Yale University

EliScholar – A Digital Platform for Scholarly Publishing at Yale

Yale Graduate School of Arts and Sciences Dissertations

Spring 2022

Calibration Instrumentation for the Hydrogen Intensity and Real-Time Analysis eXperiment

Emily Rita Kuhn
Yale University Graduate School of Arts and Sciences, emily.kuhn318@gmail.com

Follow this and additional works at: https://elischolar.library.yale.edu/gsas_dissertations

Recommended Citation

Kuhn, Emily Rita, "Calibration Instrumentation for the Hydrogen Intensity and Real-Time Analysis eXperiment" (2022). *Yale Graduate School of Arts and Sciences Dissertations*. 616.
https://elischolar.library.yale.edu/gsas_dissertations/616

This Dissertation is brought to you for free and open access by EliScholar – A Digital Platform for Scholarly Publishing at Yale. It has been accepted for inclusion in Yale Graduate School of Arts and Sciences Dissertations by an authorized administrator of EliScholar – A Digital Platform for Scholarly Publishing at Yale. For more information, please contact elischolar@yale.edu.

Abstract

Calibration Instrumentation for the Hydrogen Intensity and Real-time Analysis eXperiment

Emily R. Kuhn

2022

The Hydrogen Intensity and Real-time Analysis eXperiment (HIRAX) is a 21 cm neutral hydrogen intensity mapping experiment to be deployed in the Karoo Desert in South Africa. It aims to improve constraints on the dark energy equation of state through measurements of large-scale structure at high redshift, while doubling as a state-of-the-art fast radio burst (FRB) detector. This dissertation focuses on three aspects of the HIRAX instrument characterization: (1) measuring the antenna beam pattern, using a variety of specialized facilities, (2) optimizing the signal-to-noise of antennas, through the design and implementation of a custom test-bed for determining the noise temperature of radio antennas operating between 400-800MHz, and (3) mapping the HIRAX telescope beam pattern with a custom drone calibration system. The work described is critical to HIRAX's development, both by informing final antenna design and providing the tools to generate beam maps that will factor into all cosmological analysis.

Calibration Instrumentation for the Hydrogen Intensity and Real-time Analysis eXperiment

A Dissertation
Presented to the Faculty of the Graduate School
of
Yale University
in Candidacy for the Degree of
Doctor of Philosophy

by
Emily R. Kuhn

Dissertation Director: Laura B. Newburgh

May, 2022

Copyright © 2021 by Emily R. Kuhn

All rights reserved.

Contents

1 Background: the promise of 21cm cosmology 1

1.1 Introduction to cosmology 2

1.2 Definitions and common notation 3

1.3 ΛCDM cosmology 7

1.4 Dark energy 8

1.5 Baryon acoustic oscillations 11

1.6 The 21cm neutral hydrogen line 12

1.7 Intensity mapping 14

1.8 21cm experiments: past, present and future 15

1.9 Thesis organization 18

2 The Hydrogen Intensity and Real-time Analysis eXperiment 20

2.1 Measurement techniques in radio astronomy 20

2.1.1 Radio interferometry overview 24

2.1.2 Redundant calibration 27

2.1.3 The (u,v) plane 28

2.2 HIRAX overview 29

2.2.1 HIRAX location 31

2.3 The HIRAX science case 32

2.3.1 Dark energy 33

2.3.2 FRBs and radio transients 35

2.3.3 Other science cases 37

2.4 The HIRAX instrument 39

2.4.1 Instrument design drivers 40

2.4.2 Dish design and manufacturing 43

2.4.3 Antenna design and assembly 49

2.4.4 RF chain 50

2.4.5 Correlator 51

2.5 Summary and outlook 53

3 Antenna beam measurements 55

3.1 Introduction to antennas 55

3.2 Introduction to antenna range measurements 60

3.2.1 Types of antenna ranges 62

3.2.2 The JPL antenna range 64

3.2.3 The NCSU anechoic chamber 66

3.2.4 The GBO antenna range 68

3.3 HIRAX feed measurements 69

3.4 CHIME feed measurements 76

3.5 Biconical antenna measurements 77

3.6 Drone transmitter measurements 83

3.7 Chapter summary 85

4 Design and implementation of a noise temperature measurement system 88

4.1 The Y-factor method 89

4.2 Experiment design 90

4.2.1 Simulations 91

4.2.2 Cryogenic system construction 94
4.2.3 Faraday cage construction 99
4.3 System verification measurements 101
4.3.1 RF transparency tests of insulation materials 102
4.3.2 Return loss with RF absorber installed 103
4.3.3 Blackbody spectrum comparison 105
4.4 Systematics and uncertainties 106
4.4.1 Statistical uncertainties 106
4.4.2 Contribution of reflections to uncertainties 108
4.4.3 Estimated contribution of cavity differences to noise temperature calculations 112
4.4.4 Additional sources of error 113
4.5 Cryogenic verification measurements 116
4.5.1 Noise figure meter measurements 117
4.5.2 Passive feed noise temperature 120
4.6 HIRAX antenna results 128
4.6.1 Noise temperature 128
4.6.2 Antenna repeatability 130
4.6.3 In situ assessment of reflections 130
4.7 Summary and future work 135

5 Design of a drone calibration system 139

5.1 Introduction to drone beam calibration 139
5.2 Drone instrument constraints 143
5.2.1 Methodology for finding positional constraints 144
5.2.2 Findings for positional constraints 150
5.3 Drone system overview 153
5.3.1 The DJI Matrice 600 Pro drone 153

5.3.2 RTK GPS integration and verification 154

5.3.3 IMU sensor verification 159

5.3.4 RF chain design 160

5.4 Autopilot software 167

5.5 Travelling with the drone 168

5.6 Radio Frequency Interference (RFI) 169

5.6.1 RFI assessment 170

5.6.2 Mitigation plan 175

5.7 A test bed at Yale's Leitner Family Observatory and Planetarium (LFOP) 176

5.8 Drone flyaway 179

5.8.1 Drone loss and recovery 179

5.8.2 Failure analysis 181

5.8.3 The rebuild 182

5.9 Chapter summary 183

6 Drone calibration measurements of radio telescopes 184

6.1 Drone beam mapping data analysis 185

6.2 Owen's Valley Radio Observatory 189

6.2.1 The DSA instrument 189

6.2.2 Flight path design and considerations 190

6.2.3 Beam mapping results 192

6.2.4 Challenges and lessons learned 194

6.3 Brookhaven National Laboratory 195

6.3.1 The Baryon Mapping eXperiment (BMX) 195

6.3.2 Flight path design and considerations 195

6.3.3 Beam mapping results 198

6.3.4 Challenges and lessons learned 201

6.4 Green Bank Observatory 205

6.4.1 The CHIME FRB Outrigger array 205

6.4.2 Instrument challenges 207

6.4.3 Description of observations 208

6.4.4 Main beam mapping 210

6.4.5 Side lobe mapping 217

6.4.6 Polarization maps 219

6.5 Summary and ongoing work 223

7 Conclusion **225**

List of Figures

1.1 Visualization of the Universe's history, from the Big Bang through Present Day. Figure Credit: NASA / WMAP Science Team. 3

1.2 Constraints on the Hubble Constant H_0 from a range of experiments, spanning two decades and a variety of techniques. Tension between results from the CMB and BAO (ie Planck, WMAP, eBOSS) and from supernovae (indicated as Cepheids + SNIa) is evident. The gray vertical line is located at 67.88 km/sec/Mpc and represents the weighted average of WMAP and Planck data points. Image credit: NASA / LAMBDA Archive Team. 10

1.3 The 21cm power spectrum, with the inlaid plot showing the power spectrum divided by a smoothed power spectrum model to highlight BAO features. The x-axis is a length scale on the sky and the y-axis gives power concentrated at each length scale. In the 21cm power spectrum (shown here from eBOSS Quasars measurements), it is evident that power is concentrated at specific length scales on the sky, appearing as three different wiggles, which correspond to the expected BAO features. Figure credit: SDSS; [15, 16] are references for final results. 12

1.4 21cm brightness temperature at high redshift. This is the measurement targeted by global 21cm experiments. Figure from [26] 16

2.1 A cartoon depicting a two-element interferometer, with elements A and B. At the most basic level, radio interferometers operate by measuring the phase difference between the signals received at A and B and, through knowledge of their physical separation, mapping that phase difference into an angle on the sky. This idea can be extended to N elements of varying baselines. It can be helpful to build redundancy into the baseline separations, to boost signal on these scales and lower computational cost. Figure credit: www.hardhack.org.au/interferometry. 25

2.2 Figure and caption from [61] "(Left:) An example array configuration consisting of 8 antennas, with E and N antenna locations in meters. (Middle:) The u,v sampling that results from this array for a single time (snapshot), at zero hour angle. Each cross represents a baseline. There are N(N-1)/2 = 28 unique baselines, but there are twice this many points due to symmetry. (Right:) The u,v sampling that results after a 12 hour integration. The rotating Earth causes the projected baselines to change with time, tracing out portions of elliptical paths." 29

2.3 (Top left) Map of South Africa, with red stars indicating the primary sites for HIRAX prototyping and deployment; (Top right) The 8-element HIRAX prototype array at Hartebeesthoek Radio Astronomy Observatory, South Africa ; (Bottom) A rendering of the final 1024 dish configuration in the Karoo Desert, South Africa. 30

2.4 Several spectra from the future HIRAX site in the Karoo Desert, taken in 2017 by myself and Prof Jeff Peterson. RFI appears as spikes in the measurement, with higher spikes corresponding to higher power contamination. Thus, it is clear that the Karoo is radio quiet in the HIRAX band, with no visible sources of RFI between 400-800MHz. . 32

2.5 Left: Constraints on dark energy equation of state parameters. Right: Constraints on BAO measurements from a variety of surveys as well as the projected constraints for HIRAX 128 dish prototype and HIRAX full 1024 dish array. Note that the HIRAX prototype will now include 256 dishes. It is evident that the sensitivity of HIRAX in this redshift range is unprecedented, and not volume limited. HIRAX will use BAO as a probe of dark energy. Figures from [62] 34

2.6 Constraints on various cosmological parameters from a 4-year survey of HIRAX-1024 + Planck. HIRAX will place new constraints on several of these parameters (for comparison, see table 2.2) Visualizing constraints in this way also illuminates degeneracies in parameters–the skinnier the oval, the more degenerate parameters are. For BAO measurements in particular, it is difficult to distinguish between a larger expansion (h) and more dark energy (Ω_Λ). This degeneracy can be a feature when cross-correlating with an experiment that is degenerate but in the opposing way, as in if curves are overlapped, the cross-section is particularly small. Figure from Devin Crichton. 36

2.7 The HIRAX survey footprint, illustrated as the region enclosed by the white dashed lines, along with the food prints of related surveys. The overlap with other surveys will enable important cross correlation measurements. Figure from [62]. 38

2.8 Estimated power spectra constraints using the m-mode pipeline[81], with a foreground filter applied that assumes ideal knowledge of the instrument. Forecasts for HIRAX foreground removal indicate that HIRAX can access the $k_{\|}$ scales of interest at all frequencies, while for $k_{\perp}$ we have access to the desired scales for 3/4 of band. Figure from [62]. 42

2.9 The HIRAX dish design. The HIRAX dish is a 6m F/D 0.23 parabola, which situates the feed in line with the dish edges to reduce cross-talk (radiation from one feed element winding up measured by a neighbor). The receiver support is a fiber glass column, for structural support and improved RF performance over traditional metallic feed legs. The feed is a dual-polarization cloverleaf backed by a can, which circularlizes the beam and shields from the RF over fiber transmitter modules. Design for the foundation and mount are still underway, with the requirement that the dish can be periodically repointed in a few degree increments. Figure from [63]. 44

2.10 One of the many raw images used for photogrammetry (top left), an edited version of the left image for illustrative effect (top right), results of 3D dish model from synthesis of all images as processed by the VGOS collaboration, showing 5mm surface deviations (bottom). 48

2.11 The HIRAX RF chain. The polarized sky signal is collected by the dish and focused into the dual-polarized feed. First stage amplification of 43dB happens within the antenna structure as a noise reduction technique. The amplified signal is filtered, amplified again, and then converted from RF to optical in the RF-over-Fiber (RFoF) modules that sit behind the antenna. The optical signal travels over fiber to the correlator hut, where it is transformed back to RF, amplified and filtered once more, and finally digitized and correlated. There is active work within HIRAX on all parts of this RF chain, with different groups dedicated to antenna design and characterization (see Ch 3, 4), RFoF design and characterization, and different aspects of the correlator. All hardware in the chain should be repeatable between dishes, and the general components will be individually assessed for RFI contributions. Figure from [62]. 51

2.12 Select hardware from the HIRAX RF chain. (a) The HIRAX antenna feeds. The feed design is based on the CHIME antenna, but has the first stage amplification integrated into the antenna structure. This reduces feed noise and makes it more lightweight, but in turn prevents the noise temperature from being directly measured with a noise figure meter, which will be explored more in depth in Ch. 4. (b) The RFoF modules, developed between WVU, CMU and UKZN, will convert RF signal to optical so that it can more efficiently and cheaply travel long distances. (c,d) Different components and angles of the correlator hardware. Photos from the HIRAX collaboration. 52

3.1 Illustration of a $\lambda/2$ dipole antenna, and the distribution of surface charges in the presence of an electric field. As the electric field oscillates, the surface charges move, generating measurable currents and voltages that oscillate with the frequency of the wave. Figure from Wikipedia. 56

3.2 Illustration of a common antenna beam pattern, shown in 3D, 2D and 1D. Figure from Antenna Design and Theory by Balanis [90]. 57

3.3 Illustration of an antenna transmission measurement, in which the Friis equation is applied. P_t is the power transmitted from some transmitter antenna, and P_r is the power received. Figure from [90]. 60

3.4 Illustration of E- and H-plane cuts for co- and cross-polarization of a cartoon dipole antenna, as measured with transmitter Tx and receiver Rx. The top image shows the E and H planes for a biconical antenna, discussed later in this chapter. The black arrows indicate polarization direction. 63

3.5 The Mesa Antenna Range at Jet Propulsion Laboratory in Pasadena, CA. The background photograph is the canyon which sits beneath the signal path, and the outlined photographs are the transmitter yagi antenna (left), the range mount and antenna under test (upper right), and the reference antenna (lower right). 65

3.6 An anechoic chamber in the North Carolina State University Nanofabrication Facility (NNF), used for antenna beam measurements. Left image: a HIRAX antenna mounted in the chamber, Right image: the standard-gain horn transmitted mounted in the chamber (right photo from NNF website). 67

3.7 An antenna range at Green Bank Radio Observatory. Inlaid photos include (1) the measurement control room and (2) mounting the antenna under test, in this case a biconical antenna attached to a drone. . . . 69

3.8 The HIRAX antenna, with a can (left) and plate (right) for its ground plane. These are the main mounting configurations discussed in this section. 70

3.9 (Left:) Gain profiles for three HIRAX antennas with the integrated first stage LNAs and (Right:) gain profiles of the HIRAX and CHIME feeds normalized to the same value at 600MHz for comparison. The HIRAX feed gain is repeatable across the three measured feeds. The CHIME and HIRAX gains share certain features but are qualitatively different (despite having largely similar antenna designs), likely due to the HIRAX LNA bandpass shape. 71

3.10 Full co and cross-pol cuts in the E and H planes for the HIRAX feed mounted to a can. The data plotted in the top block was taken at North Carolina State University's anechoic chamber, and the data in the bottom plot row was taken at JPL's far field antenna range (measurements contaminated by environmental factors). In all cases, the main beam appears to grow asymmetric at higher frequencies, and develop especially high cross-pol. 72

3.11 Cuts in two planes for HIRAX antenna beam, including measurement and simulation results. 73

3.12 H-Plane beam measurements of the HIRAX antenna, repeated for the antenna flipped upside-down (rotated 180^o). Measurements are taken to verify that beam asymmetries are true features of the antenna and not artifacts of the antenna range. To do this, we complete a beam cut, then repeat it with the feed rotated by 180 degrees. The two cuts are plotted as the blue and orange curves (left plot). If the orange curve is mathematically flipped about 0^o, we find an identical signal to the blue curve (right plot), which means that the asymmetries are entirely due to the feed performance and not an artifact of reflections from the environment. 74

3.13 Cuts in two planes for the HIRAX antenna backed by a plate and a can. A direct comparison between the two measurements demonstrates the utility of the can, which is meant to shrink the backlobe, narrow the beam, and circularise the beam. The HIRAX feed and plate are drawn into the 500MHz E-co measurement to illustrate the antenna orientation for the reader. 75

3.14 The CHIME (left) and HIRAX (right) antennas are very similar in design. Apart from different material compositions, there are two main differences between the feeds that impact the beam pattern. (1) CHIME is completely passive while HIRAX integrates a first-stage LNA into the antenna structure. This is visible on the balun stems in the photographs. (2) The polarization axes for CHIME are in line with the petal gaps, while the polarization axes for HIRAX run between diagonal petals. 76

3.15 Full E and H-plane cuts for the CHIME antenna, including both co and cross-pol measurmements. 77

3.16 A comparison between the CHIME and HIRAX antennas, which share a similar design. Two cuts are shown for each antenna: the E-plane cuts are in red/orange shades and the H-plane cuts are in blue shades for HIRAX and CHIME. HIRAX is notably less symmetric, particularly at higher frequencies. This comparison is done with a plate because the CHIME feed has extra port for noise injection that prohibited attachment to the can on hand during this measurement set. . 78

3.17 Schematic of an Aaronia biconical antenna (image from Aaronia website). 79

3.18 E and H-plane co-polar cuts for the Aaronia 20300 biconical antenna, which is to be mounted to a drone for beam calibration. Both specs from the antenna datasheet and measurement results from the NCSU anechoic chamber are shown. Data sheet plots are overlaid with antenna images to illustrate antenna orientation for the different cuts. . 80

3.19 The biconical antenna beam with different backing structures, which confirm that the presence of a ground plate or can is helpful for forward-directing the beam. Here, 'free' refers to the biconical antenna free-standing with no backing; 'can' refers to the HIRAX can, which is separated from the antenna by a two inch nylon post; and 'plate' refers to a 13in plate also spaced 2in from the antenna. The top figure is antenna gain as a function of frequency, and the bottom figure includes E and H-plane cute for the different backing scenarios. The presence of a backing structure can decouple the antenna beam from metallic structures in its backlobe, which will be important when using the antenna for drone calibration, which requires precise knowledge of the transmitter beam pattern. 81

3.20 Left: Diagram of biconical antenna and ground plate, along with the post length, which is altered for this investigation. Right: Plots showing the impact of varying mounting post length. The goal is ultimately to use this antenna as a transmitter for beam measurements of telescopes, and for this particular application, the beam width variations caused by altering post length is not of great consequence. 82

3.21 (a) Antenna range beam measurements of the biconical antenna mounted to a hexacopter drone, with and without a groundplate. The presence of a ground plate considerably smooths the beam by limiting reflections off of the drone carbon fiber legs and body. These measurements were performed at the Mesa range at NASA JPL. (b) Anechoic chamber measurements of the biconical antenna and a mock-drone separated by two different backing structures. These measurements were performed at the NCSU anechoic chamber. The comparison between a mesh backing structure (which includes cutouts for drone landing legs, and has been utilized for drone calibration flights) and solid plate indicate that the solid plate is better for smoothing the beam. Ground plate performance will be revisited in future iterations of transmitter design. 84

3.22 E and H-plane beam measurements for the drone transmitter system, performed at Green Bank Radio Observatory. The top plots show a comparison of cuts for all frequencies, showing a broadening of the beam as frequency increases. The other plots show the same E and H-plane cuts, but separated by frequency, to make individual beam features more visible. Beam measurements of this sort are a critical step to removing the drone transmitter beam from telescope beam maps (see Ch. 5 for further discussion). 86

3.23 Beam cuts in 1D of a passive HIRAX feed mounted to the drone (along with a ground plate). The beam is broad and smooth at the measured frequencies, indicating that this antenna is a promising model for a drone calibration transmitter. 87

4.1 Preliminary S_{11} simulation results for the CHIME (unamplified HIRAX) feed in various cavity shapes before optimization (orange), plotted against simulation results for free space (blue). All cavities are the same height, use the same RF absorber material parameters, are made from aluminum, and share similar base dimensions (i.e. the diameter of the cylinder and cone bases and the length of the cube sides are equal in length). From these early simulations, we find little difference in RF performance between leading contenders, and settled on a cylinder to optimize for the final design. Aside from RF performance, cylinders are mechanically robust, and simple to construct. 93

4.2 S_{11} simulation results for the CHIME (unamplified HIRAX) feed. We compare free space results to results from the Y-factor measurement system hot and cold loads (the cold load contains a simulated liquid nitrogen surface), finding all profiles to be similar to well below -10dB. Their differences are computed to give a sub-1K error in noise temperature from 400-750 MHz. 94

4.3 Foam and fiberglass insert construction procedure: (1) Laying down initial fiberglass cloth layers; (2) Inflatable tub used to keep insert shape during vacuum sealing; (3) fiberglass tub with peelply and burlap layers within an open vacuum bag; (4) fiberglass insert after vacuum bagging; (5) spray-foam procedure (photograph from earlier prototype); (6) spray polyurethane foam on base of insert; (7) leveling out foam on insert base; (8)(9) After filling in our remaining foam, we have an LN2-proof insert (subject to further testing), including the RF absorber for fit-checking. 96

4.4 Various woodworking projects aid with the construction of the y-factor measurement system. 98

4.5 Top: Schematic of the noise temperature measurement system design. Bottom: Labeled photograph of the hot/cold loads for noise temperature tests, taken during an LN2 fill. The loads are covered and sealed with a steel lid for measurements. 99

4.6 A custom Faraday cage built around the Y-factor measurement system. The cage was constructed for RFI mitigation, and demonstrations show it provides 20dB of attenuation for unwanted signals. The system is built from T-slot material and contains two doors, one on the side of filling and the other at the front for personnel access. The personnel door is embedded within a larger door, sized to enable the cylinders to be removed. The cage additionally serves as an extra safety barrier between the nitrogen bath and spectators in the lab, and the structure provides a convenient way to support cables and relieve cable strain. . 101

4.7 (Top) Return loss (S_{11}) measurements of the CHIME (unamplified HIRAX) feed in a reflective cavity as various insulation components are added in the construction. (Bottom) Return loss measurements of the cavity with full insulation compared with the range of empty measurements. The empty cavity was measured twice (2019/03 and 2019/06), with the S_{11} spectra changing by up to 0.2 dB. The insulation components are shown to be RF transparent to within this range of fluctuations for most of the band. Not shown is the addition of an aluminum tape layer, which also has a negligible impact. 103

4.8 S_{11} measurements of the CHIME (unamplified HIRAX) feed in the warm cavities with RF absorber installed. We compare measurement results from the two cylinders that make up the Y-factor measurement system (both at 300 K), finding they are identical to sub-dB level, and share the same overall level as simulation results, though resonance peak locations differ. These return loss measurements match with published passive feed measurements of CHIME, which shares the HIRAX antenna design[86]. More details are described in the text. 104

4.9 Theoretical (blackbody) spectrum, measured spectrum, and residuals for the CHIME (passive HIRAX) feed + four commercial amplifier chains. The top plot shows a comparison between the expected power of the amplifier chains at 300K (dashed lines) and the corresponding spectrum measurements in the experimental system (solid lines). The expected power is computed from P = $G_{\text{chain}}(\nu)k_B(T_{\text{load}} + T_{\text{loss}} + T_{\text{LNA,1}})\Delta\nu$, where T_{load} is the thermal load temperature, T_{loss} is the assumed feed noise temperature from material loss, and $T_{\text{LNA,1}}$ is the noise temperature of the first LNA in the chain. The bottom plot shows the residuals, revealing only slight differences between experimental and theoretical values (neglecting the features in the vicinity of 750 MHz) that could be accounted for by cable loss, gain uncertainty, and other systematics, verifying that we measure a blackbody. The chains are comprised of the following Mini-circuits amplifiers: **chain 1** = ZFL-1000H+ → ZX60-P103LN+ → ZX60-P103LN+; **chain 2** = ZX60-112LN+ → ZX60-P103LN+ → ZX60-P103LN+; **chain 3** = ZX60-P103LN+ → ZX60-P103LN+ → ZX60-P103LN+; **chain 4** = ZX60-P103LN+ → ZX60-P105LN+ → ZX60-P103LN+. All amplifiers have frequency dependent gain, and all chains include 9dB attenuation. 107

4.10 S_{11} profiles (as measured by the CHIME feed) at various liquid nitrogen depths and associated uncertainty in noise temperature. (Left): The return loss S_{11} measurements of the cold load exhibit a resonance that shifts in frequency as the cavity fills with liquid nitrogen. Though the resonance shifts, the S_{11} profiles remain below -10 dB (dashed line) across the full band, which is within the design specification for the feed. (Right): The uncertainty in noise temperature measurements is computed as described in the main text (equations 4.7 to 4.11). The shaded region indicates when this uncertainty provides an upper bound on the measurements, which is the preferred regime. Outside of this regime, the uncertainties remain below approximately 3 K, which is already a cautiously high bound. 111

4.11 Deviations from the expected 30 K noise temperature due to inherent differences between the two loads (measured at 300 K). The frequency dependent feature is common across both polarizations and different feeds, and remains when absorber configurations are switched between the two systems. This offset can be removed from the final measurement results. 113

4.12 Characteristic offset in noise temperature measurements due to discrepancies between the two RF cavities, as predicted by warm verification measurements and verified by Y-factor measurements. Here, $T_N^{(1,2)}$ denotes the noise temperature measured with cavity 1 (at 300 K) as the hot load/ cavity 2 (at 77 K) as the cold load. We take $\overline{T}_N^{(1,2)}$ to denote the predicted noise temperature computed from two warm measurements (also shown in Figure 4.11) using a spectrum measurement in cavity 1 (at 300 K) as the hot power/ a spectrum measurement in cavity 2 (at 300 K) minus 4.9 dBm as the "cold" power (the 4.9 dBm offset is consistent with measuring 77 K). We plot the difference between the noise temperature measured using two different loads for hot/cold measurements and using the same load for hot/cold measurements, along with the difference in predicted noise temperature for two different loads and the same load. The recovered feature is highly consistent across different feeds and different measurement days, and as a result can be removed from final measurements. 114

4.13 Schematic of the noise figure measurement set up used for amplifier noise assessment. The set up includes an Agilent-N8975A noise figure meter, HP346B 15dB ENR calibrated noise source, amplifier under test (the DUT), and a mini-jack to relieve strain on connectors. Additional attenuators were included in the chain for the measurements utilized in this thesis, though they are not pictured in the diagram. 119

4.14 Gain [dB] and noise temperature (K) measurements of the three commercial minicircuits amplifiers used for cryogenic system verification, taken with an Agilent N8975A noise figure meter. 120

4.15 Measurement chain used for the for the cryogenic verification measurements. In this set up, we measure the noise temperature of a passive feed connected to a chain of previously characterized amplifiers. From this measurement, the feed loss can be recovered and compared to an expected value. 122

4.16 Antenna noise temperature (Tantenna) and residuals for the passive CHIME and HIRAX feeds connected to varying amplifier chains. The first stage of the amplifier chain is indicated in the legend as "amplifier model", and refers to the amplifiers listed in table 4.2. Shaded regions correspond to ±3K for CHIME and ±10K for passive HIRAX. 124

4.17 Y-factor measurements and expected results for CHIME antenna plus several amplifier chains. The expected results were determined from a noise temperature evaluation of the CHIME feed alone (shown in figure 4.16) in combination with noise figure measurements of the amplifier chain components (table 4.2). The left hand plots show the noise temperature comparisons for each of the three amplifier chains (first-stages 112, P103, P105), and the right hand plots show the difference between measured and expected along with a ±5K shaded region. The 5K bound is a comparable error performance to the commercial noise figure meter (2K error), and is the bound we require for active feed measurements. 126

4.18 Y-factor measurements and expected results for passive HIRAX antenna plus several amplifier chains. The expected results were determined from a noise temperature evaluation of the HIRAX feed alone (shown in figure 4.16) in combination with noise figure measurements of the amplifier chain components (table 4.2). The left hand plots show the noise temperature comparisons for two different amplifier chains (first-stages 112, P105), the middle-column hand plots show the difference between measured and expected along with a ±5K shaded region, and the right hand plots turn the residuals into a % error. The passive HIRAX feed plus amplifier measurements are less constrained than those made with the CHIME feed, but amplifier results are still broadly recoverable to within ±5%, including fluctuations about a mean that can be averaged down in the final active feeed measurements. 127

4.19 The CHIME feed loss and SMA connector loss, compared to investigate whether the frequency dependent features in the CHIME loss profile are a systematic from the cabling. The smooth profile of the connector lost, taken with the same RF chain and cables, indicates that the CHIME loss feature is either real or a result of reflections or some other type of interaction between antenna and cavity. We note that the slope from 400-450MHz is due to late-stage filtering in the amplifier chain that was not accounted for in the analysis. 128

4.20 The HIRAX antenna noise temperature, as measured during five different measurement sets consisting of multiple days. The individual curves on each plot represents a noise temperature for one polarization on one antenna as measured during a < 20 minute window, which may include multiple hot and cold measurements averaged together. Measurements are consistent to within ± 10K, which is shown as a gray band on each plot. 129

4.21 The HIRAX antenna noise temperature is broadly consistent to within 5K/10% over three measurement sets each consisting of multiple days, meeting the desired specifications. The shaded bands on each curve represent measurement uncertainty. Vertical gray lines mask RFI features. 131

4.22 Noise temperature consistency for HIRAX feed measurements, shown as a percent change between measurement days (left) and different feeds (right). In both scenarios, excluding RFI, noise temperatures are consistent to within 5% across measurements (light blue band). . . . 132

4.23 Cavity extenders of 6inches in length were constructed for the Y-factor system, in order to evaluate the impact of reflections and provide the capability to measure larger antennas. 132

4.24 (Top plot) Y-factor measurements of the active HIRAX antenna with both a raised lid and a lowered nitrogen level as compared with the full nitrogen level measurement. Shaded regions indicate measurement uncertainty. These measurements reveal a noise temperature profile that is overall consistent with original cavity measurements in level but differs in feature location. (Bottom plot) Y-factor measurements of the active HIRAX antenna (1) in the raised-lid system using room temperature and liquid nitrogen as hot/cold and (2) mounted to a 3m dish, using on-sky measurements of sources in beam and sources out of beam as hot/cold. The frequency-dependent features present in the raised-lid measurement track qualitatively with on-sky Y-factor measurements. For the cases of the raised-lid and on-sky measurements, there exists a 4ft characteristic length scale between antenna phase center and reflective surface (cavity floor and dish respectively), which further suggesting the presence of a reflection is modifying the measurement. This result is, however, quite encouraging, as it indicates that we have built a test bed with comparable performance to on-sky measurements, which will save considerable time and effort when evaluating future antenna models before deployment. 134

4.25 Noise temperature of the CHIME antenna plus commercial amplifier chain (first stage ZX60-P103-LN+) at two different nitrogen fill levels: full (18 inch nitrogen depth) and low (13 inch nitrogen depth). The overall noise temperature level remains the same between the two measurement sets, indicating that the black body remains at 77K even when not saturated with nitrogen. A shift of 40MHz in frequency occurs between the measurements, which is potentially indicative of changing resonances. 136

5.1 Drone calibration of a stationary telescope, illustrating measurements during a 1D slice through the telescope beam. The Yale drone system (labeled photograph shown) transmits calibrated noise through a biconical antenna, with a beam shown in blue (outlined in dashed line) and the portion of the beam received at the telescope shown in darker blue (narrower, no dashed line). The telescope under test has its own beam pattern, labelled on the diagram, that impacts how sensitively the drone transmission is received. With each pass that the drone flies over the telescope, its transmission will be received at many angles of the beam, and in turn, a beam pattern can be extracted from measurements. This concept is illustrated for a single pass over the telescope, which gives a 1D beam cut. The boxed middle panel shows the progression of measurements for 5 positions during a flight, and the bottom figure shows what these locations would look like in a flight path. A full 2D beam is built from 20+ of these 1D cuts. This picture also highlights why a broad and smooth transmitter beam is important. . 141

5.2 Cartoon sketch of 1/4 of a drone map-making flight over a telescope array. Here, the drone flies up to the single-dish far-field (200m) and then traces out a 200×200m grid in North/South passes before descending. The grid is centered over the dish under test. In practice, this style of flight takes ∼25 minutes and drains the drone batteries to 25% capacity. To complete a full single-polarization map, the drone will need three additional flights. It will: (1) repeat this style of flight but with East/West passes, (2) fly a denser 60×60m North/South grid, (3)fly a denser 60×60m East/West grid. 142

5.3 Illustration of a drone beam measurement using the Friis equation expanded to 2D. This diagram includes the coordinate system used for drone analysis. 146

5.4 Beam visualization and cuts for initial inputs of 0.1m position uncertainty and 0.8deg angle uncertainty, with telescope fwhm =6degs, drone fwhm = 60degs and 0.5 degree angular resolution. The uncertainty for this configuration, which should be attainable with many commercial drone/GPS systems, is more than sufficient. 151

5.5 Visualization of uncertainties in position/angle that will yield sub-0.1% error in FWHM (we want to be in the darkest blue region). 151

5.6 Impact of drone tipping angle (akin to flying with a constant angle offset) on acceptable uncertainty parameters. We find that we can tolerate drone tipping up to 15 degrees with angular uncertainty of 1 degree and positional certainty of 1 meter. 152

5.7 Drone calibrator and mounting plate schematic. The calibration system includes a DJI Matrice 600 Pro drone and Real Time Kinematic GPS, biconical transmitter antenna, and calibrated noise source chain. Plate and antenna are not to scale on the drone schematic. Drone diagram from DJI. 154

5.8 Block diagram for the UBlox RTK GPS, which meets the requirements for a drone calibrator GPS, including cm-level positional accuracy. It is programmable and customizable, and uses a series of low-cost, interchangeable, and well-documented parts. Design and figure credit: Annie Polish. 157

5.9 Positional accuracy of the DJI and uBlox GPS systems, as determined by recording jitter from the systems when stationary. Both GPS systems have ±2cm positional accuracy (further verified with flights testing), which is sufficiently stringent for far-field beam mapping. Figure and analysis by Annie Polish. 158

5.10 Drone roll/pitch/yaw angle definitions. Figure credit: Maile Harris. . 160

5.11 Top: Spectrum of the un-attenuated noise source plus filtering chain that is flown on the drone (typically attenuated down by ∼35-45dB). The Y-axis units are dBm/3MHz. The HIRAX band is shaded, and it is evident that the filter filters sharply around it. This is important for a number of reasons, including that broadcasting with sufficient power at 900MHz can blind the GPS to satellite signals. 163

5.12 The noise source switching and power regulation board. The board takes a TTL input from the uBlox GPS to effectively modulate the noise source on and off, allowing the telescope to measure sky background and drone RFI at each drone position so they can be removed from the beam map in analysis. The board additionally provides voltage regulation and filtering. Board design by Maile Harris [107]. . . . 165

5.13 Measurements of the RF switch board and commercial block switch outputs for inputs (1) a filtered and attenuated noise source and (2) a 50Ω terminator. In both instances there is 40dB+ isolation, which should enable clean background measurements. It is notable that the switching board outperforms the block switch in isolation. Measurements and figure by Maile Harris [107] 166

5.14 Top: Green Bank EMI Chamber set up for drone RFI assessment. Bottom: Outdoor RFI measurements at Green Bank Observatory. . . 171

5.15 Drone RFI measurements across the band regulated by MeerKAT and SKA. With the exception of intentional emitters and power supply emission below 600MHz, the drone is quiet to below the measurement instrument noise floor. Plot and analysis by Annie Polish. 173

5.16 Low frequency drone RFI measurements indicate emission in the lower portion of the band. With the exception of intentional emitters, this emission is broadly within the MeerKAT specifications. It will need to be mitigated to comply with the SKA limits. Plot and analysis by Annie Polish. 173

5.17 A spectrum at Yale's Leitner Observatory shows considerable RFI accross most of the 400-800MHz HIRAX band. We locate a clean portion of the band, from 650-700MHz (highlighted), for telescope operations. This band will be filtered around sharply in the telescope signal chain. This filtering keeps the RFI from compressing the telescope electronics and enables a functional testbed. Figure and measurements from Ry Walker. 177

5.18 The drone calibration testbed telescope at Yale's Leitner Family Observatory and Planetarium (LFOP) (left) and a preliminary 2D beam map made by the telescope (right). The axes on the beam map are East/West and North/South position in meters, and the points on the scatter plot represent drone physical position in an XY plane in the telescope far field. The color scale is power measured at the telescope in logarithmic units, and it is evident even at this preliminary stage of analysis that there is a gaussian beam feature (as expected). This telescope enables the Yale drone group to test drone hardware upgrades locally, and has been used to prototype the RF switch and provide test data to practice removing pulses in the analysis prior to higher-stakes flights. Telescope design by Ry Walker and Will Tyndall, plot by Will Tyndall. 178

6.1 The ingredients for a beam map are telescope autocorrelation data and drone positional data. Here we show the raw data sets from both the telescope and drone. The dataset includes a pulsing noise source, which enables background subtraction at each drone position, but must be fit for in the telescope data analysis. The red box in the figure is used to illustrate a common time slice for the telescope and drone data. . . 188

6.2 Photo of the DSA-110 dish mapped by the drone calibrator. 190

6.3 Drone flight path over DSA. Points on the plot representing drone position, and the colorscale representing power at the dish. The drone flew 100m throws in North/South and East/West, flying in a 2D plane in the telescope far field. Plot by Will Tyndall. 191

6.4 Left:Gaussian fits at 5 different frequencies for a single drone transit through the DSA beam center. Right: Gaussian peak amplitude (arbitrary linear units) as a function of frequency, showing a profile that qualitatively tracks the noise source spectrum used for DSA flights (and later upgraded). . Plots by Will Tyndall; noise source annotations by Emily Kuhn. 193

6.5 Full-width-half-max as a function of frequency for Gaussian fits of DSA beam data. The different curves on each plot represent different drone transits through the beam, and the two plots are for each of the two polarizations. The rippling shape is also observed in transit data, and postulated to be associated with a bounce between the feed and bottom of the dish. Plot by Will Tyndall. 193

6.6 The Baryon Mapping eXperiment (BMX). BMX is comprised of four 4m off-axis dishes, and operates in a frequency range 1100-1500MHz. The drone system made 2D beam maps of the four BMX dishes in March 2020. 196

6.7 The BMX drone flight path and measurement scheme, showing the four individual flights that make up a single polarization map (eight total flights for full mapping). The top row is the flight data for the 2D grids flown by the drone in the telescope far field. The bottom row is the flight data overlaid with a color scale associated with measured power at the telescope (log units), which gives a sense for the beam information gathered in each flight. 198

6.8 2D beam maps for all four BMX dishes, with evidence of side-lobes and polarization-dependent features. Points represent drone position, overlapped for four different flights (fig. 6.7), and color scale is log power measured at the telescope. In these maps, data from each flight is normalized to the peak of its 2D Gaussian fit prior to getting overlaid. A more rigorous calibration is under development for future analysis frameworks. 199
6.9 A single-dish BMX beam measurement (left) and simulation of the antenna beam projected onto the dish (right), which show qualitative agreement between beam shapes, including asymmetric sidelobes constrained to a single axis. Simulation by Chris Sheehy, from the BMX Collaboration. 200
6.10 Left: Full-width-half-max(FWHM) from fits of one of the BMX beams. The FWHM gives the expected behavior in frequency (above 1500MHz, a filter is present in the analog chain, causing noise). Right: Centroid locations from fits of all four BMX beams (orange dots) superimposed with dish locations (red dots) found from flight data. The centroid fits reveal the dishes to be shifted too far inwards, with the South dish additionally mispointed. 201
6.11 Top row: Example of a raw and corrected beam map for a flight with a 0.3s timing offset. Centroids in successive passes are notably out of alignment in the raw data, and the beam circularizes into the expected Gaussian shape once an offset is applied. Bottom row: Uncertainties from 2D Gaussian fit parameters (amplitude, centroid location in 2D, fwhm in 2D, and background level) as a function of timing offset show a curve that is minimized for an offset of ∼0.3 seconds. 203

6.12 Intermediate waypoints in the autopilot flight path result in the drone stopping over the telescope center, which introduces pitch (tipping) angles of up to 20 degrees. The left plot shows the drone flight path, with four passes of focus color-coded, and the right plot shows the drone pitch during these four passes. Profiles are consistent for every other pass, when the drone is flying in the same direction. 204

6.13 The CHIME Outrigger Prototype at Green Bank Radio Observatory. The 8 active dishes are marked with X's, color-coded by whether the dish is instrumented with a HIRAX or CHIME feed. The drone take off location, RTK base station location, and 'dish 0' location (which served as the coordinate system origin) are also marked. 207

6.14 **(Top:)** Telescope data with the RTK base station located nearby the array. The top left plot shows measured power (at a single frequency) as a function of time, which we expect to trace out Gaussian shapes as the drone transits (see fig. 6.1). Approximations of Gaussian features are seen in the data, though have substantial noise around the peak, indicating compression and oscillations in the electronics when the drone is over the telescope. The top right figure is a waterfall plot with frequency on the horizontal axis, time on the vertical axis, and colorscale indicating measured power at the telescope. Bright vertical stripes correspond to continuous RFI sources (power at a single frequency for all time), and bright horizontal stripes correspond to drone transits (broadband power for brief times). Zooming in around 660-740MHz, we see little intermittent bright spots, which are due to the GPS base station frequency hopping from 900-930MHz (because of of details of data taking operation, this is aliased into the HIRAX band in 670-700MHz). This power is sufficiently bright to impact with the telescope analog chain. **(Bottom:)** Telescope data with the RTK base station moved further from the array, where line-of-sight between GPS and telescope is blocked by a building. This smooths the band considerably, which is particularly evident in the waterfall plot, though the data looks more sparse than expected, which is addressed below. . . 209

6.15 Top: Beam map and time stream data when using a standard hard drive for data writing, which resulted in dropped packets and timing delays. Bottom: Beam map and time stream data when writing data to a solid state drive. 210

6.16 Flight catalog of beam maps from GBO flights on August 26, 2021, with titles indicating flight number (indexed by drone file number) and polarization. Timing challenges are evident for the first half of the dataset, after which we complete several sidelobe flights and one main beam flight. Color scale is in log units. 211

6.17 Flight catalog of beam maps from GBO flights on October 21, 2021. The first two flights are the polarization axis-finding flights described in figure 6.24, and the remainder are largely a series of main beam maps. Color scale is in log units. 212

6.18 Flight catalog of beam maps from GBO flights on October 22, 2021. These flights begin as a set of polarized main beam maps, flying grids at a range of drone alignments to maximize chances of polarization alignment between the telescope and transmitter. The remainder of the flights tackle sidelobes, and nearfield mapping (crosses in the second-to-last row). Color scale is in log units. 213

6.19 Maps of the main beam of 8 CHIME Outrigger dishes at 450MHz (one polarization), generated from the raw drone position data (expressed in meters to highlight the flight path, position will later be converted to angle on the sky). Dish 6 did not have an active channel for this measurement. 214

6.20 Upper panel: Simulated and measured 2D beam maps of a CHIME Outrigger dish for frequencies 450MHz/500MHz/660MHz/700MHz/750MHz. Lower panel: 1D slices through the 2D beam, compared between simulations (blue trace) and data (orange points). Simulations by Benjamin Saliwanchik. 215

6.21 The beam centroid drift and FWHM as a function of frequency, found from 2D fits of the main-beam maps. A main beam map illustrating the X and Y dimensions showcased in the plots is also provided. . . . 216

6.22 Simulated and measured 2D beam maps of a CHIME Outrigger dish for frequencies 450MHz/500MHz/660MHz/700MHz/750MHz. Simulations by Benjamin Saliwanchik. 218

6.23 Polarized 2D maps of a CHIME outrigger dish, generated from two orthogonal large grids. 220

6.24 (a) Flight pattern to determine the telescope polarization alignment. The drone hovers and rotates 360 degrees at each of the points marked in red. (b) Hover/rotation point locations shown within the main beams of two different telescopes, one instrumented with a HIRAX feed and the other with a CHIME feed. (c) Plots of drone yaw angle vs. received telescope power at all rotation points and for several frequencies. The two colors represent two telescope polarizations. . . 222

6.25 Simulations of rotational cuts at different locations in the beam from a dish with a HIRAX feed. Solid lines are for a co-polar source, and dashed lines are for a cross-polar source. The different colors are for different beam sampling points. This behavior is consistent with measurements at Green Bank. Simulations by Devin Crichton. 223

6.26 Simulations of the 2D cross-pol pattern, and 1D cuts through different slices of the simulated beam (marked in red). For the 1D cuts in the top row, the dashed lines are co-pol cuts and the solid cuts are cross-pol cuts. The legend indicates the angle through the beam relative to a horizontal line. At high frequencies, the cross-pol level can be comparable to the co-pol level once 5 degrees off zenith, which could result in the behaviors seen in figure 6.24. Simulations by Ben Saliwanchik . 223

List of Tables

1.1 Common ΛCDM parameters. The first six are independent parameters, used as sampling parameters by Planck 2018[2]; the rest are derived. . 8

2.1 Marginalized 68% cosmological parameter forecast constraints for the HIRAX experiment[62] compared to current eBOSS [3] results for ΛCDM, wCDM and w_0w_aCDM cosmologies. 35

2.2 Summary table of ΛCDM Cosmology parameters (from figure 2.6) to be constrained by HIRAX. Current constraints are based on eBOSS results from 2020[3]. 35

2.3 HIRAX instrument parameters for the initial 256-element array. . . . 40

2.4 Target precision values for the HIRAX mechanical design, as found from cosmological simulations. 43

2.5 Specifications for HIRAX correlator nodes. The HIRAX-256 X-engine is comprised of eight these nodes. 53

3.1 Common units in antenna measurement. 58

4.1 Experimental uncertainties and systematics. The main contributors to the error budget are spectrum analyzer fluctuations, nitrogen surface reflections, and cavity differences. These have been marked with a * and are removable or able to be averaged down, as discussed in the text. Additional sources of error are under investigation. 116

4.2 Specifications in 400-800MHz for the three commercial minicircuits amplifiers used for cryogenic system verification. These values were measured using a noise figure meter and vector network analyzer. . . 120

5.1 Uncertainty in drone instrument parameters. 144
5.2 The current weight and power budget for the drone calibrator. It is important to be cognizant of the weight of the payload, as the drone is only rated for a 6kg payload, and utilizing the full payload decreases flight time by factor of 2. 161
5.3 Sources of RFI from the DJI Matrice 600 drone Pro + RTK GPS + custom noise source chain. 175

Acknowledgements

As I have been reminded many times throughout this PhD journey, "it's really hard to do something that's never been done before." The work discussed in this dissertation is the product of very hard work by myself and many others. I have endless gratitude for the educators, mentors, and loved ones who made my dissertation possible, and to acknowledge them appropriately would take another 200+ pages and surely keep me writing past graduation. I hope that the brief remarks below will convey that they have deeply contributed to this body of work and to my scientific and personal growth.

I must first thank my committee members: Profs Laura Newburgh, Nikhil Padmanabhan, Reina Maruyama, and Steve Lamoreaux, as well as my external reader Prof Cynthia Chiang. Nikhil, Reina and Steve contributed valuable comments on this dissertation, and have been important scientific mentors for me during my time at Yale: Nikhil provided guidance in my study of the field of Cosmology, Reina inspired me as an experimental physicist and encouraged me in overcoming personal obstacles in my early years, and Steve lent his technical expertise to my projects (he was the one to suggest adding a fiberglass layer to the cavity insulation, discussed in Chapter 4). Cynthia has been an incredible collaborator and mentor through our shared involvement with HIRAX—I was fortunate to have her train me in the field in South Africa in 2017, and I'm grateful to have been able to tap into her extensive instrumentation expertise on many telecons since. Most of all, I am indebted to Laura for making me the scientist I am today, and it has been a privilege to work for someone so brilliant and kind. When Laura took me in as her (first ever) graduate student, I had never powered on a VNA or written a line of Python. Even so, Laura trusted me to tackle big, exciting projects, which have kept me on my toes and have made my PhD intellectually fulfilling. I'm also grateful to Laura for trusting me to guide my own education, even when my commitments branched out beyond traditional physics, and for always editing the heck out of my writing.

The work presented in this dissertation is a true collaborative effort, and I thus acknowledge my research group at Yale for their important contributions. In particular, Will Tyndall, Annie Polish, Maile Harris, and Ben Saliwanchik were a drone dream team—structuring the analysis code, tackling the GPS and RF-switch integration, and sharing many flight campaigns—and Ben and Maile were instrumental in putting together the Y-factor system insulation. I am also grateful to have worked with many other excellent undergraduates over the years, including Ry Walker and Ema Smith. I'm especially grateful to Ben, a patient teacher who is unselfish with his time, for spending many long hours helping me contextualize my work (even when he no longer formally worked in the Newburgh group). Ben was also my #1 travel buddy during grad school, letting me drag him to vegan buffets in South Africa and buying sneakers so we could hike PCT sections in California. Though not my direct collaborators, my other fellow Newburgh Lab graduate students were delightful to share an office with, and I thank Lauren Saunders and Sanah Bhimani in particular for being my lab support system for many years (and also for debugging my code many times).

I grew into an instrumentalist at Yale's Wright Laboratory and through my involvement with the HIRAX collaboration. The bulk of the hardware described in this thesis was built locally at the Wright Lab, using a range of shop facilities and personnel expertise. Frank, Tommy, Craig, Dave, James and Arina contributed to the construction of various pieces of my projects, with Frank in particular overseeing the liquid nitrogen safety aspects of the Y-factor system. I'm also grateful to the admins and staff who keep Wright Lab running: between Lil, Kim, Paula, Victoria and Geriana, I have had some wild travel reimbursements processed, some comical purchase orders filled, and many lovely conversations that have made my day. With respect to HIRAX, it was amazing to join on the ground floor of such an exciting experiment. I have learned so much as a member of this collaboration, and am grateful for the opportunity to contribute to its advancement. In both of these academic spaces, I have found friendship, mentorship, and collaboration, which I will always cherish.

This dissertation work was made possible by my funding sources and the generosity of external collaborations. My graduate work was funded by a NASA Space Technology Research Fellowship, which enabled travel to a multitude of important conferences and external/remote facilities, and through which I was guided by an excellent supervisor in Dr. Tzu-Ching Chang. The drone project benefited from travel to a number of observatories across the US, for which I acknowledge Vikram Ravi, Nitika Yadlapalli and Gregg Hallinan at Caltech; Anže Slosar, Justine Haupt, Paul O'Connor, and the cosmology group at BNL; and Kevin Bandura and Pranav Sanghavi at West Virginia.

I am fortunate to have been a part of many communities at Yale and to have explored many extracurricular interests, both of which gave me a sense of belonging and added dimension to my graduate education. Committees and initiatives such as the CDC, EJS, and GSAC enabled me to be a part of lasting and important departmental change, particularly in my early years. Outreach programs/groups such as GSI, CT SEED, and Pathways reminded me of my love for science while providing a platform to inspire the next generation. Groups such as (grad and undergrad) Women in Physics+ gave me a grounding support system and lifelong friends. I also thank the WiP+ members for doing the bulk of the volunteering and community building in the department; this thankless labor makes Physics better, and I hope for a more equitable distribution of it in the future. Mentorship programs like WISAY and Científico Latino introduced me to brilliant young scientists, and allowed me to pay forward some of the support I have received over the years. Teaching experiences, particularly in physics-adjacent areas such as science policy, were incredibly educational and introduced me to new avenues for applying science as well as a new passion, and I am grateful to Profs Shelly Lesher, Bonnie Flemming and Sarah Demers for giving me these opportunities. The sum of these experiences undoubtedly made me a better physicist.

I have been asked by many prospective graduate students over the years what my favorite part of Yale is, to which I always blurt out “the people”. I have had an amazing WL family of mentors, friends and office mates over the years including Supraja, Estella, Danielle, Hannah, Caitie, Samantha, Sumita, and Jeremy and Will (who recruited me to come to Yale in the first place!). I had the best graduate cohort, which gave me incredible friends in Joe, Chloe, Sam, Gwen, Emma, Kelly, Paul, Luna, Oskari, London, Jack and Margie. Many long nights of quantum homework sets, longer nights at Gryphon’s karaoke, ridiculous pranks, shots-o-clocks, rounds of Avalon, and apartment parties (where I fell asleep on the couch) later, I cannot imagine this experience without them. This final year I have also been lucky to grow my friendships with Ava and Malena through wonderful astro girls nights, and with Hannah and Chris, who worked, laughed and thesis-wrote alongside me for most nights of Winter 2022. Outside of Yale, I have been lucky to have a support system of friends from college, home, undergrad REUs, and spaces in between.

Finally, I must acknowledge my amazing family. My grandparents are the biggest supporters of my education: Lita and Grandpa Dick instilled in me a love for asking and answering life’s biggest questions and showed me how to work hard while loving what I do; and Grandma and Grandpa have asked about and remembered the small details of my research, and proudly sported a Wright Lab fleece as they traversed the country by van. I have the most wonderful siblings in Clara, Susana and Allison, who keep me grounded and laughing. My parents have made higher education possible for me through their unconditional support and encouragement, and I truly don't know how to begin to thank them. My physics education began a decade ago at Duke University, where I won the lotto by being assigned to a freshman mechanics group with Connor Hann, who has been a great study buddy and even better life partner ever since. In my undergraduate thesis, I thank Connor for "keeping me somewhat sane and somehow knowing and (explaining well) physics he's never encountered when I need someone to bounce ideas off of," which continued to hold true in the writing of this graduate thesis. I will add a thank you for inspiring me through your own incredible work in academic research and teaching, for doing all of the chores during (and outside of...) my thesis writing, for co-parenting the best puppy ever, and for loving me unconditionally. This dissertation would not have been possible without you.

Chapter 1

Background: the promise of 21cm cosmology

This dissertation discusses calibration projects for the future 21cm experiment HIRAX. In this introductory chapter we contextualize HIRAX, tracing a line from broad cosmological principles to instrument calibration challenges and reviewing related experiments.

HIRAX seeks to address questions about dark energy through measurements of baryon acoustic oscillations, conducted with 21cm intensity mapping techniques. We structure the chapter around defining and connecting these key terms: first introducing the standard model of cosmology, explaining how dark energy fits into our current understanding of the Universe, and discussing how baryon acoustic oscillations can be a useful dark energy probe; then tying in 21cm and intensity mapping as promising techniques for baryon acoustic oscillations measurements; and finally showcasing how existing 21cm experiments are revolutionizing astronomy.

1.1 Introduction to cosmology

Through careful observation of the cosmos, we can gain access to the 13.8 billion year story of our Universe's dynamic history. The field of cosmology has given us an outline of this origin story (visual companion in fig. 1.1).

The Universe begins in an inflationary phase, rapidly expanding faster than the speed of light. Inflation[1] ends after a mere 10^{-32} seconds, when the classic Big Bang picture takes over. In the ensuing minutes, the Universe undergoes big bang nucleosynthesis (BBN), during which time light elements such as deuterium and helium begin to form from protons and neutrons (elements heavier than lithium will come later from stellar physics). For the next 380,000 years, the Universe is composed of small amounts of light elements, as well as free electrons coupled to photons by the scattering cross section for Thomson scattering—together, these constituents make up the "photon-baryon plasma". Fluctuations present in the photon-baryon fluid yield observational signatures called BAOs, which will be discussed in section 1.5. The Universe steadily cools until Recombination, when free electrons are captured by hydrogen atoms, releasing photons to free-stream through the Universe (referred to as "Decoupling"). These photons make up the Cosmic Microwave Background (CMB), which is the earliest light in the Universe and is detectable today as a steady microwave glow in all directions. The CMB is studied very actively, and is the strongest piece of observational evidence for a Big Bang.

After the CMB photons are released, the Universe plunges into one billion years of 'Dark Ages'. During this time, the Universe is largely composed of neutral hydrogen, so the only light is from the CMB and the neutral hydrogen 21cm transition. Over time, the hydrogen falls into overdense regions due to gravity, which grow and form the first structures. Stars begin to form, ionize, and emit UV light, which in turn

---

1. Inflation is widely accepted as part of our cosmological model, though remains an area of active research.

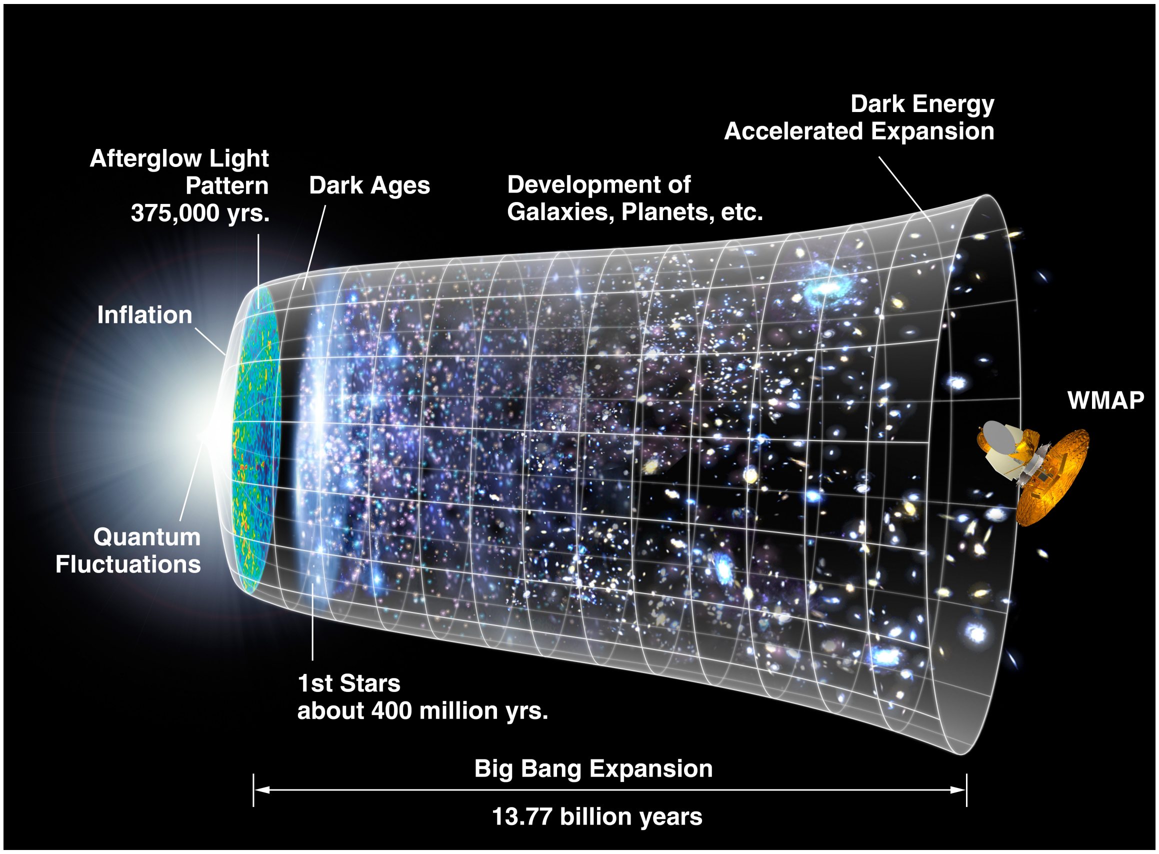
Dark Energy
Accelerated Expansion
Afterglow Light
Pattern
375,000 yrs.
Dark Ages
Development of
Galaxies, Planets, etc.
Inflation
WMAP
Quantum
Fluctuations
1st Stars
about 400 million yrs.
Big Bang Expansion
13.77 billion years

Figure 1.1: Visualization of the Universe's history, from the Big Bang through Present Day. Figure Credit: NASA / WMAP Science Team.

ionizes the neutral hydrogen, ushering in the 'Epoch of Reionization'. For the billions of years since, the Universe has continued to evolve and expand. Recently (within the last 5 billion years), the expansion of the Universe has begun to accelerate, which we attribute to something called dark energy. Dark energy will be further described in section 1.4.

1.2 Definitions and common notation

It is convenient at this stage to introduce standard definitions and common notation used in the field of cosmology.

The size of the Universe as a function of time is specified by a scale factor $a(t)$, with present day,

$$a(t_0) = 1. \tag{1.1}$$

A growing scale factor corresponds to an expanding Universe. In our Universe, the distance to celestial sources is steadily growing, and the wavelengths of light emitted by these sources expand as the Universe expands modifying the measurement wavelength, in a process termed 'redshifting'. This 'redshifting' arises in part from the Doppler shifting of the light emitted by sources. In the case where the velocity v of the sources is much less than the speed of light c, the redshift z is approximated by

$$z \approx \frac{v}{c}\,. \tag{1.2}$$

When accounting for the relativistic effects, the redshift is expressed in its more common form,

$$1 + z = \frac{1}{\sqrt{1 - v^2/c^2}}\,. \tag{1.3}$$

This redshift z can can also be expressed as a function of emitted (λ_1) and measured (λ_0) wavelengths:

$$z = \frac{\lambda_0 - \lambda_1}{\lambda_1}. \tag{1.4}$$

Redshift serves as a proxy for time in cosmology, and can be related to the scale factor $a(t)$,

$$a(t) = \frac{1}{1 + z}. \tag{1.5}$$

Scale factor, in turn, is commonly related to some parameter H,

$$H(t) = \frac{\dot{a}(t)}{a(t)} \tag{1.6}$$

termed the Hubble Parameter, which serves as a proxy for the Universe's expansion rate. The Hubble Constant H_0 is the present day value of the Hubble Parameter,

$$H_0 = \frac{\dot{a}(t_0)}{a(t_0)} = \dot{a}(t_0). \tag{1.7}$$

The expansion of the Universe depends on its constituents—such as dark energy, matter, and radiation—and their relative proportions. The expansion also depends on the Universe's geometry, described by the 4-dimensional Friedmann–Robertson–Walker (FRW) metric[1], which assumes a spatially homogeneous and isotropic Universe at all time, and distances that expand or contract as a function of time:

$$ds^2 = dt^2 - a(t)^2 \left[\frac{dr^2}{1 - \kappa r^2} + r^2 d\Omega^2 \right] \tag{1.8}$$

where κ is the Universe curvature, and r is the radial size, and $d\Omega$ is the angular volume element in spherical coordinates.

The Universe's evolution is captured by the Friedmann Equations, which are the governing equations for cosmology. They are are derived from Einstein's field equations and an assumption that spacetime is described by the FRW metric. The Friedmann Equations are commonly expressed as follows[1]:

$$\left(\frac{\dot{a}}{a}\right)^2 = H(t)^2 = \frac{8\pi G}{3}\varepsilon - \frac{\kappa}{a^2} \tag{1.9}$$

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3}(\varepsilon + 3P) \tag{1.10}$$

where G is the gravitational constant, ε is energy density, and κ is the Universe curvature. The fluid equation/continuity equation is also useful for exploring cosmological behavior, and can be derived from equations 1.9 and 1.10:

$$\dot{\varepsilon} + 3\frac{\dot{a}}{a}(\varepsilon + P) = 0/ \tag{1.11}$$

as is the equation of state, which relates pressure to energy density,

$$P = \omega \varepsilon, \tag{1.12}$$

and will be different for different Universe components. In general, we expect $\omega \approx 0$ for matter, $\omega = 1/3$ for radiation, and $\omega =$ -1 for any energy described as a cosmological constant. We can solve these three equations for $\varepsilon(t)$, $P(t)$, and $a(t)$ to track the Universe's evolution in time.

The Friedmann can be combined to describe how the energy densities of different components impact the scale factor evolution with time, and re-written as,

$$\frac{H^2}{H_0^2} = \frac{\Omega_{r,0}}{a^4} + \frac{\Omega_{m,0}}{a^3} + \Omega_{\Lambda,0} + \frac{1-\Omega_0}{a^2} \tag{1.13}$$

where $\Omega_{r,0} = \frac{\varepsilon_{r,0}}{\varepsilon_{c,0}}$, $\Omega_{m,0} = \frac{\varepsilon_{m,0}}{\varepsilon_{c,0}}$, $\Omega_{\Lambda,0} = \frac{\varepsilon_{\Lambda,0}}{\varepsilon_{c,0}}$, and $\Omega_0 = \Omega_{r,0} + \Omega_{m,0} + \Omega_{\Lambda,0}$; $\varepsilon_{r,0}$, $\varepsilon_{m,0}$, $\varepsilon_{\Lambda,0}$ are the energy densities of radiation, matter, and dark energy today; and $\varepsilon_{c,0}$ is the critical density today (the density required for a flat Universe). In this form, we see how expansion is impacted by relative densities of different components. If we know the energy density of each component ε, we can solve for the scale factor $a(t)$.

The Friedmann Equations capture the broad strokes of how the Universe has expanded given a makeup and geometry, but cannot provide the finer-grain details of its evolution. Questions about how matter aggregates and evolves over time in overdense regions, or how quantum fluctuations from inflation turn into the structure we see today, require a more detailed foundation. To trace our cosmic history, with its dramatic beginnings and varied epochs, we turn to a 6-parameter model called ΛCDM.

1.3 ΛCDM cosmology

Lambda Cold Dark Matter (ΛCDM) is the Standard Model of cosmology. It has been widely adopted by cosmologists because it describes key observables from decades of experimental work: the temperature of CMB, the expansion and evolution of large scale structure, the accelerating expansion of the Universe, and the abundance of light relics (i.e. neutrinos, quantified by N_{eff}). In the name ΛCDM, Λ refers to the cosmological constant, a negative pressure driving the Universe's expansion, and 'CDM' indicates a dark matter velocity significantly less than the speed of light, which excludes certain particles (such as neutrinos) as dark matter candidates.

The primary assumptions of ΛCDM are that the Universe is homogeneous and isotropic, and fits of observational data to ΛCDM parameters indicate that the Universe is spatially flat. At present day, in ΛCDM, the Universe is composed of: 4% baryonic matter (humans, the stars, and the matter we perceive), 26% dark matter (which we've indirectly observed in the cosmos via gravitational pull on luminous matter—though what it is exactly is still an open question and direct detection is an area of active research), and 70% dark energy (a negative pressure, filling all of space, most frequently modeled as a cosmological constant and driving the Universe's accelerating expansion).

ΛCDM involves six independent parameters, though is most frequently discussed in terms of both independent and dependent parameters, including the physical baryon density parameter; physical dark matter density parameter; scalar spectral index; curvature fluctuation amplitude; and reionization optical depth. Parameters such as dark energy density; σ_8 (RMS matter fluctuations over a $8h^{-1}$Mpc radius sphere, colloquially interpreted as 'clumpiness of matter'); and the Hubble constant (described previously) are calculated or inferred from observation. These parameters are outlined in table 1.1.

We have the broad historical outline for the Universe's evolution, but there are

	Parameter	**Definition**
Independent	$\Omega_b h^2$	baryon density
	$\Omega_c h^2$	dark matter density
	$100\theta_{MC}$	acoustic scale at recombination (angular)
	τ	Reionization optical depth
	$\ln 10^{10} A_s$	amplitude of scalar fluctuations
	n_s	power-law index of the scalar spectrum
Dependent	H_0	current expansion rate (Hubble constant)
	h	proxy for H_0; $H_0/100kms^{-1}Mpc^{-1}$
	σ_8	amplitude of matter fluctuations ($8h^{-1}$Mpc comoving scale)
	$\Omega_m h^2$	matter density ($\Omega_b h^2 + \Omega_c h^2$)
	Ω_Λ	dark energy density
	N_{eff}	Effective number of neutrino species
	ω	dark energy equation of state ($\omega = -1$) for ΛCDM
	Σm_ν	sum of neutrino masses
	t_0	age of the Universe (13.799 billion yrs)

Table 1.1: Common ΛCDM parameters. The first six are independent parameters, used as sampling parameters by Planck 2018[2]; the rest are derived.

still many details to fill in. To fully address questions of our cosmic origins, we will need large statistical data sets encompassing all of cosmic time. Presently, there exist precise maps for redshifts $z \sim 1100$, from extensive work on the CMB [2], and redshifts $0 < z \lesssim 2$ from galaxy redshift surveys (e.g. SDSS [3] and DESI [4]) and weak lensing surveys (e.g. DES [5]). There is a large swath of time in between, however, that has been studied only sparsely by a handful of experiments limited by large uncertainties. As described in the remainder of this chapter, 21cm intensity mapping of baryon acoustic oscillations will begin to fill in this gap, and help answer big questions about dark energy.

1.4 Dark energy

In ΛCDM, the accelerated expansion of the Universe is attributed to a new constituent, referred to as Λ or 'dark energy', whose energy density is constant. A cosmological constant was originally invoked by Einstein in the 1920's as a way to achieve a static Universe with the Einstein Field Equations. Not long after, Edwin Hubble made the first measurement of the Universe's expansion rate, and this constant was temporarily abandoned. It was not until 1998 that the Nobel Prize-winning teams of Saul Perlmutter, Brian P. Schmidt, and Adam G. Riess showed that the Universe was accelerating in its expansion[6]. Over the past few decades, measurements of the Hubble constant, which describes the Universe's expansion, have had uncertainties reduced from 200% down to below 5% [7]. Yet, tensions still remain regarding this expansion rate, with measurements from supernovae and the CMB in disagreement to 4.4σ [7] (fig. 1.2).

In ΛCDM, dark energy is consistent with a cosmological constant, with the equation of state,

$$\omega = -1. \tag{1.14}$$

The equation of state relates energy density and pressure, describing the pressure exerted by dark energy on spacetime. In this model, dark energy has no scale factor dependence; it is an energy density associated with space itself, such that the amount of dark energy grows as the Universe expands, further exerting outward pressure on spacetime to accelerate the expansion. A cosmological constant is consistent with modern observations, though they do not rule out a more complex equation of state. ΛCDM is also famously incompatible (to 10^{120}) with predictions of the vacuum energy from quantum field theory [8]. This tension motivates exploration of new models.

To investigate more exotic cosmological models, we can parametrize ω as a linear function. In this case, the dark energy equation of state is modelled as:

$$\omega(a) = \omega_o + (1-a)\omega_a, \tag{1.15}$$

where ω is the equation of state parameter; a is the scale factor; ω_o is the constant component of the equation of state (note that without scale factor dependence it is

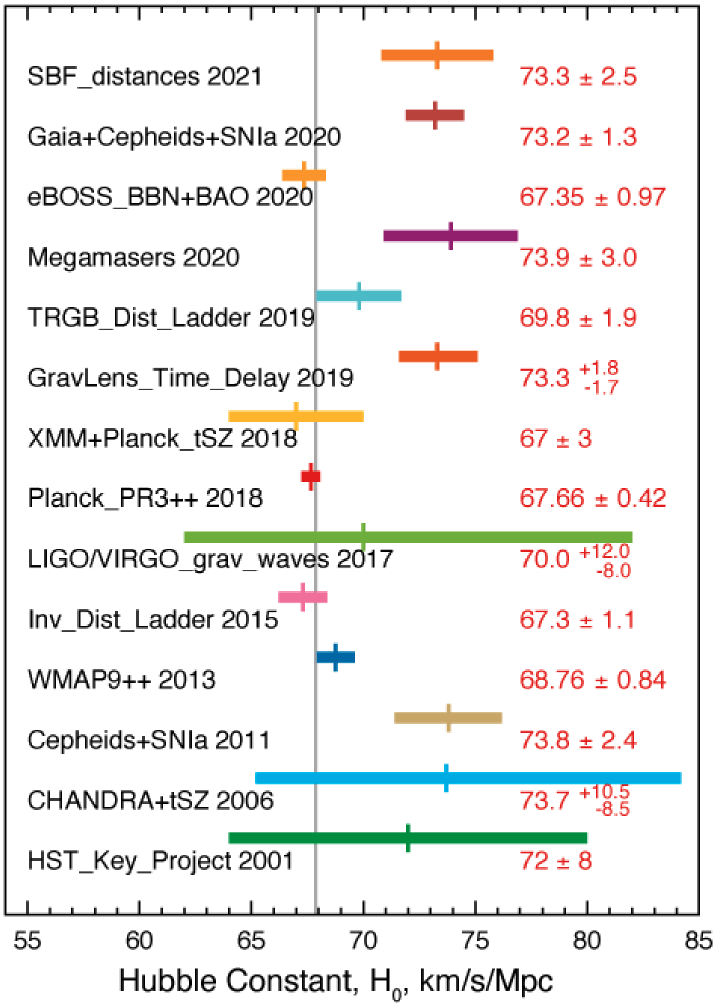

Figure 1.2: Constraints on the Hubble Constant H_0 from a range of experiments, spanning two decades and a variety of techniques. Tension between results from the CMB and BAO (ie Planck, WMAP, eBOSS) and from supernovae (indicated as Cepheids + SNIa) is evident. The gray vertical line is located at 67.88 km/sec/Mpc and represents the weighted average of WMAP and Planck data points. Image credit: NASA / LAMBDA Archive Team.

-1); and ω_a is the slope on the parameter. Presently, ω_o has been measured across multiple experiments (in the CMB[2], by supernovae[9], etc.), and is consistent with ω_o = -1, but precise measurements of ω_a have proven more difficult.

New large-volume surveys with high redshift resolution will be particularly effective at constraining these more exotic dark energy models. For the purposes of this thesis, we will focus on baryon acoustic oscillation measurements as a dark energy probe.

1.5 Baryon acoustic oscillations

To measure the Universe's expansion history, it is helpful to have a reference object that can be measured at different stages in time, referred to in cosmology as a "standard ruler". The Cosmologist's standard rulers of choice are baryon acoustic oscillations or BAOs. BAOs are characteristic length scales in the cosmos on which baryonic matter is concentrated. These scales were set in the early Universe and are still echoed today in the distributions of galaxies. The history of BAOs begins before Recombination, when interactions between the gravitational pull of overdensities in the primordial plasma and high outward pressures in these over-dense regions caused oscillations in the movement of baryonic matter, appearing as spherical shells and moving in accordance with the sound speed of the medium. When the Universe cooled sufficiently for Recombination, the oscillations subsided but their wave patterns were frozen in time, and imprinted into the Universe's matter distribution. Over our history, these BAO features have expanded with the Universe's own expansion. We know their size at Recombination from measurements of the CMB, and can measure them at a variety of redshifts as well as at present day, making BAOs an effective standard ruler.

Initial CMB measurements predict BAO features will appear at $100h^{-1}Mpc$ scales, which is confirmed by a variety of power spectrum measurements from the past decade [10–15]. The power spectrum $P(k)$ is proportional to $\langle|\delta(\mathbf{k})|^2\rangle$, where $\delta(\mathbf{k})$ is the Fourier transform of overdensities in a sky map. The power spectrum is a useful tool for revealing the amount of power concentrated at different wavenumbers k, corresponding to specific length scales in position space. In particular, the BAO power spectrum tracks how the size of the Universe has evolved in time, and from it the Universe's expansion rate can be inferred. An example of a power spectrum (taken from eBOSS) is in Figure 1.3.

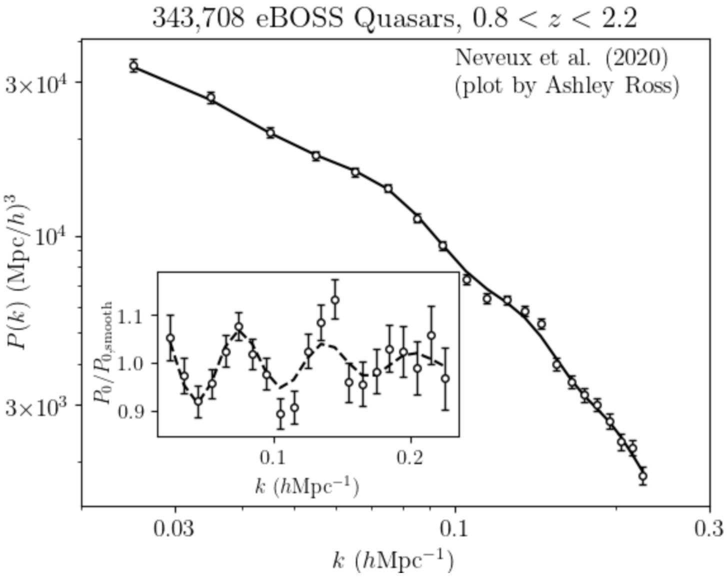

Figure 1.3: The 21cm power spectrum, with the inlaid plot showing the power spectrum divided by a smoothed power spectrum model to highlight BAO features. The x-axis is a length scale on the sky and the y-axis gives power concentrated at each length scale. In the 21cm power spectrum (shown here from eBOSS Quasars measurements), it is evident that power is concentrated at specific length scales on the sky, appearing as three different wiggles, which correspond to the expected BAO features. Figure credit: SDSS; [15, 16] are references for final results.

The first, and so far only, measurements of BAOs have been made by optical surveys, primarily with sensitivity at redshifts <0.8 [17]. For studying dark energy, it is particularly interesting to look at higher redshifts closer to when dark energy began to impact the expansion rate ($z \sim 1-2$): a window that up-and-coming 21cm experiments are setting out to explore.

1.6 The 21cm neutral hydrogen line

The 21cm line originates from a hyper-fine transition due to the electron spin-flip in neutral hydrogen. It is observed frequently in radio astronomy, since 21cm radio waves can penetrate the large clouds of interstellar cosmic dust that are opaque to visible light. This line is particularly useful for three reasons: (i) hydrogen is abundant (ii) the 21cm line is easily distinguished from other lines (iii) the long transition lifetime means that the 21cm line is quite thin, thereby giving excellent resolution in frequency [18]. Measurements in 21cm can reveal how structure in the Universe has evolved throughout its history: they are the only known way to probe the "Dark Ages," a period of time spanning from Recombination (z = 1100) to Cosmic Dawn (z = 6) when neutral hydrogen was in abundance but the first stars had not yet formed; they can shed light on the Epoch of Reionization ($6 < z < 20$), helping us to understand the energy sources and role of structure formation in that time window; they can provide a new way to measure dark energy, from measurements of BAOs, providing data at different redshifts and with different tracers than prior optical surveys.

Measuring the sky in 21cm is challenging, as the signal of interest is very small against a bright sky. The average sky brightness of the 21cm line can be estimated using:

$$T_b = 0.3 \left(\frac{\Omega_{HI}}{10^{-3}}\right) \left(\frac{\Omega_m + a^3\Omega_\Lambda}{0.29}\right)^{-1/2} \left(\frac{1+z}{1.5}\right)^{1/2} mK, \tag{1.16}$$

which corresponds to a brightness of 300μK today. The RMS variation in measured sky temperature (which encodes the noise for a given thermal signal) is:

$$\Delta T = \frac{T_{sys}}{\sqrt{\Delta\nu t_{int} N}} \tag{1.17}$$

where ΔT = the temperature in Kelvin; T_{sys} = system temperature of the instrument in Kelvin, which includes sky temperature as a contribution; N = the number of detectors/elements; $\Delta\nu$ = observing bandwidth in Hertz; and t_{int} = integration time in seconds [19]. Equation 1.16 implies sky fluctuations at the 300μK level, thus requiring $\Delta T < 300\mu$K. To accomplish this specification for a T_{sky} that can be hundreds of Kelvin in relevant radio bands [19], experiments reduce their noise by having long observing times, minimizing system temperature (Equation 1.17), and increasing the number of elements when feasible (such as in a telescope array).

The low-level cosmological signal makes foreground removal a critical step of the 21cm cosmology analysis process. There are a number of foreground mitigation techniques under development for 21cm. These techniques rely on the fact that foregrounds are smooth in frequency, while the 21cm signal is expected to contain spectral structure, so foregrounds can be fit for and removed. To fit foregrounds effectively, we need to understand instrument properties—if we mix spectral structure with instrument frequency structure, we can no longer differentiate smooth foregrounds, as they will contain instrument artifacts. Thus, for successful cosmological measurements in 21cm, having a sensitive instrument is not sufficient on its own: the instrument must also be well characterized.

1.7 Intensity mapping

Because hydrogen is abundant and traces out the structure in galaxies, 21cm is a good candidate for intensity mapping of large scale structure, which will enable high frequency resolution maps of the full radio sky.

Intensity mapping is a new technique that utilizes line emission, such as 21cm from neutral hydrogen, at low spatial and high frequency resolution to make large maps for studying redshift-dependent properties of the Universe. It can efficiently measure statistical properties of structure at higher redshift, which have been difficult to access with optical telescopes. In the case of BAO experiments, intensity mapping involves integrating many sources together to boost the telescope signal-to-noise, while retaining resolution at the BAO scale. Only in the last decade have intensity mapping experiments gained traction, with proof-of-concept measurements yielding encouraging results [20] and the advent of modern data processing techniques that allow for storage of massive amounts of data.

1.8 21cm experiments: past, present and future

Hydrogen has been the most abundant element in our Universe since 300,000 years after the Big Bang; to effectively map neutral hydrogen is to unlock astrophysical mysteries spanning all of cosmic time.

Neutral hydrogen in galaxies was first detected (at $z = 0.8$) in cross-correlation with galaxy surveys by the Green Bank Telescope, and published in Nature in 2010[20]. Over the past two decades, there have been many new 21cm experiments proposed: some are recently constructed and many others are currently under development. These experiments span the globe, and cover a range of cosmology goals. Their science is enabled by recent revolutions in computing, which allow for the storage and processing of large amounts data. They also share common calibration hurdles, several of which are addressed in this thesis.

Several of these new 21cm experiments are seeking to measure the 'global 21cm signal', which is the 21cm brightness temperature averaged spatially over the full sky. These telescopes tend to measure redshifts around the Epoch of Reionization, with some looking as far as the Dark Ages. In this context, 21cm measurements act as a thermometer for the intergalactic medium gas temperature, and give access to a breadth of exciting new physics: by measuring this global signal as a function of red shift, we can place constraints on black hole formation (as they would generate X-rays that would heat the gas), track star formation, and learn about dark matter interactions [21]. In 2018, the experiment EDGES[22] generated excitement in the cosmology community by observing a particularly large absorption feature in the 21cm brightness profile (expected brightness profile is shown in figure 1.4), possibly indicating that hydrogen had lost heat to interactions with dark matter, and setting the stage for other experiments to follow up on the results. Currently active and up-and-coming global experiments include ALBATROS, MIST and PRIZM[23] in Canada, LEDA[24] and the LWA[25] in California, among others. These experiments tend to look like arrays of large (meter-scale) dipole antennas as opposed to traditional radio telescope dishes.

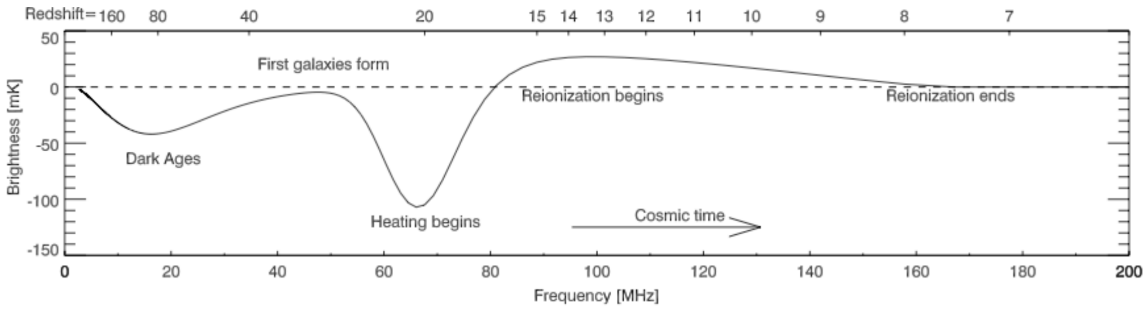

Figure 1.4: 21cm brightness temperature at high redshift. This is the measurement targeted by global 21cm experiments. Figure from [26]

Complementing the global signal experiments are a slew of low-frequency radio interferometers that are looking to shed light on the Epoch of Reionization (EOR) through measurements of large scale structure. Rather than measure a globally averaged signal, these telescopes will make sensitive sky maps and measure a power spectrum, tackling questions such as: Can we measure the first objects in the Universe? How did these first objects ionize the material around them? These EOR experiments include HERA[27, 28] and its precursor PAPER[29, 30], both located in South Africa's Karoo Desert; LOFAR[31] in the Netherlands; MWA[32] in Australia (a specifically dedicated SKA precursor); and GMRT[33] in India, among others.

Many additional experiments are using 21cm intensity mapping as a probe of dark energy, including Tianlai[34] in China, CHIME[35, 36] in Canada, and HIRAX[37] in South Africa, which is the focus of this thesis, and will be discussed in more depth in Chapter 2. These experiments span a variety of designs: CHIME consists of four

100m cylinders, HIRAX resembles a traditional array of parabolic dishes, and Tianlai is a merger of the two.

Over the coming decade, there will be further development in 21cm astronomy, building on the successes of existing experiments and innovations in calibration technologies. The Square Kilometer Array (SKA)[38] will be the world's premier radio instrument, a culmination of decades of radio science development, and will simultaneously observe the broad redshift range of $z \sim 6 - 200$[39]. In North America, an experiment called PUMA[40] is proposed that will encompass a similarly wide redshift range and probe a multitude of cosmological questions. In addition, there is a push to build observatories in increasingly quieter places, with the NASA mission concept FARSIDE[41] targeting the far-side of the moon for deployment.

Within the broader field of astrophysics, 21cm experiments can shed light on a wide-ranging set of unanswered questions. Experiments probing un-redshifted 21cm, such as DSA-2000[42], can provide high-resolution maps of the radio sky that resolve individual galaxies. Radio measurements have found fast radio bursts (FRBs), extremely bright millisecond pulses of radio light from outside of our galaxy whose origins are unknown, and large field-of-view radio interferometers are a particularly effective means of measuring them (for example, the overwhelming majority of the existing FRB catalog was compiled by the experiment CHIME [43–48]; its successor CHORD [49] will enable localization). FRBs have potential as cosmological probes, and can teach us about the locations of "missing baryons" in the Universe, shed light on the origin and evolution of cosmic magnetism, and constrain reionization history, among other things. Additionally, 21cm measurements form the backbone of the upcoming radio-based gravitational wave detector NANOGrav[50], and will enable stochastic background measurements in the next decade, novel tests of General Relativity, and general multi-messenger astrophysics.

The field of 21cm is very promising, and has the potential to address a range of astrophysical and cosmological questions, though instrumentation hurdles must first be overcome. The remainder of this thesis describes my work tackling several of these hurdles for the HIRAX experiment. Although the calibration techniques and hardware I will describe were designed for HIRAX specifically, they will broadly apply to the rest of the 21cm landscape.

1.9 Thesis organization

The work in this thesis focuses on several aspects of the HIRAX instrument characterization, and the document is structured as follows:

Chapter 2 gives an introduction to radio interferometry, and a detailed overview of the HIRAX science goals, instrument design, and current status.

Chapter 3 gives an introduction to antennas and antenna beam measurements, and details beam measurements of a variety of antennas used both in the HIRAX instrument and for calibration of the HIRAX array.

Chapter 4 focuses on assessing the HIRAX antenna contribution to overall instrument noise. It describes the design, verification, and application of a custom cryogenic system for measuring the HIRAX antenna noise temperature, which is a critical step in assessing the overall telescope sensitivity. The chapter discusses early results, as well as next steps.

Chapters 5 and 6 are companion chapters on the design and implementation of a drone calibration system for 21cm telescope beam mapping, with a focus on the HIRAX instrument. Chapter 5 focuses on the system design, which consists of a stable broadband noise source, a GPS- locked timing board, an omnidirectional transmitter antenna, and a precision differential GPS to fly on a quadcopter drone. It also details RFI characterization and key test flights. Chapter 6 discusses beam mapping results using the system, which include first beam maps of several UHF and L-Band telescopes located across the United States.

Chapter 2

The Hydrogen Intensity and Real-time Analysis eXperiment

This chapter will discuss the Hydrogen Intensity and Real-time Analysis eXperiment (hereafter, HIRAX), an up-and-coming radio interferometer based in South Africa. It begins with an overview of radio telescope arrays and radio interferometry techniques, before delving into the HIRAX science case and instrument design.

2.1 Measurement techniques in radio astronomy

Radio astronomy has produced some of the biggest revelations in science, from the discovery of a black hole at our galaxy's center[51] to indirect detection of gravitational waves from binary pulsars[52]. An abundance of astrophysical applications and techniques live under the umbrella of radio science, from imaging far-away objects like the black hole M87*, as was done successfully in 2019 by the Event Horizon Telescope[53], to tracing the Universe's expansion history through low resolution maps of large-scale structure[19].

Radio astronomy has progressed considerably since its birth in 1932, when Carl Jansky's 30m dipole array first detected 20MHz signals from the galaxy. Modern radio telescopes measure frequencies from tens of MHz through hundreds of GHz and come in all shapes and sizes, with the classic design involving a parabolic dish (a curved surface with the cross-sectional shape of a parabola) with an antenna at the focus. This is the design on which the individual HIRAX elements are based.

In a traditional radio telescope, the parabolic reflector dish collects radiation from the sky and focuses it into an antenna. This antenna couples the radiation into a waveguide, which carries the signal to some measurement device that can assess the spectrum. Single dish telescopes image by physically scanning their field-of-view over objects of interest and measuring power at each pointing, effectively building an image one pixel at a time.

All telescopes have spatial dependence to how they radiate/receive signals. The distribution of this radiation is referred to as a beam pattern (described in more detail in Ch. 3, 5, 6). A simple telescope with a circular aperture will have a beam response which is roughly an Airy pattern, where the 'main beam' region can be approximated as a Gaussian. The axis of maximum sensitivity is generally perpendicular to the aperture.

Radio instruments come in many flavors, and employ a variety of experimental techniques. There are large single dishes, such as the 100 meter Green Bank Telescope[54], which point at sources and scan over them to build an image pixel by pixel. There are even bigger single stationary dishes, such as the Five-hundred-meter Aperture Spherical Telescope (FAST) in China [55] and the former Arecibo Observatory[56] in Puerto Rico, that utilize the Earth's rotation to measure different parts of the sky, viewing sources as they drift overhead. With single dish telescopes, instruments on one focus can be easily interchanged, allowing the full telescope to be tailored to different science goals. These setups, however, can be costly to construct and maintain. An alternative to large single dishes is an array of multiple smaller dishes, which saves on cost and provides a collecting area that scales with dish number, boosting the astrophysical signal. The long baselines in arrays will resolve small scales on the sky, and collections of radio dishes can perform interferometry, which is the focus of section 2.1.1.

There are many considerations when designing a telescope or telescope array, several of which are enumerated below:

- **Collecting area**, or the physical area available to collect radio light. This is equivalent to telescope surface area, which includes the combined area of all dishes in an array. Measured power P scales linearly with total collecting area A_{tot}, which in turn scales with the number of array elements N of individual collecting area A_{element}:

  $$P \propto A_{\text{tot}} = N A_{\text{element}} \tag{2.1}$$

  Thus, larger collecting area A_{tot} increases instrument sensitivity.

- **Resolving power**, or the pixel size/image resolution, which is set by length scales in the physical system. For a circular aperture telescope that is fully illuminated, the diffraction-limited resolution θ is specified by the Rayleigh Criterion,

  $$\theta \approx 1.22 \frac{\lambda}{D} \tag{2.2}$$

  where λ is wavelength and D is telescope diameter (D is dish diameter for a single dish telescope, and the distance between elements for an array). Longer baselines between array elements and larger single-dishes will resolve smaller scales on the sky.

- **Field of view**, or the area of the sky which is imaged. In a single dish, this is directly related to resolution, while in an array, resolution and field of view can be separate specifications.

- **Geographic location**, as it is helpful to locate radio telescopes away from sources of radio frequency interference (RFI), such as cellphone towers and TV stations. RFI can contaminate the frequency band of interest, drowning out cosmological signals and saturating the analog hardware, and will thus restrict where radio telescopes can observe.

- **Low system noise**, or low instrument noise, often quantified in units of Kelvin and referred to as as a "system temperature" or T_{sys}.
- **Frequency range & resolution** will be set by science goals and impact instrumentation choices (i.e. lower frequency instruments require larger antennas and different electronics than high frequency instruments).
- **F/D**, or the focal length to diameter ratio. A lower F/D corresponds to a deeper dish which, in an array, can be helpful for blocking stray signals exchanged between neighboring antennas.

This is by no means an exhaustive list, and instead focuses on the main terminology that will be utilized when discussing HIRAX. Not all of these features are weighted equally for each instrument, as different science goals necessitate different design priorities. For instance, an all-sky survey of large scale structure will likely choose poorer angular resolution than an interferometer measuring gas properties in galaxies.

System temperature is an important consideration for radio telescopes. System temperature is an expression of instrument noise and directly related to the telescope sensitivity (for more discussion of system temperature, see Ch. 4). Sensitivity and noise temperature are related by the radiometer equation:

$$\sigma = \frac{T_{sys}}{\sqrt{\Delta\nu\tau}} \tag{2.3}$$

where σ is sensitivity, T_{sys} is system temperature, $\Delta\nu$ is bandwidth and τ is integration time[57].

When measuring small signals, a noisy instrument will necessitate longer integration times. This relationship is particularly important when measuring small signals against a bright sky, as is the case for most 21cm instruments. Thus, radio telescope designs will generally favor low-loss antennas and low-noise electronics. Many arrays will also require observing times on the order of months or years to collect their desired dataset.

2.1.1 Radio interferometry overview

To understand the HIRAX design, it is helpful to first explore how radio interferometers work, thereby enabling us to understand from where the HIRAX instrument specifications arise.

Radio telescopes with two or more elements can function as a radio interferometer. As previously discussed, by utilizing a multitude of telescopes as a single instrument, improved sensitivity and finer angular resolution can be achieved.

To explore how a radio interferometer operates, consider a two element array, with elements A and B separated by some known baseline distance (fig. 2.1). For a given celestial source, A and B will measure the same signal but with a time delay/phase difference. The source location can be determined from this phase information, which comes about mathematically in the following way (following the treatment from notes taken at the Dunlap Institute Introduction to Astronomical Instrumentation Summer School, with [58, 59] as other helpful resources).

For a point source of amplitude A_0, the electric field at a measurement plane in the far field has the form,

$$\begin{aligned} E(x,t) &= A_0 e^{i\phi(t)} \\ &= A(t), \end{aligned} \tag{2.4}$$

where $\phi(t)$ oscillates at the frequency of the source. If the point source is aligned at

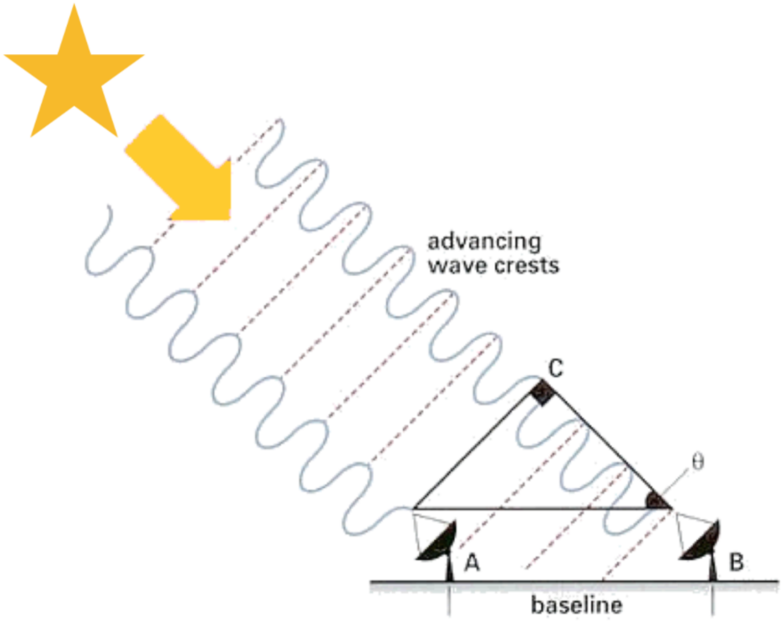

Figure 2.1: A cartoon depicting a two-element interferometer, with elements A and B. At the most basic level, radio interferometers operate by measuring the phase difference between the signals received at A and B and, through knowledge of their physical separation, mapping that phase difference into an angle on the sky. This idea can be extended to N elements of varying baselines. It can be helpful to build redundancy into the baseline separations, to boost signal on these scales and lower computational cost. Figure credit: www.hardhack.org.au/interferometry.

an angle l relative to the measurement plane,

$$\begin{aligned} E(x,t) &= A(t)e^{i2\pi x \sin l} \\ &\approx A(t)e^{i2\pi x l}, \end{aligned} \tag{2.5}$$

with E now depending on the position x on the plane.

For the case of multiple sources in the far field,

$$\begin{aligned} E(x,t) &= \sum_l A(l,t)e^{i2\pi x \sin l} \\ &= \sum_l A(l,t)e^{i2\pi x l}. \end{aligned} \tag{2.6}$$

If the sources are continuous,

$$E(x,t) = \int A(l,t) e^{i2\pi xl} dl \tag{2.7}$$

which is analogous to a spatial Fourier transform. In other words, the electric field measured by a telescope on the ground is the Fourier conjugate to the emitted field in the sky.

We are not required to measure the electric field as a continuous function of position, and can instead build an image by sampling it at discrete points, which is done with a radio interferometer. A radio interferometer measures the correlation of the electric fields at individual dish elements, called a visibility V. For the electric fields E_A and E_B at elements A and B,

$$\begin{aligned} E_A(t) &= \iint A(l,m,t) \exp\left[i\phi(l,m,t) + i2\pi(x_A l + y_A m)\right] dl dm \\ E_B(t) &= \iint A(l,m,t) \exp\left[i\phi(l,m,t) + i2\pi(x_B l + y_B m)\right] dl dm \end{aligned} \tag{2.8}$$

where (x_A, y_A) and (x_B, y_B) are the coordinates of elements A and B. The visibility V is defined as the time averaged product of the electric fields,

$$V = \left\langle E_A E_B^{\dagger} \right\rangle \tag{2.9}$$

We can use the equations in 2.8 to express the visibility as,

$$\begin{aligned}
V =& \Big\langle \iint A(l,m,t) \exp\left[i\phi(l,m,t) + i2\pi(x_A l + y_A m)\right] dl dm \times \\
& \iint A^{\dagger}(l',m',t) \exp\left[i\phi(l',m',t) - i2\pi(x_B l' + y_B m')\right] dl' dm' \Big\rangle \qquad (2.10) \\
=& \iint \left\langle |A(l,m,t)|^2 \right\rangle e^{i2\pi(\Delta x l + \Delta y m)} dl dm \qquad (2.11) \\
=& \iint I(l,m) e^{i2\pi(\Delta x l + \Delta y m)} dl dm \qquad (2.12)
\end{aligned}$$

where, $\Delta x = x_A - x_B$, $\Delta y = y_A - y_B$, and we have defined $I(l,m) = \langle |A(l,m,t)|^2 \rangle$. We have used the definition of the delta function to go from the first equation to the second. We see that the visibility is equal to the Fourier transform of $I(l,m)$.

The data output for interferometers is the visibilities for all baseline combinations, from which the phase differences in the signals can be mapped to positions on the sky to build an image.

This idea can be extended from a two-element array to an N-element array, with the telescope angular resolution set by the largest baseline spacing. Correlating all N elements can be computationally expensive, which can influence the array layout design to include redundancy.

2.1.2 Redundant calibration

The layout of elements within an array is an important consideration and can impact image quality, measurement time, and computational cost.

One technique that radio arrays can employ is baseline redundancy. Baseline redundancy occurs when telescope elements are spaced such that a multitude of pairs are separated by the same distance. By spacing array elements in this way, visibilities within groups of nominally identical baselines can be averaged, boosting signal and reducing data volume. Redundant spacing is similarly helpful for performing instrument gain calibrating, and makes correlations less expensive to compute, reducing from n^2 to nlog(n) operations[60].

There are certain trade offs for this manner of calibration. When relying on redundancy, the analysis is highly sensitive to inconsistencies between dishes, and thus design and manufacturing requirements on array elements will be rigid, as discussed later in this section. The instrument also loses (u,v) coverage (discussed below): by repeating length scales, there is less variety in baseline spacings (and missing information as a result), which will add artifacts to the maps (i.e. result in streaks and fuzziness in images).

2.1.3 The (u,v) plane

The (u,v) plane is a coordinate system in which interferometer baseline spacings are projected into angular scales on the flat sky. Specifically,

$$\mathrm{u} = \Delta x \tag{2.13}$$

$$\mathrm{v} = \Delta y \tag{2.14}$$

are the separations between different array elements (sometimes expressed in units of wavelength). For an array of N elements, there are N(N-1) pairs of (u,v) coordinates. Figure 2.2 illustrates how an array layout in physical space maps to (u,v), and how the (u,v) plane is further filled in by observing as the sky rotates overhead.

An array of few elements—or of redundantly spaced elements—will result in a sparsely sampled (u,v) plane. By contrast, variation in baseline spacing will gives access to more (u,v) coverage. Better (u,v) coverage results in higher-quality sky images.

Interferometers that are interested in imaging will randomly scatter their baselines, and will sometimes have dish locations that are re-configurable, so that elements can

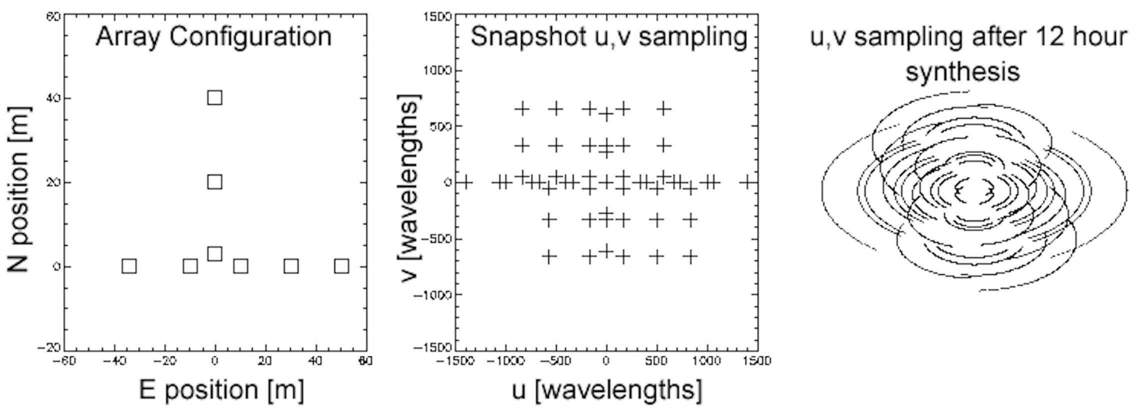

Figure 2.2: Figure and caption from [61] “(Left:) An example array configuration consisting of 8 antennas, with E and N antenna locations in meters. (Middle:) The u,v sampling that results from this array for a single time (snapshot), at zero hour angle. Each cross represents a baseline. There are N(N-1)/2 = 28 unique baselines, but there are twice this many points due to symmetry. (Right:) The u,v sampling that results after a 12 hour integration. The rotating Earth causes the projected baselines to change with time, tracing out portions of elliptical paths.”

be moved to better sample in (u,v). HIRAX, on the other hand, favors computational efficiency and large scale structure mapping over high-resolution image quality, so can forego some UV coverage in favor of boosting signal on redundant baselines.

2.2 HIRAX overview

HIRAX is a 21 cm neutral hydrogen intensity mapping experiment and radio interferometer to be deployed in the Karoo Desert in South Africa [37, 62]. It operates in the 400-800MHz frequency range, probing the redshift range $0.8 < z < 2.5$ in order to constrain the dark energy equation of state and explore fast radio burst (FRB) and radio transient science. It will consist of 1024 six-meter parabolic dishes [63], and will map much of the southern sky over the course of four years. Currently, an 8-element prototype array has been deployed at Hartebeesthoek Radio Astronomy Observatory (HartRAO), and a 256-element array is being developed at the final HIRAX site (Figure 2.3).

HIRAX is a multi-national collaboration. It is funded by the National Research Foundation (NRF) of South African, and led by the University of KwaZulu-Natal (UKZN) in Durban, South Africa. At the time of this thesis, it is made up of 26 institutions and is a training ground for over one hundred students and postdocs.

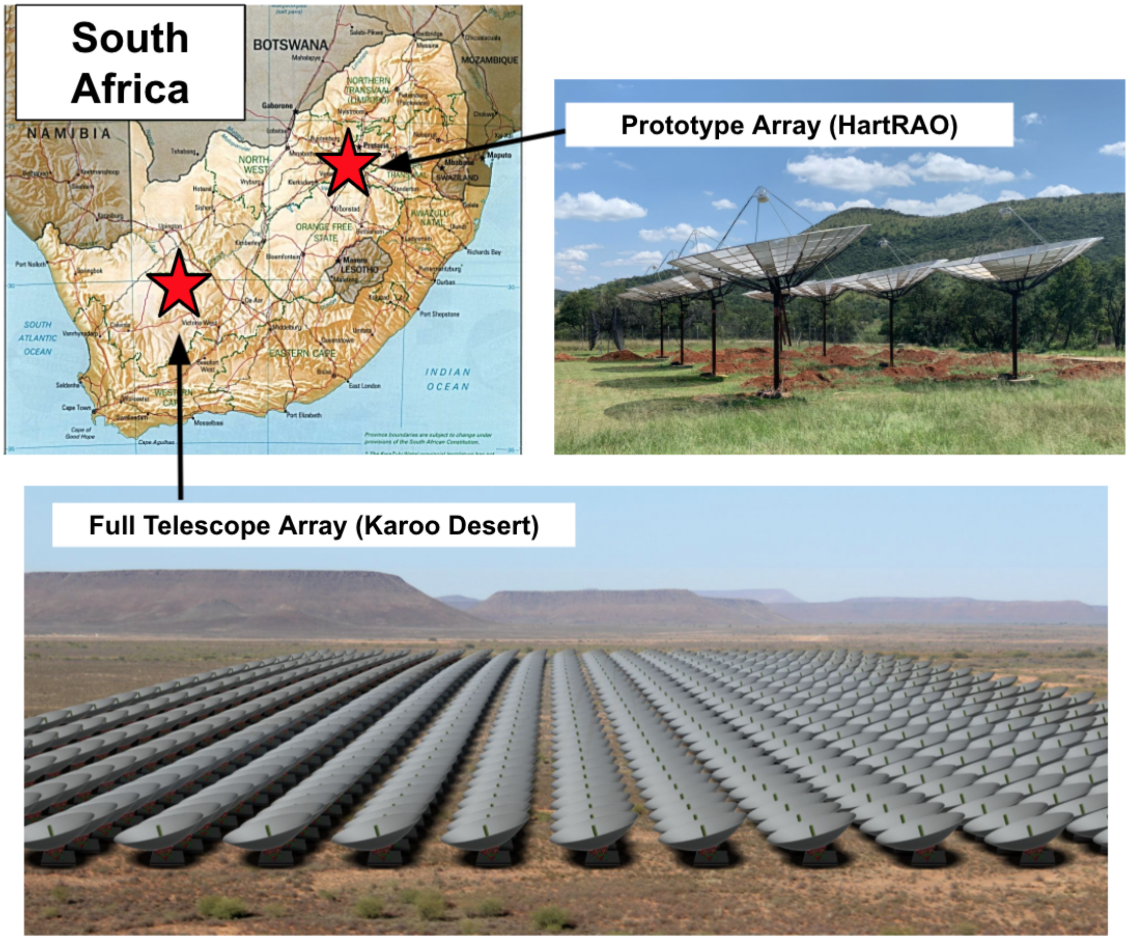
South Africa
BOTSWANA
MOZAMBIQUE
NAMIBIA
LESOTHO
SOUTH ATLANTIC OCEAN
INDIAN OCEAN
Prototype Array (HartRAO)
Full Telescope Array (Karoo Desert)

Figure 2.3: (Top left) Map of South Africa, with red stars indicating the primary sites for HIRAX prototyping and deployment; (Top right) The 8-element HIRAX prototype array at Hartebeesthoek Radio Astronomy Observatory, South Africa ; (Bottom) A rendering of the final 1024 dish configuration in the Karoo Desert, South Africa.

2.2.1 HIRAX location

HIRAX is prototyped at Hartebeesthoek Radio Astronomy Observatory (HartRAO)[1], which resides roughly 70km north west of Johannesburg, South Africa. HartRAO is home to many radio telescopes, including a 26m dish that operates as part of several intercontinental Very Long Baseline Interferometery (VLBI) networks [64], and hosts on-site lodging, machine shops, and observing facilities. The location of HartRAO relative to high-density urban areas, such as Johannesburg and Pretoria, results in significant RFI contamination. The full HIRAX array will be built 1000km South-West, in the Karoo desert.

The Karoo desert is home to the worlds largest radio-quiet site, spanning 106,306 square kilometers[2]. It hosts several world-class telescopes including HERA[65], MeerKAT[66], C-BASS[67] and, most notably, the future Square Kilometer Array (SKA)[68]. The full HIRAX array will be built on this same land, and must comply with the stringent RFI restrictions in place for the other observatories. The pristine radio environment at the Karoo is shown in figure 2.4, and the RFI restrictions with which HIRAX must comply will be discussed in Ch. 5.

A 2-element prototype will be developed at Klerefontein, in Carnarvon, South Africa, where instrument testing will take place. This deployment will be followed by commissioning of the initial 256-element array over the next several years, in preparation for a 1024-element array.

---

[1] For more on Hartebeesthoek Radio Astronomy Observatory: http://www.hartrao.ac.za/

[2] The world's second largest radio quiet site is in Green Bank, West Virginia, and spans 33,670 km^2 to the Karoo's 106,306 km^2. Information obtained from https://www.guinnessworldrecords.com/world-records/647516-largest-radio-quiet-zone

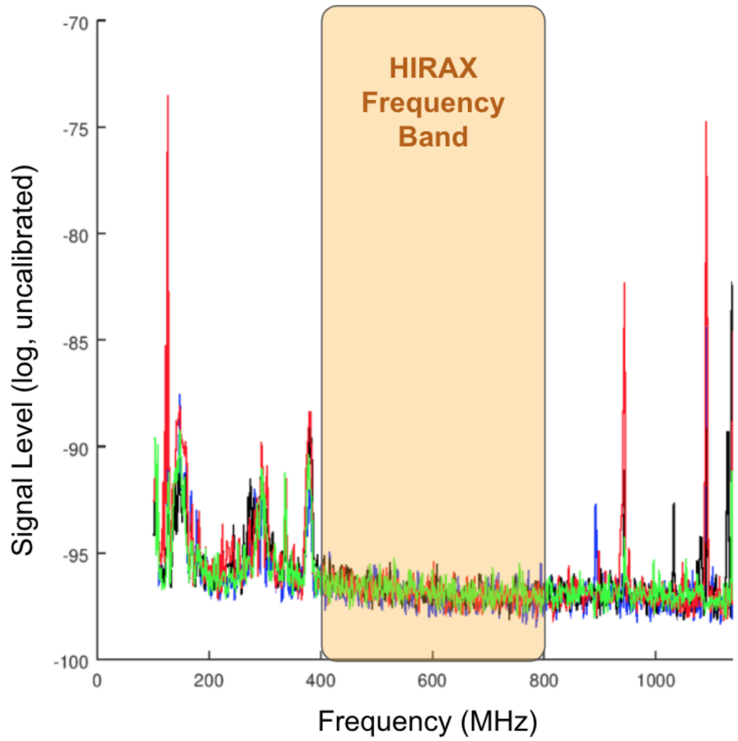

Figure 2.4: Several spectra from the future HIRAX site in the Karoo Desert, taken in 2017 by myself and Prof Jeff Peterson. RFI appears as spikes in the measurement, with higher spikes corresponding to higher power contamination. Thus, it is clear that the Karoo is radio quiet in the HIRAX band, with no visible sources of RFI between 400-800MHz.

2.3 The HIRAX science case

HIRAX has the potential to address several of the biggest open questions in astrophysics. The HIRAX science goals include:

- Probe dark energy via measurements of baryon acoustic oscillations through 21-cm intensity mapping
- Search for pulsars and radio transients including fast radio bursts (FRBs)
- Study neutral hydrogen absorbers
- Measure diffuse polarization of the Galaxy
- Cross-correlations with other southern surveys

- Demonstrate technology for the SKA

This section will expand on each of these goals.

2.3.1 Dark energy

HIRAX is designed to improve constraints on the dark energy equation of state through measurements of large scale structure at high redshift ($0.8 < z < 2.5$). It will target a measurement of the $100h^{-1}$Mpc Baryon Acoustic Oscillation (BAO) scale through 21 cm emission of neutral hydrogen contained in abundance in galaxies. Successful measurements of BAOs have been made by optical surveys at redshifts $z < 0.8$ [10–15], but for constraining the time dependent dark energy equation of state, it is particularly interesting to look at higher redshifts closer to when the universe changed from matter to dark energy dominated (prior to $z \sim 0.5$).

By measuring a higher redshift range, HIRAX can make a particularly sensitive measurement of the scale factor. Because the Universe is expanding, measuring $a(t)$ at earlier times allows a probe of a larger sky volume, and therefore gives access to more modes on the sky, which are the statistics for BAO measurements. Traditional surveys become limited at high redshift by having fewer sources to measure, which is why error bars in BAO measurements typically grow above $z \sim 1$. This trade off is not present for 21cm experiments due to the abundance of hydrogen.

The BAO power spectrum provides information about the size of the universe at different points in time, and so by measuring it, we can solve for the size of the universe as a function of time. The tight constraints that HIRAX will place on BAO measurements, as compared with existing surveys, is shown in figure 2.5.

HIRAX constraints on the dark energy equation of state parameters ω_o and ω_a are shown in figure 2.5 and table 2.1 (recall from section 1.4 that ω_o is the constant component of the equation of state and ω_a is the slope on the parameter). Table 2.1 in particular demonstrates the clear advantage HIRAX has over existing surveys

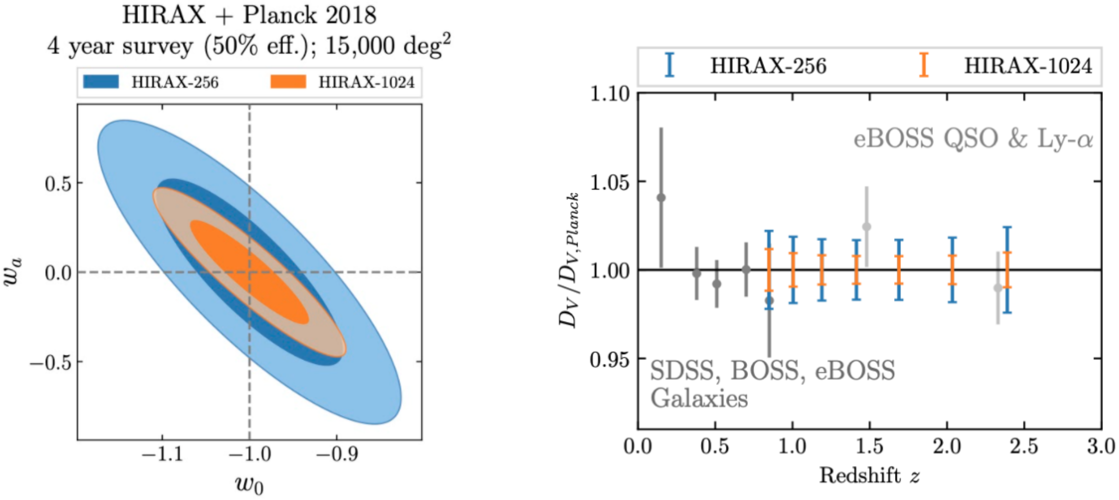

Figure 2.5: Left: Constraints on dark energy equation of state parameters. Right: Constraints on BAO measurements from a variety of surveys as well as the projected constraints for HIRAX 128 dish prototype and HIRAX full 1024 dish array. Note that the HIRAX prototype will now include 256 dishes. It is evident that the sensitivity of HIRAX in this redshift range is unprecedented, and not volume limited. HIRAX will use BAO as a probe of dark energy. Figures from [62]

in measuring non-ΛCDM models. When combined with Planck data, HIRAX can produce $\sim 7\%$ constraints on the dark energy equation of state[62].

In addition to ω_0 and ω_a, HIRAX will more tightly constrain a number of other cosmological parameters, including σ_8, n_s, Ω_Λ, and h, described in table 2.2 and figure 2.6. Projections for these parameters using fisher forecasts are shown in figure 2.6. These projections reveal the constraints that HIRAX may put on the ΛCDM parameters, and illustrate the degeneracy between parameters for the HIRAX instrument. The degeneracy is particularly strong between a larger expansion (h) and more dark energy (Ω_Λ), which BAO has difficulty distinguishing. Work in the CMB has already tightly constrained these parameters, so any improvements are notable.

HIRAX-256 + *Planck*	σ_8	Ω_Λ	w_0	w_a
ΛCDM	0.0044	0.0039	-	-
wCDM	0.0047	0.0042	0.0739	-
w_0w_aCDM	0.0053	0.0043	0.1223	0.4332
HIRAX-1024 + *Planck*				
ΛCDM	0.0027	0.0034	-	-
wCDM	0.0028	0.0036	0.0316	-
w_0w_aCDM	0.0038	0.0037	0.0506	0.1965
eBOSS + *Planck* + SNe Ia + Lens.				
ΛCDM	0.0056	0.0047	-	-
wCDM	0.0092	0.0066	0.027	-
w_0w_aCDM	0.0093	0.0069	0.073	0.5200

Table 2.1: Marginalized 68% cosmological parameter forecast constraints for the HIRAX experiment[62] compared to current eBOSS [3] results for ΛCDM, wCDM and w_0w_aCDM cosmologies.

Parameter	Definition	Current Constraint
σ_8	amplitude of matter fluctuation on $8h^{-1}$Mpc comoving scale (Proxy for clumpiness)	0.8073 ± 0.0056
n_s	power-law index of the scalar spectrum	0.9667±0.0040
Ω_Λ	density parameter of cosmological constant	0.6959 ± 0.0047
h	proxy for expansion rate ($H_0/100kms^{-1}Mpc^{-1}$)	0.6819 ± 0.0036
ω_0	dark energy equation of state	See table 2.1
ω_a	time derivative of the dark energy equation of state parameter	See table 2.1

Table 2.2: Summary table of ΛCDM Cosmology parameters (from figure 2.6) to be constrained by HIRAX. Current constraints are based on eBOSS results from 2020[3].

2.3.2 FRBs and radio transients

HIRAX will discover and monitor transients such as fast radio bursts and pulsars. Fast radio bursts (FRBs) are bright millisecond radio pulses of unknown origin. FRBs have potential as cosmological probes, providing clues to the "missing baryons" in the universe, shedding light on the origin and evolution of cosmic magnetism, and constraining reionization history[69]. The detection and localization of FRBs is rapidly becoming a priority in the field of astrophysics.

Prior to 2018, only a few dozen FRBs had ever been detected. Recently, CHIME has cataloged hundreds of events [43–48], including 18 repeating sources, in the North-

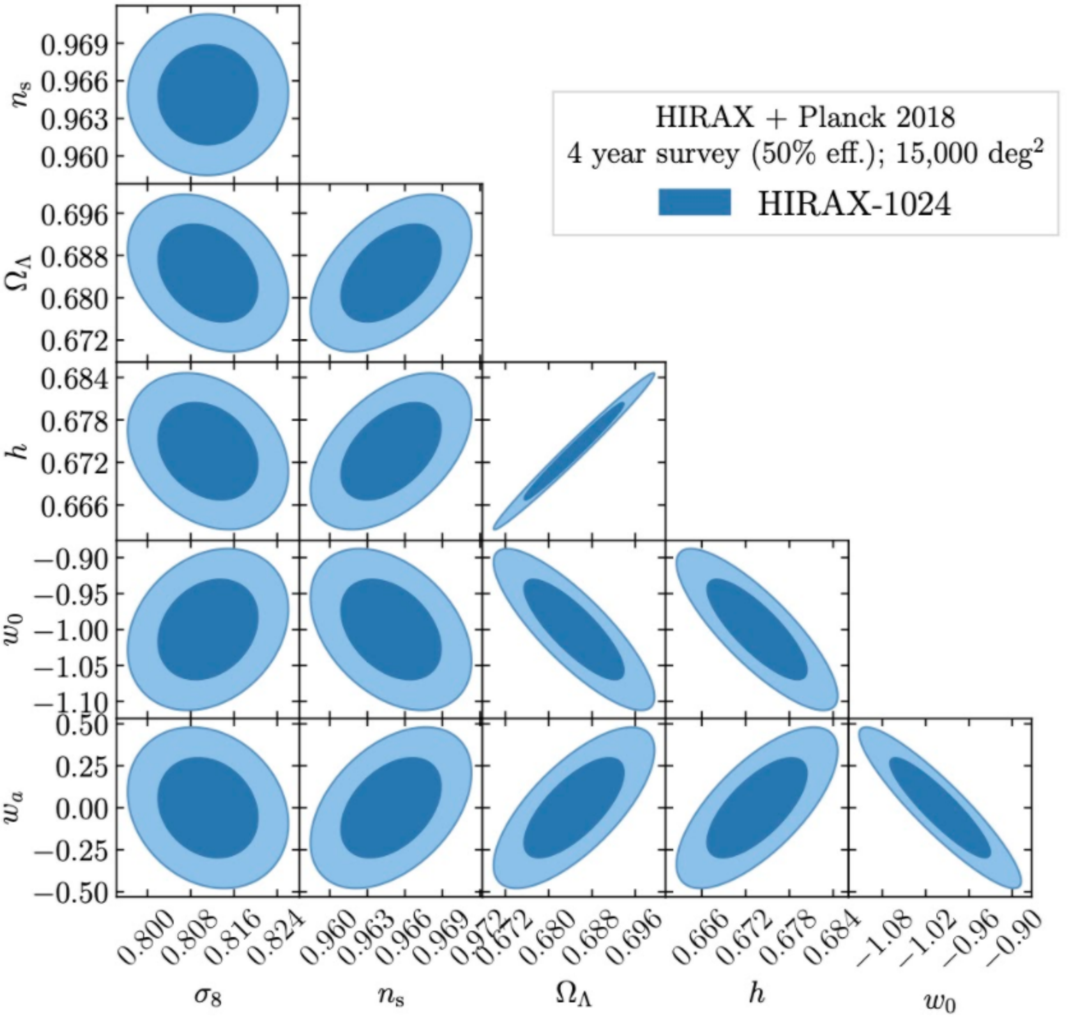

Figure 2.6: Constraints on various cosmological parameters from a 4-year survey of HIRAX-1024 + Planck. HIRAX will place new constraints on several of these parameters (for comparison, see table 2.2) Visualizing constraints in this way also illuminates degeneracies in parameters–the skinnier the oval, the more degenerate parameters are. For BAO measurements in particular, it is difficult to distinguish between a larger expansion (h) and more dark energy (Ω_Λ). This degeneracy can be a feature when cross-correlating with an experiment that is degenerate but in the opposing way, as in if curves are overlapped, the cross-section is particularly small. Figure from Devin Crichton.

ern hemisphere. Thus far, CHIME FRB has found that "bursts from repeating sources differ in burst morphology from apparent non-repeaters, found evidence for a large population of highly scattered events, and measured the fluence distribution and overall sky rate of the FRB population." [70].

HIRAX is expected to have a similar performance to CHIME, given its similar collecting area and sensitivity. As such, HIRAX anticipates finding multiple FRBs per week en route to building a catalog of Southern events. Additionally, FRB localization will be a key piece of the HIRAX FRB effort. HIRAX will build several 8-dish outrigger arrays around the African continent, creating >1000km baselines that will pinpoint event locations to within 0.1 arcseconds. In doing so, HIRAX will shed new light on the origins of FRBs. More specific details about HIRAX FRB detection can be found in [62].

2.3.3 Other science cases

Study of neutral hydrogen absorbers

In mapping the distribution of neutral hydrogen, HIRAX will observe the non-luminous parts of galaxies via their absorption of background flux (measuring 'holes' in the spectrum, akin to observing an object silhouetted by a backlight). Neutral gas is generally difficult to measure with traditional telescopes, as it does not glow optically, and most existing hydrogen absorber measurements were found through pointed observations of pre-selected targets [71]. By contrast, HIRAX surveys a large sky volume in 21cm, which opens a unique window into hydrogen absorption systems. Understanding neutral hydrogen absorbers can help answer questions about galaxy formation and thermal feedback systems in galaxies, and provide measurements to compare with existing hydrodynamic simulations of gas flow in galaxies, among other things.

Measuring diffuse polarization of the galaxy

To successfully measure BAOs, HIRAX will need to detect small signals, several orders of magnitude dimmer than their bright galactic foregrounds. Thus, measuring foregrounds is a critical (and unavoidable) part of the HIRAX operations. These measurements of foregrounds will be of interest to astronomers studying the galaxy.

Currently, state-of-the-art polarization measurements of the galaxy target narrow band regions that are sparsely distributed through the 400-800MHz HIRAX band. HIRAX will augment this existing work substantially, by providing polarized maps in 1024 frequency channels across the full 400-800MHz band. Polarized beam maps in particular will aid this effort (see Ch. 5, 6 for a discussion of telescope beam map making).

Cross-correlations with other southern surveys

HIRAX's southern location will allow for a variety of cross-correlation measurements with other cosmology surveys such as ACTPol[72], DES[73], and the Vera Rubin Observatory[74] (fig. 2.7).

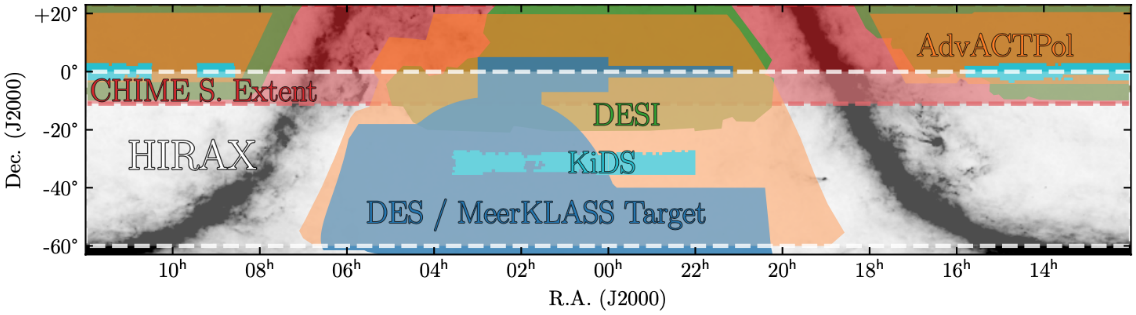

Figure 2.7: The HIRAX survey footprint, illustrated as the region enclosed by the white dashed lines, along with the food prints of related surveys. The overlap with other surveys will enable important cross correlation measurements. Figure from [62].

Implementation of cross-correlation techniques is powerful, as surveys with different systematics should have their systematics removed in a cross-correlation, allowing for more sensitive results with smaller error bars. Additionally, the high-frequency resolution of HIRAX can aid the science goals of other southern observatories through cross-correlation, which may increase redshift resolution in photometric surveys like the Vera Rubin Observatory [75].

This manner of investigation is particularly valuable given the newness of 21cm techniques. Observatories have yet to directly detect the 21cm cosmological signal, with existing measurements of the 21cm cosmological signal from cross-correlations [20, 76]. This cross correlation technique has been historically valuable, and was used for the first detection of lensing in the CMB[77, 78]. The results of a cross-correlation analysis will serve as a validation of HIRAX, which employs new techniques with different biases than existing surveys, and will augment its capabilities.

SKA technology demonstrator

The Square Kilometer Array (SKA)[68] is a next-generation radio telescope array, with first light expected in the mid-2020's. It is a massive undertaking, with billions of dollars in funding by participating nations across the globe, that will pioneer a multitude of technologies and process unprecedented amounts of data (5 zettabytes, or 106 petabytes, of data every year; for reference, global Internet traffic passed 1 zettabyte for the first time in 2016 [79]). HIRAX is an SKA pathfinder—it will be built at the SKA site and, as another telescope array in the same location with shared science goals, can provide demonstrations of RF and computing hardware, calibration techniques, and a host of other technologies in preparation for finalizing SKA designs.

2.4 The HIRAX instrument

HIRAX will consists of 1000 6m dishes of f/d 0.23 arranged in a tightly packed array, and will operate in the 400-800MHz frequency range in 1024 channels (390kHz bandwidth). The array will be built in stages, starting with a group of 256 dishes that are currently funded and soon-to-be commissioned. The cosmological analysis relies on redundancy between dish elements, so all components of the individual dishes must be designed and built with tight specifications in mind.

HIRAX is a driftscan telescope, measuring large scale structure distribution as the sky rotates overhead, rather than tracking particular sources. Because it is stationary, HIRAX will view 10° of sky at a given time, in $\sim$ 3600 square degree increments. To map the full sky, dishes will be repointed every several months, enabling a full sky map in just a few years.

The HIRAX design for the 256-element array is summarized in table 2.3 and detailed in the subsections below.

Parameter	**Value**
Number of dishes	256
Dish diameter	6 m
Dish focal ratio	0.23
Collecting area	7200 m^2
Frequency range	400–800 MHz
Frequency resolution	1024 channels, 390 kHz
Field of view	5°–10°
Resolution	0.2°–0.4°
Target system temperature	50 K
Antenna noise temperature	20-30 Kelvin

Table 2.3: HIRAX instrument parameters for the initial 256-element array.

2.4.1 Instrument design drivers

The HIRAX instrument design, outlined in 2.3 and discussed in the coming sections, is driven by the HIRAX goal to measure BAOs as a probe of dark energy:

- **Frequency range & frequency resolution** are set by the redshift range and redshift resolution. The 400-800MHz frequency range corresponds to a redshift range of $0.8 < z < 2.5$, which is when dark energy is beginning to influence the Universe's dynamics. HIRAX uses a high frequency/redshift resolution to allow a sensitive measurement of the dark energy equation of state and its evolution in time.

- **Number of dishes** and **dish diameter**—which determine **collecting area**—and **noise/system temperature** will all impact instrument sensitivity. A sensitive instrument is needed to measure the $300\mu K$ cosmological 21cm signal, which is set against foregrounds that can be hundreds of Kelvin in the relevant radio bands [19].

- **Angular resolution** is set by the overall size of array (recall, $\theta = 1.22\lambda/D$), and is driven by the requirement to resolve BAO scales (but not necessarily measure details within galaxies).
- **Field of view** is the area of sky measured at a point in time. Over the course of one day, HIRAX will observe 3600 square degrees as the sky makes a full rotation overhead($10^o \times 360^o = 3600^{o^2}$). HIRAX's large field of view is driven by a desire to measure a high volume of Fourier modes on the sky, which are the statistics for BAO. An even larger field of view could provide practical considerations for building a reliable telescope.

To place error bars on the instrument design parameters, we consider how they will impact foreground removal. Foregrounds limit the ability to measure all scales on the sky, especially at lower frequencies. This problem is commonly referred to as 'the wedge' in 21cm [80]. Forecasts for HIRAX foreground removal indicate that HIRAX can access all Fourier modes of interest along the line of sight ($k_{\parallel}$) at all frequencies; though will be restricted to 3/4 of band for modes perpendicular to the line of sight ($k_{\perp}$) (figure 2.8).

The layout and structure of HIRAX and other 21cm experiments bring about several calibration challenges. The arrays consist of hundreds of stationary radio dishes, placed in a pattern meant to exploit instrument redundancy. Additionally, it is nontrivial to measure the beams of these stationary telescopes with traditional methods, since the dishes cannot be pointed at known celestial objects and instead survey the sky through drift scans. Indeed, the Canadian experiment CHIME [82] has been limited in their cosmology effort due to imprecisely characterized beams.

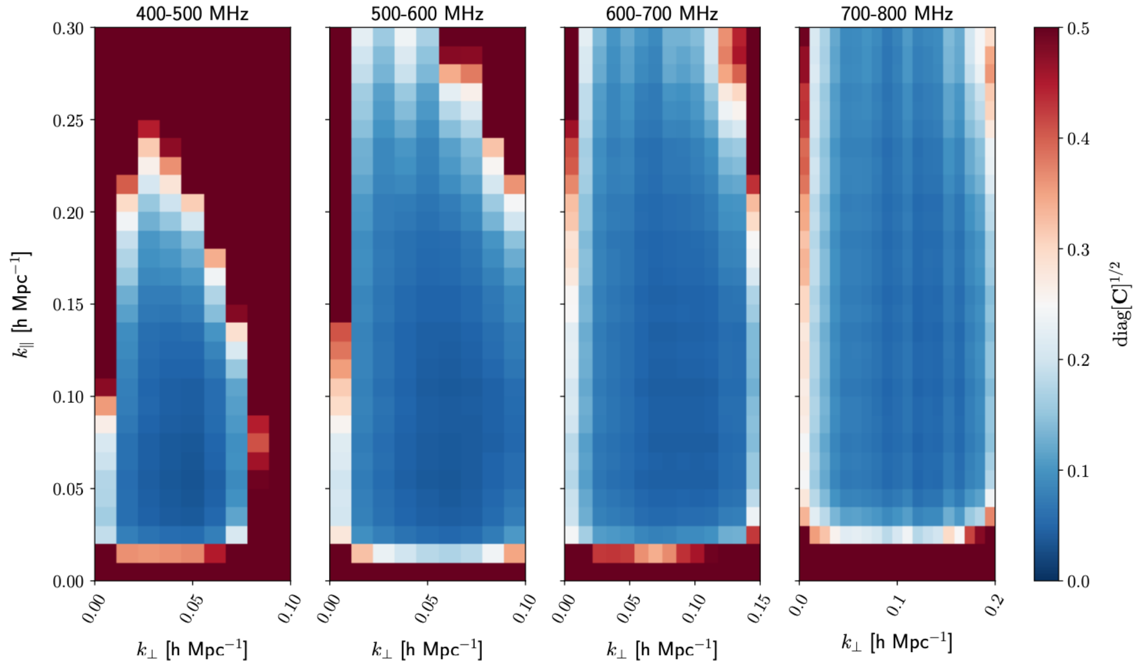

Figure 2.8: Estimated power spectra constraints using the m-mode pipeline[81], with a foreground filter applied that assumes ideal knowledge of the instrument. Forecasts for HIRAX foreground removal indicate that HIRAX can access the $k_{\parallel}$ scales of interest at all frequencies, while for $k_{\perp}$ we have access to the desired scales for 3/4 of band. Figure from [62].

Experiments like HIRAX must therefore utilize low-cost dish and antenna/amplifier instrumentation, given the sheer number of elements involved, and ensure that these elements are well understood and highly repeatable.

The layout and structure of HIRAX and other 21cm experiments bring about several calibration challenges. The arrays consist of hundreds of stationary radio dishes, placed in a pattern meant to exploit instrument redundancy. Additionally, it is nontrivial to measure the beams of these stationary telescopes with traditional methods, since the dishes cannot be pointed at known celestial objects and instead survey the sky through drift scans. Indeed, the Canadian experiment CHIME [82] has been limited in their cosmology effort due to imprecisely characterized beams.

Experiments like HIRAX must therefore utilize low-cost dish and antenna/amplifier instrumentation, given the sheer number of elements involved, and ensure that these

Telescope mechanical parameter	**Target precision (RMS)**
Receiver position relative to focus	0.5 mm
Receiver orientation relative to boresight vector	$2.5'$ polar and azimuthal
Dish surface deviations	1 mm
Dish vertex position relative to elevation axis	1 mm
Orthogonality of boresight vector and elevation axis	$1'$
Elevation axis position within the array	0.5 mm in array plane 1 mm out of array plane
Elevation axis alignment within the array	$1'$
Elevation pointing angle	$1'$

Table 2.4: Target precision values for the HIRAX mechanical design, as found from cosmological simulations.

elements are well understood and highly repeatable.

To develop instrument specifications, the following procedure is adopted: Simulate in CST a single-dish telescope instrumented with a passive HIRAX feed and monitor the far field beam pattern with 50MHz frequency resolution; perturb a parameter that must be constrained and re-simulate the beams; propagate the results from simulations through a cosmological pipeline; evaluate the resulting 21cm power spectrum errors. HIRAX cosmology sets a requirement that the "impact of the systematics under evaluation reduce the Fisher-forecasted Dark Energy Figure of Merit by no more than 20% of the fiducial forecasted level. The forecasts assume a power spectrum noise comprised of purely statistical noise with a nominal residual foreground contribution"[63]. The results of this investigation are tight but tenable instrument requirements, given in table 2.4.

2.4.2 Dish design and manufacturing

The HIRAX dish design and manufacturing requirements were developed through extensive mechanical and optical simulation work, described briefly above and detailed fully in reference [63].

HIRAX instrument requirements are summarized in table 2.4. Several design

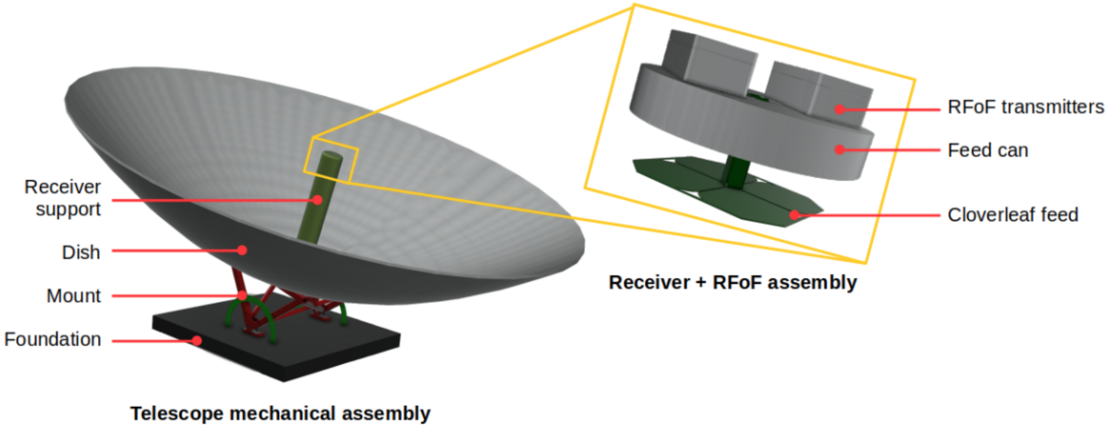

Figure 2.9: The HIRAX dish design. The HIRAX dish is a 6m F/D 0.23 parabola, which situates the feed in line with the dish edges to reduce cross-talk (radiation from one feed element winding up measured by a neighbor). The receiver support is a fiber glass column, for structural support and improved RF performance over traditional metallic feed legs. The feed is a dual-polarization cloverleaf backed by a can, which circularlizes the beam and shields from the RF over fiber transmitter modules. Design for the foundation and mount are still underway, with the requirement that the dish can be periodically repointed in a few degree increments. Figure from [63].

choices were made as a result of these specifications. They are outlined below, and shown in figure 2.9.

- **Feed support:** The antenna is mounted to fiberglass column of diameter 35cm, chosen for its structural stability and RF transparency. Mechanical simulations indicate that this mounting design increased stability under wind loading, relative to differently constructed columns and traditional metal feed legs. Electromagnetic simulations indicate that the design results in suppressed beam features at wide angles (alternate designs, such as metal feed legs, place reflectors into the beam that scatter radiation to wide angles). Built into this support structure are weather-proofing mechanisms that fully enclose the feed and RFoF modules. The feed support has stringent requirements on feed placement relative to the focus, which were developed by varying the feed position in six axes (translations and rotations), and are factored into the design. Final constraints are shown in table 2.4.

- **Cabling:** The feed support design allows for axisymmetric cabling, in which the fiber carrying the signal runs down the fiberglass column to the dish center. Electromagnetic simulations show that this cable placement symmetrizes the beam and limits sidelobe levels relative to designs in which the cable is routed over the dish lip. This finding is intuitive, as the cables contain metal and lie in the optical path, so any placement parallel to the dish surface should increase reflections.

- **Surface roughness:** Simulations of surface roughness tolerance are in progress at the time of this thesis. Presently, we adopt a limit of surface RMS $\sigma < \frac{\lambda}{16}$, derived from Ruze scattering[83].

- **Redundancy:** Dish designs under consideration include a fiberglass composite molded over metal panels, where the monolithic design and mold should yield higher redundancy than a dish assembled from multiple panels. It is important that the dish manufacturing is repeatable. As described at the start of this section, instrument redundancy plays a critical role in calibration of HIRAX and in the data analysis.

- **Focal ratio:** The HIRAX dishes will have a focal ratio of f/d = 0.23, such that the feed sits level with the aperture. This configuration will minimize the exchange of stray signals between neighboring telescopes (from reflections or radiated by the first-stage low noise amplifier (LNA)), known as cross-coupling or cross-talk. The f/d was determined in CST simulations by evaluating the relative delay performance of dishes with different focal ratios. In this exercise, the dish and feed (zenith pointed) are excited by a plane wave, and then voltages in a 50Ω termination of the feed are used to estimate the resulting foreground leakage (formalism outlined in [84]).

- **Mounting:** HIRAX dishes must be able to withstand wind loading of 52m/s (the wind speed record in the Karoo desert), maintain performance under seasonal temperature fluctuations, and generally be robust to desert conditions. These factors will influence the overall telescope mount and foundation. Preliminary designs have been verified to meet the environmental criteria using Solidworks mechanical simulations. Additionally it is important for the HIRAX measurement plan that the array be repointed to different locations in the sky periodically, and thus the ability to change elevations (to 1 degree precision) must be incorporated into the mounting.
- **Can dimensions:** Each HIRAX antenna is backed by a can of diameter 330mm to prevent ground spill and circularize the beam. A small can is expected to reduce antenna return loss, but improve aperture efficiency (as it blocks less of the beam, scaling with r^2). Can dimensions and constraints were optimized in CST simulation software.

Photogrammetry: assessment in the field

Understanding dish surface deviations in situ is a critical part of realizing instruments like HIRAX that require a high degree of sensitivity and redundancy. For this assessment, photogrammetry is a highly useful tool. Photogrammetry is a technique where digital image captures of 3D objects are stitched together to generate a 3D computer rendering. From this rendering, surface irregularities and deviations from the expected parabola shape can be assessed. Photogrammetry has been performed for several HIRAX dishes over many design iterations, including in July of 2017, during a visit I made to HartRAO.

To perform photogrammetry on early-stage HIRAX dishes, I collaborated with visiting scientists from the VLBI Global Observing System (VGOS)[85], a redesign of the VLBI. The procedure is as follows: We placed reflective stickers ∼1ft apart on the dish surface, keeping them evenly spaced throughout the dish. Then, 12 larger reflective targets were distributed on the dish surface. The dish was tilted to 40 degrees from zenith, and then we waited for the sun to go down. Once the sky was dark, I accompanied one of the VGOS scientists up in a crane and took a series of flash photos of the dish (10 photos were taken at each of 8 different positions). The process took under two hours, and therefore could realistically be repeated for a number of dishes in the array. Photos were taken on a Canon D5 camera, with the following settings: ISO 100, aperture f/18, focal length 28mm, and shutter speed 1/250s. Raw data and the reconstructed dish are shown in figure 2.10.

Photogrammetry efforts are ongoing throughout HIRAX, and these measurements have been performed several times for future dish iterations and dishes in different configurations, as well as for dish molds. These measurements elucidated important features, revealing that the dishes begin to ‘potato chip’ when pointed away from zenith, losing their circular shape and becoming slightly elliptical along the gravitational axis. It is useful to understand this behavior and evaluate its impacts on cosmology, which will place specifications on the final dish design.

F/D: assessment in the field

The focal length to diameter ratio, or F/D, is an important parameter for telescopes. A lower F/D corresponds to a deeper dish. Utilising deep design can be particularly advantageous in arrays such as HIRAX, where individual elements are densly packed. By situating a feed deep into a dish, stray signals exchanged between neighboring feeds (known as cross-talk or cross-coupling) can be mitigated.

Though the current HIRAX dish design has an F/D = 0.23, the initial prototype elements had an F/D = 0.38, making them significantly shallower. The F/D = 0.38 dishes provided the only available HIRAX data for several years, and are the dishes on which early feed designs and RF chains were tested. These dishes were similarly

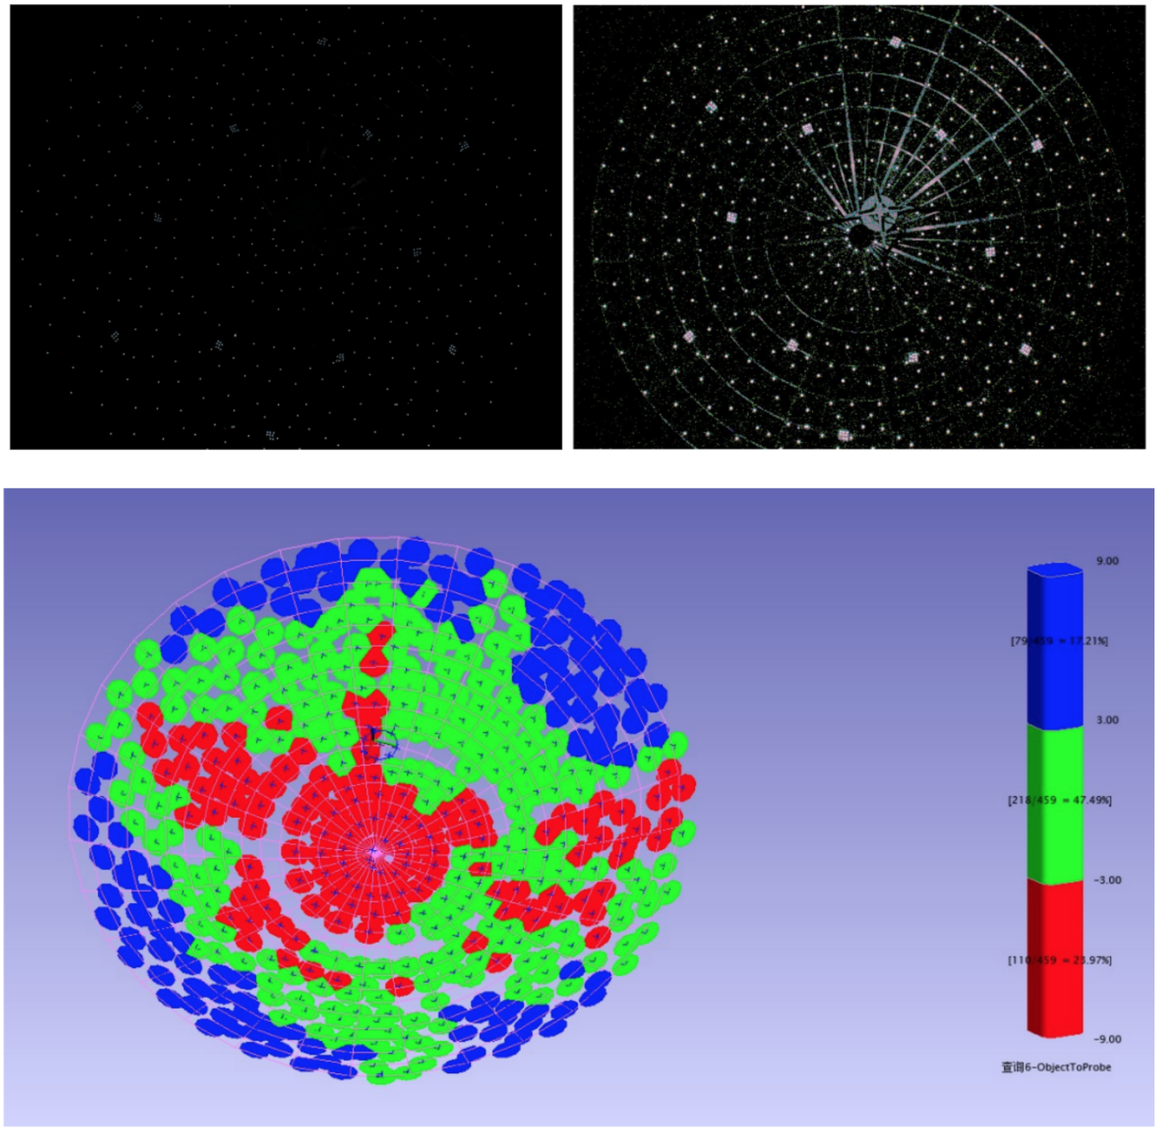

Figure 2.10: One of the many raw images used for photogrammetry (top left), an edited version of the left image for illustrative effect (top right), results of 3D dish model from synthesis of all images as processed by the VGOS collaboration, showing 5mm surface deviations (bottom).

used for the CHIME FRB outrigger prototype at Green Bank Radio Observatory, which was a useful testbed for the drone calibrator development described in this thesis (see Chapter 6). Experimental verification of the F/D was therefore necessary; I performed these measurements at HartRAO in 2017.

I performed telescope focal ratio measurements with fellow graduate student Kabelo Kesabonye. We set up a signal generator transmitting at 600MHz through an omnidirectional antenna 150 meters away from the dish under test. We adjusted the position of the dish to face the noise source, and adjusted the position of the antenna mount relative to the focus, monitoring measured power at the telescope and noting when it hit a maximum. We found the focal length to be 236cm, and noted that the signal attenuated to half power when the antenna was moved ±10cm out of the plane of the focus. The experimental result we obtained gave the expected F/D of 0.39 (236cm/600cm = 0.39).

2.4.3 Antenna design and assembly

Each of the 1024 HIRAX dishes will have a dual-polarization antenna similar in design to a clover-leaf dipole. The HIRAX feed[3] consists of FR-4 dielectric (PCB) with a metalized layer, a PCB balun, and a support board where signal summing occurs. A ring choke is used to circularize the beam and decrease crosstalk and ground spillover. The cloverleaf design is based on that of the CHIME antenna [86], with the main differences being that the HIRAX feed (1) is made from FR-4 instead of Teflon-based PCB to reduce cost, (2) integrates a first-stage LNA into the antenna structure to reduce noise, and (3) sums polarization along diagonal petals. The HIRAX feed is described in more detail in Chapter 3. Noise characterization measurements of the feed are described in Chapter 4.

When assembling a HIRAX antenna, the pieces slot in to one another but do not form a secure structure without proper soldering. It is thus effective to utilize a jig for HIRAX feed assembly. Several designs have been developed, including one at McGill University that is machined to hold the feed in place and keep the relevant elements parallel to 0.5mm. Once the feed is in the jig, it is simple to solder adjacent pads together and apply a stripe of solder at the feed/balun joints for added stability. Simulations are underway to place tolerances on assembly.

---

[3] The words ‘feed’ and ‘antenna’ are used interchangeably in this work. ‘Feed’ will generally describe an antenna mounted or to-be-mounted on a dish.

2.4.4 RF chain

The HIRAX signal chain is comprised of both custom and commercial parts. As described above, a first-stage LNA of $\sim$40dB of gain is integrated directly into the antenna structure for noise reduction. The signal is then filtered, transformed from RF to optical using an RF-over-Fiber (RFoF) module[87], amplified, and carried kilometer distances to the correlator building. There, it is amplified once more, filtered again, and channelized and digitized by an ICE board[88]. A schematic of the signal chain and photos of associated hardware components are shown in figures 2.11 and 2.12.

The RF-over-Fiber (RFoF) works in the following way: Radio frequency photons modulate an optical laser, embedding the radio wave form in optical amplitude fluctuations; the signal travels down fiber and is then demodulated back to radio before the correlator. The RFoF is a critical piece of the set up, as carrying signals over coaxial cabling (which was utilized in the early-stage HIRAX prototype) would have distinct disadvantages: coaxaial cabling such as LMR-400 attenuates by 10dB per 100m in the HIRAX band, leading to complications over km distances. Additionally, the loss in the cable has a frequency dependent slope of roughly 4dB per 100m of cable, which makes it challenging to measure the full band without compressing the analog electronics (which occurs when significant input power is provided, driving gain responses nonlinear) or losing the signal of interest in noise.

The RFoF module has a gain between 40-50dB with a flat response, and isolation from reflections at the $< -20dB$ level. The 10dBm compression point is substantially above the amplified sky signal and the correlator input requirement (-23dBm). It follows significant amplification, such that the noise contribution to the overall system

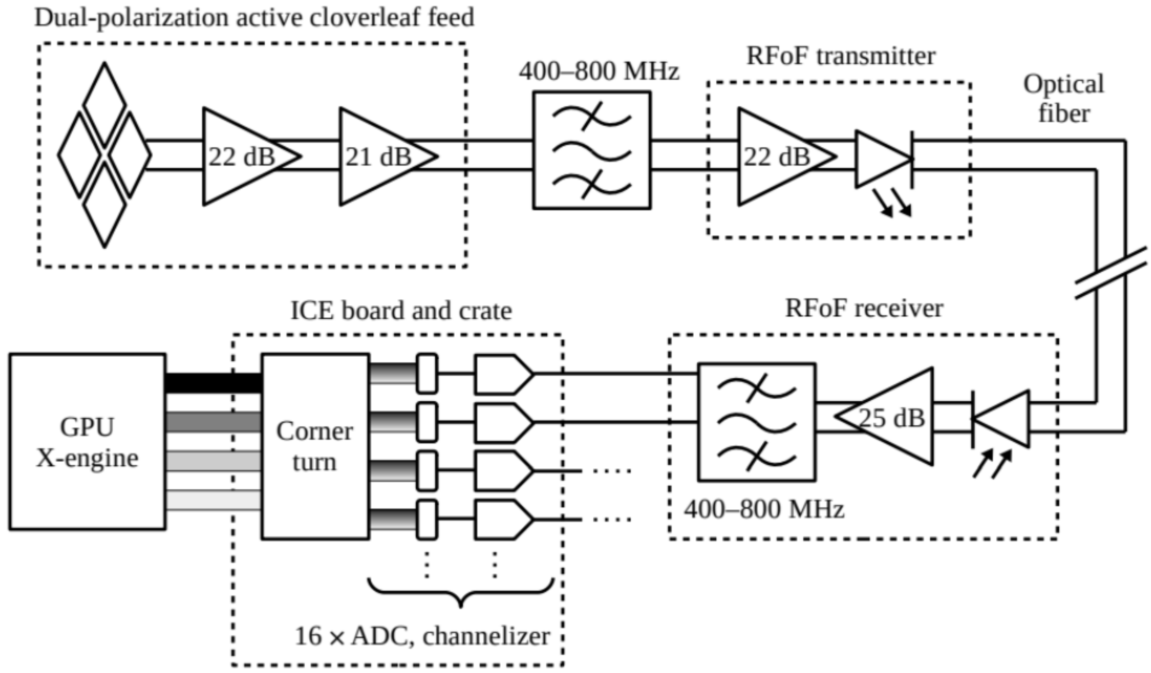
Dual-polarization active cloverleaf feed
22 dB
21 dB
400–800 MHz
RFoF transmitter
22 dB
Optical
fiber
ICE board and crate
RFoF receiver
GPU
X-engine
Corner
turn
25 dB
400–800 MHz
16 × ADC, channelizer

Figure 2.11: The HIRAX RF chain. The polarized sky signal is collected by the dish and focused into the dual-polarized feed. First stage amplification of 43dB happens within the antenna structure as a noise reduction technique. The amplified signal is filtered, amplified again, and then converted from RF to optical in the RF-over-Fiber (RFoF) modules that sit behind the antenna. The optical signal travels over fiber to the correlator hut, where it is transformed back to RF, amplified and filtered once more, and finally digitized and correlated. There is active work within HIRAX on all parts of this RF chain, with different groups dedicated to antenna design and characterization (see Ch 3, 4), RFoF design and characterization, and different aspects of the correlator. All hardware in the chain should be repeatable between dishes, and the general components will be individually assessed for RFI contributions. Figure from [62].

temperature should be percent-level.

2.4.5 Correlator

The HIRAX correlator is responsible for the telescope data taking, processing, and storage. It consists of dozens of rack-mountable components stored in a temperature-regulated area, and is based on the CHIME correlator design[88].

The correlator can be thought of in two separate pieces, commonly referred to as the F-engine and the X-engine. The F-engine is where the signals from 1000 individual

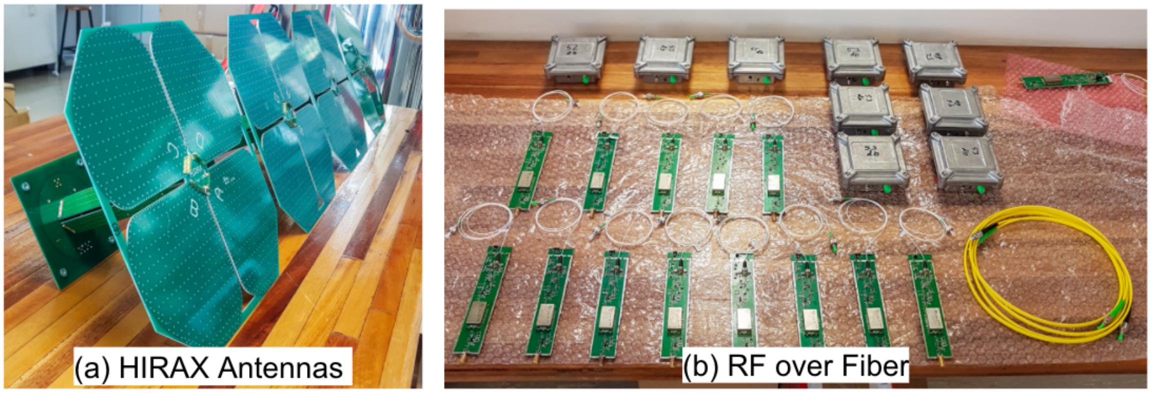

(a) HIRAX Antennas (b) RF over Fiber

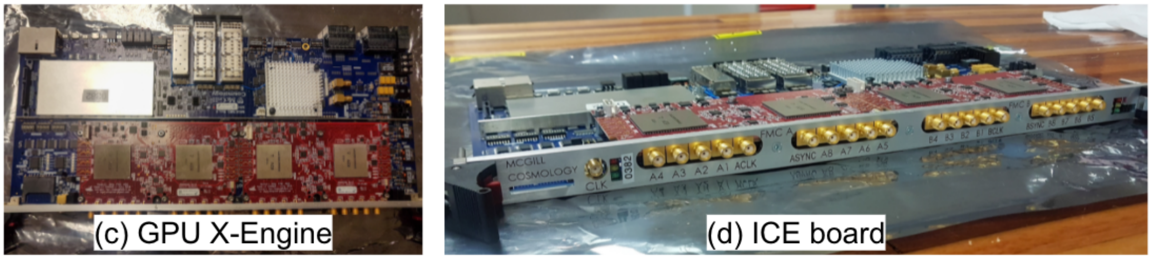

(c) GPU X-Engine (d) ICE board

Figure 2.12: Select hardware from the HIRAX RF chain. (a) The HIRAX antenna feeds. The feed design is based on the CHIME antenna, but has the first stage amplification integrated into the antenna structure. This reduces feed noise and makes it more lightweight, but in turn prevents the noise temperature from being directly measured with a noise figure meter, which will be explored more in depth in Ch. 4. (b) The RFoF modules, developed between WVU, CMU and UKZN, will convert RF signal to optical so that it can more efficiently and cheaply travel long distances. (c,d) Different components and angles of the correlator hardware. Photos from the HIRAX collaboration.

dishes are received, digitized, and channelized using an FPGA. Specific hardware specs and parameters for correlator nodes are given in table 2.5. Correlation happens on the GPUs (X-engine). HIRAX will use two separate back-ends for cosmology and FRB/transient searches, given the data volume required for high time-resolution sampling (which is needed to see the FRB signature).

The correlator operates by stacking 8-bit frames, taken at a rate of 16384 frames per 43ms period, to achieve a total number of digital counts. These counts are gain corrected to remove the instrument response. The sensitivity of the frames is set by an initial gain calibration that places the spectrum within the digital range of the correlator, in order to sensitively observe sky fluctuations. A consequence of the frame

Parameters	**Value**
Motherboard	GIGABYTE G482-Z52
Processor	2 × AMD EPYC™ 7452, 32 cores each
RAM	1TB
GPU	2 × NVIDIA A40, PCIe 4.0
F-Engine Network	4 × SILICOM PE31640G2QI71-QX4 - 2 × 40 Gbps
Outgoing Data Network	2×25 Gbps

Table 2.5: Specifications for HIRAX correlator nodes. The HIRAX-256 X-engine is comprised of eight these nodes.

stacking is that small signals that only periodically register as 1 bit out of the 8 may appear artificially suppressed, and that calibrating the instrument with bright signals can easily digitally saturate the instrument. These details were most relevant when applying the correlator to contexts outside of HIRAX science goals. Such applications include antenna noise temperature measurements (Ch. 4) and drone calibration (Ch. 7, 6).

A consequence of the correlator digitization is that, without proper filtering, frequencies outside of the HIRAX band can be aliased in, and impact the measurement results. An example of this is a GPS base station, which transmits brightly from 900-930MHz, flipping over the 800MHz correlator boundary and appearing in the measurement as bright 670-700MHz RFI. Thus, it is useful to operate in a clean RFI environment with considerable filtering.

2.5 Summary and outlook

HIRAX is a revolutionary new radio telescope to be built in South Africa. There is a planned and funded deployment of 256 dishes in the near-term, as a stepping stone towards the full 1024 element array. HIRAX will look near when the universe shifted from matter to dark energy dominated, a redshift range that is relatively unexplored; it will provide a new probe to understand BAOs, at different redshifts and with different tracers than prior optical surveys; and it will dramatically expand the detection and understanding of FRBs.

Over the course of my time in graduate school, I have worked to develop the HIRAX instrument, as summarized in this chapter. My particular contributions involved several fieldwork campaigns and participation in instrument design working groups. I fielded and debugged antenna and amplifier instrumentation in South Africa, performed the first photogrammetry measurements of HIRAX prototype dishes, and measured important telescope specifications such as dish F/D and beam width.

The majority of my PhD has been devoted to developing custom calibration hardware for HIRAX, as described in Chapters 3, 4, 5, and 6. This chapter motivated the need for such hardware, by describing the stringent instrument requirements placed on HIRAX in order to meet its lofty science goals.

Chapter 3

Antenna beam measurements

For antennas that will be used for sensitive cosmological measurements, understanding the beam pattern is a critical characterization step. This chapter discusses beam measurements of a variety of antennas used both in the HIRAX instrument and for calibration of the HIRAX array. It begins with an introduction to antennas and antenna measurement ranges, and then describes antenna range measurements taken in October 2019 at Jet Propulsion Laboratory, March 2020 at North Carolina State University, and August 2021 at Green Bank Observatory. The results reported in this section include the first measurements of the HIRAX antenna beam, which provided a critical check on simulations, and important characterization of the telescope calibrator beam.

3.1 Introduction to antennas

Antennas are devices that transmit and receive radio signals. An intuition for antenna operation can be built off of a simple qualitative picture, and it is helpful here to first consider the case of a dipole (fig. 3.1).

A dipole antenna is the simplest wire antenna model. It consists of two wire arms of total length $\lambda/2$, where λ is the target measurement wavelength. When the

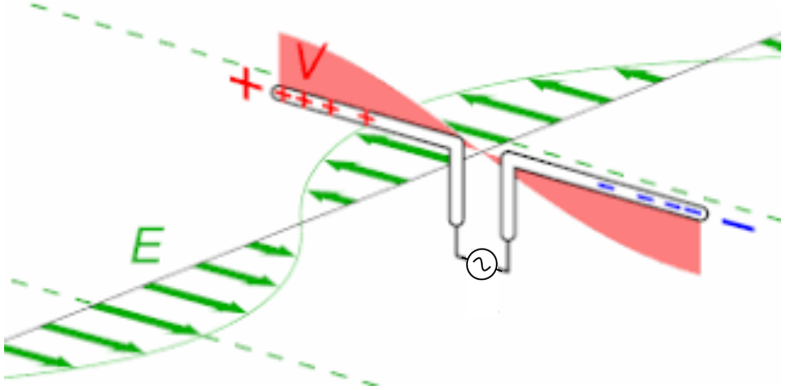

Figure 3.1: Illustration of a $\lambda/2$ dipole antenna, and the distribution of surface charges in the presence of an electric field. As the electric field oscillates, the surface charges move, generating measurable currents and voltages that oscillate with the frequency of the wave. Figure from Wikipedia.

crest of an incoming wavefront contacts a dipole antenna, a force is exerted on the surface charges that pushes them to one end of the device. The wave continues to propagate, and as it travels by a half wavelength, the wave trough contacts the dipole surface and charges are pushed to the other end of the device. This movement of charge induces time-varying voltages and currents. A measurement of the frequency with which voltage oscillates at the antenna terminals will give the frequency of the incoming wave. The same intuition holds for transmitting a signal, and we recall from fundamentals of electromagnetism [89] that electromagnetic radiation is created by accelerating charges and time varying currents. We can think of a transmitting antenna as receiving voltage oscillations which result in oscillating currents on the dipole arms, which radiate electromagnetically. An alternative way of thinking of antennas is as devices to couple freespace into a waveguide and vise versa, or impedence match between free space and a waveguide (often with an impedance of 50Ω).

Antennas and radio telescopes have spatial dependence to how they radiate and receive signals. The distribution of this radiation is referred to as a beam pattern. A generic illustration is shown in figure 3.2.

Generally, antennas have maximum sensitivity along one axis, which falls off in

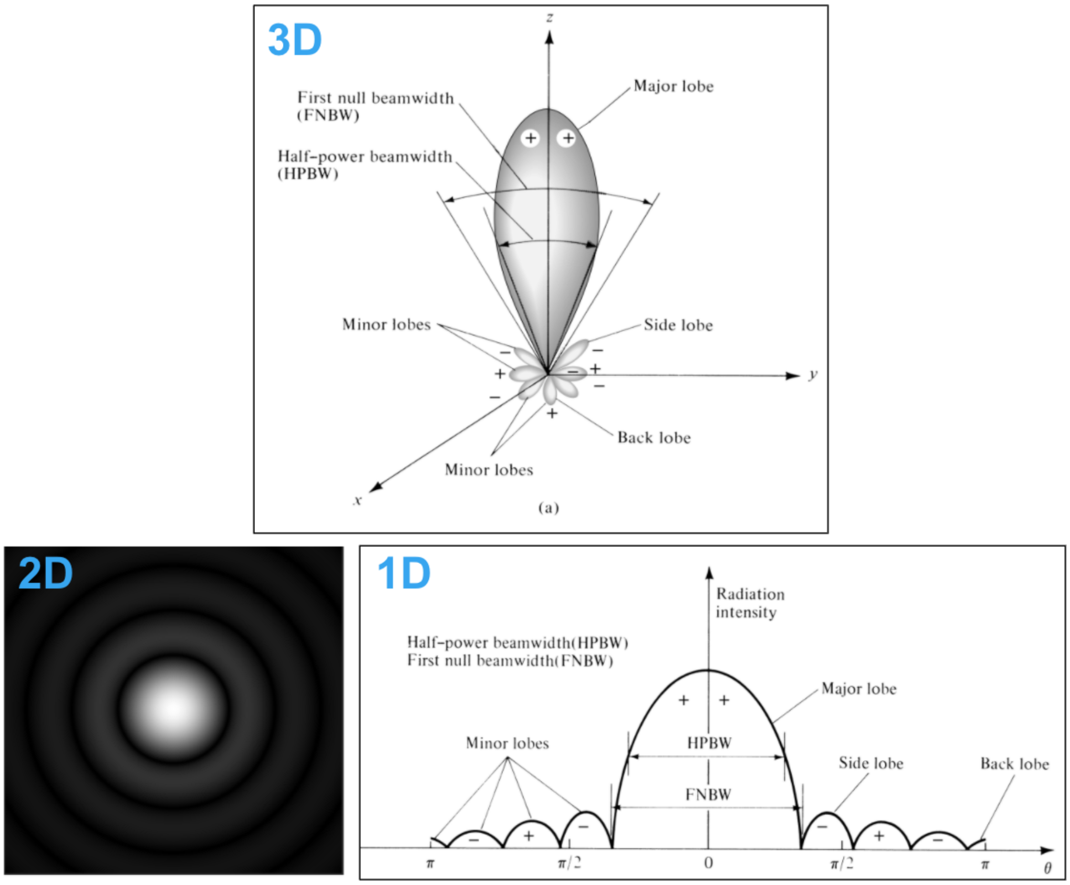
3D
z
First null beamwidth (FNBW)
Half-power beamwidth (HPBW)
Major lobe
Minor lobes
Side lobe
y
Back lobe
Minor lobes
x
(a)
2D
1D
Radiation intensity
Half-power beamwidth(HPBW)
First null beamwidth(FNBW)
Major lobe
Minor lobes
HPBW
FNBW
Side lobe
Back lobe
π
π/2
0
π/2
π
θ

Figure 3.2: Illustration of a common antenna beam pattern, shown in 3D, 2D and 1D. Figure from Antenna Design and Theory by Balanis [90].

each dimension as an approximate Gaussian of the form,

$$g(x) = A \exp\left(-\frac{(x - x_0)^2}{2\sigma^2}\right) + D \tag{3.1}$$

where A is the amplitude of the Gaussian peak, D is an offset, x_o is the curve centroid position, and σ represent beam width. This gaussian region is called the “main lobe”, and is depicted in figure 3.2 as pointing along z. It is also common to discuss the antenna or telescope beam width, which is the full-width-at-half-maximum (FWHM) of the gaussian (referred to in figure 3.2 as the half power beam width). We can relate

σ to FWHM using,

$$\text{FWHM} \approx 2.355\sigma. \tag{3.2}$$

Off of the main lobe, there will often be pockets of significant sensitivity out at wide angles, referred to as sidelobes. An antenna with high-level sidelobes could be zenith pointed and intending to measure a celestial source, but accidentally measure bright terrestrial radio signals or other sources at wide angles, with no way to distinguish where the flux is originating if proper calibration has not taken place. Thus, to measure the sky sensitively, we will need an understanding of the sidelobe patterns.

Antenna beam patterns are generally expressed in logarithmic units [dB], with the level often normalized to the peak of the beam, in order to make low-level beam features more visible. Patterns may also be expressed in units of dBi, which simply normalizes the beam pattern to that of an isotropic resonator. Common units used throughout this section and thesis include dB, dBi, and dBm, which are summarized in table 3.1.

Unit	**Description**
dB	Decibels; a way of expressing a ratio (X) in logarithmic units. $[\text{dB}] = 10\log_{10}(\text{X})$
dBi	dB relative to an isotropic resonator; used to express antenna gain.
dBm	Decibels relative to 1mW; a way of expressing power in logarithmic units. $\text{P [dBm]} = 10\log_{10}(\text{P [mW]})$

Table 3.1: Common units in antenna measurement.

Antennas are often designed to work with a ground plane. A ground plane is typically a metal sheet that sits at $\lambda/4$ behind the antenna phase center. This placement allows for destructive interference of radiation in the antenna backlobe, thereby shaping the beam into being more forward-directed, with small side and back lobes. Ground planes can be formed into different shapes—the HIRAX antenna in particular employs a can structure as the back plane, which limits the back lobe and circularizes the beam.

Common RF measurement tools

There are two main antenna measurement tools that will be utilized in this section and through much of this thesis: a spectrum analyzer and vector network analyzer (VNA). A spectrum analyzer measures input signal power as a function of frequency, while a VNA assess the reflection and/or transmission of signals through voltage measurements as a function of frequency. A common use of VNAs is to make measurements of the amount of transmission/reflection between ports on some device(s) under test (DUT). These measurements are known as S-parameter measurements, and are indexed by the ports that transmission is measured between [91]. For a two port system, with V_n^+ as the amplitude of the voltage wave incident on port n and V_n^- as the amplitude of the voltage wave reflected from port n, the scattering matrix $[S]$, is defined as follows,

$$\begin{bmatrix} V_1^- \\ V_2^- \end{bmatrix} = \begin{bmatrix} S_{11} & S_{12} \\ S_{21} & S_{22} \end{bmatrix} \begin{bmatrix} V_1^+ \\ V_2^+ \end{bmatrix} \tag{3.3}$$

S_{11}, the reflection coefficient, is a commonly employed measurement in this thesis, often discussed in logarithmic units,

$$S_{11}[dB] = 10 \log \left(\frac{V_{\text{incident}}}{V_{\text{reflected}}} \right). \tag{3.4}$$

It is a helpful assessment of antenna internal reflections, with an antenna that has $S_{11} < -10$dB in freespace considered functional. If an antenna is situated in a reflective cavity, the expected S_{11} is ∼0dB, as all signal should be returning to the antenna. S_{21} measurements, which show transmission from port 2 to port 1 of a VNA, are helpful for measuring amplifier gain, and can be used across two different antennas to measure beam pattern, as is done for the anechoic chamber beam measurements described later in this Chapter.

Antennas both transmit and receive radiation with the same angular sensitivities, as noted in Burke and Smith (1997)[92]: "An antenna can be treated either as a receiving device, gathering the incoming radiation field and conducting electrical signals to the output terminals, or as a transmitting system, launching electromagnetic waves outward. These two cases are equivalent because of time reversibility: the solutions of Maxwell's equations are valid when time is reversed." As a consequence of reciprocity, beam measurements can be performed in either transmit or receive mode. For the types of measurements performed in this section, the antenna under test is receiving.

3.2 Introduction to antenna range measurements

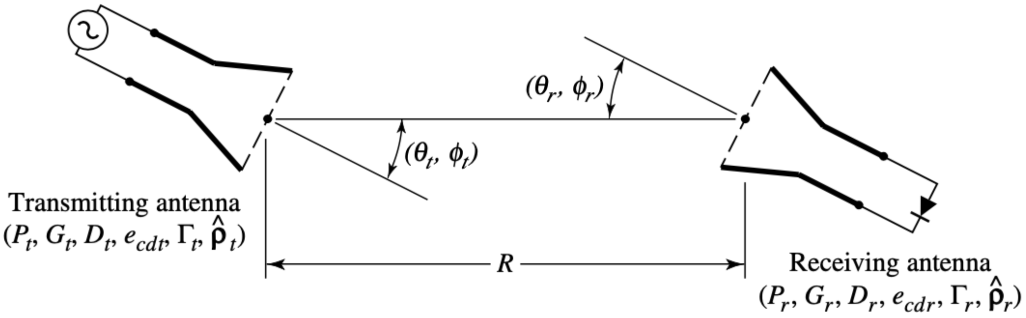
(θ_r, ϕ_r)
(θ_t, ϕ_t)
Transmitting antenna
$(P_t, G_t, D_t, e_{cdt}, \Gamma_t, \hat{\rho}_t)$
R
Receiving antenna
$(P_r, G_r, D_r, e_{cdr}, \Gamma_r, \hat{\rho}_r)$

Figure 3.3: Illustration of an antenna transmission measurement, in which the Friis equation is applied. P_t is the power transmitted from some transmitter antenna, and P_r is the power received. Figure from [90].

Before describing antenna ranges, it is helpful to explore how an antenna transmission is measured. For this, we turn to the Friis equation, a staple of radio frequency engineering[90]. The Friis equation describes how power transmitted from one antenna (P_t) is received by another (P_r), and is of the form,

$$\frac{P_r}{P_t} = D_t D_r e_t e_r \left(\frac{\lambda}{4\pi d}\right)^2 |\hat{p_t} \cdot \hat{p_r}| \tag{3.5}$$

where: D_r, D_t are antenna directivities, e_r, e_t are antenna efficiencies, λ is the measurement wavelength, d is the distance between the antenna phase centers, and $\hat{p_t}, \hat{p_r}$ are polarization unit vectors. Directivity is defined as the power density radiated by an antenna in the direction of strongest emission. Figure 3.3 illustrates an antenna measurement using the Friis equation.

Using the Friis equation can tell us directly how much of the power transmitted by some calibrated antenna should be measured at a receiving antenna under test (AUT). Antenna transmission is measured between the antennas, which enables solving for D at a variety of different antenna orientations. From these measurements, a 1D beam is constructed, giving antenna sensitivity as a function of angle from boresight. Directivity is defined as a single number (rather than a quantity that varies with angle), which makes sense historically: typical radio applications call for transmitting a signal as far as possible, for which one only must consider the most high-gain part of a beam. For the purposes of understanding antenna transmission in 2D, however, we expand this picture to all beam angles. For this thesis, we use angularly varying directivity and antenna gain interchangeably.

The full mechanics of such a measurement vary depending on the type of antenna range, but generally proceed as follows. The AUT (receiving antenna, Rx) is mounted on one end of the range opposite some transmitter (Tx). The AUT sweeps 360 degrees about its positioner and measures the received power at a variety of angles (determined by the desired angular resolution). The Friis equation is utilized to extract the antenna gain (directivity in figure 3.3) as function of angle. The overall result is a plot of received power as a function of antenna orientation. Measurements can be made at individual frequencies or made broadband, depending on the type of antenna range.

Typically, antenna ranges do not measure a full 2D beam and instead focus on two orthogonal 1D cuts through the beam. The most common cuts are the E- and H-cuts, where the E-cut is the beam inline with the electric field oscillations and the H-cut is the beam in line with the magnetic field (figure 3.4). As a rule of thumb, the E-plane will typically be parallel to the radiating element— for a dipole this is the plane that contains the two arms—and the H-plane will be orthogonal to E. We will generally expect gaussians from the dipole-style antennas used in this thesis in the E-plane, and broader beams in the H plane. These cuts will be measured for both co-polarization, where the transmitting and receiving antenna are aligned in polarization, and cross-polarization, where the transmitting and receiving antenna are orthogonal. Most of the power is transmitted in co-polarization for a polarized antenna, but measuring cross-polarization is still important, to avoid leakage of unexpected signal into polarized measurements.

3.2.1 Types of antenna ranges

There are several varieties of antenna measurement facilities, but in this chapter we will focus on only two: anechoic chambers and outdoor far field ranges, both of which measure the far-field beam pattern directly.

Anechoic chambers are indoor facilities. They typically consist of a large Faraday cage lined with absorptive foam. An antenna inside of an ideal anechoic chamber will find the conditions indistinguishable from free space at a 300 Kelvin temperature. The controlled radio environment allows measurements to be broadband, as there are no stray signals or reflections to contend with. A limitation of this facility is that the indoors can be confining, and can restrict the size of the antenna under test, particularly if its farfield exceeds the room size. Size limitations also impact the ability

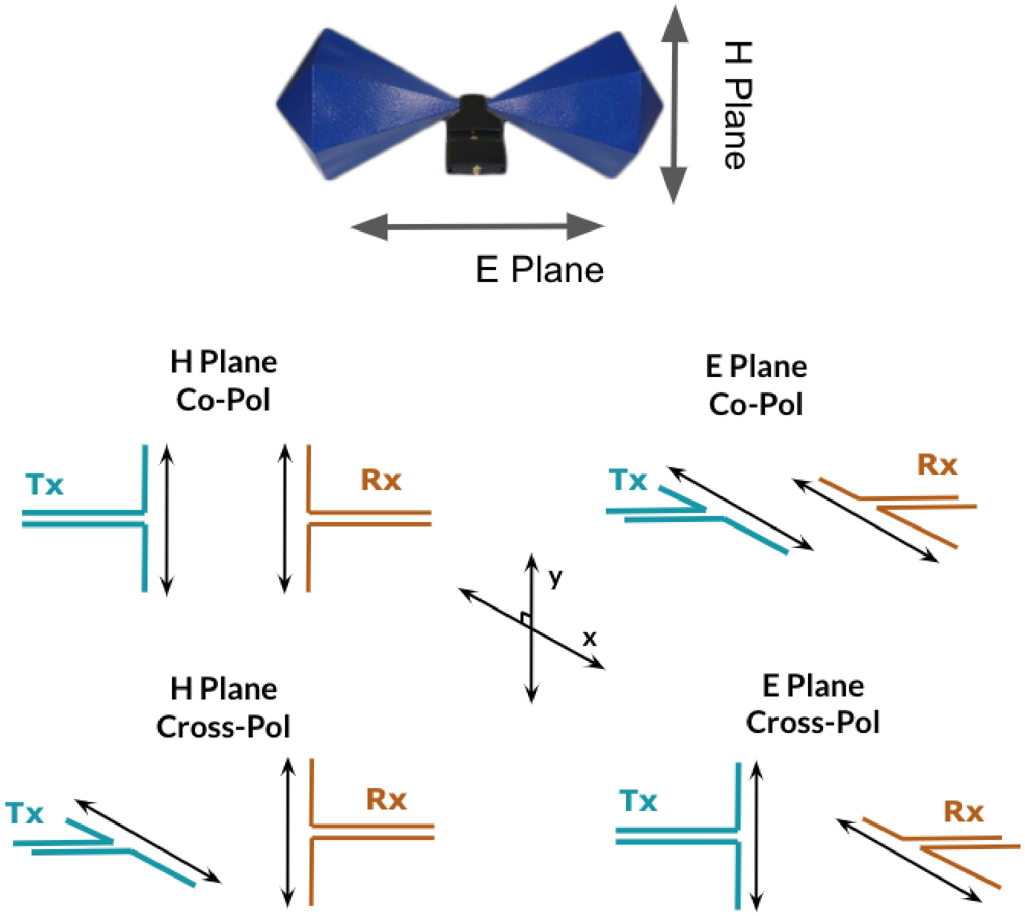
H Plane
E Plane
H Plane
Co-Pol
Tx
Rx
E Plane
Co-Pol
Tx
Rx
y
x
H Plane
Cross-Pol
Tx
Rx
E Plane
Cross-Pol
Tx
Rx

Figure 3.4: Illustration of E- and H-plane cuts for co- and cross-polarization of a cartoon dipole antenna, as measured with transmitter Tx and receiver Rx. The top image shows the E and H planes for a biconical antenna, discussed later in this chapter. The black arrows indicate polarization direction.

to characterize low-frequency hardware, whose sizes can swamp room capacity and whose wavelengths can exceed the specs of commercially available absorber, which has pyramid dimensions that scale with wavelength.

Outdoor antenna ranges, by contrast, are free from space and absorber restrictions and allow for the characterization of big equipment. Environmental conditions do, however, typically motivate a hardware setup where only one frequency is measured at a time. This narrowband signal is pulsed with a lock-in amplifier to account for

RFI contamination from surrounding areas. Outdoor measurements also have the challenges of reflections from the environment, including ground features and nearby buildings. These can be accounted for by exploiting a simple technique in which after a measurement is made, the AUT is rotated by 180^o and the measurement repeated. If asymmetries flip in the resulting beam pattern, they are from the feed. If not, they are artifacts of the range.

Beyond the environmental conditions, an additional consideration with antenna ranges is the measurement device that is used. Outdoor ranges often take power measurements directly, transmitting a modulated signal with a lock-in amplifier. Indoor ranges will often use a Vector Network Analyzer (VNA), which takes voltage measurements. It is critical to understand if your raw data is in voltage or power, especially if comparing two measurements, as

$$P \propto V^2. \tag{3.6}$$

A common error with beam measurements is overplotting measurements or simulations with different units in log space, where they will differ by a factor of 2 due to the V^2 factor.

The following sections detail measurements of several antennas used by the HIRAX collaboration and the three different antenna ranges used for characterization.

3.2.2 The JPL antenna range

Jet Propulsion Laboratory (JPL) in Pasadena, California hosts a wide variety of world-class instrumentation facilities, one of which is its Mesa Antenna Range[1]. As the name suggests, the Mesa Antenna Range sits on a flat-topped hill (known as a mesa) on the boundary of the JPL campus, overlooking the city of Los Angeles. This facility is made up of a combination of indoor and outdoor, near-field and far-field testing facilities, including the 1200ft range which I used for my first round of antenna measurements in October 2019 (figure 3.5). This was an effort that I led as part of my NSTRF Visiting Technologist Experience, with support and collaboration from Dr. Tzu-Ching Chang, Dr. Ben Saliwanchik, and many JPL staff.

---

[1] https://mesa.jpl.nasa.gov/

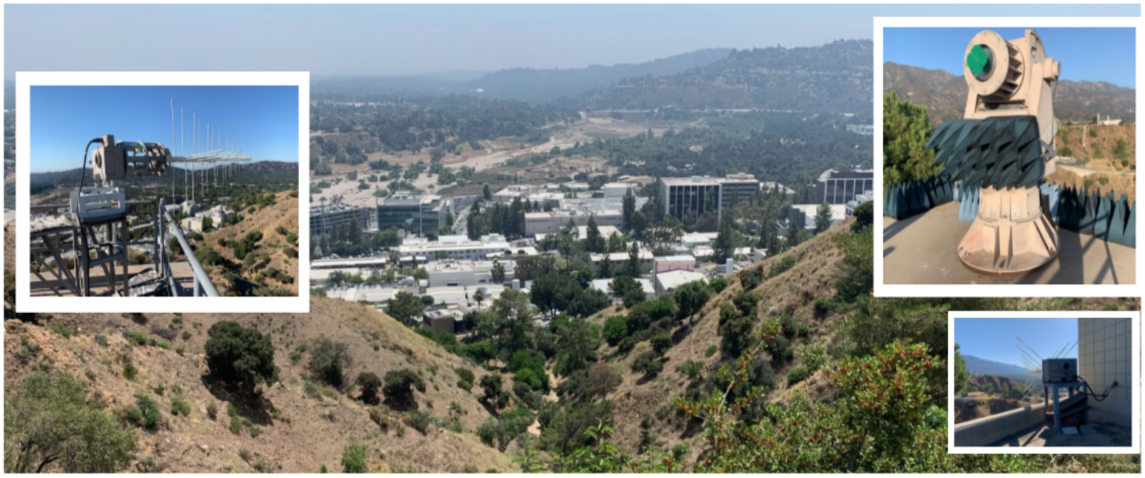

Figure 3.5: The Mesa Antenna Range at Jet Propulsion Laboratory in Pasadena, CA. The background photograph is the canyon which sits beneath the signal path, and the outlined photographs are the transmitter yagi antenna (left), the range mount and antenna under test (upper right), and the reference antenna (lower right).

The JPL outdoor range set up is as follows. There is a transmitting antenna (in our case, a log periodic antenna that is directional and broadband, covering our 400-800MHz frequency range) situated across a canyon from a reference antenna (also log periodic) and our AUT. The AUT is mounted to a three-axis test positioner (this flexibility allows for full raster scans, which have the capability to produce a 2D beam, though we restricted to measuring 1D cuts due to time considerations). The transmitter sends out amplified power from a signal generator at one frequency, which experiences ∼80dB of path loss across the 1200ft before locking in with the reference. This locking is what allows us to make measurements outdoors: the range's urban location introduces considerable RFI, from which the signal of interest must be teased out. In addition to the RFI, reflections from the canyon are a limitation of this set-up. These reflections can be better understood by repeating beam measurements that show asymmetry with the AUT rotated 180^o– if the beam shape does not change between two measurements, the asymmetric features are due to the range, but if they change in location by a sign (flip over the $\Theta = 0$ axis), they are true features of the beam.

At JPL, I primarily took measurements at three frequencies, 400MHz, 650MHz and 800MHz. Four beam slices were measured at each test frequency: two co-polarization (co-pol) cuts, and two cross-polarization (cross-pol) cuts (as described in figure 3.4). I measured four antenna configurations: (i) the passive HIRAX feed, (ii) the active HIRAX feed (with 50dB of gain), (iii) the biconical antenna mounted on the drone, and (iv) the biconical antenna on its own. These were compared to measurements of a standard gain horn for gain calibration, and were measured relative to a reference antenna to get phase. Each cut took roughly two minutes to complete, and provided resolution in half degree increments. Measurement results from this visit included too many systematics to be published, but working in this environment was incredibly instructive, and set the stage for future productive range measurements. It was through those measurements that we diagnosed that the biconical antenna transmitter had a severed wire (a common fail point with this particular brand of antenna) and that a drone fiber body is indeed quite reflective in 400-800MHz, which is relevant to later work in the thesis (section 3.6).

3.2.3 The NCSU anechoic chamber

The Nano Fabrication Facility (NNF)[2] at North Carolina State University (NCSU) hosts an anechoic chamber for UHF and L-Band antenna testing (photographs in figure 3.6). In contrast to the JPL 1200ft range, where the set up required a power measurement localized to one frequency, the NCSU set up utilizes a Vector Network Analyzer, which yields broadband measurements. This is made possible by the RF-quiet conditions within the chamber where the measurement takes place. Thus, data taking in this style of set-up simply requires a single sweep per beam slice to get broadband frequency information.

---

[2] https://nnf.ncsu.edu/

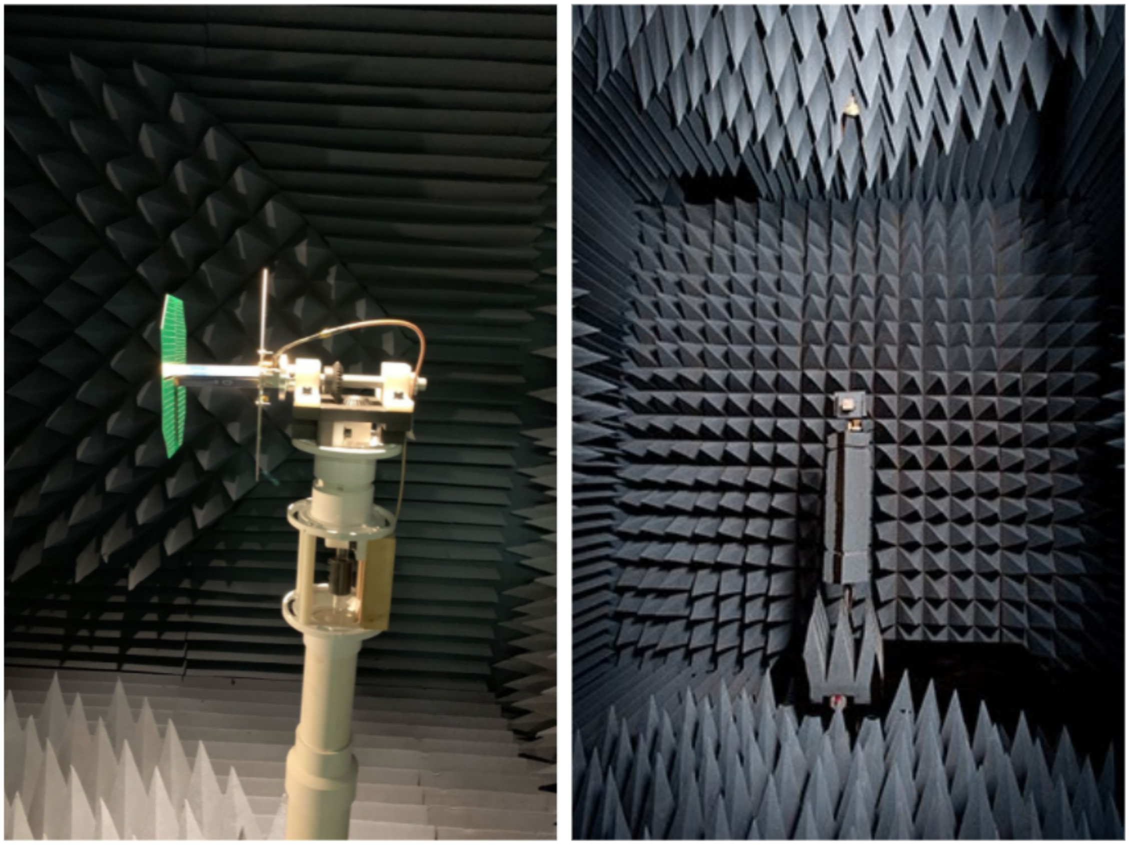

Figure 3.6: An anechoic chamber in the North Carolina State University Nanofabrication Facility (NNF), used for antenna beam measurements. Left image: a HIRAX antenna mounted in the chamber, Right image: the standard-gain horn transmitted mounted in the chamber (right photo from NNF website).

The anechoic chamber is 10ft by 16ft and lined predominantly with 18inch conical RF absorber. For the HIRAX frequency band (400-800MHz), the set up uses a standard gain horn calibrated for ~700MHz. While this produced excellent RFI conditions, and a highly useful set of measurements, this instrumentation led to certain limitations within the 400-800MHz band. The 18inch absorber used is optimized for GHz frequencies, leading to low-level reflections at the lower end of the HIRAX band. Additionally, the horn used for calibration is optimized for ∼700MHz, so calibration of the system breaks down in the lower part of the HIRAX band. This makes for noisier and less-reliable data below 500MHz. The positioner for this set up had been recently re-designed as an undergraduate project, and was made of PVC and 3D printed gears. The plastic was not sufficiently rigid to prevent visible lateral movement of the antenna during measurements. Still, the plastic mounting design is helpful in that it keeps metal, which is reflective and could impact the beam pattern, away from the antenna during tests.

At NCSU, I took more than 100 scans in total over three days of measurements. Scans had 3 degree resolution, and were focused mostly on co-pol cuts, though a full set of cuts was measured for at least one feed in each configuration. This gave the first beam measurements of the HIRAX feed, which was a critical check on simulations. Measurement results are described in more detail throughout the chapter.

3.2.4 The GBO antenna range

The far field antenna range at Green Bank Observatory (GBO) in Green Bank, West Virginia shares similar features to the JPL antenna range in that it measures cuts in individual frequencies and is located outdoors. At GBO, the transmitter and receiver are mounted atop 32ft tall towers located 48ft apart (fig. 3.7). A safety railing on the AUT tower retracts for measurements, so as not to cause reflections, which improves measurement quality but makes changing feed configurations challenging. Measurements on this range are also limited by the presence of the Jansky Lab building at the other end of a parking lot. Reflections off of this building can introduce artificial asymmetries into the measurement.

The scale of this range allows larger, heavier antennas to be measured that could not fit into an indoor anechoic chamber facility, such as a 4ft drone with an antenna attached fig. 3.7).

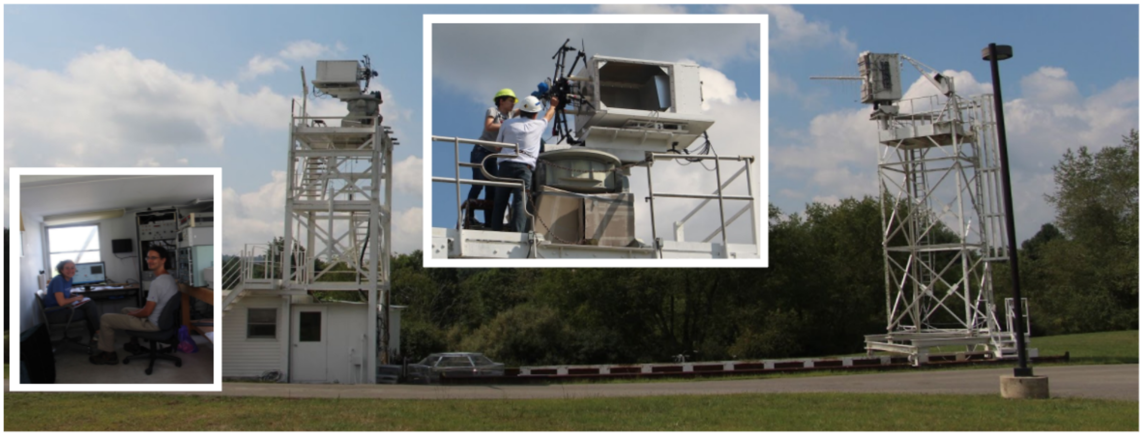

Figure 3.7: An antenna range at Green Bank Radio Observatory. Inlaid photos include (1) the measurement control room and (2) mounting the antenna under test, in this case a biconical antenna attached to a drone.

Now that we have a handle on antennas and antenna measurements, we can discuss beam pattern results for the HIRAX and CHIME antennas as well as the Aaronia biconical antenna, which is assessed in several configurations, including mounted to a drone. These results will be the focus of the remainder of the chapter.

3.3 HIRAX feed measurements

The HIRAX feed is a modified four-square antenna (photo in figure 3.8). It consists of four petals, where opposing petal pairs form a dipole-style antenna. In contrast to the 1D-nature / wire-nature of a typical dipole, the petals on the HIRAX feed have rounded edges, and this distribution of length scales gives a smooth, broadband frequency response. The HIRAX design embeds a first-stage differential low-noise amplifier (LNA, Avago MGA–16116) directly into the antenna balun for noise reduction (Figure 2.12). This design choice has been adopted to keep the total system noise of HIRAX to less than 50 K, allowing a sensitive measurement of the $\sim 100\mu$K cosmological 21 cm signal[81].

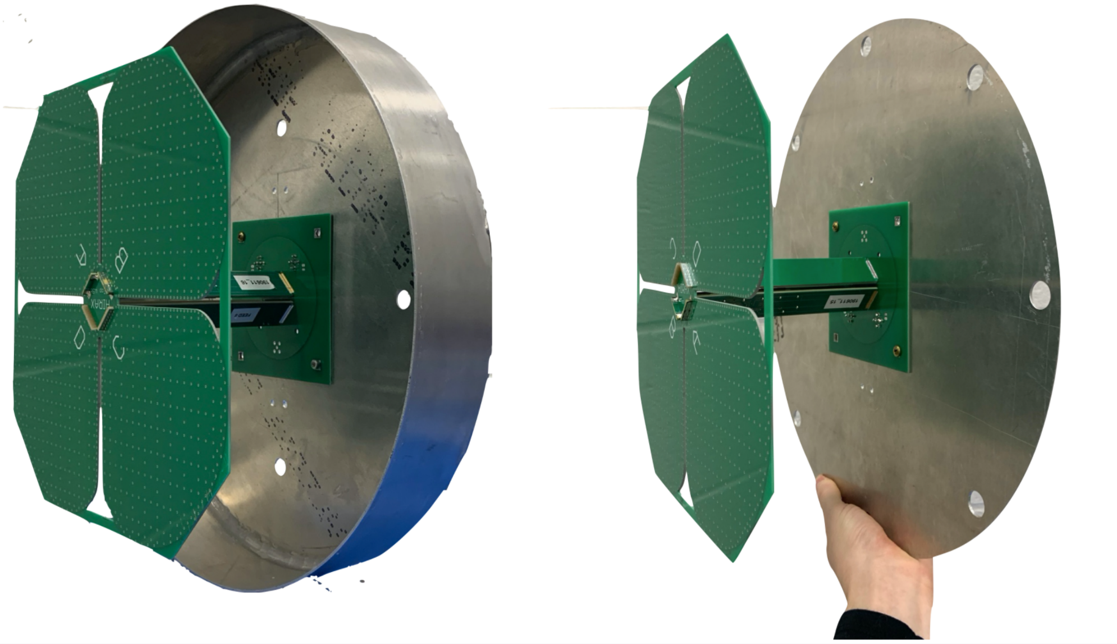

Figure 3.8: The HIRAX antenna, with a can (left) and plate (right) for its ground plane. These are the main mounting configurations discussed in this section.

Measurements of the HIRAX feeds revealed a number of interesting features that are recoverable in simulations. These measurements were taken at the JPL and NCSU ranges. Unless otherwise specified, the HIRAX feed is measured with a can.

HIRAX gain

The broadband set up at NCSU allowed for high-resolution gain measurements of HIRAX. The HIRAX feed measured gain will be a combination of antenna gain resulting from feed geometry and LNA gain. The result is as expected in the band: a roughly 40dB average gain that varies by up to 3dB in frequency but does not consistently slope. This gain measurement was repeated for three different HIRAX feeds, and found to have sub-dB variations. The CHIME feed, which shares core design features with HIRAX and is discussed later in this chapter, has a similar but not identical gain profile. Both profiles are shown in figure 3.9.

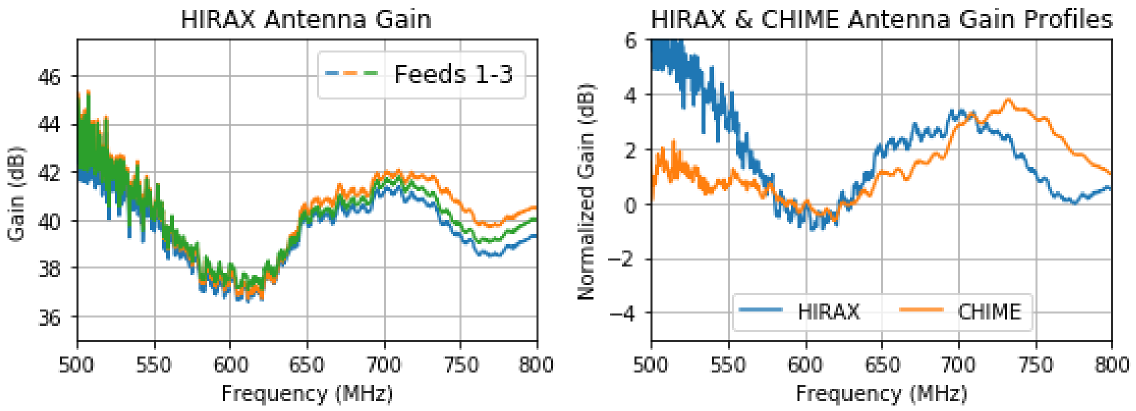

Figure 3.9: (Left:) Gain profiles for three HIRAX antennas with the integrated first stage LNAs and (Right:) gain profiles of the HIRAX and CHIME feeds normalized to the same value at 600MHz for comparison. The HIRAX feed gain is repeatable across the three measured feeds. The CHIME and HIRAX gains share certain features but are qualitatively different (despite having largely similar antenna designs), likely due to the HIRAX LNA bandpass shape.

HIRAX beam results

The HIRAX feed measurements showed frequency dependent beam asymmetries and high levels of cross-pol (figure 3.10), common across different antenna ranges. At lower frequencies ($\nu < 600$MHz), both the E and H-plane cuts appear largely symmetric, and there is roughly 20dB of isolation in the main lobe between co- and cross-pol levels. At higher frequencies, however, the H-plane centroids begin to drift, by up to 25 degrees, and a secondary peak forms within the main lobe. The cross-polarization levels also elevate considerably, and even become comparable in magnitude to co-pol within the main beam. This behavior is recovered in simulations (figure 3.11). The centroid location drift can be recovered in measurements of the antenna mounted on a dish measuring source transits.

There are many possibly drivers of this behavior. A likely explanation is that the LNAs are incorporated into the feed in an asymmetric manner, located on adjacent balun stems rather than opposing stems. The balun structure is supposed to have anti-symmetric currents, which should cancel and result in minimal radiation from the balun structure. If the balun design is flawed, there could be significant cross-pol radiation. The source of this behavior is an active investigation, and if the asymmetries are traced back to the balun geometry, there are a variety of redesigns which balance the currents that can be prototyped.

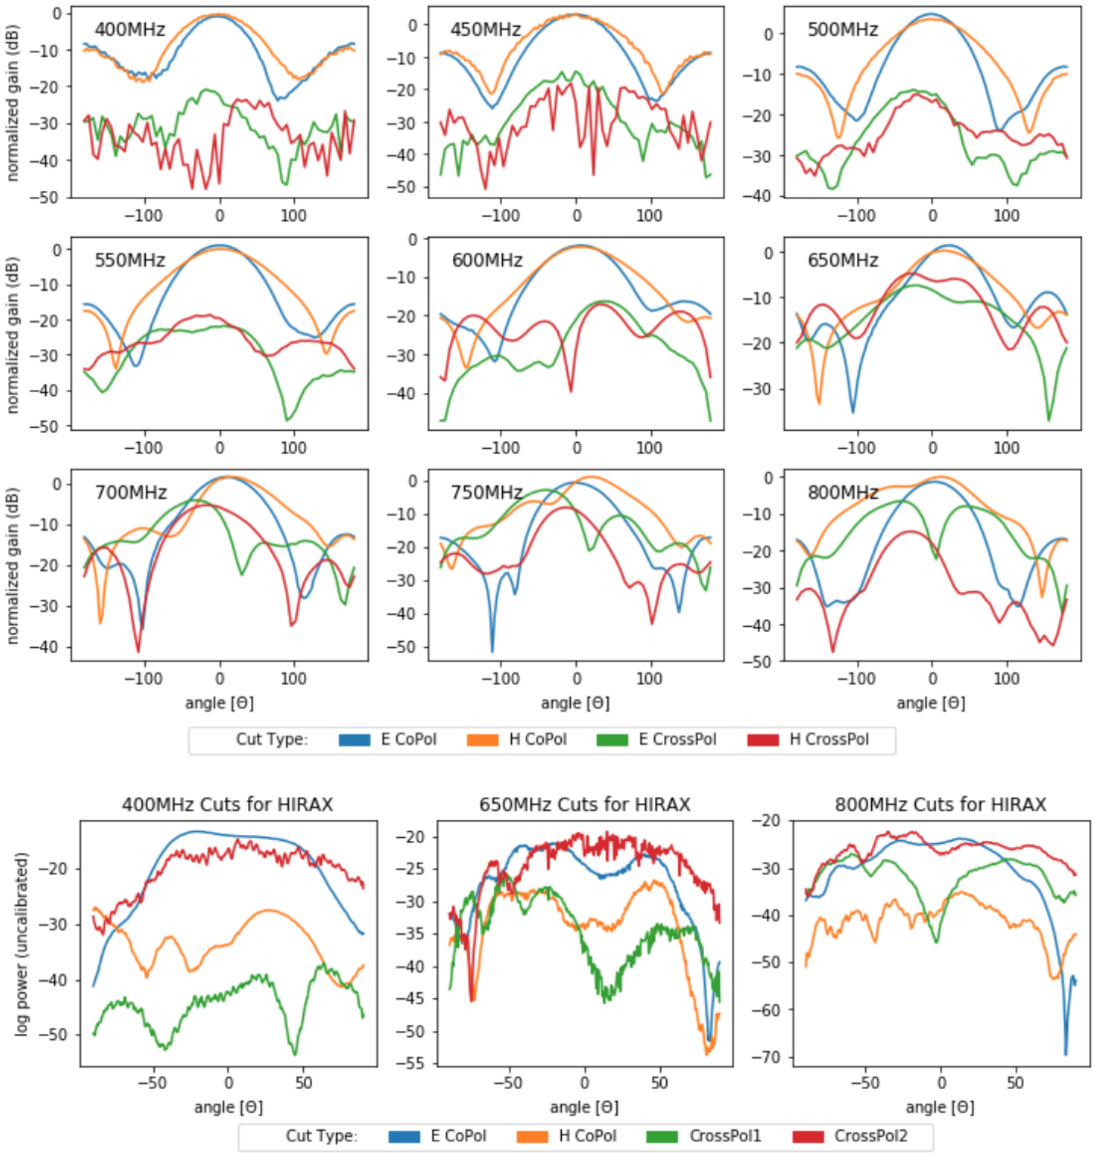

Figure 3.10: Full co and cross-pol cuts in the E and H planes for the HIRAX feed mounted to a can. The data plotted in the top block was taken at North Carolina State University's anechoic chamber, and the data in the bottom plot row was taken at JPL's far field antenna range (measurements contaminated by environmental factors). In all cases, the main beam appears to grow asymmetric at higher frequencies, and develop especially high cross-pol.

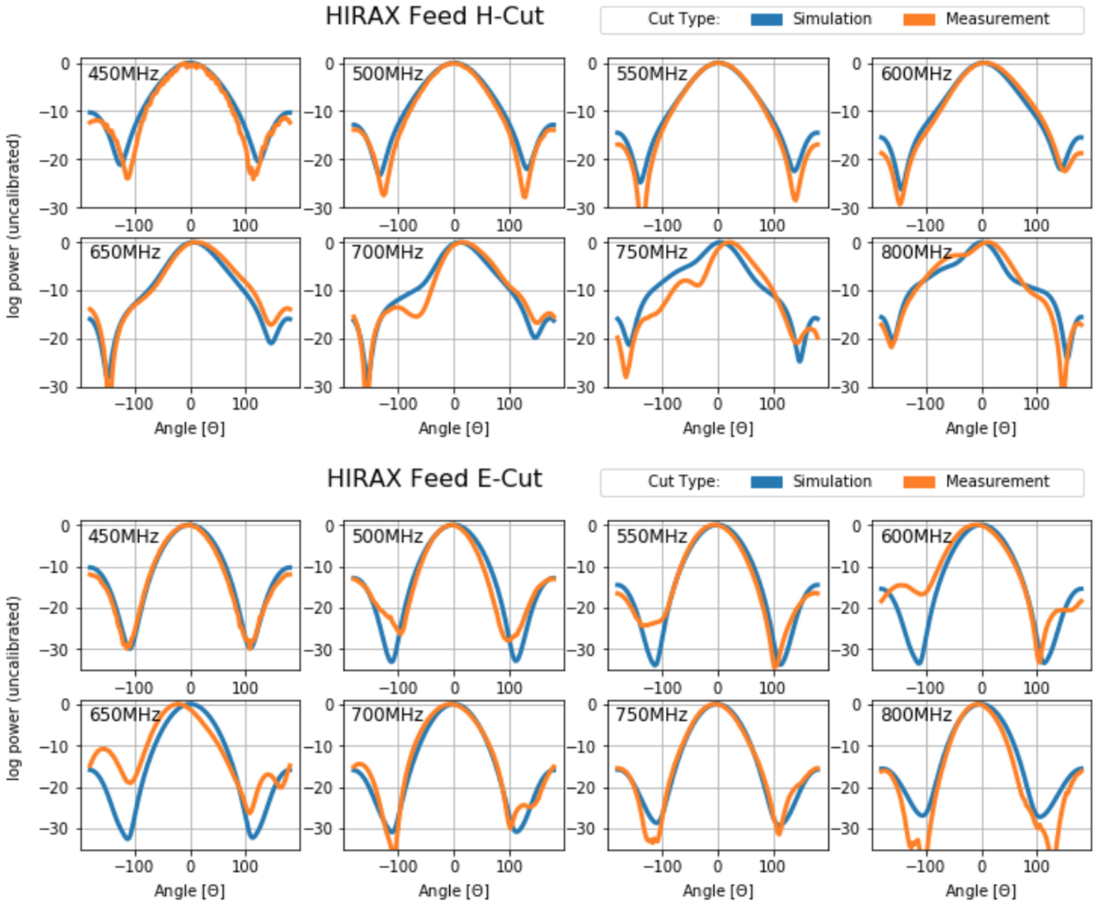

Figure 3.11: Cuts in two planes for HIRAX antenna beam, including measurement and simulation results.

Understanding asymmetric features

Measurements were additionally taken to verify the high-frequency beam asymmetries, and untangle environmental effects from real beam features. To do this, one can repeat feed beam cuts with the feed rotated by 180 degrees. If the asymmetries are due to reflections from the environment, this adjustment to the set up should not impact the results. If, however, the asymmetries are true features, they will correspondingly rotate by 180 degrees in this new measurement. We notice asymmetric features in HIRAX feed measurements, at the level of 5dB, which are shown in figure

3.12. Applying the described methodology to the HIRAX feed, measuring the beam shape in 180 degree antenna rotations and reflecting the measurements back by 180 degrees in analysis, we find alignment to the sub-dB level (right panel of fig. 3.12). This result validates that the beam is truly asymmetric and measurements are not hindered by the environment.

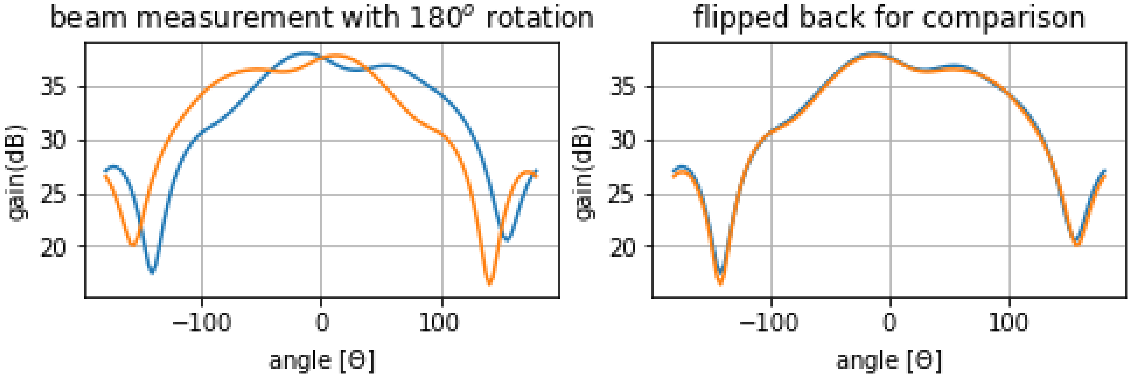

Figure 3.12: H-Plane beam measurements of the HIRAX antenna, repeated for the antenna flipped upside-down (rotated 180^o). Measurements are taken to verify that beam asymmetries are true features of the antenna and not artifacts of the antenna range. To do this, we complete a beam cut, then repeat it with the feed rotated by 180 degrees. The two cuts are plotted as the blue and orange curves (left plot). If the orange curve is mathematically flipped about 0^o, we find an identical signal to the blue curve (right plot), which means that the asymmetries are entirely due to the feed performance and not an artifact of reflections from the environment.

Measurements with a varied backing structure

The HIRAX feed design utilizes a ring-choke to circularize the beam and improve its directionality, thereby preventing spill over. The can shape and dimensions were optimized in the software CST Microwave Studio, where it was found that introducing the can narrows the beam, and smooths its shape over frequency. A smaller can is expected to reduce S_{11} (defined in section 3.1), but improves aperture efficiency by not blocking the beam (scales with r^2). The can used for HIRAX is 330mm diameter.

We can verify the can's impact and perform important checks on simulations by performing beam measurements of the HIRAX feed backed by both a can and plate (important as the HIRAX feed is designed to be mounted on some ground plate strucure), and comparing the two. Figure 3.13 shows that the introduction of the HIRAX can over a plate of comparable dimensions can shrink the antenna backlobe by as much as 6dB and improve its directionality.

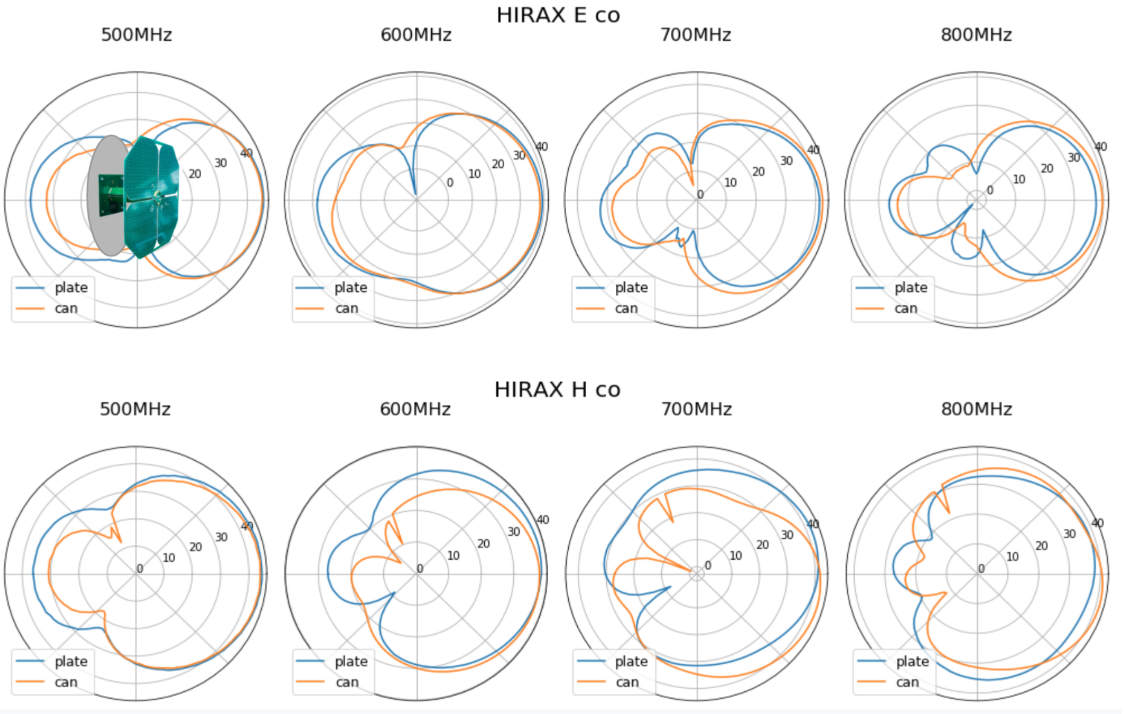

Figure 3.13: Cuts in two planes for the HIRAX antenna backed by a plate and a can. A direct comparison between the two measurements demonstrates the utility of the can, which is meant to shrink the backlobe, narrow the beam, and circularise the beam. The HIRAX feed and plate are drawn into the 500MHz E-co measurement to illustrate the antenna orientation for the reader.

Finally, beam measurements were taken for four different HIRAX feeds, which showed excellent repeatability between feeds. This is critical to realizing the HIRAX design, which relies on highly redundant dish instrumentation.

3.4 CHIME feed measurements

Because this is a thesis largely focused on HIRAX, the following CHIME feed results will be largely discussed in the context of their similarities to and differences from the HIRAX feed. Their similar designs makes for a highly instructive basis of comparison.

The CHIME feed is based on the Molonglo Telescope feeds[93], and was the design inspiration for the HIRAX antenna. It is similarly broadband and optimized for performance in 400-800MHz frequency band. In contrast to HIRAX, it is made from teflon, a lower loss (but higher cost) material than FR4, which allows the first stage amplifier to be placed after the antenna structure while maintaining an acceptable noise temperature. Polarization is summed along an axis rotated 45 degrees with respect to the HIRAX polarization, after the signal makes its way down tapered striplines on the antenna balun, which impedance matches freespace to 50Ω (figure 3.14).

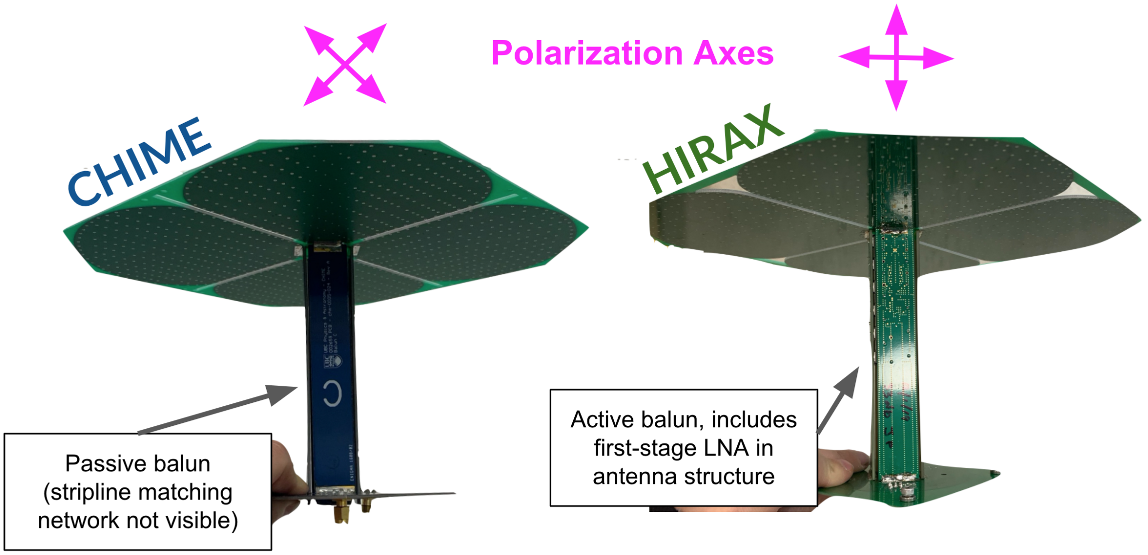
Polarization Axes
CHIME
HIRAX
Passive balun (stripline matching network not visible)
Active balun, includes first-stage LNA in antenna structure

Figure 3.14: The CHIME (left) and HIRAX (right) antennas are very similar in design. Apart from different material compositions, there are two main differences between the feeds that impact the beam pattern. (1) CHIME is completely passive while HIRAX integrates a first-stage LNA into the antenna structure. This is visible on the balun stems in the photographs. (2) The polarization axes for CHIME are in line with the petal gaps, while the polarization axes for HIRAX run between diagonal petals.

The full cuts of the CHIME feed (figure 3.15) show expected behavior from simulations and prior beam characterization: a broad, smooth, gaussian beam that is generally symmetric, with low-level cross-polarization.

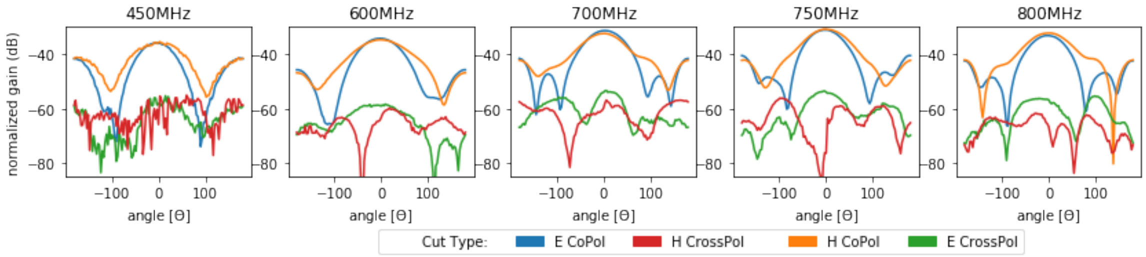

Figure 3.15: Full E and H-plane cuts for the CHIME antenna, including both co and cross-pol measurmements.

In measuring full cuts of the CHIME feed, there are revealing contrasts with the HIRAX feed performance (a comparison of co-pol cuts for HIRAX and CHIME is shown in figure 3.16). For one, while asymmetries do arise at high frequencies, they are not nearly as stark as in HIRAX, and the centroid drift(not shown) is much less substantial (fewer than 5 degrees for CHIME, while up to 25 degrees for HIRAX). Just as critical is the level of cross pol, which remains over 20dB down from the co-pol for the majority of the main beam.

Observing substantially more isolation between co and cross pol on CHIME feeds over HIRAX suggests that the modifications between the two designs change more than initially expected. This comparison is also encouraging for HIRAX, as it suggests that with minimal design tweaks it can accomplish an improved antenna performance.

3.5 Biconical antenna measurements

A biconical antenna is comprised of two conical shapes which are fed at their meeting point. It is similar to a dipole in design, with the increased surface area giving an improved impedance match over a broad frequency range, thus making it a good

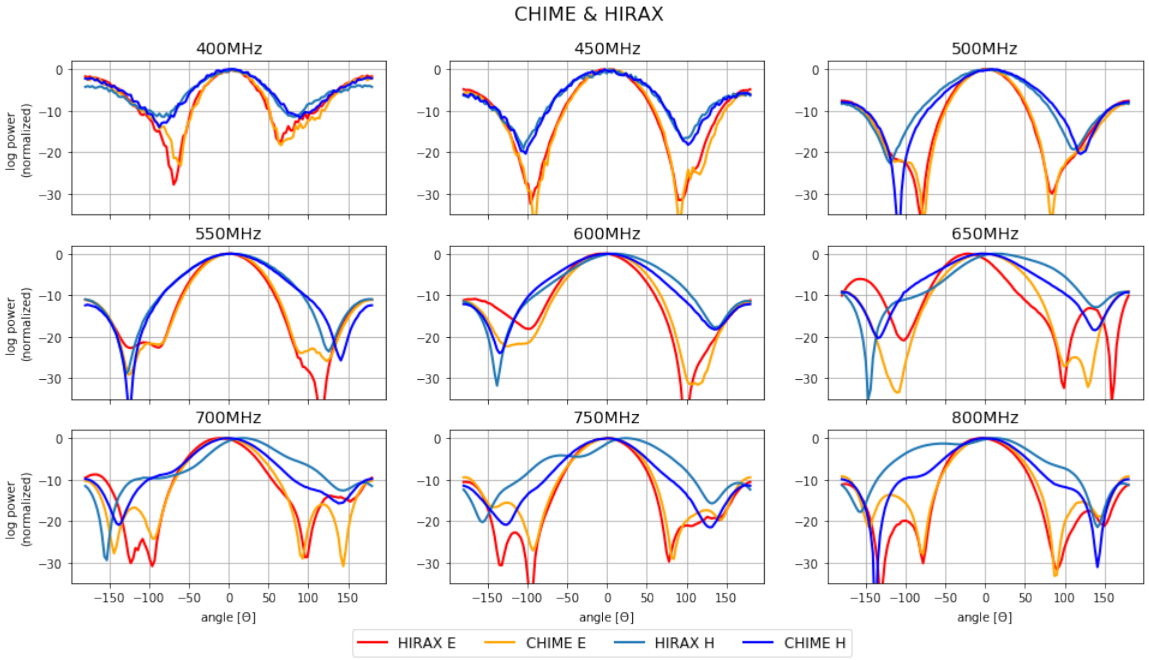

Figure 3.16: A comparison between the CHIME and HIRAX antennas, which share a similar design. Two cuts are shown for each antenna: the E-plane cuts are in red/orange shades and the H-plane cuts are in blue shades for HIRAX and CHIME. HIRAX is notably less symmetric, particularly at higher frequencies. This comparison is done with a plate because the CHIME feed has extra port for noise injection that prohibited attachment to the can on hand during this measurement set.

choice for antenna applications that require a broad beam and a broad band. As discussed in Chapter 5, the drone developed at Yale for HIRAX beam calibration utilizes a biconical antenna as its transmitter, and thus beam characterization of biconical antennas for drone characterization has been critical. The biconical antennas will be used for L-band (1-2GHz band) applications (in addition to measurements in the HIRAX band), and thus requires characterization across a 400MHz-1.5GHz frequency range.

The Aaronia biconical antennas (from the BicoLOG[3] series, referred to as 'bicologs' during this dissertation) are composed of two hexagonal cones arranged from a series of metallic panels each (fig. 3.17). They are lightweight antennas, weighing 362 grams, and compact, measuring only 34cm in length, despite functionality down to 20MHz. Measurements of the three different Aaronia biconical antennas (one model 20300, operating between 20-300MHz and two 30100 models, operating between 30-100MHz) were performed. Each antenna was measured as a standalone, and then measured again and with different backing structures. All measurements of the biconical antennas were taken in the NCSU anechoic chamber, with results described below.

---

[3] https://aaronia.com/antennas/bicolog-series-biconical/

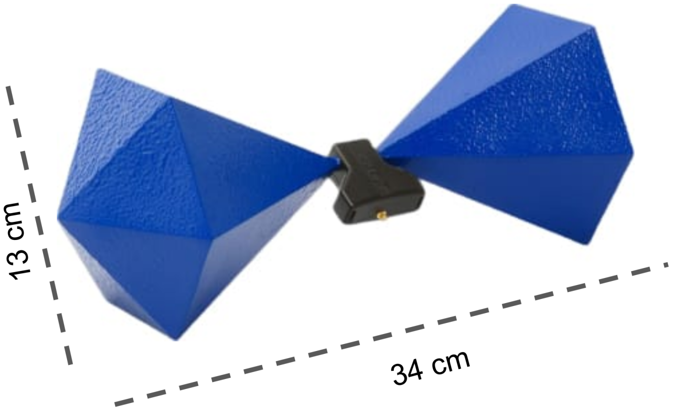

Figure 3.17: Schematic of an Aaronia biconical antenna (image from Aaronia website).

Bicolog beam measurements

The biconical antenna beam pattern was in line with expected profile from the datasheet. It exhibits classic dipole beam features: there is a dual-lobed pattern in the E plane, with nulls where the bicolog panels come to a point, and the beam is omnidirectional in the H plane. Measurements are smooth to within ∼2dB (fig. 3.18), and the antenna is particularly high performing in the 400-800MHz HIRAX band.

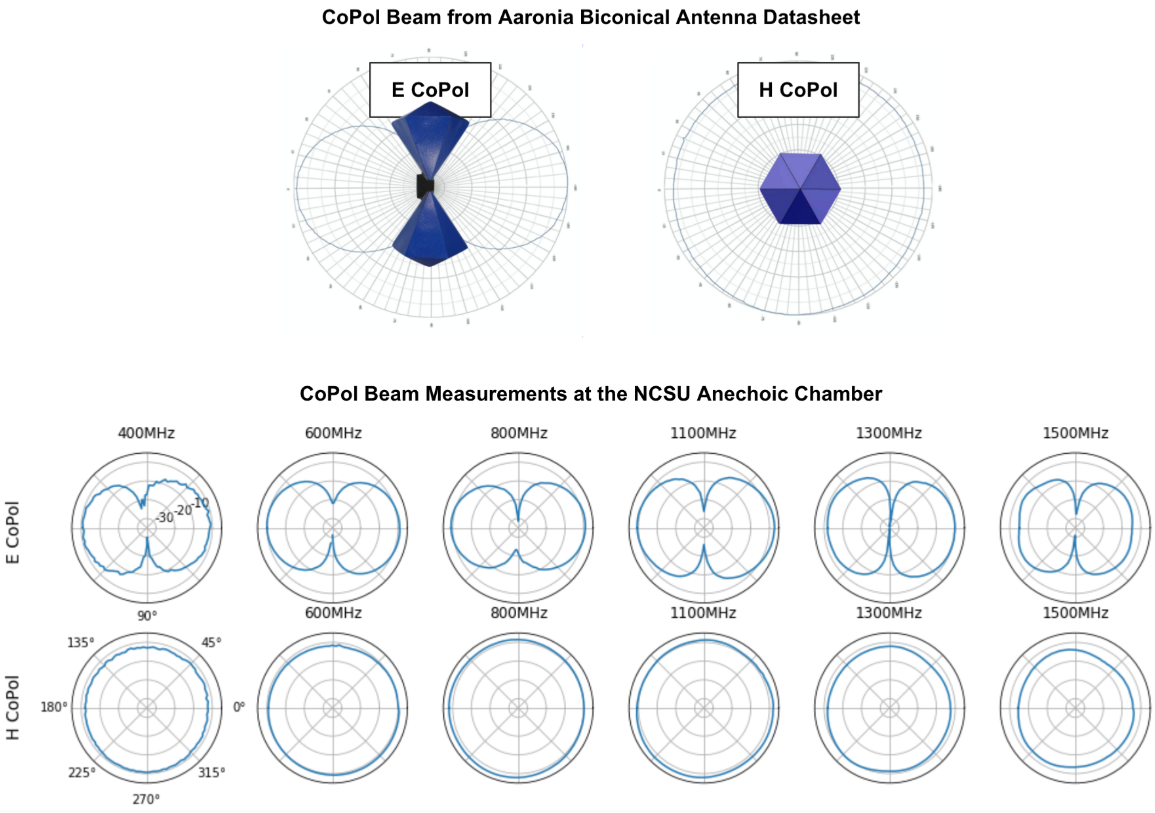

Figure 3.18: E and H-plane co-polar cuts for the Aaronia 20300 biconical antenna, which is to be mounted to a drone for beam calibration. Both specs from the antenna datasheet and measurement results from the NCSU anechoic chamber are shown. Data sheet plots are overlaid with antenna images to illustrate antenna orientation for the different cuts.

Measurements with varied backing structures

The biconical antenna's smooth beam does indeed change shape if a reflective object (for instance, a carbon fiber drone), is placed nearby. Because we plan to fly the biconical antenna on a drone, it is thus necessary to explore and characterize the bicolog antennas with backing fixtures and mounting styles that can shield from structures in the backlobe. The existence of a plate or can does indeed shrink the back lobe, which will be important for decoupling the feed from drone. These measurements are shown in figure 3.19, in a gain plot that reveals how plate and can structures will form a more directional beam, and in a series of polar plots showing full 360 degree 1D cuts. At higher frequencies, a can is particularly effective at limiting the backlobe, which we keep in mind for future drone mounting designs. The next section will delve into the effectiveness of backing the antenna with a ground plate as part of the full drone system.

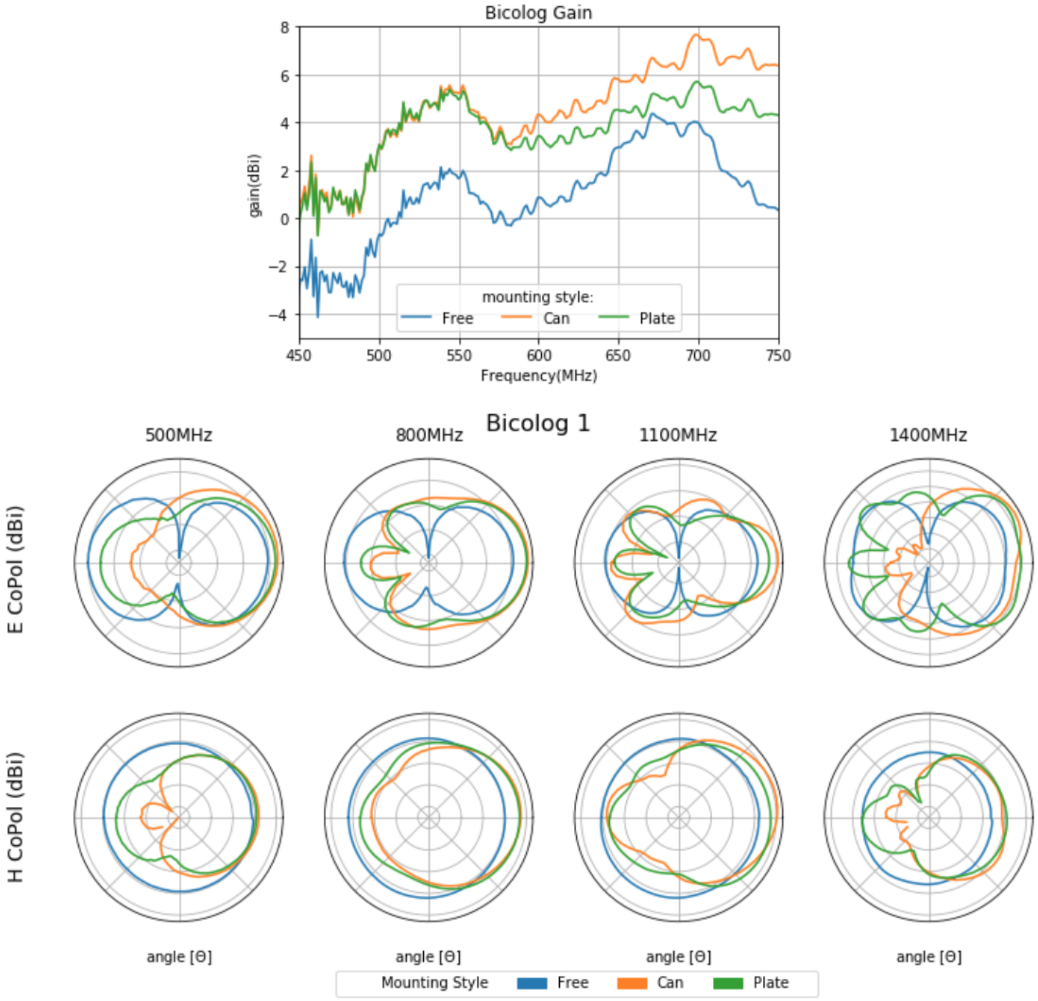

Figure 3.19: The biconical antenna beam with different backing structures, which confirm that the presence of a ground plate or can is helpful for forward-directing the beam. Here, 'free' refers to the biconical antenna freestanding with no backing; 'can' refers to the HIRAX can, which is separated from the antenna by a two inch nylon post; and 'plate' refers to a 13in plate also spaced 2in from the antenna. The top figure is antenna gain as a function of frequency, and the bottom figure includes E and H-plane cute for the different backing scenarios. The presence of a backing structure can decouple the antenna beam from metallic structures in its backlobe, which will be important when using the antenna for drone calibration, which requires precise knowledge of the transmitter beam pattern.

Spacing the antenna relative to a ground plane

A companion question to what to shield the bicolog with is how to space the bicolog with respect to shielding. Thus, it is also important to examine the effect of mounting post length on the beam. The mounting post is made from RF-transparent nylon material, so the inclusion of a post should only impact the beam insofar as spacing it from the backing structure, and should not itself induce reflections. In changing the post, we expect that a shorter post will give a wider beam, as the outgoing wave from the antenna is less planar as it interacts with the reflective surface, and thus will reflect at wider angles. Indeed, when comparing post lengths ranging from 2-3in, which are the sides practical for mounting in the drone setup, you will find as expected that a shorter post widens the beam, though only by a few percent. This is helpful information for drone applications, because the bicolog is trying to approximate a point source, as described in the later chapters. The bicolog beam appears to approximate that of a point source for all available spacings, as designed.

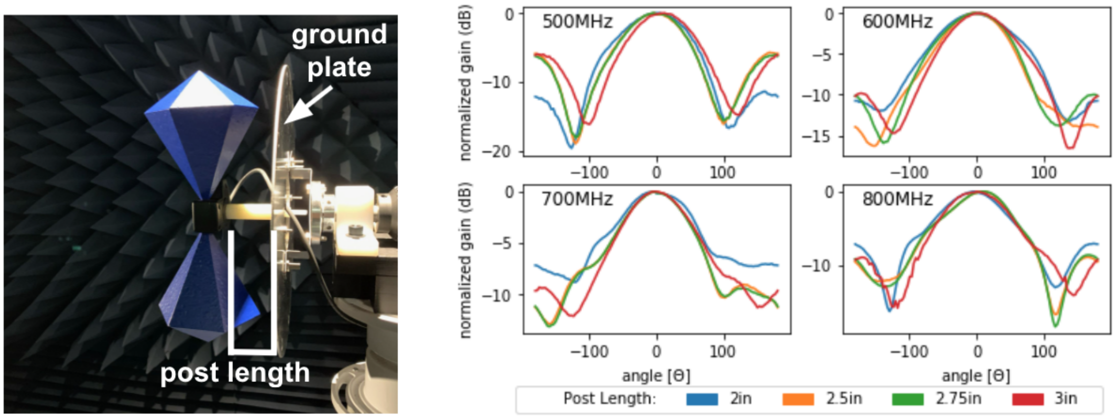

Figure 3.20: Left: Diagram of biconical antenna and ground plate, along with the post length, which is altered for this investigation. Right: Plots showing the impact of varying mounting post length. The goal is ultimately to use this antenna as a transmitter for beam measurements of telescopes, and for this particular application, the beam width variations caused by altering post length is not of great consequence.

3.6 Drone transmitter measurements

As mentioned, and detailed in later chapters, the biconical antenna measurements are critical to this thesis for the part they play in drone calibration of telescopes. Thus, their measurements are most relevant for HIRAX applications in the context of being attached to a drone. In the section above, I describe the biconical antenna design and report measurements attached to different forms of mounting structures. In this section, I will explain how the full antenna-and-drone system performs, as measured on the JPL and GBO outdoor antenna ranges and approximated in the NCSU anechoic chamber.

The drone used for calibration, described in Chapters 5 and 6, has long, thin carbon fiber legs, which are known to act as reflectors. As a result, the smooth beam of the biconical antenna is substantially modified when mounted on the drone. This behavior was first measured at the JPL range in October 2019, and corrected preliminarily by inserting a 13in. metallic mesh plate behind the antenna. The comparison between the beam patterns with and without the plate ('ground plane') is shown in figure 3.21(a). Here, the orange curve shows the drone transmitter beam pattern with a ground plane between the transmitter and drone, and the blue curve shows the drone transmitter beam pattern without a ground plane present. The presence of a ground plane smooths the beam, giving it an expected shape that peaks at boresight and falls off as an approximate Gaussian. Without a ground plane, the beam has several main lobes and does not peak with the antenna at boresight, which is a hindrance for our measurement scheme.

The beam of a bicolog backed by a plate in the presence of a drone was similarly explored in the NCSU anechoic chamber. The bicolog was mounted on a mesh plate with a two-inch mounting post and measured as described in the prior section. This measurement was then repeated with a mock 'drone', with 6 legs fashioned from 24inch copper strips and affixed behind the mount (the drone was approximated

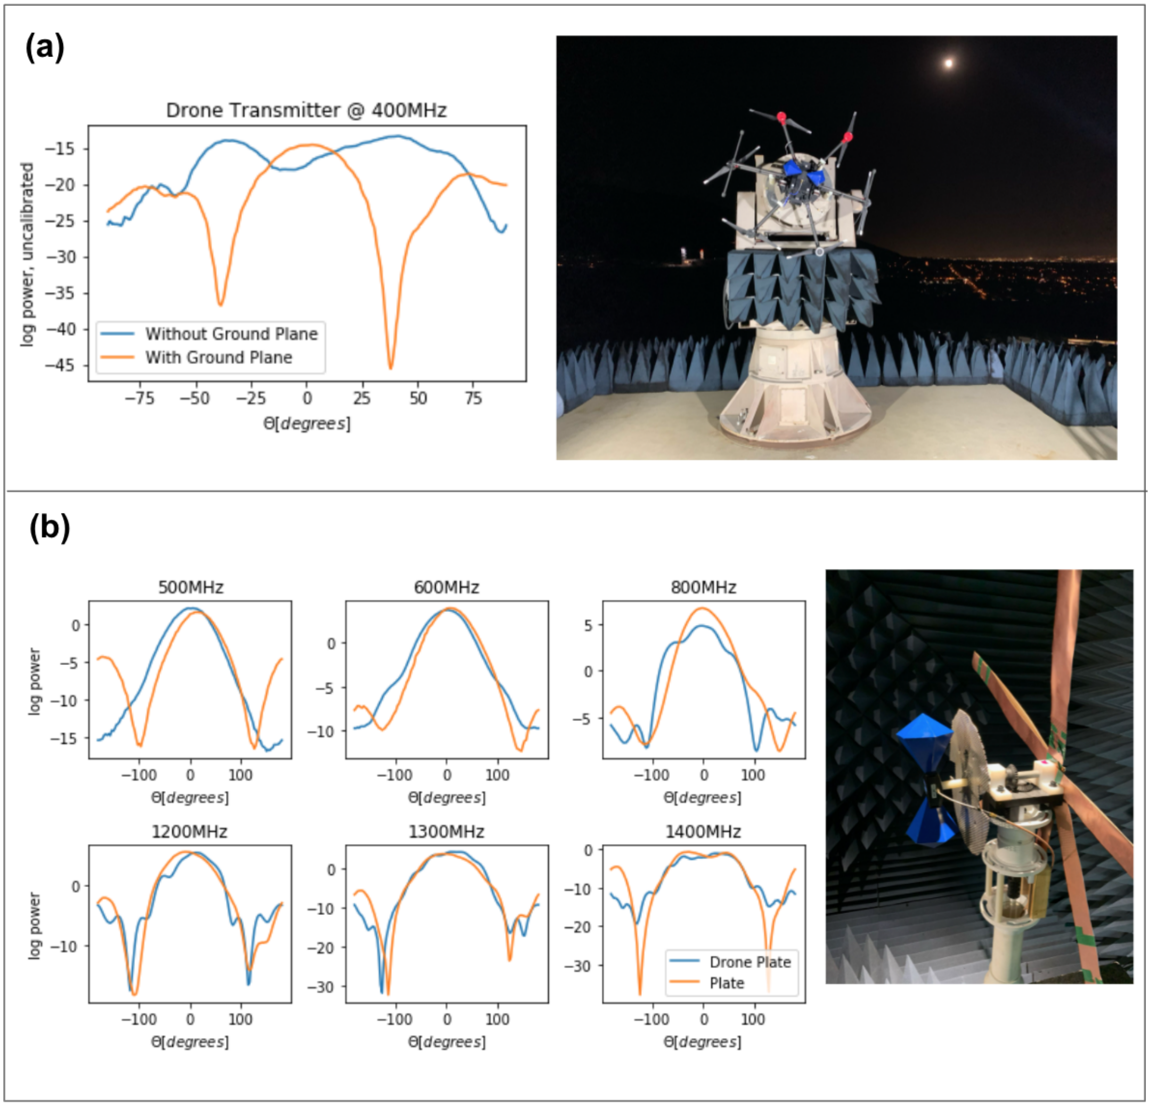

Figure 3.21: (a) Antenna range beam measurements of the biconical antenna mounted to a hexacopter drone, with and without a groundplate. The presence of a ground plate considerably smooths the beam by limiting reflections off of the drone carbon fiber legs and body. These measurements were performed at the Mesa range at NASA JPL. (b) Anechoic chamber measurements of the biconical antenna and a mock-drone separated by two different backing structures. These measurements were performed at the NCSU anechoic chamber. The comparison between a mesh backing structure (which includes cutouts for drone landing legs, and has been utilized for drone calibration flights) and solid plate indicate that the solid plate is better for smoothing the beam. Ground plate performance will be revisited in future iterations of transmitter design.

because the mount at NCSU could not support its weight) and results are shown in figure 3.21(b). In this figure, the drone mock measurements have some additional features, particularly at high frequencies outside of the HIRAX band, but still retain a broad shape. This is further evidence that the plate substantially shields the bicolog from the drone at a wide variety of frequencies.

A more detailed and reliable set of drone transmitter beam measurements was taken at the Green Bank Observatory outdoor antenna range, where we were able to use the actual drone with the bicolog mounted for measurements. These measurements validated observations that our drone transmitter set up has a smooth, broad beam. The beam is particularly broad for frequencies above 500MHz, when the beam full-width-half-max (FWHM) is greater than 60 degrees for both E- and H-plane cuts. The full set of cuts at all measured frequencies is shown in fig 3.22, which also illustrates the way in which the beam broadens with frequency, particularly at the lower end of the band.

The requirement for a drone calibration transmitter is that it be broadband, lightweight, and physically mountable on a drone. A passive version of the HIRAX feed fits this criteria. Thus, this feed has been measured while mounted on the drone and backed by a plate, in order to assess it as a viable candidate for late upgrades. It was only measured at three frequencies (400/600/800MHz) due to time constraints, but appeared broad and smooth, as shown in figure 3.23.

3.7 Chapter summary

Over the course of several years and several facilities, I have measured a number of antennas useful to the HIRAX experiment.

In characterizing the HIRAX antenna, I found that it exhibits the expected gain and has a beam pattern in agreement with simulations. I also found the design to be repeatable, through comparisons of repeated beam and gain measurements across four different antennas. I additionally discovered an asymmetry in the co-polar beam,

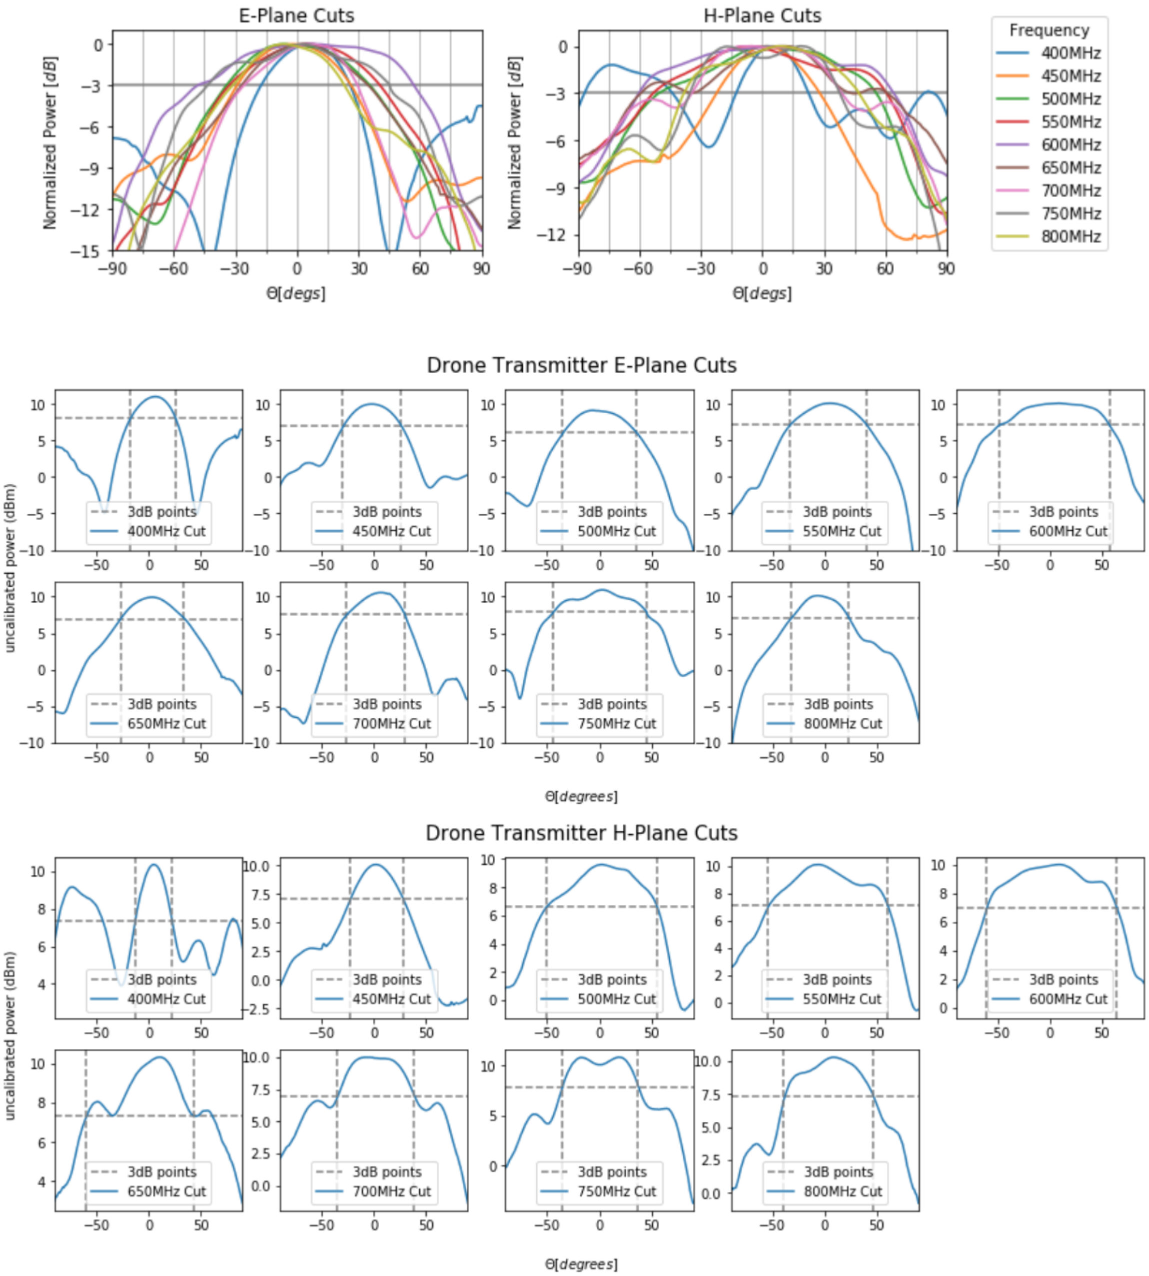

Figure 3.22: E and H-plane beam measurements for the drone transmitter system, performed at Green Bank Radio Observatory. The top plots show a comparison of cuts for all frequencies, showing a broadening of the beam as frequency increases. The other plots show the same E and H-plane cuts, but separated by frequency, to make individual beam features more visible. Beam measurements of this sort are a critical step to removing the drone transmitter beam from telescope beam maps (see Ch. 5 for further discussion).

and notably high levels of cross-pol, which should be corrected in the future antenna versions. As a result, this design is currently under assessment by the collaboration.

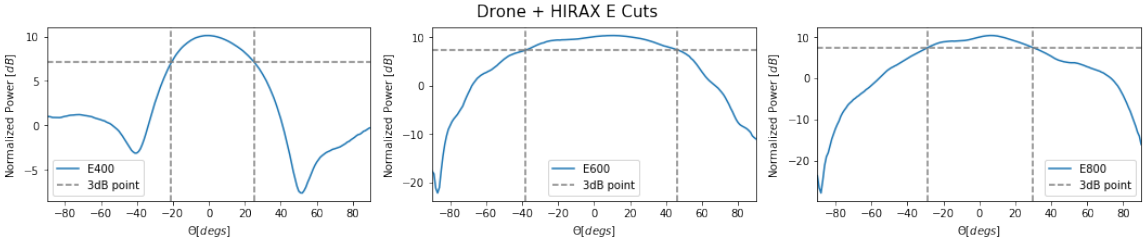

Figure 3.23: Beam cuts in 1D of a passive HIRAX feed mounted to the drone (along with a ground plate). The beam is broad and smooth at the measured frequencies, indicating that this antenna is a promising model for a drone calibration transmitter.

To complement HIRAX characterizations and provide a design comparison, I measured one of the CHIME antennas. For CHIME, I found a more symmetrized beam and lower cross pol, which motivates further study of the design. This comparison suggests that with minimal design tweaks HIRAX can accomplish an improved antenna performance.

Finally, I investigated a biconical antenna, to be used as a transmitter in a drone calibration system. The bicolog has, as expected, a broad-band and omni-directional beam when measured on its own. The beam pattern is, however, easily influenced by the presence of additional structures, such as a drone carbon fiber body. Thus, for our purposes the bicolog should be backed with a structure that decouples its beam features from the carbon fiber drone, in order to preserve its broad, smooth beam features.

This work is continued in later chapters, including Chapter 4 and Chapter 5, which discuss the HIRAX and bicolog antennas at length.

Chapter 4

Design and implementation of a noise temperature measurement system

The focus of this chapter is on assessing the HIRAX feed contribution to overall system noise ($T_{\rm sys}$) (the early sections were published as a SPIE proceedings in 2021 [94]). To measure cosmological emission efficiently we require the total system noise to be kept below 50K, of which at most 30K can come from the feed itself. The feed design is optimized to reduce system noise by removing particular sources of loss in the analog chain, primarily by moving the first stage amplification into the antenna balun. The choice to embed amplification directly into the balun stem makes noise temperature characterization challenging, as the noise temperature, gain, and stability of the amplifier cannot be measured directly with a noise figure meter or related laboratory equipment. To measure the noise temperature of the feed and amplifier, we must inject a broad-band signal into the HIRAX antenna. We have built and implemented a test-bed built at Yale University that allows us to measure signals from radio absorber at two known temperatures, and infer the noise temperature of the antenna and amplifier via the Y-factor method (see e.g. Pozar's *Microwave Engineering* [91]).

4.1 The Y-factor method

The Y-factor method allows one to determine the noise temperature of an antenna by comparing its output power at two different temperatures. This method can be derived from the observation that an antenna in a cavity in thermal equilibrium at temperature T behaves like a resistor of temperature T, where the antenna internal noise can be likened to that of Johnson noise in the resistor[95]. The Rayleigh-Jeans approximation of the black-body law at microwave frequencies can be used to approximate power[91]:

$$P = GkT\Delta\nu \tag{4.1}$$

for a black-body of temperature T, where G is the measurement device gain, k is Boltzmann's constant, and $\Delta\nu$ is the bandwidth in Hz. When two identical loads of different temperatures are available (labeled here as "hot" and "cold"), measurements can be performed at both temperatures,

$$P_{\text{hot}} = Gk(T_{\text{hot}} + T_{\text{noise}})\Delta\nu \tag{4.2}$$

$$P_{\text{cold}} = Gk(T_{\text{cold}} + T_{\text{noise}})\Delta\nu \tag{4.3}$$

resulting in a system of equations that can be solved for the noise temperature T_{noise}. This method is most commonly expressed through the following set of equations:

$$Y = \frac{P_{\text{hot}}}{P_{\text{cold}}}. \tag{4.4}$$

$$T_{\text{noise}} = \frac{T_{\text{hot}} - YT_{\text{cold}}}{Y - 1}. \tag{4.5}$$

with Y referred to as the "Y Factor".

This type of measurement is commonly taken using the sky as a cold load, but this is impractical for our experiment due to the size of the beam, and for us locally due to the RFI rich environments around universities. We instead designed and constructed an experimental system with hot/cold loads of 300K/77K (corresponding to room temperature and liquid nitrogen temperature) to be used in a laboratory setting. We aim to understand the antenna noise temperature to within 5K, or 10% of the expected system noise.

4.2 Experiment design

The noise temperature measurement system has been designed as a pair of reflective closed cylindrical cavities with radio-frequency (RF) absorber in the bottom. One cavity is kept at ambient temperatures such that the absorber emits at $T_{\text{hot}} = 300\,\text{K}$, and is designated the 'warm' load. The second cavity is filled with liquid nitrogen such that the absorber emits at $T_{\text{cold}} = 77\,\text{K}$, and is designated the 'cold' load. The feed is attached to the lid of each cavity such that its beam terminates in the absorber. There were several constraints on the design, including:

- Cost: We use commercially available materials with minimal labor to assemble.
- Size: We required that the cavities are smaller than $\sim$1.5 m, to stay within a reasonable footprint in the lab space, and ensure standard material stock is suitable for construction.
- Cavity material: We use steel because it is easily weldable (important for containing liquid nitrogen as a safeguard) and inexpensive relative to other mate-

rials.

- Absorber: We use commercially available AEP-18 series pyramidal foam[1] for the RF absorber, which has ∼30 dB absorption in our band.

- Shape: We optimize the cavity dimensions to minimize resonances, and to be structurally stable and capable of containing 600 L of liquid nitrogen.

- Insulation: We line all surfaces of the chamber with insulation to reduce liquid nitrogen boil-off and limit the accumulation of water vapor. We also add a foam lid to encourage a nitrogen vapor layer and isolate the feed from the cold vapor.

- Reflectivity: We coat the inside of the cavities with aluminum tape to increase the reflectivity, primarily to improve the hold-time of the liquid nitrogen in the cavity.

- Indistinguishable: The two cavities must be sufficiently similar, or their differences sufficiently repeatable and characterized, to keep systematic errors less than 5 K. Cavity differences will be quantified in more detail in later sections.

As described in this chapter, the design of the test-bed was optimized using the simulation software CST, and materials were chosen to allow one of the cavities to hold liquid nitrogen.

4.2.1 Simulations

We optimized the design of the cavities using CST Microwave Studio[2]. CST was a natural choice for modeling, as the HIRAX and CHIME collaborations had previously used it to construct feed models, which were leveraged for this work. We began by simulating the HIRAX feed in free space and monitoring the S_{11} (as described in Ch. 3). Because the cavity should be an absorptive blackbody, the primary figure of merit was to match the cavity reflections, parameterized by S_{11}, to those of the feed in free space (such that dominant reflections are internal to the feed itself). The simulations were performed and optimized with a passive HIRAX feed attached to the lid of the cavity, with the lid functioning as a ground plane. HIRAX plans to employ a circular choke to circularize the beam, which was not included in these simulations.

---

[1] https://www.mvg-world.com/en/products/absorbers/standard-absorbers/pyramidal-absorbers-aep-series

[2] https://www.3ds.com/products-services/simulia/products/cst-studio-suite/solvers/

The company from which we obtained RF absorber did not supply relevant material parameters, such as dielectric constant (ϵ) and loss tangent ($\tan\delta$), so we created a user-defined absorptive material through the CST optimization function with guidance from literature[96]. This optimization was performed by placing a slab of RF absorber in front of the feed in software, and sweeping material parameters until S_{11} was minimized. We ultimately simulated the RF absorber as 18 in tall pyramidal cones (numbering 16 to a 2 ft $\times$ 2 ft block, to match dimensions of those available commercially) with density 159 kg/m^3, ϵ=2.7 and $\tan\delta$=1.

Once the RF absorber material was modeled, we set it within several steel cavities of various shapes to determine which to pursue for the design. This determination was made by comparing the feed S_{11} profile of different cavity options with the free space profile. We ultimately settled on a cylinder, as it matched free space as well as the other options, is known to be the most mechanically robust shape, and is simplest to construct (Figure 4.1). We then optimized cavity dimensions, finding a diameter of 129.5 cm and a height of 70.8 cm.

We similarly estimated the reflection off of the liquid nitrogen surface by specifying $\epsilon = 1.44$ and $\tan\delta = 5 \times 10^{-5}$, which are measured parameters for 18-26 GHz[97, 98] (parameters were not available in our band). We determined reflections to be below -14 dB. As described in more detail in Section 4.4, the nitrogen is sufficiently transparent to continue with the design, but may have to be estimated and accounted

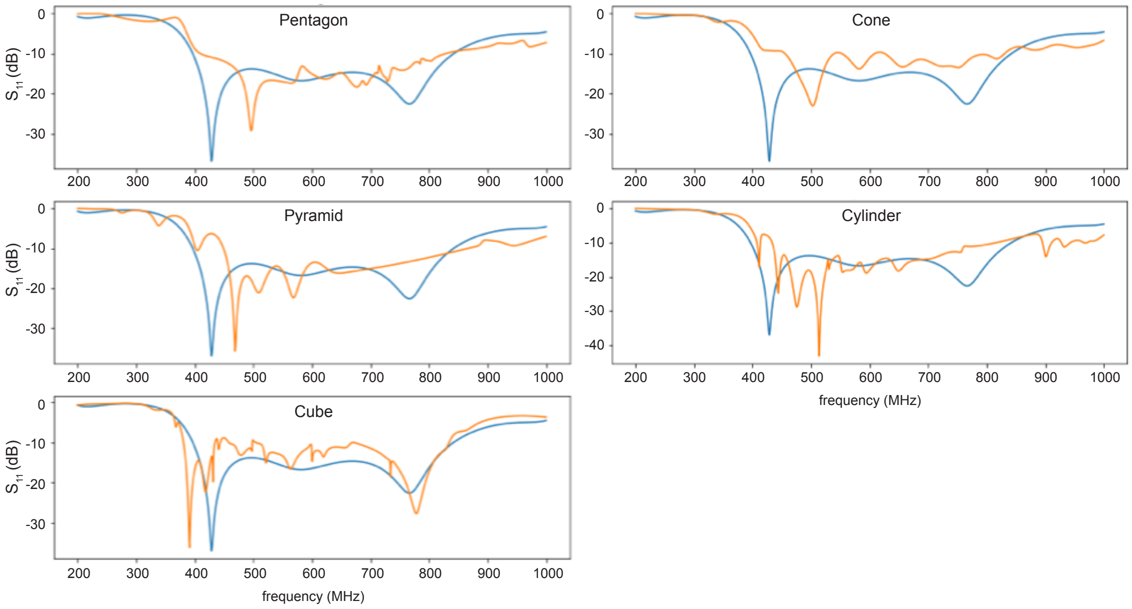

Figure 4.1: Preliminary S_{11} simulation results for the CHIME (unamplified HIRAX) feed in various cavity shapes before optimization (orange), plotted against simulation results for free space (blue). All cavities are the same height, use the same RF absorber material parameters, are made from aluminum, and share similar base dimensions (i.e. the diameter of the cylinder and cone bases and the length of the cube sides are equal in length). From these early simulations, we find little difference in RF performance between leading contenders, and settled on a cylinder to optimize for the final design. Aside from RF performance, cylinders are mechanically robust, and simple to construct.

for during later analysis.

Finally, we investigated tolerances on cavity dimensions by sweeping through a variety of length parameters centered about the optimal dimensions and solving for S_{11}. The resulting deviations in S_{11} would represent possible differences between the two cavities, which we required to be less than 1 dB. The results indicated that the cavity dimensions must be constructed with tolerance of $\sim$1 cm. We followed a similar procedure to set a tolerance for the insulation thickness, accomplished by assuming total RF transparency of insulation, and leaving an air gap of corresponding volume in simulation. This procedure yielded an insulating layer of thickness up to 80 mm.

The simulated S_{11} for the final, optimized cavity design is shown in Figure 4.2. The resulting S_{11} is similar to the free-space reflection, and below the -10 dB reflection requirement for HIRAX. Also shown is the simulation result for reflection including the liquid nitrogen surface, which modestly changes the S_{11} but remains well below -10 dB within the 400-800 MHz band.

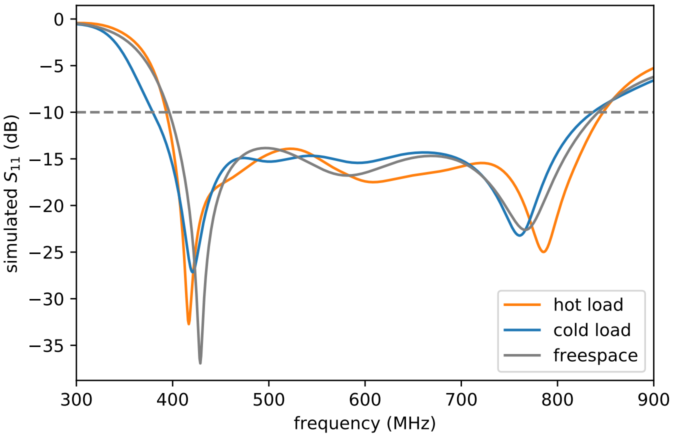

Figure 4.2: S_{11} simulation results for the CHIME (unamplified HIRAX) feed. We compare free space results to results from the Y-factor measurement system hot and cold loads (the cold load contains a simulated liquid nitrogen surface), finding all profiles to be similar to well below -10dB. Their differences are computed to give a sub-1K error in noise temperature from 400-750 MHz.

4.2.2 Cryogenic system construction

From the simulation work described above, we determined that cylinders, 129.5 ±1 cm in diameter and 70.8 ±1 cm in height, were the optimal cavity shape for the noise temperature test-bed loads. The cylinders arrived at Yale in December 2018[3]. We initially measured the RFI environment inside the enclosures, as well as the S_{11} of each box relative to the other. We found that most (but not all) RFI is blocked, and that the boxes are remarkably consistent–to within a small fraction of a dB. We also investigated metallic coatings for minimizing radiative heat loss in the liquid nitrogen system, ultimately settling on aluminum tape.

Once the boxes were shown to be consistent to within a dB, insulation construction began. Initially, we planned to heat bond two layers of HD30 zotefoam to form a fully welded interior insert. It was quickly apparent, however, that a heat bond would not hold as the zotefoam bent to the cylinder radius. We tried several approaches, including building a wooden frame of the necessary radius and forming the zotefoam around the outside with ratcheting tie straps. We tried instead tensioning the foam into place within the box and sealing the seams with RTV, but this system failed initial nitrogen testing.

---

[3] The cylinders were fabricated by the company Welding Work located in Madison, Connecticut.

With multiple zotefoam tests repeatedly failing at the box seams, we elected to try a monolithic insulating layer. After considerable trial and error, we determined that two-component polyurethane spray foam insulation was a viable option if used with an additional liquid-proofing layer. For this layer, we elected to use fiberglass cloth soaked with resin, which provided both liquid proofing and additional structural stability (there is precedent for using this material in cryostats).

The final nitrogen-tight insulating insert is composed of a fiberglass inner layer bonded to 1/4in foamular foam, followed by a 2in layer of cryogenically-rated spray foam, all encased in an outer layer of the same 1/4 in foamular. The insert was constructed in multiple phases, shown in figure 4.3.

The inner-most insulation layer was constructed first. It is comprised of a fiberglass tub of height 24in and diameter 46in, and the construction procedure is as follows. An outline of the tub was made from 1/4in foamular foam cut to size and connected together with packing tape to form a cylinder. This was then filled with an inflated round swimming pool connected to a rigid foam piece (for stability), and set on a 50ft 3mil polyethylene sheet. We then painted on a coat of epoxy using standard paint brushes, over which we laid fiberglass cloth and then evenly distributed the remaining resin with a rubber scraper. Once the fiberglass cloth was fully saturated and smoothed we covered the entire surface with a coat of peelply, a material that allows resin to seep through while retaining the ability to peel away from the surface

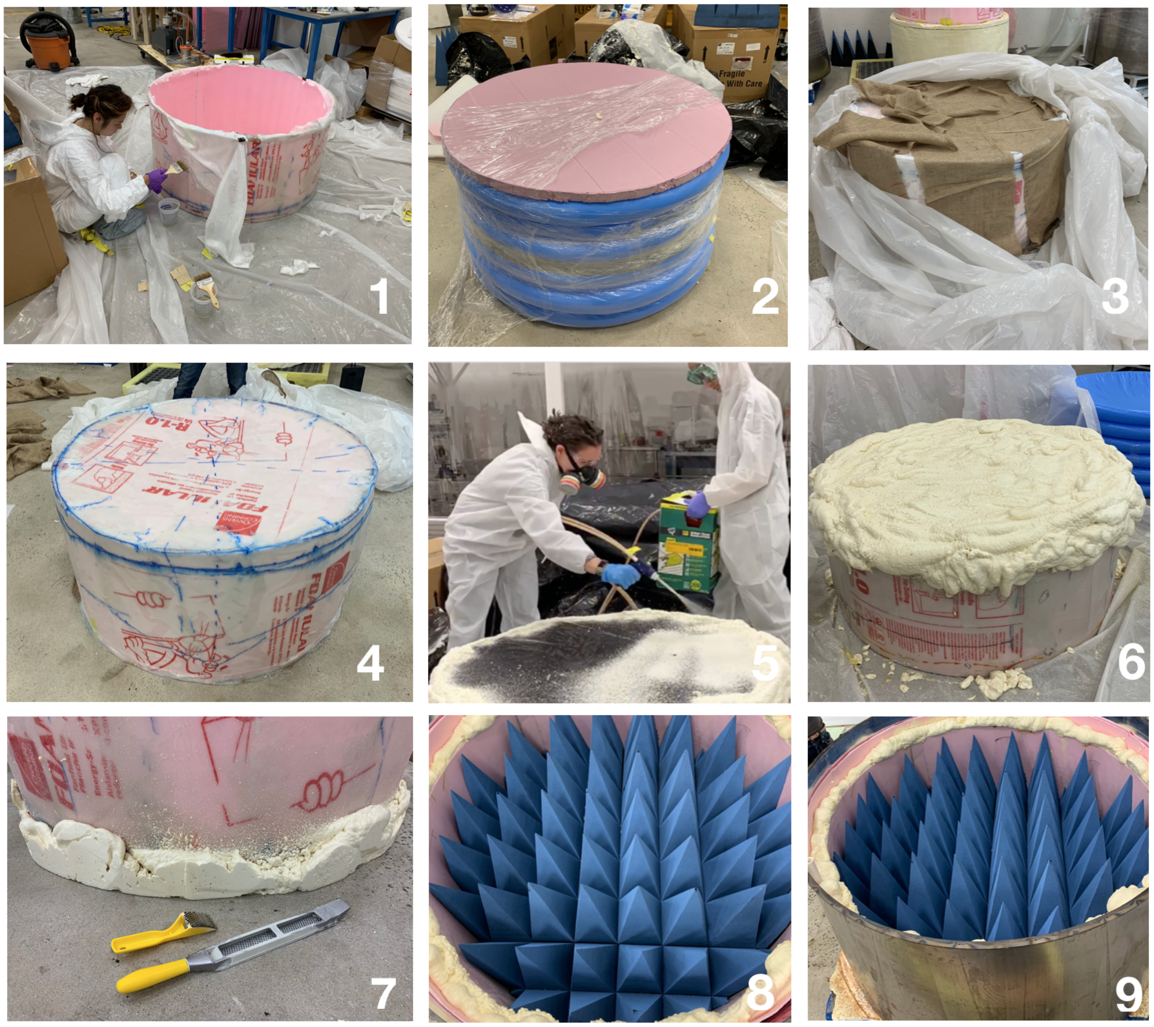

Figure 4.3: Foam and fiberglass insert construction procedure: (1) Laying down initial fiberglass cloth layers; (2) Inflatable tub used to keep insert shape during vacuum sealing; (3) fiberglass tub with peelply and burlap layers within an open vacuum bag; (4) fiberglass insert after vacuum bagging; (5) spray-foam procedure (photograph from earlier prototype); (6) spray polyurethane foam on base of insert; (7) leveling out foam on insert base; (8)(9) After filling in our remaining foam, we have an LN2-proof insert (subject to further testing), including the RF absorber for fit-checking.

once dry. We laid burlap sacks over the peelply as well, for soaking up additional resin. The entire system was then bundled in the plastic, which was ratcheted at the top with a zip tie to fully seal the inside. A vacuum pump and resin trap were attached to the bag, and the inner contents pumped out. Taken together, this set up allowed the extra resin to be absorbed, and prevented it from pooling in ways that would encourage cracking and structural instability. At this stage, we overlaid thicker layers on both the inside and outside surfaces, for stability and extra security, applied in smaller patches that could easily smooth down. With the additional layers of fiberglass cloth, the tub was able to hold nitrogen, which we tested over many thermal cycles.

With the fiberglass layer complete, we next sprayed the insert bottom with the cryogenically-rated foam and leveled it off with a wood scraper. We then placed it in the pink foam outer layer, which we held flush against the metallic tub (with a layer of thin plastic in between). We included a pvc pipe against the wall, so as not to draw on vacuum when pulling the insert out of the cavity. Then the cryogenic foam was sprayed evenly into the gap (process photos in figure 4.3). This full insert held nitrogen successfully over multiple thermal cycles.

The other key component of system construction was fitting the RF absorber blocks flush into the insulation insert. The construction of the RF absorbing insert involved forming blocks of RF absorber to the cylinder curvature. This required building a custom jig and foam cutter, as all of the available saws in the Yale Wright Lab facilities were not large enough to accommodate the 18in absorber height (fig. 4.4).

A schematic of the two cylinders is shown in Figure 4.5 (upper panel), and a photo of the constructed cylinders is shown in Figure 4.5 (lower panel). We have filled the cold cylinder on six occasions with $\sim$550 L of liquid nitrogen for a total of $>$ one month of measurement time, and it has maintained its structural integrity. It takes three 230 L nitrogen dewars to fill the cavity insert, and one dewar per day for refilling. The boil-off rate is $\sim$5.5 L/hour, with slight variations depending on seasonal climate conditions.

Over time, alterations were made to the system to improve performance and accommodate social distancing restrictions during the COVID-19 pandemic. The

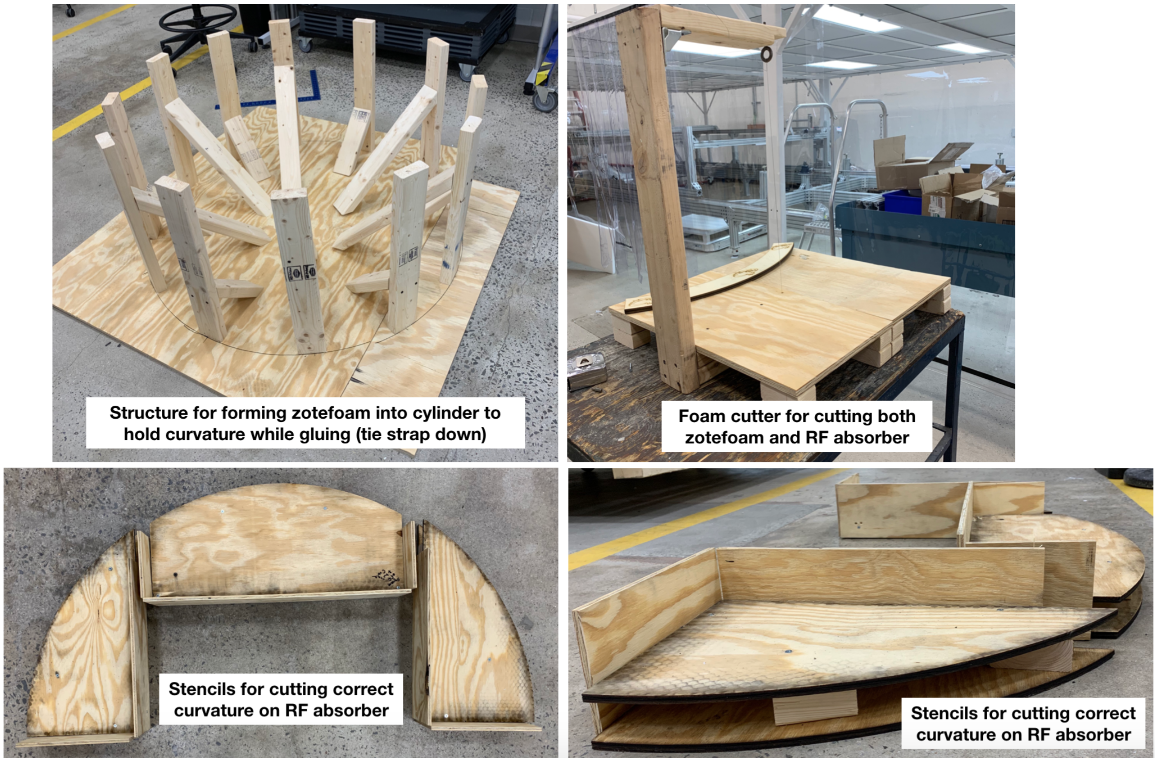
Structure for forming zotefoam into cylinder to hold curvature while gluing (tie strap down)
Foam cutter for cutting both zotefoam and RF absorber
Stencils for cutting correct curvature on RF absorber
Stencils for cutting correct curvature on RF absorber

Figure 4.4: Various woodworking projects aid with the construction of the y-factor measurement system.

cavity lids were stiffened by bolting them to angle iron, and a 12in x 12in cut out was added for the feed and its mounting plate to slot into. This second modification allowed measurements to more easily be performed by one person. We also painted both steel containers in order to minimize rusting, which had occurred at an alarming rate during the humid summer months.

There were several things that worked well for the current construction process. The fiber glassing was done with two people, which was critical for working within the resin cure time and preventing fatigue (more people would have been preferable). Extra strands of fiber were stripped off ahead of epoxying, to avoid creating sharp spikes on the insert surface. The vacuum bagging method worked well for quickly sealing the system to outside airflow, which was important given the cure time of the resin. Finally, using small bits of cloth to patch up holes and reinforce seams helped

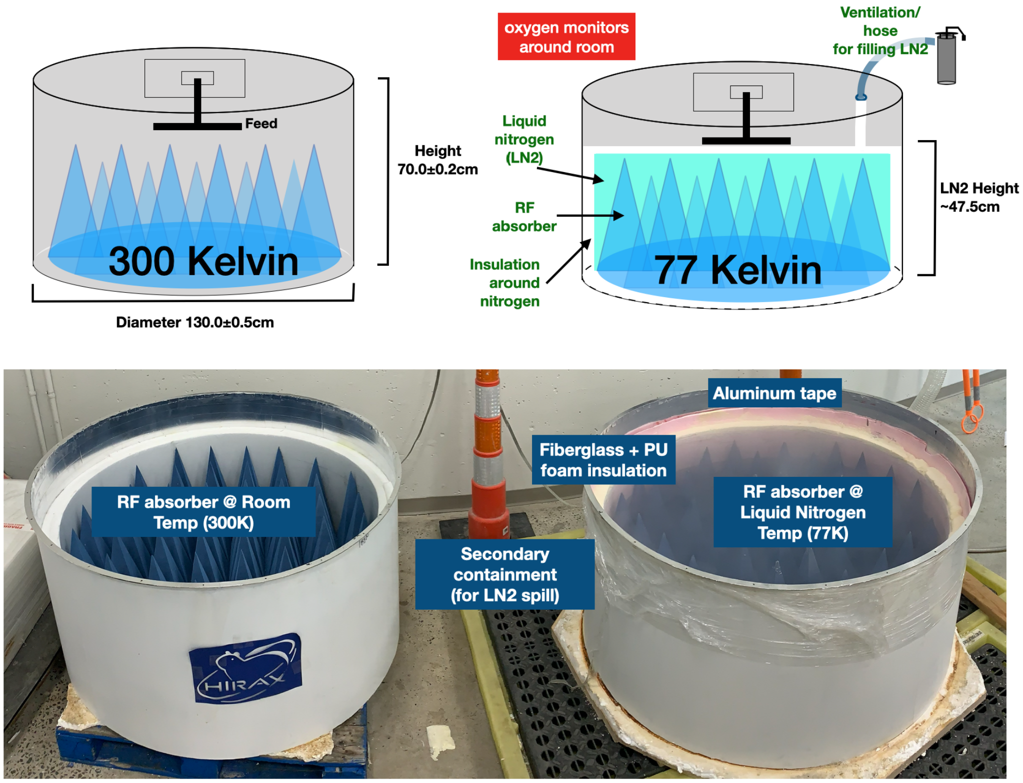
oxygen monitors around room
Ventilation/ hose for filling LN2
Feed
Height 70.0±0.2cm
300 Kelvin
Diameter 130.0±0.5cm
Liquid nitrogen (LN2)
RF absorber
Insulation around nitrogen
77 Kelvin
LN2 Height ~47.5cm
Aluminum tape
Fiberglass + PU foam insulation
RF absorber @ Room Temp (300K)
RF absorber @ Liquid Nitrogen Temp (77K)
Secondary containment (for LN2 spill)
HIRAX

Figure 4.5: Top: Schematic of the noise temperature measurement system design. Bottom: Labeled photograph of the hot/cold loads for noise temperature tests, taken during an LN2 fill. The loads are covered and sealed with a steel lid for measurements.

strengthen the structure. Still, there are many ways to improve on this procedure. We could involve additional personnel in the construction or use slower resin, as resin started to cure as we were applying it. We could also use a more powerful vacuum, as ours was slow to pump down, and some of the resin cured before the system was fully compressed.

4.2.3 Faraday cage construction

Early characterization measurements indicate that RFI is a contaminant to the data stream, with the potential to significantly bias the noise temperature measurement. Thus, a Faraday cage was constructed to enclose the measurement setup, as shown in figure 4.6. The cage additionally serves as an extra safety barrier between the nitrogen bath and spectators in the lab, and the structure provides a convenient way to support cables and relieve cable strain.

The Faraday cage took roughly two months to complete. It is 12ft long, 8ft high, and 8ft wide, with a frame made of aluminum T-slot material. Panels of 0.5in steel mesh are bolted onto the frame and then electrically connected by layers of aluminum tape with conductive adhesive. An adaptor plate is slotted into the frame and populated with with SMA and BNC connectors, for carrying the RF signal and powering amplifiers during measurement. The cage has three doors: one 4ft wide door situated by the building ventilation for use when filling liquid nitrogen, another 4ft door for main access during measurement, and one 8ft door (which the prior 4ft door is nested into) that largely remains latched, though can be opened for moving large objects into and out of the enclosure as needed. The doors are latched with draw latches and sealed with conductive gasket from digikey (part number 903-1099-ND). Flexible copper mesh is swen into the steel mesh over the hinges with 20AWG copper wire as an RF sealant. All seams in the cage are aluminum taped or otherwise covered with conductive material, forming an RF-tight enclosure.

Several verification measurements were performed to demonstrate that the RF cage is indeed RF-tight. For the main test, signal from a high-power calibrated noise source was transmitted by a directional antenna located outside of the cage, and the signal received was measured by an omnidirectional biconical antenna within the cage. The transmitter was moved to a variety of positions, maintaining a distance of 2m from the receiving antenna, and pointed directly at both the receiver and cage seams. In each instance, the measured power was recorded with the door both open and closed. From these measurements, it was determined that the cage provides 20dB isolation from stray RF signals, allowing for a cleaner noise temperature measurement.

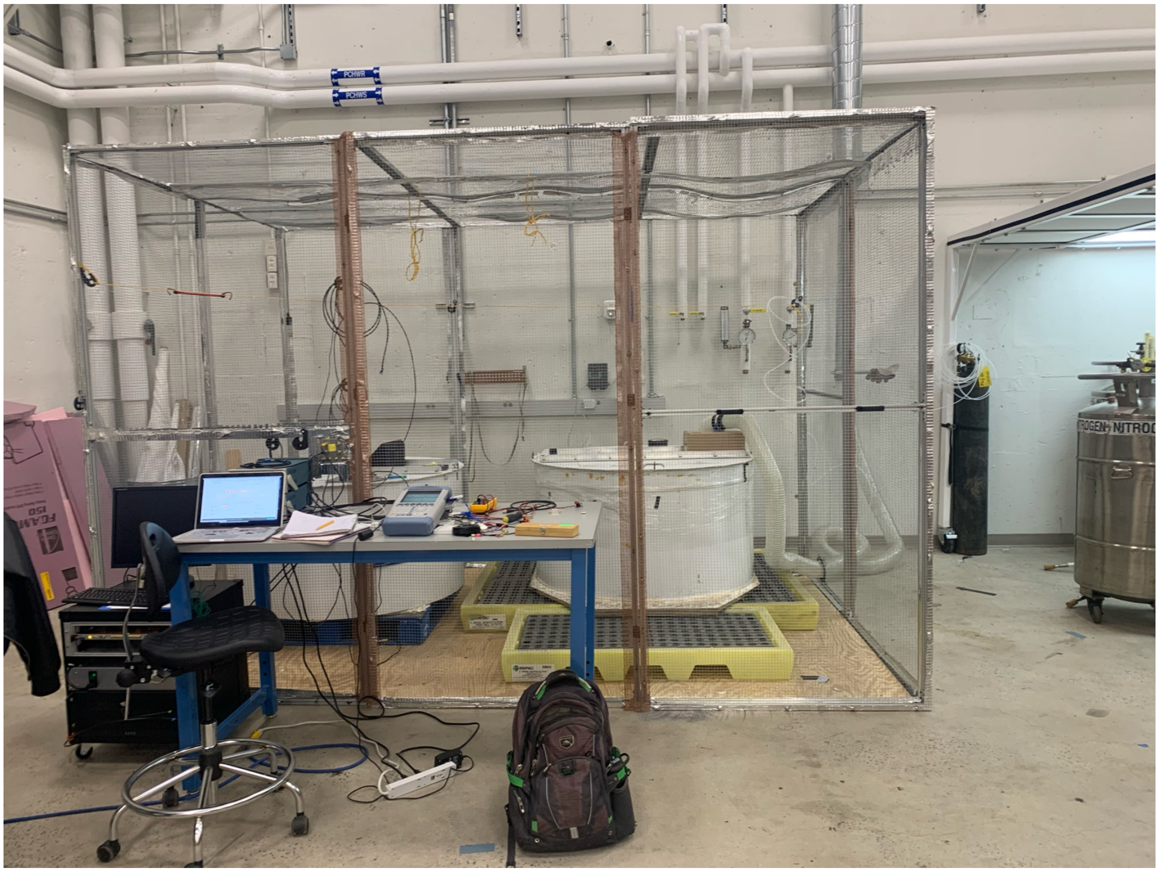

Figure 4.6: A custom Faraday cage built around the Y-factor measurement system. The cage was constructed for RFI mitigation, and demonstrations show it provides 20dB of attenuation for unwanted signals. The system is built from T-slot material and contains two doors, one on the side of filling and the other at the front for personnel access. The personnel door is embedded within a larger door, sized to enable the cylinders to be removed. The cage additionally serves as an extra safety barrier between the nitrogen bath and spectators in the lab, and the structure provides a convenient way to support cables and relieve cable strain.

4.3 System verification measurements

A variety of measurements were performed to characterize experimental uncertainties and verify that the cavities met specifications. As described below, these measurements include verifying the radio-frequency transparency of the foam materials, quantifying the degree of similarity between the two cavities, and assessing the RF spectrum of the absorbers in each cavity. The verification measurements were primarily performed with a CHIME feed (which for this purpose can be considered an unamplified HIRAX feed) to allow S_{11} measurements, which are not possible with the amplified HIRAX feed.

4.3.1 RF transparency tests of insulation materials

As described in Section 4.2.2, several layers of insulation were added into the experimental system to successfully contain the 550 L of liquid nitrogen required for this experiment. These layers must be RF transparent, as any absorption in the insulating foam will add an unquantified warm temperature component to the cold temperature measurement and bias the calculated noise temperature. To assess material transparency and inform the final design, we took S_{11} measurements[4] of the cavity at different stages of insulation construction, which occurred over the span of several months in 2019. These measurements were performed without RF absorber such that the cavity should be highly reflective, forming a sensitive measurement of absorption in the insulating foam.

The S_{11} measurements used to verify insulation transparency are shown in Figure 4.7. As noted, there is no RF absorber present in the cavity for these measurements, so the S_{11} value should be near 0.0 dB (indicating a fully reflective system). The median value across the spectrum is 0.3 dB, which can be attributed to losses in the feed. The lines marked ‘empty’ are measurements of the cavity devoid of any insulating foam. There are strong negative features in the S_{11} spectrum which are not present when RF absorber is added (see Figure 4.8), so we attribute these features to destructive interference from standing waves at a variety of characteristic distances inside of the cavity. The foam inserts were constructed over the course of six months, in order of: the cryogenic polyurethane foam (2019/03), the fiberglass insert (2019/06), and the full insulation layer (2019/08). In addition, the empty cavity was measured twice (2019/03 and 2019/06), with the S_{11} spectra changing by up to 0.2 dB. The typical change between the empty cylinder and the full insulation is <0.1 dB, within the range of fluctuations between the two empty measurements. This indicates that the insulating foam layer has no discernible absorption within measurement errors.

---

[4] A R&S FSH4 multi-purpose analyzer was used for these measurements, www.rohde-schwarz.com/us/product/fsh-productstartpage_63493-8180.html

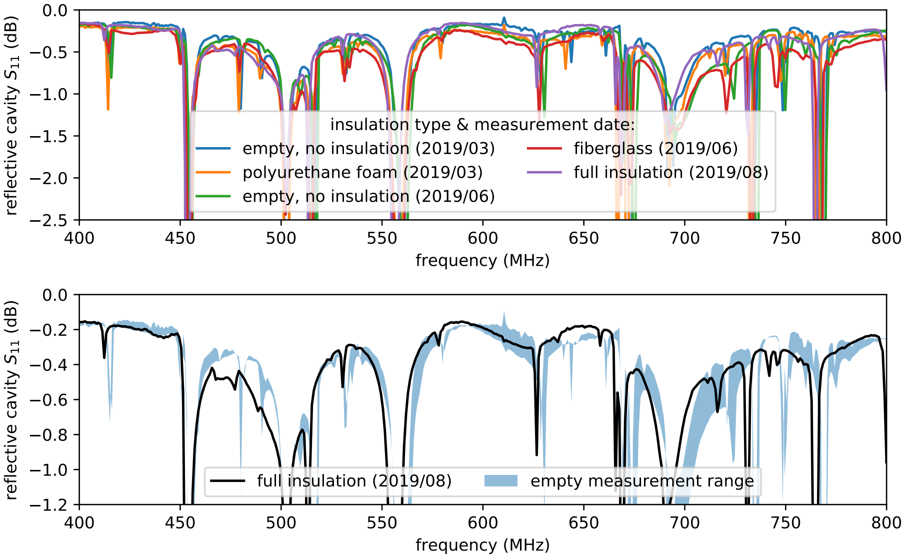

Figure 4.7: (Top) Return loss (S_{11}) measurements of the CHIME (unamplified HIRAX) feed in a reflective cavity as various insulation components are added in the construction. (Bottom) Return loss measurements of the cavity with full insulation compared with the range of empty measurements. The empty cavity was measured twice (2019/03 and 2019/06), with the S_{11} spectra changing by up to 0.2 dB. The insulation components are shown to be RF transparent to within this range of fluctuations for most of the band. Not shown is the addition of an aluminum tape layer, which also has a negligible impact.

4.3.2 Return loss with RF absorber installed

For the second set of verification measurements, we performed a series of S_{11} measurements of the full system (including RF absorber) at ambient temperature to verify: (1) that the system does indeed mimic free space and (2) that the two cavities are sufficiently similar to one another, within the -10 dB design specification for the antenna. These S_{11} measurements were taken over several months using a CHIME feed (similar to a passive HIRAX feed).

The S_{11} measured for a single polarization in both cylinders is shown in Figure 4.8. Also shown is the simulated feed in free-space for comparison, reproduced from Figure 4.2. The S_{11} profiles evolved slightly in time, but consistently remained at or below -12dB across the full band, indicating a well-matched antenna viewing a system simulating free space. These measurements matched simulations in overall S_{11} level, but had different resonance locations. Although the measurements do not fully agree with the free-space simulations, they appear consistent with measured free-space values shown for a similar feed built for the CHIME experiment[86], and so we attribute differences in resonance locations to differences between the modelled feed in CST and the as-built feed. Despite their differences, these measurements indicate that the cavities are similar to within specifications, with cavity differences accounting for sub-Kelvin uncertainty across our band. This will be explored in Section 4.4.

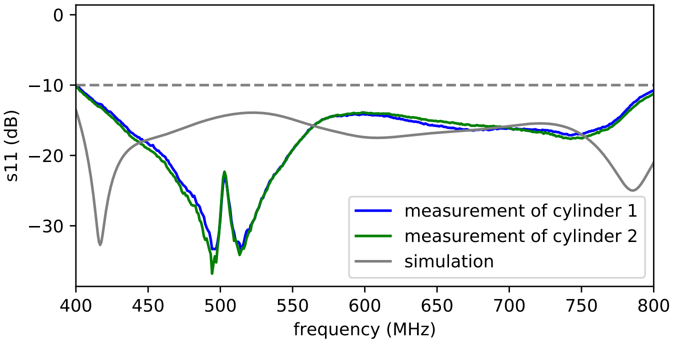

Figure 4.8: S_{11} measurements of the CHIME (unamplified HIRAX) feed in the warm cavities with RF absorber installed. We compare measurement results from the two cylinders that make up the Y-factor measurement system (both at 300 K), finding they are identical to sub-dB level, and share the same overall level as simulation results, though resonance peak locations differ. These return loss measurements match with published passive feed measurements of CHIME, which shares the HIRAX antenna design[86]. More details are described in the text.

4.3.3 Blackbody spectrum comparison

We can perform an additional check to verify that the RF absorber in the cavity is indeed functioning as a blackbody. The RF absorber should emit thermal radiation, and hence have a blackbody thermal spectrum. This aspect of system performance is verified by installing the unamplified feed in one of the cavities, amplifying the resulting signal with commercial amplifiers of known gain and noise temperature, and measuring the resulting spectrum with a spectrum analyzer[5]. For a ∼300 K absorber in the frequency range 400-800 MHz, the low-frequency approximation to the blackbody spectrum is valid, providing an estimated power of:

$$P = GkT\Delta\nu \tag{4.6}$$

where G is the device gain, k is Boltzmann's constant, T is the sum of the physical temperature and device noise temperature, and $\Delta\nu$ is the bandwidth in Hz. We assume a temperature T containing contributions from the temperature of the RF absorber (300 K), the estimated feed loss (20 K), and the noise temperature of the first amplifier in the amplifier chain (determined by the data sheets).

We compared this theoretical power with the measured power in dBm from the spectrum analyzer, using a variety of amplification chains. The amplifiers used in this measurement are commercially obtained from Mini-Circuits[6], and have available data sheets reporting gain and noise temperature. These amplifiers were chosen because they had good gain in the HIRAX band and have high enough compression points to ensure the measurements would be linear with input power. Using multiple amplification stages brought the signal well above the noise floor of the spectrum analyzer, and 3 dB attenuators were placed between the amplifiers to reduce reflections and oscillations in the amplified signal.

---

[5] Measurements are made with an R&S FSH4 Multi-purpose analyzer

[6] www.minicircuits.com/

The results comparing the inferred power to the expected blackbody spectrum are shown in Figure 4.9. The different amplifier chains agree with the expected spectrum to within 2 dBm, which is consistent with contributions we have not taken into account, such as estimated losses from the cable (<1 dB), temperature variations in the amplifier gain, and systematic errors in the absolute power measurement from the spectrum analyzer (<1 dBm, typ. <0.5 dBm). The additional features in the spectrum between 725-775 MHz occur near communication bands, which is evidence that we have not eliminated RFI in these measurements using the closed cylinder as the only RFI protection (the Faraday cage was not used for this measurement, but was employed as RFI mitigation for future measurements).

4.4 Systematics and uncertainties

In this section we describe and quantify the contributions of statistical and systematic errors to the noise temperature measurement error budget.

4.4.1 Statistical uncertainties

The noise temperature measurements are limited in integration time, as the antennas under test can cool while attached to the cold cylinder, thereby changing the noise temperature of the LNAs (LNA noise is temperature dependent and lower at lower temperatures). For initial data taking, we averaged 50 samples with sweep time 0.02 s for a total integration time of 1 s. This averaging takes <10 s, which is more than adequate to keep the feed from cooling (temperature effects still to be characterized). For an integration over 50 samples in 3 MHz frequency bins, statistical fluctuations in the individual frequency bins are limited to ±0.1 dBm. These fluctuations can be extrapolated to an error in noise temperature by standard error propagation for the

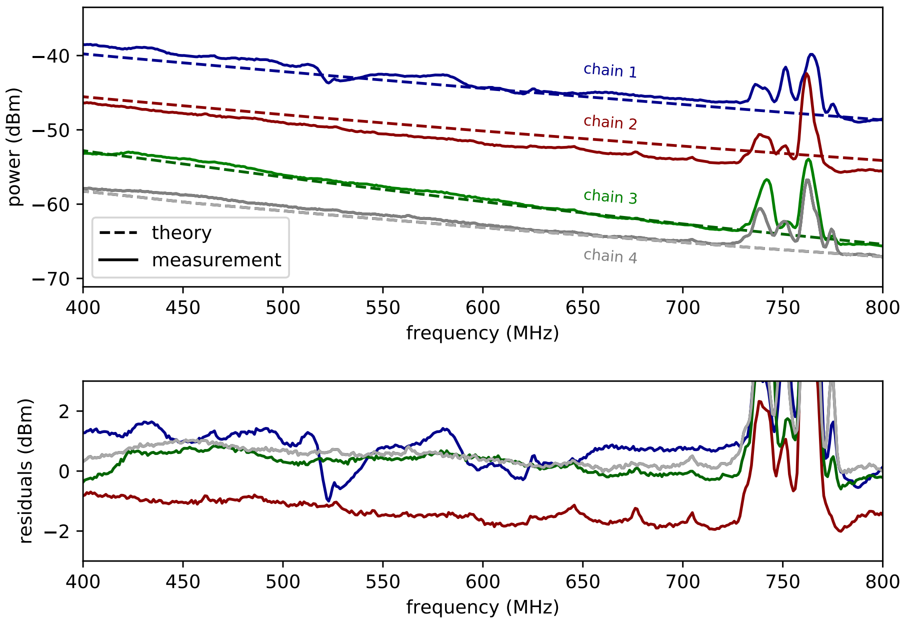

Figure 4.9: Theoretical (blackbody) spectrum, measured spectrum, and residuals for the CHIME (passive HIRAX) feed + four commercial amplifier chains. The top plot shows a comparison between the expected power of the amplifier chains at 300K (dashed lines) and the corresponding spectrum measurements in the experimental system (solid lines). The expected power is computed from $P = G_{\text{chain}}(\nu) k_B (T_{\text{load}} + T_{\text{loss}} + T_{\text{LNA,1}})\Delta\nu$, where T_{load} is the thermal load temperature, T_{loss} is the assumed feed noise temperature from material loss, and $T_{\text{LNA,1}}$ is the noise temperature of the first LNA in the chain. The bottom plot shows the residuals, revealing only slight differences between experimental and theoretical values (neglecting the features in the vicinity of 750 MHz) that could be accounted for by cable loss, gain uncertainty, and other systematics, verifying that we measure a blackbody. The chains are comprised of the following Mini-circuits amplifiers: **chain 1** = ZFL-1000H+ → ZX60-P103LN+ → ZX60-P103LN+; **chain 2** = ZX60-112LN+ → ZX60-P103LN+ → ZX60-P103LN+; **chain 3** = ZX60-P103LN+ → ZX60-P103LN+ → ZX60-P103LN+; **chain 4** = ZX60-P103LN+ → ZX60-P105LN+ → ZX60-P103LN+. All amplifiers have frequency dependent gain, and all chains include 9dB attenuation.

Y-factor linear calculation, summing errors in measured P_{hot} and P_{cold} in quadrature, to obtain ± 4.82 K for a 30 K assumed noise temperature. This noise can be smoothed through further binning, and fluctuations in noise temperature can be integrated down between successive noise temperature measurements to further reduce noise. We also observe slight fluctuations in the average spectrum level to within ±0.01 dBm over the course of a one-hour measurement, corresponding to a ± 0.48 K uncertainty.

4.4.2 Contribution of reflections to uncertainties

A variety of systematics must be considered for our measurements, including reflections from various system components such as the liquid nitrogen and zotefoam insulation lid. We can use passive feed S_{11} measurements to bound how cavity differences will impact the noise temperature measurements. The measured S_{11} contains three contributions: (i) signal reflected within the feed structure, (ii) losses in the feed (i, ii inherent to the feed), and (iii) signal reflected back to the feed within the cavity (which depends on the cavity and its interplay with the feed). Here we use the measured S_{11} without distinguishing between these contributions since they cannot be differentiated within our test setup. We follow the scheme detailed below:

The amount of energy radiated from a load can be expressed as an equivalent brightness temperature[90],

$$T_B = (1 - \Gamma^2)T \tag{4.7}$$

where Γ is the load reflection coefficient and T is load temperature. Supposing a noise temperature T_{noise}, the power measured from this load during a Y-factor measurement would be:

$$P_B = Gk[T_B + T_{\text{noise}}]\Delta\nu = Gk[(1 - \Gamma^2)T + T_{\text{noise}}]\Delta\nu. \tag{4.8}$$

For the purposes of estimating uncertainty, the reflection coefficient Γ can be obtained from S_{11} via

$$\Gamma = 10^{S_{11}[\text{dB}]/20}. \tag{4.9}$$

S_{11} measurements of the Y-factor system provide an upper bound on load reflections, as Γ includes reflection contributions from the feed itself as well as from the RF absorber. The following set of equations shows the noise temperature computation with and without corrections for the reflections

$$T_{\text{noise}}^{\text{true}} = \frac{(1-\Gamma_{\text{hot}}^2)T_{\text{hot}} - Y(1-\Gamma_{\text{cold}}^2)T_{\text{cold}}}{Y-1} \tag{4.10}$$

$$T_{\text{noise}}^{\text{uncorrected}} = \frac{T_{\text{hot}} - YT_{\text{cold}}}{Y-1} \tag{4.11}$$

where $T_{\text{noise}}^{\text{true}}$ is the true noise temperature and $T_{\text{noise}}^{\text{uncorrected}}$ is the noise temperature one would naively calculate by applying a Y-factor computation without accounting for reflections. When conducting a noise temperature measurement, Y is measured directly, but for the purposes of error assessment we compute a theoretical Y value from expected noise temperature and measured reflections:

$$Y = \frac{P_{B,\text{hot}}}{P_{B,\text{cold}}} = \frac{(1-\Gamma_{\text{hot}}^2)T_{\text{hot}} + T_{\text{noise}}}{(1-\Gamma_{\text{cold}}^2)T_{\text{cold}} + T_{\text{noise}}} \tag{4.12}$$

We estimate uncertainties from reflections as the difference between the "true" and "uncorrected" cases (eqs. 4.10, 4.11) for various measurement scenarios. We take T_{hot} = 300 K, T_{cold} = 77 K, and T_{noise} = 30 K, and compute Γ_{hot} and Γ_{cold} directly from S_{11} measurements in the ambient temperature/cryogenic cavities.

Using this description, we estimate the impact of reflections in two different regimes:

- When $S_{11} = S_{11,\text{hot}} = S_{11,\text{cold}}$, the noise temperature is modified by $T_{\text{uncorrected}} = T_{\text{true}}/(1-\Gamma^2)$, so $T_{\text{uncorrected}} > T_{\text{true}}$ (this is the case where the two cylinders are identical and the liquid nitrogen surface is completely sub-dominant). In this regime, we can use the Y-factor measurements without temperature modifications to get an upper bound on noise temperature.
- If $S_{11,\text{hot}} \neq S_{11,\text{cold}}$, and they are sufficiently discrepant, we are no longer guaranteed an upper bound on noise temperature. For a noise temperature of 30 K, and $S_{11,\text{hot}}$ = -15 dB, the expected level of the free space HIRAX feed profile,

we require $S_{11,\mathrm{cold}} < -13.5\,\mathrm{dB}$ to maintain this upper bound.

Figure 4.10 (right) shows the biases we incur when not accounting for reflections in these two cases. Without liquid nitrogen, we are in the regime where $S_{11,\mathrm{hot}} = S_{11,\mathrm{cold}}$, and find the uncertainty at most $\sim$1 K (black curve). When liquid nitrogen is introduced, there are new reflections ($S_{11,\mathrm{hot}} \neq S_{11,\mathrm{cold}}$), giving at most ± 3.5 K uncertainty. These uncertainties provide a lower bound at some frequencies and upper bound at others, and remain within 2 K for much of the frequency band. Although a 2 K error is within the specification, it is significant enough that removing or accounting for some of the effect of these reflections is desirable.

We measured the S_{11} of the system at various nitrogen depths, as shown in Figure 4.10 (left), to assess the reflections off of the liquid nitrogen layer. This measurement revealed that as the distance from the liquid nitrogen surface changes, the location of the 'dip' feature, which we associate with constructive interference in the cavity, shifts by $\sim$30 MHz. The profile continued to change when the cavity was almost full of nitrogen, at a stage when material contractions due to cooling would be complete and RF absorber was nearly submerged, so we suspect this feature is due to reflections from the liquid surface layer and not contraction or altered RF properties from cooling. Still, it is possible that the RF absorber shrinks in cold conditions, and we are currently investigating the impact of perturbations in absorber size in CST.

In addition to reflections from liquid nitrogen, it is important to consider the reflections that might occur off of the insulation lid. This lid is made from zotefoam, which is known for its RF transparent properties. The insulating lid sits directly below the antenna under test, shielding it from the cold nitrogen gas. It is important to understand the lid transparency, as any absorption in the zotefoam could introduce a >77 K component into the 77 K measurement in a way that is difficult to quantify. To assess the lid RF reflections, we measured the S_{11} of the cold system with and without the zotefoam lid on, finding the measurements identical to $\sim$0.2 dB for values

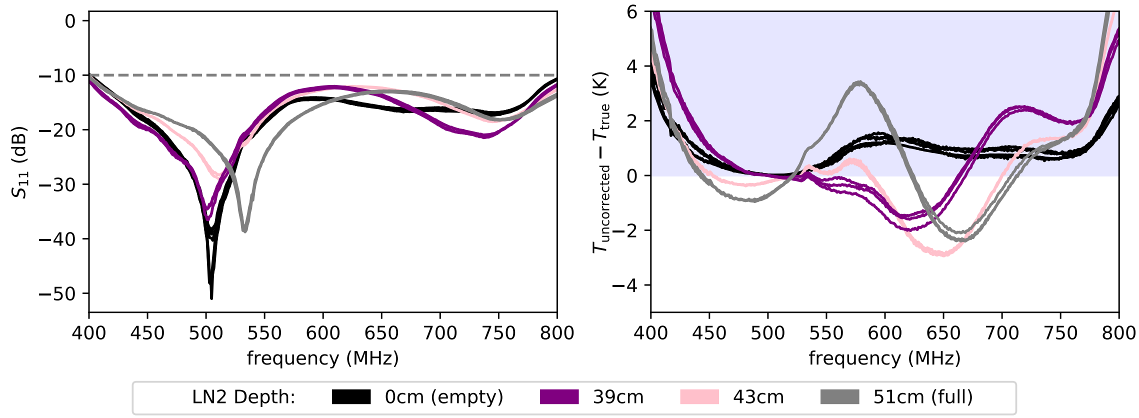

Figure 4.10: S_{11} profiles (as measured by the CHIME feed) at various liquid nitrogen depths and associated uncertainty in noise temperature. (Left): The return loss S_{11} measurements of the cold load exhibit a resonance that shifts in frequency as the cavity fills with liquid nitrogen. Though the resonance shifts, the S_{11} profiles remain below -10 dB (dashed line) across the full band, which is within the design specification for the feed. (Right): The uncertainty in noise temperature measurements is computed as described in the main text (equations 4.7 to 4.11). The shaded region indicates when this uncertainty provides an upper bound on the measurements, which is the preferred regime. Outside of this regime, the uncertainties remain below approximately 3 K, which is already a cautiously high bound.

above -15dB, indicating these reflections provide a negligible contribution to the uncertainty.

We also consider the impact that asymmetries in the system might have on reflections, and how that propagates into noise temperature. To assess system symmetry, we took S_{11} measurements at a series of feed rotations (rotating in polarization angle), considering angles 0, 30, 60, and 90 degrees. Comparing these measurements, we see no discernible difference in S_{11} (order <0.25 dB discrepancies for S_{11} >-17 dB), and determine system asymmetries to have negligible contributions to uncertainty. To further support this determination, we also performed Y-factor measurements at the same series of angles, and found no discernible difference in the corresponding noise temperature results.

4.4.3 Estimated contribution of cavity differences to noise temperature calculations

Differences between cavities can cause systematic errors in the measurements, and because only one cavity is filled with liquid nitrogen the cavities cannot be interchanged. We can assess the magnitude of these systematic errors through a series of tests at both ambient and cryogenic temperatures.

For the first of the analyses, we use a HIRAX feed to measure spectra from both cavities when warm, and apply a -4.89 dB offset to the measurement from the cylinder that will be filled with nitrogen. This result corresponds to the level we would expect from a cold spectrum measurement, assuming a 30 K noise temperature. We take the warm cavity result as P_{hot} and the offset result as P_{cold} and use the Y-factor method with $T_{\text{hot}} = 300\,\text{K}$, $T_{\text{cold}} = 77\,\text{K}$ to compute a noise temperature. The resulting 'mock' noise temperature should have a mean value of 30 K, with spectral features that are generated only by differences between the spectra in the two cylinders. The results are shown in Figure 4.11, where the mean noise temperature of 30 K has been removed, showing only the expected variations from differences between the cylinders. This feature, a few Kelvin in size, indicates a discrepancy between the RF properties of the two cavities. It is consistent across polarizations and repeatable across feeds and measurement days. It dampens slightly, though remains, when RF absorber is swapped between two systems (demonstrated in Figure 4.11). These results suggest there could be slight geometric differences in cylinder construction or a difference in insulation transparency (one cylinder has different insulation than the other, for ease of construction).

This sinusoidal feature is recovered in real noise temperature data from hot/cold measurements using a cryogenic load. We take a Y-factor measurement using the same cavity for both hot and cold measurements, where the hot measurement is taken just before the liquid nitrogen fill and the cold measurement is taken directly

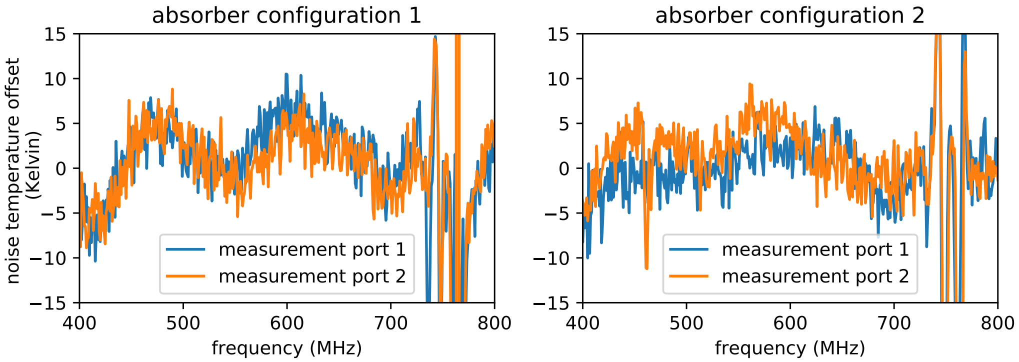

Figure 4.11: Deviations from the expected 30 K noise temperature due to inherent differences between the two loads (measured at 300 K). The frequency dependent feature is common across both polarizations and different feeds, and remains when absorber configurations are switched between the two systems. This offset can be removed from the final measurement results.

after. Subtracting this measurement from a measurement taken using the two different cavities at two different temperatures removes any spectral features common to both measurements (e.g. from the noise temperature) and yields features that result from cavity differences. This subtraction is shown in Figure 4.12 and reveals the same frequency-dependent profile seen in the warm measurement comparison detailed above. Their strong agreement indicates that the systematics from cavity differences can be subtracted out from measurements to obtain the final noise temperature results.

4.4.4 Additional sources of error

We anticipate additional sources of error in our measurements that we are working to mitigate. These include amplifier drifts, feed temperature fluctuations, cable effects, airflow over the nitrogen surface, and RFI. The full error budget is shown in Table 4.1.

We have observed long timescale drifts and fluctuations in amplifier gain. Over the course of several months, the amplifier gain appears to drift up to a few dB, presum-

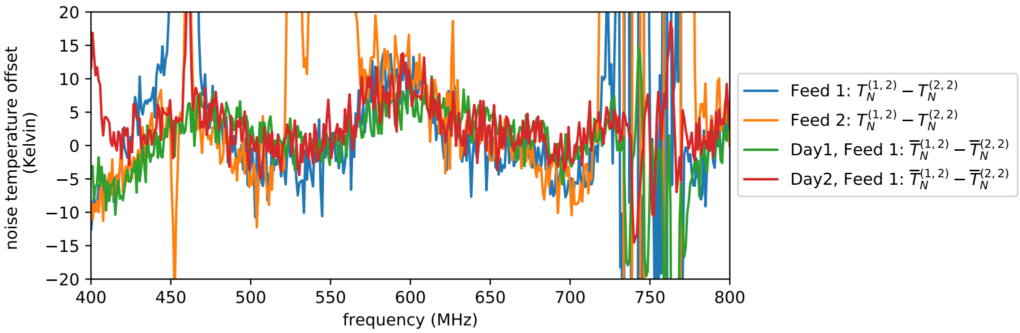

Figure 4.12: Characteristic offset in noise temperature measurements due to discrepancies between the two RF cavities, as predicted by warm verification measurements and verified by Y-factor measurements. Here, $T_N^{(1,2)}$ denotes the noise temperature measured with cavity 1 (at 300 K) as the hot load/ cavity 2 (at 77 K) as the cold load. We take $\overline{T}_N^{(1,2)}$ to denote the predicted noise temperature computed from two warm measurements (also shown in Figure 4.11) using a spectrum measurement in cavity 1 (at 300 K) as the hot power/ a spectrum measurement in cavity 2 (at 300 K) minus 4.9 dBm as the “cold” power (the 4.9 dBm offset is consistent with measuring 77 K). We plot the difference between the noise temperature measured using two different loads for hot/cold measurements and using the same load for hot/cold measurements, along with the difference in predicted noise temperature for two different loads and the same load. The recovered feature is highly consistent across different feeds and different measurement days, and as a result can be removed from final measurements.

ably in accordance with temperature and humidity conditions in the lab environment. Although we plan to measure noise temperature on far shorter time scales (typically no more than 5 minutes between a hot and cold measurement), these observations highlight the importance of documenting ambient temperature and humidity during the measurements in case absolute power measurements must to be compared across longer time scales at a later point.

The long-term gain drifts suggest that lab conditions like temperature and humidity impact antenna gain. These effects are critical to consider in the measurement scheme, as we cannot completely isolate the antenna from the cool nitrogen gas leaking out of the insulation when making a cold measurement, and therefore cannot guarantee it measures both loads at the same physical temperature. We have a preliminary data set to characterize this effect, where we continuously measured the spectrum from each load for 10 minutes (the largest timescale we'd consider for one measurement, and ample time for the feed temperature to equilibrate), and then compared the relative drift of spectrum medians in that time frame. Although these measurements suggest no discernible drift, if further analysis indicates this is problematic we can measure the cold load for multiple, shorter time periods. We seek to further mitigate temperature and humidity effects by enclosing the front of the antenna in a zotefoam box through which we flow room-temperature nitrogen gas. This precaution will help keep the antenna at 300K for the full measurement and reduce frost build up.

In addition to long-term gain drifts, we notice that the HIRAX amplifiers and the spectrum analyzer both take up to an hour to warm up and stabilize in gain and readout when first plugged in. These effects are mitigated by keeping all amplifiers under test powered for at least one hour prior to measurements.

We also consider the effects of using lossy cables and measuring too close to the amplifier noise floor. The noise temperature of our set up propagates as,

$$T_{\text{noise}} = T_{\text{LNA}} + \frac{T_{\text{cable}}}{G_{\text{LNA}}} + \frac{T_{\text{LNA, 2}}}{G_{\text{LNA}} G_{\text{cable}}} + ... \tag{4.13}$$

(following noise of cascaded devices treatment from *Microwave Engineering* [91]). From this equation, we see that noise temperature contributions from cable loss should be negligible in the measurement, as the cable sits behind $G_{\text{LNA}} = 40\,\text{dB}$ of amplification in the signal chain. Assuming cautiously that $T_{\text{cable}} = 300\,\text{K}$, $G_{\text{cable}} <$ -3 dB and $T_{\text{LNA,2}} < 100\,\text{K}$, cable loss and second-stage amplifier effects would contribute $<0.05\,\text{K}$ to the final result. Similarly, measuring near the spectrum analyzer noise floor is not of concern. While measuring close to the noise floor will bias a noise temperature measurement, this is avoidable by adding appropriate second-stage LNAs into the measurement chain.

Previous experiments have found that rapid airflow over a nitrogen surface can artificially raise the nitrogen temperature by 2 degrees[99]. This was initially a concern for this experiment—for safety reasons, we draw the nitrogen gas expelled from the tank into the laboratory's HVAC system, and were worried about drawing O2 through the tank in the process. However, this is very likely not an issue in the set up, as we are careful to only remove the gas once it has already been vented, and the zotefoam insulation lid forms a seal that prevents rapid O2 draw. As a worst case, if we were to mistake T_{cold} as off by 2K, we would see a 3.5K error in the noise temperature. This error is more than we can tolerate, given existing sources of uncertainty from the nitrogen surface and cavity discrepancies.

Category	Source of Error	Projected Uncertainty
statistical uncertainty	spectrum analyzer fluctuations	$<5\,$K*
	spectrum analyzer & amplifier gain drift	$<0.5\,$K
characterized together	liquid nitrogen surface reflections	$<3.5\,$K*
	ln2 fill cavity contractions	$\ll 3.5\,$K
	insulation reflections	$\ll 3.5\,$K
characterized separately	cavity RF differences	$<8\,$K*
	airflow over ln2 surface	$\ll 3.5\,$K
	cavity asymmetry	$<0.1\,$K
	cable effects	$<0.1\,$K
	amplifier compression	$<0.1\,$K
characterization in progress	temperature/humidity effects	—
	RFI	—
	back lobe size	—

Table 4.1: Experimental uncertainties and systematics. The main contributors to the error budget are spectrum analyzer fluctuations, nitrogen surface reflections, and cavity differences. These have been marked with a * and are removable or able to be averaged down, as discussed in the text. Additional sources of error are under investigation.

4.5 Cryogenic verification measurements

The set of verification measurements described in section 4.3 have focused on quantifying sources of systematic error and assessing whether the cylinders are similar to within specifications. We can now perform noise temperature verification measurements with the cryogenic system as a check on our methods.

For cryogenic verification measurements, we measure the noise temperature of the HIRAX passive feed connected to commercial amplifiers of known noise temperature and gain as the first stage LNAs (amplifiers listed amplifiers are from the set of four used in Figure 4.9, and are detailed in table 4.2). The noise temperature of the commercial amplifiers is independently verified with a commercial noise figure meter, and the gain is measured with both a noise figure meter and VNA. The noise figure meter measurements described in this section were performed by Ema Smith.

By comparing the noise figure and Y-factor measurements of the same amplifiers, we can assess whether the noise temperature system has met required specifications. From the Y-factor measurements, we expect to recover the amplifier noise temperature plus some consistent contribution from the antenna loss, which we estimate at 10-20 K from simulations.

4.5.1 Noise figure meter measurements

Noise figure is a measure of how the the signal to noise ratio is degraded by a device. It is directly related to noise temperature through the following equation:

$$\mathrm{NF}\ [\mathrm{dB}] = 10\log_{10}\left(\frac{T_{\mathrm{noise}}}{T_{\mathrm{ref}}} + 1\right). \tag{4.14}$$

Conversely, noise temperature can be converted to noise figure:

$$T_{\mathrm{noise}}[K] = T_{\mathrm{Ref}} \times (10^{\mathrm{NF[dB]}/10} - 1). \tag{4.15}$$

In the most cases, $T_{ref} = 290\mathrm{K}$, a convention we take for this chapter.

Noise figure is commonly measured with a device called a “noise figure meter”. A noise figure meter works by performing a Y-factor measurement with 50Ω as a cold load and a calibrated noise source as a hot load. The device requires a calibration that removes its own noise contribution as well as the noise from associated cabling and other elements in the measurement chain. Bandwidth and measurement averaging can be adjusted to reduce uncertainty.

The calibrated noise source output is usually given in "excess noise ratio" or ENR. ENR can be thought of as the amount of noise above a 290K thermal radiator. At high ENRs (>15dB), it can be translated to temperature with the following equation:

$$T_{\text{noise}} = -174\text{dBm/Hz} + ENR \tag{4.16}$$

where -174dBm/Hz is the amount of power in a one hertz bandwidth radiated by a 290K black body. Once the device ENR is understood, it can be input into the noise figure meter as a calibration table, which is referenced for the calibration procedure.

In our particular set up, we use the Agilent-N8975A noise figure meter[7] and HP346B calibrated noise source[8]. The noise source has ENR ranging from 15.05-15.75dB across its 0.01-18.00GHz operating band, and is powered directly by the noise figure meter. Noise source output power is fed directly to the input of the amplifier under test, which outputs amplified noise to the noise figure meter. The RF chain is supported by mini lab jacks to eliminate strain on connectors. The full set up is shown in fig. 4.13.

The noise figure meter can measure noise temperature to 2K precision. This specification is verified experimentally through repeated measurements of a constant set up over multiple days, to understand statistical deviations and identify systematics. This particular noise figure meter also simultaneously provides gain measurements,

---

[7] https://www.keysight.com/us/en/product/N8975A/noise-figure-analyzer-10-mhz-to-265-ghz.html

[8] https://www.keysight.com/us/en/product/346B/noise-source-10-mhz-18-ghz-nominal-enr-15-db.html

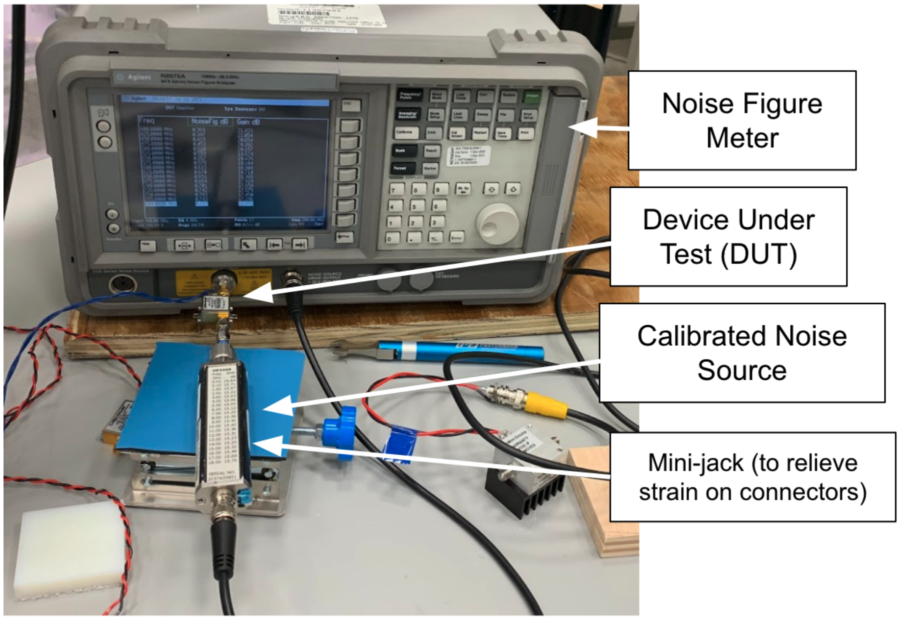

Figure 4.13: Schematic of the noise figure measurement set up used for amplifier noise assessment. The set up includes an Agilent-N8975A noise figure meter, HP346B 15dB ENR calibrated noise source, amplifier under test (the DUT), and a mini-jack to relieve strain on connectors. Additional attenuators were included in the chain for the measurements utilized in this thesis, though they are not pictured in the diagram.

which will be needed for the noise temperature verification, and were checked separately with a VNA. The Agilent-N8975A is an old model, so the data is read out to a floppy disc, which then is converted to CSV. The data read-in and measurements were performed by Yale undergraduate Ema Smith.

In addition to standard noise figure meter operating procedure, there were certain measurement best practices that were adopted over the course of several months, including adding attenuators of differing values to the front of the RF chain, between the noise source and amplifier. This change was made because initial measurements showed ripples in the noise figure curve when plotted with frequency, which we trace to an impedence mismatch between noise source and amplifier introducing reflections. Adding in an attenuator dampens those reflections.

We completed noise figure measurements of the three commercial amplifier models ZX60-112LN+, ZX60-P103LN+, and ZX60-P105LN+, measuring at least two amplifiers per model to develop specs on gain, noise figure, and noise temperature. Statistical and systematic uncertainties in noise temperature were determined to be 2K. Results were in line with data sheet values to within uncertainty for the ZX60-P103LN+ and ZX60-P105LN+ amplifiers, and showed improvement over the data sheet for the ZX60-112LN+ amplifier. Measurements are shown in table 4.2 and figure 4.14.

Amplifier Part #	**Gain (dB)**	**Noise Figure (dB)**	**Noise Temp (K)**
ZX60-112LN+	26.40 - 26.68	1.03 - 1.11	78- 84 ± 2K
ZX60-P103LN+	16.84 - 21.44	0.49 - 0.68	34 - 49 ± 2K
ZX60-P105LN+	14.75	1.81 - 1.96	150-165 ± 2K

Table 4.2: Specifications in 400-800MHz for the three commercial minicircuits amplifiers used for cryogenic system verification. These values were measured using a noise figure meter and vector network analyzer.

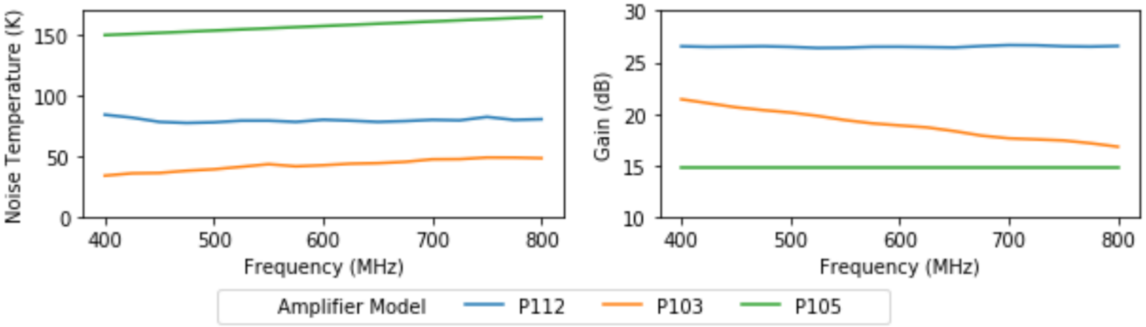

Figure 4.14: Gain [dB] and noise temperature (K) measurements of the three commercial minicircuits amplifiers used for cryogenic system verification, taken with an Agilent N8975A noise figure meter.

4.5.2 Passive feed noise temperature

The noise temperature T_1 of a passive element at ambient temperature T_{ref} with loss L can be expressed as,

$$T_1 = T_{\text{ref}}(L - 1). \tag{4.17}$$

As before, we take $T_{\text{ref}} = 290$. Loss (L) and gain (G) are related by,

$$G = \frac{1}{L}. \tag{4.18}$$

With this information, we can rewrite the equation for noise temperature of cascading devices (eq. 4.13) for the case of a lossy element as the first stage device (temperature T_1, loss L, where the quantities are related per equation 4.17) followed by a chain of amplifiers (T_{chain}).

$$\begin{aligned} T_{\text{noise}} &= T_1 + \frac{T_2}{G_1} + \frac{T_3}{G_1 G_2} + ... \\ &= 290(L-1) + L\left(T2 + \frac{T_3}{G_2} + ...\right) \\ &= 290(L-1) + LT_{\text{chain}} \end{aligned} \tag{4.19}$$

We can then solve for the first-stage element loss (L) as a function of T_{noise}, which is measured directly with the Y-factor system, and T_{chain}, which is computed from the noise figure and gain measurements of the individual amplifiers.

$$L = \frac{T_{\text{noise}} + 290}{T_{\text{chain}} + 290} \tag{4.20}$$

We assume ohmic losses as the primary source of loss, and will consider the impacts of impedance mismatch and lower-level system reflections in future work.

For our particular set of verification measurements, the first stage is a passive feed, and a schematic of the measurement set up is shown in fig. 4.15. The amplifiers used are the set previously characterized by the noise figure meter (table 4.2, figure 4.14).

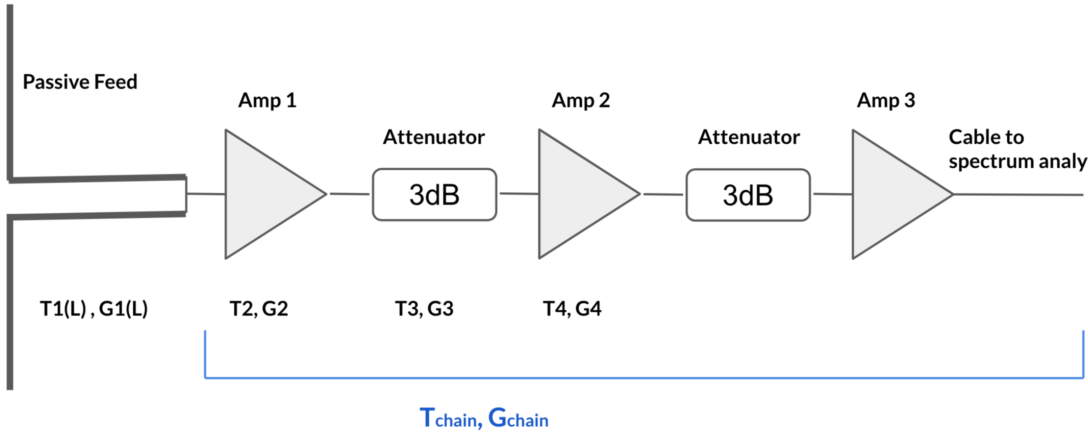

Figure 4.15: Measurement chain used for the for the cryogenic verification measurements. In this set up, we measure the noise temperature of a passive feed connected to a chain of previously characterized amplifiers. From this measurement, the feed loss can be recovered and compared to an expected value.

Y-factor measurement procedure

The measurement procedure is the same for the passive feed verification measurements discussed in this section and the active feed measurement described in the subsequent section. It is as follows: We first mount the antenna under test on a 12 in × 12 in ground plate that slots into the cavity lid, and (when applicable) power the antenna through a bias-t along an SMA cable. We then bolt the mounting plate into place on the "hot" cylinder lid, so the antenna is directly facing the RF absorber. We initiate data taking on the spectrum analyzer, and leave the antenna passively taking data (integrating with 3 MHz bandwidth for 1 s total integration time), saving a file every 10 seconds for 5 minutes. At this stage, we move the antenna and plate over to the cold cylinder and bolt the plate onto its lid. We leave it passively data taking for 5 minutes. This process is repeated three times.

These verification measurements were undertaken for two different antennas: a passive HIRAX feed made of FR4 and a CHIME feed made of teflon. These two antennas share design, dimensions and basic functionality, with the primary difference being the material makeup (a more detailed description of the other differences can be found in Ch. 3).

Verification results

We can solve for the CHIME antenna noise temperature from a Y-factor measurement of the CHIME feed and known amplifier chain (T_{noise}), a noise figure meter measurement of the amplifier chain alone (T_{chain}), and an assumed ambient temperature (T_{ref}), following methods above (fig. 4.15). We combine equations 4.20 and 4.17 to solve for the CHIME antenna noise temperature T_{antenna} (taking $T_{\text{ref}} = 290\text{K}$),

$$T_{\text{antenna}} = T_{\text{ref}} \left(\frac{T_{\text{noise}} + 290}{T_{\text{chain}} + 290} - 1 \right). \tag{4.21}$$

We solve for T_{antenna} for three different amplifier chains, and find the CHIME antenna noise temperature recoverable to within 3K for each chain. We repeat this procedure with the passive HIRAX antenna measurements and recover the result to within 10K. Figure 4.16 shows the antenna noise temperature and the residuals extracted from measurements of the antenna plus amplifier chain (amplifiers from table 4.2 and figure 4.14), with the legend indexed by the first LNA in the chain.

There are a few features to note in the passive antenna noise temperatures (fig. 4.16). The average level of CHIME is as expected: it is comparable in temperature to an SMA connector, which should dominate the CHIME noise given that it is make up of extremely low loss material. There is also a sizable ripple in the noise temperature—a behaviour we similarly observe with the HIRAX active feed (discussed later in the chapter). This warrants further investigation, paritcularly as the ripple in the noise temperature approaches zero at some frequencies, which may be unphysical and suggestive of destructive interference.

There is one more step that will help further quantify the Y-factor system per-

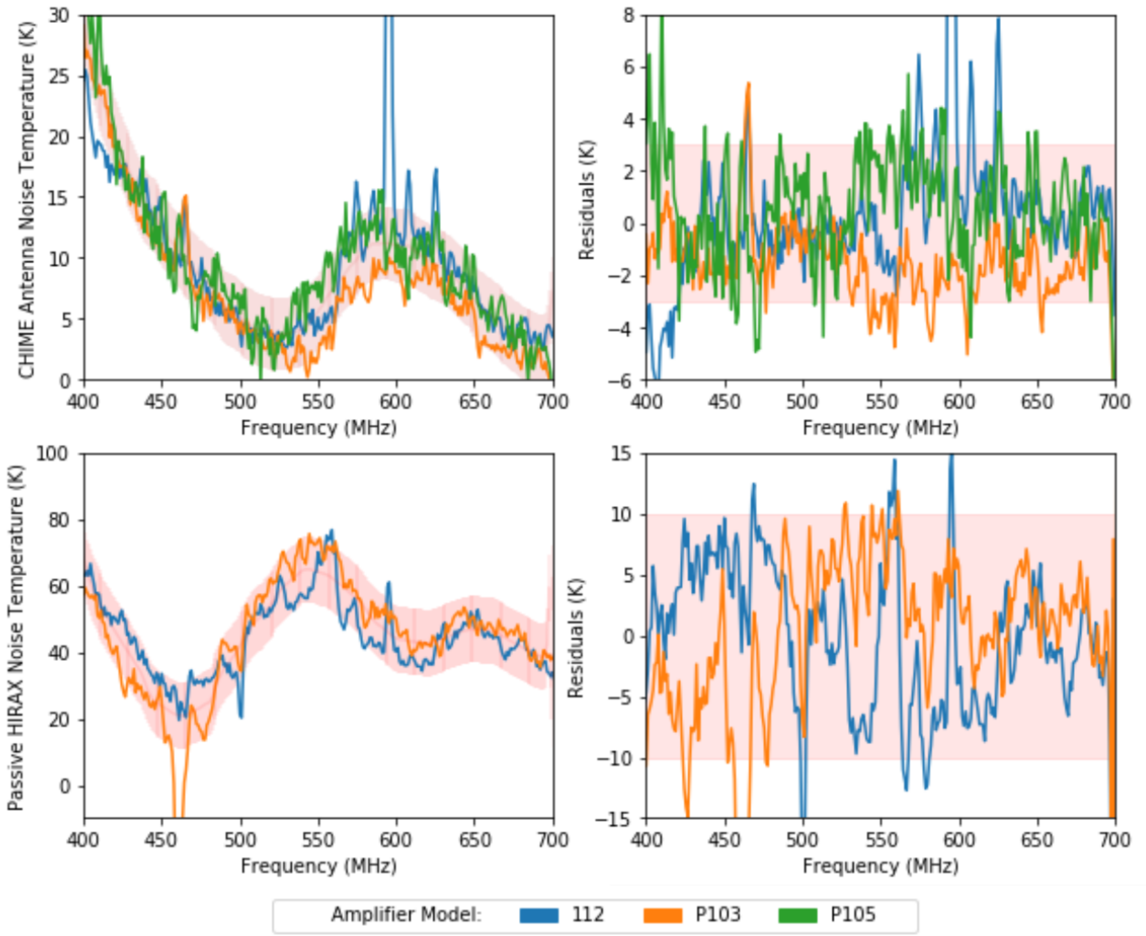

Figure 4.16: Antenna noise temperature (Tantenna) and residuals for the passive CHIME and HIRAX feeds connected to varying amplifier chains. The first stage of the amplifier chain is indicated in the legend as "amplifier model", and refers to the amplifiers listed in table 4.2. Shaded regions correspond to ±3K for CHIME and ±10K for passive HIRAX.

formance. We infer a CHIME antenna noise temperature, $T_{\text{antenna, inferred}}$, computed as the average of measurements in fig 4.16. We can use $T_{\text{antenna, inferred}}$ in conjunction with the Y-factor measurements of the full CHIME and amplifier chains T_{noise} to solve for the amplifier chain noise temperature ($T_{\text{chain, inferred}}$):

$$T_{\text{chain, inferred}} = (T_{\text{noise}} - T_{\text{antenna, inferred}}) \left(\frac{T_{\text{ref}}}{T_{\text{antenna, inferred}} + T_{\text{ref}}} \right) \tag{4.22}$$

As discussed, the amplifier noise temperatures were previously characterized with a noise figure meter, so we can compare the inferred result ($T_{\text{chain, inferred}}$) with the original amplifier chain measurement (T_{chain}) to determine how tightly we have constrained noise temperature. Main results are shown in figure 4.17 for CHIME and in figure 4.18 for passive HIRAX.

With the CHIME feed measurements, we are able to recover noise temperature to within ±5K of the expected values for all three measured amplifier chains (fig. 4.17). The 5K bound is a comparable error performance to the commercial noise figure meter (2K error), and is the bound we require (recall: we expect a 50K noise temperature from the HIRAX active feed, and desire to measure it within 10%). We also note that two of the chains have noise temperatures in excess of 100K, so the 5K error is a constant offset that does not grow with amplifier temperature.

The passive HIRAX antenna measurements (fig. 4.18) show a greater antenna noise temperature than CHIME, which is expected due to its composition from lossier material. Amplifier recovery measurements show a greater uncertainty than with CHIME, but the noise temperature is still recovered to within 10%, which is the active HIRAX antenna specification (active HIRAX is the target device to measure).

In an attempt to disentangle the passive antenna noise temperature rippling features from the cabling set up, a noise temperature measurement was performed for the 112-amplifier chain terminated at 50Ω (figure 4.19). The thermal noise of the terminator (along with an SMA connector), amplified by the amplifier chain, was measured with the connector at room temperature (290K) and with the connector dunked in a container of liquid nitrogen (77K). The connector loss inferred from the measurement is roughly 0.15dB across the band, which was independently verified with the VNA ‘cable loss’ setting. This value is 0.1dB higher than published data sheet values for a standard connector[9], and we expect with the inclusion of cable loss and terminator noise the value would agree. The CHIME loss is expected to be in

---

[9] https://www.pasternack.com/images/ProductPDF/PE9505.pdf

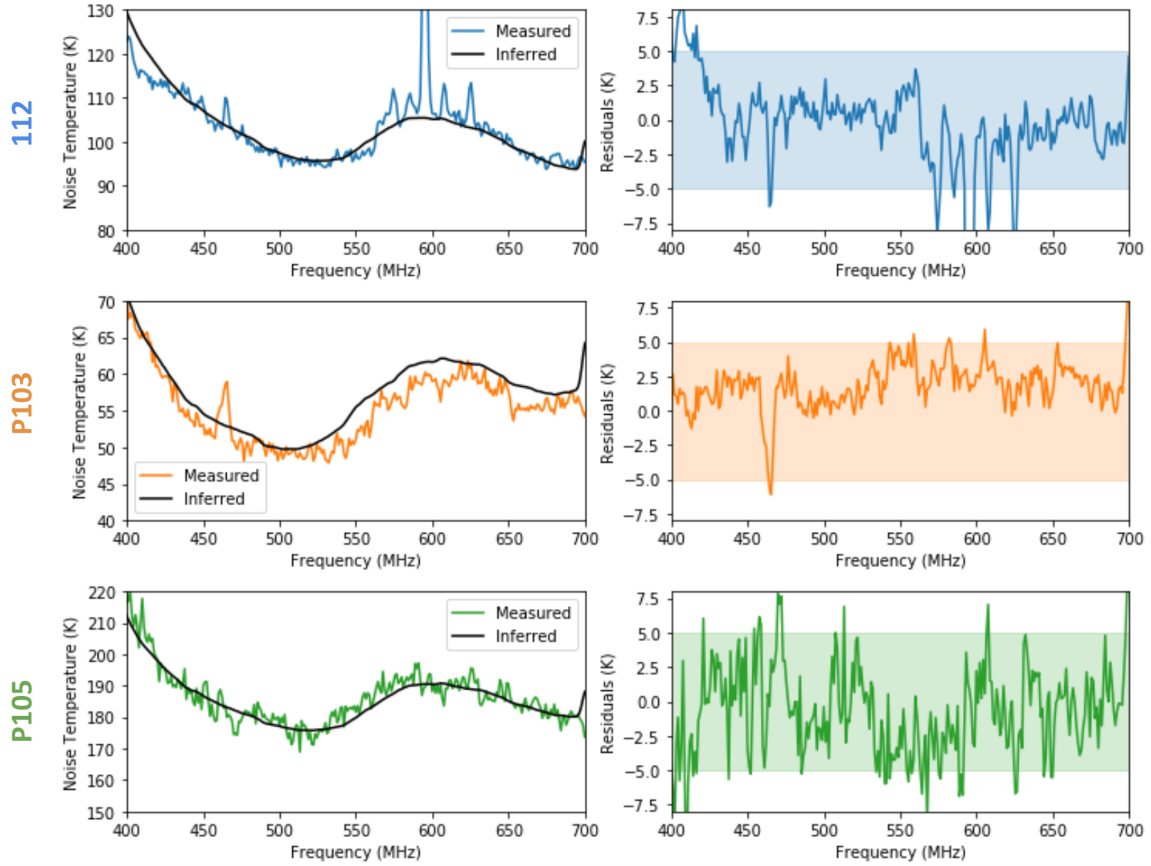

Figure 4.17: Y-factor measurements and expected results for CHIME antenna plus several amplifier chains. The expected results were determined from a noise temperature evaluation of the CHIME feed alone (shown in figure 4.16) in combination with noise figure measurements of the amplifier chain components (table 4.2). The left hand plots show the noise temperature comparisons for each of the three amplifier chains (first-stages 112, P103, P105), and the right hand plots show the difference between measured and expected along with a ±5K shaded region. The 5K bound is a comparable error performance to the commercial noise figure meter (2K error), and is the bound we require for active feed measurements.

this same regime, as it is made of a low-loss material, and we see it averages to the same general value but with higher amplitude features. This result indicates that the features are due to either the antenna or the antenna/system interaction, so we can eliminate the cabling as a dominant systematic.

That we are able to effectively measure amplifier noise temperature to within 5K (utilizing a passive feed noise temperature determined from averaged measurements) is important verification for the Y-factor system. The precision of the measurement

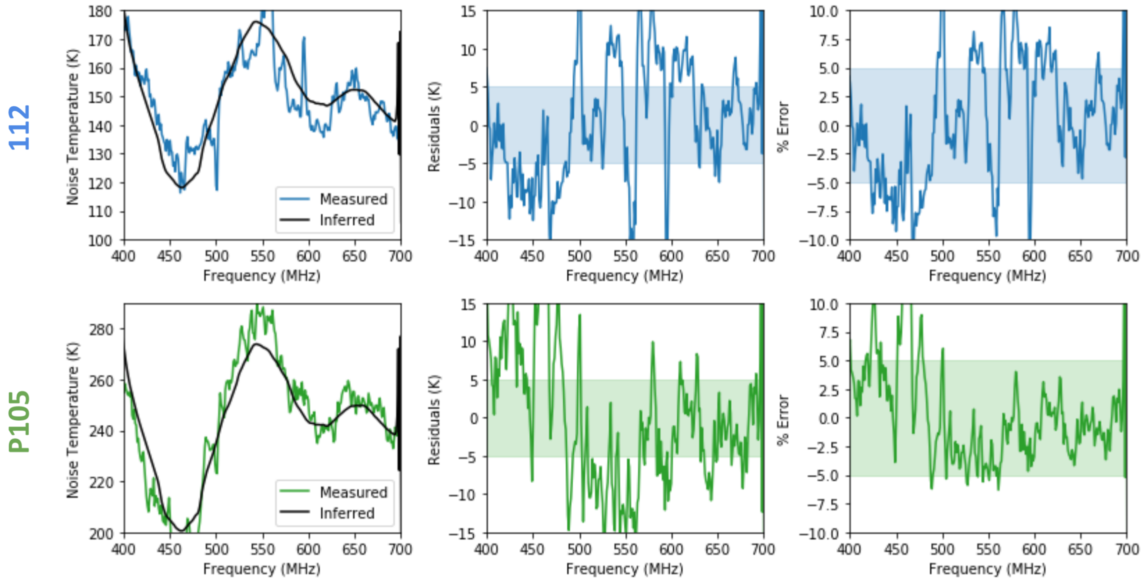

Figure 4.18: Y-factor measurements and expected results for passive HIRAX antenna plus several amplifier chains. The expected results were determined from a noise temperature evaluation of the HIRAX feed alone (shown in figure 4.16) in combination with noise figure measurements of the amplifier chain components (table 4.2). The left hand plots show the noise temperature comparisons for two different amplifier chains (first-stages 112, P105), the middle-column hand plots show the difference between measured and expected along with a ±5K shaded region, and the right hand plots turn the residuals into a % error. The passive HIRAX feed plus amplifier measurements are less constrained than those made with the CHIME feed, but amplifier results are still broadly recoverable to within ±5%, including fluctuations about a mean that can be averaged down in the final active feeed measurements.

means that at a minimum, this chamber is useful and reliable for comparing noise across antennas to assess repeatability, and can effective perform a relative comparison between different antenna models, which will be key for evaluating future antenna upgrade prototypes.

Overall, the passive verification measurements were able to recover noise temperature to within 10% for both antennas. This gives us confidence in measuring the active HIRAX noise temperature.

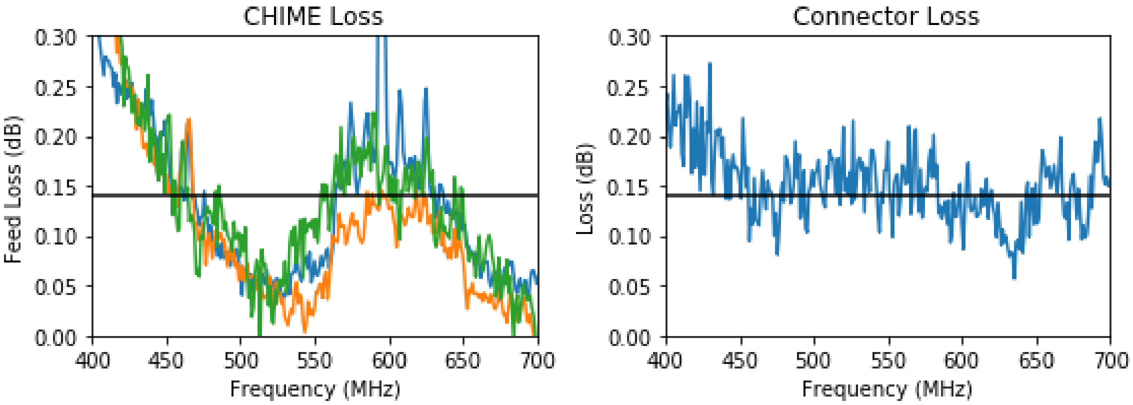

Figure 4.19: The CHIME feed loss and SMA connector loss, compared to investigate whether the frequency dependent features in the CHIME loss profile are a systematic from the cabling. The smooth profile of the connector lost, taken with the same RF chain and cables, indicates that the CHIME loss feature is either real or a result of reflections or some other type of interaction between antenna and cavity. We note that the slope from 400-450MHz is due to late-stage filtering in the amplifier chain that was not accounted for in the analysis.

4.6 HIRAX antenna results

We have taken Y-factor measurement data for four different HIRAX feeds over the course of five separate measurement trials, from December 2019 through December 2021.

Though we established that feed orientation does not impact noise temperature results in a measurable way, we keep the orientation consistent for all measurements. We are also careful to ensure that cables are positioned to reduce strain on the feed SMA connector joints.

4.6.1 Noise temperature

We find the noise temperature of the HIRAX antenna to be frequency dependent, oscillating about ∼50Kelvin (Figures 4.20, 4.21), with a ripple of peak-to-peak height 40K. These measurements were repeated for four different feeds with two working polarizations, as well as with CHIME and dipole feeds.

Figure 4.20 shows the noise temperature of HIRAX as computed during five different measurement sets. Each measurement set consists of multiple days. The individual curves on each plot represents a noise temperature for one polarization on one antenna as measured during a < 20 minute window, which may include multiple hot and cold measurements averaged together (for later data sets). These curves are plotted within a gray band that represents the average of the curves ± 10K. Early measurements, in December 2019 and August 2020, involved recording one hot spectra and one cold spectra per antenna and solving for noise temperature. Data taking methods were improved since to include dozens of files per reported measurement averaged together, and thus these later measurement sets are more rigorously examined.

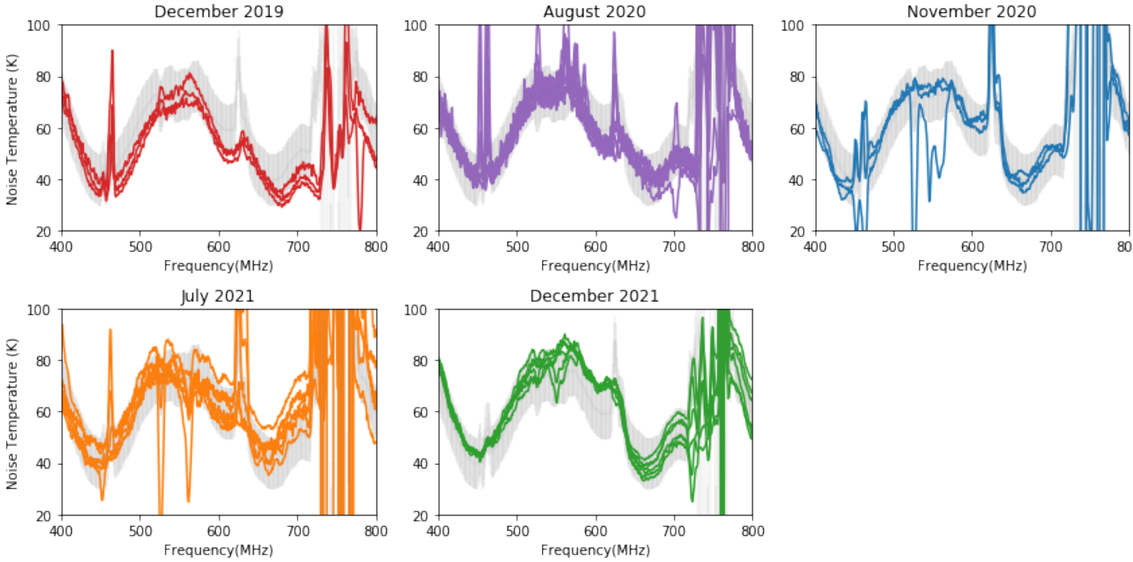

Figure 4.20: The HIRAX antenna noise temperature, as measured during five different measurement sets consisting of multiple days. The individual curves on each plot represents a noise temperature for one polarization on one antenna as measured during a < 20 minute window, which may include multiple hot and cold measurements averaged together. Measurements are consistent to within ±10K, which is shown as a gray band on each plot.

Looking more closely at the latter three measurement sets, we compute a noise temperature per measurement set and associated uncertainty. The results between measurement sets are broadly consistent to within 5K and within 10% of the measurement mean (per frequency), excluding RFI-prone areas in the band (Fig. 4.21). Excluding these RFI areas, we also find the measurements broadly consistent within the measurement statistical uncertainty, which is computed using standard error propagation:

$$\delta T_{\text{noise}}(x, y) = \sqrt{(\frac{\partial T_{\text{noise}}}{\partial P_h}\delta P_h)^2 + (\frac{\partial T_{\text{noise}}}{\partial P_c}\delta P_c)^2}, \tag{4.23}$$

with P_h, P_c as the "hot" and "cold" power measurements and $\delta P_h, \delta P_c$ as the standard deviations of the respective power measurements taken within a 20 minute measurement time window. For all data sets, there are dozens of hot/cold measurements included in the statistics.

4.6.2 Antenna repeatability

We investigate relative measurements between different HIRAX feeds and different measurement days within a measurement set in addition to absolute noise temperature values. We find the HIRAX feed noise temperature results are repeatable across all three feeds and three separate measurement days to within ± 5% (shown in Figure 4.22).

4.6.3 In situ assessment of reflections

The frequency dependent features in both passive and active feed measurements are unexpected, and prompted the additional investigations described in this section. We first consider system asymmetries by repeating Y-factor measurements with different antenna rotation angles. System asymmetries had previously been examined with S_{11} measurements, but we re-verified in noise temperature measurements to be thorough. In measuring the active HIRAX antenna at 15 degree rotational increments (lid turns, cavity remains stationary) from its typical position, we found no perceptible change in noise temperature features.

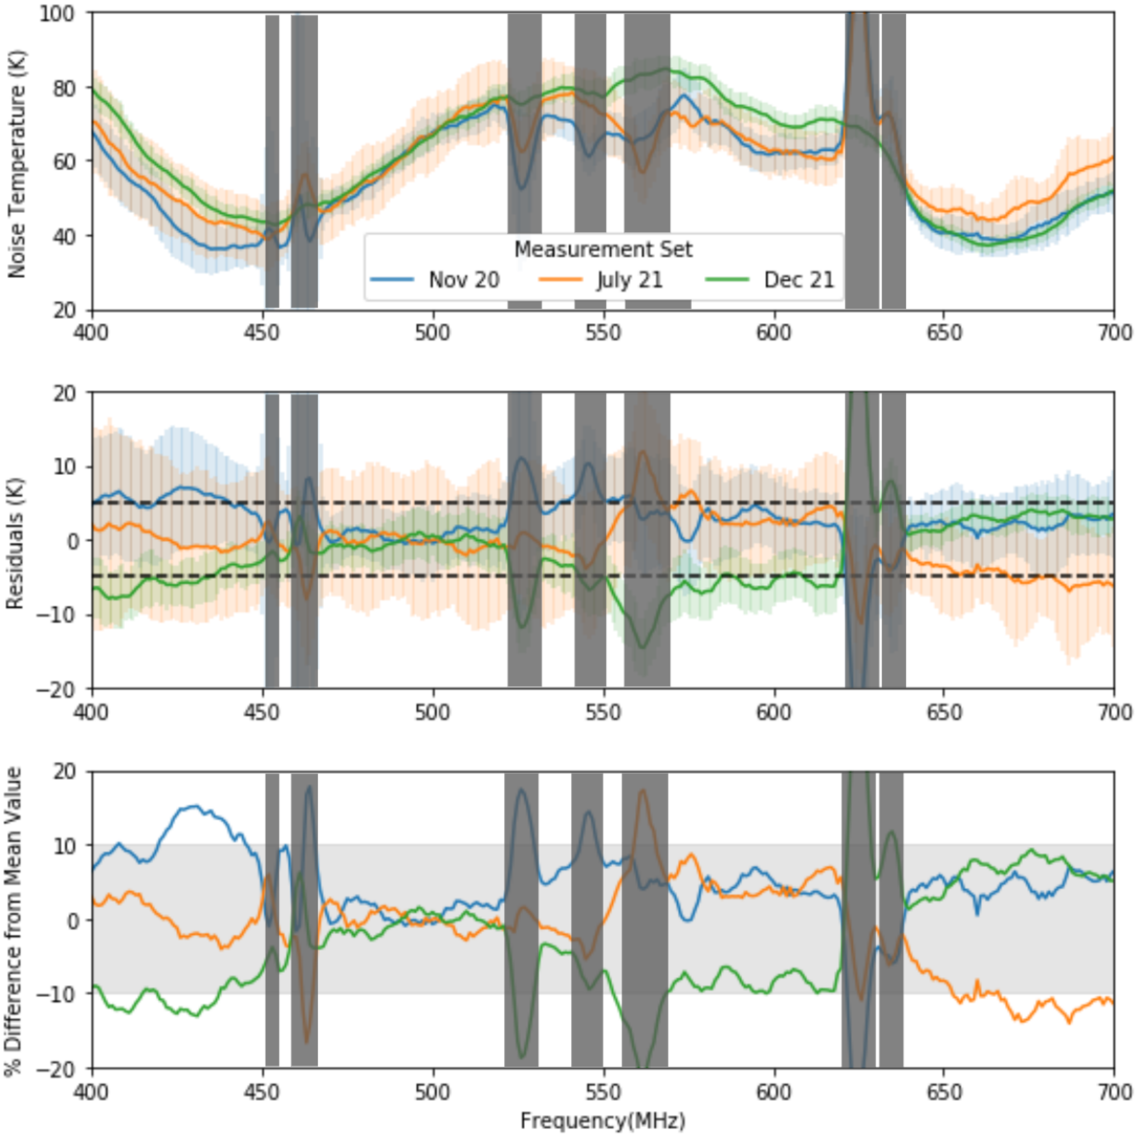

Figure 4.21: The HIRAX antenna noise temperature is broadly consistent to within 5K/10% over three measurement sets each consisting of multiple days, meeting the desired specifications. The shaded bands on each curve represent measurement uncertainty. Vertical gray lines mask RFI features.

We also consider that features may arise from a reflection in the system, which we investigated in two ways: (1) by raising the system lid, thereby increasing the cavity length and distance to nitrogen and (2) by lowering the liquid nitrogen level, thereby changing the distance between the antenna phase center and liquid nitrogen surface.

To elongate the cavity, we had to construct a cavity lid extender. This construction

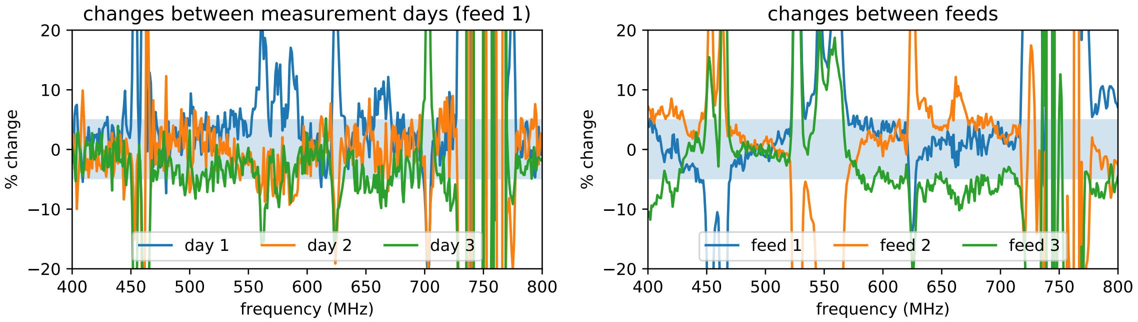

Figure 4.22: Noise temperature consistency for HIRAX feed measurements, shown as a percent change between measurement days (left) and different feeds (right). In both scenarios, excluding RFI, noise temperatures are consistent to within 5% across measurements (light blue band).

consisted of rolling 6inch wide aluminum pieces into the cavity circumference, and welding them together to form a ring. We then attached 24 L-brackets to the ring with pop rivets that mounted to a companion series of L-brackets bolted to the original cylinder. A system photo is in fig. 4.23.

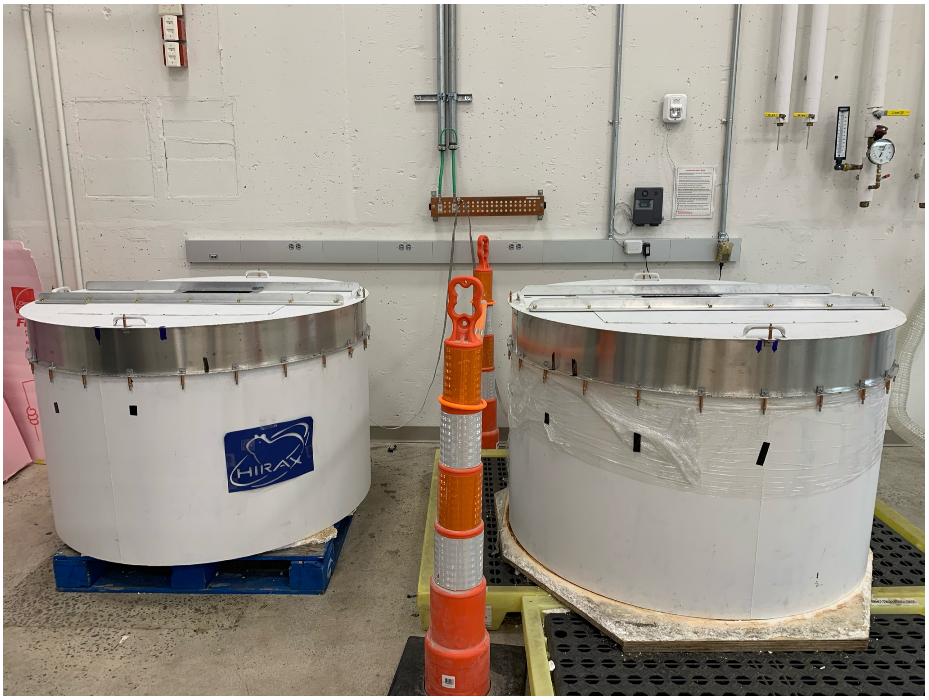

Figure 4.23: Cavity extenders of 6inches in length were constructed for the Y-factor system, in order to evaluate the impact of reflections and provide the capability to measure larger antennas.

Raising the lid has the dual benefit of allowing space for measurements of larger feeds, including a vivaldi antenna in the style prototyped for the CHIME successor CHORD[49]. By evaluating a different feed style, we can determine whether the frequency dependent features in the HIRAX and CHIME noise temperatures are specific to HIRAX-style feeds. Measuring vivaldi, biconical, and dipole antennas do not reveal the HIRAX/CHIME frequency dependence, so we can rule out more general systematic from the cavity. It is worth noting that in raising the lid, we are deviating from the system dimensions that were optimized for in simulations, and a predictable regression of system performance is revealed in S_{11} measurements.

Y-factor measurements of the HIRAX antenna with both a raised lid and a lowered nitrogen level revealed a noise temperature profile that was overall consistent with original cavity measurements in level. The features were, however, shifted and warped in frequency (fig. 4.24), and this work is remains under investigation. Given the consistency of the overall level, however, the Y-factor measurement system has been verified to function as a method of comparison between different antenna models, which will be key as HIRAX continues to iterate on its antenna design.

The raised lid measurements were compared with on-sky Y-factor measurements, and shown to have close qualitative agreement. The on-sky measurements were taken by an active HIRAX feed mounted to a 3m dish belonging to the Deep Dish Development array (D3A)[100] at the Dominion Radio Astronomy Observatory in Penticton, British Columbia, Canada. For the Y-factor, hot measurements were taken with a known source transiting overhead, and cold measurements were taken when no known sources were present in the beam. A comparison between the two measurements is shown in the bottom plot of figure 4.24, revealing that features track in frequency location and amplitude across the full band. We recognize that both systems share a 4ft characteristic length scale between antenna phase center and metalic reflective surface (cavity bottom and dish parabola), further suggesting the presence of a reflec-

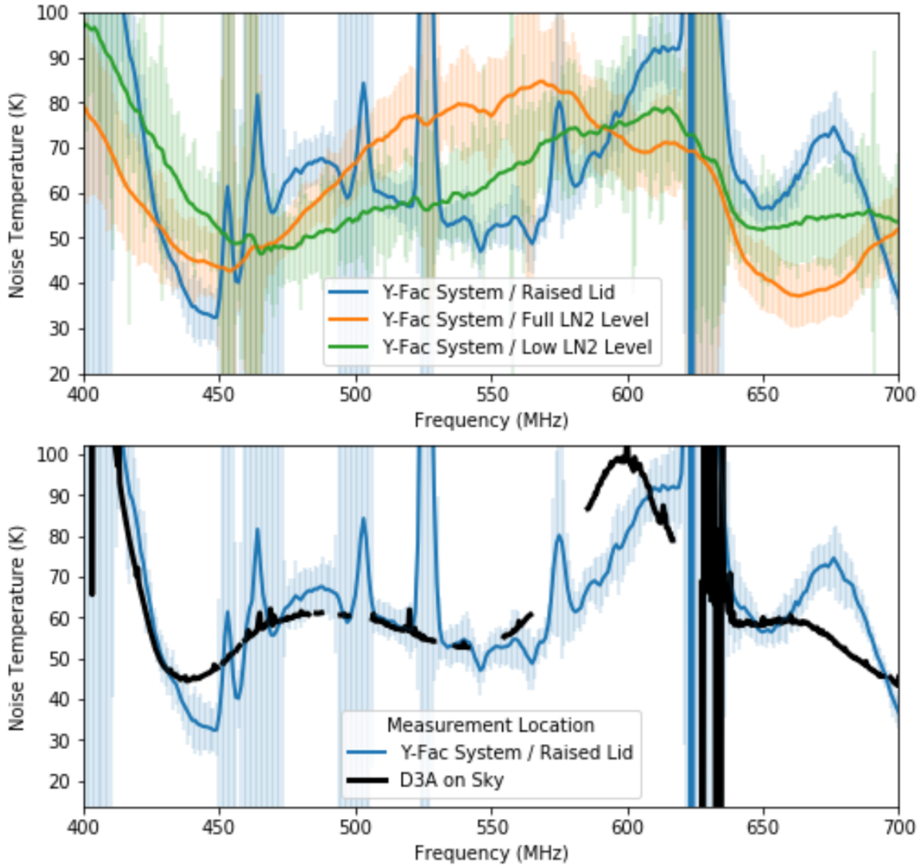

Figure 4.24: (Top plot) Y-factor measurements of the active HIRAX antenna with both a raised lid and a lowered nitrogen level as compared with the full nitrogen level measurement. Shaded regions indicate measurement uncertainty. These measurements reveal a noise temperature profile that is overall consistent with original cavity measurements in level but differs in feature location. (Bottom plot) Y-factor measurements of the active HIRAX antenna (1) in the raised-lid system using room temperature and liquid nitrogen as hot/cold and (2) mounted to a 3m dish, using on-sky measurements of sources in beam and sources out of beam as hot/cold. The frequency-dependent features present in the raised-lid measurement track qualitatively with on-sky Y-factor measurements. For the cases of the raised-lid and on-sky measurements, there exists a 4ft characteristic length scale between antenna phase center and reflective surface (cavity floor and dish respectively), which further suggesting the presence of a reflection is modifying the measurement. This result is, however, quite encouraging, as it indicates that we have built a test bed with comparable performance to on-sky measurements, which will save considerable time and effort when evaluating future antenna models before deployment.

tion is modifying the measurement. The mounting evidence for reflections will focus future investigations. In the meantime, we remain encouraged by the close qualitative alignment between on-sky and lab-based measurements. Having a testbed with comparable performance to on-sky measurements will save considerable time and effort when evaluating future antenna models before deployment.

A final interesting probe of reflections involved measuring the CHIME feed and commercial amplifier noise temperature at two different nitrogen levels (one with the nitrogen full, the other with it 5inches down from full). The results are shown in figure 4.25. The overall noise temperature level remains the same between the two measurement sets, but a shift of 40MHz in frequency occurs. This is notable because it is potentially indicative of changing resonances (i.e. from reflections). The constant overall level is encouraging as it indicates the absorber cones (our black body) remain 77K even as the nitrogen boils off: if they warmed, we would see the overall level raise between the two measurements. We confirm that the expected reflection frequency from an under 10inch cavity, as would exist between antenna phase center and nitrogen level is does not correspond to characteristic features in the curve:

$$\nu = \frac{c}{\lambda} \sim 1GHz. \tag{4.24}$$

This indicates that reflection should present on the 1GHz scale, which exceeds the measurement range. We are still considering the source of this feature.

4.7 Summary and future work

This chapter details a system designed and optimized to measure the noise temperature of the HIRAX feeds. In addition, we report results from verification measurements performed to assess systematic and statistical errors. Because HIRAX uses an embedded amplifier in the antenna, a system like the one we have developed, which can inject an optical signal, is the only way to measure noise temperature in the lab. This measurement will be critical to verifying that the HIRAX feed design meets the noise specifications required to detect the faint cosmological signal of interest.

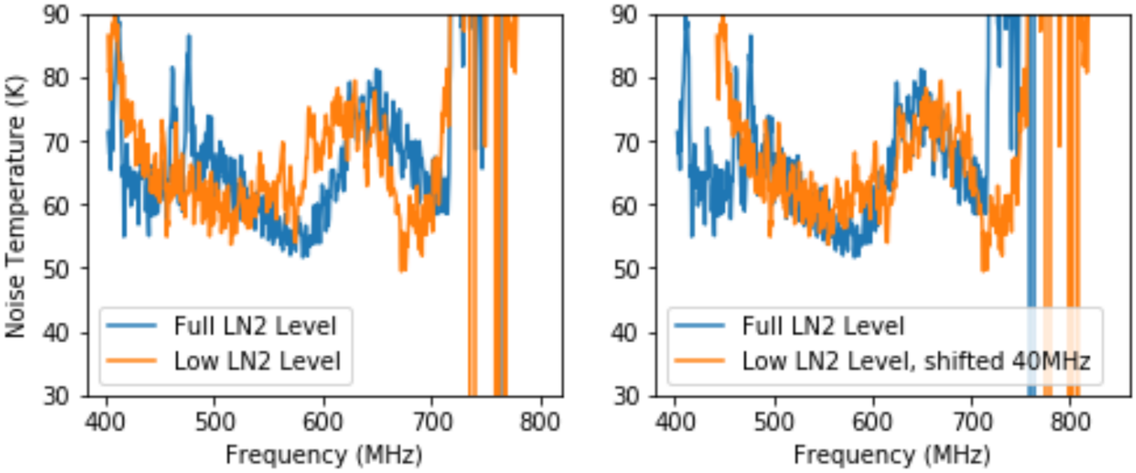

Figure 4.25: Noise temperature of the CHIME antenna plus commercial amplifier chain (first stage ZX60-P103-LN+) at two different nitrogen fill levels: full (18 inch nitrogen depth) and low (13 inch nitrogen depth). The overall noise temperature level remains the same between the two measurement sets, indicating that the black body remains at 77K even when not saturated with nitrogen. A shift of 40MHz in frequency occurs between the measurements, which is potentially indicative of changing resonances.

This measurement system consists of two identical cavities held at different temperatures to allow for a 'Y-factor' measurement. It was designed in the software CST, and constructed over several months in 2019. The construction process included building a cryogenic cavity able to safely and effectively contain over 550 L of liquid nitrogen. The verification procedure utilized both passive and active HIRAX antennas to measure return loss and spectra of the two cavities and quantify systematic and statistical errors. As a secondary verification, Y-factor measurements were performed with several previously-characterized passive feed and amplifier chains, to assess how accurately the test bed recovers a known quantity. Results indicate that this system is within tolerances, with the main sources of error able to be removed or averaged down (Table 4.1).

We perform the first measurements of the HIRAX antenna noise temperature, finding it broadly consistent with the 50K expected value. The noise temperature profile is frequency dependent, oscillating about ∼50K (Figures 4.20, 4.21), with a ripple of peak-to-peak height 40K. These measurements were repeated for four different feeds with two working polarizations, and are shown to be consistent across antennas and measurement days.

There will be limited additional analysis and an additional measurement set required to fully characterize the system, with a focus on understanding reflections and the antenna backlobe.

- Assessing Reflections: Several of the verification measurements described in this chapter modulated nitrogen depth or cavity height, which shifted features in the noise temperature profile. Signal processing techniques may be useful for gauging reflections and remove them from the result. For instance, we could perform a Fourier transform on the noise temperature, remove high-amplitude features suggestive of reflections from the time-domain result, and then Fourier transform back to frequency space. This analysis requires special filtering algorithms, RFI removal, and other techniques, and is an important future extension of the work described.

- Understanding Backlobe: Any backlobe in this system will be pointing towards the laboratory ceiling, a $\sim$300 K source. As a result, additional power is added during the cold measurement equivalent to 300 K times the ratio of the backlobe to the total beam. Backlobe effects are a significant contributor to noise temperature ($>$5 K) if the backlobe makes up more than 1.5% of the total beam, which sets a requirement that the backlobe remain on average below -17 dB relative to the peak. Simulations have indicated that the HIRAX feed may be in this regime, such that we will need to account for the additional power in the cold temperature computation. We do note, however, that the presence of a back lobe will bias the noise temperature high, so we can already place an upper bound on the noise.

Once the final verification work is completed and systematics are fully characterized, we expect to have a fully operational Y-factor measurement system with sensitivity approaching that of a commercial noise figure meter. This system will be used to measure all 256 feeds used for the initial HIRAX deployment, as a spot check on production quality and consistency. Though the work in this chapter has demonstrated that the current feed model meets HIRAX noise specifications, additional work from this thesis indicates non-idealities in other aspects of antenna performance, such as high levels of cross-polarization and beam asymmetry (described in Ch. 3, 6). Thus, there is motivation to design a new antenna for the full HIRAX array, and the Y-factor system will be a critical part of characterizing all prototypes.

Chapter 5

Design of a drone calibration system

Chapters 5 and 6 focus on the design and implementation of a drone calibration system for 21cm beam mapping, with a particular focus on HIRAX. This chapter describes the drone system design, which underwent many iterations from 2018 through the writing of this thesis in 2022, with several upgrades still in progress. It also touches on drone RFI characterization, development of a Yale-based drone calibration testbed, and debugging a critical flight mishap. Characterization of the drone transmitter system beam is described in Chapter 3. For a review of telescope and antenna beams, we refer the reader to Chapter 3.

5.1 Introduction to drone beam calibration

For a telescope to make sensitive cosmological measurements, its response to incoming radiation as a function of angle on the sky must be understood. This response is known as a "beam pattern", and is described in detail in Chapter 3. For experiments like HIRAX, the beam must be understood to 1 part in 10000 [81]. Beam mapping is difficult, however, as the HIRAX dishes are stationary, and cannot scan celestial sources to fully calibrate the beam pattern in 2D, as is standard procedure for calibrating a steerable telescope. In lieu of a movable dish, we can create a mobile source by transmitting broadband noise from a drone. Drone calibration has been demonstrated on some stationary radio dishes, but the existing literature is limited to a subset of techniques, not all of which apply to HIRAX [101–103]. At the most basic level, drone calibration involves flying a drone in some pre-determined grid pattern over the instrument, making amplitude measurements for each polarization. Afterwards, drone pointing and data sets are used to analyze and generate beam maps, which factor into the cosmological analysis. An illustration of drone beam mapping is found in figure 5.1, and an illustration of a sample flight path is found in figure 5.2.

This thesis work focuses on far-field beam mapping, in which the drone flies in the telescope far field. The telescope far-field is the distance at which its radiation can be approximated as planar. It can be found mathematically with the following expression:

$$\text{Far Field} = \frac{2D^2}{\lambda} \tag{5.1}$$

where D is the telescope diameter and λ is the operating frequency [90]. All astrophysical measurements are done in the far-field, and thus analysis requires a far-field beam pattern. It is, unsurprisingly, most convenient to obtain the far-field beam pattern from far-field beam measurements—not only does this get a direct measurement of the quantity of interest; there is also a 1:1 correspondence at a given coordinate between the amplitude measurement and beam response. This is in contrast to near-field measurements, where amplitude and phase measurements at multiple positions are needed to determine the beam response at a single far-field coordinate. Additionally, measurements in the near-field necessitate densely packed grids, of minimum spacing ($\Delta x = \Delta x = \lambda/2 \approx 20$cm for HIRAX[90]), and a transformation that incorporates phase in order to make a beam map.

The HIRAX single-element far-field ranges from 96m to 192m, corresponding to

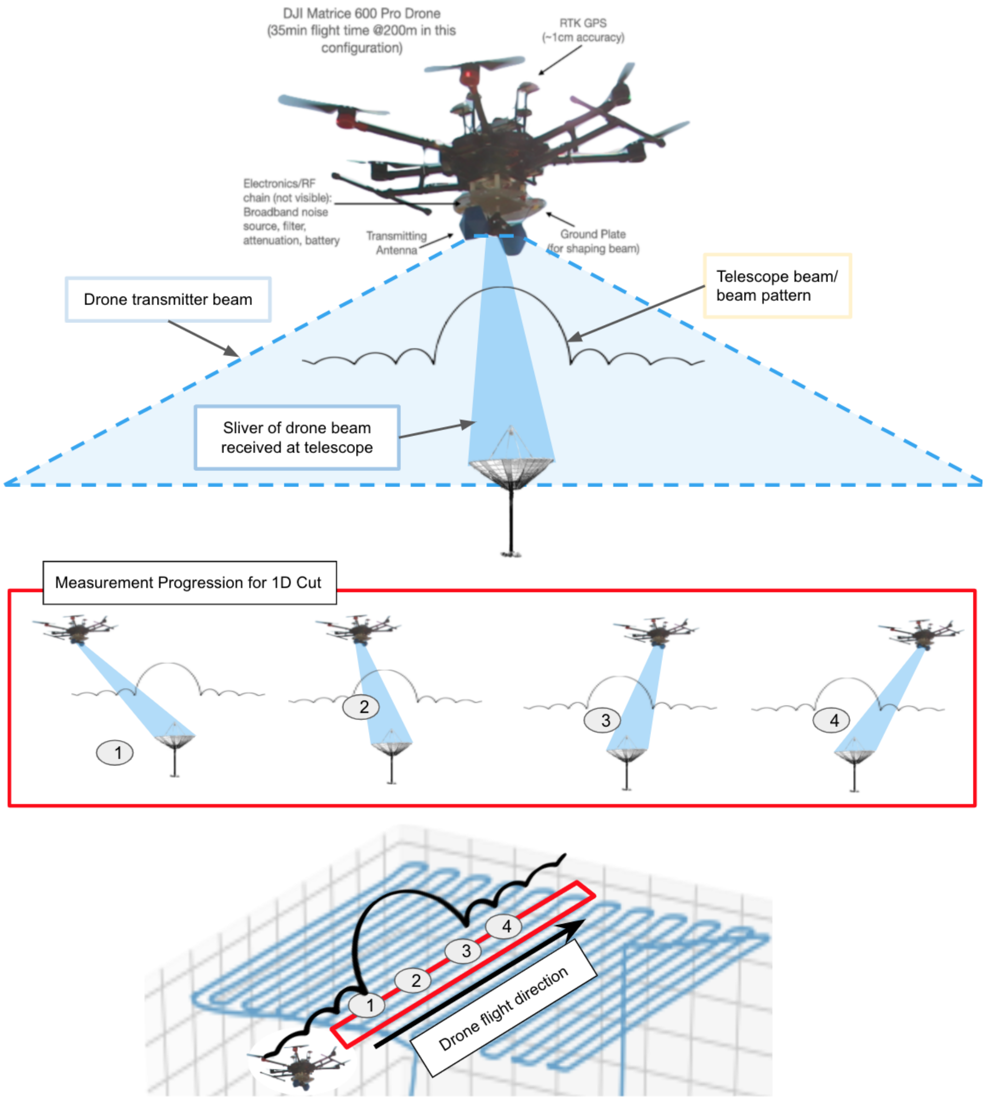
DJI Matrice 600 Pro Drone (35min flight time @200m in this configuration)
RTK GPS (~1cm accuracy)
Electronics/RF chain (not visible): Broadband noise source, filter, attenuation, battery
Transmitting Antenna
Ground Plate (for shaping beam)
Drone transmitter beam
Telescope beam/ beam pattern
Sliver of drone beam received at telescope
Measurement Progression for 1D Cut
1
2
3
4
Drone flight direction

Figure 5.1: Drone calibration of a stationary telescope, illustrating measurements during a 1D slice through the telescope beam. The Yale drone system (labeled photograph shown) transmits calibrated noise through a biconical antenna, with a beam shown in blue (outlined in dashed line) and the portion of the beam received at the telescope shown in darker blue (narrower, no dashed line). The telescope under test has its own beam pattern, labelled on the diagram, that impacts how sensitively the drone transmission is received. With each pass that the drone flies over the telescope, its transmission will be received at many angles of the beam, and in turn, a beam pattern can be extracted from measurements. This concept is illustrated for a single pass over the telescope, which gives a 1D beam cut. The boxed middle panel shows the progression of measurements for 5 positions during a flight, and the bottom figure shows what these locations would look like in a flight path. A full 2D beam is built from 20+ of these 1D cuts. This picture also highlights why a broad and smooth transmitter beam is important.

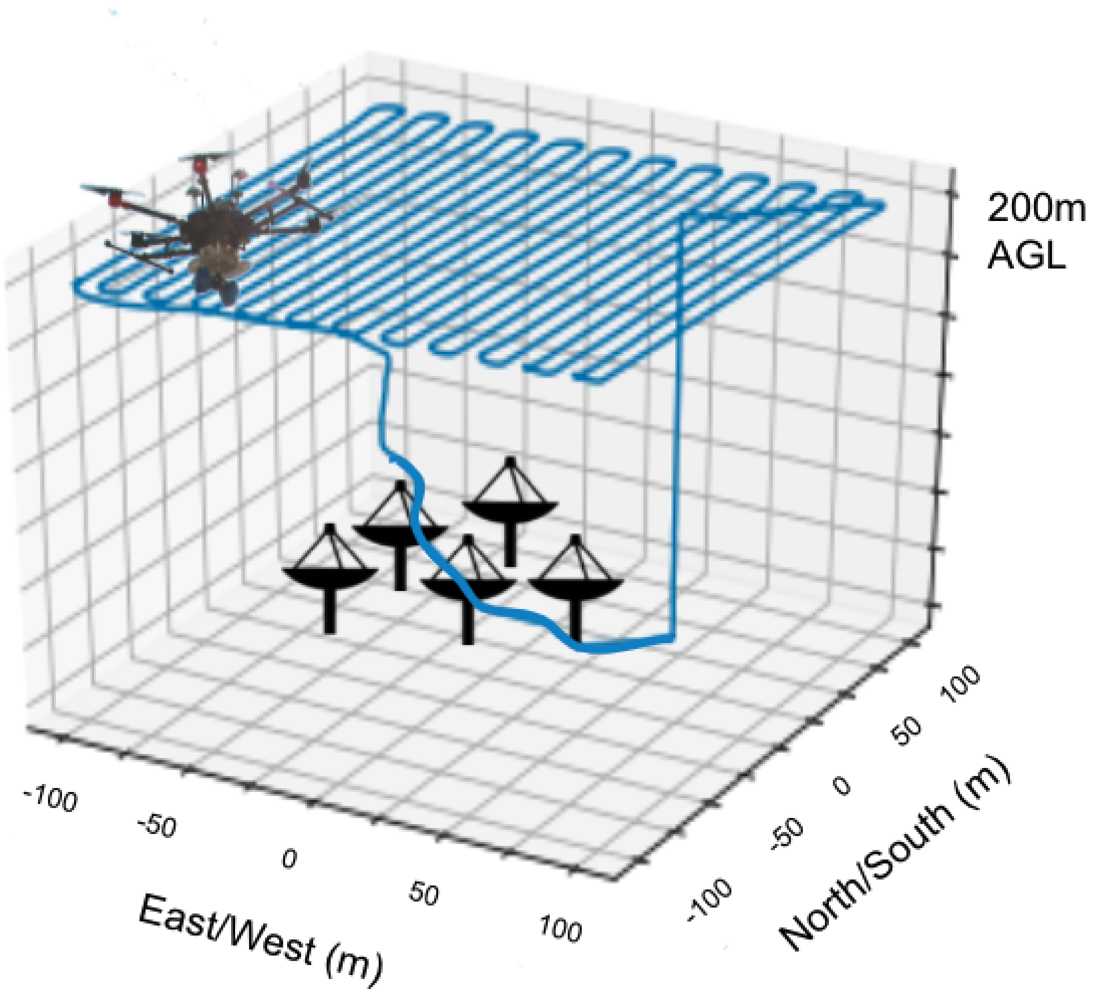

Figure 5.2: Cartoon sketch of 1/4 of a drone map-making flight over a telescope array. Here, the drone flies up to the single-dish far-field (200m) and then traces out a 200×200m grid in North/South passes before descending. The grid is centered over the dish under test. In practice, this style of flight takes ~25 minutes and drains the drone batteries to 25% capacity. To complete a full single-polarization map, the drone will need three additional flights. It will: (1) repeat this style of flight but with East/West passes, (2) fly a denser 60×60m North/South grid, (3)fly a denser 60×60m East/West grid.

the 400-800MHz frequency range observed with a 6m dish. This altitude is easily reachable with amateur drones. Investigating the beam formed from multiple array elements, or from larger telescopes such as cylinders, may require near-field mapping, which will not be covered in this thesis.

Drone beam mapping is a powerful tool for relative calibration of stationary dishes (absolute calibration will be done with sky measurements). The technique can also be effectively expanded to non-stationary dishes, and is particularly useful for generating polarized maps which is not possible with standard calibration from celestial source transits. The noise source used for drone beam mapping is transmitted by a linearly polarized antenna, which yields polarized beam maps in both co- and cross-polarization (as defined in Ch. 3). The ability to generate polarized beam maps gives drone beam mapping an advantage over related beam-mapping techniques, and will be important for extracting sensitive cosmological signals. In 21cm, we expect polarized foregrounds that will need to be removed in analysis, and anticipate these foregrounds will have spectral structure that further complicates their removal, necessitating a well-calibrated instrument.

5.2 Drone instrument constraints

In order to make successful BAO measurements, we need to understand the telescope beam full-width half max to $< 0.1\%$ [81]. This telescope beam constraint imposes instrument constraints on the drone calibrator system that we use to guide the system design.

We constrain drone position in six parameters, all of which can introduce uncertainty into the beam mapping measurements:

- position: x, y, z (unit = meters)
- angle: θ_1, θ_2, ϕ (unit = degrees)

Here, θ_1, θ_2 are the drone tip/tilt angles (called roll and pitch by pilots), and ϕ is the angular rotation/polarization angle (called yaw). Tolerances on these parameters are expressed in table 5.1 with methodology detailed in this section. The constraints analysis made certain approximations when inputting errors as a convenience during

Parameter	Tolerance
x, y, z	± 1 meter
roll, pitch, yaw	± 1 degree

Table 5.1: Uncertainty in drone instrument parameters.

early analysis—we thus expect the results shown in this section to be an approximation of the true result, giving a maximally stringent picture of the tolerances we must comply with.

5.2.1 Methodology for finding positional constraints

Background

Recall from Chapter 3 that Friis equation tells how power transmitted from one antenna (P_t) is received by another (P_r),

$$\frac{P_r}{P_t} = D_t D_r e_t e_r \left(\frac{\lambda}{4\pi d}\right)^2 |\hat{p_t} \cdot \hat{p_r}| \tag{5.2}$$

where: D_r, D_t are antenna directivities, e_r, e_t are antenna efficiencies, λ is the measurement wavelength, d is the distance between the antenna phase centers, and $\hat{p_t}, \hat{p_r}$ are polarization unit vectors. Directivity is defined as the power density radiated by an antenna in the direction of strongest emission (accompanying diagram in figure 3.3).

Using the Friis equation can tell us directly how the power transmitted by the drone antenna should be measured at the receiving telescope, and in turn can help us back out the original telescope beam from our measurements. Directivity is traditionally defined in 1D, though for the purposes of understanding antenna transmission in 2D, we expand this picture to all beam angles, with $D(\theta_1, \theta_2)$ as the antenna beam pattern.

In the context of drone beam measurements, we will consider λ effects because we want to make sure we can measure the full HIRAX frequency band within the correlator dynamic range. Over the HIRAX band of 400-800MHz, λ^2 provides a factor of 4 difference in power. While understanding this power difference is important for avoiding compression issues, we do not need to consider it for full calibration purposes, as goal of our drone measurements is to make a normalized beam map and measure beam full-width-half-max and the relative side lobe features. Absolute gain calibration will happen in a separate measurement. So, for the subsequent analysis, we can choose to consider only parameters that are functions of x,y,z,angles, and roll independent parameters such as wavelength and efficiency into a constant C to worry about in later calibration.

$$\frac{P_r}{P_t} = C \cdot \frac{D_t D_r}{d^2} |\hat{p_t} \cdot \hat{p_r}| \tag{5.3}$$

In the case of our drone measurements, the ingredients for Friis are as follows:

$$D_r = T(x, y, z, \sigma_{\theta_x}, \sigma_{\theta_y}) = \text{telescope beam pattern}$$

$$D_t = D(\theta_1, \theta_2, \sigma_{\theta_1}, \sigma_{\theta_2}) = \text{drone beam pattern}$$

$$|\hat{p_t} \cdot \hat{p_r}| = P(\phi, \phi_0) = \text{polarization angle contribution} = \cos^2(\phi)$$

$$d = d(x, y, z) = \text{radial distance contribution} = 1/\sqrt{x^2 + y^2 + z^2}$$

Here, x,y,z are drone position variables relative to the telescope phase center, and θ_1, θ_2, ϕ are drone tipping, tilting and polarization angles (known as roll, pitch, and yaw), and the σs are related to the telescope and drone beam widths in two dimensions, as we will approximate them later as 2D gaussians. The coordinate system is shown in figure 5.3.

These ingredients can be propagated through the Friis formula to illustrate what power a telescope of radiation pattern T can expect to measure when a transmitter

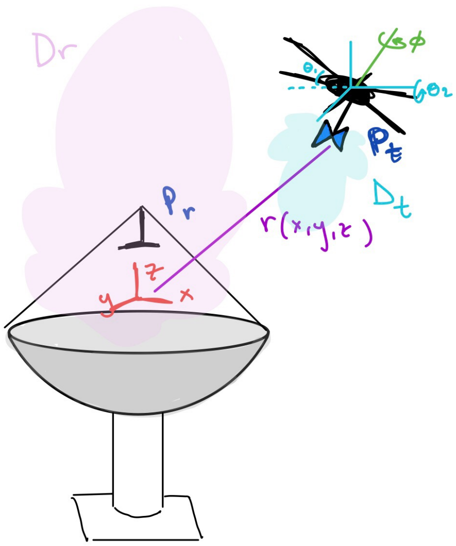

Figure 5.3: Illustration of a drone beam measurement using the Friis equation expanded to 2D. This diagram includes the coordinate system used for drone analysis.

of radiation pattern D is flying on a drone at any given spatial position (x,y,z).

$$P_r = C \cdot \frac{T(x, y, z, \sigma_{\theta_x}, \sigma_{\theta_y}) D(\theta_1, \theta_2, \sigma_{\theta_1}, \sigma_{\theta_2})}{d(x, y, z)^2} \cos\phi \cdot P_t \tag{5.4}$$

P_r is the data we will take with the HIRAX telescopes during drone flights. We are able to measure P_t with a spectrum analyzer pre-flight, should have an understanding of D from prior range measurements, and can find $\cos\phi$ and $d(x, y, z)$ from drone sensor data. To recover the telescope beam pattern T, we simply modify this equation by rearranging variables.

$$T(x, y, z, \sigma_{\theta_x}, \sigma_{\theta_y}) = \frac{1}{C} \cdot \frac{d(x, y, z)^2}{D(\theta_1, \theta_2, \sigma_{\theta_1}, \sigma_{\theta_2}) \cos\phi} \cdot \frac{P_r}{P_t} \tag{5.5}$$

Note that this analysis will require knowledge of drone position in (x,y,z), as well as its tipping and tilting angles $(\theta_1, \theta_2, \phi)$, and an understanding of drone beam width $(\sigma_{\theta_1}, \sigma_{\theta_2})$.

The modifications in analysis will be subject to uncertainties in x, y, z, θ_1, θ_2, ϕ from drone flight data, which are important to characterize.

Procedure

The structure of the drone instrument constraints analysis is as follows. We begin with an expression for the power measured at the telescope (based on eq 5.4). This expression will include the true telescope beam pattern T_{true},

$$P_r = C \cdot \frac{T_{\text{true}}(x, y, z, \sigma_{\theta_x}, \sigma_{\theta_y}) D(\theta_1, \theta_2, \sigma_{\theta_1}, \sigma_{\theta_2})}{d(x, y, z)^2} \cos\phi \cdot P_t, \tag{5.6}$$

with T_{true} approximated as a gaussian for the purposes of this analysis. The power measured at the telescope P_r is dependent on drone spatial parameters, and subject to correlator measurement uncertainties that we assume to be systematic. Uncertainties on drone GPS/sensor data do not impact this measured quantity.

We now propagate the above equation into the expression for the telescope beam pattern that is recovered after analysis $T_{\text{recovered}}$(eq 5.5). This is the desired output of this whole drone project. It involves the measured power at the telescope P_r modified by the drone position and beam.

$$\begin{aligned} T_{\text{recovered}} &= \frac{1}{C} \cdot \frac{d(x - \delta x, y - \delta y, z - \delta z)^2}{D(\theta_1 - \delta\theta_1, \theta_2 - \delta\theta_2, \sigma_{\theta_1}, \sigma_{\theta_2}) \cos\phi - \delta\phi} \cdot \frac{P_r}{P_t} \\ &= \left[\frac{d(x - \delta x, y - \delta y, z - \delta z)^2}{D(\theta_1 - \delta\theta_1, \theta_2 - \delta\theta_2, \sigma_{\theta_1}, \sigma_{\theta_2}) \cos\phi - \delta\phi} \right] \times \\ & \left[\frac{T_{\text{true}}(x, y, z, \sigma_{\theta_x}, \sigma_{\theta_y}) D(\theta_1, \theta_2, \sigma_{\theta_1}, \sigma_{\theta_2}) \cos\phi}{d(x, y, z)^2} \right] \end{aligned} \tag{5.7}$$

In this equation, the first bracketed term is constructed from drone beam information as well as drone flight/GPS data, and is therefore subject to uncertainties. The second term is a model for what is measured by the telescope, which we assume is only subject to systematic errors that are independent of drone position (so constant throughout the flight) and can therefore be calibrated out. As a gut check, if all uncertainties = 0, $T_{\text{recovered}} = T_{\text{true}}$.

We now have an expression for the measured telescope beam pattern as a function of drone positional/beam data. We can now add random noise to each uncertainty parameter, and fit a gaussian to the resulting telescope beam $T_{\text{recovered}}$. We then check to see whether $T_{\text{recovered}}$ is within 0.1% of T_{true}. Our drone GPS/sensor tolerances are when we hit that limit.

To recap, we:

1. Start with an expression for telescope beam pattern T recovered from drone measurements (eq 5.8)
2. This expression should include one term with uncertainty that is a correction using drone flight data, and a second term without uncertainty that approximates the power measured at the telescope.
3. Add random noise to each parameter
4. Fit a gaussian, and monitor how the full-width at half-maximum (FWHM) changes.
5. We need to know FWHM to within 0.1% [81], so tolerances are when we hit that limit.

Note that in this analysis, we just cancel out transmitter power P_t. In reality, we want to check the power to make sure we will not compress the correlator, but do not have to be so precise in the measurement because it should be independent of positional variables and constant through the whole flight (at each frequency).

As mentioned, we model the telescope beam and drone beam as gaussians. Our plug ins to the expressions above are as follows.

$$T(\theta_x, \theta_y, \sigma_{\theta_x}, \sigma_{\theta_y}) = \text{telescope beam pattern} = \exp\left[-\left(\frac{\theta_x^2}{2\sigma_{\theta_x}^2} + \frac{\theta_y^2}{2\sigma_{\theta_y}^2}\right)\right] \tag{5.8}$$

$$D(\theta_1, \theta_2, \sigma_{\theta_1}, \sigma_{\theta_2}) = \text{drone beam pattern} = \exp\left[-\left(\frac{\theta_1^2}{2\sigma_{\theta_1}^2} + \frac{\theta_2^2}{2\sigma_{\theta_2}^2}\right)\right] \tag{5.9}$$

$$P(\phi, \phi_0) = \text{polarization angle contribution} = \cos^2(\phi) \tag{5.10}$$

$$d(x, y, z) = \text{radial distance contribution} = 1/\sqrt{x^2 + y^2 + z^2} \tag{5.11}$$

Where,

- $T(\theta_x, \theta_y, \sigma_{\theta_x}, \sigma_{\theta_y})$ models the telescope beam as a 2D Gaussian of sigmas = σ_x, σ_y, with sigma related to beam FWHM through,

  $$\sigma \approx \text{FWHM}/2.355 \tag{5.12}$$

  We can express θ_x, θ_y in distance coordinates, projecting the beam onto an xy plane at height z, using $\theta_x = \arctan(x/z)$ and $\theta_y = \arctan(y/z)$. T will be written as $T(x, y, z, \sigma_{\theta_x}, \sigma_{\theta_y})$ in later equations.

- $D(\theta_1, \theta_2, \sigma_{\theta_1}, \sigma_{\theta_2})$ models the drone transmitter beam as a 2D Gaussian of sigmas = $\sigma_{\theta_1}, \sigma_{\theta_2}$, and gives its radiation pattern as a function of angular coordinates θ_1, θ_2. As analysis progresses, we will measure the true drone beam and be able to incorporate the exact expression into this pipeline.

- $P(\phi, \phi_0)$ models the transmitter beam polarization as linear with a non-existent cross-polarization component. We plan to fly with $\phi = 0$ for the entire flight.

- $R(x, y, z)$ models how transmitter power falls off radially between drone and telescope

The final expression used in our drone constraints analysis is:

$$T_{\text{recovered}} = \left[\frac{(x-\delta x)^2+(y-\delta y)^2+(z-\delta z)^2}{\exp[-(\frac{(\theta_1-\delta\theta_1)^2}{2\sigma_{\theta_1}^2}+\frac{(\theta_2-\delta\theta_2)^2}{2\sigma_{\theta_2}^2})]\cos{(\phi-\delta\phi)}}\right]\times$$
$$\left[\frac{\exp\left[-\left(\frac{(\arctan(x/z))^2}{2\sigma_x^2}+\frac{\arctan(y/z)^2}{2\sigma_y^2}\right)\right]\exp\left[-\left(\frac{\theta_1^2}{2\sigma_{\theta_1}^2}+\frac{\theta_2^2}{2\sigma_{\theta_2}^2}\right)\right]\cos\phi}{x^2+y^2+z^2}\right] \tag{5.13}$$

Once again, we confirm that with all uncertainties = 0, $T_{\text{recovered}} = T_{\text{true}}$.

This expression is solved in python on a 2D grid with random noise added to each parameter corresponding to parameter uncertainties. Fits are performed 100 times, the percent uncertainty on sigma is calculated for each fit, and all results are averaged to determine whether we meet the 0.1% specification.

5.2.2 Findings for positional constraints

We randomly sample perturbations to the drone positional and angular variables from the following intervals:

- $\delta x, \delta y, \delta z \in [-0.1, 0.1]$ meters
- $\delta\theta_1, \delta\theta_2, \delta\phi \in [-0.8, 0.8]$ degrees

Additionally, we fix the following parameters:

- $\sigma_{\theta_x} = \sigma_{\theta_y} = 2.55$ degrees (corresponding to a 6.0 degrees telescope FWHM)
- $\sigma_{\theta_1} = \sigma_{\theta_2} = 25.5$ degrees (corresponding to a 60 degree transmitter FWHM, which was characterized on antenna ranges as described in Ch. 3)
- angular resolution = 0.5 degrees (look over -30 to 30 degrees in 120 steps)

With these initial inputs, which are achievable with commercial drones, we find an uncertainty of <0.01 percent. This is visualized in figure 5.4.

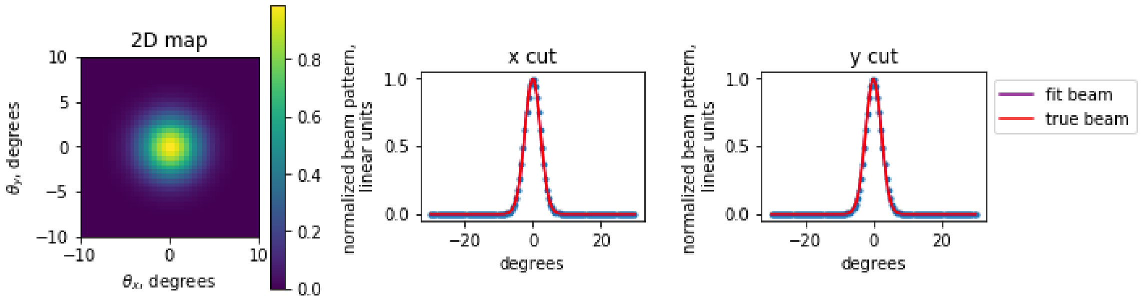

Figure 5.4: Beam visualization and cuts for initial inputs of 0.1m position uncertainty and 0.8deg angle uncertainty, with telescope fwhm =6degs, drone fwhm = 60degs and 0.5 degree angular resolution. The uncertainty for this configuration, which should be attainable with many commercial drone/GPS systems, is more than sufficient.

We can also go through a series of sanity checks to make sure the code is working as expected. These include fitting with 0 uncertainty in position/angle (should recover input parameters, which we do), fitting for large uncertainty (should see smearing in the imshow plot and noisy gaussians in the 1D cuts, which we do), and generating a beam with different FWHMs in x,y (which yields the expected result).

For the expected HIRAX FWHM (6 degrees) and expected drone beam FWHM (60 degrees), we can also generate fits for a variety of different uncertainties and explore the parameter space to figure out the easiest restrictions to comply with, shown in figure 5.5.

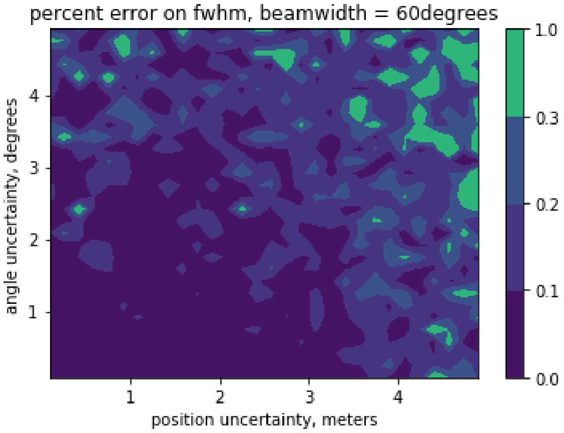

Figure 5.5: Visualization of uncertainties in position/angle that will yield sub-0.1% error in FWHM (we want to be in the darkest blue region).

Impact of drone beam width

If the drone antenna does not remain level during flight, (tipping up by a few degrees) beam width starts to become relevant. For instance, a drone flying level can tolerate a 15 degree beam width and still measure fwhm to 0.1% (assuming 0.1m error in position and 0.8 degree error in angle). If the drone tips up 10 degrees in flight, we will require a 60 degree beam width to maintain that 0.1% error comfortably.

Impact of systematic vs statistical offsets

If the drone flies tipping up at some angle (which we can expect without a gimble, approximated as assuming a constant angular offset in θ_1), the space of acceptable uncertainty parameters narrows. However, assuming a transmitter beam width of 60 degrees and a drone flying at a 15 degree tilted angle produces constraints of 1 degree uncertainty in angle and 1 meter uncertainty in position, which many commercial drones can comply with (figure 5.6).

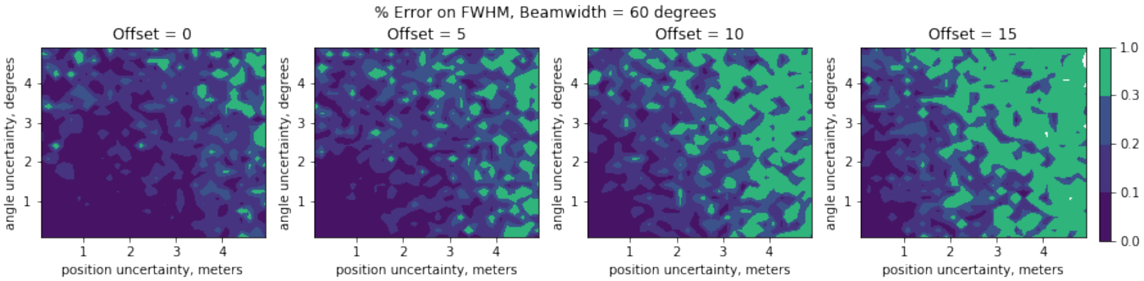

Figure 5.6: Impact of drone tipping angle (akin to flying with a constant angle offset) on acceptable uncertainty parameters. We find that we can tolerate drone tipping up to 15 degrees with angular uncertainty of 1 degree and positional certainty of 1 meter.

Practically, once we have assessed the uncertainties of the drone system, we can use that information to figure out how much angular resolution will be need for satisfactory fits, which will in turn impact flight speed/time. We can also constrain drone beam width requirements, and assess the impact of systematic offsets.

The analysis completed thus far gives us confidence that many commercially available drones and GPSs will be able to deliver the precision we require for far field mapping. We now describe the calibration system constructed at Yale, which abides by these constraints.

5.3 Drone system overview

The drone system used by the Yale radio group consists of a DJI Matrice 600 Pro drone instrumented with two real time kinematic (RTK) GPSs and a series of on-board inertial measurement unit (IMU) sensors, which include accelerometers and gyroscopes to detect linear acceleration rate and angular motions. This section details the motivation for using our specific drone and GPS and describes verification tests of GPS and IMU performance.

5.3.1 The DJI Matrice 600 Pro drone

The DJI Matrice 600 Pro drone[1] has been the Yale drone group's UAV since 2018. It was chosen for its long flight times (38 minutes unloaded, 18 minutes fully loaded), manageable size (10kg unloaded, physical dimensions 1668 mm × 1518 mm × 727 mm), moderate price point, and compatibility with existing differential GPS systems. The drone has proved incredibly stable in flight, and capable of flying planar grids with the noise source and transmitter payload attached for up to 30min at a time. The drone has been robust to hundreds of flights, weathering difficult environmental conditions and surviving a crash landing into a tree (section 5.8).

The commercial availability of the DJI Matrice 600 Pro means that replacement parts have been straightforward to procure when needed for repairs. The flip side of this is that we are beholden to DJI continuing to sell the Matrice 600 and stock Matrice 600 replacement parts. Another downside is that the drone runs proprietary software that makes the data read-in a challenge, and it does not interface easily with external autopilot software.

---

[1] https://www.dji.com/matrice600-pro/info

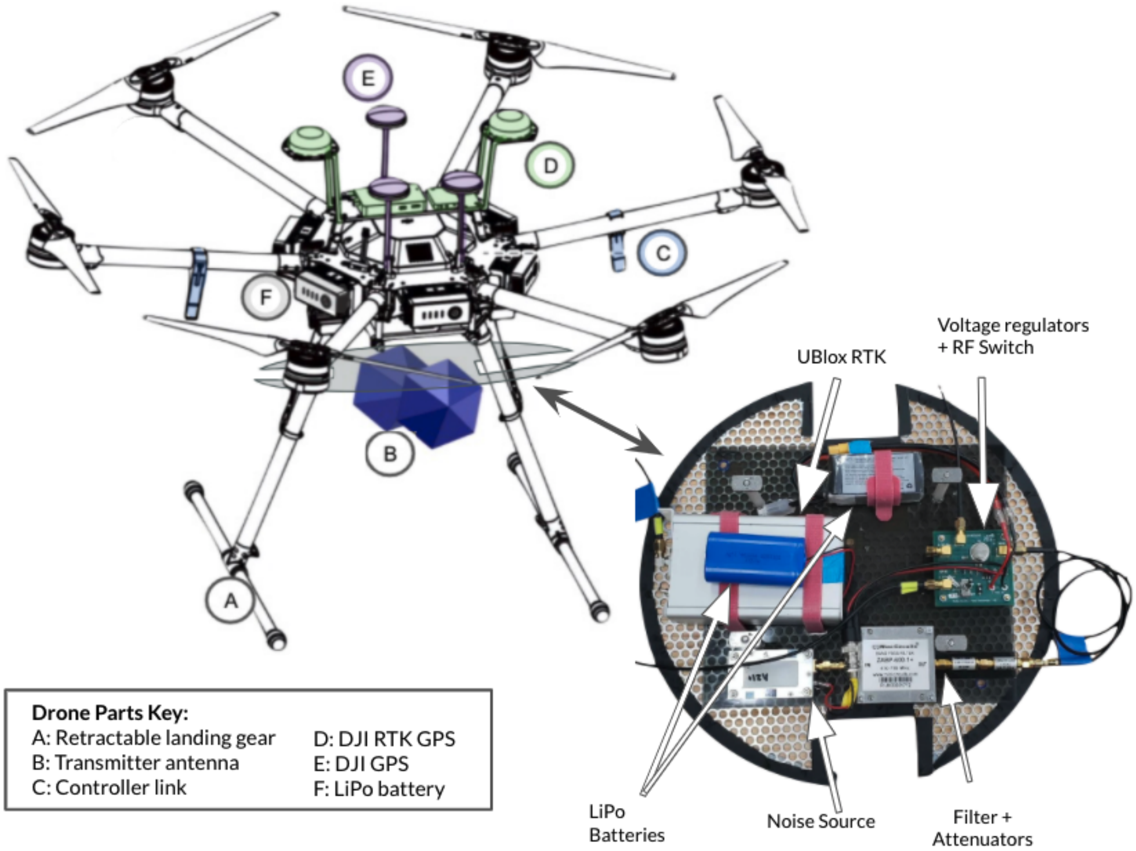
E
D
C
F
B
A
UBlox RTK
Voltage regulators
+ RF Switch
LiPo
Batteries
Noise Source
Filter +
Attenuators
Drone Parts Key:
A: Retractable landing gear
B: Transmitter antenna
C: Controller link
D: DJI RTK GPS
E: DJI GPS
F: LiPo battery

Figure 5.7: Drone calibrator and mounting plate schematic. The calibration system includes a DJI Matrice 600 Pro drone and Real Time Kinematic GPS, biconical transmitter antenna, and calibrated noise source chain. Plate and antenna are not to scale on the drone schematic. Drone diagram from DJI.

5.3.2 RTK GPS integration and verification

The Global Positioning System (GPS) is administered by the US government[2] and consists of a constellation of 24 satellites transmitting one-way signals that give their position and time, as recorded by a precise atomic clock. A GPS receiver measures the time difference between its local clock and the clocks on visible satellites to determine its distance to the different satellites, and from that information can triangulate its location. For the drone GPS to operate effectively and without communicating warnings, it must see a minimum of eight satellites. A standard GPS can give up to meter-scale locational accuracy depending on the number of visible satellites. We target cm-level accuracy for understanding the drone location, which will be necessary for future near-field beam mapping efforts, and requires a higher-precision GPS.

---

[2] https://www.gps.gov/

A Real Time Kinematic (RTK) GPS is a particularly accurate type of GPS that can give the cm-level location accuracy we desire. RTKs involve a ‘base station’ in addition to the GPS receiver, which remains stationary for the duration of measurements. RTKs work in a similar manner to ordinary GPSs, with the receiving GPS and base station both measuring their respective positions from satellite communications. The RTK then exploits the fact that the base station is stationary, looking at variations in its constructed location (as determined by satellite communications), flagging variations as positional errors (since the RTK knows it is stationary), and streaming corrections to the receiving GPS that can be applied in real-time to the position data as it is collected.

The Yale drone system utilizes two RTK GPSs:

1. A DJI RTK GPS that is specifically compatible with the Matrice 600 Pro, and is mounted to the top of the drone. This RTK was integrated in 2019 by Yale undergraduate Ry Walker [104].
2. A UBlox RTK GPS, which is programmable and customizable, and uses a series of low-cost, interchangeable, and well-documented parts. It is mounted on the electronics plate that sits below the drone, alongside the noise source RF chain. This RTK was integrated in 2020 by Yale undergraduate Annie Polish as a senior thesis [105].

The DJI RTK GPS was easy to integrate, and was verified to perform to our cm-level accuracy specification. There were, however, several issues with the DJI RTK, which are summarized by Annie Polish in her senior undergraduate thesis [105]: "The D-RTK costs almost as much as the drone itself, has been discontinued, and is only compatible with our (also discontinued) model of drone. The flight logs, which include the RTK GPS data, are encrypted, and we are only able to access them using an additional tool called DatCon[3]. Accessing the logs is a tedious and time-consuming process, and if DatCon ceased development, we would no longer be able to access our drone's data at all. In addition, the DJI D-RTK is missing several important features. It does not provide estimates of its own error, nor does it indicate which GPS correction methods are active at any given time. It also does not have a GPS-locked Pulse-Per-Second (PPS) output. The PPS output is an easy way to access an extremely accurate clock, and all GPS PPS outputs are in sync with each other to within a few nanoseconds. We require a PPS output in order to pulse our noise source in sync with a telescope correlator's integration period."

To make up for the DJI GPS shortcomings and build redundancy into the system, we integrated a new RTK, the uBlox RTK GPS, which meets the following criteria, outlined in [105]:

- Produces data in a format that is both human- and computer-readable
- Built mostly from off-the-shelf parts for lower cost and to make the system replicable by ourselves or other groups
- Provides a configurable Pulse-Per-Second output
- Sufficiently rugged, small, and lightweight to fly on our drone
- At least 8 hours of battery life
- Communication radio range of at least 2 kilometers

---

[3] Available at https://datfile.net/

- RF-shielded for use on radio-quiet sites

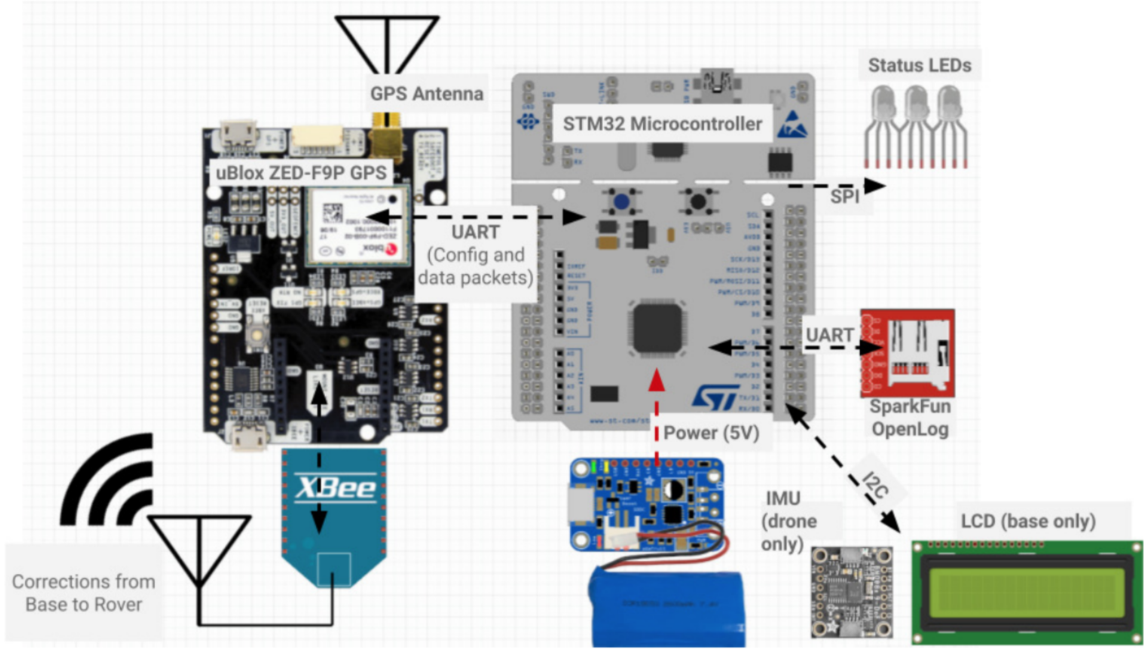
GPS Antenna
Status LEDs
STM32 Microcontroller
uBlox ZED-F9P GPS
SPI
UART
(Config and
data packets)
UART
SparkFun
OpenLog
Power (5V)
XBee
IMU
(drone
only)
I2C
LCD (base only)
Corrections from
Base to Rover

Figure 5.8: Block diagram for the UBlox RTK GPS, which meets the requirements for a drone calibrator GPS, including cm-level positional accuracy. It is programmable and customizable, and uses a series of low-cost, interchangeable, and well-documented parts. Design and figure credit: Annie Polish.

The UBlox GPS consists of an Arduino-compatible microcrontroller (STM32) programmed in C that reads a configuration file from an SD card. Communications are via an XBee radio[4], which currently operates around 900MHz though can be easily moved in frequency. Data is saved to a SparkFun OpenLog[5] data logger and saved on an SD card. A schematic is shown in fig 5.8.

Both the DJI RTK GPS and the UBlox GPS were evaluated for positional accuracy in a similar series of ground-based tests. In these tests, the drone plus GPS system was left stationary for 30 minutes, which is the anticipated duration of a drone flight. During that time, GPS data was recorded, and drifts or jitter from stationary position were logged. For both GPS systems, the jitter remained within a ±2cm box, with sub-cm RMS values in both dimensions for the DJI RTK and in one dimension for the uBlox. Results for one minute of data (which captures the extent of the jitter) are shown in fig. 5.9.

---

[4] https://www.digi.com/xbee

[5] https://www.sparkfun.com/products/13712

Data from routine test flights confirms the results from the dedicated positional accuracy testing. During flights, we follow a procedure of locating the drone on a marked location for take off and returning it to that location after landing. The drone sits stationary at that position, recording GPS data for at least 1 minute. From this data we find the GPS positional errors to be within the 2cm human placement error. We have similarly compared the recorded flight paths from both GPSs for a series of flights and found the absolute location to vary between $\sim 1-5$ meters depending on the flight (inconsistency in absolute location is not a problem, particularly given that we fly at telescope sites that have GPS survey markers able to provide a standard reference for location calibration).

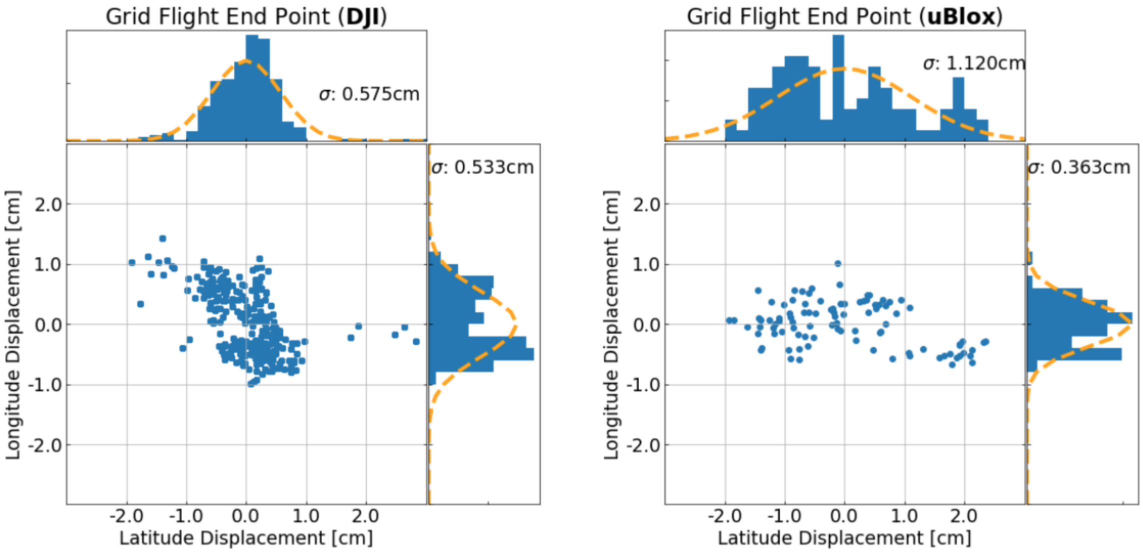

Figure 5.9: Positional accuracy of the DJI and uBlox GPS systems, as determined by recording jitter from the systems when stationary. Both GPS systems have ±2cm positional accuracy (further verified with flights testing), which is sufficiently stringent for far-field beam mapping. Figure and analysis by Annie Polish.

For the DJI GPS in particular, we also used on-site survey markers at Owen's Valley Radio Observatory in Bishop, CA to investigate the accuracy of the drone's on-board GPS. We did this by attaching a laser to the bottom of the drone, positioning it exactly over the survey marker (to within sub-mm), waiting for roughly one minute, and rotating the drone 90 degrees to repeat the measurement. After a first pass of measurements, we determined that the drone's GPS center is not at the center of the drone but rather beneath one of the RTK antennas 33mm offset from center. Repeating measurements with this knowledge, we determined that we have good relative positional accuracy, though further testing at verified GPS stationary markers is required for evaluating absolute accuracy. Because we expect any offset to remain constant within a flight, this error could be mitigated in the future by always landing at a GPS-marked location and calibrating off of that reference.

We have also performed height tests with the differential GPS installed. This was done using a vertical frame made from PVC, which allowed the drone to sit at various known heights for several minutes while data taking. We determined the drone's relative positional accuracy to be ±1cm, but still need to assess the absolute positional accuracy, and repeat similar tests in flight.

5.3.3 IMU sensor verification

In addition to understanding drone locational accuracy, we must be able to quantify the way the drone tips, tilts, and rotates during flight, as these angles will impact the drone transmission to the telescope. These angles are referred to as roll, pitch, and yaw, defined in the fig. 5.10 diagram. We have designed and built angle and position testing props to assess these specs.

Our first accuracy test was for pitch and roll angular sensitivity, performed separately by Maile Harris and Ry Walker. We bolted the drone to a sine plate which locks at angles 0, 5, 10, and 15 degrees. We powered the drone and held it at each

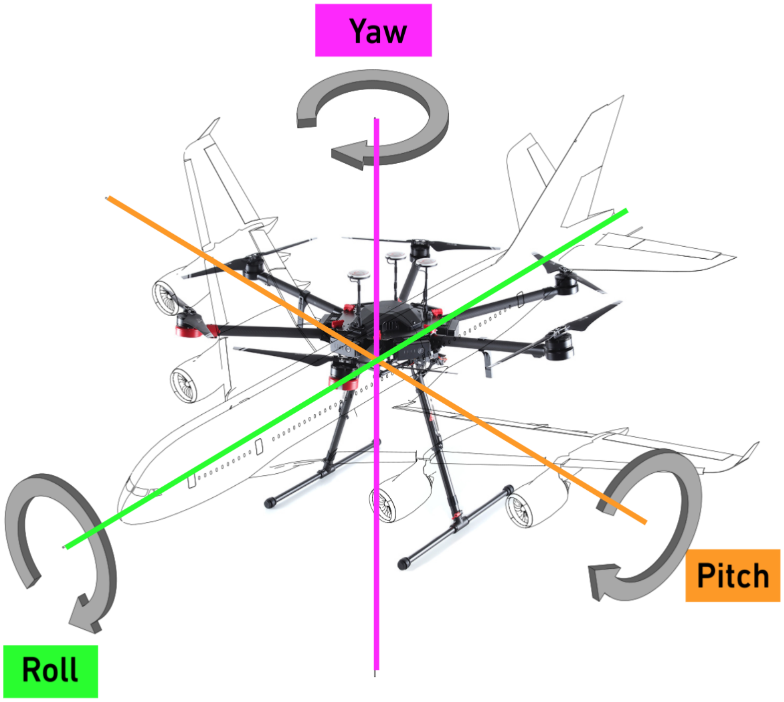

Figure 5.10: Drone roll/pitch/yaw angle definitions. Figure credit: Maile Harris.

respective angle for one minute. Looking at the resulting data, we see that the sensor reads out the expected value with fluctuations up to 0.2 degrees, which is within our requirement.

Accuracy tests for polarization angle were conducted by sitting the drone on a rotation plate with resolution of arcminutes. Again, the drone was held at a variety of known angles for one minute. The data showed a drift over time at several different angles, despite the drone sitting completely still. This measurement was performed without the RTK integrated, and should be repeated with the full set up. If angular drifts are still recorded, we can add additional redundant internal sensors.

5.3.4 RF chain design

For the drone calibration technique to succeed, a stable, broadband noise must be broadcast from a calibrated source during flight. This is accomplished with a 95dB

ENR calibrated noise source filtered to 400-800MHz and attenuated to an appropriate output power (which will vary based on telescope); an RF switch to toggle transmission between a noise source and 50Ω, allowing regular sky background measurements; and a biconical antenna transmitter, mounted on a ground plate for shaping the beam. The electronics are powered by two different 15V LiPo batteries, which last for a full day of flights. The individual chain components are discussed below, and detailed in table 5.2. A schematic of the full chain is shown in figure 5.7.

Component	**Weight (g)**	**Current Draw (mA)**	**Input Voltage (V)**
Noise source	64	173	12
Switching board	76	10	16.8
UBlox GPS	314	300	4.2
Filter	108	–	–
LiPo 1	149	–	–
LiPo 2	95	–	–
Attenuators	18	–	–
Bicolog Transmitter	362	–	–
Mounting Hardware	300	–	–
Ground Plate	116	–	–
Total Weight (g): 1602		**Total Power Draw (W): 3.504**	

Table 5.2: The current weight and power budget for the drone calibrator. It is important to be cognizant of the weight of the payload, as the drone is only rated for a 6kg payload, and utilizing the full payload decreases flight time by factor of 2.

Noise source

There are several requirements for the drone noise source:

- The spectrum should be flat to within 1dB across the HIRAX band, enabling a measurement of the full band at once. If the spectrum were sufficiently sloped, we would have to choose between (1) digitally saturating the telescope at the high-signal end of the band to accurately measure the low-signal end of the band and (2) accurately measuring the high-signal end of the band while taking low-sensitivity measurements in the low-signal end.

- The power should be broadband, encompassing 300MHz-1.5GHz. This is because we desire to calibrate telescopes that measure the un-redshifted 21cm signal (1-1.5GHz) in addition to HIRAX (400-800MHz). Chapter 6 details measurements of such telescopes, which were incredibly useful for drone calibrator development.

We adopt the noise source originally used by the Tianlai collaboration for their drone calibration system [106]. It has 95dB ENR across the HIRAX band, and is operational from 10MHz-1500MHz. We have tested four of these noise sources, finding the spectra flat to dB-level across the HIRAX band, and similar to within 2dB (shown in figure 5.11).

Filtering, shielding, and attenuation

Filtering, shielding, and attenuation are all critical parts of the drone calibrator RF chain. We use a very bright noise source, which can interfere with GPS communication frequencies if it is not appropriately filtered and attenuated. In particular, we must filter out 900MHz, which is where the GPS operates, and attenuate by at least 30dB when operating in L-band (1-1.5GHz)—it was determined experimentally during test flights that the drone will not link to the GPS if too much L-band power is broadcast. The filter we use provides > 40dB attenuation outside of the HIRAX frequency band. It lapses in performance above 1.5GHz, though the noise source does not transmit brightly in this regime so the overall power-level remains appropriate. Filter and noise source performance are shown in figure (5.11).

Balancing attenuation levels is a key piece of the flight plan strategy. We have to be careful not to saturate receiver electronics by transmitting a signal that is too bright. Additionally, we typically fly two different modes of flights—one out to wider angles with higher power to capture low-level sidelobe features (this will saturate in the main beam), and another as a denser grid at lower power to capture main beam

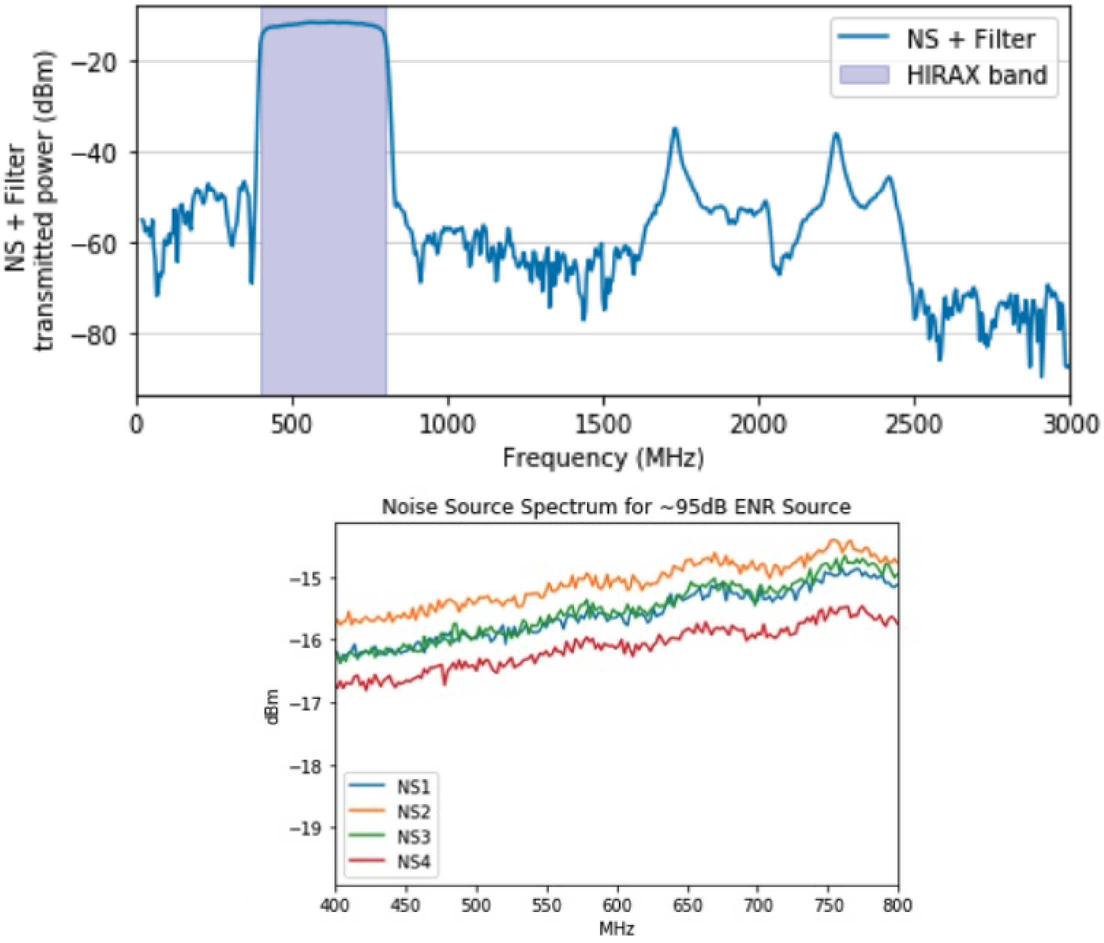

Figure 5.11: Top: Spectrum of the un-attenuated noise source plus filtering chain that is flown on the drone (typically attenuated down by ∼35-45dB). The Y-axis units are dBm/3MHz. The HIRAX band is shaded, and it is evident that the filter filters sharply around it. This is important for a number of reasons, including that broadcasting with sufficient power at 900MHz can blind the GPS to satellite signals.

features—so control over attenuation levels is important. We use chains of standard minicircuits attenuators for this power regulation. We are additionally working on integrating a variable attenuator into the chain to simplify this process and eliminate the need for multiple RF chains.

We additionally have to attenuate the noise source transmission to within the broadcast limits outlined by the SKA, who set RFI restrictions around the HIRAX site, as outlined later in this chapter. These restrictions impose 30dB attenuation automatically for the noise source. A critical part of keeping the noise level controlled is appropriately SMA wrench-tightening the component connectors, as loose connectors can radiate considerable power. The current RF chain plate design discourages (and eliminates the need for) unscrewing the noise source-to-filter connection.

Switching board

During beam mapping flights, the telescope records continuous spectra as the drone flies a calibrated noise source overhead. Each of these spectra is primarily a measurement of drone transmitter power, but will also include contributions from the drone's own RFI as well as a shifting sky background. To make the most sensitive maps possible, we should remove the background signals from the measurement, and favor a method of removing backgrounds at each drone location. This background subtraction can be achieved by pulsing the drone noise source on and off in flight, with a cadence that allows for multiple on/off signals within 1 degree of the beam. The noise source modulation is achieved with timing board, which was designed by Maile Harris as an undergraduate senior thesis project [107].

The switching board design handles both RF and DC inputs and outputs. It is based around a TOSW-230+ RF switch from minicircuits[6], and includes voltage regulators and DC filtering components (i.e. capacitors, pi filters). The board takes a 5V TTL input (delivered by the uBox GPS), which dictates the switching cadence, and a 15V DC input (from a LiPo battery), which is regulated down to 12V to power the noise source and 5V to power the switch. The board also takes two inputs to the switch—(1) the filtered and unmodulated noise source signal and (2) a 50Ω terminator—and alternatingly passes each signal to the output port, which is connected to the transmitter antenna. A diagram of the switching board is shown in fig. 5.12

---

[6] https://www.minicircuits.com/pdfs/TOSW-230+.pdf

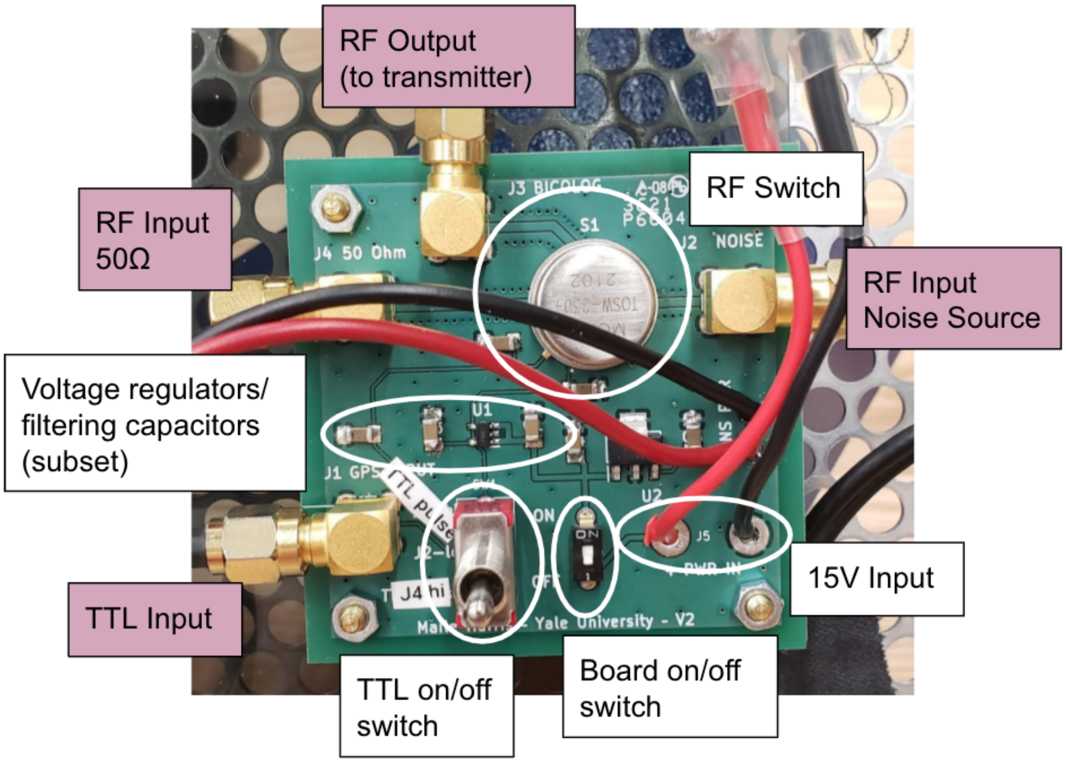
RF Output
(to transmitter)
RF Input
50Ω
RF Switch
RF Input
Noise Source
Voltage regulators/
filtering capacitors
(subset)
TTL Input
15V Input
TTL on/off
switch
Board on/off
switch

Figure 5.12: The noise source switching and power regulation board. The board takes a TTL input from the uBlox GPS to effectively modulate the noise source on and off, allowing the telescope to measure sky background and drone RFI at each drone position so they can be removed from the beam map in analysis. The board additionally provides voltage regulation and filtering. Board design by Maile Harris [107].

For the switch to work successfully, it needs a high-degree of isolation between the noise source input port and the RF output port when the switch is passing 50Ω. Otherwise, the "background measurement", taken with the assumption that the noise source is not broadcasting, will include additional signal beyond the drone RFI and sky temperature. Isolation was measured in multiple ways, including with both a VNA and spectrum analyzer, and found to be 40dB+ between the two ports as tested with a 50dB+ ENR noise source. This isolation exceeds even the performance of the block version of the RF switch (tested separately, model number ZSDR-230+ from minicircuits, denoted ZSDR in fig. 5.13). The power draw of board and noise source is sufficiently low that the battery can last full flight day (projected 17 hours from lab tests). The switch isolation results are shown in 5.13.

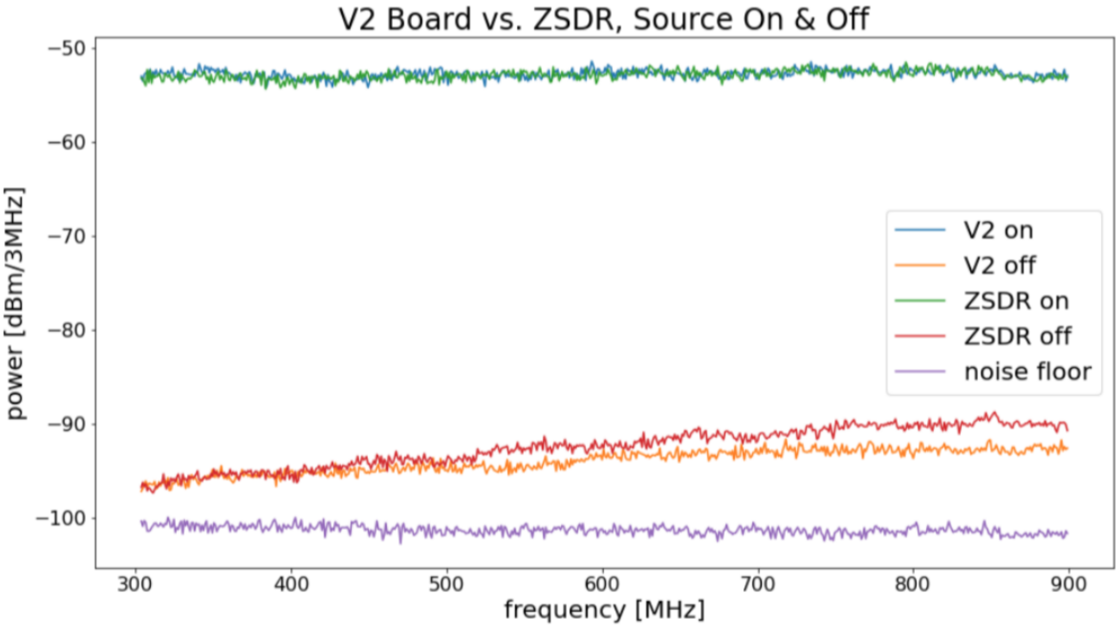

Figure 5.13: Measurements of the RF switch board and commercial block switch outputs for inputs (1) a filtered and attenuated noise source and (2) a 50Ω terminator. In both instances there is 40dB+ isolation, which should enable clean background measurements. It is notable that the switching board outperforms the block switch in isolation. Measurements and figure by Maile Harris [107]

Work is underway to develop a two-output switch with a multitude of pulsing options and improved functionality. The new design will allow for a dual-polarization transmitter, which in turn will enable simultaneous measurements of the two HIRAX telescope polarizations with a single flight. This upgrade would cut total flight time for a single dish in half, and enable measurements of more elements and exploration of wider beam angles. A dual-output switching board has been designed and is currently being fabricated and readied for testing.

Transmitter antenna

The calibrated noise from the noise source RF chain is transmitted with a biconical antenna from Aaronia (models 20300 and 30100). This style of antenna is chosen because it is lightweight and broadband and has a broad beam. Extensive characterization of the beam properties is detailed in Chapter 3.

The biconical antennas from Aaronia have a common fail point worth noting. The bicolog is mounted from a 1/4-20 threaded socket, as is standard for tripods, which in the case of several of the antennas has no stop for the mounting screw. This is problematic as there are circuit components positioned directly under the socket with a hole behind them, such that if a screw or rod is threaded too far, it can push on the components and snap their soldered wire connections. Connections were severed on 2/2 of the bicologs with no stop used in our lab, and caused poor antenna performance ($S_{11} > -3dB$). Repairs were trivial—a simple resoldering of wires—and after discovering this fail point, we were able to get different mounting plates with a stop from Aaronia.

5.4 Autopilot software

Future 21cm cosmology pipelines will require high-resolution 2D beam maps. For drone calibration to achieve this, the drone must be capable of flying a controlled, gridded pattern (planar or spherical). Flying this sort of pattern is best done with an autopilot, so finding a suitable software is a critical part of developing the overall calibration system. There are several considerations for an autopilot:

- Ability to program a 2D grid
- Can be activated reliably without internet
- Has a clear manual override
- Interfaces with DJI (specific to the Yale system, which utilizes a DJI drone)

The autopilot software we settled on was UGCS flight controller[7]. It has smooth start up operations and an intuitive interface for planar gridded flights. Most critically, it can interface with DJI, which is not the case for many of the commonly used drone autopilot software options. The software, in its current version, may not be suitable for more complex flight plans intended in the future, such as spherical grids.

---

[7] https://www.ugcs.com/

5.5 Travelling with the drone

We intend for this drone system to serve as a calibrator for internationally-based telescopes including HIRAX in South Africa. It must therefore be manageable for airplane travel. We assembled a carrying case for the DJI Matrice drone that successfully transported it from Connecticut to California and back, and has kept it safe on multiple 10+ hour car rides. This case consists of a $24 \times 24 \times 24$ inch pelican case[8] in which the original foam packaging from DJI is embedded. The DJI foam, which is custom-fitted to the drone, is surrounded with packing foam for added padding and structural support. When in its case, the drone can be checked to the plane cargo area.

We travel with at least three sets of drone batteries, which is the minimum to maintain continuous flight operations. These batteries are transported in a military-spec hard-shell carrying case[9], and must be held as a carry-on per airline restrictions on LiPo batteries. Smaller LiPos, for powering the uBlox GPS and noise source, are carried in fire-proof bags for safety[10].

There have been additional pieces of equipment added to the packing list over time after proving especially useful for remote operation. These include a low power-draw windows laptop, for running the drone autopilot software and the DJI Assistant software (which is needed to perform in-field diagnoses on the drone). A portable power station such as a jackery[11] has also been critical for keeping the laptop, flight controller, and GPS batteries charged. We additionally have built a transport backpack, that houses a complete toolkit and backup electronics.

---

[8] `https://www.pelican.com/us/en/product/cases/cube-case/protector/0370`

[9] `https://www.bhphotovideo.com/c/product/1306343-REG/go_professional_cases_gpc_dji_m600_btry_12_dji_matrice_600_battery.html`

[10] `https://www.amazon.com/ENGPOW-Fireproof-Explosionproof-Resistant-260x130x150mm/dp/B07BN7WLQ5?th=1`

5.6 Radio Frequency Interference (RFI)

Radio quiet sites, such as the SKA site in the Karoo Desert where HIRAX is based, are explicitly located away from cities, and regulate the areas around them to limit power from cell towers, TV stations, and other radio contaminants. Typically, Ethernet is the source of internet connection, as WiFi is prohibited, and even microwave ovens are not allowed unless contained within a Faraday cage. For the drone calibrator to be suitable for use at such radio sites, it must also comply with stringent RFI limits. As a result, RFI characterization and mitigation are critical pieces of drone calibrator development.

We anticipate several sources of RFI from the drone—it has intentional emitters for communications, as well as clocks, power supplies, and motors—which are corroborated by FCC documentation[12].

To obtain permission for flights in the Karoo, the drone calibrator must comply with the following restrictions, designed for the MeerKAT array[66]:

- Broadband sources of RFI must be under-64dBm (at 10MHz) to -35dBm (at 1GHz) EIRP ("Effective Isotropic Radiated Power")
- Narrow band sources of RFI, such as intentional emitters, must be under -50dBm (at 10MHz) to -20dBm (at 1GHz) EIRP

---

[11] `https://www.jackery.com/`

[12] `https://fccid.io/SS3-PM8201512`

Within the next decade, restrictions will be tightened to accommodate the SKA requirements. This will be during the lifetime of HIRAX, and so the drone should ideally abide by them. The SKA-mid restrictions are as follows:

- Broadband sources of RFI must be under ~ -90dBm EIRP
- Narrow band sources of RFI, such as intentional emitters, must be under ~ -75dBm EIRP

The requirements for MeerKAT are plotted alongside HIRAX RFI measurements in figure 5.15.

This section describes several modes of RFI assessment and current mitigation strategies.

5.6.1 RFI assessment

RFI measurements of the drone were taken in a variety of different environments:

- An anechoic chamber/EMI chamber at the Green Bank Observatory, which is used for Green Bank's own RFI compliance testing. This set up allowed for sensitive measurements of the stationary drone, but did not measure intentional emitters or in-flight source of RFI such as motors.
- Outdoors at Green Bank Observatory, where the pristine RFI-environment gave the opportunity to evaluate intentional emitters and in-flight RFI.
- A Faraday cage at Yale University, for assessing mitigation effectiveness.

Set ups in the Green Bank EMI chamber and outdoors at the Green Bank Observatory are shown in figure 5.14.

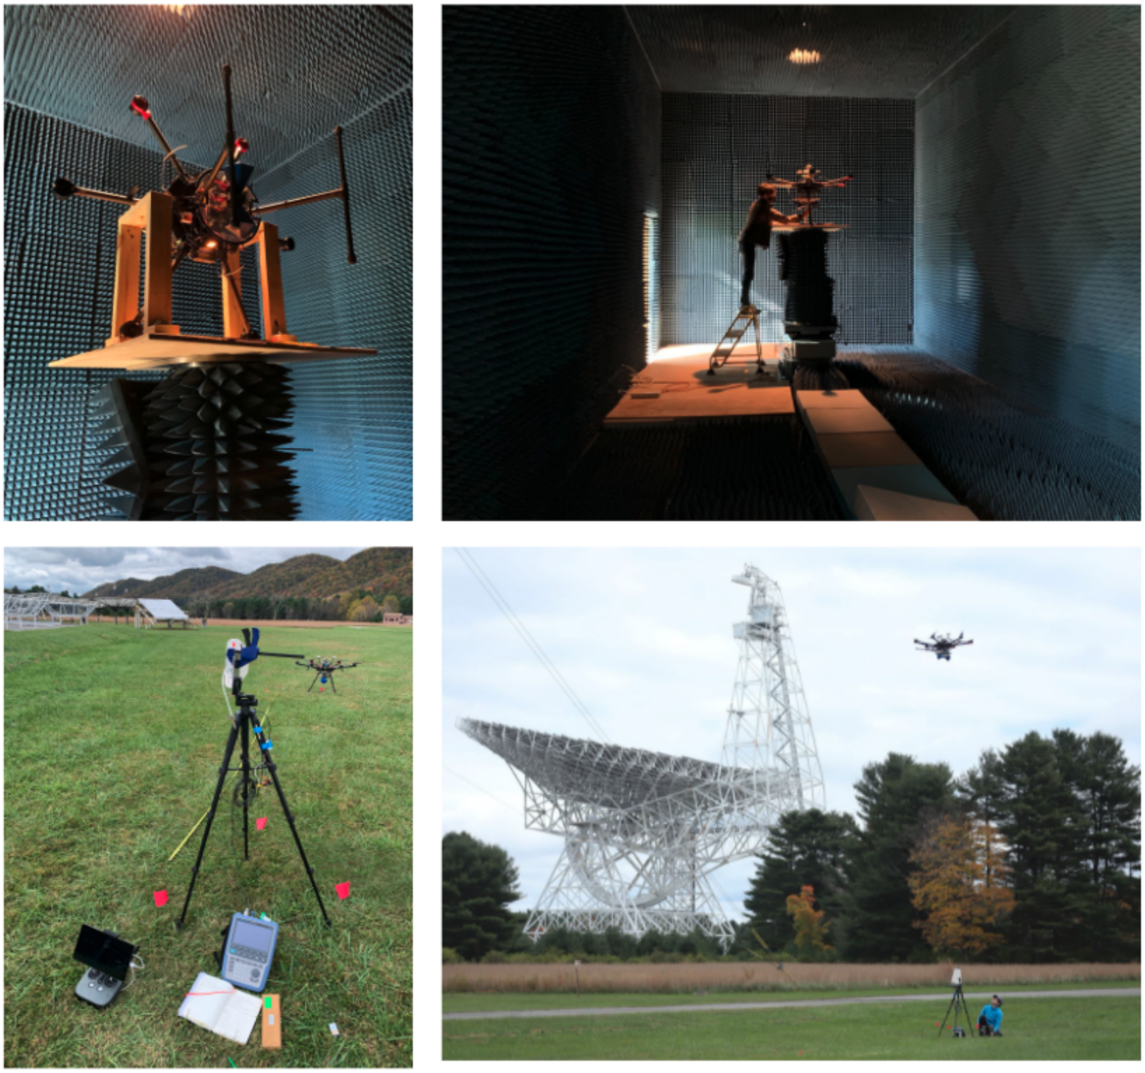

Figure 5.14: Top: Green Bank EMI Chamber set up for drone RFI assessment. Bottom: Outdoor RFI measurements at Green Bank Observatory.

In an EMI chamber

We utilized the anechoic chamber at Green Bank Observatory to assess broadband drone RFI at a variety of drone orientations, looking to understand both its magnitude and directionality. Drone orientations are shown in fig. 5.14.

Measurements were made from 80MHz-12GHz in 2GHz sections and with 10kHz resolution. In each measurement block, a ‘quiet room’ background spectrum was taken for reference. A log periodic receiver was used for measurements below 1GHz, and a horn antenna was used for all other measurements. These receivers were located 7m from the drone, as is standard for compliance testing. It is worth noting that the absorber in the chamber is not rated for measurements below 1GHz, so we are likely obtaining an upper bound in that regime (due to measuring radiation reflected from around the chamber). The chamber measurements should capture RFI sources such as the power supply and digital electronics, though intentional emitters will require additional characterization.

Results from chamber measurements are shown in figure 5.15, which shows the RFI across the full 80MHz-12GHz that SKA restricts, and in figure 5.16, which zooms in on 80MHz-1GHz where much of the RFI is concentrated. Fig. 5.15 also indicates the expected location of radiation from intentional emitters. Both figures overlay measurements with MeerKAT RFI restrictions. Table 5.3 gives a breakdown of RFI sources and emissions. Taken together, these results reveal that the drone is RFI-quiet above 1GHz to within the noise floor of the measurement device (excluding intentional emitters), though has low frequency emission below 600MHz. By rotating the drone in 15 degree increments, we determine that the RFI is not meaningfully directional. We additionally used near-field probe measurements to diagnose the low-frequency RFI as from the power supply, and notice RF leakage through transmitter chain electronics. We can now proceed with shielding both emission sources.

It is worth noting that the SKA restrictions are substantially tighter than MeerKAT restrictions, and the Green Bank Anechoic Chamber was not sensitive to below the SKA limits. Thus, we may need to pursue additional RFI characterization in the coming years.

In the field

The EMI chamber measurements must be augmented with measurement sets that can assess motors and intentional emitters. We perform these measurements outdoors at

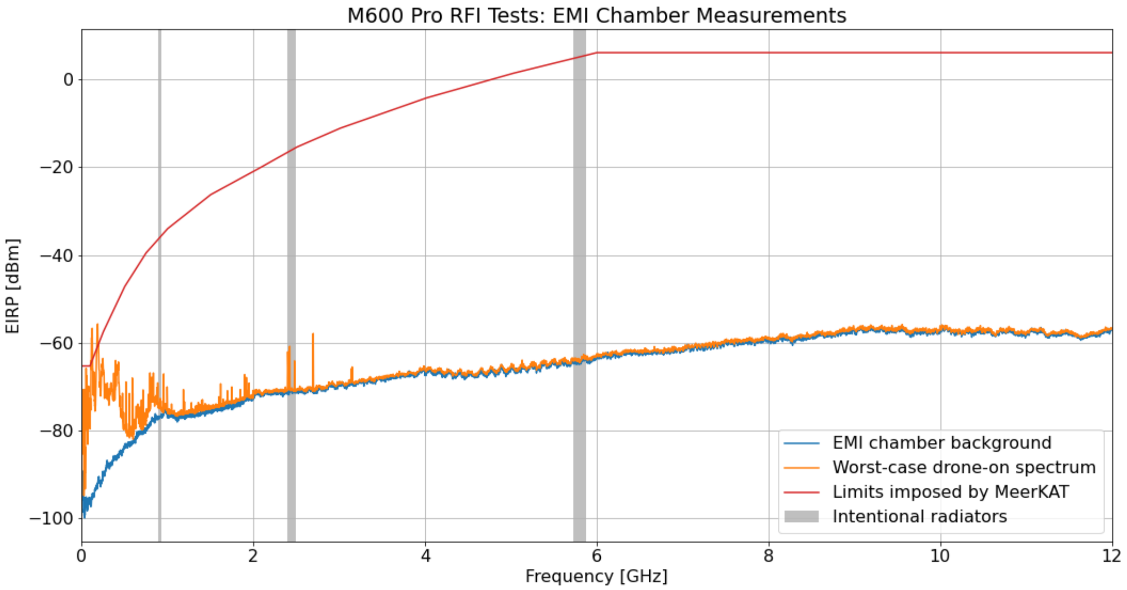

Figure 5.15: Drone RFI measurements across the band regulated by MeerKAT and SKA. With the exception of intentional emitters and power supply emission below 600MHz, the drone is quiet to below the measurement instrument noise floor. Plot and analysis by Annie Polish.

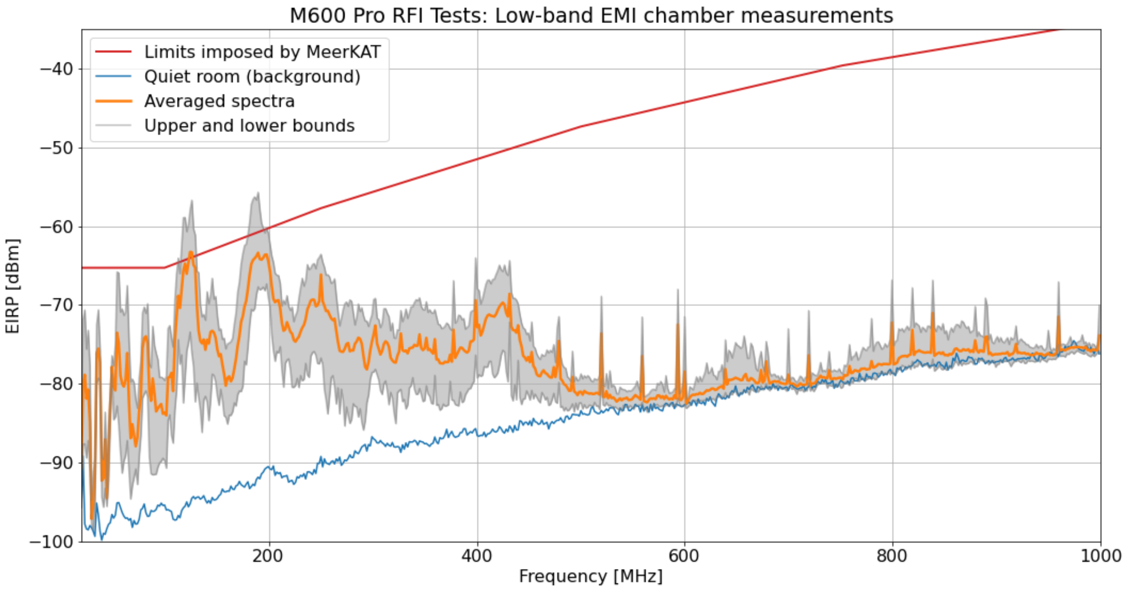

Figure 5.16: Low frequency drone RFI measurements indicate emission in the lower portion of the band. With the exception of intentional emitters, this emission is broadly within the MeerKAT specifications. It will need to be mitigated to comply with the SKA limits. Plot and analysis by Annie Polish.

Green Bank, using a two different yagi antennas (for different measurement bands) as the receivers. The spectrum is amplified with a commercial minicircuits LNA and a max-hold measurement is recorded by a Rohde & Schwarz spectrum analyzer (this is because we are primarily interested in an upper bound on the drone RFI to ensure compliance). The spectrum analyzer limits measurements to below 3GHz.

We again perform the RFI measurements at a variety of drone orientations, which capture the drone emissions from its side and from its bottom, in order ensure we are not missing directional emissions. The drone is hovered 8.5m above ground level corresponding to a distance of 7m from the receiver, which is standard practice for this manner of measurement and the same distance used in the EMI chamber.

By measuring spectra from the drone stationary on the ground (motors off) and the drone hovering a few cm above the ground (motors on), we can assess the RFI from the drone motors, which were a source of concern. We find they are undetectably quiet. We also used in-field measurements to characterize intentional emitters, finding bright emission in the communications bands, and leakage around 1GHz from a faulty cable (which underscores the importance of regular equipment assessment).

In field drone RFI spectra do not fully align with EMI chamber spectra, likely due to environmental factors. Still, the measurements are instructive, and a good jumping off point for future work.

In the lab

We are presently comparing the spectra taken at Green Bank to spectra taken in a Faraday Cage at Yale University. The set-up is similar to the EMI chamber at Green Bank in that the drone is situated in several orientations and measured by a receiver several meters away in the far field (in this case, the receiver is a biconical antenna). Pieces of pyramidal RF absorber are strategically placed around the cage to limit reflections. The shielding inherent in a faraday cage should block stray signals and

Component	**Band**	**MeerKAT Compliant?**
Intentional Emitters		
Noise Source	400-800MHz	✓
RTK GPS	900-930MHz	X
Controller Comms 1	2.4GHz	X
Controller Comms 2	5.8GHz	X
Unintentional Emitters		
Power Distribution Board	$<$ 500MHz	X
Motors	–	✓
Digital Communications	1-3GHz (sporadic)	✓

Table 5.3: Sources of RFI from the DJI Matrice 600 drone Pro + RTK GPS + custom noise source chain.

ensure that everything recorded by the spectrum analyzer is drone emission. If we can recover the RFI spectral features measured at Green Bank in our local environment, we can test mitigation strategies locally and eliminate the need for repeated travel.

5.6.2 Mitigation plan

We are currently exploring several RFI mitigation strategies. For intentional emitters, we are investigating whether communications can be moved up in frequency to 5.8GHz, which is less tightly regulated by the SKA. Within the drone, we are prototyping different shielding methods for the power supply, and designing a mesh box to enclose on-board electronics. We hope that after careful mitigation, we will be under the Karoo RFI limits, with the exception of intentional emitters. At this stage, we can consider moving communications into less-regulated bands, as well as developing on-ground structures for blocking line-of-site between transmitters and the nearest SKA dishes. If complete mitigation proves untenable, we can plan flight campaigns for telescope maintenance periods.

5.7 A test bed at Yale's Leitner Family Observatory and Planetarium (LFOP)

Access to a local testbed in the 400-800MHz HIRAX band will be an important part of iterating on the drone design. Previously, hardware upgrades were tested at other institutions, but the travel required limited real-time feedback on new prototypes. The importance of a local test bed was underscored in 2020, when pandemic travel restrictions limited research travel and hampered our ability to fly at telescopes for over one year. As a result, a testbed was developed at Yale's Leitner Observatory, with work led initially by Ry Walker (as a Yale undergradaute senior thesis project[104]) and later taken over by graduate student William Tyndall.

Building a telescope at Yale comes with several challenges. Yale's location in New Haven, Connecticut means contending with a high-RFI environment, which is shown in Figure 5.17. The RFI renders much of the HIRAX band unusable for science, and will compress the analog and digital telescope electronics without proper filtering. Thus, we identify a "clean" portion of the band, from 650-700MHz, for the telescope operations; a concession that will enable us to make beam maps locally. The urban location also places us in proximity to airports and thus within restricted airspace zones. The Yale team had to obtain clearance from the FCC to fly a drone at LFOP, which was successful though limited operations to the observatory grounds.

LFOP is a public observatory that is regularly visited by community members, with building safety restrictions that prohibit students from working on the roof. As a result, we could not install a permanent telescope on the site grounds and could not mount it on the roof. Essentially, we had to create a 'pop up telescope' that can be moved into a field for measurements and then retracted once work is done. The testbed that was ultimately designed and built consists of a 3m dish that sits within a frame made from aluminum T-slot material. The frame sits atop wheels, enabling the

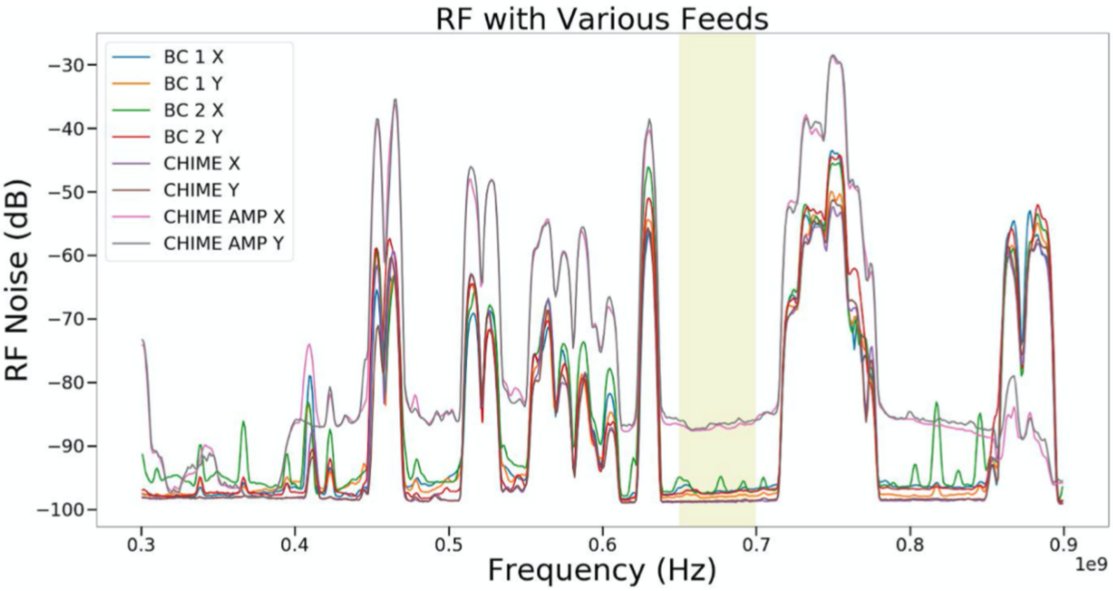

Figure 5.17: A spectrum at Yale's Leitner Observatory shows considerable RFI accross most of the 400-800MHz HIRAX band. We locate a clean portion of the band, from 650-700MHz (highlighted), for telescope operations. This band will be filtered around sharply in the telescope signal chain. This filtering keeps the RFI from compressing the telescope electronics and enables a functional testbed. Figure and measurements from Ry Walker.

telescope to be rolled out to a field for operations and then stored behind a building. There are breaks on the wheels for stabilization, and sandbags to rest on the frame during storage, to further affix the telescope in place. Bubble levels are attached to the telescope frame for alignment. A vivaldi-style antenna is used as the receiver, which was designed and fabricated at the University of Toronto as a prototype for the future CHORD telescope array[49]. It is attached with metallic feed legs, which are expected to impact the beam. A photograph of the testbed when readied for measurements is shown in fig. 5.18.

The signal received at the antenna travels down 100ft of LMR400 cable, and is then amplified and filtered by an RF chain consisting of a series of commercial minicircuits LNAs, attenuators, and custom 650-700MHz SAW filters designed to cut power sharply outside of the band. At this stage of the measurement chain, the signal power should be -23dBm accross the 400-800MHz HIRAX band, which is the

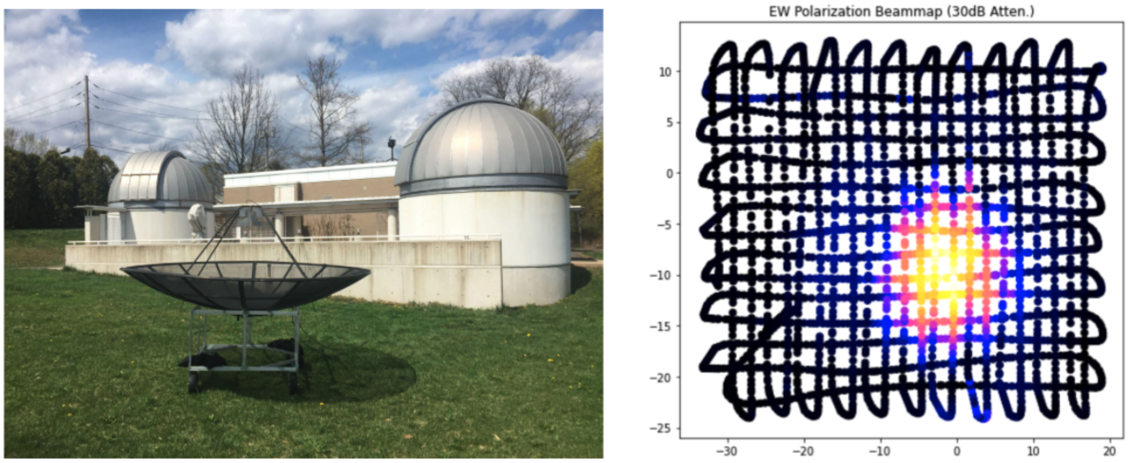

Figure 5.18: The drone calibration testbed telescope at Yale's Leitner Family Observatory and Planetarium (LFOP) (left) and a preliminary 2D beam map made by the telescope (right). The axes on the beam map are East/West and North/South position in meters, and the points on the scatter plot represent drone physical position in an XY plane in the telescope far field. The color scale is power measured at the telescope in logarithmic units, and it is evident even at this preliminary stage of analysis that there is a gaussian beam feature (as expected). This telescope enables the Yale drone group to test drone hardware upgrades locally, and has been used to prototype the RF switch and provide test data to practice removing pulses in the analysis prior to higher-stakes flights. Telescope design by Ry Walker and Will Tyndall, plot by Will Tyndall.

appropriate level for the correlator input. The signal is then digitized and channelized by the ICE board used for CHIME and HIRAX, and autocorrelation data is saved for duo-polarization map making.

The decision to have the antenna connect directly to LNR400 was made due to power considerations, as it seemed simpler to keep active components in the observatory building rather than mount them on the telescope. This design choice results in a high system temperature > 1000K, which prevents the telescope from performing cosmological measurements. To be an effective testbed, however, the telescope is only required to pick up bright signals from drone, so the design can be less flexible in, for instance, the location of amplification, and can afford to have a high T_{sys}. The high T_{sys} does mean that adjustments will have to be made to the drone signal attenuation and the data taking set-up when mapping a more sensitive instrument.

Several characterization measurements were performed on this telescope. S_{11} measurements of the antenna mounted on the telescope show < -10dB within the HIRAX band, indicating we have a usable antenna. Each element of the RF chain was independently characterized by a VNA, to ensure stable gain performance from the amplifiers, and appropriate filtering from the filters.

With the completed testbed, we were finally able to test the drone calibration system upgrades locally. It is on the LFOP testbed that we tested the RF switching in flight for the first time, which enabled us to implement a pulse removal algorithm before heading to the field. Early 2D beam maps illustrate the functionality of the telescope, revealing gaussian features as expected, and one is shown in fig. 5.18.

5.8 Drone flyaway

On April 14 2021, during a routine gridded autopilot test flight at Yale's Leitner Observatory (LFOP), the drone lost connection with the controller and drifted away over the city of New Haven. It ultimately took a hard landing in a tree two miles away, which required retrieval by a professional tree-climber. This section will recount the process of recovering the drone, diagnosing the failure modes, and rebuilding for the successful flights outlined in Ch. 6.

5.8.1 Drone loss and recovery

There were several unusual occurrences prior to the flyaway.

- Prior to takeoff, the drone had trouble connecting to RTK GPS, and required power cycling several times to connect. This was not alarming, as it had been a routine occurence during previously successful flights. The two devices ultimately found a link when the noise source was powered off. It was powered again for flights.

- The autopilot had trouble finding satellites and indicated this roughly 10 times. The autopilot also lost connection at least once during flight, requiring the flight plan to be re-upload while in the air. After the first re-upload, the drone orientation changed unexpectedly by 90 degrees (determined visually from ground observation). Right before the drone flew off, the autopilot disconnected and the flight plan would no longer upload.

Finally, the drone stopped responding to any UGCS directions and began drifting northwest. The team on the ground tried overriding the autopilot with the controller, but the drone would only respond to yaw/altitude commands and would not move in lateral directions. The drone flew two-miles to the base of a public park and parked itself in a tree.

It was later revealed that during the uncontrolled part of the flyaway flight, the drone had switched into “Attitude” mode, in which the GPS and other positioning systems are disabled, so flights must be done completely manually (in this mode, the drone cannot hover stably). Aside from intermittent periods of reconnection with the controller, during which attempts were made to pilot the drone back, the flight data indicated that the aircraft moved with the speed and direction of the wind that day.

The Yale drone team was able to obtain DJI GPS coordinates for a point of last connection, roughly 1 mile from LFOP. We went by car to this site and looked for signs of the drone, ultimately finding a weak GPS connection with the uBlox on board the drone. From this connection, we were able to retrieve the drone’s coordinates, which were over a mile from that location.

The drone was found in a tree on the outskirts of West Rock Park in Hamden, CT. Landing gear was deployed, suggesting the drone chose this spot to try and land (which would automatically happen at 10%battery, a status later corroborated in the flight data). A professional tree climber was able to retrieve the drone by hand within 24 hours. The newly-rescued drone had been rained on, had multiple damaged propellers, and two arms in need of replacement. All active components and batteries were placed in rice for several days and then successfully powered on. We could then retrieve the drone data and build a more comprehensive picture of what went wrong.

5.8.2 Failure analysis

A series of tests revealed two major failure modes in the prior drone set up:

- The uBlox antenna had a magnet embedded in it, which caused problems for the drone internal IMUs (angular sensors). The IMUS rely on magnetic sensors for orientation, and the presence of a magnet rendered them unable to determine heading or pitch to change flight direction.
- A loose SMA connection was likely present between noise source and filter, which would have radiate enough power to completely jam the DJI GPS. In this configuration, the uBlox GPS will drop off in performance but retain more satellites than the DJI GPS, explaining why we were able to successfully link to the uBlox for the drone recovery.

We completed a series of diagnostic tests as part of the 4/14 flyaway investigation. These tests recreated possible GPS and noise source conditions from 4/14 and used real-time monitoring through the DJI Assistant software to understand how different conditions impact the number of satellites seen by the GPS, track the IMU errors, and generally monitor drone internal sensor performance. The tests are as follows:

- Test 1: Drone on its own. We monitored satellite count for the different drone GPS antennas while we individually shielded them with mesh. This dropped satellite counts from 16-18 down to ∼8. This test was intended to confirm our intuition for GPS operations and verify the monitoring was sensible.
- Test 2: Drone + uBlox. We monitored satellite count and IMU errors while incorporating the uBlox GPS in its typical flight configuration (mounted above

the RTK box). We found that satellite counts do not change in the presence of the GPS, but that the IMU performance is completely degraded. This held when the uBlox was both on and off, and this error was eventually traced to an embedded magnet.

- Test 3: Drone + transmitting antenna. We observe satellite count for different transmission from the drone transmitter antenna, varying the filtering, attenuation levels, and connector tightness. We find that an unfiltered signal from the noise source completely jams the GPS, a properly filtered signal does not, and a loose connector before a filtered signal radiates enough power to jam the GPS as if a filter were not present.
- Test 4: Drone + transmitting antenna + UBlox. We repeated the loose connector tests done in Test 3 with the UBlox GPS mounted on the drone. We find that the UBlox GPS has a higher threshold for jamming, and continues to see 12 satellites from its old mounting spot under conditions in which the DJI GPS sees 0. Locating the UBlox GPS closer to the transmitter will drop the satellite count to $\sim$8.

5.8.3 The rebuild

The drone structure was rebuilt without issue. As part of the rebuild, we modified the on-board electronics to eliminate the fail points from the flyaway. The magnet was removed from on-board GPS through heat treatments, and the electronics plate was redesigned to discourage loosening the noise source-to-filter connection (which was properly wrench tightened).

The on-board hardware was reintegrated slowly, and the team built back up to drone flights through a series of grounded tests. We power cycled the drone on the ground many times, looking for stable GPS connection, and monitoring numbers of satellites and compass errors. We also spun propellers up without initiating liftoff to build confidence in the reconstruction, and performed firmware updates. Once multiple on-ground tests had been successfully passed, we resumed flight operations without any complications.

With this rebuilt drone system, we were able to make the first maps of a 400-800MHz HIRAX-style telescope.

5.9 Chapter summary

This chapter outlined the construction and characterization of a drone calibration system for 21cm telescope beam mapping. It is important for 21cm experiments such as HIRAX to understand the frequency dependence of their beams to 1 part in 100000, making calibration a critical piece of their realization.

The drone system described in this chapter consists of a stable broadband noise source, a GPS-locked timing board, an omnidirectional transmitter antenna, and a precision differential GPS flown on a hexacopter drone. It has iterated through multiple designs, and has been characterized on multiple far field antenna ranges and in an anechoic chamber. Implementation of the calibration system will be described in the coming chapter.

Chapter 6

Drone calibration measurements of radio telescopes

This chapter steps through drone calibration of three different instruments: DSA-110 at Owen's Valley Radio Observatory, BMX at Brookhaven National Laboratory, and the CHIME Outrigger Prototype at Green Bank Radio Observatory. The work in this chapter is shared chronologically, so the reader will see a progression in flight strategy, data quality, and beam mapping results as the calibrator system and analysis procedure are refined: starting from 1D cuts through a main beam and working towards detailed 2D polarized maps. The chapter also touches on the drone data analysis procedure, and describes ongoing work and next steps.

The beam mapping campaigns discussed in this Chapter were instrumental in informing the drone design, which was iterated over several years to converge on the instrument outlined in Chapter 5. Thus, only the Green Bank Observatory section uses the 'final' drone calibrator design, with the predecessor flights at Owen's Valley and Brookhaven National Lab using earlier instances of the hardware. The hardware details, and reasons for upgrade as informed by flight data, will be described in the relevant sections.

6.1 Drone beam mapping data analysis

We have developed an analysis pipeline that can quickly produce 2D beam maps from flight and telescope data. Here we describe the most recent version, used for flights at Green Bank Observatory.

The beam map has two main ingredients: autocorrelation data from the telescope and GPS/IMU data from the drone. Separate processing is done for the two datasets in individualized classes, before they are combined into a beam map, and all analysis is written in Python.

Telescope data

We begin analysis with raw telescope data, saving the auto correlations in all channels, all times and all frequencies (we do not process the cross-correlation data for now, though do save it for later explorations of array effects). The telescope data is read in as a series of HDF5 files, with the time/frequency/visibility axes populating a 3D numpy array. The telescope data is gain corrected (when applicable), and a channel map is applied that organizes the autocorrelation data from different channels into a sensible format (i.e. have indices 0,1 correspond to polarizations 0,1 on dish 0; have indices 2,3 correspond to polarizations 0,1 on dish 1; etc.).

The raw telescope data is sampled with time resolution on the order of 10ms. For gridded drone flights, the power measured from one co-pol channel on the telescope, viewed for one frequency as a function of time, will include several Gaussian features (up to 22, from the 22 drone transits over the telescope), corresponding to 1D cuts through the beam. The peaks of the individual Gaussians will themselves trace out a Gaussian, which corresponds to a 1D beam cut in the direction orthogonal to flight. These features are depicted in the telescope data panel of figure 6.1.

For the recent flights at Green Bank, the raw data included a square wave feature from the noise source pulsing (implemented to allow for sky background/RFI removal). Therefore, the next analysis step is to fit a square wave of the appropriate period to the raw telescope data. We add a timing offsets to the square wave and use scipy's pearsonr function to calculate the Pearson-R correlation coefficients r between the square wave (x, with mean m_x) and data (y, with mean m_y) for each offset,

$$r = \frac{\Sigma(x - m_x)\Sigma(y - m_y)}{\sqrt{\Sigma(x - m_x)^2 \Sigma(y - m_y)^2}} \tag{6.1}$$

to find the offset where the correlation is maximized. This functionality was encorporated by grad student Will Tyndall. Once the square wave is properly fit to the data, we can tag points when the source is on, off, and transitioning. We throw away the transition points, and subtract the neighboring off-points from the on-points to obtain background-subtracted data. This completes the telescope data processing.

Drone data

For successful beam maps, we need to know the drone position, orientation, and tipping/tilting angle for all timestamps, which come from the drone GPS/IMU data. We read in this data as a CSV (more details on conversion from DJI proprietary format to CSV in Ch. 5). The drone data class takes inputs of the telescope coordinates to transform latitude and longitude coordinates from the drone GPS into positions in meters from some reference point and angle on the sky from the telescope zenith. We plan to use both the DJI and UBlox RTK GPS systems in our analysis, though the reported results in this chapter utilize only the DJI GPS data. The DJI GPS samples with 100ms time resolution, though the timestamps we utilize only report with second-level precision, so some interpolation is required.

Making a beam map

Once the drone and telescope datasets have been separately read in and processed, the drone timestamps are interpolated to the higher-cadence telescope timestamps, and the data is combined in a Python class. A timing offset is applied to the drone timestamps at this stage, if found to be necessary (procedure described later in this chapter, in section 6.3.4). The analysis procedure is shown with example data in figure 6.1.

Once the drone and telescope data have been combined, we can examine the telescope measured power as a function of drone position. As a standard first analysis step, we perform Gaussian fits on the data sets and extract relevant beam parameters.

The 2D Gaussian we fit for has parameters $[A, xo, yo, \sigma_x, \sigma_y, \theta, D]$ and is of the form,

$$F(x, y) = A \exp\left[-a(x-xo)^2 - 2b(x-xo)(y-yo) - c(y-yo)^2\right] + D \tag{6.2}$$

where A is an amplitude factor (we do not worry about other parameters within A because the beam maps will be normalized), D is an offset, xo and yo are centroid locations, and:

$$\begin{aligned} a &= \frac{\cos^2\theta}{2\sigma_x^2} + \frac{\sin^2(\theta)}{2\sigma_y^2} \\ b &= -\frac{\sin(2\theta)}{4\sigma_x^2} + \frac{\sin(2\theta)}{4\sigma_y^2} \\ c &= \frac{\sin^2\theta}{2\sigma_x^2} + \frac{\cos^2\theta}{2\sigma_y^2} \end{aligned} \tag{6.3}$$

where, σ_x and σ_y represent beam width, and θ is the rotation of the ellipsoid.

We can relate σ_x and σ_y to beam full-width-half maximum (FWHM) using,

$$\mathrm{FWHM} \approx 2.355\sigma. \tag{6.4}$$

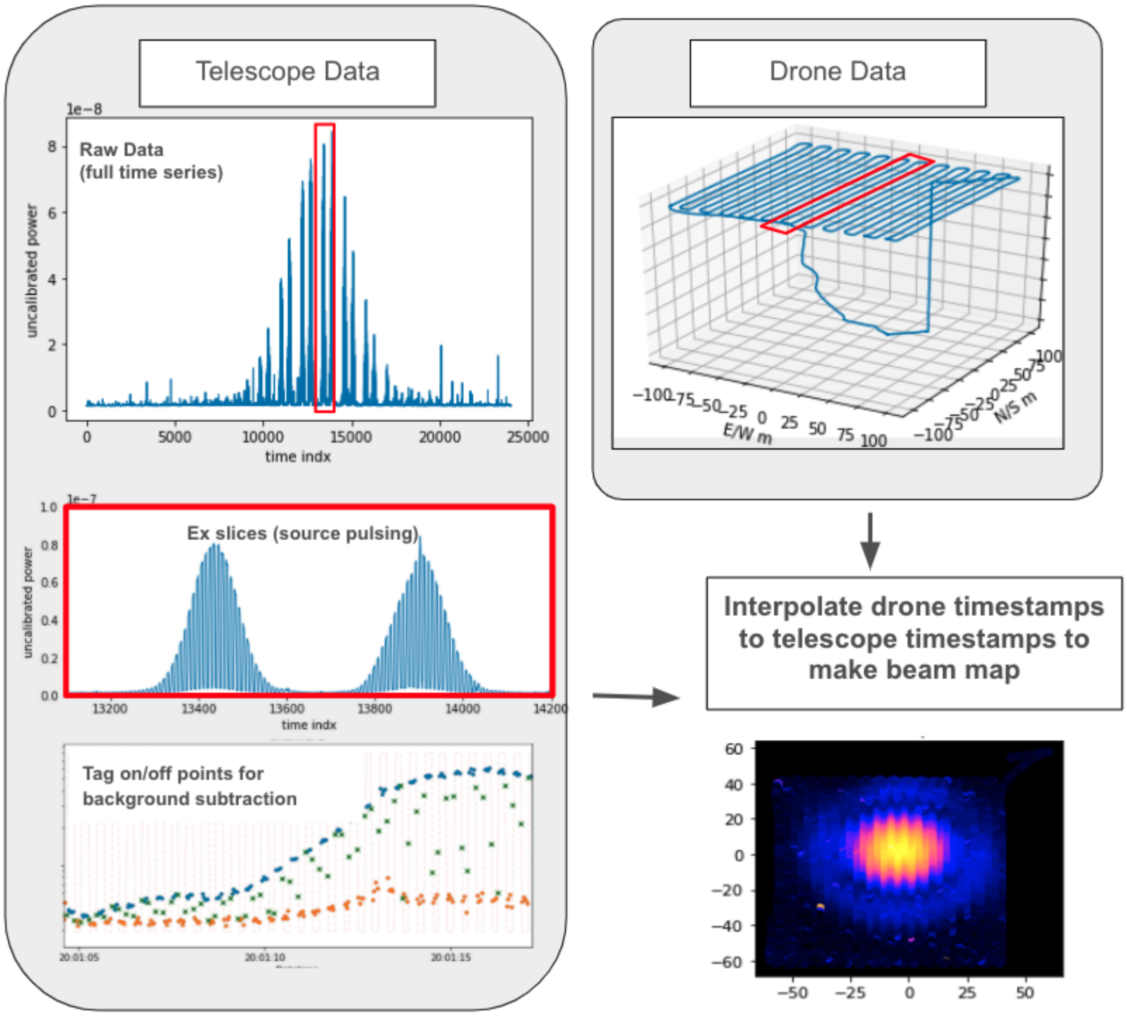
Telescope Data
1e−8
Raw Data
(full time series)
uncalibrated power
time indx
1e−7
Ex slices (source pulsing)
uncalibrated power
time indx
Tag on/off points for
background subtraction
20:01:05
20:01:10
20:01:15
Drone Data
E/W m
N/S m
Interpolate drone timestamps
to telescope timestamps to
make beam map

Figure 6.1: The ingredients for a beam map are telescope autocorrelation data and drone positional data. Here we show the raw data sets from both the telescope and drone. The dataset includes a pulsing noise source, which enables background subtraction at each drone position, but must be fit for in the telescope data analysis. The red box in the figure is used to illustrate a common time slice for the telescope and drone data.

The beam pattern of the drone transmitter, which has been characterized (as described in Chapter 3) will be factored into the analysis.

Variants on this analysis scheme were used for maps at Owen's Valley Radio Observatory, Brookhaven National Lab, and Green Bank. With an understanding of the analysis procedure, we can now walk through flight data results.

6.2 Owen's Valley Radio Observatory

In October 2019, the Yale drone team visited Owens' Valley Radio Observatory (OVRO) in Bishop, CA to measure the DSA-110 telescope[42][1]. This trip marked our first drone flights over a telescope, and results were key to informing the currently implemented hardware designs. This section will discuss the flight plans, results, and lessons learned from our first-ever beam mapping flights at DSA.

6.2.1 The DSA instrument

The Deep Synoptic Array (DSA) is a state of the art radio camera under development at Owens Valley Radio Observatory in California. Once built, DSA will consist of 2000 steerable five meter parabolic dishes spanning a frequency range of 700MHz-2GHz. It will will be an important survey instrument, increasing the number of known radio sources by a factor of 1000, will be a piece of the NANOGrav gravitational wave detector, and will continue to catalog and localize FRBs [42]. DSA-2000 is the third iteration of the DSA concept, following the successful DSA-10[108] and DSA-110 FRB detectors that consisted of 10 and 110 4.65m dishes respectively and cover a frequency range 1280-1530MHz. For our first beam mapping campaign, we mapped a single DSA-110 dish, pictured in figure 6.2.

DSA-110 operates in a higher frequency band than HIRAX, which required a small adjustment to the noise source filtering, and has a higher F/D, which results in different expected beam features. The instrument was still able to provide a useful test bed for drone calibrator development. Adapting HIRAX-focused calibration instrumentation to the DSA was little added effort, which shows the potential for expanding the drone system to be a broadband calibrator for all 21cm.

---

[1] We acknowledge the OVRO staff, Caltech radio group for being incredibly generous, accommodating and helpful during our OVRO visit, in particular Vikram Ravi, Nitika Yadlapalli, and Gregg Hallinan.

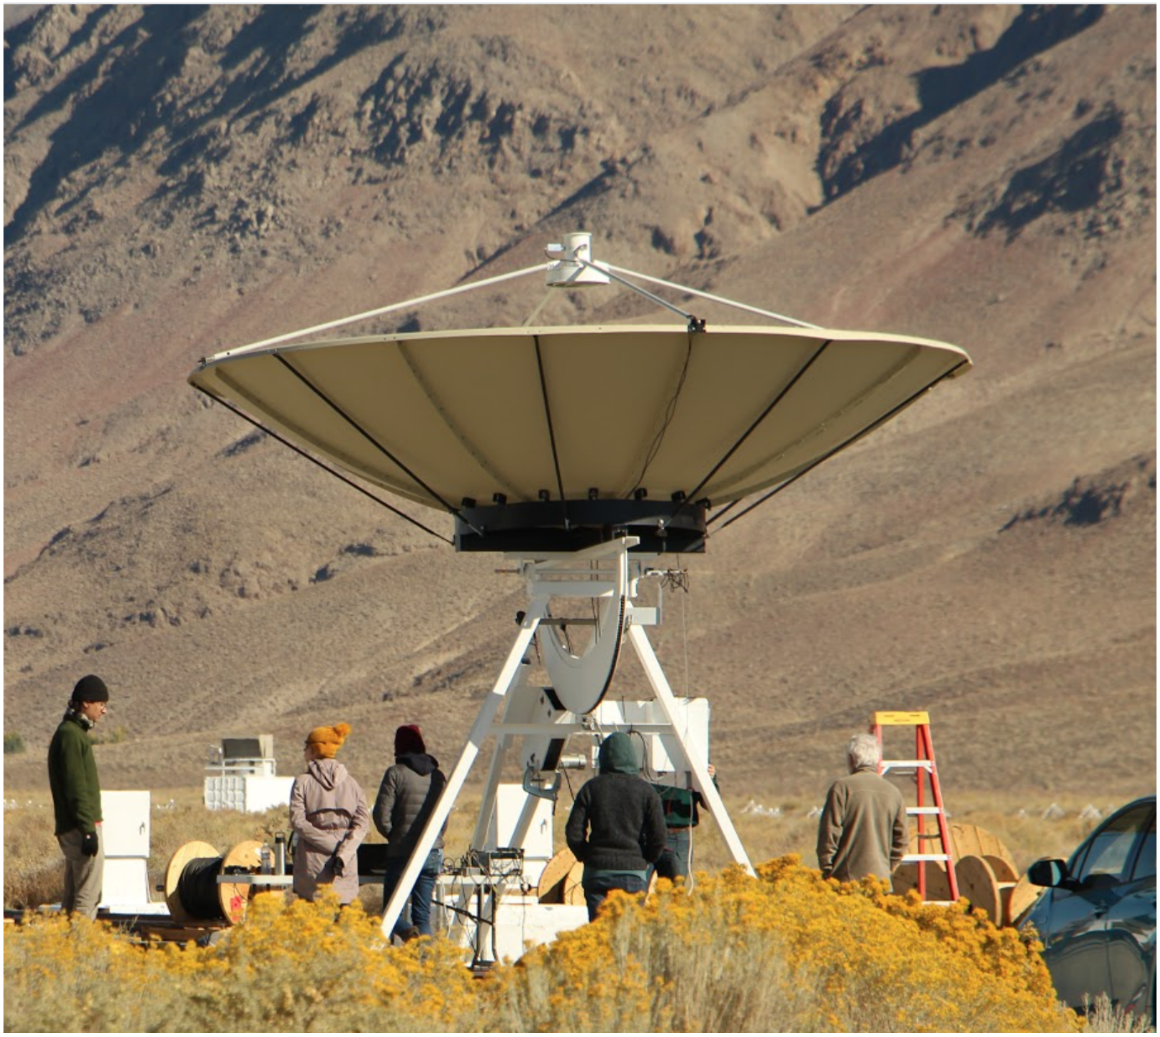

Figure 6.2: Photo of the DSA-110 dish mapped by the drone calibrator.

6.2.2 Flight path design and considerations

At the time of flights over DSA, we had not yet integrated an autopilot into the drone calibrator. Thus, the flight path was limited to two orthogonal 1D beam cuts, which could be performed with manual flight control. We were able to locate the dish center and mark it in the flight control app, then fly orthogonal North/South and East/West throws over the mark. We fly this cross in a plane at a constant height above ground level (h_{AGL}) in the telescope far-field,

$$h_{\text{AGL}} = \frac{2D^2}{\lambda} = \frac{2 \times (4.65m)^2}{0.2m} = 207\text{m}, \tag{6.5}$$

transmitting power from a 50dB ENR calibrated noise source through a biconical antenna. The path is flown twice to measure the two polarizations.

The flight path in 2D is shown in figure 6.3, with the points on the plot representing drone position and the colorscale representing power measured at the dish. We see the measured power peak when the drone is over the telescope center, and then fall off as the drone finishes the pass.

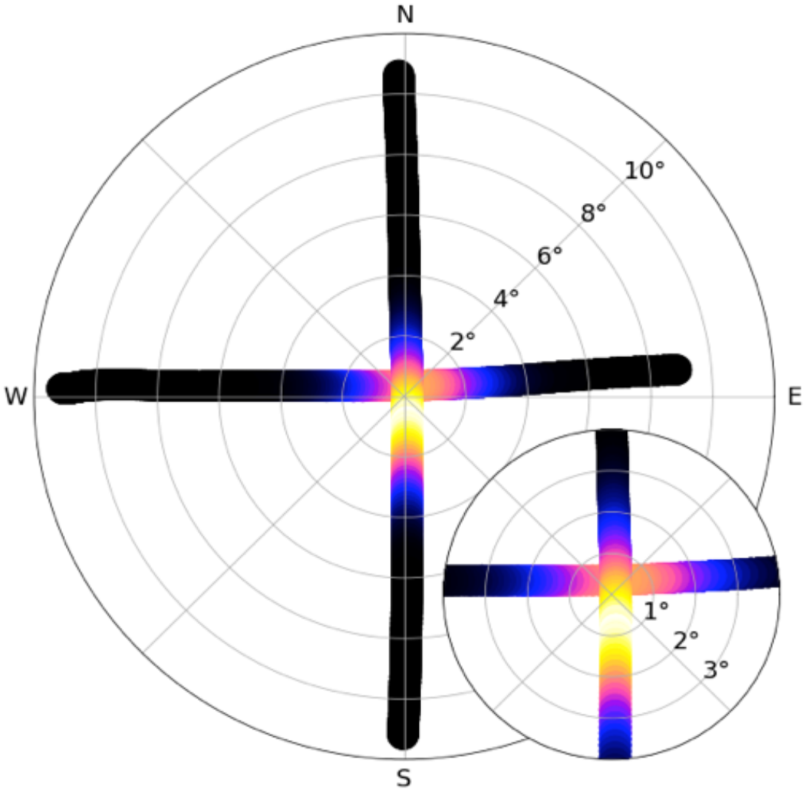

Figure 6.3: Drone flight path over DSA. Points on the plot representing drone position, and the colorscale representing power at the dish. The drone flew 100m throws in North/South and East/West, flying in a 2D plane in the telescope far field. Plot by Will Tyndall.

6.2.3 Beam mapping results

We combine telescope auto-correlation data and drone positional data to understand measured power as a function of drone location on the sky. We isolate individual passes over the telescope, and plot measured power as a function of zenith angle, finding the expected Gaussian beam shape. Data and fits from one of the drone transits at a variety of frequencies are shown in figure 6.4. We note that a Gaussian fits the data effectively, which is an encouraging check on our methods. We also observe that the Gaussian peak amplitude varies as a function of frequency by over 6dB. This amplitude profile is qualitatively similar to that of the noise source used for these flights (later upgraded), so we deduce it is not an artifact of the DSA instrument or the transmitter antenna (figure 6.4). We desire a flatter frequency response, sloped to within 1dB. As described in Chapter 5, with a sufficiently sloped spectrum we cannot see low-level beam features across the full band at once—the high-power part of the band will digitally saturate while the low-power section will be lacking sensitivity.

From the 1D Gaussian fits, we pull out beam parameters and compare to expected results (Figure 6.5). The full-width-half-max of the beam is consistent with old DSA-10 published numbers (3.5 degrees) though wider than expected for the DSA-110 design (3.1 degrees)[108]. The full-width-half-max profile in frequency follows the expected $1.22\lambda/D$ slope for a circular aperture telescope, though exhibits a ripple with a phase offset between polarizations. This same behavior is observed in DSA transit data, and similar behavior is seen by the CHIME instrument in holography measurements[109], associating the ripple with a bounce between the feed and the bottom of the dish.

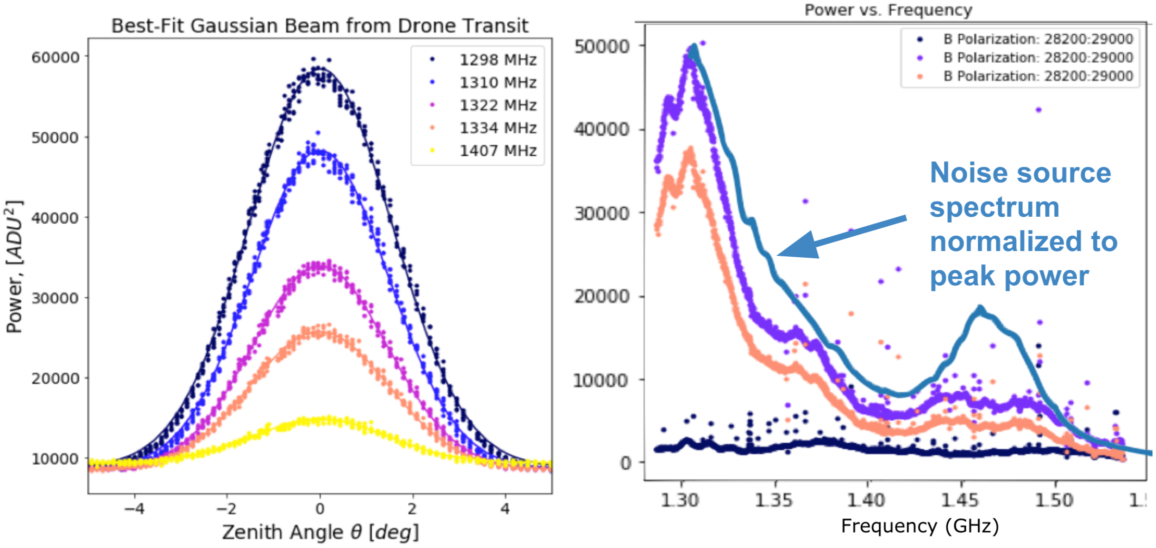

Figure 6.4: Left:Gaussian fits at 5 different frequencies for a single drone transit through the DSA beam center. Right: Gaussian peak amplitude (arbitrary linear units) as a function of frequency, showing a profile that qualitatively tracks the noise source spectrum used for DSA flights (and later upgraded). . Plots by Will Tyndall; noise source annotations by Emily Kuhn.

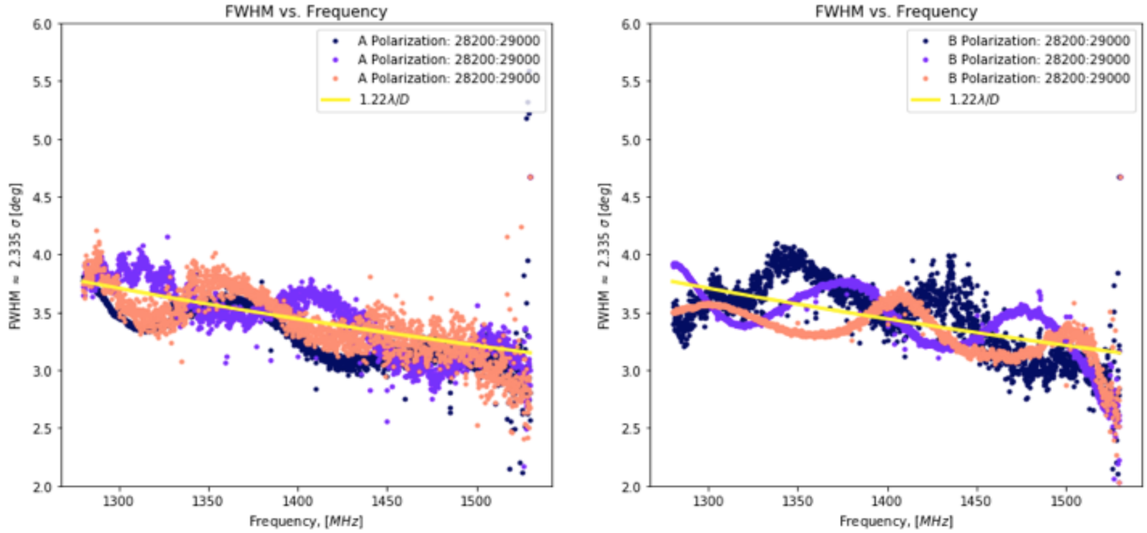

Figure 6.5: Full-width-half-max as a function of frequency for Gaussian fits of DSA beam data. The different curves on each plot represent different drone transits through the beam, and the two plots are for each of the two polarizations. The rippling shape is also observed in transit data, and postulated to be associated with a bounce between the feed and bottom of the dish. Plot by Will Tyndall.

6.2.4 Challenges and lessons learned

The DSA beam results were the first data obtained with the Yale drone system. As such, the flights were highly instructive and results encouraging. There were several successes in this beam mapping campaign, most notably that we delivered a first beam map, which corroborated beam features expected from models and observed from source transits. We also set up the drone for travel and completed GPS system verification measurements at the site, which are discussed in Chapter 5.

We left OVRO with a clear path forward for hardware upgrades. The noise source used for mapping DSA had a sloped transmission in frequency, and was not sufficiently bright to capture low-level side lobe features. We thus learned the importance of a flat, bright noise source, and upgraded to a more suitable hardware for future flights (described in section 5.3.4).

After successfully measuring 1D beam slices, the next step was to measure a full 2D beam shape. The extra coverage is critical for ensuring we capture the beam centroid which crosses in 1D can feasibly miss (resulting in artificially narrow Gaussians). Additionally, we would like to measure cross-polarization maps, which depend sensitively on beam location– there is a characteristic dip in the cross-pol beam center that can be easily missed without 2D coverage. Furthermore, 2D maps give access to off-axis beam features and allow assessment of asymmetries.

Autopilot and noise source upgrades were implemented between the October 2019 flights at OVRO and the next flight set, in March 2020 at Brookhaven National Laboratory.

6.3 Brookhaven National Laboratory

In March of 2020, we visited Brookhaven National Laboratory in Long Island, New York to measure the BMX telescope[2]. This was another incredibly useful set of measurements, as we were able to deploy an upgraded system after DSA that included a new noise source and autopilot, which are described in Chapter 5. The new noise source has few gain variations across a broadband and is sufficiently bright to measure sidelobes (95dB ENR). The autopilot allows for 2D gridded flights.

This section describes: the BMX instrument; the flight path design and measurement scheme used for 2D mapping; the beam maps we generated with this campaign, which include main beam and side lobe results; and lessons learned.

6.3.1 The Baryon Mapping eXperiment (BMX)

The Baryon Mapping eXperiment (BMX)[110] is a telescope based at Brookhaven National Laboratory in Long Island, NY. It is a pathfinder for future intensity mapping experiments, and in particular is developing technology for the upcoming PUMA array[111]. BMX is comprised of four 4m off-axis dishes, which utilize dual-polarization pyramidal horn antennas as receivers, and operates in a frequency range 1100-1500MHz. A photo of the instrument is shown in figure 6.6.

6.3.2 Flight path design and considerations

Our main goal for the BMX flight set was to produce 2D telescope beam maps in two polarizations that included both main beam features and sidelobes. To accomplish this map, we decided to do two different types of flights:

1. Flights transmitting higher power (∼60dB ENR), which will capture side-lobe

---

[2] We acknowledge Anže Slosar, Justine Haupt, Paul O'Connor, & the BNL Cosmology group, who hosted us for a successful and informative visit to the BMX telescope.

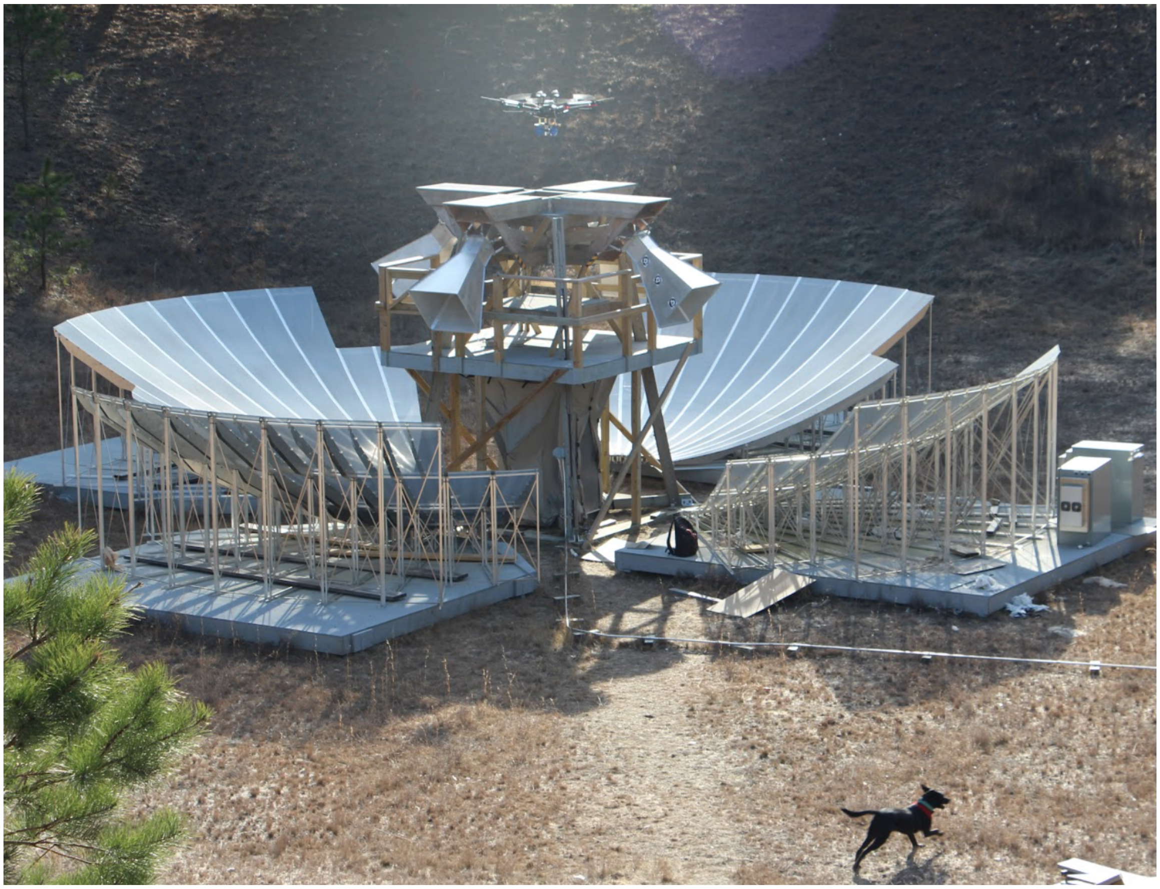

Figure 6.6: The Baryon Mapping eXperiment (BMX). BMX is comprised of four 4m off-axis dishes, and operates in a frequency range 1100-1500MHz. The drone system made 2D beam maps of the four BMX dishes in March 2020.

features but saturate the main beam (the result will look like a flat-topped Gaussian)

2. Flights transmitting lower power (∼47dB ENR), which give resolution in the main beam but do not illuminate side lobes above the background.

We can then match amplitude levels between the two styles of flight to build a complete map. To ensure beam coverage and for ease of autopilot programming, we fly the drone in a planar grid as described in Chapter 5.

There are a few additional constraints on our mapping scheme:

- We must fly in the telescope far field (which is 150m at BMX)

- We are limited to 25 minute flight times due to drone battery capacity and our payload weight
- We aim for < 1 degree main beam resolution and < 3 degree side lobe resolution
- For later mapping campaigns, we will add the additional constraint of measure out to sufficiently wide angles that we capture beam features with amplitudes above -40dB (relative to the peak), as determined from simulations.

All of these factors will dictate the number passes in the grid, the length of each pass, the distance between neighboring passes, the drone flight speed (at later points, we will consider how flight speed impacts drone tipping angles), and the total number of flights to complete a map.

At BMX, we flew two high-power and two lower-power flights at each polarization, giving a total of eight flights required for a full polarized beam map. The 2D flight data showing full coverage for a single polarization is shown in figure 6.7. This figure also shows an example of the data extracted from each flight, as the 2D flight paths are plotted with a color scale corresponding to measured power at the telescope (in logarithmic units, uncalibrated). For the first two flights, we aim to measure side lobes, so we cover an expansive area and transmit down high power. In the first flight (left most plot), we fly North/South passes covering a 120×120 meter area ($\sim 20 \times 20$ degrees). This flight takes the full 25min flight time and drains the drone batteries. After batteries are changed, we fly back up to far field and continue with East/West passes, again covering a 120×120 meter area (plot second from left) and completing the grid from flight 1. The third and fourth plot are the flight paths done at lower transmitted power, where we aim to measure the main beam. We fly denser passes for the main beam flights, and when taken together they cover a 60×60 meter area. The rightmost plot shows the total coverage from all four flights that make up one polarized map.

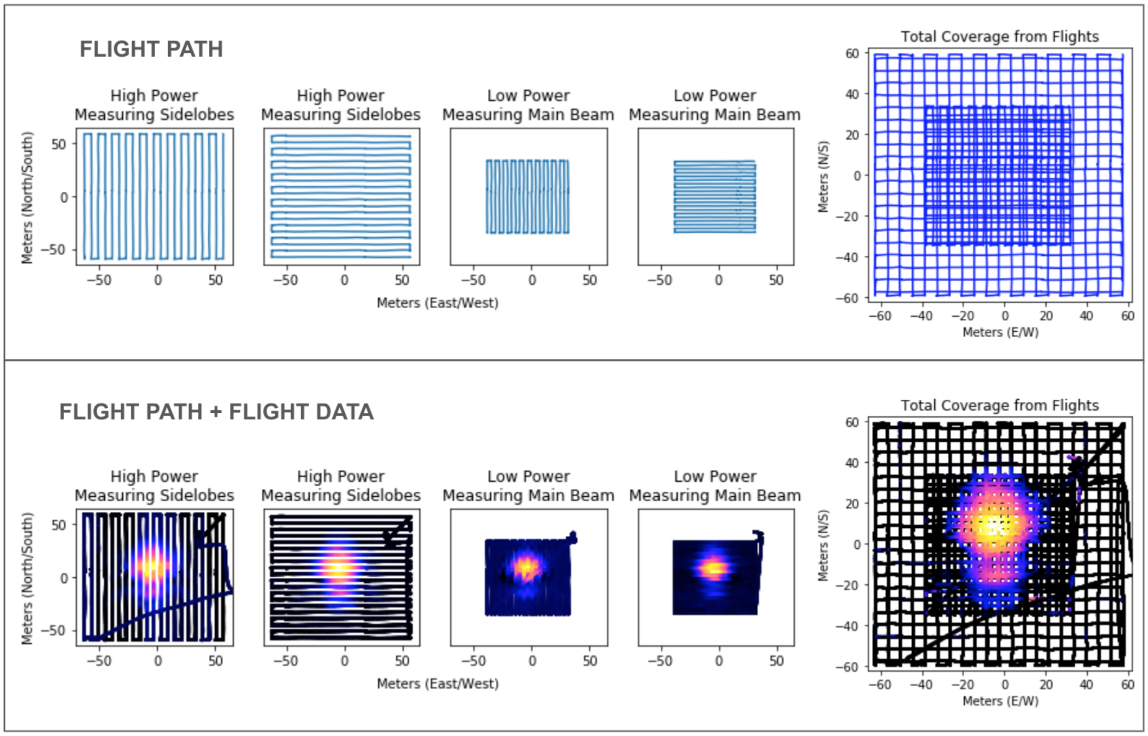

Figure 6.7: The BMX drone flight path and measurement scheme, showing the four individual flights that make up a single polarization map (eight total flights for full mapping). The top row is the flight data for the 2D grids flown by the drone in the telescope far field. The bottom row is the flight data overlaid with a color scale associated with measured power at the telescope (log units), which gives a sense for the beam information gathered in each flight.

6.3.3 Beam mapping results

We made preliminary 2D beam maps for the full range of BMX frequencies, for all four dishes and two polarizations. Results for all dishes at one polarization at 1230MHz are shown in figure 6.8, where the origin of the coordinate system is the telescope center. The larger grid is incomplete in Polarization 2 due to running out of time to complete flights.

We measure the telescope main beam and first sidelobes, observing that the sidelobe structure is not radially symmetric (i.e. we do not observe an airy pattern, which is the expectation for traditional circular aperture telescope), and rotates with polarization. These sorts of feature can arise if transmitter beam is narrower than

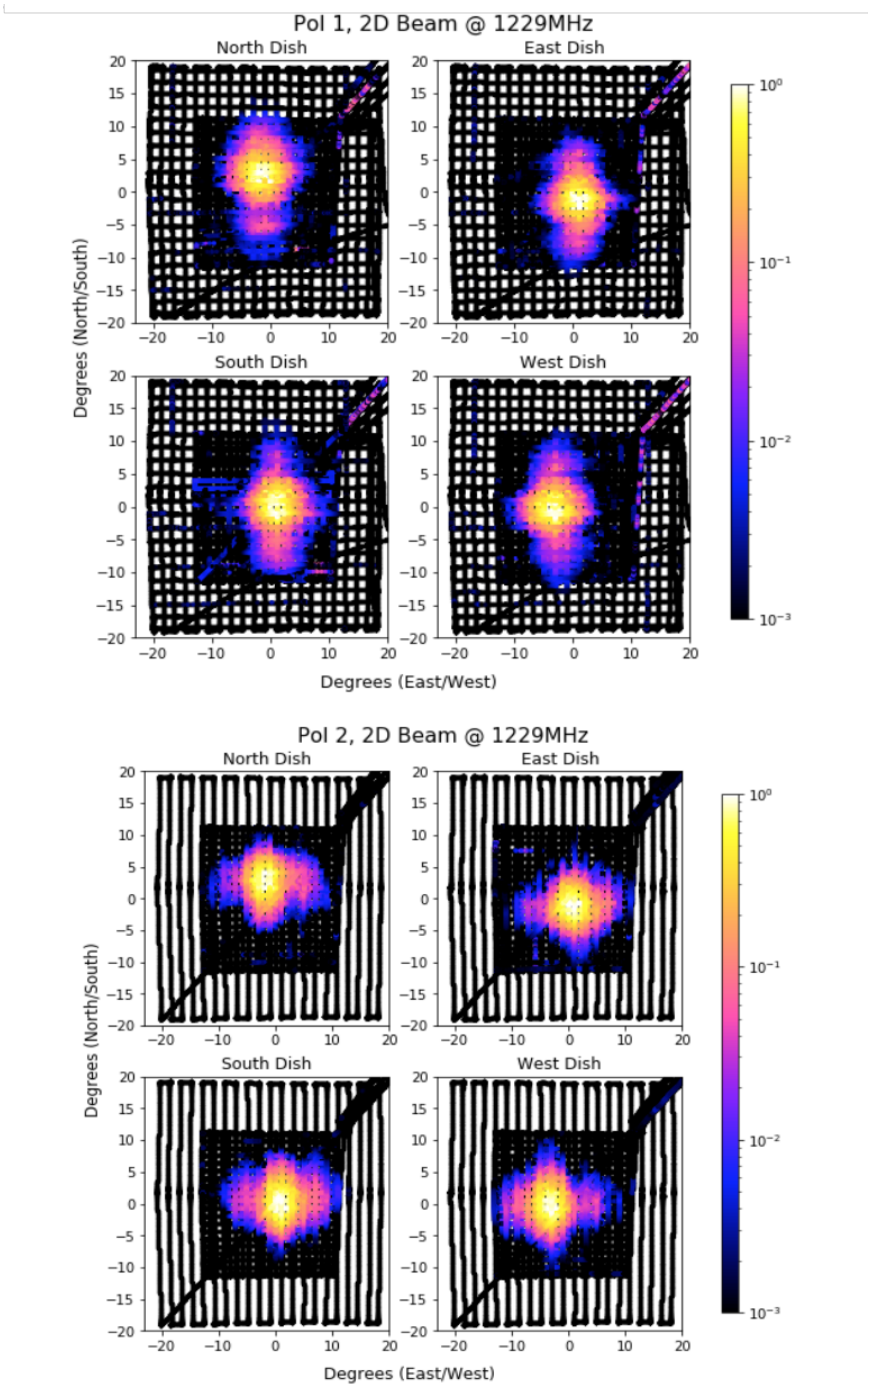

Figure 6.8: 2D beam maps for all four BMX dishes, with evidence of side-lobes and polarization-dependent features. Points represent drone position, overlapped for four different flights (fig. 6.7), and color scale is log power measured at the telescope. In these maps, data from each flight is normalized to the peak of its 2D Gaussian fit prior to getting overlaid. A more rigorous calibration is under development for future analysis frameworks.

the features it is trying to measure in one dimension, though we confirm this is not the case from range measurements of the transmitter system (Ch. 3). Instead, these features are likely due to the fact that the BMX design utilizes horn antennas, which themselves tend to have higher-level sidelobes in one dimension. These antenna features are reflected in the overall telescope beam, and may additionally be compounded with artifacts of the off-axis dish design, as off-axis dish design makes for more complex beam patterns than a traditional circular aperture. The single-axis, asymmetric sidelobes appear in instrument simulations, which serve as a validation of the drone measurements (figure 6.9).

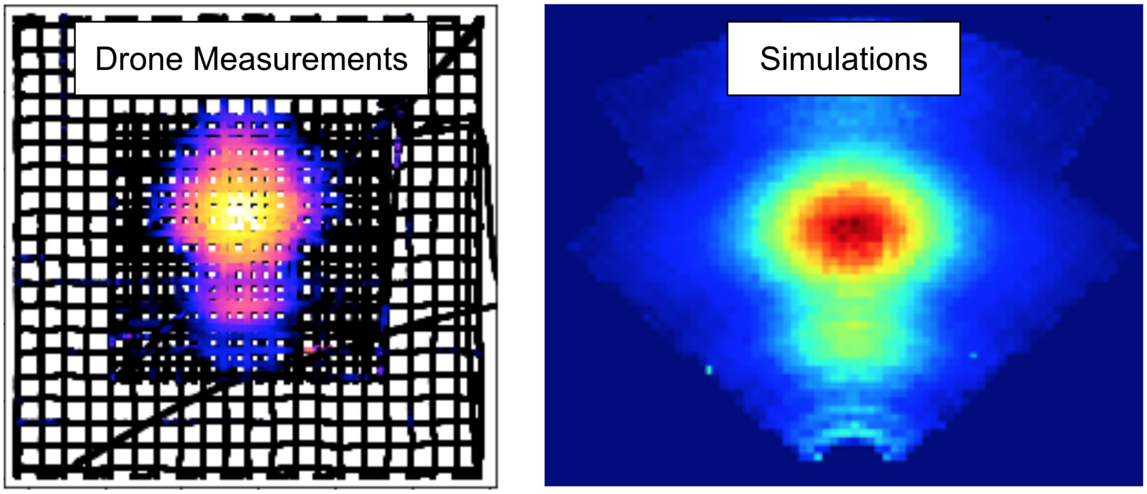

Figure 6.9: A single-dish BMX beam measurement (left) and simulation of the antenna beam projected onto the dish (right), which show qualitative agreement between beam shapes, including asymmetric sidelobes constrained to a single axis. Simulation by Chris Sheehy, from the BMX Collaboration.

The main beam was sufficiently Gaussian to perform 2D fits at all frequencies, from which a host of useful parameters were extracted, shown in figure 6.10. We verified that the telescope has a full-width-half-max of ∼5 degrees in one dimension and ∼4.5 in the other, as expected from modelling by the collaboration, and that the FWHM falls off with the expected $1.22\lambda/D$ in frequency. We also fit for the centroid locations of the four different dishes and compare results with the of each of the four dish locations to understand telescope pointing errors. The dish locations were found by flying the drone at low altitude over the telescope and hovering over each of the different dish centers (determined visually) for over one minute (the path of this flight is overlaid in blue in the right plot of figure 6.10). GPS coordinates of the four hover points were taken as the dish location. We observe in the flight data that the beam centroids do not align with the physical location of the dishes in the XY plane, and in fact are spread further outward, indicating the dishes are mis-pointed and shifted too far inwards.

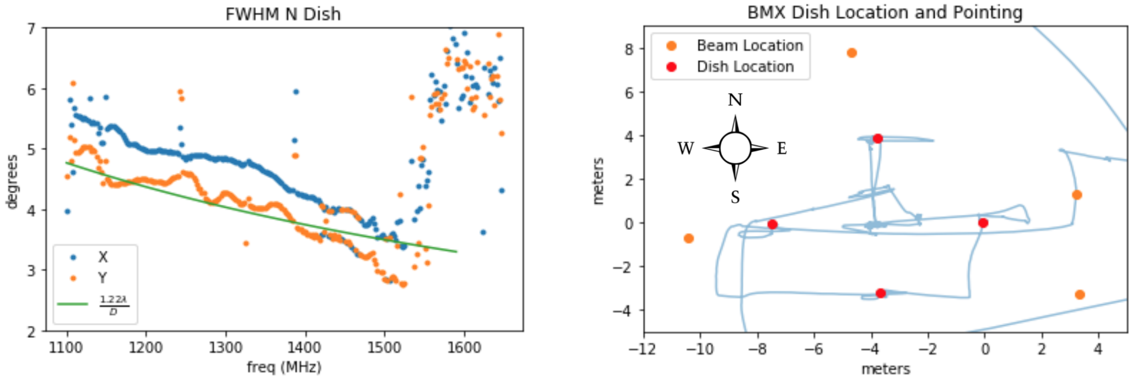

Figure 6.10: Left: Full-width-half-max(FWHM) from fits of one of the BMX beams. The FWHM gives the expected behavior in frequency (above 1500MHz, a filter is present in the analog chain, causing noise). Right: Centroid locations from fits of all four BMX beams (orange dots) superimposed with dish locations (red dots) found from flight data. The centroid fits reveal the dishes to be shifted too far inwards, with the South dish additionally mispointed.

The BMX analysis yielded the first 2D beam maps with our drone calibrator, and gave important instrument feedback to the BMX collaboration. This work also provided key feedback on the drone analysis and hardware that was critical to informing upgrades for future flights.

6.3.4 Challenges and lessons learned

There were two main challenges to contend with in the analysis, that informed updates for future flights: Timing offsets between the drone GPS clock and the telescope clock, and waypoints in the autopilot flight plan that caused the drone to stop mid-pass over the telescope.

Timing offsets

In the raw data, there are clear timing offsets between the drone GPS timestamps and the telescope timestamps. These timing offsets range up to 0.8 seconds, depending on the flight data set. This was initially discovered by looking at beam maps, where it was apparent that the 1D Gaussian centroids for successive drone transits through the beam are out of alignment with one another (visuals in the left plot of figure 6.11). This would occur if drone timestamps and telescope timestamps are out of sync.

These timing offsets can be fit for in the following way (demonstrated in figure 6.11). We apply timing corrections to the drone data, adding values between -1 and 1s in 0.1s increments to the GPS timestamp. Beam maps are generated for each of these timing offsets, and a 2D Gaussian is fit at each iteration. From the Gaussian fits, a series of fit parameters can be extracted as well as the covariance matrix for each parameter, which gives information about parameter uncertainty and therefore the quality of the fit. These uncertainties can be plotted as a function of timing offset, where the minima should give the correct timing offset. This procedure should be done with co-pol data, and should still be effective even in higher power flights with flat-topped Gaussians. We note that the data shown for BMX flights was all time-corrected.

The potential for timing inconsistencies motivated us to integrate a second RTK GPS into the drone hardware, which provides a second check on the DJI drone timing as well as positional redundancy. Going through the timing-offset fitting exercises also provided an analysis infrastructure for handling timing discrepancies, which proved useful for flight data on later beam mapping campaigns.

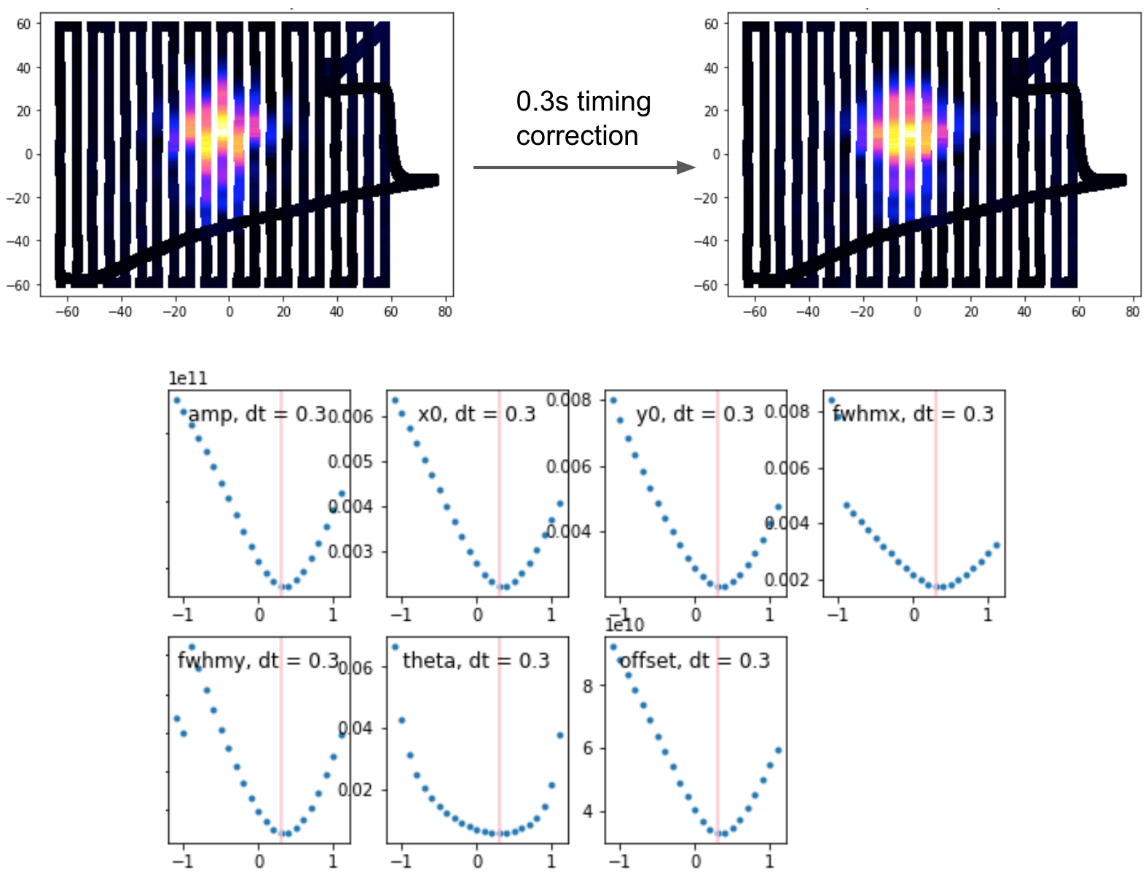

Figure 6.11: Top row: Example of a raw and corrected beam map for a flight with a 0.3s timing offset. Centroids in successive passes are notably out of alignment in the raw data, and the beam circularizes into the expected Gaussian shape once an offset is applied. Bottom row: Uncertainties from 2D Gaussian fit parameters (amplitude, centroid location in 2D, fwhm in 2D, and background level) as a function of timing offset show a curve that is minimized for an offset of ∼0.3 seconds.

Waypoints in flight

For flight at BMX, we integrated an autopilot to enable 2D mapping. The autopilot was previously tested on small grids (20-50 meter side lengths) at Yale, where it delivered smooth flight control. However, at BMX we expanded our test grids to 120 meter passes. For these larger gridded flights, the autopilot introduced ‘waypoints’ in the center of each pass, causing the drone to stopped during abruptly over the telescope. This behavior introduced wild swings in the drone tipping angle (up to 20 degrees), as the aircraft tried to brake, which interfered with usable flight time and transmitter pointing.

The impacts of intermediate waypoints on the drone tipping angle is shown in figure 6.12. The left figure shows the drone flight path, and the right plot shows the drone tipping angle during four different passes as a function of distance along the pass. There is color-coding between the two plots as a visual aid, to show which tipping profiles correspond to which passes. The drone flies successive passes in opposing directions (i.e. flies from north to south, then moves over, then flies south to north) to trace the grid, and tips generally during flight. When the drone stops midway to brake, it tips considerably in the other direction, rocks to balance out, then tips up and proceeds. This rocking occurs over the beam center, as shown in the plots.

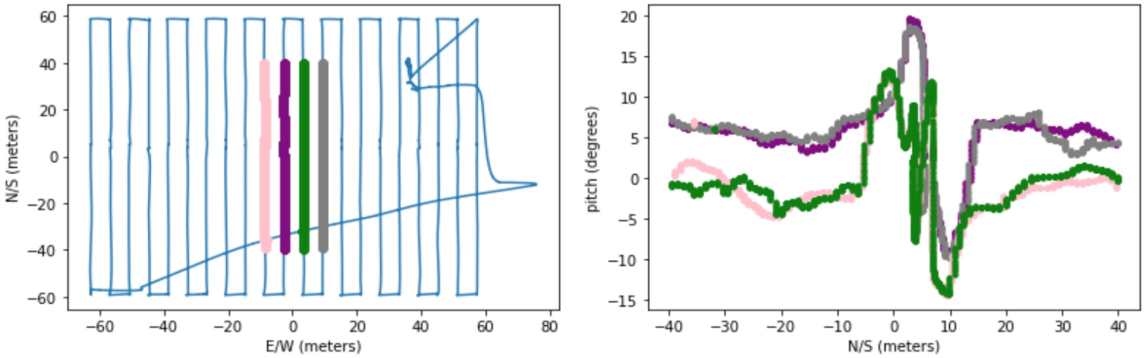

Figure 6.12: Intermediate waypoints in the autopilot flight path result in the drone stopping over the telescope center, which introduces pitch (tipping) angles of up to 20 degrees. The left plot shows the drone flight path, with four passes of focus color-coded, and the right plot shows the drone pitch during these four passes. Profiles are consistent for every other pass, when the drone is flying in the same direction.

The large swings of the drone cause swings in the transmitter pointing. Movement of the transmitter pointing introduces noise into the beam maps. For BMX analysis, these artifacts remain in the data. In principle, a sensitive measurement of the transmitter beam can be incorporated into the drone maps to correct for the impacts of drone tipping. This analysis pipeline has been built, and initial measurements of the transmitter beam have been taken (reported in Chapter 3), though not integrated into the BMX analysis. The flight path was corrected for subsequent flights through small tweaks in the software (changing the turn type to adaptive bank, for instance).

6.4 Green Bank Observatory

In August and October 2021, we visited Green Bank Observatory in West Virginia[3]. The focus for these trips was mapping our first 400-800MHz telescope and implementing the new hardware developed during a 1.5 year fieldwork hiatus (due to COVID-19 pandemic related travel restrictions). This was the first flight set to implement the noise source switching board

This section describes the CHIME FRB Outrigger prototype, which was the telescope measured on this expedition; shares instrument challenges and resolutions; reports early results from main beam and sidelobe flights, which showed excellent agreement with simulations; and explores polarized properties of the telescope, and what the results imply for the HIRAX feed design. This thesis is written during the early stages of the data analysis, so we focus on qualitative results.

6.4.1 The CHIME FRB Outrigger array

The CHIME Outrigger Prototype is designed to help CHIME localize FRBs by providing a >4000km baseline between British Columbia and West Virginia to resolve small angles on the sky. This telescope consists of eight 6m parabolic dishes, and operates in the 400-800MHz frequency band, providing a first test of drone calibration in the HIRAX band.

The individual dishes are instrumented with both CHIME and HIRAX feeds, which provides an interesting mode of comparison between similar feed models. Additionally, the dishes share a design with the first iteration HIRAX prototype, and the data-taking is performed using the same correlator hardware as HIRAX, making this telescope an excellent testbed for our calibration system before taking it to South Africa.

---

[3] We acknowledge Kevin Bandura, Pranav Sanghavi & GBO staff, and thank them for supporting our beam mapping campaign over the Chime Outrigger Prototype.

The Chime Outrigger Prototype array layout with the drone beam mapping set up, including the drone take off and RTK base station locations, is shown in figure 6.13. We chose the drone take-off location using a few different criteria: we desired visibility from the shade of the building; a flat surface for landing stability, which we affix marker tape to to ensure take off and landing in the same spot (to within <2cm, as a check on GPS positional drifts); and line of sight to the RTK base station. The RTK base station was initially located closer to the array, but was later moved further away (and to a location where the building blocked line-of-sight between RTK and telescope receiver) after its high-power transmission at 900MHz was compressing the telescope measurement chain.

We utilize a similar flight plan to the one for BMX, described in section 6.3.2, though this time fly 200×200m grids for the sidelobe flights at height 180m (the instrument far field). This allows beam coverage out to 30 degrees from zenith, which we expect from simulations to capture beam features above -30dB. We distribute efforts into main beam mapping and side lobe mapping.

As a preparatory step, we manually aligned the expected dish polarization axes with the cardinal directions. This alignment is standard for most telescopes, and in the case of drone calibration, allows for easy alignment between the transmitter polarization and the receiver polarization, as North/South or East/West flight heading alignments are simplest to program into the autopilot.

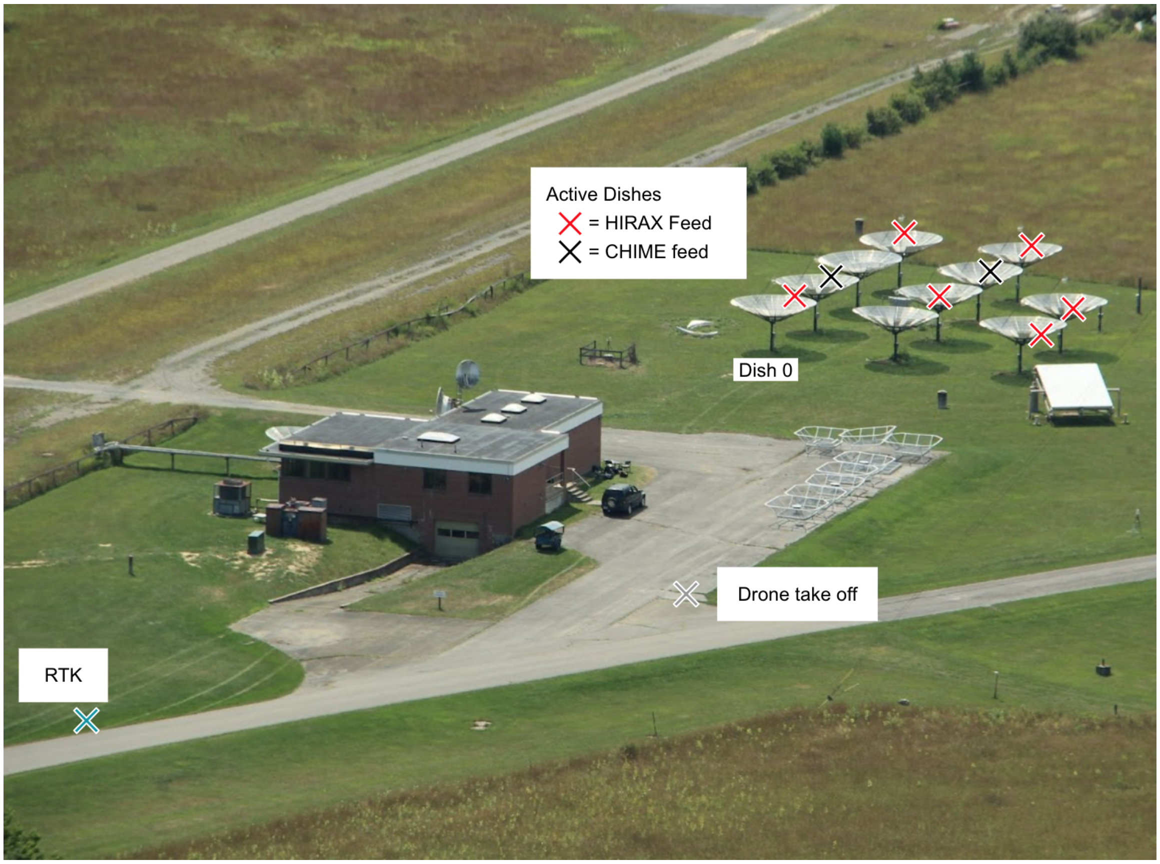

Figure 6.13: The CHIME Outrigger Prototype at Green Bank Radio Observatory. The 8 active dishes are marked with X's, color-coded by whether the dish is instrumented with a HIRAX or CHIME feed. The drone take off location, RTK base station location, and 'dish 0' location (which served as the coordinate system origin) are also marked.

6.4.2 Instrument challenges

There were a number of instrument challenges to contend with at Green Bank. They are as follows:

- Digital saturation: The HIRAX correlator is used for datataking for the CHIME Outrigger prototype dishes. When the correlator initiates datataking, it performs a gain calibration to set its digital range. If this gain solution is found in the typical manner, with the drone grounded, it will set digital sensitivity to an appropriate level for sky measurements. The digital range will then saturate when the drone transits overhead, as the noise source transmitted by the drone is substantially brighter than the sky. One way to overcome this is to instruct the correlator to iterate only once to set a gain "solution". This makes for a less sensitive instrument, which can work in our favor, as it allows us to broadcast more noise without saturating. In this configuration, the telescope is also less sensitive to the sun and other RFI, allowing for cleaner beam maps, with backgrounds that vary less in time.

- Analog Saturation: In our initial set up, the bright RTK transmission compressed the telescope analog electronics, leading to spurts of oscillating behavior. Additionally, harmonics were appearing in the measured spectrum at 670-700MHz (an alias of the 900-930MHz RTK transmission, depicted in figure 6.14). To solve this problem, we moved the RTK location further from dishes, placing it behind a building to obscure its line-of-sight path to the telescope receivers.

- Write speed: We were initially writing data to a standard hard drive, but realized that packets were dropping and timing issues were occurring. We later switched over to using solid state drives, which resolved this issue.

These challenges would all occur at the HIRAX instrument, so it was instructive to develop a path for tackling them more locally.

6.4.3 Description of observations

Once the instrument challenges were overcome, we were able to have several productive days of beam mapping flights. Figures 6.16, 6.17, 6.18 catalog the measurements taken over three different flight days (one in August, two in October), and give a sense of what can be accomplished in a day of flights. These catalogs also give context to explain where the results presented throughout the rest of the section are drawn from.

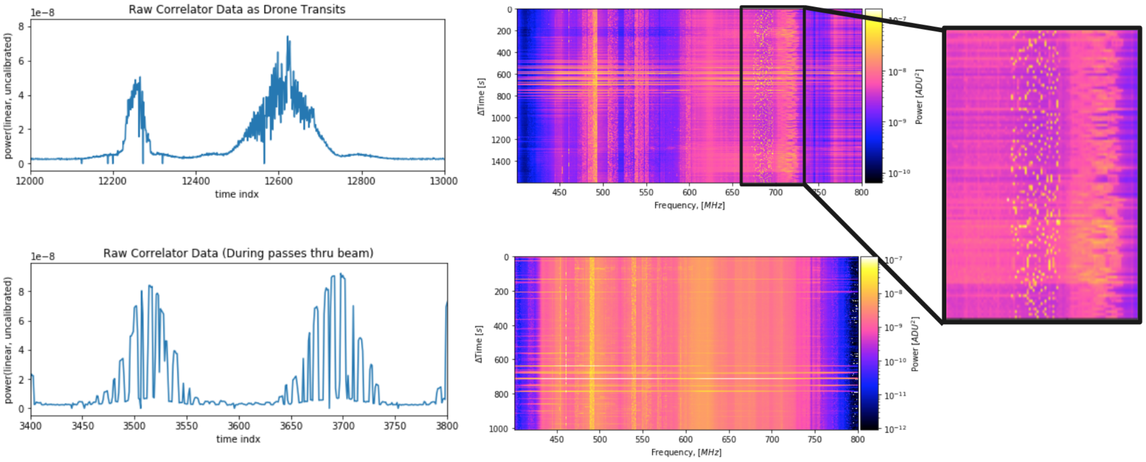

Figure 6.14: **(Top:)** Telescope data with the RTK base station located nearby the array. The top left plot shows measured power (at a single frequency) as a function of time, which we expect to trace out Gaussian shapes as the drone transits (see fig. 6.1). Approximations of Gaussian features are seen in the data, though have substantial noise around the peak, indicating compression and oscillations in the electronics when the drone is over the telescope. The top right figure is a waterfall plot with frequency on the horizontal axis, time on the vertical axis, and colorscale indicating measured power at the telescope. Bright vertical stripes correspond to continuous RFI sources (power at a single frequency for all time), and bright horizontal stripes correspond to drone transits (broadband power for brief times). Zooming in around 660-740MHz, we see little intermittent bright spots, which are due to the GPS base station frequency hopping from 900-930MHz (because of of details of data taking operation, this is aliased into the HIRAX band in 670-700MHz). This power is sufficiently bright to impact with the telescope analog chain. **(Bottom:)** Telescope data with the RTK base station moved further from the array, where line-of-sight between GPS and telescope is blocked by a building. This smooths the band considerably, which is particularly evident in the waterfall plot, though the data looks more sparse than expected, which is addressed below.

We used a similar flight scheme to BMX (figure 6.7), flying separate main beam and side lobe flights, and tackling that analysis separately. The August mapping campaign focused on the general main beam/side lobe measurement set (seen in the latter half of figure 6.16), and the October campaign finished measurement set, and then added an exploration of main beam asymmetries (seen in figure 6.17), sensitive polarization maps (seen in figure 6.18), and nearfield measurements (seen as crosses in figure 6.18).

Most beam maps shown in this thesis are from August data, with polarization

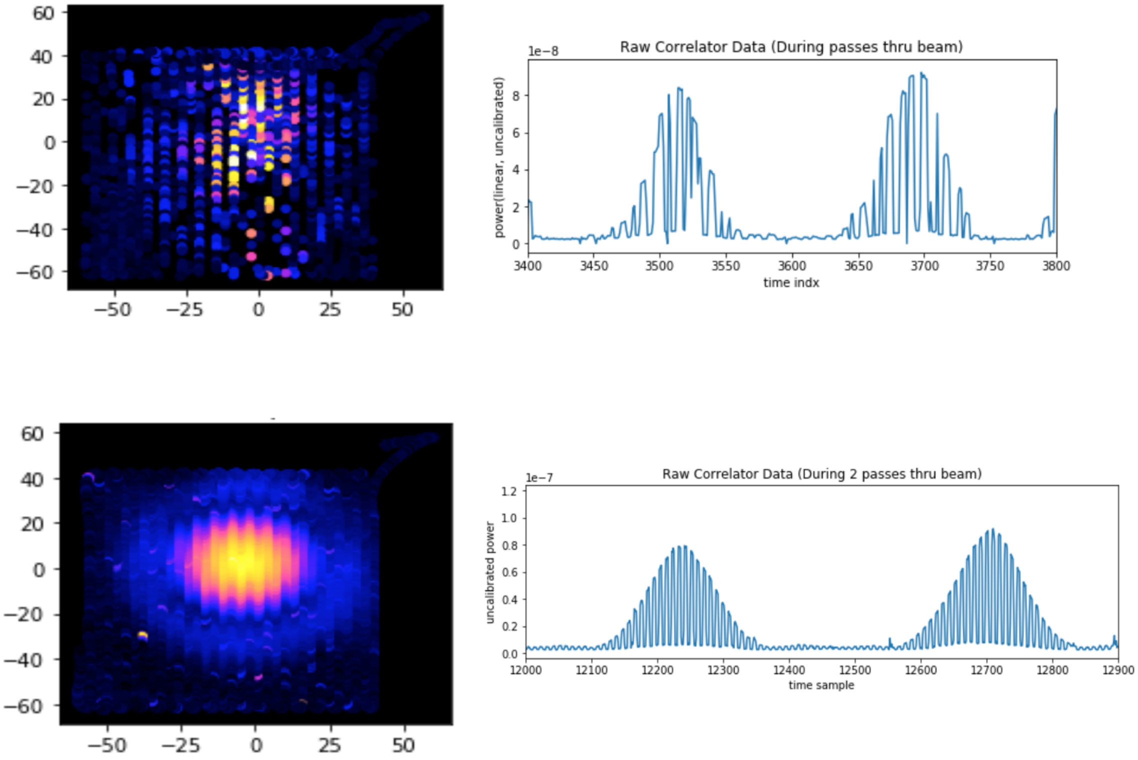

Figure 6.15: Top: Beam map and time stream data when using a standard hard drive for data writing, which resulted in dropped packets and timing delays. Bottom: Beam map and time stream data when writing data to a solid state drive.

results from October.

6.4.4 Main beam mapping

For the main beam mapping, we flew the drone in 120 × 120 meter grids. These flights were centered over dish 0 (labelled in figure 6.13), though the drone flies through the beam of all eight dishes, resulting in beam maps for all dishes (since all dishes take data at once). We had a concrete set of goals for main beam mapping: to measure the main beam Gaussian, perform fits and compare with simulations, and investigate any asymmetries.

Main beam maps from a single North/South grid flight are shown in figure 6.19. The points on the plot correspond to 2D drone position during flight (the N/S move-

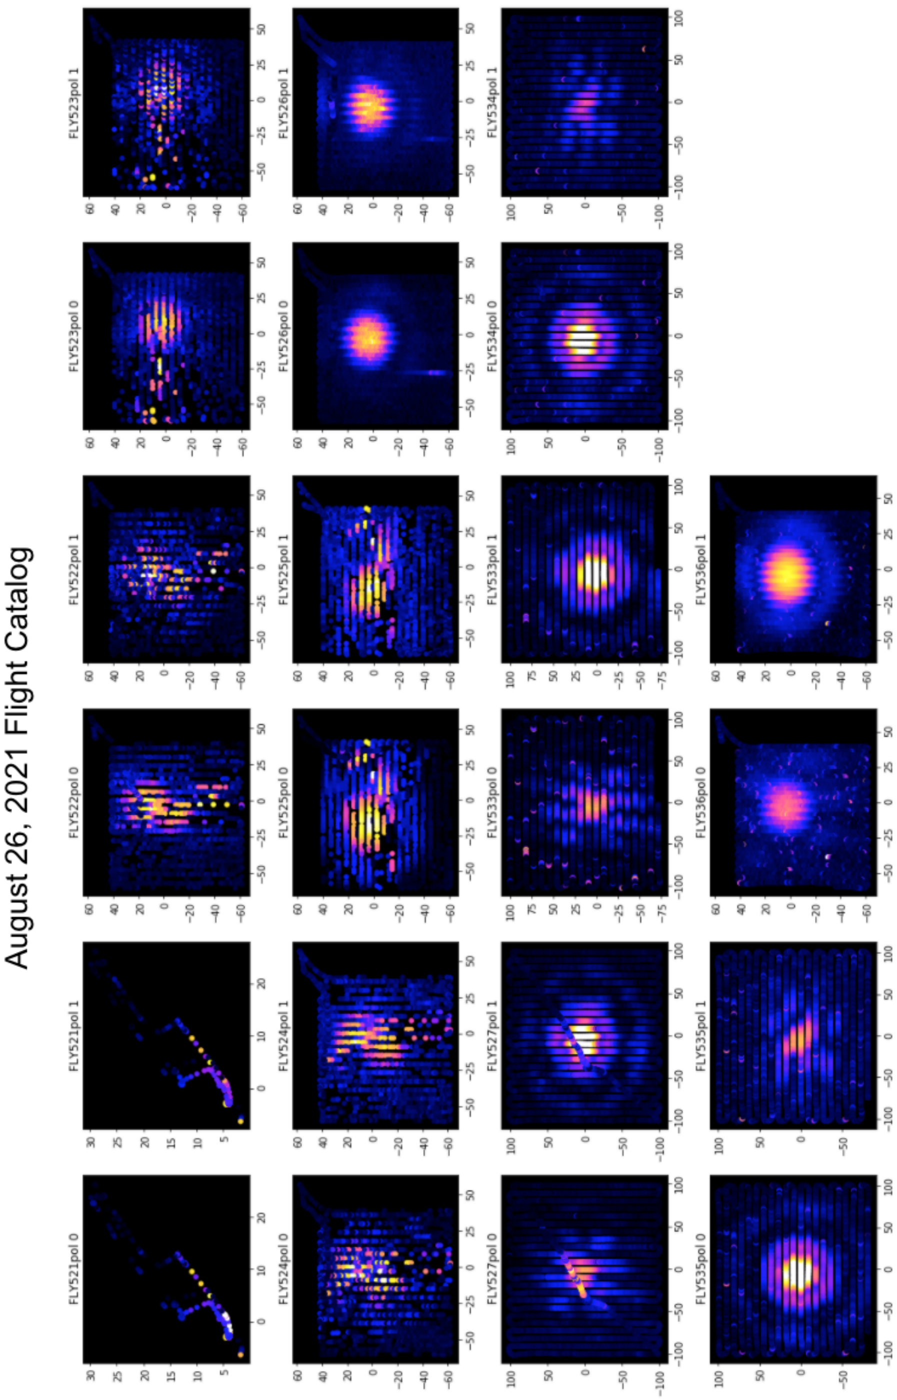

Figure 6.16: Flight catalog of beam maps from GBO flights on August 26, 2021, with titles indicating flight number (indexed by drone file number) and polarization. Timing challenges are evident for the first half of the dataset, after which we complete several sidelobe flights and one main beam flight. Color scale is in log units.

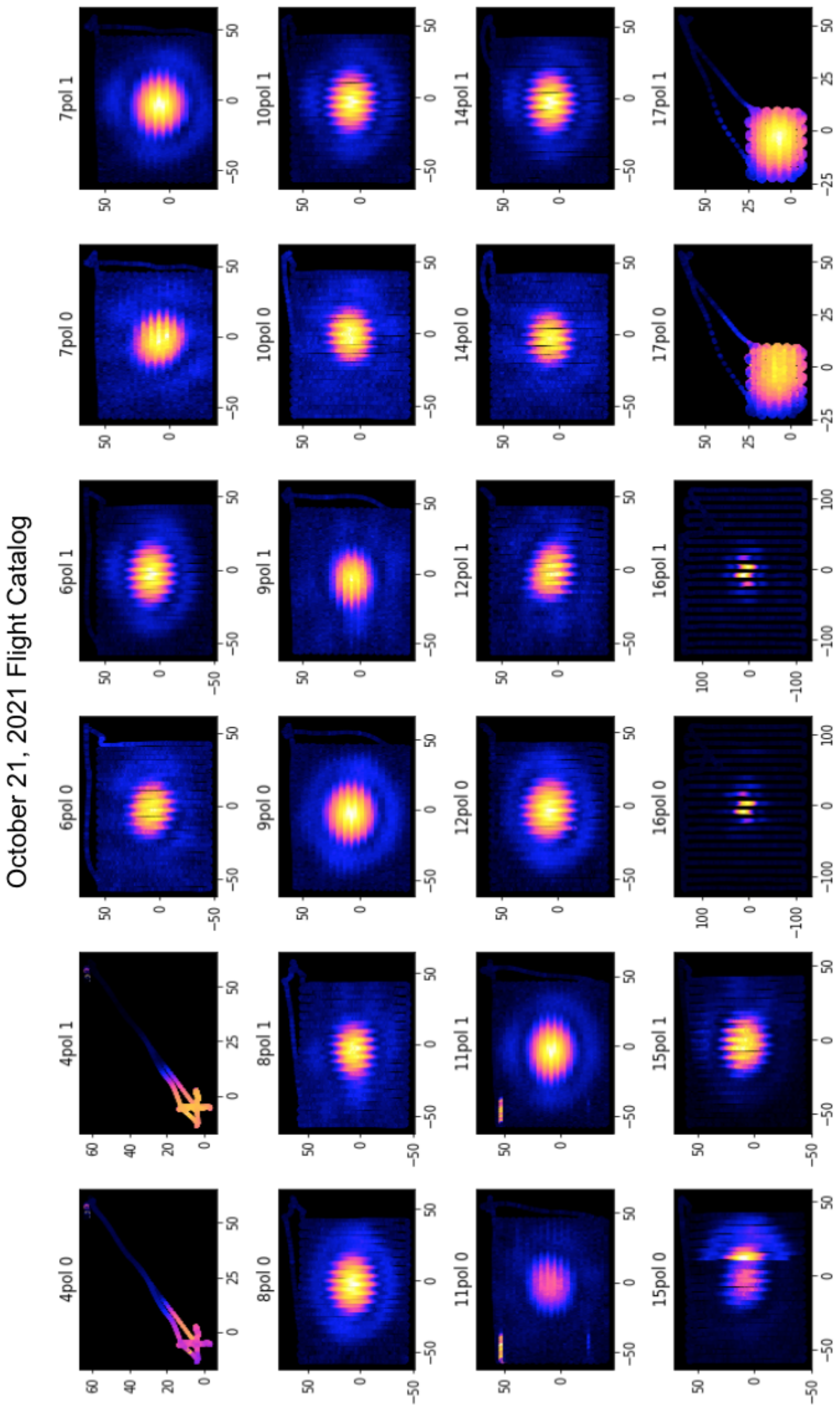

Figure 6.17: Flight catalog of beam maps from GBO flights on October 21, 2021. The first two flights are the polarization axis-finding flights described in figure 6.24, and the remainder are largely a series of main beam maps. Color scale is in log units.

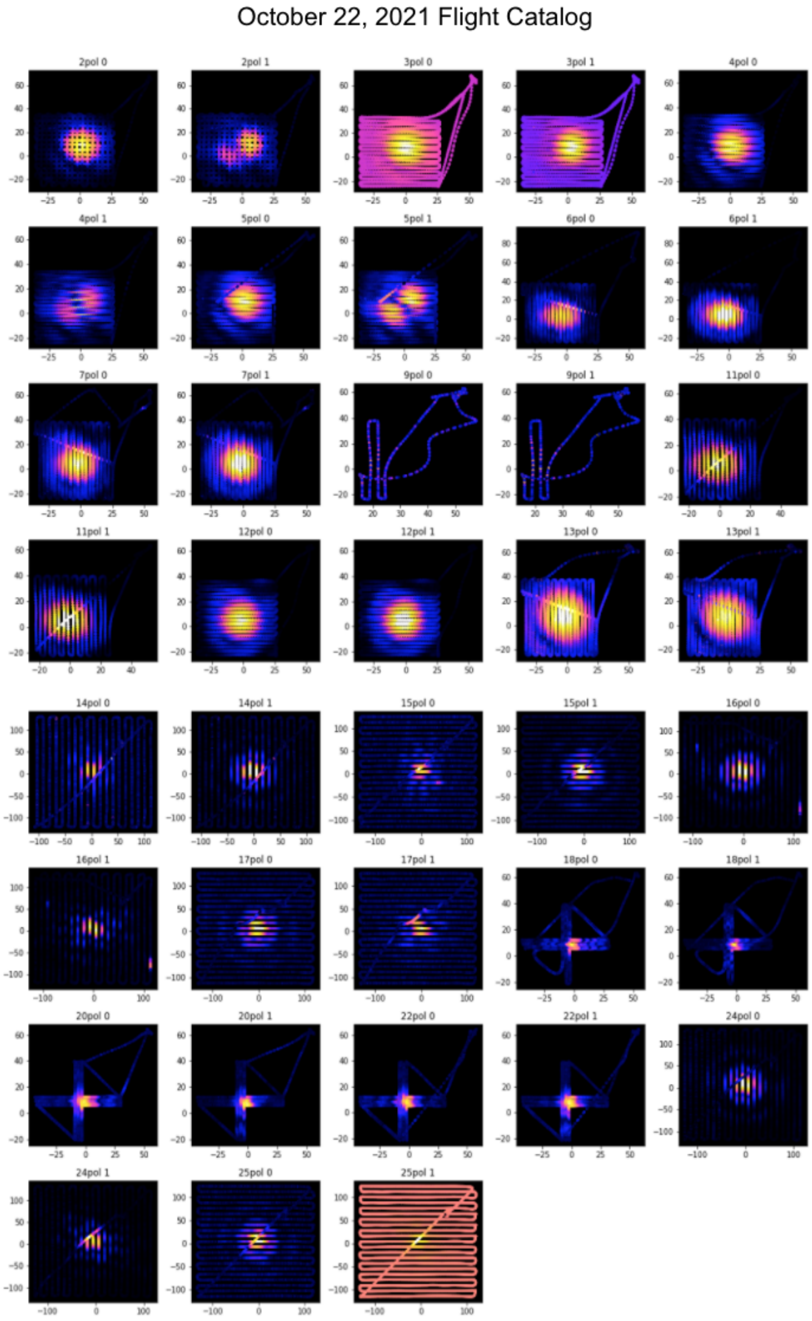
October 22, 2021 Flight Catalog

2pol 0 2pol 1 3pol 0 3pol 1 4pol 0
4pol 1 5pol 0 5pol 1 6pol 0 6pol 1
7pol 0 7pol 1 9pol 0 9pol 1 11pol 0
11pol 1 12pol 0 12pol 1 13pol 0 13pol 1
14pol 0 14pol 1 15pol 0 15pol 1 16pol 0
16pol 1 17pol 0 17pol 1 18pol 0 18pol 1
20pol 0 20pol 1 22pol 0 22pol 1 24pol 0
24pol 1 25pol 0 25pol 1

Figure 6.18: Flight catalog of beam maps from GBO flights on October 22, 2021. These flights begin as a set of polarized main beam maps, flying grids at a range of drone alignments to maximize chances of polarization alignment between the telescope and transmitter. The remainder of the flights tackle sidelobes, and nearfield mapping (crosses in the second-to-last row). Color scale is in log units.

ments are noticeable), and the color scale indicates the power measured at the telescope for each position. Figure 6.19 includes beam maps for all 8 dishes (dish 6 had an inactive channel) in one polarization at 450MHz.

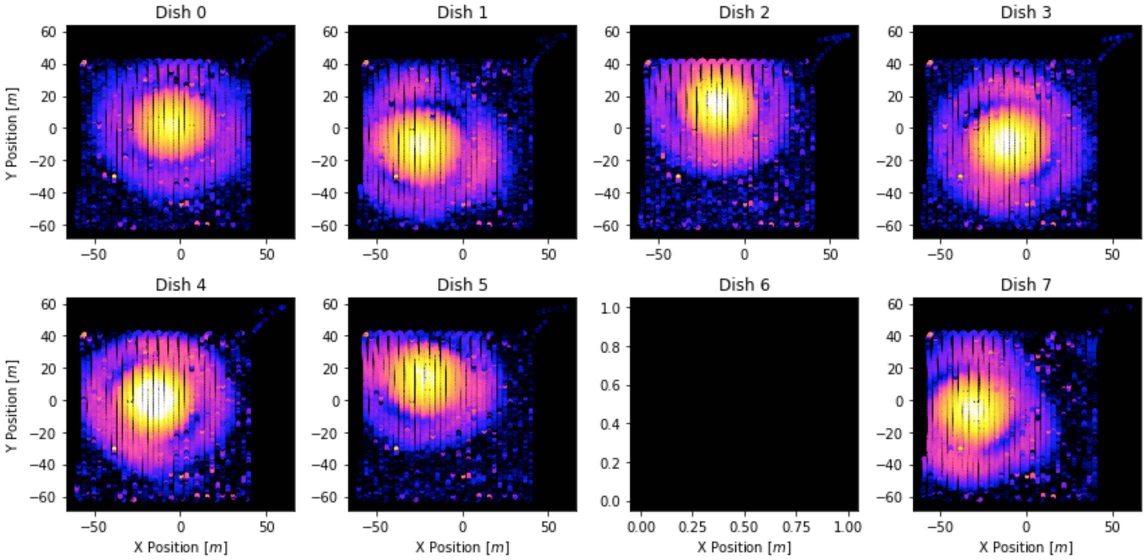

Figure 6.19: Maps of the main beam of 8 CHIME Outrigger dishes at 450MHz (one polarization), generated from the raw drone position data (expressed in meters to highlight the flight path, position will later be converted to angle on the sky). Dish 6 did not have an active channel for this measurement.

There are several features to note in the main beam maps. We can see the main beam Gaussian clearly (rounded yellow feature in the plot center), and can see first side lobes. We also observe that the sidelobe structures are not radially symmetric, with null features along a single axis that appears to rotate depending on the dish. Broadly, these nulls appear as if clustered in two positions, which are separated by 90 degrees. If this feature is from the antenna beam initially, and then propagated forward into the telescope beam, it could give an indication of how the feeds are aligned with respect to the dishes (some could have been mounted 90 degrees off of the expected orientation).

It is instructive to compare our mapping results to simulations. In figure 6.20, we show simulations of an individual dish and results from measurements of one of the CHIME Outrigger dishes at a variety of frequencies, as well as 1D cuts through the simulated and measured beams. We note that the off-axis nulls present in the data are also reflected in the simulations, and are particularly evident at 450 MHz. There also exists a rotation in features between the measured and simulated beams, of up to 50 degrees. The evolution of the beam with frequency is also consistent between measurement and simulation, where in both instances the main-beam loses its Gaussian shape and overall symmetries as it smears in one dimension.

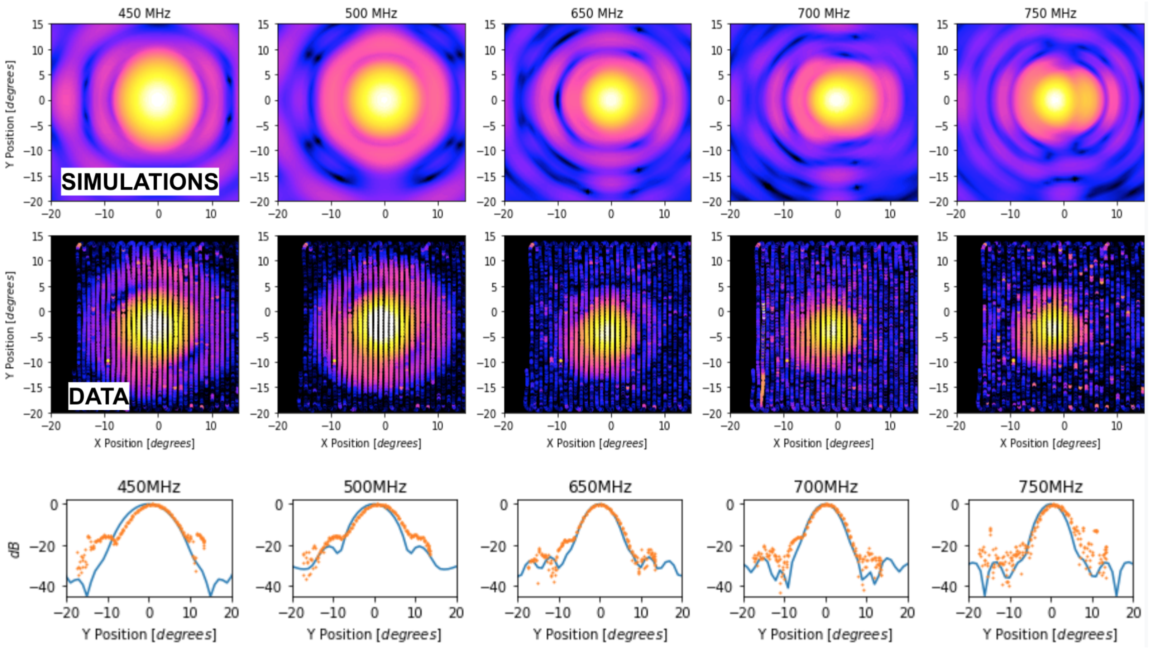

Figure 6.20: Upper panel: Simulated and measured 2D beam maps of a CHIME Outrigger dish for frequencies 450MHz/500MHz/660MHz/700MHz/750MHz. Lower panel: 1D slices through the 2D beam, compared between simulations (blue trace) and data (orange points). Simulations by Benjamin Saliwanchik.

We can also take 1D slices of the beam, to better visualize relative sidelobe levels and patterns. This is most convenient to do along the direction of the drone flight passes, to preserve angular resolution (in the orthogonal direction, resolution will be set by grid spacing in which points are sparser, though still < 0.5 degrees apart). The bottom row of figure 6.20 shows 1D cuts along the Y-dimension, corresponding to the central pass over the telescope (fig 5.1 outlines this procedure). In these cuts we see that main beam widths are consistent between simulations and data, and sidelobe locations in similar angular locations and of similar magnitudes. Further analysis is required to check the sidelobe levels and make a more quantitative comparison. It is also likely that we are comparing different slices of the beam. The simulation and data beam profiles appear rotated with respect to one another, and if we are truly out of alignment in this way, it would be necessary to rotate the profiles into the same coordinate system before comparing cuts along the Y-dimension.

From fits at all frequencies, we can learn additional information about the telescope. We learned that the beam is more circular than simulations predicted (they predicted a 1 degree difference in FWHM between orthogonal beam cuts), and that the full width half max agrees generally with simulations (figure 6.21).

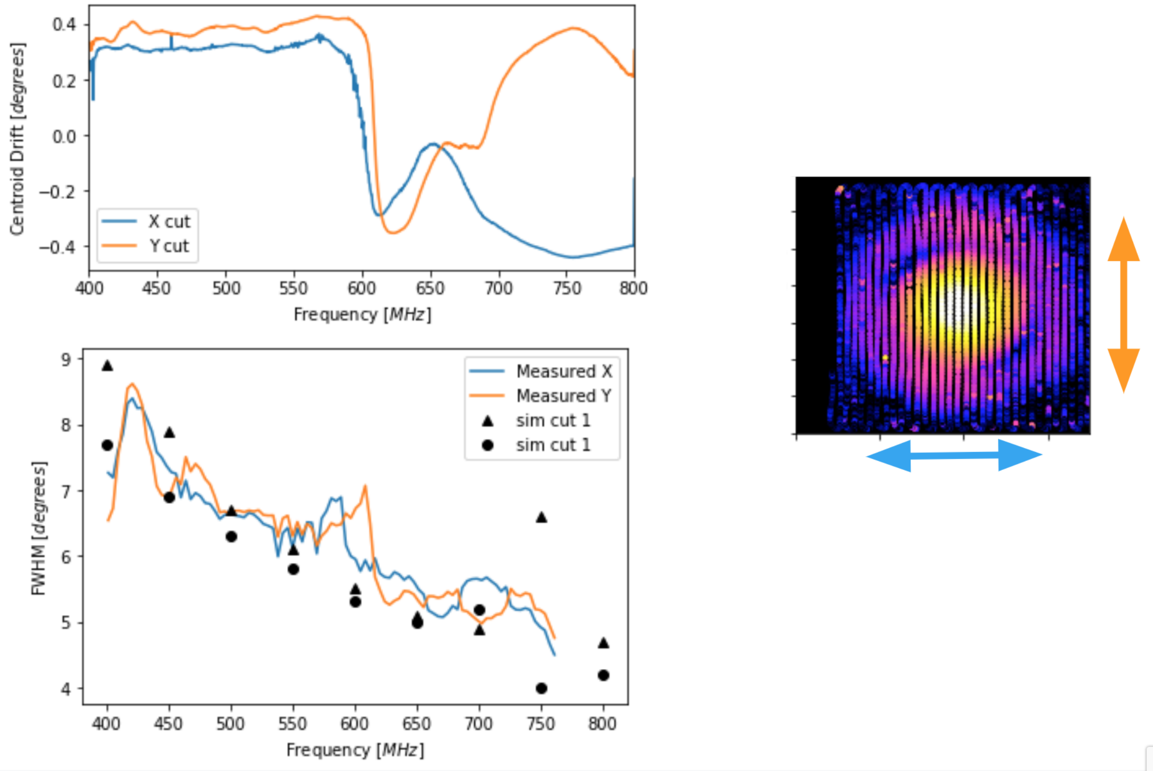

Figure 6.21: The beam centroid drift and FWHM as a function of frequency, found from 2D fits of the main-beam maps. A main beam map illustrating the X and Y dimensions showcased in the plots is also provided.

We can also observe a centroid drift with frequency. This was important, as we know from antenna range measurements of the HIRAX feed (Chapter 3) that the H-plane centroids drift in frequency by up to 25 degrees, and at high frequencies morph into a secondary peak within the main lobe (figure 3.11). A centroid location drift of up to 3 degrees was also observed in separate measurements of source transits by the HIRAX antenna mounted on a 3m dish. We observe a less substantial drift in the CHIME outrigger dishes, up to ±0.4 degrees. The drift does, however, follow the same general profile at the HIRAX feed measurements, dipping between 600-650MHz in both dimensions. This investigation gives information about how antenna non-idealities propagate into the telescope beam, and is a demonstration of how different dish designs will mask such features to varying degrees. More analysis on this front is in progress. Fit results are shown in figure 6.21.

To round out the main beam flight set, we performed flights with the drone rotated 180 degrees from its prior alignments. As in antenna range measurements, this rotation of transmitter can distinguish whether asymmetric features in the beam map are due to the drone beam pattern or the telescope beam (discussion in Chapter 3). We know from antenna range measurements that the drone beam is smooth and symmetric in two orthogonal 1D cuts, so in principle should not produce prominent asymmetries, but this early analysis is a useful second check. We find that the rotated drone orientation does not change main-beam feature locations or beam map information qualitatively.

6.4.5 Side lobe mapping

Side lobe mapping data sets provide a broader angular range in which to compare between data and simulations. For these flight sets, the drone flew 200 x 200 meter grids, capturing multiple side lobes out to 30 degrees from zenith. This flight plan was informed by simulations, which show side lobe expected levels to drop below -30dB out of this range, and predicts we will see at least the initial two sidelobes at most frequencies.

A comparison between simulations and measurements at several frequencies is shown in figure 6.22. There are several qualitative features in common between simulation and measurement:

- At 450MHz, the main beam appears Gaussian, and sidelobes display a petal-like pattern.
- At 650MHz, the main beam remains Gaussian though shrinks in width (as expected), and a cross-pattern of higher intensity patches emerges in the sidelobes.
- Above 700MHz, the main beam loses its gaussian shape and smears into two asymmetric lobes (also seen in the flights dedicated to main beam mapping), while the sidelobes become more densely concentrated.

The next analysis phase will involve combining these sidelobe results with the main beam results to obtain complete maps.

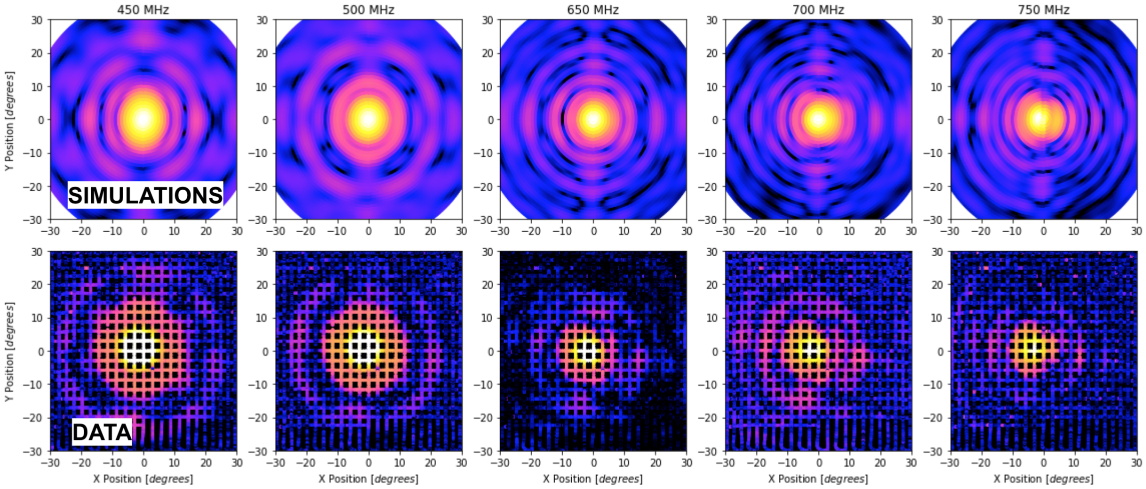

Figure 6.22: Simulated and measured 2D beam maps of a CHIME Outrigger dish for frequencies 450MHz/500MHz/660MHz/700MHz/750MHz. Simulations by Benjamin Saliwanchik.

In addition to the frequency-dependent features, we notice a rotation between simulations and data, of ∼15 degrees. This is suggestive of the telescope having an unexpected physical rotation, but it is too large of a discrepancy to be accounted for by dish alignment errors (estimated at 2 degrees) or by an accidental misalignment to magnetic rather than true north (9 degrees discrepant at Green Bank). We continue to investigate the source of this rotation, which is consistent with observations from polarization axis alignment, as discussed in the coming section.

6.4.6 Polarization maps

A main advantage of using drones over traditional calibration methods, even for steerable dishes, is that drone calibration can assess polarization dependent properties of a telescope. The capability to measure cross-polarization is important for making truly polarized maps, which will minimize leakage of unexpected signal into sky measurements.

The drone's transmitter antenna is singularly polarized, so multiple flights at orthogonal transmitter orientations have to be undertaken to make polarized maps. Figure 6.23 shows example large-grid maps for two polarizations on one dish, which required four total flights (2 per polarization). We see a 90 degree rotation between beam features in the two polarizations. This is still preliminary work, shown as a demonstration of what is possible with drone calibration.

When making a co-pol map of one channel on a dual-polarization telescope, we automatically obtain cross-pol of the other channel (also in figure 6.23). In early flights, we see a petal-structure in the cross-pol pattern, which could be due to the telescope's metallic feed legs giving structure to the beam. Additionally, in cross-pol we expect for there to be a dip in beam center, while instead we see a peak. The pattern is difficult to interpret at this time, but there could be a few factors at play: the dishes may not be perfectly aligned with the polarizaton axes, which could

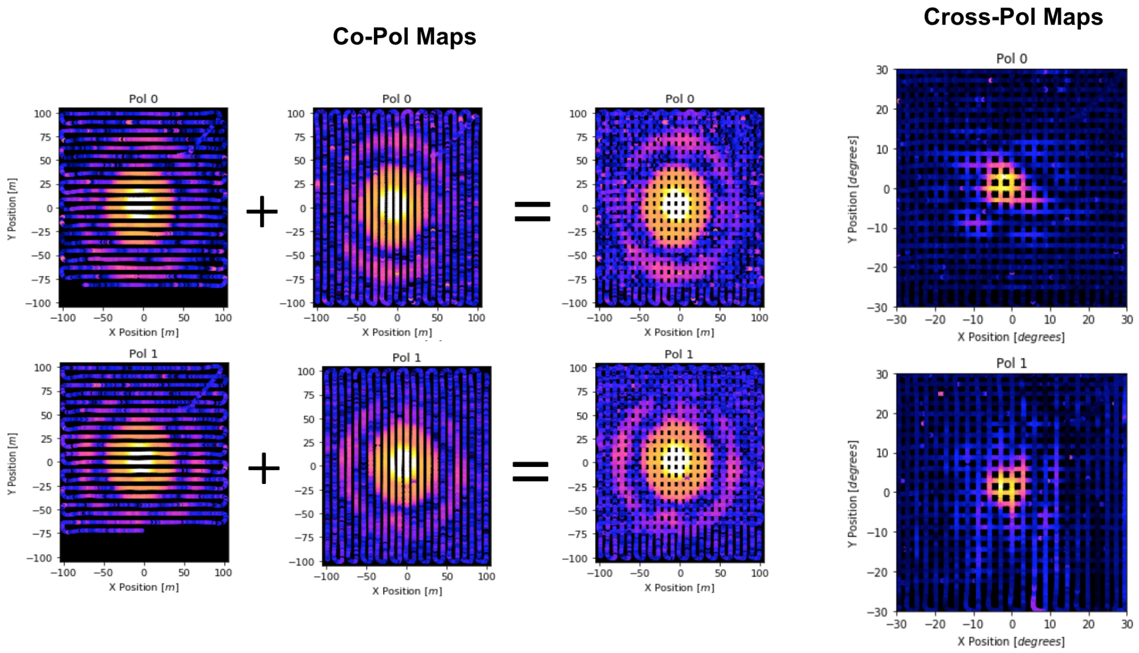

Figure 6.23: Polarized 2D maps of a CHIME outrigger dish, generated from two orthogonal large grids.

result in co-pol power showing up in the cross-pol pattern (even small amounts can be enough to obscure the dip); the maps have sky backgrounds subtracted from the data, which in practice involves subtracting 2 small numbers, meaning small fluctuations could have large impacts.

In August flights (which includes maps from fig. 6.23), the drone was oriented to align the transmitter polarization North/South, which assumes the same alignment from the telescope. In October, a main focus was on making sensitively polarized maps, for which the first step is determining the true polarization axes of the telescope receiver.

To determine the orientation of the telescope polarization, we perform a flight where: the drone flies in the telescope far field to the telescope beam center, rotates 360 degrees, moves to a new location in the beam, rotates 360 degrees, and repeats for five points in the beam. The power measured during the rotations should peak when the drone transmitter polarization and telescope polarization are aligned, and the peak locations should remain consistent across the rotation points. When measuring two polarizations, we expect peaks 90 degrees apart. We look at data for all dishes, and focus particularly on a neighboring pair of dishes, where one is instrumented with a HIRAX feed and the other instrumented with a CHIME feed. Set up and results shown in figure 6.24.

In the telescope data from such a flight, there is unusual behaviors for dishes instrumented with HIRAX feeds. When plotting measured power vs drone yaw angle at high frequencies, the peak locations were dependent on drone location in the beam, and spread more broadly as frequency increases. Additionally, the two polarizations appear orthogonal at low frequencies but nearly aligned in the middle of the band ($\sim$600MHz). Looking at a neighboring dish that is instrumented with a CHIME feed, we see all of the expected behaviors. The HIRAX frequency dependent behavior is recovered in simulations, shown in figure 6.25.

High levels of cross-pol could account for the apparent polarization axis shift at different positions. The effects would be particularly pronounced at high frequencies, when the main beam narrows, so positional shifts within the beam will place the drone further off of the cross-pol null. Indeed, simulations and antenna range measurements (described in Chapter 3) both indicate high levels of cross-pol for the HIRAX feed. In these simulations, it appears that cross-pol levels are comparable to co-pol levels at high frequencies at 5 degrees off of boresight. This is shown in figure 6.26, which visualizes the cross-pol beam pattern at 400MHz, 600MHz, and 800MHz, and shows 1D cuts through different slices of the beam at each of these frequencies. It is clear that the cross-pol null narrows, and the overall level rises, at high frequencies. This analysis is ongoing, but has already sparked discussions of upgrading the HIRAX antenna design among the collaboration. If/when the design is upgraded, the noise temperature measurement system described in Chapter 4 will be one of the first places the new antenna undergoes testing.

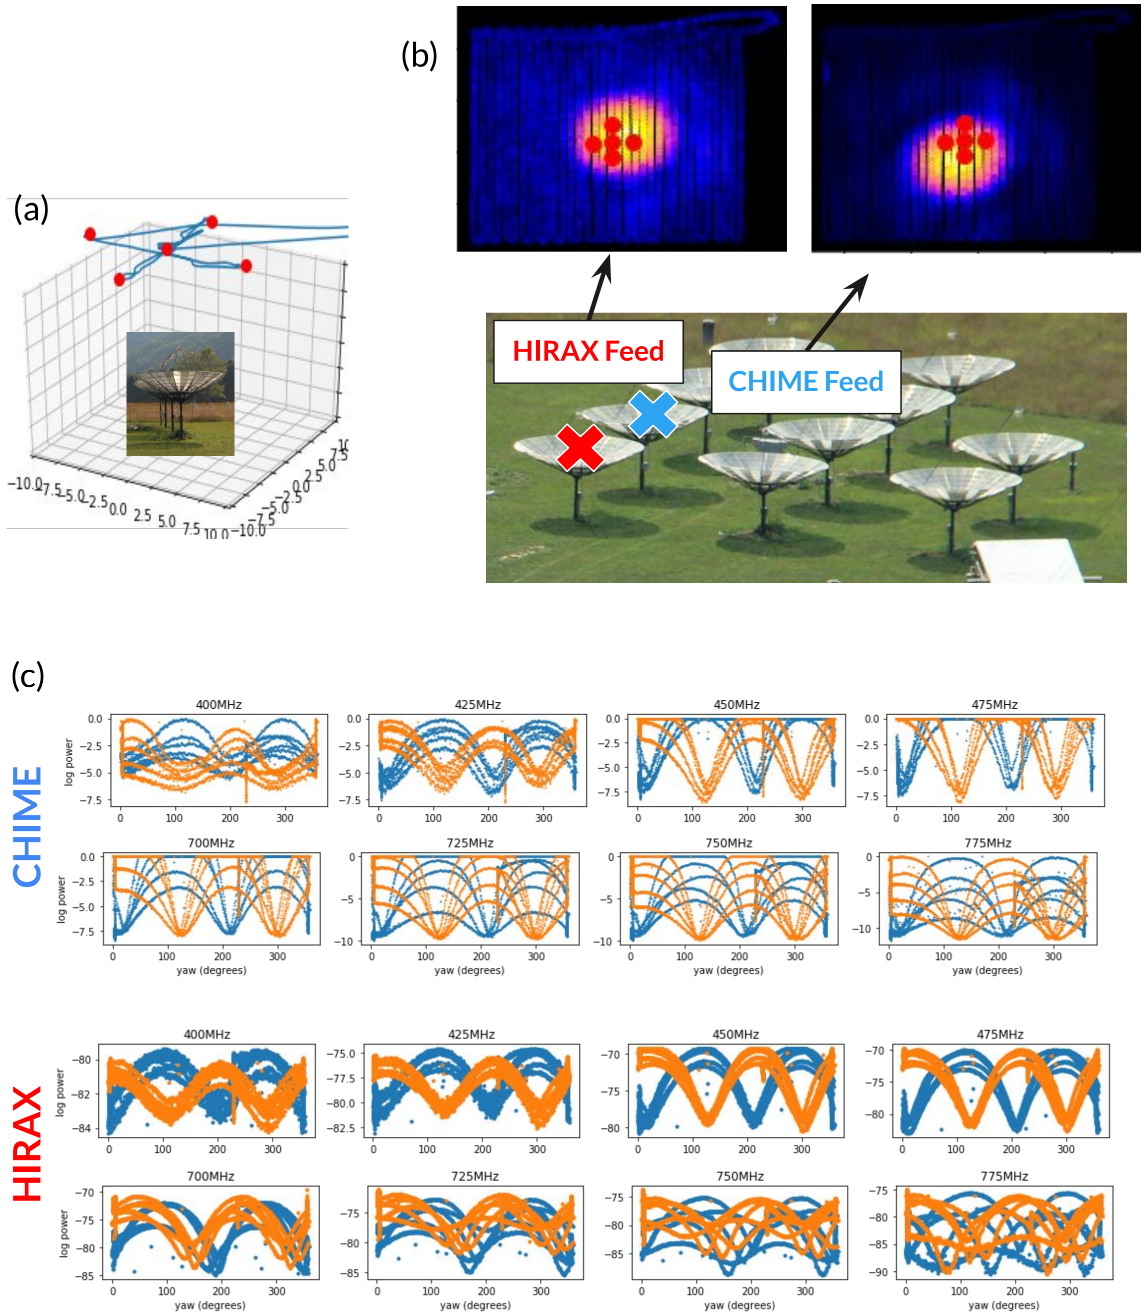

Figure 6.24: (a) Flight pattern to determine the telescope polarization alignment. The drone hovers and rotates 360 degrees at each of the points marked in red. (b) Hover/rotation point locations shown within the main beams of two different telescopes, one instrumented with a HIRAX feed and the other with a CHIME feed. (c) Plots of drone yaw angle vs. received telescope power at all rotation points and for several frequencies. The two colors represent two telescope polarizations.

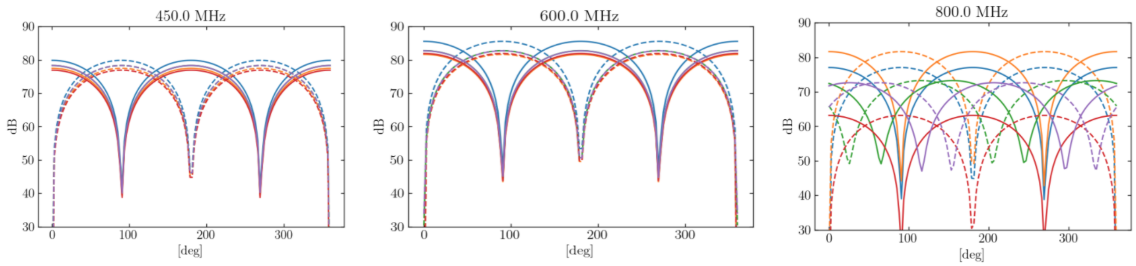

Figure 6.25: Simulations of rotational cuts at different locations in the beam from a dish with a HIRAX feed. Solid lines are for a co-polar source, and dashed lines are for a cross-polar source. The different colors are for different beam sampling points. This behavior is consistent with measurements at Green Bank. Simulations by Devin Crichton.

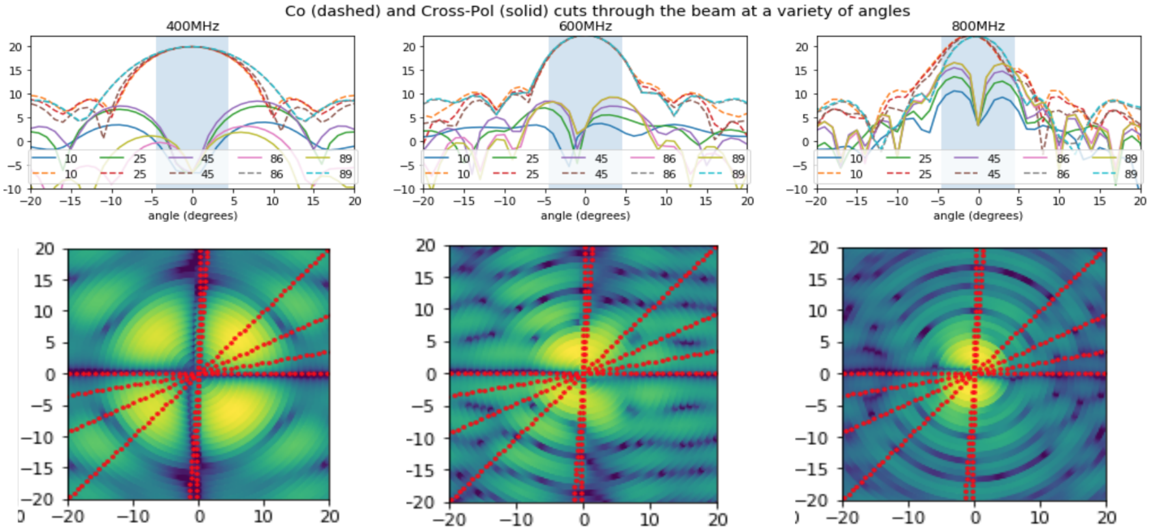

Figure 6.26: Simulations of the 2D cross-pol pattern, and 1D cuts through different slices of the simulated beam (marked in red). For the 1D cuts in the top row, the dashed lines are co-pol cuts and the solid cuts are cross-pol cuts. The legend indicates the angle through the beam relative to a horizontal line. At high frequencies, the cross-pol level can be comparable to the co-pol level once 5 degrees off zenith, which could result in the behaviors seen in figure 6.24. Simulations by Ben Saliwanchik

6.5 Summary and ongoing work

This chapter described the implementation of a drone-based beam calibration system, reporting results from flight campaigns at three different telescopes: DSA at Owen's Valley Radio Observatory, BMX at Brookhaven National Laboratory, and the CHIME Outrigger Prototype at Green Bank Observatory. The chapter also highlights how flight data informed the drone design progression (with the current design described fully in Chapter 5).

We will use the dataset from Green Bank to build the full analysis pipeline required for beam maps of HIRAX and other 21cm instruments. The work presented in this thesis demonstrates important progress towards this goal, but there is still more to be done. To finalize the analysis pipeline, and to provide further feedback on the CHIME Outrigger system, we will provide a more rigorous comparison between simulations and data, incorporate the drone transmitter beam pattern into the telescope beam maps, develop a scheme for amplitude-matching high and low power flights, and complete cross-pol investigations.

There are several extensions of this work ongoing at Yale as well. The drone group is looking to develop near-field mapping capabilities, to enable use with telescopes such as CHIME that have unreachable far-fields. A digital noise source, which would enable phase measurements in addition to amplitude measurements, is under development for this purpose. We are also expanding to a dual-polarization transmitter, and exploring autopilots that will allow for mapping with spherical grids.

Looking further ahead, we can extend the drone hardware to encompass lower frequencies, and establish it a standard calibrator for all 21cm instruments.

Chapter 7

Conclusion

The work in this dissertation focused on custom calibration hardware developed for the Hydrogen Intensity and Real-time Analysis eXperiment (HIRAX), a revolutionary new telescope array that will utilize intensity mapping techniques to probe dark energy. Because 21cm experiments like HIRAX aim to detect small signals set against bright foregrounds, calibration will be a key piece of their realization. The layout and structure of these experiments bring about several calibration challenges, including beam mapping and antenna sensitivity and redundancy. In this thesis, I described three projects relating to the characterization and calibration of HIRAX: beam measurements of the HIRAX antennas, noise temperature characterization of the HIRAX antennas, and drone calibration of HIRAX-style arrays.

For the first project (Chapter 3), I used a range of facilities to take beam measurements of a number of antennas relevant to the HIRAX experiment, including several HIRAX antennas, a CHIME antenna, and a biconical antenna (to use as a transmitter for drone calibration). In characterizing multiple HIRAX antennas, I found the designs to be repeatable, and in agreement with simulations, providing an important check on the collaboration simulation pipeline. I additionally discovered an asymmetry in the co-polar beam, and notably high levels of power in cross-polarization (this was re-established in drone calibration measurements)—as a result, this design is currently under assessment by the collaboration. In comparing the HIRAX antenna design with the CHIME antenna design, I found the CHIME antenna to have a more symmetrized beam and lower cross pol, which motivates further study of the design, and suggests that HIRAX can accomplish an improved antenna performance with minimal design changes. Finally, I investigated a biconical antenna, finding it has the expected broad-band and omni-directional beam shape when measured on its own, but that the beam pattern is easily influenced by the presence of additional structures, such as a drone carbon fiber body. Thus, for our purposes the bicolog should be backed with a structure that decouples its beam features from the carbon fiber drone, in order to preserve its broad, smooth beam features. This work is continued in later chapters, including Chapter 4 and Chapter 5.

For the second project (Chapter 4), I designed and implemented a custom cryogenic system for taking Y-factor measurements of radio antennas, going from preliminary electromagnetic modelling to a full experimental apparatus. I fabricated a variety of custom instrumentation, including a cryogenic RF chamber capable of containing > 550L of liquid nitrogen, a $12 \times 8 \times 8$ft faraday cage for RFI shielding, and tools for shaping RF pyramidal absorber. I performed hundreds of detailed VNA and Spectrum Analyzer verification measurements to assess system RF performance and characterize systematics. I found the HIRAX antenna noise temperature to meet its 50K spec and the HIRAX antennas repeatable to within 10%. With this work, I delivered the first noise temperature analysis of the HIRAX antennas and the first assessment of antenna repeatability. This work also served as a critical validation of the experiment antenna design before fielding the first 256-dish array. This characterization scheme is cheap, repeatable, widely relevant, and of use to all researchers wanting to better characterize RF equipment. It opens the door for lighter-weight, lower-noise, lower-cost antennas to be cheaply characterized, which could be helpful for a range of research and industrial applications.

The third project in my dissertation (Chapters 5, Chapter 6) focuses on drone calibration of radio telescopes. Since 2018, I have built and characterized a transmitter system consisting of a stable broadband noise source, a GPS- locked timing board, an omnidirectional transmitter antenna, and a precision differential GPS to fly on a hexacopter drone for 21cm beam mapping. This drone system has been deployed to telescopes at both the DSA-100 telescope at Owens Valley Radio Observatory in Bishop, CA in 2019, the BMX telescope at Brookhaven National Lab in Long Island, NY in early 2020, and the CHIME Outrigger Prototype at Green Bank Observatory in August and October of 2021. This work provided important instrument feedback to the respective telescope teams, on such specifications as dish pointing and polarization alignment, and enabled a sensitive measurement of the HIRAX antenna polarization properties, thereby validating results found at antenna test ranges in Chapter 3 and prompting discussions about HIRAX feed model alterations. The beam maps generated through drone calibration measurements will be crucial to foreground removal from cosmology data, without which the HIRAX science goals cannot be met. The development of novel drone techniques will be useful to collaborations ranging from other 21cm experiments like HERA, to Simons Observatory in CMB, to space missions like NASA's Europa, all of whom share beam mapping challenges, and are independently investigating their own drone-based solutions.

Bibliography

[1] B. Ryden, *Introduction to cosmology* (Cambridge University Press, 2017).

[2] N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. Banday, R. Barreiro, N. Bartolo, S. Basak, *et al.*, Astronomy & Astrophysics **641**, A6 (2020).

[3] S. Alam, M. Aubert, S. Avila, C. Balland, J. E. Bautista, M. A. Bershady, D. Bizyaev, M. R. Blanton, A. S. Bolton, J. Bovy, *et al.*, Physical Review D **103**, 083533 (2021).

[4] T. D. Collaboration, arXiv preprint arXiv:1611.00036 (2016).

[5] D. Collaboration, Monthly Notices of the Royal Astronomical Society **480**, 3879 (2018).

[6] A. G. Riess, A. V. Filippenko, P. Challis, A. Clocchiatti, A. Diercks, P. M. Garnavich, R. L. Gilliland, C. J. Hogan, S. Jha, R. P. Kirshner, *et al.*, The Astronomical Journal **116**, 1009 (1998).

[7] W. L. Freedman, B. F. Madore, D. Hatt, T. J. Hoyt, I.-S. Jang, R. L. Beaton, C. R. Burns, M. G. Lee, A. J. Monson, J. R. Neeley, *et al.*, arXiv preprint arXiv:1907.05922 (2019).

[8] R. J. Adler, B. Casey, and O. C. Jacob, American Journal of Physics **63**, 620 (1995).

[9] T. Abbott, S. Allam, P. Andersen, C. Angus, J. Asorey, A. Avelino, S. Avila, B. Bassett, K. Bechtol, G. Bernstein, *et al.*, The Astrophysical Journal Letters **872**, L30 (2019).

[10] M. Ata, F. Baumgarten, J. Bautista, F. Beutler, D. Bizyaev, M. R. Blanton, J. A. Blazek, A. S. Bolton, J. Brinkmann, J. R. Brownstein, *et al.*, Monthly Notices of the Royal Astronomical Society **473**, 4773 (2017).

[11] T. Delubac, J. E. Bautista, J. Rich, D. Kirkby, S. Bailey, A. Font-Ribera, A. Slosar, K.-G. Lee, M. M. Pieri, J.-C. Hamilton, *et al.*, Astronomy & Astrophysics **574**, A59 (2015).

[12] E. A. Kazin, J. Koda, C. Blake, N. Padmanabhan, S. Brough, M. Colless, C. Contreras, W. Couch, S. Croom, D. J. Croton, *et al.*, Monthly Notices of the Royal Astronomical Society **441**, 3524 (2014).

[13] L. Samushia, B. A. Reid, M. White, W. J. Percival, A. J. Cuesta, G.-B. Zhao, A. J. Ross, M. Manera, É. Aubourg, F. Beutler, *et al.*, Monthly Notices of the Royal Astronomical Society **439**, 3504 (2014).

[14] C. Blake, E. A. Kazin, F. Beutler, T. M. Davis, D. Parkinson, S. Brough, M. Colless, C. Contreras, W. Couch, S. Croom, *et al.*, Monthly Notices of the Royal Astronomical Society **418**, 1707 (2011).

[15] J. Hou, A. G. Sánchez, A. J. Ross, A. Smith, R. Neveux, J. Bautista, E. Burtin, C. Zhao, R. Scoccimarro, K. S. Dawson, *et al.*, Monthly Notices of the Royal Astronomical Society **500**, 1201 (2021).

[16] R. Neveux, E. Burtin, A. de Mattia, A. Smith, A. J. Ross, J. Hou, J. Bautista, J. Brinkmann, C.-H. Chuang, K. S. Dawson, *et al.*, Monthly Notices of the Royal Astronomical Society **499**, 210 (2020).

[17] D. Eisenstein *et al.*, The Astrophysical Journal **633**, 560 (2005).

[18] M. Zaldarriaga, S. R. Furlanetto, and L. Hernquist, The Astrophysical Journal **608**, 622 (2004).

[19] T.-C. Chang, U.-L. Pen, J. B. Peterson, and P. McDonald, Physical Review Letters **100**, 091303 (2008).

[20] T.-C. Chang, U.-L. Pen, K. Bandura, and J. B. Peterson, Nature **466**, 463 (2010).

[21] J. R. Pritchard and A. Loeb, Physical Review D **82**, 023006 (2010).

[22] J. D. Bowman, A. E. Rogers, R. A. Monsalve, T. J. Mozdzen, and N. Mahesh, Nature **555**, 67 (2018).

[23] L. Philip, Z. Abdurashidova, H. Chiang, N. Ghazi, A. Gumba, H. Heilgendorff, J. Jáuregui-García, K. Malepe, C. Nunhokee, J. Peterson, *et al.*, Journal of Astronomical Instrumentation **8**, 1950004 (2019).

[24] D. Price, L. Greenhill, A. Fialkov, G. Bernardi, H. Garsden, B. Barsdell, J. Kocz, M. Anderson, S. Bourke, J. Craig, *et al.*, Monthly Notices of the Royal Astronomical Society **478**, 4193 (2018).

[25] M. W. Eastwood, M. M. Anderson, R. M. Monroe, G. Hallinan, M. Catha, J. Dowell, H. Garsden, L. J. Greenhill, B. C. Hicks, J. Kocz, *et al.*, The Astronomical Journal **158**, 84 (2019).

[26] A. Liu, J. R. Pritchard, M. Tegmark, and A. Loeb, Physical Review D **87**, 043002 (2013).

[27] D. R. DeBoer, A. R. Parsons, J. E. Aguirre, P. Alexander, Z. S. Ali, A. P. Beardsley, G. Bernardi, J. D. Bowman, R. F. Bradley, C. L. Carilli, *et al.*, Publications of the Astronomical Society of the Pacific **129**, 045001 (2017).

[28] Z. Abdurashidova, J. E. Aguirre, P. Alexander, Z. S. Ali, Y. Balfour, R. Barkana, A. P. Beardsley, G. Bernardi, T. S. Billings, J. D. Bowman, *et al.*, The Astrophysical Journal **924**, 51 (2022).

[29] C. Cheng, A. R. Parsons, M. Kolopanis, D. C. Jacobs, A. Liu, S. A. Kohn, J. E. Aguirre, J. C. Pober, Z. S. Ali, G. Bernardi, *et al.*, The Astrophysical Journal **868**, 26 (2018).

[30] A. R. Parsons, A. Liu, J. E. Aguirre, Z. S. Ali, R. F. Bradley, C. L. Carilli, D. R. DeBoer, M. R. Dexter, N. E. Gugliucci, D. C. Jacobs, *et al.*, The Astrophysical Journal **788**, 106 (2014).

[31] M. P. van Haarlem, M. W. Wise, A. Gunst, G. Heald, J. P. McKean, J. W. Hessels, A. G. de Bruyn, R. Nijboer, J. Swinbank, R. Fallows, *et al.*, Astronomy & astrophysics **556**, A2 (2013).

[32] S. J. Tingay, R. Goeke, J. D. Bowman, D. Emrich, S. M. Ord, D. A. Mitchell, M. F. Morales, T. Booler, B. Crosse, R. B. Wayth, *et al.*, Publications of the Astronomical Society of Australia **30** (2013).

[33] Y. Gupta, B. Ajithkumar, H. Kale, S. Nayak, S. Sabhapathy, S. Sureshkumar, R. Swami, J. Chengalur, S. Ghosh, C. Ishwara-Chandra, *et al.*, Current Science , 707 (2017).

[34] X. Chen, in *International Journal of Modern Physics: Conference Series*, Vol. 12 (World Scientific, 2012) pp. 256–263.

[35] M. Amiri, K. Bandura, T. Chen, M. Deng, M. Dobbs, M. Fandino, S. Foreman, M. Halpern, A. S. Hill, G. Hinshaw, *et al.*, arXiv preprint arXiv:2202.01242 (2022).

[36] M. Amiri, K. Bandura, A. Boskovic, T. Chen, J.-F. Cliche, M. Deng, N. Denman, M. Dobbs, M. Fandino, S. Foreman, *et al.*, arXiv preprint arXiv:2201.07869 (2022).

[37] L. Newburgh, K. Bandura, M. Bucher, T.-C. Chang, H. Chiang, J. Cliche, R. Davé, M. Dobbs, C. Clarkson, K. Ganga, *et al.*, in *Ground-based and Airborne Telescopes VI*, Vol. 9906 (International Society for Optics and Photonics, 2016) p. 99065X.

[38] P. E. Dewdney, P. J. Hall, R. T. Schilizzi, and T. J. L. Lazio, Proceedings of the IEEE **97**, 1482 (2009).

[39] L. Koopmans, J. Pritchard, G. Mellema, F. Abdalla, J. Aguirre, K. Ahn, R. Barkana, I. van Bemmel, G. Bernardi, A. Bonaldi, *et al.*, arXiv preprint arXiv:1505.07568 (2015).

[40] A. Slosar, Z. Ahmed, D. Alonso, M. A. Amin, R. Ansari, E. J. Arena, K. Bandura, N. Battaglia, J. Blazek, P. Bull, *et al.*, Journal of Environmental Sciences (China) English Ed (2019).

[41] J. O. Burns, G. Hallinan, *et al.*, arXiv:1911.08649 (2019).

[42] G. Hallinan, V. Ravi, S. Weinreb, J. Kocz, Y. Huang, D. Woody, J. Lamb, L. D'Addario, M. Catha, J. Shi, *et al.*, arXiv preprint arXiv:1907.07648 (2019).

[43] P. Boyle, C. Collaboration, *et al.*, ATel **11901**, 1 (2018).

[44] M. Amiri, K. Bandura, M. Bhardwaj, P. Boubel, M. M. Boyce, P. J. Boyle, C. Brar, M. Burhanpurkar, T. Cassanelli, P. Chawla, J. F. Cliche, D. Cubranic, M. Deng, N. Denman, M. Dobbs, M. Fandino, E. Fonseca, B. M. Gaensler, A. J. Gilbert, A. Gill, U. Giri, D. C. Good, M. Halpern, D. S. Hanna, A. S. Hill, G. Hinshaw, C. Höfer, A. Josephy, V. M. Kaspi, T. L. Landecker, D. A. Lang,

H. H. Lin, K. W. Masui, R. Mckinven, J. Mena-Parra, M. Merryfield, D. Michilli, N. Milutinovic, C. Moatti, A. Naidu, L. B. Newburgh, C. Ng, C. Patel, U. Pen, T. Pinsonneault-Marotte, Z. Pleunis, M. Rafiei-Ravandi, M. Rahman, S. M. Ransom, A. Renard, P. Scholz, J. R. Shaw, S. R. Siegel, K. M. Smith, I. H. Stairs, S. P. Tendulkar, I. Tretyakov, K. Vanderlinde, P. Yadav, and T. C. Collaboration, Nature **566**, 235 (2019).

[45] M. Amiri, K. Bandura, M. Bhardwaj, P. Boubel, M. M. Boyce, P. J. Boyle, C. Brar, M. Burhanpurkar, P. Chawla, J. F. Cliche, D. Cubranic, M. Deng, N. Denman, M. Dobbs, M. Fandino, E. Fonseca, B. M. Gaensler, A. J. Gilbert, U. Giri, D. C. Good, M. Halpern, D. Hanna, A. S. Hill, G. Hinshaw, C. Höfer, A. Josephy, V. M. Kaspi, T. L. Landecker, D. A. Lang, K. W. Masui, R. Mckinven, J. Mena-Parra, M. Merryfield, N. Milutinovic, C. Moatti, A. Naidu, L. B. Newburgh, C. Ng, C. Patel, U. Pen, T. Pinsonneault-Marotte, Z. Pleunis, M. Rafiei-Ravandi, S. M. Ransom, A. Renard, P. Scholz, J. R. Shaw, S. R. Siegel, K. M. Smith, I. H. Stairs, S. P. Tendulkar, I. Tretyakov, K. Vanderlinde, P. Yadav, and T. C. Collaboration, Nature **566**, 230 (2019).

[46] B. Andersen, K. Bandura, M. Bhardwaj, P. Boubel, M. Boyce, P. Boyle, C. Brar, T. Cassanelli, P. Chawla, D. Cubranic, *et al.*, The Astrophysical Journal Letters **885**, L24 (2019).

[47] M. Amiri, B. C. Andersen, K. M. Bandura, M. Bhardwaj, P. J. Boyle, C. Brar, P. Chawla, T. Chen, J. F. Cliche, D. Cubranic, M. Deng, N. T. Denman, M. Dobbs, F. Q. Dong, M. Fandino, E. Fonseca, B. M. Gaensler, U. Giri, D. C. Good, M. Halpern, J. W. T. Hessels, A. S. Hill, C. Höfer, A. Josephy, J. W. Kania, R. Karuppusamy, V. M. Kaspi, A. Keimpema, F. Kirsten, T. L. Landecker, D. A. Lang, C. Leung, D. Z. Li, H. H. Lin, B. Marcote, K. W. Masui, R. Mckinven, J. Mena-Parra, M. Merryfield, D. Michilli, N. Milutinovic, A. Mirhosseini,

A. Naidu, L. B. Newburgh, C. Ng, K. Nimmo, Z. Paragi, C. Patel, U. L. Pen, T. Pinsonneault-Marotte, Z. Pleunis, M. Rafiei-Ravandi, M. Rahman, S. M. Ransom, A. Renard, P. Sanghavi, P. Scholz, J. R. Shaw, K. Shin, S. R. Siegel, S. Singh, R. J. Smegal, K. M. Smith, I. H. Stairs, S. P. Tendulkar, I. Tretyakov, K. Vanderlinde, H. Wang, X. Wang, D. Wulf, P. Yadav, A. V. Zwaniga, and T. C. Collaboration*, Nature **582**, 351 (2020).

[48] B. C. Andersen, K. M. Bandura, M. Bhardwaj, A. Bij, M. M. Boyce, P. J. Boyle, C. Brar, T. Cassanelli, P. Chawla, T. Chen, J. F. Cliche, A. Cook, D. Cubranic, A. P. Curtin, N. T. Denman, M. Dobbs, F. Q. Dong, M. Fandino, E. Fonseca, B. M. Gaensler, U. Giri, D. C. Good, M. Halpern, A. S. Hill, G. F. Hinshaw, C. Höfer, A. Josephy, J. W. Kania, V. M. Kaspi, T. L. Landecker, C. Leung, D. Z. Li, H. H. Lin, K. W. Masui, R. Mckinven, J. Mena-Parra, M. Merryfield, B. W. Meyers, D. Michilli, N. Milutinovic, A. Mirhosseini, M. Münchmeyer, A. Naidu, L. B. Newburgh, C. Ng, C. Patel, U. L. Pen, T. Pinsonneault-Marotte, Z. Pleunis, B. M. Quine, M. Rafiei-Ravandi, M. Rahman, S. M. Ransom, A. Renard, P. Sanghavi, P. Scholz, J. R. Shaw, K. Shin, S. R. Siegel, S. Singh, R. J. Smegal, K. M. Smith, I. H. Stairs, C. M. Tan, S. P. Tendulkar, I. Tretyakov, K. Vanderlinde, H. Wang, D. Wulf, A. V. Zwaniga, and T. C. Collaboration, Nature **587**, 54 (2020).

[49] K. Vanderlinde, K. Bandura, L. Belostotski, R. Bond, P. Boyle, J. Brown, H. Chiang, M. Dobbs, B. Gaensler, G. Hinshaw, *et al.*, arXiv preprint arXiv:1911.01777 (2019).

[50] Z. Arzoumanian, A. Brazier, S. Burke-Spolaor, S. J. Chamberlin, S. Chatterjee, B. Christy, J. M. Cordes, N. J. Cornish, K. Crowter, P. B. Demorest, *et al.*, The Astrophysical Journal **821**, 13 (2016).

[51] B. Balick and R. L. Brown, The Astrophysical Journal **194**, 265 (1974).

[52] R. A. Hulse and J. H. Taylor, The Astrophysical Journal **195**, L51 (1975).

[53] K. Akiyama, A. Alberdi, W. Alef, K. Asada, R. Azulay, A.-K. Baczko, D. Ball, M. Baloković, J. Barrett, D. Bintley, *et al.*, The Astrophysical Journal Letters **875**, L4 (2019).

[54] K. Masui, E. Switzer, N. Banavar, K. Bandura, C. Blake, L.-M. Calin, T.-C. Chang, X. Chen, Y.-C. Li, Y.-W. Liao, *et al.*, The Astrophysical Journal Letters **763**, L20 (2013).

[55] A. Offringa, B. McKinley, N. Hurley-Walker, F. Briggs, R. Wayth, D. Kaplan, M. Bell, L. Feng, A. Neben, J. Hughes, *et al.*, Monthly Notices of the Royal Astronomical Society **444**, 606 (2014).

[56] M. P. Haynes and R. Giovanelli, The Astronomical Journal **89**, 758 (1984).

[57] J. W. Baars, *The paraboloidal reflector antenna in radio astronomy and communication*, Vol. 348 (Springer, 2007).

[58] J. J. Condon and S. M. Ransom, *Essential radio astronomy* (Princeton University Press, 2016).

[59] A. R. Thompson, J. M. Moran, and G. W. Swenson, *Interferometry and synthesis in radio astronomy* (Springer Nature, 2017).

[60] A. Liu, M. Tegmark, S. Morrison, A. Lutomirski, and M. Zaldarriaga, Monthly Notices of the Royal Astronomical Society **408**, 1029 (2010).

[61] D. E. Gary, "Lecture (njit radio astronomy); fourier synthesis imaging; radio astronomy," .

[62] D. Crichton, M. Aich, A. Amara, K. Bandura, B. A. Bassett, C. Bengaly, P. Berner, S. Bhatporia, M. Bucher, T.-C. Chang, *et al.*, arXiv preprint arXiv:2109.13755 (2021).

[63] B. Saliwanchik, K. Bandura, H. C. Chiang, D. Chichton, A. Ewall-Wice, E. Kuhn, V. MacKay, K. Moodley, L. Newburgh, J. Peterson, *et al.*, in *Ground-based and Airborne Telescopes VIII*, Vol. 11445 (International Society for Optics and Photonics, 2020) p. 114455O.

[64] H. Schuh and D. Behrend, Journal of Geodynamics **61**, 68 (2012).

[65] D. R. DeBoer, A. R. Parsons, J. E. Aguirre, P. Alexander, Z. S. Ali, A. P. Beardsley, G. Bernardi, J. D. Bowman, R. F. Bradley, C. L. Carilli, *et al.*, Publications of the Astronomical Society of the Pacific **129**, 045001 (2017).

[66] R. Booth and J. Jonas, African Skies **16**, 101 (2012).

[67] M. E. Jones, A. C. Taylor, M. Aich, C. Copley, H. C. Chiang, R. Davis, C. Dickinson, R. Grumitt, Y. Hafez, H. M. Heilgendorff, *et al.*, Monthly Notices of the Royal Astronomical Society **480**, 3224 (2018).

[68] P. E. Dewdney, P. J. Hall, R. T. Schilizzi, and T. J. L. Lazio, Proceedings of the IEEE **97**, 1482 (2009).

[69] V. Ravi, N. Battaglia, S. Burke-Spolaor, S. Chatterjee, J. Cordes, G. Hallinan, C. Law, T. J. W. Lazio, K. Masui, M. McQuinn, *et al.*, arXiv preprint arXiv:1903.06535 (2019).

[70] M. Amiri, B. C. Andersen, K. Bandura, S. Berger, M. Bhardwaj, M. M. Boyce, P. Boyle, C. Brar, D. Breitman, T. Cassanelli, *et al.*, arXiv preprint arXiv:2106.04352 (2021).

[71] R. Morganti, E. M. Sadler, S. J. Curran, *et al.*, arXiv preprint arXiv:1501.01091 (2015).

[72] M. D. Niemack, P. A. Ade, J. Aguirre, F. Barrientos, J. Beall, J. Bond, J. Britton, H. Cho, S. Das, M. Devlin, *et al.*, in *Millimeter, Submillimeter, and Far-*

Infrared Detectors and Instrumentation for Astronomy V, Vol. 7741 (International Society for Optics and Photonics, 2010) p. 77411S.

[73] T. Abbott, F. B. Abdalla, J. Aleksić, S. Allam, A. Amara, D. Bacon, E. Balbinot, M. Banerji, K. Bechtol, A. Benoit-Lévy, *et al.*, Monthly Notices of the Royal Astronomical Society **460**, 1270 (2016).

[74] S. J. Thomas, J. Barr, S. Callahan, A. W. Clements, F. Daruich, J. Fabrega, P. Ingraham, W. Gressler, F. Munoz, D. Neill, *et al.*, in *Ground-based and Airborne Telescopes VIII*, Vol. 11445 (International Society for Optics and Photonics, 2020) p. 114450I.

[75] D. Alonso, P. G. Ferreira, M. J. Jarvis, and K. Moodley, Physical Review D **96**, 043515 (2017).

[76] M. Amiri, K. Bandura, T. Chen, M. Deng, M. Dobbs, M. Fandino, S. Foreman, M. Halpern, A. S. Hill, G. Hinshaw, *et al.*, arXiv preprint arXiv:2202.01242 (2022).

[77] K. M. Smith, O. Zahn, and O. Dore, Physical Review D **76**, 043510 (2007).

[78] D. Hanson, S. Hoover, A. Crites, P. Ade, K. Aird, J. Austermann, J. Beall, A. Bender, B. Benson, L. Bleem, *et al.*, Physical review letters **111**, 141301 (2013).

[79] R. Pool, “Drowning in data,” .

[80] J. C. Pober, A. R. Parsons, J. E. Aguirre, Z. Ali, R. F. Bradley, C. L. Carilli, D. DeBoer, M. Dexter, N. E. Gugliucci, D. C. Jacobs, *et al.*, The Astrophysical Journal Letters **768**, L36 (2013).

[81] J. R. Shaw, K. Sigurdson, M. Sitwell, A. Stebbins, and U.-L. Pen, Physical Review D **91**, 083514 (2015).

[82] K. Bandura *et al.*, in *Ground-based and Airborne Telescopes V*, Vol. 9145 (International Society for Optics and Photonics, 2014) p. 914522.

[83] J. Ruze, Proceedings of the IEEE **54**, 633 (1966).

[84] A. Ewall-Wice *et al.*, The Astrophysical Journal **831**, 196 (2016), arXiv:1602.06277 [astro-ph.CO] .

[85] H. Hase, D. Behrend, C. Ma, B. Petrachenko, H. Schuh, and A. Whitney, in *Launching the Next-Generation IVS Network", IVS 2012 General Meeting Proceedings, NASA/CP-2012-217504* (2012) pp. 8–12.

[86] M. Deng, D. Campbell-Wilson, and the CHIME Collaboration (IEEE, 2014).

[87] J. Mena, K. Bandura, J.-F. Cliche, M. Dobbs, A. Gilbert, and Q. Y. Tang, Journal of Instrumentation **8**, T10003 (2013).

[88] K. Bandura, A. Bender, J. Cliche, T. de Haan, M. Dobbs, A. Gilbert, S. Griffin, G. Hsyu, D. Ittah, J. M. Parra, *et al.*, Journal of Astronomical Instrumentation **5**, 1641005 (2016).

[89] J. D. Jackson, "Classical electrodynamics," (1999).

[90] C. A. Balanis, *Antenna theory: analysis and design* (John wiley & sons, 2016).

[91] D. M. Pozar, *Microwave Engineering* (John Wiley & Sons, 2009).

[92] B. F. Burke, F. Graham-Smith, and P. N. Wilkinson, *An introduction to radio astronomy* (Cambridge University Press, 2019).

[93] J. Robertson, Australian Journal of Physics **44**, 729 (1991).

[94] E. R. Kuhn, B. R. Saliwanchik, M. Harris, M. Aich, K. Bandura, T.-C. Chang, H. C. Chiang, D. Crichton, A. Ewall-Wice, A. A. Gumba, *et al.*, arXiv preprint arXiv:2101.06337 (2021).

[95] R. H. Dicke, Review of Scientific Instruments **17**, 268 (1946), https://doi.org/10.1063/1.1770483 .

[96] J. W. Lamb, International Journal of Infrared and Millimeter Waves **17**, 1997 (1996).

[97] G. Reesor, I. Dagg, and S. Mohabir, Canadian Journal of Physics **53**, 2611 (1975).

[98] P. Smith, L. Davis, T. Button, and N. M. Alford, Superconductor Science and Technology **4**, 128 (1991).

[99] S. N. Paine, D. D. Turner, and N. Küchler, Journal of Atmospheric and Oceanic Technology **31**, 647 (2014).

[100] M. N. Islam and D. Ölçek, in *Ground-based and Airborne Telescopes VIII*, Vol. 11445 (International Society for Optics and Photonics, 2020) p. 114454J.

[101] C. Chang, C. Monstein, *et al.*, Publications of the Astronomical Society of the Pacific **127**, 1131 (2015).

[102] D. C. Jacobs, J. Burba, L. Turner, and B. Stinnett, in *2016 IEEE Conference on Antenna Measurements & Applications (CAMA)* (IEEE, 2016) pp. 1–3.

[103] G. Virone, F. Paonessa, E. Capello, O. Peverini, G. Addamo, R. Tascone, R. Orta, M. Orefice, A. Lingua, M. Piras, *et al.*, in *2016 IEEE Conference on Antenna Measurements & Applications (CAMA)* (IEEE, 2016) pp. 1–3.

[104] R. Walker, "Dishes, drones, &amp; dark energy: The process of building a testbed in new haven for the calibration of hirax and chime," (2020).

[105] A. R. Polish, "High-precision gps for drone-based calibration of radio telescopes," (2021).

[106] J. Zhang, J. Liu, F. Wu, X. Chen, J. Li, P. T. Timbie, S. Das, R. Yan, J. He, and O. Calvin, IEEE Antennas and Propagation Magazine **63**, 98 (2021).

[107] M. Harris, "Development of a noise source pulse modulation board for drone calibration of radio telescopes," (2021).

[108] J. Kocz, V. Ravi, M. Catha, L. D'Addario, G. Hallinan, R. Hobbs, S. Kulkarni, J. Shi, H. Vedantham, S. Weinreb, *et al.*, Monthly Notices of the Royal Astronomical Society **489**, 919 (2019).

[109] P. Berger, L. B. Newburgh, M. Amiri, K. Bandura, J.-F. Cliche, L. Connor, M. Deng, N. Denman, M. Dobbs, M. Fandino, *et al.*, in *Ground-based and Airborne Telescopes VI*, Vol. 9906 (International Society for Optics and Photonics, 2016) p. 99060D.

[110] P. O'Connor, A. Slosar, M. Harris, J. Haupt, J. Kuczewski, E. Kuhn, L. Newburgh, A. Polish, B. Saliwanchik, C. Sheehy, *et al.*, in *Ground-based and Airborne Telescopes VIII*, Vol. 11445 (SPIE, 2020) pp. 1326–1341.

[111] A. Slosar, Z. Ahmed, D. Alonso, M. A. Amin, R. Ansari, E. J. Arena, K. Bandura, N. Battaglia, J. Blazek, P. Bull, *et al.*, Journal of Environmental Sciences (China) English Ed (2019).