gaussian distribution with variance corresponding to the per-pixel aleatoric uncertainty estimates, which is added to the CMB map realization. The mean and variance of this set of map realizations are then calculated to produce the final output.

9.2 Training data

In order to train a neural network model, many training samples are required. Since we have only one sky to observe, and since we do not know the ground truth for that sky, simulations must be used. For this, the pre-launch *Planck Sky Model* was used (Delabrouille et al. 2013). This mostly template-based model includes diffuse synchrotron, thermal dust, free–free, spinning dust, and CO foreground components, as well as the Sunyaev–Zel′dovich effect for clusters and the cosmic infrared background (CIB), using the knowledge of these components that was available when the Planck spacecraft was launched. Although the Planck Sky Model can also simulate strong and weak point sources, these were excluded

from the simulations that were run, since we found the inclusion of point sources to cause issues with the neural network training; a catalog of point sources can be used to mask the final output map before angular power spectrum estimation instead, if desired, but we did not find this masking to be necessary. The Planck Sky Model generates randomized constrained realizations of these foreground components, which match prior observations to within observational errors. As the uncertainty on the knowledge of these components is included in the simulations, this can be used to help better calibrate the uncertainties in the Bayesian neural network model. Furthermore, it allows for creating an independent foreground cleaning technique, since no Planck data are used in the simulations, except for derived instrument parameters. For reasons that will be discussed in Section 9.4.1, the thermal dust simulation in the Planck Sky Model code was modified to use Model 8 of Finkbeiner et al. (1999) instead of the default of Model 7.[2]

To generate the CMB used with the simulations, the CLASS Boltzman code (Blas et al. 2011) was used with parameters drawn from the WMAP9 flat ΛCDM Monte Carlo chains (Hinshaw et al. 2013) to produce theory angular power spectra, some of which deviate significantly from current constraints; thus, the simulated spectra cover the range of plausible spectra. These spectra were then modified to account for CMB lensing, and the lensed spectra were used to generate sky map realizations. The simulated sky was then observed using the Planck 2018 reduced instrument model (Planck Collaboration I 2020), without including instrument noise or instrument beams, using temperature data from both the 70 GHz low frequency instrument (LFI) band and all six high frequency instrument (HFI) bands. A resolution of NSIDE=512 was used,[3] and the sky was smoothed to a common 13.1′ Gaussian beam, corresponding to the largest Planck beam used, that

[2]Contrary to the description in Delabrouille et al. (2013), the pre-launch Planck Sky Model (v1.7.8) does not include Model 8 as an option.

[3]Early experimentation was performed with maps with resolution NSIDE=64 or NSIDE=256, since this was considerably faster.

of the 70 GHz channel. Finally, the *TT* variance maps included with the published Planck 2018 frequency maps were used to add white instrument noise to the simulated frequency maps. In addition to the observed frequency maps, the reference CMB map was also saved, to be used as the reference during model training. A set of 1000 simulations was created. Using the full simulation set, a set of normalization values were derived and applied multiplicatively to keep each frequency map between −1 and 1 and each CMB map between 0 and 1, with the CMB maps centered at 0.5.

9.3 Training procedure

For training, the set of 1000 simulations was subdivided into a training set consisting of 800 simulations and a test set containing the remaining 200 simulations. As mentioned in the previous section, these simulations differ in their randomized constrained realizations of foreground components, CMB components (including ΛCDM parameters), and instrument noise realizations. The network architecture was implemented using *TensorFlow* (Abadi et al. 2016), and the training procedure was performed on an NVIDIA TITAN RTX GPU. For training, the *AMSGrad* optimizer (Reddi et al. 2018) was used with default parameters. Six filters were used for each convolution layer (parameter F in Figure 9.1), with a polynomial order of nine; increasing either parameter was found to only provide small improvements, while thus increasing the network complexity unnecessarily. For Concrete dropout, a length scale of 10^{-4} was used. The final model includes 50 729 trainable parameters.

The model was trained for 400 epochs,[4] with a batch size of one,[5] a process that took approximately nine days. The *TensorFlow Large Model Support* library (Le et al. 2019)[6] was used to allow portions of the model to be swapped out of GPU memory during the training

[4] An epoch refers to a single pass through the training dataset during the training procedure.
[5] The batch size denotes how many examples are used simultaneously during the training procedure.
[6] `https://github.com/IBM/tensorflow-large-model-support`

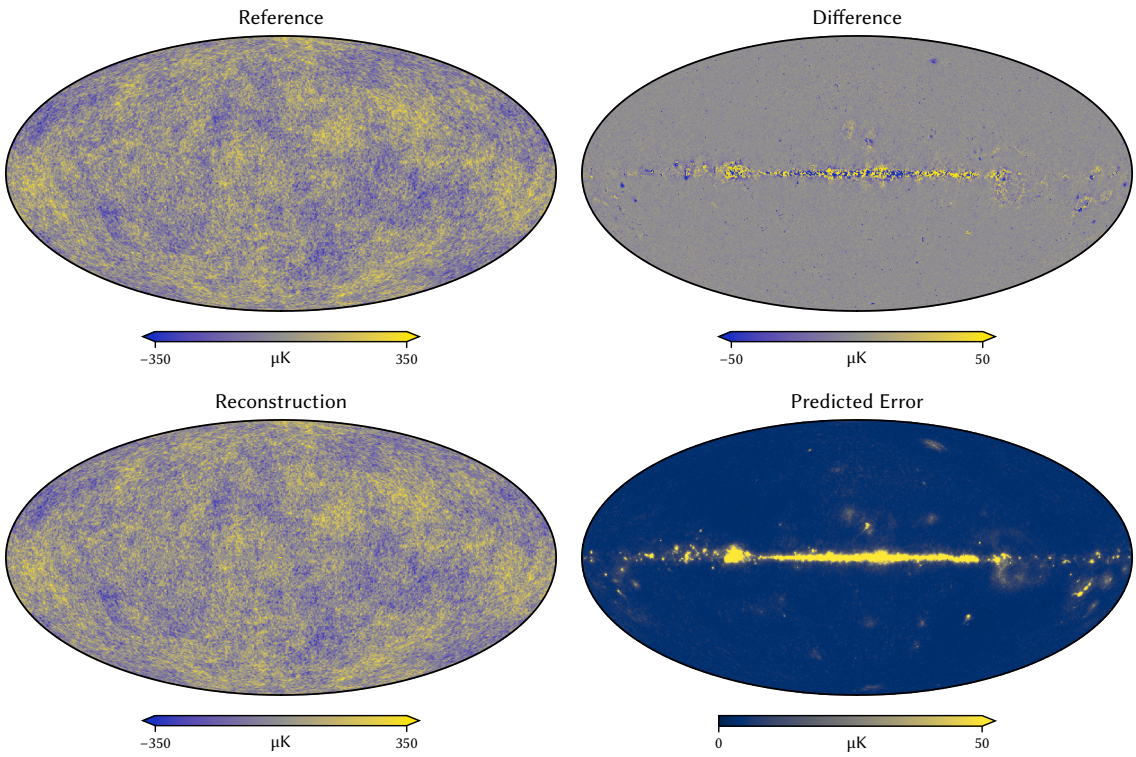

Figure 9.2: Simulated maps from one example in the test set. The top left panel shows the reference CMB map, the bottom left panel shows the CMB as reconstructed by the neural network model, the top right panel shows the difference between the reference and reconstructed CMB maps, and the bottom right panel shows the error predicted by the neural network model. Note that the color scales for the difference and predicted error saturate on portions of the Galactic plane.

process, since the model could not otherwise be trained, even with a batch size of one, due to GPU memory limitations. After training, the average loss [Equation (9.1)] on the training simulations was −7.22, while the average loss on the test simulations was also −7.22. This shows that the model was not overfit to the training examples, although this does not eliminate the possibility of fitting to shortcomings in the simulations relative to reality.

The results of evaluating the trained model for one set of maps from the test set are shown in Figure 9.2, utilizing the average of 100 evaluations to predict the uncertainty using the previously mentioned procedure of Kendall and Gal (2017). The reconstructed

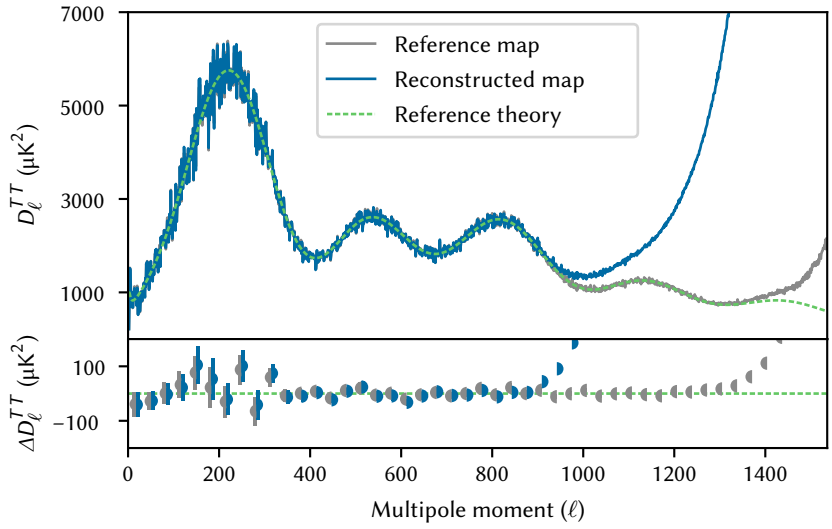

Figure 9.3: Temperature angular power spectra of one example of simulated maps from the test set. In the top panel, the spectrum calculated from the reference map is shown in gray, the spectrum calculated from the reconstructed map is shown in blue, and the theory spectrum used to create the simulations is shown in green. The reference map spectrum is calculated without a mask or weighting, while the reconstructed map spectrum is calculated using the mask and weighting scheme described in the text. The bottom panel shows the mean and standard deviation of the difference between the map-derived spectra and the theory spectrum, using a bin width of 33, for both maps.

map shows only small residuals away from the Galactic plane (<4 μK, on average), and the larger residual directly on the plane (~40 μK, on average) is accurately captured by the predicted error. For this particular example, the mean predicted standard error and mean absolute difference over the full sky are 5.4 μK and 3.8 μK, respectively, after masking map pixels with a predicted standard error >50 μK; for 75% of pixels, the actual difference was within the predicted standard error. These statistics show that the predicted error is close, although slightly overestimated.

The temperature angular power spectrum of the reconstructed map was also evaluated using the *PolSpice* estimator (Chon et al. 2004), with pixels with a predicted standard error >50 μK masked out and inverse variance weighting used for the remaining pixels, using the inverse of their predicted errors. While the reconstructed map shown in Figure 9.2 is

the pixel-space average of the 100 individual maps output by the model, including noise drawn from their predicted errors, the angular power auto-spectrum was instead calculated for each individual map after the predicted noise was added;[7] the resulting spectra were then averaged. This procedure was used as it produced improved results at smaller angular scales, when compared to the auto-spectrum calculated from the averaged map or to the average of cross-spectra calculated between the individual maps output by the model. The recovered spectrum matches up until $\ell \approx 900$, after which beam and noise effects start causing the spectra to diverge; the Planck 70 GHz channel becomes noise-dominated for $\ell \gtrsim 800$ (Planck Collaboration II 2020), so this is a likely cause of the lack of signal at these smaller angular scales. Beam effects also affect the reference map, since it is also smoothed with the common 13.1′ beam, but to a lesser degree. The above procedure was also applied to 100 map sets from the training set, and the resulting spectra were compared to spectra computed from the reference maps, to check for bias. A small multiplicative bias ($\lesssim 3\%$) was found; this was corrected by using the results of a quadratic fit to the bias for $2 < \ell < 900$, which resulted in a mean absolute residual bias of ∼0.3% for $2 < \ell < 900$. The need for this correction is thought to be primarily due to the use of a pixel-space loss function during model training; however, it also effectively removes any noise bias in the ℓ range it is fit to. The resulting angular power spectrum of the test set example is shown in Figure 9.3, along with the spectrum of the reference map and the theory spectrum used to generate the reference map and the simulated frequency maps.

9.4 Application to Planck observations

With the neural network model trained, it was then applied to the Planck 2018 intensity frequency maps, which were first corrected to use the same CIB monopole value as used

[7] A common mask and weighting were used for all the maps.

in the simulations and normalized with the same multiplicative constants used for the simulations. The same evaluation procedure previously described for the map set from the simulation test set was used, with one caveat. The resulting reconstructed map has a slight monopole bias of ∼1 μK, which is corrected for by masking ±10° from the Galactic plane, calculating the median, and subtracting the median value from the map. This correction was implemented during development of the technique, when the bias was noticeably larger; while no longer necessary, the correction was kept for the sake of robustness. The bias is thought to result from the model using slightly different weights for each of the Planck frequency maps than when the simulated maps are used. When the frequency maps do not share a common beam window function, the power spectrum is also biased for the Planck data, unlike for the simulations, which is thought to be due to the same effect. This led to the decision to use a common beam window function for the frequency maps; the 30 GHz and 44 GHz LFI maps were thus excluded, due to their larger beam sizes. Although additionally excluding the 70 GHz LFI map would allow for use of an even smaller beam, this benefit was found to be outweighed by the reduced ability to exclude synchrotron foregrounds.

The resulting reconstructed CMB map is shown in Figure 9.4, with comparison to the Planck Commander foreground-cleaned map (Planck Collaboration IV 2020), smoothed with the same beam window function. A residual structure like that of a diffuse Galactic foreground component is apparent, although it cannot be conclusively attributed to a specific foreground cleaning technique, since the ground truth is unknown. The mean predicted standard error is 5.5 μK, after masking map pixels with a predicted standard error >50 μK. With the same masking, the mean absolute difference between the reconstructed map and the Commander map is 6.2 μK, which is similar to the differences between the various Planck foreground cleaning techniques; the Commander–NILC, Commander–SEVEM, and Commander–SMICA mean absolute differences are 5.5 μK, 6.6 μK, and 4.1 μK,

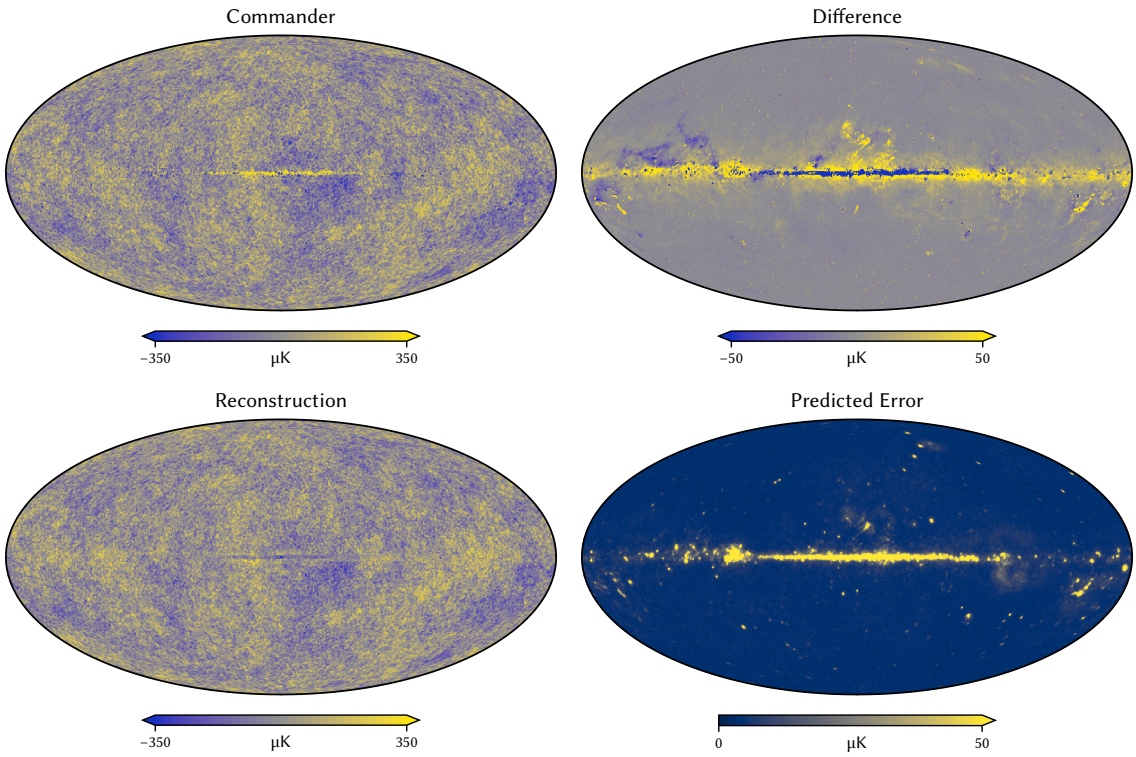

Figure 9.4: Planck maps. The top left panel shows the Planck Commander CMB map, the bottom left panel shows the CMB as reconstructed by the neural network model from the Planck frequency maps, the top right panel shows the difference between the Commander map and the neural network reconstructed CMB map, and the bottom right panel shows the error predicted by the neural network model. Note that the color scales for the difference and predicted error saturate on portions of the Galactic plane.

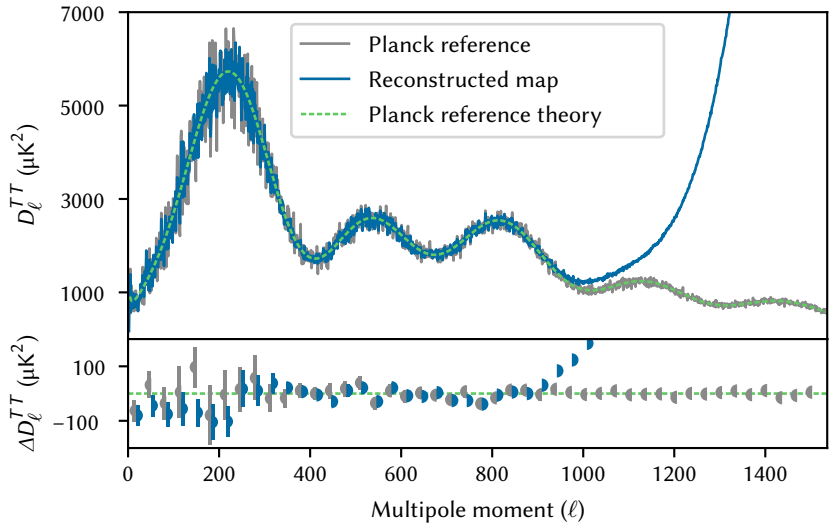

Figure 9.5: Planck temperature angular power spectra. In the top panel, the published Planck spectrum is shown in gray, the spectrum calculated from the neural network reconstructed map is shown in blue, and the published Planck best-fit theory spectrum is shown in green. The reconstructed map spectrum is calculated using the mask and weighting scheme described in the text. The bottom panel shows the mean and standard deviation of the difference between the map-derived spectrum and the published Planck best-fit theory spectrum, as well as the difference between the published Planck spectrum and the published Planck best-fit theory spectrum, using a bin width of 33.

respectively, when smoothed with the previously used beam window function. Since the predicted standard error matches the mean absolute difference with the Commander map, the error may again be overestimated, as the Commander map also has errors. However, since some of the errors in the two maps might be correlated, a definitive conclusion on the error estimation accuracy cannot be reached.

The temperature angular power spectrum of the reconstructed Planck map was also evaluated using the same technique that was used for the map set from the simulation test set, results of which are shown in Figure 9.5, with comparison to the published Planck spectrum and the published Planck best-fit theory spectrum (Planck Collaboration V 2020). The recovered spectrum roughly matches up until $\ell \approx 900$, after which beam effects start causing the spectrum to diverge. It also has slightly less scatter relative to the theory

spectrum than the published Planck spectrum and has some minor disagreement for $\ell \lesssim 250$. This minor disagreement is not manifested when the model is used to clean Planck FFP10 simulated frequency maps (Planck Collaboration III 2020), so the source of the disagreement is not captured by these simulations. Furthermore, using frequency maps produced by the *SRoll 2.2* mapmaking pipeline (Delouis et al. 2019) produces no significant difference in the recovered spectrum when compared to that recovered from the Planck 2018 legacy maps, so the minor disagreement is likely not due to the instrumental effects that are better corrected for by the SRoll 2.2 pipeline. The exact cause of the minor disagreement remains unknown.

9.4.1 Comparison of dust models

The neural network model described above was originally trained using Planck Sky Model simulations that modeled thermal dust using Model 7 of Finkbeiner et al. (1999). When the neural network model was evaluated on the Planck frequency maps, this resulted in a large residual when compared to the Commander map. As qualitative visual comparisons showed a strong resemblance between this residual and maps of Galactic thermal dust foregrounds, the Planck Sky Model code was modified to use Model 8 of Finkbeiner et al. (1999) for its thermal dust simulation. The angular power spectrum resulting from evaluation of the Planck frequency maps with the neural network model did not significantly change when the dust model used for the training simulations was updated, suggesting that the minor discrepancy with the published Planck spectrum for $\ell \lesssim 250$ is not related to the dust model used in the simulations. Due to the long training process, an updated simulation was first evaluated with the existing neural network model before the model was retrained. The residual for this evaluation appeared similar to the difference between the results of the evaluation of Planck frequency maps and the Commander map, suggesting that the new simulations better matched it. The original Commander difference, the updated

simulation residual, and the difference between the two are shown in Figure 9.6.

The difference between the two residual maps strongly resembles the actual difference between the retrained model evaluated on the Planck frequency maps and the Commander map, shown in Figure 9.4. This demonstrates the ability of the neural network model to evaluate revised simulations. Although retraining the model is computationally expensive, evaluating the trained model for new data is not. Thus, adjustments to simulations can be quickly checked against observations by comparing the residuals produced when the trained model is used to evaluate both the observations and the adjusted simulations; the closer the residuals match, the closer the adjusted simulations match the observations. Once the residuals are sufficiently close, the neural network model can be retrained on the adjusted simulations and then be used to reevaluate the observations. This process can be repeated as necessary. By only revising the simulations instead of both the simulations and the model, one can iterate on simulation differences and improvements much more quickly than would otherwise be possible. Although comparing the residuals requires an external foreground cleaning method to reference to, it has the advantage of providing hints as to what needs to be changed in the simulations. The neural network model alone also provides a metric by which to evaluate the fidelity of the simulations, its mean predicted error. This went from 11.2 μK to 7.4 μK without any sky regions masked, again showing the improvement of using the revised thermal dust model.

9.5 Conclusions

Using a Bayesian spherical convolutional neural network, a machine-learning foreground cleaning technique was developed for CMB observations. The model was trained on full-sky temperature simulations that incorporated knowledge of the millimeter and submillimeter sky that was available at the time of the launch of the Planck satellite. After being shown

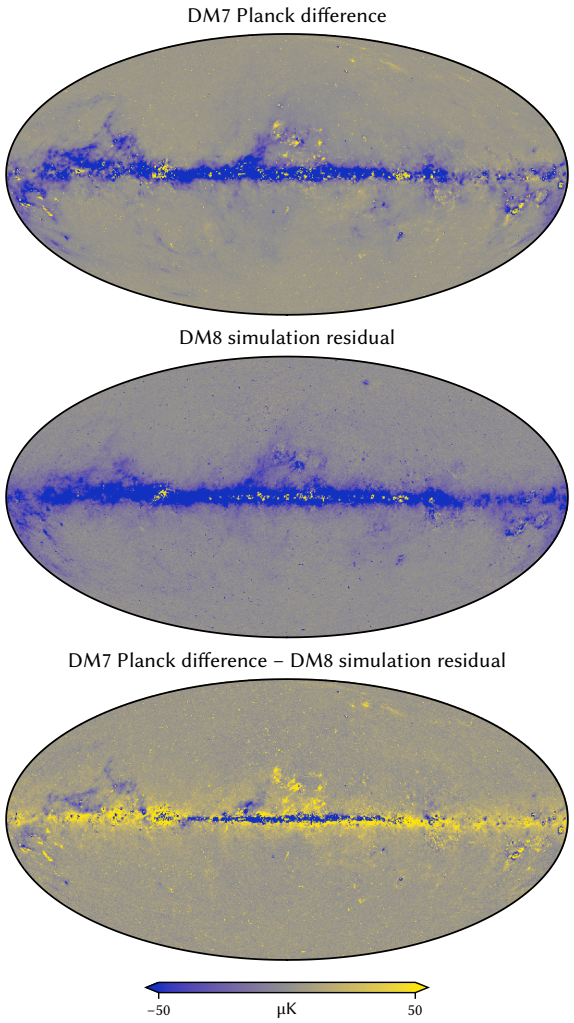

Figure 9.6: Dust model comparison. The top panel shows the difference between the Planck Commander map and the CMB as reconstructed by the neural network model when trained on simulations that use Model 7 of Finkbeiner et al. (1999) for thermal dust. The middle panel shows the residual when the same neural network model is used to evaluate a simulation that uses Model 8 of Finkbeiner et al. (1999) for thermal dust. The bottom panel shows the difference between the top two panels. Note that it strongly resembles the difference map shown in Figure 9.4, which was produced using the retrained model. All three panels share the same color scale; note that it saturates on portions of the Galactic plane, particularly for the top two panels.

to work on simulations, the trained model was applied to Planck observations. The neural network code, trained weights, and results have been made available (Petroff 2020).

As the current neural network model is trained on simulations, the effectiveness of the foreground cleaning is limited by the accuracy of the simulations. Thus, the effectiveness of the cleaning should improve if the training simulations are improved. While the Planck Sky Model is a template-based phenomenological method, the use of more physically motivated simulations, e.g., along the lines of what Puglisi et al. (2017) developed for diffuse Galactic CO emission, could potentially be an improvement, although the development of such foreground simulations is outside the scope of the present work. The demonstrated ability of the current neural network model to evaluate how well different simulations match observations or other simulations could potentially aid in the development of these improved simulations, as it can quickly provide a metric for evaluating how well the simulation matches and a residual map that can hint at what parts of the simulation still need to be improved.

In this work, only temperature maps were considered, so one avenue of future work is to extend the presented technique to polarization maps. In early experiments, the presented technique was found to perform poorly on Planck-like polarization maps. However, as the technique did work when the noise levels on the simulated polarization maps were reduced, the deficiency seems to be due to the inability to handle the higher noise levels in the Planck polarization maps; this caused the loss function to be best minimized by smoothing the maps and eliminating small angular scale features. Potential solutions to this issue include incorporating the angular power spectrum difference into the loss function or using generative adversarial networks (GANs) (Isola et al. 2017).

References

Abadi, M., A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mane, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viegas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng (Mar. 2016). "TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems". In: *arXiv e-prints*. arXiv: 1603.04467 [cs.DC].

Abitbol, M. H., Z. Ahmed, D. Barron, R. Basu Thakur, A. N. Bender, B. A. Benson, C. A. Bischoff, S. A. Bryan, J. E. Carlstrom, C. L. Chang, D. T. Chuss, K. T. Crowley, A. Cukierman, T. de Haan, M. Dobbs, T. Essinger-Hileman, J. P. Filippini, K. Ganga, J. E. Gudmundsson, N. W. Halverson, S. Hanany, S. W. Henderson, C. A. Hill, S.-P. P. Ho, J. Hubmayr, K. Irwin, O. Jeong, B. R. Johnson, S. A. Kernasovskiy, J. M. Kovac, A. Kusaka, A. T. Lee, S. Maria, P. Mauskopf, J. J. McMahon, L. Moncelsi, A. W. Nadolski, J. M. Nagy, M. D. Niemack, R. C. O'Brient, S. Padin, S. C. Parshley, C. Pryke, N. A. Roe, K. Rostem, J. Ruhl, S. M. Simon, S. T. Staggs, A. Suzuki, E. R. Switzer, O. Tajima, K. L. Thompson, P. Timbie, G. S. Tucker, J. D. Vieira, A. G. Vieregg, B. Westbrook, E. J. Wollack, K. W. Yoon, K. S. Young, and E. Y. Young (June 2017). "CMB-S4 Technology Book, First Edition". In: *arXiv e-prints*. arXiv: 1706.02464 [astro-ph.IM].

Addison, G. E., D. J. Watts, C. L. Bennett, M. Halpern, G. Hinshaw, and J. L. Weiland (Feb. 2018). "Elucidating ΛCDM: Impact of Baryon Acoustic Oscillation Measurements on the Hubble Constant Discrepancy". In: *Astrophys. J.* 853.2, 119. DOI: 10.3847/1538-4357/aaa1ed. arXiv: 1707.06547 [astro-ph.CO].

Ahmad, Q. R. et al. (Aug. 2001). "Measurement of the Rate of $\nu_e + d \rightarrow p + p + e^-$ Interactions Produced by ^{8}B Solar Neutrinos at the Sudbury Neutrino Observatory". In: *Phys. Rev. Lett.* 87.7, 071301. DOI: 10.1103/PhysRevLett.87.071301. arXiv: nucl-ex/0106015.

Albrecht, A., P. J. Steinhardt, M. S. Turner, and F. Wilczek (May 1982). "Reheating an Inflationary Universe". In: *Phys. Rev. Lett.* 48.20, pp. 1437–1440. DOI: 10.1103/PhysRevLett.48.1437.

Allison, R., P. Caucal, E. Calabrese, J. Dunkley, and T. Louis (Dec. 2015). "Towards a cosmological neutrino mass detection". In: *Phys. Rev. D* 92.12, 123535. DOI: 10.1103/PhysRevD.92.123535. arXiv: 1509.07471 [astro-ph.CO].

Alpher, R. A. and R. Herman (Nov. 1948). "Evolution of the Universe". In: *Nature* 162.4124, pp. 774–775. DOI: 10.1038/162774b0.

Appel, J. W., Z. Xu, I. L. Padilla, K. Harrington, B. P. Marquez, A. Ali, C. L. Bennett, M. K. Brewer, R. Bustos, M. Chan, D. T. Chuss, J. Cleary, J. Couto, S. Dahal, K. Denis, R. Dünner, J. R. Eimer, T. Essinger-Hileman, P. Fluxa, D. Gothe, G. C. Hilton, J. Hubmayr, J. Iuliano, J. Karakla, T. A. Marriage, N. J. Miller, C. Núñez, L. Parker, M. Petroff, C. D. Reintsema, K. Rostem, R. W. Stevens, D. A. N. Valle, B. Wang, D. J. Watts, E. J. Wollack, and L. Zeng (May 2019). "On-sky Performance of the CLASS Q-band Telescope". In: *Astrophys. J.* 876.2, p. 126. DOI: 10.3847/1538-4357/ab1652. arXiv: 1811.08287 [astro-ph.IM].

Armstrong, B. (Jan. 1967). "Spectrum line profiles: The Voigt function". In: *J. Quant. Spectrosc. Radiat. Transf.* 7.1, pp. 61–88. DOI: 10.1016/0022-4073(67)90057-x.

Astropy Collaboration, T. P. Robitaille, E. J. Tollerud, P. Greenfield, M. Droettboom, E. Bray, T. Aldcroft, M. Davis, A. Ginsburg, A. M. Price-Whelan, W. E. Kerzendorf, A. Conley, N. Crighton, K. Barbary, D. Muna, H. Ferguson, F. Grollier, M. M. Parikh, P. H. Nair, H. M. Günther, C. Deil, J. Woillez, S. Conseil, R. Kramer, J. E. H. Turner, L. Singer, R. Fox, B. A. Weaver, V. Zabalza, Z. I. Edwards, K. A. Bostroem, D. J. Burke, A. R. Casey, S. M. Crawford, N. Dencheva, J. Ely, T. Jenness, K. Labrie, P. L. Lim, F. Pierfederici, A. Pontzen, A. Ptak, B. Refsdal, M. Servillat, and O. Streicher (Sept. 2013). "Astropy: A community Python package for astronomy". In: *Astron. Astrophys.* 558, A33. DOI: 10.1051/0004-6361/201322068. arXiv: 1307.6212 [astro-ph.IM].

Astropy Collaboration et al. (Aug. 2018). "The Astropy Project: Building an Open-science Project and Status of the v2.0 Core Package". In: *Astron. J.* 156.3, p. 123. DOI: 10.3847/1538-3881/aabc4f. arXiv: 1801.02634 [astro-ph.IM].

Auld, T., M. Bridges, and M. P. Hobson (July 2008). "COSMONET: fast cosmological parameter estimation in non-flat models using neural networks". In: *Mon. Not. Roy. Astron. Soc.* 387.4, pp. 1575–1582. DOI: 10.1111/j.1365-2966.2008.13279.x. arXiv: astro-ph/0703445.

Aylor, K., M. Haq, L. Knox, Y. Hezaveh, and L. Perreault-Levasseur (Jan. 2021). "Cleaning our own dust: simulating and separating galactic dust foregrounds with neural networks". In: *Mon. Not. Roy. Astron. Soc.* 500.3, pp. 3889–3897. DOI: 10.1093/mnras/staa3344. arXiv: 1909.06467 [astro-ph.IM].

Bai, M., D. Xia, and M. Jin (Sept. 2017). "Effects of Coating Material Properties on the Wideband Reflectivity Performance of Microwave Calibration Targets". In: *IEEE Trans. Antennas Propag.* 65.9, pp. 4909–4913. ISSN: 0018-926X. DOI: 10.1109/TAP.2017.2731383.

Ball, N. M. and R. J. Brunner (Jan. 2010). "Data Mining and Machine Learning in Astronomy". In: *Int. J. Mod. Phys. D* 19.7, pp. 1049–1106. DOI: 10.1142/S0218271810017160. arXiv: 0906.2173 [astro-ph.IM].

Barkana, R. and A. Loeb (July 2001). "In the beginning: the first sources of light and the reionization of the universe". In: *Phys. Rep.* 349.2, pp. 125–238. DOI: 10.1016/S0370-1573(01)00019-9. arXiv: astro-ph/0010468.

Basak, S. and J. Delabrouille (Jan. 2012). "A needlet internal linear combination analysis of WMAP 7-year data: estimation of CMB temperature map and power spectrum". In: *Mon.*

Not. Roy. Astron. Soc. 419.2, pp. 1163–1175. DOI: 10.1111/j.1365-2966.2011.19770.x. arXiv: 1106.5383 [astro-ph.CO].

Battistelli, E. S., M. Amiri, B. Burger, M. Halpern, S. Knotek, M. Ellis, X. Gao, D. Kelly, M. Macintosh, K. Irwin, and C. Reintsema (May 2008). "Functional Description of Read-out Electronics for Time-Domain Multiplexed Bolometers for Millimeter and Sub-millimeter Astronomy". In: *J. Low Temp. Phys.* 151.3-4, pp. 908–914. DOI: 10.1007/s10909-008-9772-z.

Bennett, C. L., R. S. Hill, G. Hinshaw, M. R. Nolta, N. Odegard, L. Page, D. N. Spergel, J. L. Weiland, E. L. Wright, M. Halpern, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, G. S. Tucker, and E. Wollack (Sept. 2003). "First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Foreground Emission". In: *Astrophys. J. Suppl. Ser.* 148.1, pp. 97–117. DOI: 10.1086/377252. arXiv: astro-ph/0302208.

Bennett, C. L., D. Larson, J. L. Weiland, N. Jarosik, G. Hinshaw, N. Odegard, K. M. Smith, R. S. Hill, B. Gold, M. Halpern, E. Komatsu, M. R. Nolta, L. Page, D. N. Spergel, E. Wollack, J. Dunkley, A. Kogut, M. Limon, S. S. Meyer, G. S. Tucker, and E. L. Wright (Oct. 2013). "Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results". In: *Astrophys. J. Suppl. Ser.* 208.2, 20. DOI: 10.1088/0067-0049/208/2/20. arXiv: 1212.5225 [astro-ph.CO].

Berestetskii, V., E. Lifshitz, and L. Pitaevskii (1982). *Quantum Electrodynamics*. 2nd ed. London: Pergamon Press. ISBN: 0080265030.

BICEP2 Collaboration, Keck Array Collaboration, P. A. R. Ade, Z. Ahmed, R. W. Aikin, K. D. Alexander, D. Barkats, S. J. Benton, C. A. Bischoff, J. J. Bock, R. Bowens-Rubin, J. A. Brevik, I. Buder, E. Bullock, V. Buza, J. Connors, J. Cornelison, B. P. Crill, M. Crumrine, M. Dierickx, L. Duband, C. Dvorkin, J. P. Filippini, S. Fliescher, J. Grayson, G. Hall, M. Halpern, S. Harrison, S. R. Hildebrandt, G. C. Hilton, H. Hui, K. D. Irwin, J. Kang, K. S. Karkare, E. Karpel, J. P. Kaufman, B. G. Keating, S. Kefeli, S. A. Kernasovskiy, J. M. Kovac, C. L. Kuo, N. A. Larsen, K. Lau, E. M. Leitch, M. Lueker, K. G. Megerian, L. Moncelsi, T. Namikawa, C. B. Netterfield, H. T. Nguyen, R. O'Brient, R. W. Ogburn, S. Palladino, C. Pryke, B. Racine, S. Richter, A. Schillaci, R. Schwarz, C. D. Sheehy, A. Soliman, T. St. Germaine, Z. K. Staniszewski, B. Steinbach, R. V. Sudiwala, G. P. Teply, K. L. Thompson, J. E. Tolan, C. Tucker, A. D. Turner, C. Umiltà, A. G. Vieregg, A. Wandui, A. C. Weber, D. V. Wiebe, J. Willmert, C. L. Wong, W. L. K. Wu, H. Yang, K. W. Yoon, and C. Zhang (Nov. 2018). "Constraints on Primordial Gravitational Waves Using Planck, WMAP, and New BICEP2/Keck Observations through the 2015 Season". In: *Phys. Rev. Lett.* 121.22, 221301. DOI: 10.1103/PhysRevLett.121.221301. arXiv: 1810.05216 [astro-ph.CO].

Bischoff, C., A. Brizius, I. Buder, Y. Chinone, K. Cleary, R. N. Dumoulin, A. Kusaka, R. Monsalve, S. K. Næss, L. B. Newburgh, G. Nixon, R. Reeves, K. M. Smith, K. Vanderlinde, I. K. Wehus, M. Bogdan, R. Bustos, S. E. Church, R. Davis, C. Dickinson, H. K. Eriksen, T. Gaier, J. O. Gundersen, M. Hasegawa, M. Hazumi, C. Holler, K. M. Huffenberger, W. A. Imbriale, K. Ishidoshiro, M. E. Jones, P. Kangaslahti, D. J. Kapner, C. R. Lawrence, E. M. Leitch, M. Limon, J. J. McMahon, A. D. Miller, M. Nagai, H. Nguyen, T. J. Pearson,

L. Piccirillo, S. J. E. Radford, A. C. S. Readhead, J. L. Richards, D. Samtleben, M. Seiffert, M. C. Shepherd, S. T. Staggs, O. Tajima, K. L. Thompson, R. Williamson, B. Winstein, E. J. Wollack, and J. T. L. Zwart (May 2013). "The Q/U Imaging ExperimenT Instrument". In: *Astrophys. J.* 768.1, 9. DOI: 10.1088/0004-637X/768/1/9. arXiv: 1207.5562 [astro-ph.IM].

Blas, D., J. Lesgourgues, and T. Tram (July 2011). "The Cosmic Linear Anisotropy Solving System (CLASS). Part II: Approximation schemes". In: *J. Cosmol. Astropart. Phys.* 2011.7, 034. DOI: 10.1088/1475-7516/2011/07/034. arXiv: 1104.2933 [astro-ph.CO].

Bobin, J., J. .-L. Starck, J. M. Fadili, Y. Moudden, and D. L. Donoho (Nov. 2007). "Morphological Component Analysis: An Adaptive Thresholding Strategy". In: *IEEE Trans. Image Process.* 16.11, pp. 2675–2681. DOI: 10.1109/TIP.2007.907073.

Bond, J. R., L. Kofman, and D. Pogosyan (Apr. 1996). "How filaments of galaxies are woven into the cosmic web". In: *Nature* 380.6575, pp. 603–606. DOI: 10.1038/380603a0. arXiv: astro-ph/9512141.

Bondi, H. (1952). *Cosmology*. Cambridge: Cambridge University Press.

Born, M. and E. Wolf (1959). *Principles of Optics*. London: Pergamon Press.

Bowman, J. C. and A. Hammerlindl (2008). "Asymptote: A vector graphics language". In: *TUGboat* 29.2, pp. 288–294. URL: https://www.tug.org/TUGboat/tb29-2/tb92bowman.pdf.

Caldeira, J., W. L. K. Wu, B. Nord, C. Avestruz, S. Trivedi, and K. T. Story (July 2019). "DeepCMB: Lensing reconstruction of the cosmic microwave background with deep neural networks". In: *Astron. Comput.* 28, 100307. DOI: 10.1016/j.ascom.2019.100307. arXiv: 1810.01483 [astro-ph.CO].

Cardoso, J.-F., M. L. Jeune, J. Delabrouille, M. Betoule, and G. Patanchon (Oct. 2008). "Component Separation With Flexible Models—Application to Multichannel Astrophysical Observations". In: *IEEE J. Sel. Top. Signal Process.* 2.5, pp. 735–746. DOI: 10.1109/jstsp.2008.2005346. arXiv: 0803.1814.

Carleo, G., I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová (Oct. 2019). "Machine learning and the physical sciences". In: *Rev. Mod. Phys.* 91.4, 045002. DOI: 10.1103/RevModPhys.91.045002. arXiv: 1903.10563 [physics.comp-ph].

Carlstrom, J. E., P. A. R. Ade, K. A. Aird, B. A. Benson, L. E. Bleem, S. Busetti, C. L. Chang, E. Chauvin, H. .-M. Cho, T. M. Crawford, A. T. Crites, M. A. Dobbs, N. W. Halverson, S. Heimsath, W. L. Holzapfel, J. D. Hrubes, M. Joy, R. Keisler, T. M. Lanting, A. T. Lee, E. M. Leitch, J. Leong, W. Lu, M. Lueker, D. Luong-Van, J. J. McMahon, J. Mehl, S. S. Meyer, J. J. Mohr, T. E. Montroy, S. Padin, T. Plagge, C. Pryke, J. E. Ruhl, K. K. Schaffer, D. Schwan, E. Shirokoff, H. G. Spieler, Z. Staniszewski, A. A. Stark, C. Tucker, K. Vanderlinde, J. D. Vieira, and R. Williamson (May 2011). "The 10 Meter South Pole Telescope". In: *Publ. Astron. Soc. Pac.* 123.903, p. 568. DOI: 10.1086/659879. arXiv: 0907.4445 [astro-ph.IM].

Carroll, S. M., G. B. Field, and R. Jackiw (Feb. 1990). "Limits on a Lorentz- and parity-violating modification of electrodynamics". In: *Phys. Rev. D* 41.4, pp. 1231–1240. DOI: 10.1103/PhysRevD.41.1231.

Cheshire, S. and M. Baker (Apr. 1999). "Consistent overhead byte stuffing". In: *IEEE-ACM Trans. Netw.* 7.2, pp. 159–172. ISSN: 1063-6692. DOI: 10.1109/90.769765.

Choi, J., H. Ishitsuka, S. Mima, S. Oguri, K. Takahashi, and O. Tajima (Nov. 2013). "Radio-transparent multi-layer insulation for radiowave receivers". In: *Rev. Sci. Instrum.* 84.11, p. 114502. DOI: 10.1063/1.4827081. URL: https://doi.org/10.1063/1.4827081.

Chon, G., A. Challinor, S. Prunet, E. Hivon, and I. Szapudi (May 2004). "Fast estimation of polarization power spectra using correlation functions". In: *Mon. Not. Roy. Astron. Soc.* 350.3, pp. 914–926. DOI: 10.1111/j.1365-2966.2004.07737.x. arXiv: astro-ph/0303414.

Chulliat, A., P. Alken, M. Nair, A. Woods, and S. Maus (2015). *The Enhanced Magnetic Model 2015–2020*. DOI: 10.7289/v56971hv.

Chuss, D. T., K. Rostem, E. J. Wollack, L. Berman, F. Colazo, M. DeGeorge, K. Helson, and M. Sagliocca (2017). "A cryogenic thermal source for detector array characterization". In: *Rev. Sci. Instrum.* 88.10, p. 104501. ISSN: 1089-7623. DOI: 10.1063/1.4996751. arXiv: 1707.05879 [astro-ph.IM].

Chuss, D. T., E. J. Wollack, R. Henry, H. Hui, A. J. Juarez, M. Krejny, S. H. Moseley, and G. Novak (Jan. 2012). "Properties of a variable-delay polarization modulator". In: *Appl. Optics* 51.2, p. 197. DOI: 10.1364/ao.51.000197. arXiv: 1106.5984 [astro-ph.IM].

Ciuca, R. and O. F. Hernández (Aug. 2017). "A Bayesian framework for cosmic string searches in CMB maps". In: *J. Cosmol. Astropart. Phys.* 2017.8, 028. DOI: 10.1088/1475-7516/2017/08/028. arXiv: 1706.04131 [astro-ph.CO].

Clevert, D.-A., T. Unterthiner, and S. Hochreiter (2016). "Fast and Accurate Deep Network Learning by Exponential Linear Units (ELUs)". In: *Int. Conf. Learn. Represent.* arXiv: 1511.07289 [cs.LG].

Clough, S., F. Kneizys, and R. Davies (Oct. 1989). "Line shape and the water vapor continuum". In: *Atmos. Res.* 23.3, pp. 229–241. DOI: 10.1016/0169-8095(89)90020-3.

Cogollos, S., J. Vague, V. E. Boria, and J. D. Martinez (Apr. 2018). "Novel Planar and Waveguide Implementations of Impedance Matching Networks Based on Tapered Lines Using Generalized Superellipses". In: *IEEE Trans. Microw. Theory Tech.* 66.4, pp. 1874–1884. DOI: 10.1109/tmtt.2018.2791952.

Cohen, T. S., M. Geiger, J. Köhler, and M. Welling (2018). "Spherical CNNs". In: *Int. Conf. Learn. Represent.* arXiv: 1801.10130 [cs.LG]. URL: https://openreview.net/forum?id=Hkbd5xZRb.

Cooray, A., A. Melchiorri, and J. Silk (Feb. 2003). "Is the cosmic microwave background circularly polarized?" In: *Phys. Lett. B* 554.1-2, pp. 1–6. DOI: 10.1016/s0370-2693(02)03291-4. arXiv: astro-ph/0205214.

Dahal, S., M. Amiri, J. W. Appel, C. L. Bennett, L. Corbett, R. Datta, K. Denis, T. Essinger-Hileman, M. Halpern, K. Helson, G. Hilton, J. Hubmayr, B. Keller, T. Marriage, C. Nunez, M. Petroff, C. Reintsema, K. Rostem, K. U-Yen, and E. Wollack (Jan. 2020). "The

CLASS 150/220 GHz Polarimeter Array: Design, Assembly, and Characterization". In: *J. Low Temp. Phys.* 199.1-2, pp. 289–297. DOI: 10.1007/s10909-019-02317-0. arXiv: 1908.00480 [astro-ph.IM].

Dahal, S., A. Ali, J. W. Appel, T. Essinger-Hileman, C. Bennett, M. Brewer, R. Bustos, M. Chan, D. T. Chuss, J. Cleary, F. Colazo, J. Couto, K. Denis, R. Dünner, J. Eimer, T. Engelhoven, P. Fluxa, M. Halpern, K. Harrington, K. Helson, G. Hilton, G. Hinshaw, J. Hubmayr, J. Iuliano, J. Karakla, T. Marriage, J. McMahon, N. Miller, C. Nuñez, I. Padilla, G. Palma, L. Parker, M. Petroff, B. Pradenas, R. Reeves, C. Reintsema, K. Rostem, M. Sagliocca, K. U-Yen, D. Valle, B. Wang, Q. Wang, D. Watts, J. Weiland, E. Wollack, Z. Xu, Z. Yan, and L. Zeng (July 2018). "Design and characterization of the Cosmology Large Angular Scale Surveyor (CLASS) 93 GHz focal plane". In: *Proc. SPIE*. Ed. by J. Zmuidzinas and J.-R. Gao. Vol. 10708. DOI: 10.1117/12.2311812. arXiv: 1807.03927 [astro-ph.IM].

Datta, R., S. Aiola, S. K. Choi, M. Devlin, J. Dunkley, R. Dünner, P. A. Gallardo, M. Gralla, M. Halpern, M. Hasselfield, M. Hilton, A. D. Hincks, S.-P. P. Ho, J. Hubmayr, K. M. Huffenberger, J. P. Hughes, A. Kosowsky, C. H. López-Caraballo, T. Louis, M. Lungu, T. Marriage, L. Maurin, J. McMahon, K. Moodley, S. K. Naess, F. Nati, M. D. Niemack, L. A. Page, B. Partridge, H. Prince, S. T. Staggs, E. R. Switzer, E. J. Wollack, and G. Farren (July 2019). "The Atacama Cosmology Telescope: two-season ACTPol extragalactic point sources and their polarization properties". In: *Mon. Not. Roy. Astron. Soc.* 486.4, pp. 5239–5262. DOI: 10.1093/mnras/sty2934. arXiv: 1811.01854 [astro-ph.CO].

Defferrard, M., M. Milani, F. Gusset, and N. Perraudin (2020). "DeepSphere: a graph-based spherical CNN". In: *Int. Conf. Learn. Represent.* URL: https://openreview.net/forum?id=B1e3OlStPB.

Delabrouille, J., M. Betoule, J.-B. Melin, M.-A. Miville-Deschênes, J. Gonzalez-Nuevo, M. L. Jeune, G. Castex, G. de Zotti, S. Basak, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, J.-P. Bernard, F. R. Bouchet, D. L. Clements, A. da Silva, C. Dickinson, F. Dodu, K. Dolag, F. Elsner, L. Fauvet, G. Faÿ, G. Giardino, S. Leach, J. Lesgourgues, M. Liguori, J. F. Macías-Pérez, M. Massardi, S. Matarrese, P. Mazzotta, L. Montier, S. Mottet, R. Paladini, B. Partridge, R. Piffaretti, G. Prezeau, S. Prunet, S. Ricciardi, M. Roman, B. Schaefer, and L. Toffolatti (May 2013). "The pre-launch *Planck* Sky Model: a model of sky emission at submillimetre to centimetre wavelengths". In: *Astron. Astrophys.* 553, A96. DOI: 10.1051/0004-6361/201220019. arXiv: 1207.3675 [astro-ph.CO].

Delouis, J. .-M., L. Pagano, S. Mottet, J. .-L. Puget, and L. Vibert (Sept. 2019). "SRoll2: an improved mapmaking approach to reduce large-scale systematic effects in the Planck High Frequency Instrument legacy maps". In: *Astron. Astrophys.* 629, A38. DOI: 10.1051/0004-6361/201834882. arXiv: 1901.11386 [astro-ph.CO].

Denis, K. L., N. T. Cao, D. T. Chuss, J. Eimer, J. R. Hinderks, W. .-T. Hsieh, S. H. Moseley, T. R. Stevenson, D. J. Talley, K. U. -yen, and E. J. Wollack (Dec. 2009). "Fabrication of an Antenna-Coupled Bolometer for Cosmic Microwave Background Polarimetry". In: *AIP Conf. Proc.* Ed. by B. Young, B. Cabrera, and A. Miller. Vol. 1185, pp. 371–374. DOI: 10.1063/1.3292355.

Deutsch, A.-S., E. Dimastrogiovanni, M. C. Johnson, M. Münchmeyer, and A. Terrana (Dec. 2018). "Reconstruction of the remote dipole and quadrupole fields from the kinetic Sunyaev Zel'dovich and polarized Sunyaev Zel'dovich effects". In: *Phys. Rev. D* 98.12, 123501. DOI: 10.1103/PhysRevD.98.123501. arXiv: 1707.08129 [astro-ph.CO].

Dicke, R. H. and P. J. E. Peebles (1979). "The big bang cosmology — enigmas and nostrums". In: *General relativity: An Einstein centenary survey*. Cambridge: Cambridge University Press, pp. 504–517.

Dicke, R. H., P. J. E. Peebles, P. G. Roll, and D. T. Wilkinson (July 1965). "Cosmic Black-Body Radiation." In: *Astrophys. J.* 142, pp. 414–419. DOI: 10.1086/148306.

Eimer, J. R., C. L. Bennett, D. T. Chuss, T. Marriage, E. J. Wollack, and L. Zeng (Sept. 2012). "The Cosmology Large Angular Scale Surveyor (CLASS): 40 GHz optical design". In: *Proc. SPIE*. Ed. by W. S. Holland. Vol. 8452. SPIE, pp. 619–633. DOI: 10.1117/12.925464. arXiv: 1211.0041 [astro-ph.IM].

Emerson, W. (July 1973). "Electromagnetic wave absorbers and anechoic chambers through the years". In: *IEEE Trans. Antennas Propag.* 21.4, pp. 484–490. ISSN: 0018-926X. DOI: 10.1109/TAP.1973.1140517.

Eriksen, H. K., J. B. Jewell, C. Dickinson, A. J. Band ay, K. M. Górski, and C. R. Lawrence (Mar. 2008). "Joint Bayesian Component Separation and CMB Power Spectrum Estimation". In: *Astrophys. J.* 676.1, pp. 10–32. DOI: 10.1086/525277. arXiv: 0709.1058.

Essinger-Hileman, T., A. Ali, M. Amiri, J. W. Appel, D. Araujo, C. L. Bennett, F. Boone, M. Chan, H.-M. Cho, D. T. Chuss, F. Colazo, E. Crowe, K. Denis, R. Dünner, J. Eimer, D. Gothe, M. Halpern, K. Harrington, G. C. Hilton, G. F. Hinshaw, C. Huang, K. Irwin, G. Jones, J. Karakla, A. J. Kogut, D. Larson, M. Limon, L. Lowry, T. Marriage, N. Mehrle, A. D. Miller, N. Miller, S. H. Moseley, G. Novak, C. Reintsema, K. Rostem, T. Stevenson, D. Towner, K. U-Yen, E. Wagner, D. Watts, E. J. Wollack, Z. Xu, and L. Zeng (July 2014). "CLASS: the Cosmology Large Angular Scale Surveyor". In: *Proc. SPIE*. Ed. by W. S. Holland and J. Zmuidzinas. Vol. 9153. SPIE, pp. 491–513. DOI: 10.1117/12.2056701. arXiv: 1408.4788 [astro-ph.IM].

Esteves, C., C. Allen-Blanchette, A. Makadia, and K. Daniilidis (2018). "Learning SO(3) Equivariant Representations with Spherical CNNs". In: *Eur. Conf. Comput. Vis.* Ed. by V. Ferrari, M. Hebert, C. Sminchisescu, and Y. Weiss. Springer International Publishing, pp. 54–70. ISBN: 978-3-030-01261-8. DOI: 10.1007/978-3-030-01261-8_4. arXiv: 1711.06721 [cs.CV].

Faddeyeva, V. N. and N. M. Terent'ev (1961). *Mathematical Tables of Values of the Function* $w(z) = e^{-z^2}\left(1 + \frac{2i}{\sqrt{\pi}}\int_0^z e^{t^2}dt\right)$ *for Complex Argument*. Ed. by V. A. Fok. Trans. by D. G. Fry and B. A. Hons. Oxford: Pergamon Press.

Fairchild, M. D. (2013). *Color Appearance Models*. 3rd. Wiley-IS&T Series in Imaging Science and Technology. Chichester, West Sussex: Wiley. ISBN: 978-1-119-96703-3.

Filippini, J. P., P. A. R. Ade, M. Amiri, S. J. Benton, R. Bihary, J. J. Bock, J. R. Bond, J. A. Bonetti, S. A. Bryan, B. Burger, H. C. Chiang, C. R. Contaldi, B. P. Crill, O. Doré, M. Farhang, L. M. Fissel, N. N. Gandilo, S. R. Golwala, J. E. Gudmundsson, M. Halpern, M. Hasselfield, G. Hilton, W. Holmes, V. V. Hristov, K. D. Irwin, W. C. Jones, C. L. Kuo,

C. J. MacTavish, P. V. Mason, T. E. Montroy, T. A. Morford, C. B. Netterfield, D. T. O'Dea, A. S. Rahlin, C. D. Reintsema, J. E. Ruhl, M. C. Runyan, M. A. Schenker, J. A. Shariff, J. D. Soler, A. Trangsrud, C. Tucker, R. S. Tucker, and A. D. Turner (July 2010). "SPIDER: a balloon-borne CMB polarimeter for large angular scales". In: *Proc. SPIE*. Ed. by W. S. Holland and J. Zmuidzinas. Vol. 7741. DOI: 10.1117/12.857720. arXiv: 1106.2158 [astro-ph.CO].

Finkbeiner, D. P., M. Davis, and D. J. Schlegel (Oct. 1999). "Extrapolation of Galactic Dust Emission at 100 Microns to Cosmic Microwave Background Radiation Frequencies Using FIRAS". In: *Astrophys. J.* 524.2, pp. 867–886. DOI: 10.1086/307852. arXiv: astro-ph/9905128.

Fixsen, D. J. (Dec. 2009). "The Temperature of the Cosmic Microwave Background". In: *Astrophys. J.* 707.2, pp. 916–920. DOI: 10.1088/0004-637X/707/2/916. arXiv: 0911.1955 [astro-ph.CO].

Fixsen, D. J., E. S. Cheng, J. M. Gales, J. C. Mather, R. A. Shafer, and E. L. Wright (Dec. 1996). "The Cosmic Microwave Background Spectrum from the Full COBE FIRAS Data Set". In: *Astrophys. J.* 473, p. 576. DOI: 10.1086/178173. arXiv: astro-ph/9605054.

Fluke, C. J. and C. Jacobs (Dec. 2019). "Surveying the reach and maturity of machine learning and artificial intelligence in astronomy". In: *Wiley Interdiscip. Rev.-Data Mining Knowl. Discov.* 10.2, e1349. DOI: 10.1002/widm.1349. arXiv: 1912.02934 [astro-ph.IM].

Folkner, W. M., J. G. Williams, D. H. Boggs, R. S. Park, and P. Kuchynka (2014). *The Planetary and Lunar Ephemerides DE430 and DE431*. IPN Progress Report 42-196. Jet Propulsion Laboratory, California Institute of Technology. URL: https://ipnpr.jpl.nasa.gov/progress_report/42-196/196C.pdf.

Friedmann, A. (Jan. 1922). "Über die Krümmung des Raumes". In: *Z. Phys.* 10, pp. 377–386. DOI: 10.1007/BF01332580.

— (Dec. 1924). "Über die Möglichkeit einer Welt mit konstanter negativer Krümmung des Raumes". In: *Z. Phys.* 21.1, pp. 326–332. DOI: 10.1007/BF01328280.

Fukuda, H., M. Shimizu, and G. Nakamura (2001). "New Gosper Space Filling Curves". In: *Proc. Int. Conf. Comput. Graph. Imaging*, pp. 34–38. URL: https://kilin.clas.kitasato-u.ac.jp/museum/gosperex/343-024.pdf.

Fukuda, Y. et al. (Aug. 1998). "Evidence for Oscillation of Atmospheric Neutrinos". In: *Phys. Rev. Lett.* 81.8, pp. 1562–1567. DOI: 10.1103/PhysRevLett.81.1562. arXiv: hep-ex/9807003.

Gal, Y. and Z. Ghahramani (June 2016). "Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning". In: *Proc. Int. Conf. Mach. Learn.* Ed. by M. F. Balcan and K. Q. Weinberger. Vol. 48. New York: PMLR, pp. 1050–1059. URL: https://proceedings.mlr.press/v48/gal16.html.

Gal, Y., J. Hron, and A. Kendall (2017). "Concrete Dropout". In: *Adv. Neural Inf. Process. Syst.* Ed. by I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett. Curran Associates, Inc., pp. 3581–3590. arXiv: 1705.07832 [stat.ML]. URL: https://papers.nips.cc/paper/6949-concrete-dropout.

Gardner, M. (Dec. 1976). "Mathematical Games. In which 'monster' curves force redefinition of the word 'curve'". In: *Sci. Am.* 235.6, pp. 124–133. DOI: 10.1038/scientificamerican1276-124.

Gilbert, S. and N. Lynch (June 2002). "Brewer's Conjecture and the Feasibility of Consistent, Available, Partition-tolerant Web Services". In: *SIGACT News* 33.2, pp. 51–59. ISSN: 0163-5700. DOI: 10.1145/564585.564601.

Giovannini, M. (Dec. 2009). "V-mode polarization of the cosmic microwave background". In: *Phys. Rev. D* 80.12. DOI: 10.1103/physrevd.80.123013. arXiv: 0909.3629 [astro-ph.CO].

Gordon, I., L. Rothman, C. Hill, R. Kochanov, Y. Tan, P. Bernath, M. Birk, V. Boudon, A. Campargue, K. Chance, B. Drouin, J.-M. Flaud, R. Gamache, J. Hodges, D. Jacquemart, V. Perevalov, A. Perrin, K. Shine, M.-A. Smith, J. Tennyson, G. Toon, H. Tran, V. Tyuterev, A. Barbe, A. Császár, V. Devi, T. Furtenbacher, J. Harrison, J.-M. Hartmann, A. Jolly, T. Johnson, T. Karman, I. Kleiner, A. Kyuberis, J. Loos, O. Lyulin, S. Massie, S. Mikhailenko, N. Moazzen-Ahmadi, H. Müller, O. Naumenko, A. Nikitin, O. Polyansky, M. Rey, M. Rotger, S. Sharpe, K. Sung, E. Starikova, S. Tashkun, J. V. Auwera, G. Wagner, J. Wilzewski, P. Wcisło, S. Yu, and E. Zak (Dec. 2017). "The HITRAN2016 molecular spectroscopic database". In: *J. Quant. Spectrosc. Radiat. Transf.* 203, pp. 3–69. DOI: 10.1016/j.jqsrt.2017.06.038.

Gorski, K. M., E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann (Apr. 2005). "HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere". In: *Astrophys. J.* 622.2, pp. 759–771. DOI: 10.1086/427976. arXiv: astro-ph/0409513.

Gothe, D. (Aug. 2015). "Measuring Inflation with CLASS". PhD thesis. Johns Hopkins University. URL: https://jhir.library.jhu.edu/handle/1774.2/38721.

Gray, F. (1947). "Pulse code communication". U.S. pat. US2632058.

Guth, A. H. (Jan. 1981). "Inflationary universe: A possible solution to the horizon and flatness problems". In: *Phys. Rev. D* 23.2, pp. 347–356. DOI: 10.1103/PhysRevD.23.347.

Halley, E. (1720). "Of the Infinity of the Sphere of Fix'd Stars". In: *Philos. Trans. R. Soc. Lond.* 31.364, pp. 22–24. DOI: 10.1098/rstl.1720.0006.

Hanany, S., P. Ade, A. Balbi, J. Bock, J. Borrill, A. Boscaleri, P. de Bernardis, P. G. Ferreira, V. V. Hristov, A. H. Jaffe, A. E. Lange, A. T. Lee, P. D. Mauskopf, C. B. Netterfield, S. Oh, E. Pascale, B. Rabii, P. L. Richards, G. F. Smoot, R. Stompor, C. D. Winant, and J. H. P. Wu (Dec. 2000). "MAXIMA-1: A Measurement of the Cosmic Microwave Background Anisotropy on Angular Scales of 10′-5°". In: *Astrophys. J. Lett.* 545.1, pp. L5–L9. DOI: 10.1086/317322. arXiv: astro-ph/0005123.

Hanany, S. and P. Rosenkranz (Dec. 2003). "Polarization of the atmosphere as a foreground for cosmic microwave background polarization experiments". In: *New Astron. Rev.* 47.11-12, pp. 1159–1165. DOI: 10.1016/j.newar.2003.09.017. arXiv: astro-ph/0307052.

Hansen, F. K., A. J. Banday, H. K. Eriksen, K. M. Górski, and P. B. Lilje (Sept. 2006). "Foreground Subtraction of Cosmic Microwave Background Maps Using WI-FIT (Wavelet-

Based High-Resolution Fitting of Internal Templates)". In: *Astrophys. J.* 648.2, pp. 784–796. DOI: 10.1086/506015. arXiv: astro-ph/0603308.

Harrington, K., A. Ali, J. W. Appel, C. L. Bennett, M. Brewer, R. Bustos, M. Chan, D. T. Chuss, J. Cleary, J. D. Couto, S. Dahal, K. L. Denis, R. Dünner, J. R. Eimer, T. M. Essinger-Hileman, P. Fluxa, M. Halpern, G. C. Hilton, G. F. Hinshaw, J. Hubmayr, J. Iuliano, J. Karakla, T. Marriage, N. J. Miller, J. J. McMahon, C. Nuñez, I. Padilla, G. Palma, L. Parker, M. Petroff, B. P. Márquez, R. Reeves, C. D. Reintsema, K. Rostem, D. A. N. Valle, T. V. Engelhoven, B. Wang, Q. Wang, D. J. Watts, J. Weiland, E. J. Wollack, Z. L. Xu, Z. Yan, L. Zeng, M. DeGeorge, and T. W. Grunberg (July 2018). "Variable-delay polarization modulators for the CLASS telescopes". In: *Proc. SPIE*. Ed. by J. Zmuidzinas and J.-R. Gao. Vol. 10708. SPIE, pp. 369–390. DOI: 10.1117/12.2313614. arXiv: 1807.03807 [astro-ph.IM].

Harrington, K., R. Datta, K. Osumi, A. Ali, J. W. Appel, C. L. Bennett, M. K. Brewer, R. Bustos, M. Chan, D. T. Chuss, J. Cleary, J. Denes Couto, S. Dahal, R. Dünner, J. R. Eimer, T. Essinger-Hileman, J. Hubmayr, F. R. Espinoza Inostroza, J. Iuliano, J. Karakla, Y. Li, T. A. Marriage, N. J. Miller, C. Núñez, I. L. Padilla, L. Parker, M. A. Petroff, B. Pradenas Márquez, R. Reeves, P. Fluxá Rojas, K. Rostem, D. A. Nunes Valle, D. J. Watts, J. L. Weiland, E. J. Wollack, and Z. Xu (Jan. 2021). "Two Year Cosmology Large Angular Scale Surveyor (CLASS) Observations: Long Timescale Stability Achieved with a Front-End Variable-delay Polarization Modulator at 40 GHz". In: *Astrophys. J., submitted*. arXiv: 2101.00034 [astro-ph.IM].

Harrington, K., T. Marriage, A. Ali, J. W. Appel, C. L. Bennett, F. Boone, M. Brewer, M. Chan, D. T. Chuss, F. Colazo, S. Dahal, K. Denis, R. Dünner, J. Eimer, T. Essinger-Hileman, P. Fluxa, M. Halpern, G. Hilton, G. F. Hinshaw, J. Hubmayr, J. Iuliano, J. Karakla, J. McMahon, N. T. Miller, S. H. Moseley, G. Palma, L. Parker, M. Petroff, B. Pradenas, K. Rostem, M. Sagliocca, D. Valle, D. Watts, E. Wollack, Z. Xu, and L. Zeng (July 2016). "The Cosmology Large Angular Scale Surveyor". In: *Proc. SPIE*. Ed. by W. S. Holland and J. Zmuidzinas. Vol. 9914. SPIE, pp. 380–400. DOI: 10.1117/12.2233125. arXiv: 1608.08234 [astro-ph.IM].

Harris, C. R., K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett, A. Haldane, J. F. del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, and T. E. Oliphant (Sept. 2020). "Array programming with NumPy". In: *Nature* 585.7825, pp. 357–362. DOI: 10.1038/s41586-020-2649-2. arXiv: 2006.10256 [cs.MS].

He, S., S. Ravanbakhsh, and S. Ho (2018). "Analysis of Cosmic Microwave Background with Deep Learning". In: *Int. Conf. Learn. Represent.* URL: https://openreview.net/forum?id=B15uoOyvz.

Hedman, M. M. (May 2002). "The Princeton IQU Experiment and constraints on the polarization of the cosmic microwave background at 90 GHz". PhD thesis. Princeton University. BIBCODE: 2002PhDT.....13H.

Herbert, F. (Sept. 1974). “Spectrum line profiles: A generalized Voigt function including collisional narrowing”. In: *J. Quant. Spectrosc. Radiat. Transf.* 14.9, pp. 943–951. DOI: 10.1016/0022-4073(74)90021-1.

Hilbert, D. (1891). “Ueber die stetige Abbildung einer Linie auf ein Flächenstück”. In: *Math. Ann.* 38.3, pp. 459–460. ISSN: 1432-1807. DOI: 10.1007/BF01199431.

Hill, R. J. (May 1986). “Water vapor-absorption line shape comparison using the 22-GHz line: The Van Vleck-Weisskopf shape affirmed”. In: *Radio Sci.* 21.3, pp. 447–451. DOI: 10.1029/rs021i003p00447.

Hiltgen, A., K. Paterson, and M. Brandestini (1996). “Single-track Gray codes”. In: *IEEE Trans. Inf. Theory* 42.5, pp. 1555–1561. DOI: 10.1109/18.532900.

Hinshaw, G., D. Larson, E. Komatsu, D. N. Spergel, C. L. Bennett, J. Dunkley, M. R. Nolta, M. Halpern, R. S. Hill, N. Odegard, L. Page, K. M. Smith, J. L. Weiland, B. Gold, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, G. S. Tucker, E. Wollack, and E. L. Wright (Oct. 2013). “Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Parameter Results”. In: *Astrophys. J. Suppl. Ser.* 208.2, 19. DOI: 10.1088/0067-0049/208/2/19. arXiv: 1212.5226 [astro-ph.CO].

Hirsch, M. (2019). *space-physics/msise00: MSISE-00 in Python and Matlab*. Version 1.6.1. DOI: 10.5281/zenodo.3379163.

Hogg, D. W. (May 1999). “Distance measures in cosmology”. In: *arXiv e-prints*. arXiv: astro-ph/9905116.

Holloway, C. L., R. R. DeLyser, R. F. German, P. McKenna, and M. Kanda (Feb. 1997). “Comparison of electromagnetic absorber used in anechoic and semi-anechoic chambers for emissions and immunity testing of digital devices”. In: *IEEE Trans. Electromag. Compat.* 39.1, pp. 33–47. ISSN: 0018-9375. DOI: 10.1109/15.554693.

Houtz, D. A., W. Emery, D. Gu, K. Jacob, A. Murk, D. K. Walker, and R. J. Wylde (Aug. 2017). “Electromagnetic Design and Performance of a Conical Microwave Blackbody Target for Radiometer Calibration”. In: *IEEE Trans. Geosci. Remote Sens.* 55.8, pp. 4586–4596. ISSN: 0196-2892. DOI: 10.1109/TGRS.2017.2694319.

Hu, W. and T. Okamoto (Aug. 2002). “Mass Reconstruction with Cosmic Microwave Background Polarization”. In: *Astrophys. J.* 574.2, pp. 566–574. DOI: 10.1086/341110. arXiv: astro-ph/0111606.

Hu, W. and M. White (Oct. 1997). “A CMB polarization primer”. In: *New Astron.* 2.4, pp. 323–344. DOI: 10.1016/s1384-1076(97)00022-5. arXiv: astro-ph/9706147.

Hubble, E. (Mar. 1929). “A Relation between Distance and Radial Velocity among Extra-Galactic Nebulae”. In: *Proc. Natl. Acad. Sci. U.S.A.* 15.3, pp. 168–173. DOI: 10.1073/pnas.15.3.168.

Hunter, J. D. (2007). “Matplotlib: A 2D Graphics Environment”. In: *IEEE Comput. Sci. Eng.* 9.3, pp. 90–95. DOI: 10.1109/mcse.2007.55.

Iafolla, L., M. Filipozzi, S. Freund, A. Zam, G. Rauter, and P. C. Cattin (Sept. 2020). “Proof of concept of a novel absolute rotary encoder”. In: *Sens. Actuator A-Phys.* 312, 112100. DOI: 10.1016/j.sna.2020.112100.

Ichiki, K., H. Kanai, N. Katayama, and E. Komatsu (Mar. 2019). "Delta-map method of removing CMB foregrounds with spatially varying spectra". In: *Prog. Theor. Exp. Phys.* 2019.3, 033E01. DOI: 10.1093/ptep/ptz009. arXiv: 1811.03886 [astro-ph.CO].

Inomata, K. and M. Kamionkowski (Feb. 2019). "Circular polarization of the cosmic microwave background from vector and tensor perturbations". In: *Phys. Rev. D* 99.4. DOI: 10.1103/physrevd.99.043501. arXiv: 1811.04957 [astro-ph.CO].

Isola, P., J.-Y. Zhu, T. Zhou, and A. A. Efros (July 2017). "Image-to-Image Translation with Conditional Adversarial Networks". In: *IEEE Conf. Comput. Vis. Pattern Recognit.* IEEE. DOI: 10.1109/cvpr.2017.632. arXiv: 1611.07004 [cs.CV].

Iuliano, J., J. Eimer, L. Parker, G. Rhoades, A. Ali, J. W. Appel, C. Bennett, M. Brewer, R. Bustos, D. Chuss, J. Cleary, J. Couto, S. Dahal, K. Denis, R. Dünner, T. Essinger-Hileman, P. Fluxa, M. Halpern, K. Harrington, K. Helson, G. Hilton, G. Hinshaw, J. Hubmayr, J. Karakla, T. Marriage, N. Miller, J. J. McMahon, C. Nuñez, I. Padilla, G. Palma, M. Petroff, B. Pradenas Márquez, R. Reeves, C. Reintsema, K. Rostem, D. Augusto Nunes Valle, T. Van Engelhoven, B. Wang, Q. Wang, D. Watts, J. Weiland, E. J. Wollack, Z. Xu, Z. Yan, and L. Zeng (July 2018). "The Cosmology Large Angular Scale Surveyor receiver design". In: *Proc. SPIE*. Ed. by J. Zmuidzinas and J.-R. Gao. Vol. 10708. DOI: 10.1117/12.2312954. arXiv: 1807.04167 [astro-ph.IM].

Jackson, J. D. (1998). *Classical Electrodynamics*. 3rd. New York: John Wiley & Sons. ISBN: 978-0-471-30932-1.

Jaki, S. L. (Mar. 1967). "Olbers', Halley's, or Whose Paradox?" In: *Am. J. Phys.* 35.3, pp. 200–210. DOI: 10.1119/1.1974001.

Janaswamy, R. (Feb. 1992). "Oblique scattering from lossy periodic surfaces with applications to anechoic chamber absorbers". In: *IEEE Trans. Antennas Propag.* 40.2, pp. 162–169. ISSN: 0018-926X. DOI: 10.1109/8.127400.

Janssen, M. A., ed. (1993). *Atmospheric Remote Sensing by Microwave Radiometry*. New York: John Wiley & Sons. ISBN: 0471628913.

Jiang, C. M., J. Huang, K. Kashinath, Prabhat, P. Marcus, and M. Nießner (2019). "Spherical CNNs on Unstructured Grids". In: *Int. Conf. Learn. Represent.* arXiv: 1901.02039 [cs.CV]. URL: https://openreview.net/forum?id=Bkl-43C9FQ.

Jones, R., P. Haufe, E. Sells, P. Iravani, V. Olliver, C. Palmer, and A. Bowyer (2011). "RepRap — the replicating rapid prototyper". In: *Robotica* 29.1, pp. 177–191. DOI: 10.1017/S026357471000069X.

Kamionkowski, M., A. Kosowsky, and A. Stebbins (June 1997). "Statistics of cosmic microwave background polarization". In: *Phys. Rev. D* 55.12, pp. 7368–7388. DOI: 10.1103/PhysRevD.55.7368. arXiv: astro-ph/9611125.

Katayama, N. and E. Komatsu (Aug. 2011). "Simple Foreground Cleaning Algorithm for Detecting Primordial B-mode Polarization of the Cosmic Microwave Background". In: *Astrophys. J.* 737.2, 78. DOI: 10.1088/0004-637X/737/2/78. arXiv: 1101.5210 [astro-ph.CO].

Keating, B., P. Timbie, A. Polnarev, and J. Steinberger (Mar. 1998). "Large Angular Scale Polarization of the Cosmic Microwave Background Radiation and the Feasibility of Its

Detection". In: *Astrophys. J.* 495.2, pp. 580–596. DOI: 10.1086/305312. arXiv: astro-ph/9710087.

Kendall, A. and Y. Gal (2017). "What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision?" In: *Adv. Neural Inf. Process. Syst.* Ed. by I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett. Curran Associates, Inc., pp. 5574–5584. arXiv: 1703.04977 [cs.CV]. URL: https://papers.nips.cc/paper/7141-what-uncertainties-do-we-need-in-bayesian-deep-learning-for-computer-vision.

King, S. and P. Lubin (July 2016). "Circular polarization of the CMB: Foregrounds and detection prospects". In: *Phys. Rev. D* 94.2. DOI: 10.1103/physrevd.94.023501. arXiv: 1606.04112 [astro-ph.CO].

Kondor, R., Z. Lin, and S. Trivedi (2018). "Clebsch–Gordan Nets: a Fully Fourier Space Spherical Convolutional Neural Network". In: *Adv. Neural Inf. Process. Syst.* Ed. by S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett. Curran Associates, Inc., pp. 10117–10126. arXiv: 1806.09231 [stat.ML]. URL: https://papers.nips.cc/paper/8215-clebschgordan-nets-a-fully-fourier-space-spherical-convolutional-neural-network.

Koshelev, M. A., I. N. Vilkov, and M. Y. Tretyakov (Mar. 2015). "Pressure broadening of oxygen fine structure lines by water". In: *J. Quant. Spectrosc. Radiat. Transf.* 154, pp. 24–27. DOI: 10.1016/j.jqsrt.2014.11.019.

— (Jan. 2016). "Collisional broadening of oxygen fine structure lines: The impact of temperature". In: *J. Quant. Spectrosc. Radiat. Transf.* 169, pp. 91–95. DOI: 10.1016/j.jqsrt.2015.09.018.

Krachmalnicoff, N. and M. Tomasi (Aug. 2019). "Convolutional neural networks on the HEALPix sphere: a pixel-based algorithm and its application to CMB data analysis". In: *Astron. Astrophys.* 628, A129. DOI: 10.1051/0004-6361/201935211. arXiv: 1902.04083 [astro-ph.IM].

Krizhevsky, A., I. Sutskever, and G. E. Hinton (2012). "ImageNet Classification with Deep Convolutional Neural Networks". In: *Adv. Neural Inf. Process. Syst.* Ed. by F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger. Curran Associates, Inc., pp. 1097–1105. URL: https://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.

Kronberger, R. and P. Soboll (Oct. 2016). "New 3D printed microwave metamaterial absorbers with conductive printing materials". In: *Eur. Microw. Conf.* Pp. 596–599. DOI: 10.1109/EuMC.2016.7824413.

Kunkee, D. B., G. A. Poe, D. J. Boucher, S. D. Swadley, Y. Hong, J. E. Wessel, and E. A. Uliana (Apr. 2008). "Design and Evaluation of the First Special Sensor Microwave Imager/Sounder". In: *IEEE Trans. Geosci. Remote Sensing* 46.4, pp. 863–883. DOI: 10.1109/tgrs.2008.917980.

Kusaka, A., J. Appel, T. Essinger-Hileman, J. A. Beall, L. E. Campusano, H.-M. Cho, S. K. Choi, K. Crowley, J. W. Fowler, P. Gallardo, M. Hasselfield, G. Hilton, S.-P. P. Ho, K. Irwin, N. Jarosik, M. D. Niemack, G. W. Nixon, M. Nolta, J. Page Lyman A., G. A. Palma, L. Parker,

S. Raghunathan, C. D. Reintsema, J. Sievers, S. M. Simon, S. T. Staggs, K. Visnjic, and K.-W. Yoon (Sept. 2018). "Results from the Atacama B-mode Search (ABS) experiment". In: *J. Cosmol. Astropart. Phys.* 2018.9, 005. DOI: 10.1088/1475-7516/2018/09/005. arXiv: 1801.01218 [astro-ph.CO].

Lam, S. K., A. Pitrou, and S. Seibert (2015). "Numba". In: *LLVM '15*. ACM Press. DOI: 10.1145/2833157.2833162. URL: https://doi.org/10.1145/2833157.2833162.

Lang, D., D. W. Hogg, K. Mierle, M. Blanton, and S. Roweis (May 2010). "Astrometry.net: Blind Astrometric Calibration of Arbitrary Astronomical Images". In: *Astrophys. J.* 139.5, pp. 1782–1800. DOI: 10.1088/0004-6256/139/5/1782. arXiv: 0910.2233 [astro-ph.IM].

Larsson, R., S. A. Buehler, P. Eriksson, and J. Mendrok (Jan. 2014). "A treatment of the Zeeman effect using Stokes formalism and its implementation in the Atmospheric Radiative Transfer Simulator (ARTS)". In: *J. Quant. Spectrosc. Radiat. Transf.* 133, pp. 445–453. DOI: 10.1016/j.jqsrt.2013.09.006.

Larsson, R., B. Lankhaar, and P. Eriksson (Feb. 2019). "Updated Zeeman effect splitting coefficients for molecular oxygen in planetary applications". In: *J. Quant. Spectrosc. Radiat. Transf.* 224, pp. 431–438. DOI: 10.1016/j.jqsrt.2018.12.004.

Le, T. D., H. Imai, Y. Negishi, and K. Kawachiya (2019). "Automatic GPU memory management for large neural models in TensorFlow". In: *Proc. ACM SIGPLAN Int. Symp. Mem. Manag.* ACM Press. DOI: 10.1145/3315573.3329984.

Lemaître, G. (Jan. 1927). "Un Univers homogène de masse constante et de rayon croissant rendant compte de la vitesse radiale des nébuleuses extra-galactiques". In: *Ann. Soc. Sci. Brux.* 47, pp. 49–59. BIBCODE: 1927ASSB...47...49L.

Lenoir, W. B. (Dec. 1967). "Propagation of Partially Polarized Waves in a Slightly Anisotropic Medium". In: *J. Appl. Phys.* 38.13, pp. 5283–5290. DOI: 10.1063/1.1709315.

— (Jan. 1968). "Microwave spectrum of molecular oxygen in the mesosphere". In: *J. Geophys. Res.* 73.1, pp. 361–376. DOI: 10.1029/ja073i001p00361.

Lewis, A. and S. Bridle (Nov. 2002). "Cosmological parameters from CMB and other data: A Monte Carlo approach". In: *Phys. Rev. D* 66.10, 103511. DOI: 10.1103/PhysRevD.66.103511. arXiv: astro-ph/0205436.

Lewis, A., A. Challinor, and A. Lasenby (Aug. 2000). "Efficient Computation of Cosmic Microwave Background Anisotropies in Closed Friedmann-Robertson-Walker Models". In: *Astrophys. J.* 538.2, pp. 473–476. DOI: 10.1086/309179. arXiv: astro-ph/9911177.

Liebe, H., G. Hufford, and M. Cotton (Nov. 1993). "Propagation modeling of moist air and suspended water/ice particles at frequencies below 1000 GHz". In: *AGARD Conference Proceedings 542, Atmospheric Propagation Effects Through Natural and Man-Made Obscurants for Visible to MM-Wave Radiation*. Neuilly sur Seine, France: North Atlantic Treaty Organization. ISBN: 9283507274. URL: https://apps.dtic.mil/dtic/tr/fulltext/u2/a276919.pdf.

Liebe, H., P. Rosenkranz, and G. Hufford (Nov. 1992). "Atmospheric 60-GHz oxygen spectrum: New laboratory measurements and line parameters". In: *J. Quant. Spectrosc. Radiat. Transf.* 48.5-6, pp. 629–643. DOI: 10.1016/0022-4073(92)90127-p.

Liebe, H. J. (Nov. 1981). "Modeling attenuation and phase of radio waves in air at frequencies below 1000 GHz". In: *Radio Sci.* 16.6, pp. 1183–1199. DOI: 10.1029/rs016i006p01183.

— (June 1989). "MPM—An atmospheric millimeter-wave propagation model". In: *Int. J. Infrared Millimeter Waves* 10.6, pp. 631–650. DOI: 10.1007/bf01009565.

Linde, A. D. (Feb. 1982). "A new inflationary universe scenario: A possible solution of the horizon, flatness, homogeneity, isotropy and primordial monopole problems". In: *Phys. Lett. B* 108.6, pp. 389–393. DOI: 10.1016/0370-2693(82)91219-9.

Lubin, P., P. Melese, and G. Smoot (Oct. 1983). "Linear and circular polarization of the cosmic background radiation". In: *Astrophys. J. Lett.* 273, pp. L51–L54. DOI: 10.1086/184128.

Machta, L. and E. Hughes (June 1970). "Atmospheric Oxygen in 1967 to 1970". In: *Science* 168.3939, pp. 1582–1584. DOI: 10.1126/science.168.3939.1582.

Mackay, A. (Nov. 1989). "Proof of polarisation independence and nonexistence of crosspolar terms for targets presenting n-fold ($n > 2$) rotational symmetry with special reference to frequency-selective surfaces". In: *Electron. Lett.* 25.24, pp. 1624–1625. ISSN: 0013-5194. DOI: 10.1049/el:19891088.

Mainini, R., D. Minelli, M. Gervasi, G. Boella, G. Sironi, A. Baú, S. Banfi, A. Passerini, A. D. Lucia, and F. Cavaliere (Aug. 2013). "An improved upper limit to the CMB circular polarization at large angular scales". In: *J. Cosmol. Astropart. Phys.* 2013.08, pp. 033–033. DOI: 10.1088/1475-7516/2013/08/033. arXiv: 1307.6090 [astro-ph.CO].

Makarov, D., M. Tretyakov, and P. Rosenkranz (June 2011). "60-GHz oxygen band: Precise experimental profiles and extended absorption modeling in a wide temperature range". In: *J. Quant. Spectrosc. Radiat. Transf.* 112.9, pp. 1420–1428. DOI: 10.1016/j.jqsrt.2011.02.018.

Maltoni, M., T. Schwetz, M. Tórtola, and J. W. F. Valle (Sept. 2004). "Status of global fits to neutrino oscillations". In: *New J. Phys.* 6.1, p. 122. DOI: 10.1088/1367-2630/6/1/122. arXiv: hep-ph/0405172.

Mandelbrot, B. B. (1977). *Fractals: form, chance, and dimension*. San Francisco: W.H. Freeman. ISBN: 0716704730.

Martínez-González, E., J. M. Diego, P. Vielva, and J. Silk (Nov. 2003). "Cosmic microwave background power spectrum estimation and map reconstruction with the expectation-maximization algorithm". In: *Mon. Not. Roy. Astron. Soc.* 345.4, pp. 1101–1109. DOI: 10.1046/j.1365-2966.2003.06885.x. arXiv: astro-ph/0302094.

Mather, J. C., E. S. Cheng, J. Eplee R. E., R. B. Isaacman, S. S. Meyer, R. A. Shafer, R. Weiss, E. L. Wright, C. L. Bennett, N. W. Boggess, E. Dwek, S. Gulkis, M. G. Hauser, M. Janssen, T. Kelsall, P. M. Lubin, J. Moseley S. H., T. L. Murdock, R. F. Silverberg, G. F. Smoot, and D. T. Wilkinson (May 1990). "A Preliminary Measurement of the Cosmic Microwave Background Spectrum by the Cosmic Background Explorer (COBE) Satellite". In: *Astrophys. J. Lett.* 354, p. L37. DOI: 10.1086/185717.

Meeks, M. L. and A. E. Lilley (Mar. 1963). "The microwave spectrum of oxygen in the Earth's atmosphere". In: *J. Geophys. Res.* 68.6, pp. 1683–1703. DOI: 10.1029/jz068i006p01683.

Melchiorri, A., P. A. R. Ade, P. de Bernardis, J. J. Bock, J. Borrill, A. Boscaleri, B. P. Crill, G. De Troia, P. Farese, P. G. Ferreira, K. Ganga, G. de Gasperis, M. Giacometti, V. V. Hristov,

A. H. Jaffe, A. E. Lange, S. Masi, P. D. Mauskopf, L. Miglio, C. B. Netterfield, E. Pascale, F. Piacentini, G. Romeo, J. E. Ruhl, and N. Vittorio (June 2000). "A Measurement of Ω from the North American Test Flight of Boomerang". In: *Astrophys. J. Lett.* 536.2, pp. L63–L66. DOI: 10.1086/312744. arXiv: astro-ph/9911445.

Melsheimer, C., C. Verdes, S. A. Buehler, C. Emde, P. Eriksson, D. G. Feist, S. Ichizawa, V. O. John, Y. Kasai, G. Kopp, N. Koulev, T. Kuhn, O. Lemke, S. Ochiai, F. Schreier, T. R. Sreerekha, M. Suzuki, C. Takahashi, S. Tsujimaru, and J. Urban (Feb. 2005). "Intercomparison of general purpose clear sky atmospheric radiative transfer models for the millimeter/submillimeter spectral range". In: *Radio Sci.* 40.1. DOI: 10.1029/2004rs003110.

Miller, N. J., D. T. Chuss, T. A. Marriage, E. J. Wollack, J. W. Appel, C. L. Bennett, J. Eimer, T. Essinger-Hileman, D. J. Fixsen, K. Harrington, S. H. Moseley, K. Rostem, E. R. Switzer, and D. J. Watts (Feb. 2016). "Recovery of Large Angular Scale CMB Polarization for Instruments Employing Variable-delay Polarization Modulators". In: *Astrophys. J.* 818.2, 151. DOI: 10.3847/0004-637X/818/2/151. arXiv: 1509.04628 [astro-ph.IM].

Milligan, M., P. Ade, F. Aubin, C. Baccigalupi, C. Bao, J. Borrill, C. Cantalupo, D. Chapman, J. Didier, M. Dobbs, W. Grainger, S. Hanany, S. Hillbrand, J. Hubmayr, P. Hyland, A. Jaffe, B. Johnson, T. Kisner, J. Klein, A. Korotkov, S. Leach, A. Lee, L. Levinson, M. Limon, K. MacDermid, T. Matsumura, A. Miller, E. Pascale, D. Polsgrove, N. Ponthieu, K. Raach, B. Reichborn-Kjennerud, I. Sagiv, H. Tran, G. S. Tucker, Y. Vinokurov, A. Yadav, M. Zaldarriaga, and K. Zilic (July 2010). "Software systems for operation, control, and monitoring of the EBEX instrument". In: *Proc. SPIE*. Ed. by N. M. Radziwill and A. Bridger. Vol. 7740. DOI: 10.1117/12.857583. arXiv: 1006.5256 [astro-ph.IM].

Mohammadi, R. (Oct. 2014). "Evidence for cosmic neutrino background from CMB circular polarization". In: *Eur. Phys. J. C* 74.10, p. 3102. DOI: 10.1140/epjc/s10052-014-3102-1. arXiv: 1312.2199 [astro-ph.CO].

Moon, B., H. V. Jagadish, C. Faloutsos, and J. H. Saltz (Jan. 2001). "Analysis of the clustering properties of the Hilbert space-filling curve". In: *IEEE Trans. Knowl. Data Eng.* 13.1, pp. 124–141. ISSN: 1041-4347. DOI: 10.1109/69.908985.

Münchmeyer, M. and K. M. Smith (May 2019). "Fast Wiener filtering of CMB maps with Neural Networks". In: *arXiv e-prints*. arXiv: 1905.05846 [astro-ph.CO].

Næss, S., M. Madhavacheril, and M. Hasselfield (Feb. 2021). *Pixell: Rectangular pixel map manipulation and harmonic analysis library*. ASCL: 2102.003.

Nagy, J. M., P. A. R. Ade, M. Amiri, S. J. Benton, A. S. Bergman, R. Bihary, J. J. Bock, J. R. Bond, S. A. Bryan, H. C. Chiang, C. R. Contaldi, O. Doré, A. J. Duivenvoorden, H. K. Eriksen, M. Farhang, J. P. Filippini, L. M. Fissel, A. A. Fraisse, K. Freese, M. Galloway, A. E. Gambrel, N. N. Gandilo, K. Ganga, J. E. Gudmundsson, M. Halpern, J. Hartley, M. Hasselfield, G. Hilton, W. Holmes, V. V. Hristov, Z. Huang, K. D. Irwin, W. C. Jones, C. L. Kuo, Z. D. Kermish, S. Li, P. V. Mason, K. Megerian, L. Moncelsi, T. A. Morford, C. B. Netterfield, M. Nolta, I. L. Padilla, B. Racine, A. S. Rahlin, C. Reintsema, J. E. Ruhl, M. C. Runyan, T. M. Ruud, J. A. Shariff, J. D. Soler, X. Song, A. Trangsrud, C. Tucker, R. S. Tucker, A. D. Turner, J. F. V. D. List, A. C. Weber, I. K. Wehus, D. V. Wiebe, and E. Y. Young (Aug. 2017). "A New Limit on CMB Circular Polarization from SPIDER".

In: *Astrophys. J.* 844.2, p. 151. DOI: 10.3847/1538-4357/aa7cfd. arXiv: 1704.00215 [astro-ph.CO].

National Institute of Standards and Technology (Aug. 2015). *FIPS PUB 180-4: Secure Hash Standard*. Gaithersburg, MD: National Institute of Standards and Technology. DOI: 10.6028/NIST.FIPS.180-4.

Navas-Guzmán, F., N. Kämpfer, A. Murk, R. Larsson, S. A. Buehler, and P. Eriksson (Apr. 2015). "Zeeman effect in atmospheric O_2 measured by ground-based microwave radiometry". In: *Atmos. Meas. Tech.* 8.4, pp. 1863–1874. DOI: 10.5194/amt-8-1863-2015.

Nørgaard-Nielsen, H. U. (Sept. 2010). "Foreground removal from WMAP 5 yr temperature maps using an MLP neural network". In: *Astron. Astrophys.* 520, A87. DOI: 10.1051/0004-6361/201014288. arXiv: 1010.1634 [astro-ph.CO].

O'Dea, D., A. Challinor, and B. R. Johnson (Apr. 2007). "Systematic errors in cosmic microwave background polarization measurements". In: *Mon. Not. Roy. Astron. Soc.* 376.4, pp. 1767–1783. DOI: 10.1111/j.1365-2966.2007.11558.x. arXiv: astro-ph/0610361.

Olbers, W. (1826). "Über die Durchsichtigkeit des Weltraums". In: *Astronomisches Jahrb.*, pp. 110–121.

Padilla, I. L., J. R. Eimer, Y. Li, G. E. Addison, A. Ali, J. W. Appel, C. L. Bennett, R. Bustos, M. K. Brewer, M. Chan, D. T. Chuss, J. Cleary, J. Couto, S. Dahal, K. Denis, R. Dünner, T. Essinger-Hileman, P. Fluxá, D. Gothe, S. K. Haridas, K. Harrington, J. Iuliano, J. Karakla, T. A. Marriage, N. J. Miller, C. Núñez, L. Parker, M. A. Petroff, R. Reeves, K. Rostem, R. W. Stevens, D. A. N. Valle, D. J. Watts, J. L. Weiland, E. J. Wollack, and Z. Xu (Jan. 2020). "Two-year Cosmology Large Angular Scale Surveyor (CLASS) Observations: A Measurement of Circular Polarization at 40 GHz". In: *Astrophys. J.* 889.2, p. 105. DOI: 10.3847/1538-4357/ab61f8. arXiv: 1911.00391 [astro-ph.CO].

Pardo, J., L. Pagani, M. Gerin, and C. Prigent (Dec. 1995). "Evidence of the Zeeman splitting in the $2_1 \rightarrow 0_1$ rotational transition of the atmospheric $^{16}O^{18}O$ molecule from ground-based measurements". In: *J. Quant. Spectrosc. Radiat. Transf.* 54.6, pp. 931–943. DOI: 10.1016/0022-4073(95)00129-9.

Partridge, R. B., J. Nowakowski, and H. M. Martin (Jan. 1988). "Linear polarized fluctuations in the cosmic microwave background". In: *Nature* 331.6152, pp. 146–147. DOI: 10.1038/331146a0.

Partridge, R. B. and D. T. Wilkinson (Apr. 1967). "Isotropy and Homogeneity of the Universe from Measurements of the Cosmic Microwave Background". In: *Phys. Rev. Lett.* 18.14, pp. 557–559. DOI: 10.1103/PhysRevLett.18.557.

Peano, G. (Mar. 1890). "Sur une courbe, qui remplit toute une aire plane". In: *Math. Ann.* 36.1, pp. 157–160. ISSN: 1432-1807. DOI: 10.1007/BF01199438.

Peebles, P. J. E. (July 1968). "Recombination of the Primeval Plasma". In: *Astrophys. J.* 153, p. 1. DOI: 10.1086/149628.

Penndorf, R. (Mar. 1949). "The vertical distribution of atomic oxygen in the upper atmosphere". In: *J. Geophys. Res.* 54.1, pp. 7–38. DOI: 10.1029/jz054i001p00007.

Penzias, A. A. and R. W. Wilson (July 1965). "A Measurement of Excess Antenna Temperature at 4080 Mc/s." In: *Astrophys. J.* 142, pp. 419–421. DOI: 10.1086/148307.

Perlmutter, S., S. Gabi, G. Goldhaber, A. Goobar, D. E. Groom, I. M. Hook, A. G. Kim, M. Y. Kim, J. C. Lee, R. Pain, C. R. Pennypacker, I. A. Small, R. S. Ellis, R. G. McMahon, B. J. Boyle, P. S. Bunclark, D. Carter, M. J. Irwin, K. Glazebrook, H. J. M. Newberg, A. V. Filippenko, T. Matheson, M. Dopita, and W. J. Couch (July 1997). "Measurements of the Cosmological Parameters Ω and Λ from the First Seven Supernovae at $z \geq 0.35$". In: *Astrophys. J.* 483.2, pp. 565–581. DOI: 10.1086/304265. arXiv: astro-ph/9608192.

Perraudin, N., M. Defferrard, T. Kacprzak, and R. Sgier (Apr. 2019). "DeepSphere: Efficient spherical convolutional neural network with HEALPix sampling for cosmological applications". In: *Astron. Comput.* 27, pp. 130–146. DOI: 10.1016/j.ascom.2019.03.004. arXiv: 1810.12186 [astro-ph.CO].

Petroff, M. (July 2018). *3D-printable broadband millimeter wave absorber geometries*. DOI: 10.5281/zenodo.1322839. URL: https://doi.org/10.5281/zenodo.1322839.

Petroff, M., J. Appel, K. Rostem, C. L. Bennett, J. Eimer, T. Marriage, J. Ramirez, and E. J. Wollack (Feb. 2019). "A 3D-printed broadband millimeter wave absorber". In: *Rev. Sci. Instrum.* 90.2, 024701. DOI: 10.1063/1.5050781. arXiv: 1808.00820 [astro-ph.IM].

Petroff, M. A. (Nov. 2019). *Polarized millimeter-wave atmospheric emission model*. Version 1.0. DOI: 10.5281/zenodo.3526652.

— (Apr. 2020). *Supplement to* Full-sky Cosmic Microwave Background Foreground Cleaning Using Machine Learning. DOI: 10.5281/zenodo.3963028.

Petroff, M. A., G. E. Addison, C. L. Bennett, and J. L. Weiland (Nov. 2020a). "Full-sky Cosmic Microwave Background Foreground Cleaning Using Machine Learning". In: *Astrophys. J.* 903.2, 104. DOI: 10.3847/1538-4357/abb9a7. arXiv: 2004.11507 [astro-ph.CO].

Petroff, M. A., J. W. Appel, C. L. Bennett, M. K. Brewer, M. Chan, D. T. Chuss, J. Cleary, J. Denes Couto, S. Dahal, J. R. Eimer, T. Essinger-Hileman, P. Fluxá Rojas, K. Harrington, J. Iuliano, T. A. Marriage, N. J. Miller, D. Augusto Nunes Valle, D. J. Watts, and Z. Xu (Dec. 2020b). "Control and systems software for the Cosmology Large Angular Scale Surveyor (CLASS)". In: *Proc. SPIE*. Vol. 11452. DOI: 10.1117/12.2561609. arXiv: 2012.08433 [astro-ph.IM].

Petroff, M. A., J. R. Eimer, K. Harrington, A. Ali, J. W. Appel, C. L. Bennett, M. K. Brewer, R. Bustos, M. Chan, D. T. Chuss, J. Cleary, J. D. Couto, S. Dahal, R. Dünner, T. Essinger-Hileman, P. F. Rojas, D. Gothe, J. Iuliano, T. A. Marriage, N. J. Miller, C. Núñez, I. L. Padilla, L. Parker, R. Reeves, K. Rostem, D. A. Nunes Valle, D. J. Watts, J. L. Weiland, E. J. Wollack, and Z. Xu (Feb. 2020c). "Two-year Cosmology Large Angular Scale Surveyor (CLASS) Observations: A First Detection of Atmospheric Circular Polarization at Q band". In: *Astrophys. J.* 889.2, 120. DOI: 10.3847/1538-4357/ab64e2. arXiv: 1911.01016 [astro-ph.IM].

Picone, J. M., A. E. Hedin, D. P. Drob, and A. C. Aikin (Dec. 2002). "NRLMSISE-00 empirical model of the atmosphere: Statistical comparisons and scientific issues". In: *J. Geophys. Res.* 107.A12, SIA 15–1–SIA 15–16. DOI: 10.1029/2002ja009430.

Pietranera, L., S. A. Buehler, P. G. Calisse, C. Emde, D. Hayton, V. Oommen John, B. Maffei, L. Piccirillo, G. Pisano, G. Savini, and T. R. Sreerekha (Apr. 2007). "Observing cosmic microwave background polarization through ice". In: *Mon. Not. Roy. Astron. Soc.* 376.2, pp. 645–650. DOI: 10.1111/j.1365-2966.2007.11464.x. arXiv: astro-ph/0611678.

Planck Collaboration XXX (2014). "Planck 2013 results. XXX. Cosmic infrared background measurements and implications for star formation". In: *Astron. Astrophys.* 571, A30. DOI: 10.1051/0004-6361/201322093. arXiv: 1309.0382.

Planck Collaboration I (2016). "Planck 2015 results. I. Overview of products and results". In: *Astron. Astrophys.* 594, A1. DOI: 10.1051/0004-6361/201527101. arXiv: 1502.01582.

Planck Collaboration XXV (2016). "Planck 2015 results. XXV. Diffuse, low-frequency Galactic foregrounds". In: *Astron. Astrophys.* 594, A25. DOI: 10.1051/0004-6361/201526803. arXiv: 1506.06660.

Planck Collaboration I (Sept. 2020). "Planck 2018 results. I. Overview and the cosmological legacy of Planck". In: *Astron. Astrophys.* 641, A1. DOI: 10.1051/0004-6361/201833880. arXiv: 1807.06205 [astro-ph.CO].

Planck Collaboration II (Sept. 2020). "Planck 2018 results. II. Low Frequency Instrument data processing". In: *Astron. Astrophys.* 641, A2. DOI: 10.1051/0004-6361/201833293. arXiv: 1807.06206 [astro-ph.CO].

Planck Collaboration III (Sept. 2020). "Planck 2018 results. III. High Frequency Instrument data processing and frequency maps". In: *Astron. Astrophys.* 641, A3. DOI: 10.1051/0004-6361/201832909. arXiv: 1807.06207 [astro-ph.CO].

Planck Collaboration IV (Sept. 2020). "Planck 2018 results. IV. Diffuse component separation". In: *Astron. Astrophys.* 641, A4. DOI: 10.1051/0004-6361/201833881. arXiv: 1807.06208 [astro-ph.CO].

Planck Collaboration V (Sept. 2020). "Planck 2018 results. V. CMB power spectra and likelihoods". In: *Astron. Astrophys.* 641, A5. DOI: 10.1051/0004-6361/201936386. arXiv: 1907.12875 [astro-ph.CO].

Planck Collaboration VI (2020). "Planck 2018 results. VI. Cosmological parameters". In: *Astron. Astrophys.* 641, A6. DOI: 10.1051/0004-6361/201833910. arXiv: 1807.06209.

Puglisi, G., G. Fabbian, and C. Baccigalupi (Aug. 2017). "A 3D model for carbon monoxide molecular line emission as a potential cosmic microwave background polarization contaminant". In: *Mon. Not. Roy. Astron. Soc.* 469.3, pp. 2982–2996. DOI: 10.1093/mnras/stx1029. arXiv: 1701.07856 [astro-ph.CO].

Puglisi, G. and X. Bai (Dec. 2020). "Inpainting Galactic Foreground Intensity and Polarization Maps Using Convolutional Neural Networks". In: *Astrophys. J.* 905.2, 143. DOI: 10.3847/1538-4357/abc47c. arXiv: 2003.13691 [astro-ph.CO].

Randa, J., J. Lahtinen, A. Camps, A. J. Gasiewsk, M. Hallikainen, D. M. Le Vine, M. Martin-Neira, J. Piepmeier, P. W. Rosenkranz, C. S. Ruf, J. Shiue, and N. Skou (Aug. 2008). *Recommended Terminology for Microwave Radiometry*. Technical Note 1551. NIST. DOI: 10.6028/NIST.TN.1551. URL: https://www.nist.gov/document/nist-tn-1551pdf.

Reddi, S. J., S. Kale, and S. Kumar (2018). "On the Convergence of Adam and Beyond". In: *Int. Conf. Learn. Represent.* arXiv: 1904.09237 [cs.LG]. URL: https://openreview.net/forum?id=ryQu7f-RZ.

Remazeilles, M., J. Delabrouille, and J.-F. Cardoso (Nov. 2011). "Foreground component separation with generalized Internal Linear Combination". In: *Mon. Not. Roy. Astron. Soc.* 418.1, pp. 467–476. DOI: 10.1111/j.1365-2966.2011.19497.x. arXiv: 1103.1166 [astro-ph.CO].

Ren, J. and J. Y. Yin (2018). "3D-Printed Low-Cost Dielectric-Resonator-Based Ultra-Broadband Microwave Absorber Using Carbon-Loaded Acrylonitrile Butadiene Styrene Polymer". In: *Materials* 11.7, p. 1249. ISSN: 1996-1944. DOI: 10.3390/ma11071249.

Riess, A. G., S. Casertano, W. Yuan, J. B. Bowers, L. Macri, J. C. Zinn, and D. Scolnic (Feb. 2021). "Cosmic Distances Calibrated to 1% Precision with Gaia EDR3 Parallaxes and Hubble Space Telescope Photometry of 75 Milky Way Cepheids Confirm Tension with ΛCDM". In: *Astrophys. J. Lett.* 908.1, L6. DOI: 10.3847/2041-8213/abdbaf. arXiv: 2012.08534 [astro-ph.CO].

Riess, A. G., A. V. Filippenko, P. Challis, A. Clocchiatti, A. Diercks, P. M. Garnavich, R. L. Gilliland, C. J. Hogan, S. Jha, R. P. Kirshner, B. Leibundgut, M. M. Phillips, D. Reiss, B. P. Schmidt, R. A. Schommer, R. C. Smith, J. Spyromilio, C. Stubbs, N. B. Suntzeff, and J. Tonry (Sept. 1998). "Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant". In: *Astron. J.* 116.3, pp. 1009–1038. DOI: 10.1086/300499. arXiv: astro-ph/9805201.

Rindler, W. (Jan. 1956). "Visual horizons in world models". In: *Mon. Not. Roy. Astron. Soc.* 116, p. 662. DOI: 10.1093/mnras/116.6.662.

Robertson, H. P. (Jan. 1933). "Relativistic Cosmology". In: *Rev. Mod. Phys.* 5.1, pp. 62–90. DOI: 10.1103/RevModPhys.5.62.

— (Nov. 1935). "Kinematics and World-Structure". In: *Astrophys. J.* 82, p. 284. DOI: 10.1086/143681.

Ronneberger, O., P. Fischer, and T. Brox (2015). "U-Net: Convolutional Networks for Biomedical Image Segmentation". In: *Med. Image Comput. Comput.-Assist. Interv.* Springer International Publishing, pp. 234–241. DOI: 10.1007/978-3-319-24574-4_28. arXiv: 1505.04597 [cs.CV].

Rosenkranz, P. W. and D. H. Staelin (Sept. 1988). "Polarized thermal microwave emission from oxygen in the mesosphere". In: *Radio Sci.* 23.5, pp. 721–729. DOI: 10.1029/rs023i005p00721.

Rosenkranz, P. W. (July 1998). "Water vapor microwave continuum absorption: A comparison of measurements and models". In: *Radio Sci.* 33.4, pp. 919–928. DOI: 10.1029/98rs01182.

Russell, S. J. and P. Norvig (2010). *Artificial Intelligence: A Modern Approach*. 3rd ed. Prentice Hall series in artificial intelligence. Upper Saddle River: Prentice Hall. ISBN: 978-0-13-604259-4.

Sachs, R. K. and A. M. Wolfe (Jan. 1967). "Perturbations of a Cosmological Model and Angular Variations of the Microwave Background". In: *Astrophys. J.* 147, p. 73. DOI: 10.1086/148982.

Sagan, H. (1994). *Space-Filling Curves*. New York: Springer-Verlag. ISBN: 0387942653.

Sanz-Izquierdo, B. and E. A. Parker (Dec. 2014). "3-D Printing of Elements in Frequency Selective Arrays". In: *IEEE Trans. Antennas Propag.* 62.12, pp. 6060–6066. ISSN: 0018-926X. DOI: 10.1109/TAP.2014.2359470.

Sawyer, R. (Jan. 2015). "Photon-photon interactions as a source of cosmic microwave background circular polarization". In: *Phys. Rev. D* 91.2. DOI: 10.1103/physrevd.91.021301.

Schwartz, M., W. Read, and W. V. Snyder (May 2006). "EOS MLS forward model polarized radiative transfer for Zeeman-split oxygen lines". In: *IEEE Trans. Geosci. Remote Sensing* 44.5, pp. 1182–1191. DOI: 10.1109/tgrs.2005.862267.

Seljak, U. and M. Zaldarriaga (Mar. 1997). "Signature of Gravity Waves in the Polarization of the Microwave Background". In: *Phys. Rev. Lett.* 78.11, pp. 2054–2057. DOI: 10.1103/PhysRevLett.78.2054. arXiv: astro-ph/9609169.

Sihvola, A. (1999). *Electromagnetic mixing formulas and applications*. Electromagnetic Waves. Institution of Electrical Engineers. ISBN: 9780852967720.

Smoot, G. F., C. L. Bennett, A. Kogut, E. L. Wright, J. Aymon, N. W. Boggess, E. S. Cheng, G. de Amici, S. Gulkis, M. G. Hauser, G. Hinshaw, P. D. Jackson, M. Janssen, E. Kaita, T. Kelsall, P. Keegstra, C. Lineweaver, K. Loewenstein, P. Lubin, J. Mather, S. S. Meyer, S. H. Moseley, T. Murdock, L. Rokke, R. F. Silverberg, L. Tenorio, R. Weiss, and D. T. Wilkinson (Sept. 1992). "Structure in the COBE Differential Microwave Radiometer First-Year Maps". In: *Astrophys. J. Lett.* 396, p. L1. DOI: 10.1086/186504.

Spergel, D. N., L. Verde, H. V. Peiris, E. Komatsu, M. R. Nolta, C. L. Bennett, M. Halpern, G. Hinshaw, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, L. Page, G. S. Tucker, J. L. Weiland, E. Wollack, and E. L. Wright (Sept. 2003). "First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Determination of Cosmological Parameters". In: *Astrophys. J. Suppl. Ser.* 148.1, pp. 175–194. DOI: 10.1086/377226. arXiv: astro-ph/0302209.

Spinelli, S., G. Fabbian, A. Tartari, M. Zannoni, and M. Gervasi (May 2011). "A template of atmospheric O2 circularly polarized emission for cosmic microwave background experiments". In: *Mon. Not. Roy. Astron. Soc.* 414.4, pp. 3272–3280. DOI: 10.1111/j.1365-2966.2011.18625.x. arXiv: 1103.0160 [astro-ph.IM].

Srivastava, N., G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov (2014). "Dropout: A Simple Way to Prevent Neural Networks from Overfitting". In: *J. Mach. Learn. Res.* 15, pp. 1929–1958. URL: http://jmlr.org/papers/v15/srivastava14a.html.

Starobinskiĭ, A. A. (Dec. 1979). "Spectrum of relict gravitational radiation and the early state of the universe". In: *Sov. J. Exp. Theor. Phys. Lett.* 30, p. 682. BIBCODE: 1979JETPL..30..682S.

Swetz, D. S., P. A. R. Ade, M. Amiri, J. W. Appel, E. S. Battistelli, B. Burger, J. Chervenak, M. J. Devlin, S. R. Dicker, W. B. Doriese, R. Dünner, T. Essinger-Hileman, R. P. Fisher, J. W. Fowler, M. Halpern, M. Hasselfield, G. C. Hilton, A. D. Hincks, K. D. Irwin, N. Jarosik, M. Kaul, J. Klein, J. M. Lau, M. Limon, T. A. Marriage, D. Marsden, K. Martocci, P. Mauskopf, H. Moseley, C. B. Netterfield, M. D. Niemack, M. R. Nolta, L. A. Page, L. Parker, S. T. Staggs, O. Stryzak, E. R. Switzer, R. Thornton, C. Tucker, E. Wollack, and Y. Zhao (June 2011). "Overview of the Atacama Cosmology Telescope: Receiver, Instrumentation, and Telescope Systems". In: *Astrophys. J. Suppl. Ser.* 194.2, 41. DOI: 10.1088/0067-0049/194/2/41. arXiv: 1007.0290 [astro-ph.IM].

Switzer, E. R., C. Allen, M. Amiri, J. W. Appel, E. S. Battistelli, B. Burger, J. A. Chervenak, A. J. Dahlen, S. Das, M. J. Devlin, S. R. Dicker, W. B. Doriese, R. Dünner, T. Essinger-Hileman, X. Gao, M. Halpern, M. Hasselfield, G. C. Hilton, A. D. Hincks, K. D. Irwin, S. Knotek, R. P. Fisher, J. W. Fowler, N. Jarosik, M. Kaul, J. Klein, J. M. Lau, M. Limon, R. H. Lupton, T. A. Marriage, K. L. Martocci, S. H. Moseley, C. B. Netterfield, M. D. Niemack, M. R. Nolta, L. Page, L. P. Parker, B. A. Reid, C. D. Reintsema, A. J. Sederberg, J. L. Sievers, D. N. Spergel, S. T. Staggs, O. R. Stryzak, D. S. Swetz, R. J. Thornton, E. J. Wollack, and Y. Zhao (July 2008). "Systems and control software for the Atacama Cosmology Telescope". In: *Proc. SPIE*. Ed. by A. Bridger and N. M. Radziwill. Vol. 7019. DOI: 10.1117/12.790006.

Takakura, S., M. A. O. Aguilar-Faúndez, Y. Akiba, K. Arnold, C. Baccigalupi, D. Barron, D. Beck, F. Bianchini, D. Boettger, J. Borrill, K. Cheung, Y. Chinone, T. Elleflot, J. Errard, G. Fabbian, C. Feng, N. Goeckner-Wald, T. Hamada, M. Hasegawa, M. Hazumi, L. Howe, D. Kaneko, N. Katayama, B. Keating, R. Keskitalo, T. Kisner, N. Krachmalnicoff, A. Kusaka, A. T. Lee, L. N. Lowry, F. T. Matsuda, A. J. May, Y. Minami, M. Navaroli, H. Nishino, L. Piccirillo, D. Poletti, G. Puglisi, C. L. Reichardt, Y. Segawa, M. Silva-Feaver, P. Siritanasak, A. Suzuki, O. Tajima, S. Takatori, D. Tanabe, G. P. Teply, and C. Tsai (Jan. 2019). "Measurements of Tropospheric Ice Clouds with a Ground-based CMB Polarization Experiment, POLARBEAR". In: *Astrophys. J.* 870.2, 102. DOI: 10.3847/1538-4357/aaf381. arXiv: 1809.06556 [astro-ph.IM].

Thielbeer, B. (1997). "Mit Codierung versehenes Markierungssystem sowie mit einer Codierung versehene Markierung". German pat. DE19733466A1.

Thomson, J. J. (1903). *Conduction of electricity through gases*. Cambridge: Cambridge University Press, pp. 268–273.

Thorne, B., J. Dunkley, D. Alonso, and S. Næss (May 2017). "The Python Sky Model: software for simulating the Galactic microwave sky". In: *Mon. Not. Roy. Astron. Soc.* 469.3, pp. 2821–2833. DOI: 10.1093/mnras/stx949. arXiv: 1608.02841 [astro-ph.CO]. URL: https://doi.org/10.1093/mnras/stx949.

Tizchang, S., S. Batebi, M. Haghighat, and R. Mohammadi (Aug. 2016). "Cosmic microwave background polarization in non-commutative space-time". In: *Eur. Phys. J. C* 76.9, p. 478. DOI: 10.1140/epjc/s10052-016-4312-5.

Torbet, E., M. J. Devlin, W. B. Dorwart, T. Herbig, A. D. Miller, M. R. Nolta, L. Page, J. Puchalla, and H. T. Tran (Aug. 1999). "A Measurement of the Angular Power Spectrum

of the Microwave Background Made from the High Chilean Andes". In: *Astrophys. J. Lett.* 521.2, pp. L79–L82. DOI: 10.1086/312197. arXiv: astro-ph/9905100.

Townes, C. H. and A. L. Schawlow (1955). *Microwave Spectroscopy*. New York: McGraw-Hill.

Tretyakov, M. (Oct. 2016). "Spectroscopy underlying microwave remote sensing of atmospheric water vapor". In: *J. Mol. Spectrosc.* 328, pp. 7–26. DOI: 10.1016/j.jms.2016.06.006.

Tufte, E. R. (2006). *Beautiful Evidence*. Cheshire, CT: Graphics Press. ISBN: 0961392177.

Vafaei Sadr, A. and F. Farsian (Apr. 2020). "Inpainting via Generative Adversarial Networks for CMB data analysis". In: *arXiv e-prints*. arXiv: 2004.04177 [astro-ph.CO].

Van Vleck, J. H. and V. F. Weisskopf (Apr. 1945). "On the Shape of Collision-Broadened Lines". In: *Rev. Mod. Phys.* 17.2-3, pp. 227–236. DOI: 10.1103/revmodphys.17.227.

Varghese, P. L. and R. K. Hanson (July 1984). "Collisional narrowing effects on spectral line shapes measured at high resolution". In: *Appl. Optics* 23.14, p. 2376. DOI: 10.1364/ao.23.002376.

Virtanen, P., R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, İ. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors (Feb. 2020). "SciPy 1.0: fundamental algorithms for scientific computing in Python". In: *Nat. Methods* 17.3, pp. 261–272. DOI: 10.1038/s41592-019-0686-2. arXiv: 1907.10121 [cs.MS].

Wainwright, J. and S. W. Goode (Oct. 1980). "Some exact inhomogeneous cosmologies with equation of state $p = \gamma\mu$". In: *Phys. Rev. D* 22.8, pp. 1906–1909. DOI: 10.1103/PhysRevD.22.1906.

Walker, A. G. (Jan. 1937). "On Milne's Theory of World-Structure". In: *Proc. London Math. Soc.* 42, pp. 90–127. DOI: 10.1112/plms/s2-42.1.90.

Walt, S. van der, S. C. Colbert, and G. Varoquaux (Mar. 2011). "The NumPy Array: A Structure for Efficient Numerical Computation". In: *IEEE Comput. Sci. Eng.* 13.2, pp. 22–30. DOI: 10.1109/mcse.2011.37. arXiv: 1102.1523 [cs.MS].

Watts, D. J., D. Larson, T. A. Marriage, M. H. Abitbol, J. W. Appel, C. L. Bennett, D. T. Chuss, J. R. Eimer, T. Essinger-Hileman, N. J. Miller, K. Rostem, and E. J. Wollack (Dec. 2015). "Measuring the Largest Angular Scale CMB B-mode Polarization with Galactic Foregrounds on a Cut Sky". In: *Astrophys. J.* 814.2, 103. DOI: 10.1088/0004-637X/814/2/103. arXiv: 1508.00017 [astro-ph.CO].

Watts, D. J., B. Wang, A. Ali, J. W. Appel, C. L. Bennett, D. T. Chuss, S. Dahal, J. R. Eimer, T. Essinger-Hileman, K. Harrington, G. Hinshaw, J. Iuliano, T. A. Marriage, N. J. Miller, I. L. Padilla, L. Parker, M. Petroff, K. Rostem, E. J. Wollack, and Z. Xu (Aug. 2018). "A Projected Estimate of the Reionization Optical Depth Using the CLASS Experiment's Sample Variance Limited E-mode Measurement". In: *Astrophys. J.* 863.2, 121. DOI: 10.3847/1538-4357/aad283. arXiv: 1801.01481 [astro-ph.CO].

Weinberg, S. (2008). *Cosmology*. Oxford / New York: Oxford University Press. ISBN: 978-0-19-852682-7.

Wiebe, D. V., C. B. Netterfield, and T. S. Kisner (Dec. 2015). *GetData: A filesystem-based, column-oriented database format for time-ordered binary data*. ASCL: 1512.002.

Wollack, E. (Nov. 1995). *TCHEB$_X$: Homogenous Stepped Waveguide Transformers*. Electronics Division Technical Note 176. National Radio Astronomy Observatory. URL: https://library.nrao.edu/public/memos/edtn/EDTN_176.pdf.

Wollack, E. J., D. J. Fixsen, R. Henry, A. Kogut, M. Limon, and P. Mirel (2008). "Electromagnetic and Thermal Properties of a Conductively Loaded Epoxy". In: *Int. J. Infrared Millimeter Waves* 29.1, pp. 51–61. ISSN: 1572-9559. DOI: 10.1007/s10762-007-9299-4.

Wollack, E. J., R. E. Kinzer, and S. A. Rinehart (Apr. 2014). "A cryogenic infrared calibration target". In: *Rev. Sci. Instrum.* 85.4, 044707. DOI: 10.1063/1.4871108. arXiv: 1403.6667 [physics.ins-det].

Xu, Z., M. K. Brewer, P. F. Rojas, Y. Li, K. Osumi, B. Pradenas, A. Ali, J. W. Appel, C. L. Bennett, R. Bustos, M. Chan, D. T. Chuss, J. Cleary, J. D. Couto, S. Dahal, R. Datta, K. L. Denis, R. Dünner, J. R. Eimer, T. Essinger-Hileman, D. Gothe, K. Harrington, J. Iuliano, J. Karakla, T. A. Marriage, N. J. Miller, C. Núñez, I. L. Padilla, L. Parker, M. A. Petroff, R. Reeves, K. Rostem, D. A. Nunes Valle, D. J. Watts, J. L. Weiland, E. J. Wollack, and CLASS Collaboration (Mar. 2020). "Two-year Cosmology Large Angular Scale Surveyor (CLASS) Observations: 40 GHz Telescope Pointing, Beam Profile, Window Function, and Polarization Performance". In: *Astrophys. J.* 891.2, 134. DOI: 10.3847/1538-4357/ab76c2. arXiv: 1911.04499 [astro-ph.IM].

Yadav, J., S. Bharadwaj, B. Pandey, and T. R. Seshadri (Dec. 2005). "Testing homogeneity on large scales in the Sloan Digital Sky Survey Data Release One". In: *Mon. Not. Roy. Astron. Soc.* 364.2, pp. 601–606. DOI: 10.1111/j.1365-2966.2005.09578.x. arXiv: astro-ph/0504315.

Yi, K., Y. Guo, Y. Fan, J. Hamann, and Y. G. Wang (Jan. 2020). "CosmoVAE: Variational Autoencoder for CMB Image Inpainting". In: *arXiv e-prints*. arXiv: 2001.11651 [eess.IV].

Zaldarriaga, M. (Nov. 2001). "Nature of the E-B decomposition of CMB polarization". In: *Phys. Rev. D* 64.10, p. 103001. DOI: 10.1103/PhysRevD.64.103001. arXiv: astro-ph/0106174.

Zarei, M., E. Bavarsad, M. Haghighat, R. Mohammadi, I. Motie, and Z. Rezaei (Apr. 2010). "Generation of circular polarization of the CMB". In: *Phys. Rev. D* 81.8. DOI: 10.1103/physrevd.81.084035. arXiv: 0912.2993 [hep-th].

Zeldovich, Y. B. and R. A. Sunyaev (July 1969). "The Interaction of Matter and Radiation in a Hot-Model Universe". In: *Astrophys. Space Sci.* 4.3, pp. 301–316. DOI: 10.1007/BF00661821.

Zonca, A., L. Singer, D. Lenz, M. Reinecke, C. Rosset, E. Hivon, and K. Gorski (Mar. 2019). "healpy: equal area pixelization and spherical harmonics transforms for data on the sphere in Python". In: *J. Open Source Softw.* 4.35, p. 1298. DOI: 10.21105/joss.01298.

Vita

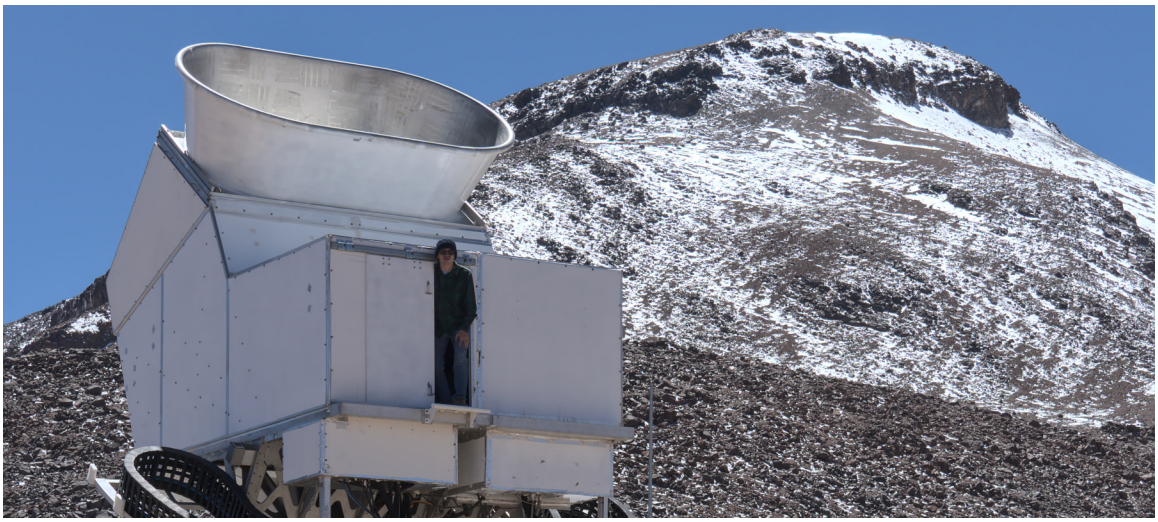

Matthew Aidan Petroff (ORCID: 0000-0002-4436-4215) was born and raised in Connecticut, attending St. Bridget School in Cheshire through the eighth grade and The Taft School in Watertown for high school, with his twin brother, Christopher. He then matriculated to Johns Hopkins University in Baltimore, earning a Bachelor of Science degree in Physics with a second major in Computer Science and a minor in Mathematics in 2015. As an undergraduate, he began working on the Cosmology Large Angular Scale Surveyor (CLASS) project, and he stayed at Johns Hopkins University to continue working on the CLASS project and studying the cosmic microwave background (CMB) for his Ph.D. research. He spent one year in Chile—over the course of ten trips—doing field work for CLASS. For the future, he has accepted a post-doctoral position at Harvard University to continue to work on CMB instrumentation and analysis.

Selected publications

Petroff, M., J. Appel, K. Rostem, C. L. Bennett, J. Eimer, T. Marriage, J. Ramirez, and E. J. Wollack (Feb. 2019). "A 3D-printed broadband millimeter wave absorber". In: *Rev. Sci. Instrum.* 90.2, 024701. DOI: 10.1063/1.5050781. arXiv: 1808.00820 [astro-ph.IM].

Petroff, M. A. (Aug. 2019). "Pannellum: a lightweight web-based panorama viewer". In: *J. Open Source Softw.* 4.40, 1628. DOI: 10.21105/joss.01628.

Petroff, M. A., J. R. Eimer, K. Harrington, A. Ali, J. W. Appel, C. L. Bennett, M. K. Brewer, R. Bustos, M. Chan, D. T. Chuss, J. Cleary, J. Denes Couto, S. Dahal, R. Dünner, T. Essinger-Hileman, P. Fluxá Rojas, D. Gothe, J. Iuliano, T. A. Marriage, N. J. Miller, C. Núñez, I. L. Padilla, L. Parker, R. Reeves, K. Rostem, D. A. Nunes Valle, D. J. Watts, J. L. Weiland, E. J. Wollack, and Z. Xu (Jan. 2020a). "Two-year Cosmology Large Angular Scale Surveyor (CLASS) Observations: A First Detection of Atmospheric Circular Polarization at Q band". In: *Astrophys. J.* 889.2, 120. DOI: 10.3847/1538-4357/ab64e2. arXiv: 1911.01016 [astro-ph.IM].

Petroff, M. A., G. E. Addison, C. L. Bennett, and J. L. Weiland (Nov. 2020b). "Full-sky Cosmic Microwave Background Foreground Cleaning Using Machine Learning". In: *Astrophys. J.* 903.2, 104. DOI: 10.3847/1538-4357/abb9a7. arXiv: 2004.11507 [astro-ph.CO].

Petroff, M. A., J. W. Appel, C. L. Bennett, M. K. Brewer, M. Chan, D. T. Chuss, J. Cleary, J. Denes Couto, S. Dahal, J. R. Eimer, T. Essinger-Hileman, P. Fluxá Rojas, K. Harrington, J. Iuliano, T. A. Marriage, N. J. Miller, D. A. Nunes Valle, D. J. Watts, and Z. Xu (Dec. 2020c). "Control and systems software for the Cosmology Large Angular Scale Surveyor (CLASS)". In: *Proc. SPIE* 11452, 114521O. DOI: 10.1117/12.2561609. arXiv: 2012.08433 [astro-ph.IM].

Petroff, M. A. (July 2021a). "A Square Equal-area Map Projection with Low Angular Distortion, Minimal Cusps, and Closed-form Solutions". In: *ACM Trans. Spat. Algorithms Syst.* 7.4, 21. DOI: 10.1145/3460521. arXiv: 2008.13670 [cs.GR].

Petroff, M. A. (July 2021b). "Accessible Color Cycles for Data Visualization". In: *arXiv e-prints*. arXiv: 2107.02270 [cs.GR].

This document is primarily set in 12 pt Libertinus Serif, with Libertinus Sans used for sans-serif text and Inconsolata used for monospaced text. It was typeset using pdfTeX and the LaTeX packages included in TeX Live 2019, specifically those distributed with Ubuntu 20.04.