

©Copyright 2021
Raphael Cervantes

A Search for Wavelike Dark Matter with Dielectrically-loaded Multimode Cavities

Raphael Cervantes

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2021

Reading Committee:

Gray Rybka, Chair

Leslie Rosenberg

Masha Baryakhtar

Program Authorized to Offer Degree:

Physics

University of Washington

Abstract

A Search for Wavelike Dark Matter with Dielectrically-loaded Multimode Cavities

Raphael Cervantes

Chair of the Supervisory Committee:
Associate Professor Gray Rybka
Physics

Dark matter makes up 85% of the matter in the universe and 27% of its energy density, but we don't know what comprises dark matter. There are several compelling candidates for dark matter that have wavelike properties, including axions and dark photons. Wavelike dark matter can be detected using ultra-sensitive microwave cavities. The ADMX experiment uses a cylindrical cavity operating at the fundamental mode to search for axions in the few μeV mass range. However, the ADMX search technique becomes increasingly challenging with increasing axion mass. This is because higher masses require smaller-diameter cavities, and a smaller cavity volume reduces the signal strength. Thus, there is interest in developing more sophisticated resonators to overcome this problem.

The ADMX-Orpheus experiment uses a dielectric-loaded Fabry-Perot cavity to search for axions and dark photons with masses approaching $100 \mu\text{eV}$. Orpheus maintains a large volume by operating at a higher-order mode, and the dielectrics shape the electric field so that the mode couples more strongly to the axion and dark photon.

This thesis describes the development and commissioning of ADMX-Orpheus to search for dark photons with masses between $65.5 \mu\text{eV}$ and $69.3 \mu\text{eV}$. This thesis includes

- Motivation for why dielectric cavities are suitable for detecting axions and dark photons around $100 \mu\text{eV}$.

- The design and characterization of such a cavity, including the resonant frequency, quality factor, and detection volume of the mode of interest.
- The mechanical design, electronics, and data acquisition system for the Orpheus experiment.
- The inaugural search for dark photons $65.5 \mu\text{eV}$ and $69.3 \mu\text{eV}$ and the resulting excluded parameter space.
- Plans for upgrading Orpheus to search for axions in a similar mass range.

TABLE OF CONTENTS

| | Page |
|----------------------------------------------------------------------------------|------|
| List of Figures | iii |
| Chapter 1: Introduction | 1 |
| 1.1 Dark Matter Problem | 1 |
| 1.2 Axions as Dark Matter | 2 |
| 1.3 ALPs and Dark Photons | 11 |
| Chapter 2: Dielectric Haloscopes for Detecting Axions and Dark Photons | 15 |
| 2.1 Haloscopes to search for wavelike dark matter | 15 |
| 2.2 Haloscopes for Dark Photons | 19 |
| 2.3 Scaling the Haloscope Concept to Higher Frequencies | 20 |
| 2.4 Multimode Dielectric Haloscopes and Orpheus Conceptual Design | 21 |
| 2.5 Other related experiments | 23 |
| Chapter 3: Orpheus RF Design and Characterization | 24 |
| 3.1 The Empty Fabry-Perot Cavity | 24 |
| 3.2 Adding dielectrics to a Fabry-Perot Cavity | 37 |
| 3.3 Coupling to Orpheus | 51 |
| Chapter 4: Orpheus Cryogenic and Mechanical Design | 59 |
| 4.1 Orpheus Cavity Mechanical Design | 59 |
| 4.2 Insert | 63 |
| 4.3 Issues with the Tuning Mechanism | 69 |
| 4.4 Thermal Gradient Issues | 71 |
| Chapter 5: Electronics, Data Acquisition System, and Operations | 72 |
| 5.1 Cold Electronics | 72 |

| | | |
|--------------------------------------------------|-------------------------------------------------------------------------------------|-----|
| 5.2 | Warm Electronics | 73 |
| 5.3 | System Noise Temperature | 77 |
| 5.4 | Receiver Noise Temperature | 78 |
| 5.5 | Software Stack | 80 |
| 5.6 | Data-taking cadence | 81 |
| 5.7 | September 1st Commissioning Narrative | 82 |
| Chapter 6: Dark Photon Search Analysis | | 89 |
| 6.1 | Confirmation that the DAQ Followed the Correct Mode | 91 |
| 6.2 | Confirmation of a Healthy Receiver | 94 |
| 6.3 | Parameter extraction | 97 |
| 6.4 | Estimating expected sensitivity | 103 |
| 6.5 | Processing Individual Raw Spectrum | 104 |
| 6.6 | Combined Spectrum | 112 |
| 6.7 | Placing 90% Exclusion Limit | 113 |
| Chapter 7: Future Direction | | 119 |
| 7.1 | Upcoming Axion Run | 119 |
| 7.2 | More Science with Same Apparatus | 122 |
| 7.3 | How to reach KSVZ and DFSZ | 125 |
| Chapter 8: Conclusion | | 127 |
| Appendix A: Reflection Fitting | | 135 |
| A.1 | First pass at a reflection fit | 135 |
| A.2 | Deriving the reflection off of the cavity | 139 |
| A.3 | Model of a resonator connected to directional coupler and amplifier | 142 |
| A.4 | Axion coupling sensitivity to uncertainty in antenna coupling coefficient | 147 |
| A.5 | Magnet Design | 148 |

LIST OF FIGURES

| Figure Number | Page |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|
| 1.1 A neutron with an electric dipole moment. This system violates both time-reversal symmetry and parity symmetry, and therefore violates charge-parity reversal symmetry.[10] | 3 |
| 1.2 The effective potential for the axion field. It resembles a tilted Mexican hat.[12] | 5 |
| 1.3 The Feynmann interaction for axions coupling to photons. | 5 |
| 1.4 Left: PQ symmetry is broken before inflation and never restored. At inflation, each “patch” of initial axion value is blown up. Our observable universe sits in one of those patches. Thus our observable universe has only experienced one value for the axion (with some quantum perterbations). Right: PQ symmetry was broken after inflation. The axion field starts with random values throughout the observable universe. Our universe will contain topological defects like axion strings[13]. | 8 |
| 1.5 A figure from 2010 that shows how astrophysical sources place limits on axion mass. [14] | 10 |
| 1.6 The excluded parameter space for axions, zoomed into microwave cavity haloscopes. The bottom plot shows different theoretical predictions for post-inflationary axions. A lot of the tensions have to do with how much do topological defects contribute to the axion population. Figure courtesy of Ciaran O’Hare[16]. | 12 |
| 2.1 Haloscope search for axions. An axion, when it sees a strong magnetic field, can convert to a photon. If the photon frequency matches the resonant frequency of the cavity, the axion signal power is resonantly enhanced. A haloscope search for dark photons works the same way, except that a dark photon can convert to an SM photon without the presence of a magnetic field. Figure courtesy of Christian Boutan. | 17 |
| 2.2 Traditional haloscope signal power scale poorly with frequency. (a) Small cavity operating at the fundamental mode, e.g., TM_{001} mode. The volume of such a cavity scales as $\sim f^{-3}$ for the same aspect ratio. (b) Operating a large cavity at higher frequencies requires operating at higher order modes. The effective volume $V_{eff} \propto \int dV \mathbf{E}_a \cdot \mathbf{B}_o \approx 0$ | 21 |

| | | |
|-----|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----|
| 2.3 | A multimode dielectrically-loaded cavity. Dielectrics are placed where the electric field is anti-aligned with the external magnetic field. Dielectrics suppress electric fields, so $V_{eff} \propto \int dV \vec{E} \cdot \vec{B} > 0$ | 23 |
| 3.1 | (a) The Hermite-Gaussian modes of a Fabry-Perot cavity. Figure from Kogelnik[40]. (b) FEA simulation of the TEM ₀₀₋₁₈ mode in an empty Fabry-Perot cavity. | 25 |
| 3.2 | A comparison of simulated quality factor of the TEM ₀₀₋₁₈ mode with different sets of loss mechanisms. The blue curve corresponds to a cavity with perfectly-conducting mirrors. The orange curve corresponds to a cavity with resistive aluminum mirrors. They are nearly identical, suggesting that the dominant loss is diffraction. | 31 |
| 3.3 | The empty Fabry-Perot cavity consists of a flat mirror and a curved mirror. The cavity coupling scheme shown in this picture is out of date. This picture is just meant to show the cavity without dielectric plates. | 32 |
| 3.4 | The reflection coefficient of the cavity's strongly-coupled port is measured using a directional coupler. S_{11} can't be measured directly because the amplifier is a non-reciprocal device. The directional coupler allows an injected signal to bypass the amplifier to reach the cavity's strongly-coupled port. The signal from the VNA travels from port 3 of the directional coupler to port 1 (with a 20 dB attenuation), bounces off the cavity's strongly-coupled port, then reaches port 2 of the directional coupler. The signal is amplified and then reaches the VNA's port 2. | 33 |
| 3.5 | The measured scattering parameters S_{21} and S_{11} of an empty Fabry-Perot cavity as a function of frequency for two different cavity lengths. The prominent peaks are Lorentzian and correspond to different Gaussian modes. | 34 |
| 3.6 | (a) Measured mode map for an empty Fabry-Perot Cavity. This is the measured scattering parameter S_{21} as a function of frequency and cavity length. The dark lines correspond to the modes of the cavity where transmission is highest. (b) The analytical TEM modes and the simulated modes are overlaid on the measurement. The agreement between the analytical formula, FEA simulation resonances, and measured resonances are within a few MHz. | 35 |
| 3.7 | S_{21} and S_{11} of an empty Fabry-Perot cavity near the TEM ₀₀₋₁₈ mode for two different cavity lengths. The transmitted power is a Lorentzian. The reflected power is approximately Lorentzian, but effects discussed in Section A.3 distort the Lorentzian. Because of the skew in the data, the fitted Lorentzian in (a) predicts a small negative reflected power at resonance, which is unphysical | 36 |

| | | |
|------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----|
| 3.8 | The loaded quality factor and cavity coupling coefficient of the TEM_{00-18} mode as a function of cavity length. | 36 |
| 3.9 | The simulated electric field magnitude of the TEM_{00-18} mode with different cavity lengths. There are mode crossings around 15.4 GHz and 16.3 GHz. . . . | 39 |
| 3.10 | The simulated electric field magnitude of the TEM_{00-18} mode with different cavity lengths. The field profile starts to deviate from what is expected from physical optics. One intuit from physical optics that the waist of the beam would be on the flat mirror instead of the curved mirror. But that is not the case. | 40 |
| 3.11 | Intruder mode while it's separated from the TEM_{00-18} mode. Here, the cavity length is $L = 14.4219$ cm. The TEM_{00-18} resonance is at $f_{00-18} = 17.07$ GHz while the intruder mode frequency is at 16.86 GHz. They will cross around $L \approx 14.95$ cm, $f_{00-18} \approx 16.4$ GHz. Left: Magnitude of the intruder mode electric field. Right: Complicated vector field superimposed on the magnitude of the electric field. The field is complicated, but it looks like some mixture of a TM mode (based on the field pointing into the mirror) and a TEM mode (based on the vector field being transverse to the axis for much of the interior of the cavity). TM modes don't tune with the changing cavity length, which may explain why the intruder mode tunes slower than the TEM_{00-18} mode. | 41 |
| 3.12 | The simulated frequency and loaded quality factor for the TEM_{00-18} mode. (a) The simulations struggles to simulate the mode crossing, hence the discontinuity in resonant frequency at around 16.4 GHz. Perhaps we need higher accuracy requirements. (b) The simulated loaded quality factor. In our simulations, the cavity is severely undercoupled, so this is close to the unloaded quality factor. | 42 |
| 3.13 | The effective volume and form factor for the TEM_{00-18} mode. The sharp dip at 16.4 GHz corresponds to the mode crossing. | 43 |
| 3.14 | The effect of dielectric plate position error on the electric field, simulated at about 15.82 GHz. (a) No position error. (b) $\delta_{top} = -0.52$ mm, $\delta_{bottom} = -2$ mm. Note the δ_{bottom} corresponds to the bottom dielectric plate closest to the curved mirror, so it appears as the uppermost plate in this figure. | 45 |
| 3.15 | The effect of position error on the effective volume. This study suggests that it would actually be more optimal to deviate from the evenly-spaced configuration. At 15.82 GHz, it appears that the effective volume can increase by more than 8% by moving the dielectrics further away from the curved mirror. | 46 |
| 3.16 | The effect of position error on the form factor. | 46 |
| 3.17 | The tabletop Orpheus cavity. | 47 |

| | | |
|------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----|
| 3.18 | The measured scattering parameters S_{21} and S_{11} of an dielectrically-loaded Fabry-Perot cavity as a function of frequency for a fixed cavity length. The prominent peaks are Lorentzian and correspond to different Gaussian modes. | 48 |
| 3.19 | Measured mode map for dielectrically-loaded Fabry-Perot cavity. This is the measured scattering parameter S_{21} as a function of frequency and cavity length. The dark lines correspond to the modes of the cavity where transmission is highest. (b) The simulated modes are overlaid on the measurement. | 49 |
| 3.20 | The measured transmission wide scans are compared to the simulated wide scans. The general shape between the two matches. There is some discrepancy in the location of the resonances. This may be improved by demanding more accuracy from the wideband simulation, but this improving the agreement doesn't have much bearing on the science results. | 50 |
| 3.21 | The measured loaded quality factor and cavity coupling coefficient of the TEM ₀₀₋₁₈ mode as a function of cavity length. Measurements are done in air at room temperature. | 51 |
| 3.22 | The simulated Q_u is compared to the measured Q_u . The quality factor is notoriously difficult to simulate, so the values don't match. But the general shape between the two curves are similar. | 52 |
| 3.23 | Two dipole antennas were inserted into the cavity through the flat mirror and bent so they are transverse to the axis of the cavity. The antenna in the center of the mirror corresponds to the strongly-coupled port; the antenna further from the center corresponds to the weakly-coupled port. | 53 |
| 3.24 | The mode map of an empty Fabry-Perot cavity and dielectrically-loaded Fabry-Perot cavity that's coupled using coaxial antennas. (a) The doubly-degenerate TEM modes are split into two. The modes with the polarization that most closely lines up with the orientation of the dipole antennas are pushed lower in frequency by about 100 MHz compared to where they are expected to be from the analytical formula. The other polarization is pushed to higher frequencies than the analytical expectation and couples more weakly to the antennas. (b) The TEM ₀₀₋₁₈ mode is lower than the simulated prediction by about 100 MHz. | 54 |
| 3.25 | The waveguide fits into a mirror pocket shown in (a). The mirror thickness at the pocket is 0.15 in. (b) The thin mirror surface improves the cavity coupling, but resulted in warping of the mirror surface around the beam axis. This likely reduced the Q. | 55 |

| | | |
|------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----|
| 3.26 | Having the strongly-coupled port and weakly-coupled port right next to each other resulted in significant crosstalk. The crosstalk interferes with the cavity signal, which causes the Lorentzian peaks to distort. In the most perverse case, the crosstalk and cavity signal interfere perfectly destructively, causing a dip in the S_{21} measurement where a Lorentzian peak is expected. | 56 |
| 3.27 | (a) S_{21} can be plotted in the complex plane. A cavity without any crosstalk or phase shift will have the S_{21} in the canonical position. Crosstalk will cause the q-circle to be shifted from this canonical position[44]. (b) This shift is observed for the measured S_{21} for the configuration when the strongly-coupled and weakly-coupled ports were next to each other. | 56 |
| 3.28 | Moving the weakly-coupled port to the other side of the cavity mitigated crosstalk. | 57 |
| 3.29 | Waveguide tuner used to attempt to create an impedance matching network. I was unsuccessful in using it in that manner. | 58 |
| 4.1 | CAD model for the Orpheus cavity. | 60 |
| 4.2 | CAD model for one of the dielectric plates. | 61 |
| 4.3 | The dielectric holder assembly without the dielectrics. There is a lip inside the octagon where the dielectrics freely rest. Note that the thrust bearings in this photo are 0.25in long are shorter than the final bearing length of 0.615in. | 62 |
| 4.4 | Orpheus cavity attached to the cryogenic insert. | 63 |
| 4.5 | Orpheus cryogenic insert. | 65 |
| 4.6 | Cryogenic insert in real life. | 66 |
| 4.7 | Orpheus gearbox | 67 |
| 4.8 | Insert top flange with all the vacuum ports, cryogen ports, and motor stage. | 68 |
| 4.9 | (a)The vacuum ports for the rotary feedthroughs are misaligned. This contributed to a cascade of effects that caused the motors to stall frequently after the cooldown. (b) The vacuum shaft consists of a stainless steel tube connected to a G10 fiberglass tube connected by a spiral shaft coupling with a 5/8" outer diameter. The vacuum housing tube has a 3/4" inner diameter, and the misalignment from the top flange caused the shaft coupler to touch the vacuum walls. The point of contact is made obvious by the localized frost marks. | 70 |

| | | |
|------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----|
| 5.1 | A diagram of how transmission, reflection, and power measurements are taken. Room temperature Teledyne switches are used to switch between transmission and reflection measurements, and from VNA measurements to power measurements. | 73 |
| 5.2 | Orpheus measurement scheme. There are 3 paths: transmission, reflection, and digitization. | 74 |
| 5.3 | RF/IF electronics box. Contains warm RF stage, IF stage, and coaxial switches to switch between a VNA transmission measurement, VNA reflection measurement, and a power spectrum measurement. | 74 |
| 5.4 | A demonstration of image noise after mixing. The negative image frequency folds into the positive IF frequency, and the positive image frequency folds into the negative IF frequency. Figure from [46]. | 75 |
| 5.5 | The derived receiver noise temperature. The Friis equation is used to derive the receiver noise temperature (d) from the noise temperature (a) and gain (b) of the 1st stage cryogenic LNA and the noise temperature of the 2nd stage room-temperature amplifier (c). These values are obtained from data sheets from Low Noise Factory. The cavity temperature can be added to the receiver noise temperature to obtain the system noise temperature, as referenced to the cavity. | 79 |
| 5.6 | Dripline software architecture. Figure from [48]. | 81 |
| 5.7 | Flat mirror and curved mirror temperatures throughout the entire commissioning of the experiment, including warmup and cooldown. | 83 |
| 5.8 | Flat mirror and curved mirror temperatures throughout the data-taking run. Physics data-taking stopped because the flat mirror temperature was deviating significantly from the curved mirror temperature, causing a significant increase in the uncertainty in the cavity temperature. | 84 |
| 5.9 | The TEM ₀₀₋₁₈ resonant frequency vs time. | 86 |
| 5.10 | The motor steps as recorded by the encoder. From the motor steps, one derives the distance traveled by the curved mirror, bottom dielectric plate, and top dielectric plate. The top dielectric plate motor stalled after cooldown, preventing the top dielectric plate from being tuned remotely and in a continuous manner. I needed to turn the motor feedthroughs by hand periodically (explaining the discrete tuning steps for the top dielectric plate. | 87 |
| 5.11 | Due to a cascade of mechanical failures, software design flaws, and operation errors, the top dielectric plate and bottom dielectric plate deviated from their intended positions throughout the data-taking run. | 88 |
| 5.12 | Number of spectrum containing a particular RF bin. | 88 |

| | | |
|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|
| 6.1 | Wide scan transmission measurements with the simulated mode overlaid for different cavity lengths. Because the resonator is encased in a stainless steel can, there are low-Q non-tuning modes that make the transmission measurement messier than in Figure 3.18. | 92 |
| 6.2 | Piecing together wide scan measurements to form a mode map. The discontinuities near 15 cm and 15.3 cm are from adjusting the top dielectric plate every once in a while. (b) is the same as (a), but the simulated frequency is shifted over by 0.7 mm to show the systematic error in measuring cavity length. | 92 |
| 6.3 | The TEM ₀₀₋₁₈ mode frequency vs the cavity length. The cavity length is determined from the stepper motor encoder values. | 93 |
| 6.4 | Use the VNA to inject a tone into the cavity. Both the tone frequency and the resonant frequency was at 16.07 GHz. The loaded Q was about 10,000. . . | 96 |
| 6.5 | Plotted integrated digitizer power as the cavity is warming up. For reasons outlined in the text, this measurement is difficult to interpret as a Y-factor measurement to obtain the receiver noise temperature. But the digitized power rises more-or-less linearly with the cavity temperature. This suggests a healthy, low-noise receiver chain. | 96 |
| 6.6 | Cavity temperature as a function of mode frequency. | 98 |
| 6.7 | Loaded quality factor as a function of mode frequency. | 99 |
| 6.8 | Cavity coupling as a function of mode frequency. | 99 |
| 6.9 | Unloaded quality factor as a function of frequency, as derived from the measured loaded quality factor and cavity coupling. The results from the simulation are overlaid for comparison and to identify the mode crossing in the data. The notch in the simulated Q is much wider and steeper than the measured Q. This discrepancy suggests the mode crossing in the simulation is much larger than the mode crossing of the physical cavity. | 100 |
| 6.10 | The simulated effective volume a function of mode frequency. V_{eff} is plotted with and without the mode crossing because I will have to take care of how I handle the mode crossing in the data analysis. | 101 |
| 6.11 | The deviation of the plate positions from their targeted positions. This position error was used to estimate the uncertainty in V_{eff} | 103 |
| 6.12 | Raw spectrum as measured by the digitizer. Assuming the spectrum is just noise, the raw spectrum is the noise power of the system. | 105 |

| | | |
|------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|
| 6.13 | Raw spectrum with a hypothetical dark photon signal of arbitrary power superimposed. (a) The dark photon signal is much narrower than the large-scale structure of the noise background. Thus a low-pass filter can be safely applied to each power spectrum. (b) Same as (a), but zoomed into the dark photon signal. The hypothetical dark photon signal covers about 8 RF bin widths. | 106 |
| 6.14 | The average of all the recorded spectrum. This is thought of as an good estimation of the gain variation caused by the IF electronics. The estimation is less good if you don't have even and continuous tuning through your frequency range. | 106 |
| 6.15 | The raw spectra are all divided by the IF baseline to form some normalized spectrum. A Savitzky-Golay filter is then applied (orange) to obtain gain variation from the RF electronics. | 107 |
| 6.16 | The power excess in units of system noise power. | 108 |
| 6.17 | The power excess rescaled to units of single-bin dark photon power. | 109 |
| 6.18 | A diagram of how convolution can enhance the signal of a multi-bin dark photon signal. Figure courtesy of Christian Boutan [53]. | 110 |
| 6.19 | The rescaled power excess after the dark matter halo lineshape filter has been applied. | 111 |
| 6.20 | The (a) power excess δ_p and (b) rescaled power excess δ_s has a Gaussian distribution. The filtered power excess has a smaller standard deviation than what is expected of Gaussian noise. This can be attributed to correlations induced by the convolution. | 112 |
| 6.21 | The combined spectrum. | 113 |
| 6.22 | The probability distribution of finding a dark photon signal, given a measured power excess, is represented as a truncated Gaussians. The truncated Gaussians are plotted for (a) $\delta_{fc}/\sigma_{\delta_{fc}} = 1$ and (b) $\delta_{fc}/\sigma_{\delta_{fc}} = -1$. The 90% confidence interval is shaded in the plots. The value representing the upper limit of the confidence interval is plotted for different values of $\delta_{fc}/\sigma_{\delta_{fc}}$. The confidence interval for a full normal distribution is plotted in orange for comparison. | 115 |
| 6.23 | The 90% confidence limit on the dark photon power, normalized to dark photons signal power of mixing angle $\chi = 1$ | 116 |
| 6.24 | A 90% exclusion on the mixing angle parameter space. | 117 |
| 6.25 | Orpheus science reach in the context of other microwave cavity results. The limits are for the uniformly polarized case. Figure adapted from [16]. | 118 |
| 7.1 | A 95% exclusion on the axion-photon coupling $g_{a\gamma\gamma}$ had there been a 1.5 T dipole magnet. | 120 |

| | | |
|-----|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|
| 7.2 | The projected axion limits in the context of other microwave cavity searches. Figure adapted from [16]. | 121 |
| 7.3 | (a) It is likely that we can achieve sensitivity with the same insert and electronics by swapping out dielectrics of different thicknesses and mirrors with different radius of curvature. (b) DFSZ can be reached with a similar cavity design at similar frequency ranges if the magnetic field was increased to 10 T, the effective volume increased by a factor of 10, the quality factor doubled, and the noise temperature reduced by a factor of 10. Figures adapted from [16]. | 126 |
| A.1 | The process of fitting a Lorentzian to the reflection coefficient to extract the resonant frequency and quality factor. The blue curve is the square of the measured reflected coefficient. The orange line is the initial guess of the Lorentzian. The green line is the curve fit, which tries to minimize the χ^2 | 136 |
| A.2 | The phase change of the reflection coefficient for the undercoupled case (blue), critically-coupled case (green), and overcoupled case (orange). The derivative on resonance is positive for the undercoupled case and negative for the overcoupled case. | 137 |
| A.3 | The reflection measurement off of the Orpheus cavity. The VNA sends a signal to port 3 of the directional coupler. The signal travels to port 1 and reflects off the cavity. Then the signal travels from port 1 to port 2, and then through the amplifier and back to the VNA. | 142 |
| A.4 | A simplified mathematical model for a reflection measurement. A signal is injected into the cavity. V_1 is the signal reflected off the cavity heading towards the direction of the amplifier. V_1 reaches the amplifier and is partially reflected back to the cavity. That signal is then reflected again off the cavity and is added to V_1 . There is an additional V_x term that represents the complications from the injected signal interacting with the four different ports in the directional coupler. | 143 |
| A.5 | The fitted transmitted (standard Lorentzian fit) and reflected power (perturbed Lorentzian fit, Equation A.8). | 146 |
| A.6 | The fitted loaded quality factors for transmission (standard Lorentzian fit) and reflection (perturbed Lorentzian fit, Equation A.8). | 146 |
| A.7 | The derivative of the axion-photon coupling constant vs antenna coupling coefficient. This relationship determines how the uncertainty in β affects the uncertainty in $g_{a\gamma\gamma}$ | 147 |
| A.8 | Calculated magnetic field of the planned Orpheus magnet. For this calculation, Z is the axis of the magnet, X is along the long side of the dipole magnet, and Y is along the short side of the dipole magnet. | 149 |

A.9 Practicing Orpheus magnet winding with copper wire. 150

ACKNOWLEDGMENTS

I thank my advisor, Gray Rybka, for six years of continuous and accomodating advising. He was a well of knowledge, experience, intuition, creativity, resourcefulness, and equanimity that I drew from as I built a dark matter detector from scratch. He also put me on various exciting projects in both ADMX and Project 8 before I settled on ADMX-Orpheus. He tolerated my ambitious timelines that I could never meet.

I thank everyone who has worked on Orpheus. I thank Parashar Mohapatra and James Sinnis for letting me advise and supervise them. There was always a lot to do and a lot of problems to solve. Most of the time, I didn't know what I was doing and guessed my way to a functional haloscope. They took all of this in stride. I thank Jihee Yang, Charles Hanretty, and Grant Leum for helping me put together the vacuum and cryogenic systems; Nick Du for helping me with temperature sensors, Yujin Park for designing the first version of the superheterodyne; Ben LaRoque and Noah Oblath for helping me set up and develop the DAQ; Eric Smith and Gary Holman for facility support; Rich Ottens and Seth Kimes for designing the dipole magnet and much of the mechanics; the CENPA and Physics machine shop for fabricating many components for the Orpheus resonator and insert; and Ryan Rohnelt for his excellent mechanical advice as I commissioned the insert.

I thank the ADMX group for creating a welcoming and fun work environment. I have many fond memories, in and out of the lab, that I will always cherish. In particular, I will always remember the night at the Unicorn in Capitol Hill and 1UP Arcade in Denver. At the risk of favoritism, I thank Nick Du and Chelsea Bartram for years of friendship and emotional support. And I thank them for the wonderful hikes out in the Pacific Northwest. I valued everyone else's friendship, but our time together was more brief. Finally, I thank

Leslie Rosenberg for his advice and for understanding realistic timelines. There was a time when I was falling apart because an Indium seal surface wasn't machined to spec. Leslie sat me down and gave me a homegrown orange. I will always remember that orange.

I thank Masha Baryakhtar for agreeing to be on the reading committee, for her helpful feedback, and for her patience while explaining dark photon cosmology.

I thank CENPA for being a great place to work. I enjoyed learning about all the cool tabletop experiments probing fundamental physics.

I thank the author, Anne Lamott, for giving sound life advice. I repeated the mantra "bird by bird" as problems mounted while commissioning Orpheus. Her essay about first drafts helped me power through this thesis.

I thank my roommates for being cool (especially during quarantine), for taking care of my dog while I was on travel, and for letting me have a dog. I thank my dog, Seraphine, for her enthusiastic companionship these past five years. She made me a better person and taught me that love is more of a verb than a noun.

I thank my hometown friends. In particular, I thank Rose Soto, Jesse Muñoz, and Karen Quirarte for helping me through a family crisis.

I thank my family for their continual love and support. They keep asking me how my classes are going or when I'll get married, but that's ok.

I thank my dad for a good 28 years of love and support. I miss you, and I wish you could have seen me get my Ph.D.

DEDICATION

To my father Ruben Cervantes.

Chapter 1

INTRODUCTION

1.1 Dark Matter Problem

Dark matter is the non-luminous, non-absorbing matter that makes up about 84.4% of the matter of the universe and 26.4% of its mass[1]. The Lambda cold dark matter (Λ CDM) model describes dark matter as feebly interacting, non-relativistic, non-interacting, and stable on cosmological timescales. The evidence for dark matter is abundant and includes the galactic rotation curves[2, 3], gravitational lensing[4, 5], the bullet cluster[6], and the cosmic microwave background[7]. However, despite the overwhelming evidence for the abundance of non-luminous matter in our universe, not much is known about the nature of dark matter. We don't know what makes up dark matter. Many hypothetical particles could make up dark matter. These candidates include WIMPS, sterile neutrinos, chameleons, axions, axion-like particles, and dark photons. Dark matter may be made up of a combination of many of these particles (a dark sector).

This thesis focuses on a set of dark matter candidates that have wavelike properties. For a particle to be wavelike, the de Broglie wavelength λ_{dB} is much greater than the inter-particle spacing. In other words, if the number of particles N_{dB} inside a de Broglie volume λ_{dB}^3 is very large, the set of particles is best described as a classical wave[8]. The dark matter density in our halo is fixed to $\rho \sim 0.45 \text{ GeV}/\text{cm}^3$. So the smaller the dark matter mass, the larger N_{dB} . Wavelike dark matter must be bosonic because fermions cannot occupy the same phase space.

A subset of dark matter candidates can be wavelike: axions, axion-like particles (ALPs), and dark photons. Axions are a particularly compelling candidate because they solve an outstanding problem in particle physics known as the Strong CP problem. ALPs and dark

photos aren't as motivated as the axion but are still well-motivated for other reasons and can be searched for with the same axion detector and even the same data.

I will first describe axions because Orpheus was developed in the context of an axion search. But much of the discussion applies readily to ALPs and dark photons. After my exposition on axions, I will focus on the other two candidates, emphasizing how they are different from axions.

1.2 Axions as Dark Matter

The axion is a hypothetical particle that, if shown to exist, would solve both the strong CP problem and the dark matter problem. This section explains the motivation for axions, axion physics, and axion cosmology. To summarize, the axion is a result of extending the Standard Model with a field that makes the strong force conserve CP symmetry. The new field results in a new particle called the axion, which would have been produced copiously in the early universe. If the axion is light enough, it would survive on cosmological time scales and interact feebly with normal matter. This makes the axion a compelling dark matter candidate. Astrophysical observations has suggested that the axion mass can range from 10^{-22} eV to 10^{-2} eV.

1.2.1 Strong CP Problem

The Standard Model of particle physics describes the interactions of elementary particles through the strong force, weak force, and electromagnetic force. Underlying Standard Model physics is the conservation of symmetries. For example, all standard model processes must obey charge, parity, and time reversal (CPT) symmetry, meaning that the physics of a system remains the same if one were to simultaneously invert the charges, invert the spatial coordinates, and reverse the direction of time. Even though CPT symmetry is conserved in the Standard Model, CP symmetry (charge conjugation and parity symmetry) is violated. This CP violation is caused by the complex quark masses that contain arbitrary phases. It's also expected that the strong force also violates CP symmetry. The CP violation is described

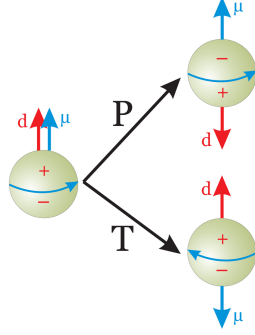


Figure 1.1: A neutron with an electric dipole moment. This system violates both time-reversal symmetry and parity symmetry, and therefore violates charge-parity reversal symmetry.[10]

by a term in the QCD Lagrangian, $L_{\bar{\Theta}} = -\bar{\Theta} \left(\frac{\alpha_s}{8\pi} \right) G^{\mu\nu a} \tilde{G}_{\mu\nu}^a$, where α_s is the strong coupling constant, $G^{\mu\nu a}$ is the gluonic field strength, $\tilde{G}_{\mu\nu}^a$ is its dual, and $\bar{\Theta}$ is an angular parameter that ranges from $-\pi$ to π . More precisely, $\bar{\Theta} = \theta - \arg \det \mathcal{M}$, where θ is a phase from the QCD vacuum and \mathcal{M} is the quark mass matrix. It's expected that $\bar{\Theta}$ be of $\mathcal{O}(1)$.

Strong CP violation would manifest itself as a neutron electric dipole moment (nEDM). Figure 1.1 demonstrates how neutron with an EDM would violate CP symmetry. If time is reversed, the magnetic dipole moment would be flipped, but the electric dipole moment would stay the same. If parity is inverted, the electric dipole moment would flip, but the magnetic dipole moment would stay the same. Thus a neutron with an EDM would violate P symmetry and T symmetry. Because CPT symmetry holds, a neutron EDM would violate CP symmetry. However, the neutron dipole moment has not been observed. The upper bound on a possible nEDM is measured to be $|d_n| < 1.8 \cdot 10^{-26} e \cdot cm$ [9]. This limits any possible strong CP violation to $\bar{\Theta} < 10^{-10}$. The Standard Model doesn't explain why $\bar{\Theta}$ is unnaturally small. This is the strong CP problem.

One might struggle to understand why the strong CP problem is actually a problem. Why can't $\bar{\Theta} = 0$? Isn't the idea of naturalness just a human construct? The absence of CP violation is troubling because $\bar{\Theta} \ll 1$ would require that the random phases in the quark mass matrix conspire to make $\theta - \arg \det \mathcal{M} \ll 1$. One should keep in mind that the random phases

in the quark mass matrix (and the neutrino mixing matrix) are responsible for CP violation in the weak force, and that the weak force is maximally CP violating. So $\theta - \arg \det \mathcal{M} = 0$ would require that the CP violation from the strong force (the θ contribution) and weak force (the $\arg \det \mathcal{M}$ contribution) cancel each other exactly, even though the strong force and weak force have no relation to each other. The absence of a neutron dipole moment is also strange since this would require that the neutron's valence quarks and ever-fluctuating sea quarks are always arranged to not have an electric dipole moment. One ought to demand a mechanism for this coincidental arrangement, but the Standard Model does not provide one.

1.2.2 Peccei-Quinn Mechanism to Solve Strong CP problem

A few solutions have been proposed to solve the Strong CP problem. One elegant solution is to make the up quark massless. A massless up quark would mean $\det \mathcal{M} = 0$, implying that $\arg \det \mathcal{M}$ is unphysical. Quantum fields can then be redefined, and θ can be rotated to remove the CP violating term in the QCD Lagrangian. However, measurements and lattice QCD simulations have determined the mass of the up quark to be $m_u \sim 2 \text{ MeV}$, making the massless up quark solution untenable. Another proposed solution is the Nelson-Barr mechanism, but it requires a fine-tuning more perverse than $\bar{\Theta} = 0$ [11].

The most popular solution to the strong CP problem is the Peccei-Quinn mechanism. In this solution, one extends the standard model by adding a complex scalar field ψ with a $U_{PQ}(1)$ symmetry (meaning that the physics of this field is invariant under a rotational transformation $\psi \rightarrow \psi \exp(i\alpha)$). When the system cools below a critical temperature, the field undergoes a phase transition and the $U_{PQ}(1)$ symmetry is spontaneously broken. The potential associated with this field becomes a Mexican hat potential and the angular mode is a Goldstone boson called the axion¹. Once the temperature approaches the QCD scale, the field starts interacting with gluons and the $U_{PQ}(1)$ symmetry is explicitly broken. The Mexican

¹It might help to know that particles in QFT are excitations of a quantum field's eigenmodes.

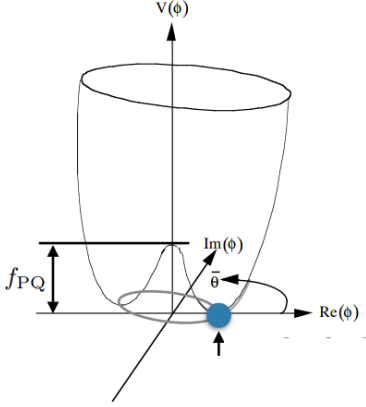


Figure 1.2: The effective potential for the axion field. It resembles a tilted Mexican hat.[12]

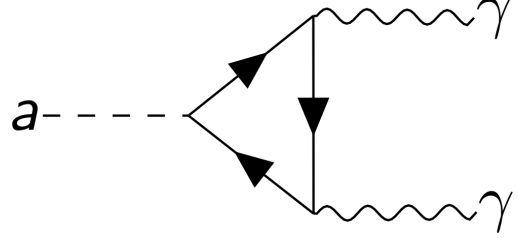


Figure 1.3: The Feynmann interaction for axions coupling to photons.

hat potential becomes tilted and now has a minimum potential (as shown in Figure 1.2). The axion now has mass. The CP violating term in the Lagrangian is modified to $L_\Theta = \left(\frac{a}{f_A} - \bar{\Theta}\right) \left(\frac{\alpha_S}{8\pi}\right) G^{\mu\nu a} \tilde{G}_{\mu\nu}^a$, where f_a is the Peccei-Quinn energy scale, and a is the axion field. The axion particle naturally relaxes to the potential minimum where the CP violation goes away. The effective potential for the axion field is $V(a) = f_a^2 m_a^2(T) \left(1 - \cos \frac{a}{f_a}\right)$. The mass of the axion is formally defined as the Taylor expansion of the effective potential, $m_a^2 = \left.\frac{\partial^2 V(a)}{\partial a^2}\right|_{a=f\bar{\Theta}}$. QCD calculations show that $m_a = 5.691(51) \left(\frac{10^9 \text{ GeV}}{f_a}\right) \text{ meV}$.

Even though axions were invented to solve the strong CP problem, they happen to also be an excellent candidate for dark matter. A proper dark matter candidate would have existed before the epoch of recombination, would be stable on cosmological timescales, and would be feebly interacting with normal matter. As will be explained in Section 1.2.3, cold (i.e., non-relativistic) axions would have been produced copiously in the early universe. Also, if the axion mass is small enough, then axions from the early universe would still be around today. That is because the axion oscillation period is inversely proportional to the axion mass², so a light axion would still be oscillating today, even in the presence of a

²Near the potential minimum, $V \propto m_a^2 a^2$. A spring's potential is $V \propto kx^2$, and the oscillation period is $T \propto \frac{1}{\sqrt{k}}$. m_a is analogous to k , so the axion oscillation period is $T \propto \frac{1}{m_a}$.

damping mechanism. Axions also have a feeble interaction with normal matter, especially when they are light. For axions interacting with photons (Figure 1.3), the coupling constant is $g_{a\gamma\gamma} = \left(0.203(3)\frac{E}{N} - 0.39(1)\right)\frac{m_a}{\text{GeV}}$, where $\frac{E}{N} \sim \mathcal{O}(1)$. That means if $m_a \sim 1\mu\text{eV}$, then $g_{a\gamma\gamma} \sim 10^{-15}$.

To recap, the axion field is hypothesized to solve the strong CP problem by enabling the CP violating term to naturally relax to zero. It turns out that the axion is an excellent dark matter candidate because it would have been produced copiously in the early universe, would be stable on cosmological timescales, and would have feeble interactions with normal matter.

1.2.3 Axion Cosmology

Sources of Cosmological Axions

There are three methods of generating cosmological axions: thermal production, the misalignment mechanism, and radiation from topological defects.

Thermal (or relativistic) axions would have been produced through the scattering of the cosmic plasma. One dominant process would be $\pi + \pi \leftrightarrow \pi + a$. The relic axion number density and temperature would be comparable to that of one species of neutrinos and would contribute to hot dark matter in the universe. However, the dark matter primarily responsible for driving galactic structure is known to be cold. Thus axions must be produced athermally through the misalignment mechanism and radiation of topological defects.

The misalignment mechanism describes how axions are produced from the oscillation of the axion field about its potential minimum. The following is a history of how these oscillations started. The early universe was hot ($\sim 10^{32}$ K). When the universe cools below the PQ energy scale, the $U_{PQ}(1)$ symmetry is spontaneously broken. The effective potential takes the shape of a Mexican hat and the axion field takes on random initial values. As the universe cools down further and approaches the QCD energy scale, the $U_{PQ}(1)$ symmetry is

explicitly broken. The axion field evolves according to the equation

$$\ddot{\theta} + 3H(T)\dot{\theta} + m_a^2(T)\theta = 0 \quad (1.1)$$

where $\theta \equiv a/f_a$ is the defined misalignment angle, $H(T)$ is the Hubble constant as a function of temperature³, and $m_a(T)$ is the axion mass as a function of temperature.

The temperature dependence of the axion mass is described by

$$m_a(T) = \begin{cases} C m_a(T=0) (\Lambda_{QCD}/T)^4 & \text{for } T \gtrsim \Lambda_{QCD} \\ m_a(T=0) & \text{for } T \lesssim \Lambda_{QCD} \end{cases} \quad (1.2)$$

The mass dependence on temperature makes sense because the effective potential is also temperature dependent. The mass of a field can be thought of as describing the steepness of the paraboloid near the potential minimum. The paraboloid becomes steeper as the universe approaches the QCD scale and the effective potential begins to tilt.

The axion field evolves like a harmonic oscillator with damping due to the Hubble expansion⁴. Thus the axion oscillates coherently around the potential minimum. This coherent oscillation forms a zero momentum condensate that forms cold dark matter. The mass-energy density of axions today related to the misalignment production is $\Omega_{a,mis} \sim 0.15 \left(\frac{f_a}{10^{12} \text{ GeV}}\right)^{7/6} \theta_i^2$, where θ_i is the initial misalignment angle.

Axions can also be produced through topological defects in the axion field, such as axion strings and domain walls. A cosmic string is a topological defect formed during a symmetry breaking phase transition when the topology of a vacuum manifold is not simply connected. In particular, an axion string is a linear structure where the misalignment angle θ changes by 2π when physically traversing the string. A domain wall is a surface where $\theta = \pi$, and each axion string has one domain wall ending on it. This is all terribly abstract. To give a more

³equivalently as a function of time since universe's temperature evolves with time.

⁴More technically and possibly more correct, the field evolved according to the Klein-Gordon equation on a Friedmann-Robertson-Walker metric.

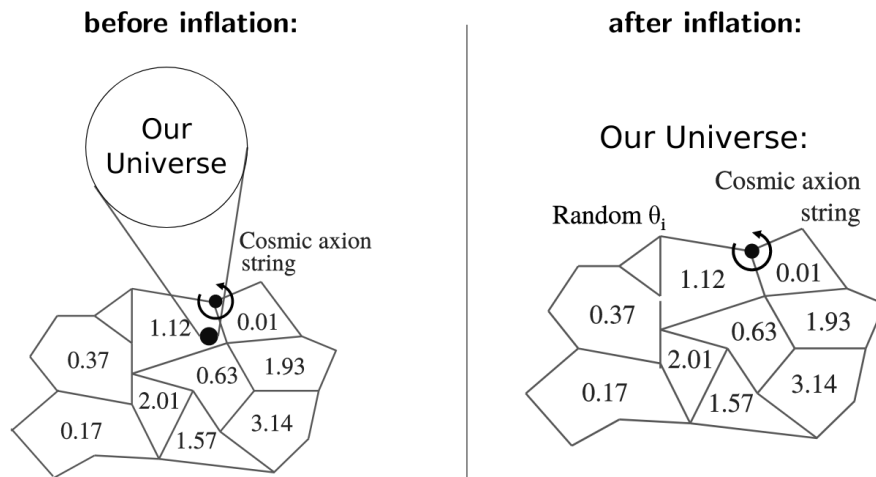


Figure 1.4: Left: PQ symmetry is broken before inflation and never restored. At inflation, each “patch” of initial axion value is blown up. Our observable universe sits in one of those patches. Thus our observable universe has only experienced one value for the axion (with some quantum perturbations). Right: PQ symmetry was broken after inflation. The axion field starts with random values throughout the observable universe. Our universe will contain topological defects like axion strings[13].

intuitive explanation of these defects, axion strings are where the axion field has wrapped around itself in a way that has caused a discontinuity in the axion field. Quantum fields do not like discontinuities, so the axion field will want to time-evolve in a way that removes these discontinuities. But removing these discontinuities around the strings would require changing the axion field values throughout the entire universe. That would require an infinite amount of energy and therefore would be untenable. So axion strings are metastable and will disappear once they radiate away axions.

PQ Symmetry Breaking and Inflation

There are two scenarios to consider for cosmic axion production: PQ symmetry is broken before inflation and is never restored, and PQ symmetry is broken after inflation (Figure 1.4).

If PQ symmetry was broken before symmetry breaking, then each random patch of theta is inflated. Our observable universe would be inside a single value of theta. The

axion strings would inflate away and so the only sources of axions are either through thermal production or through the misalignment mechanism. The energy density would be $\Omega_{a,mis} \sim 0.15 \left(\frac{f_a}{10^{12} \text{GeV}}\right)^{7/6} \theta_i^2$. If one posits that the axions make up all of dark matter, then $\Omega_{dm} \sim 0.15 \left(\frac{f_a}{10^{12} \text{GeV}}\right)^{7/6} \theta_i^2$. Ω_{dm} is known, so the axion mass would be fixed by the initial misalignment angle. There will however be a variance in the initial angle due to quantum fluctuations, $\langle \sigma_\theta^2 \rangle = \frac{H_I}{2\pi f_a}$, where H_I is the Hubble constant at the time of inflation. Quantum fluctuations of the axion field during inflation lead to isocurvature density fluctuations which get imprinted as temperature fluctuations in the cosmic microwave background⁵. The resulting isocurvature power spectrum is $\Delta_a(k) = \frac{H_I^2}{\pi^2} \theta_i^2 f_a^2$.

If PQ symmetry was broken after inflation, then at the time of symmetry breaking, the observable universe would be filled with different random values of theta. The axion energy density will be averaged over the Hubble volume, such that $\Omega_{a,mis} \sim 0.15 \left(\frac{f_a}{10^{12} \text{GeV}}\right)^{7/6} \langle \theta_i^2 \rangle \sim 2.07 \left(\frac{f_a}{10^{12} \text{GeV}}\right)^{7/6}$. If one posits that axions make up all of dark matter, then this suggests that $m_a > \mu\text{eV}$. The initial state of the axion field doesn't affect present-day cosmology. However, axion strings and domain walls will contribute to the axion matter in the universe.

1.2.4 Axion Mass Predictions

Direct detection requires looking for axions with a specific range of mass and coupling strengths. A wide range of axion masses would solve the strong CP problem. Furthermore, axions are difficult to detect because of their feeble interaction with normal matter. The most sensitive axion searches use resonators that enhance the signal power at the cost of only being able to search a very narrow range of parameter space. Thus it behooves the physics community to use whatever clues nature has provided to reduce the search parameter space.

Astrophysical sources, such as white dwarves, neutron stars, red giants, and supernovas,

⁵The CMB is remarkably uniform but has tiny fluctuations in its power spectrum. Adiabatic fluctuations are density variations in all forms of matter and energy, which have equal over/under number densities. Isocurvature fluctuations are density variations in which the number density variations for one component do not correspond to the number density variations in other components.

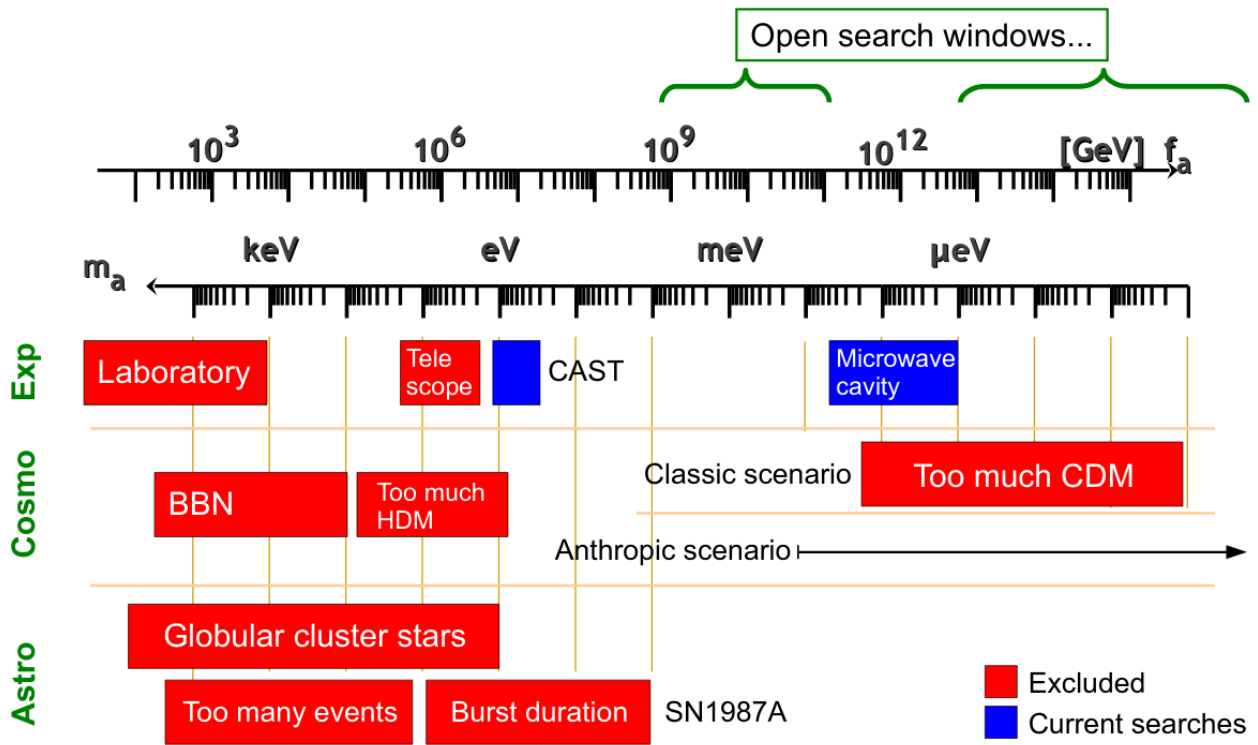


Figure 1.5: A figure from 2010 that shows how astrophysical sources place limits on axion mass. [14]

can be used to limit the axion mass parameter space. The lower limit on the axion mass is $m_a > 10^{-22}$ eV [15]. Otherwise, the de Broglie wavelength would be bigger than galaxies, which would have a noticeable effect on galactic structure. Astrophysical objects would also produce axions thermally, and this production is a source of energy loss. For axions to efficiently transport energy, their mean free path must be comparable to that of the astrophysical object. If the axion coupling strength is too large, then the axions have a hard time escaping these bodies. If the axion coupling strength is too small, then astrophysical bodies cannot produce them efficiently. If astrophysical objects cool at the rate predicted by standard model processes, one can place an upper limit on the axion coupling strength. The nuclear physics is tedious for the purposes of this thesis, so I will just state limits. Stars in the horizontal branch limit the axion-photon coupling strength to $|g_{a\gamma\gamma}| < 6.6 \times 10^{-11}$. Red giants on the horizontal branch limit the axion-electron coupling to $|g_{aee}| < 4.3 \times 10^{-13}$. The best limit is from the detection of neutrinos from supernova SN1987a. Sufficiently light axions would be copiously produced by this supernova and would steal energy away from neutrinos, reducing the duration of the neutrino burst. The burst duration limits the strength of axions coupling to normal matter to $g_{ann}^2 + 0.29g_{app}^2 + 0.27g_{ann}g_{app} < 3.25 \times 10^{-18}$. This translates to a conservative mass limit of $m_a < 16$ meV. Figure 1.5 shows how astrophysics has limited the possible axion mass ranges. If PQ symmetry breaking happened before inflation, then 10^{-22} eV $< m_a < 10^{-2}$ eV. However, if we posit that PQ symmetry breaking happened after inflation, the lower limit becomes far more stringent. However, there are topological complications such as axion strings and domain walls. Lattice QCD is used to predict the contributions of axion string radiation. Figure 1.6 shows the range of published mass predictions.

1.3 ALPs and Dark Photons

Axion-like particles (ALPs) have the same properties as axions, except their photon coupling $g_{a\gamma\gamma}$ is now independent of the ALP mass. ALPs do not solve the strong CP problem, but they can make up all of the dark matter. ALPs are predicted to arise generically from

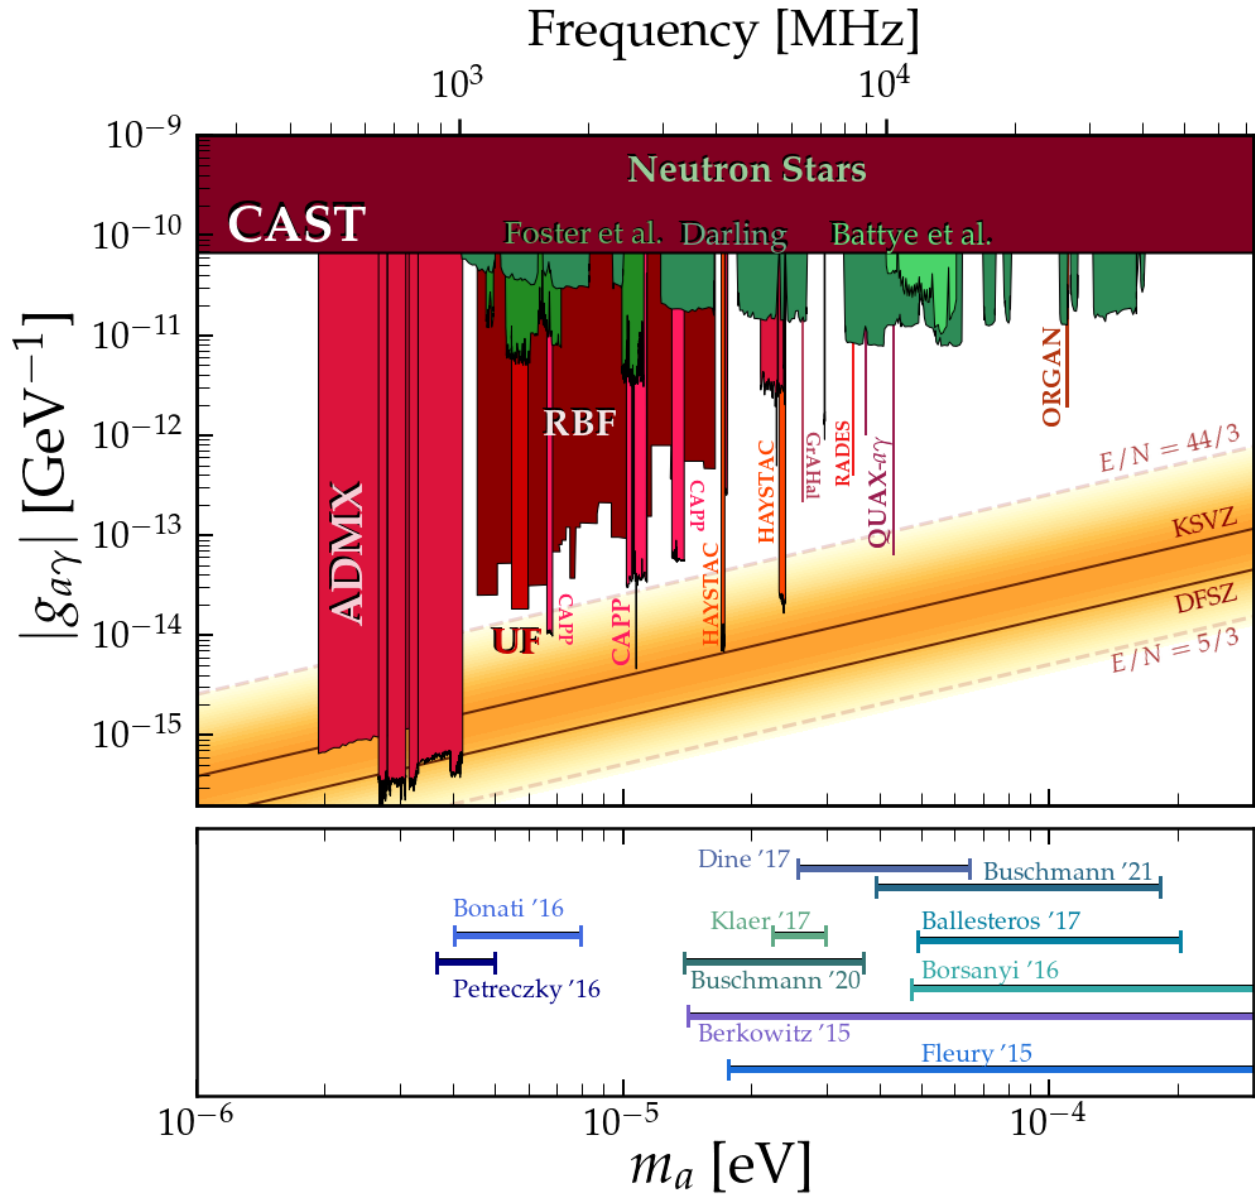


Figure 1.6: The excluded parameter space for axions, zoomed into microwave cavity haloscopes. The bottom plot shows different theoretical predictions for post-inflationary axions. A lot of the tensions have to do with how much do topological defects contribute to the axion population. Figure courtesy of Ciaran O’Hare[16].

low-energy effective field theories emerging from string theory[1].

The dark photon⁶ (DP) story is similar to the axion story. The dark photon is a vector boson associated with an added Abelian U(1) symmetry to the Standard Model[17, 18, 19]. The dark photon is analogous to the Standard Model (SM) photon in that the SM photon is also a vector boson associated with an Abelian U(1) gauge symmetry. The dark photon interacts with the SM photon through kinetic mixing [20, 21]

$$\mathcal{L} = -\frac{1}{4}F_1^{\mu\nu}F_{1\mu\nu} - \frac{1}{4}F_2^{\mu\nu}F_{2\mu\nu} + \frac{1}{2}\chi F_1^{\mu\nu}F_{2\mu\nu} + \frac{1}{2}m_{A'}^2 A'^2 \quad (1.3)$$

where $F_1^{\mu\nu}$ is the electromagnetic field tensor, $F_2^{\mu\nu}$ is the dark photon field tensor, χ is the kinetic mixing angle, $m_{A'}$ is the DP mass, and A' is the DP gauge field. The consequence is that the dark photon and SM photon can oscillate into each other (reminiscent of neutrino oscillations). Unlike axions, dark photons and SM photons do not need the presence of a magnetic field to mix. Also, unlike axions, the dark photon kinetic mixing angle is independent of the dark photon mass.

If the dark photon has a small kinetic mixing, i.e., small coupling to the SM particles, then it is stable on cosmological timescales and makes for a compelling dark matter candidate. The rate that dark photons decay into three photons is

$$\Gamma_{A' \rightarrow 3\gamma} = (4.70 \times 10^{-8}) \times \alpha^3 \alpha' \frac{m_{A'}^9}{m_e^9} \quad (1.4)$$

where α is the fine structure constant, m_e is the mass of the electron, and $\alpha' \equiv \frac{(e\chi)^2}{4\pi}$ is the dark photon counterpart to the fine structure constant[22]. The dark photon lifetime is about the same as the age of the universe if $m_{A'}(\alpha')^{1/9} < 1$ keV. For dark photons with $m_{A'} \approx 10^{-4}$ eV and $\chi < 10^{-12}$, this condition is easily met.

Several mechanisms would produce cosmic dark photons. One of these mechanisms is a misalignment mechanism that is similar to that for axions. However, the dark photon,

⁶Also known as a hidden photon or paraphoton.

due to its vector nature, will have its energy density redshifted as the universe expands⁷ and wouldn't contribute to the energy density of matter Ω_m . For the misalignment mechanism to generate the correct relic abundance, a nonminimal coupling to gravity needs to be invoked. This leads to instabilities in the longitudinal DP mode, which would lead to DP with a fixed polarization within a cosmological horizon[19]. Dark photons may also be produced by quantum fluctuations during inflation[23]. These quantum fluctuations seed excitations in the dark photon field that result in the cold dark matter observed today in the form of coherent oscillation of this field. The predicted mass from this mechanism is $m_{A'} \approx 10^{-5} \text{ eV} \times \left(\frac{10^{14} \text{ GeV}}{H_I} \right)^4$, where H_I is the Hubble constant during inflation. Measurements of the CMB tensor to scalar ratio constrain $H_I < 10^{14} \text{ GeV}$ [24]. This mechanism is the most minimal model for producing dark photon dark matter and makes the search for $m_{A'} > 10^{-5} \text{ eV}$ well-motivated. Dark photons can also be produced from topological defects like cosmic strings. Dark photons may also be produced thermally through processes like $e^- + \gamma \rightarrow e^- + A'$ [25]. However, due to the small coupling to SM particles, this process would be inefficient and not produce the observed abundance of dark matter.

⁷In the same way SM photons redshift during cosmic expansion.

Chapter 2

DIELECTRIC HALOSCOPES FOR DETECTING AXIONS AND DARK PHOTONS

This chapter describes how ADMX uses microwave cavities to search for $\mathcal{O}(1 \mu\text{eV})$ axions, how the same technique gets harder to apply at higher axion masses, and how dielectric haloscopes can be used for higher mass searches.

This section focuses heavily on axions because that's the context for which Orpheus was developed. Most of the time, the following discussion applies to dark photons. The main difference between detecting either is in the polarization of the detected photons. I will talk about the detection of axions and remark about dark photons when appropriate.

This chapter assumes the reader understands resonant cavities.

2.1 Haloscopes to search for wavelike dark matter

An axion can be detected by detecting its feeble coupling to the SM photon. Axions mix with photons in a magnetic field (similar to how neutrinos of different flavors mix in free space). If this mixing happens inside a resonant cavity, and if that resonant frequency matches the photon frequency, the axion signal power is resonantly enhanced[26].

The Lagrangian that describes an axion coupling to two photons is

$$\mathcal{L}_{a\gamma\gamma} = \frac{1}{4}g_{a\gamma\gamma}aF_{\mu\nu}\tilde{F}_{\mu\nu} = -g_{a\gamma\gamma}a\mathbf{E}\cdot\mathbf{B} \quad (2.1)$$

From the Lagrangian, the derived axion modified Maxwell's equations are [27, 28]

$$\nabla \cdot \mathbf{E} = \rho - g_{a\gamma\gamma} \mathbf{B} \cdot \nabla a, \quad (2.2)$$

$$\nabla \cdot \mathbf{B} = 0, \quad (2.3)$$

$$\nabla \times \mathbf{E} = -\dot{\mathbf{B}}, \quad (2.4)$$

$$\nabla \times \mathbf{B} = \dot{\mathbf{E}} + \mathbf{J} - g_{a\gamma\gamma} (\mathbf{E} \times \nabla a - \dot{a} \mathbf{B}), \quad (2.5)$$

$$\ddot{a} - \nabla^2 a + m_a^2 a = g_{a\gamma\gamma} \mathbf{E} \cdot \mathbf{B}. \quad (2.6)$$

This can be simplified by noting that $\nabla a \approx 0$. This approximation works because the de Broglie wavelength of the axion is much larger than the size of the detector. For an axion with $m_a = 80 \mu\text{eV}$ and $\langle v^2 \rangle^{1/2} = 270 \text{ km/s}$, the de Broglie wavelength is $\lambda_{\text{dB}} \approx 17 \text{ m}$. Orpheus is about 15 cm, so the axion field is approximately constant throughout the detector volume.

Also note that $g_{a\gamma\gamma} \dot{a} \mathbf{B}$ is on equal footing with \mathbf{J} . So the axion field can act like a current source that excites electromagnetic fields. Treatments of currents exciting cavity modes can be found in Jackson Chapter 8[29] and Pozar 4th ed. Chapter 6[30].

This coupling is feeble, but the resulting signal can be made detectable if the interacting magnetic field is very strong and if the signal is resonantly enhanced by a high-Q cavity¹. This signal can be detected with low noise electronics as a spectrally narrow power excess over the noise background, as shown in Figure 2.1.

I will outline the steps to deriving the equation for the axion signal power inside a cavity. First, derive the wave equation from Maxwell's equations. Assume harmonic time dependence and write \mathbf{E} as the sum of cavity eigenmodes, i.e., rewrite the wave equation in k-space such that $\mathbf{E} = \sum_i \mathbf{E}_i \exp(-i(\omega t - \mathbf{k}_i \cdot \mathbf{x}))$. The k-space wave equation should have a damping term that's related to the cavity Q. Solve for the electric field amplitude of the relevant cavity mode and find the time-averaged energy stored $U = \int dV \frac{1}{2} \epsilon |E_0|^2$. For a signal

¹The cavity allows for the axion signal to build up (think of a photon of having about Q bounces before it dissipates).

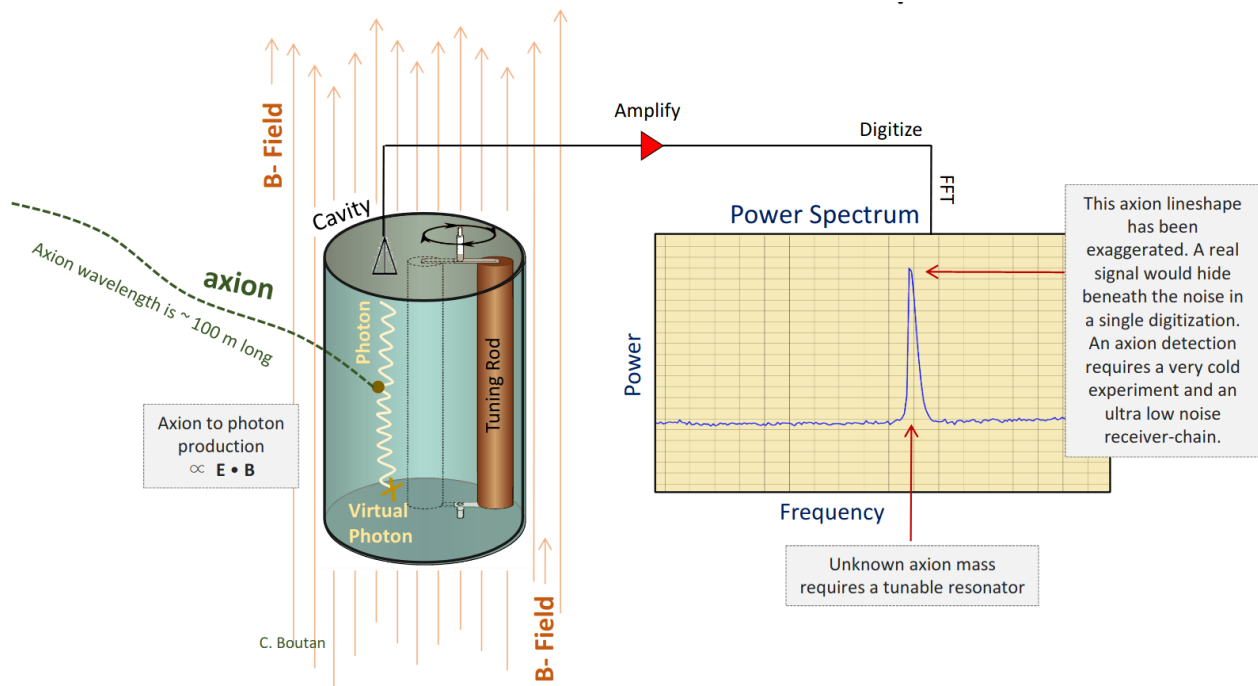


Figure 2.1: Haloscope search for axions. An axion, when it sees a strong magnetic field, can convert to a photon. If the photon frequency matches the resonant frequency of the cavity, the axion signal power is resonantly enhanced. A haloscope search for dark photons works the same way, except that a dark photon can convert to an SM photon without the presence of a magnetic field. Figure courtesy of Christian Boutan.

to be detected, the cavity has to be coupled to an external receiver. From the energy stored in the cavity, the unloaded Q of the cavity, and the coupling of the receiver to the cavity, one can derive the power dissipated in the receiver, i.e., the signal power. Let

$$\frac{1}{Q_L} = \frac{1}{Q_0} + \frac{1}{Q_e} \quad (2.7)$$

where Q_L is the loaded Q of the cavity, Q_0 is the unloaded Q , and Q_e characterizes how much energy gets dissipated in the external load. The cavity coupling is $\beta \equiv \frac{Q_0}{Q_e}$. The power drawn into the external load is $P_a = \omega \frac{U}{Q_e}$. Putting all this together, the resulting signal power is (in natural units)

$$P_a = \frac{g_{a\gamma\gamma}^2 m_a \rho_a B_0^2 V_{eff} Q_L \beta}{m_a (\beta + 1)} L(f, f_0, Q_L) \quad (2.8)$$

$$L(f, f_0, Q_L) = \frac{1}{1 + 4\Delta^2} \quad \Delta \equiv Q_L \frac{f - f_0}{f_0} \quad (2.9)$$

$$V_{eff} = \frac{|\int dV \mathbf{B}_0 \cdot \mathbf{E}_a|^2}{B_0^2 \int dV \epsilon_r |E_a|^2} \quad (2.10)$$

The derivation was merely verbally outlined, and the math is left as an exercise to the reader. The reader is encouraged to read other sources like [31] for a more thorough treatment.

Axion searches look for a power excess over a thermal noise floor P_n . P_n is the average of thousands of power spectrum measurements and consequently follows the Central Limit Theorem, so $\sigma_{P_n} = \frac{P_n}{\sqrt{N}}$, where N is the number of averaged spectra. The number of spectra is $N = b\Delta t$, where b is the frequency bin width and Δt is the total digitization time. The SNR of a haloscope signal is

$$\text{SNR} = \frac{P_s}{\sigma_{P_n}} = \frac{P_s}{P_n} \sqrt{N} = \frac{P_s}{P_n} \sqrt{b\Delta t} \quad (2.11)$$

Keeping in mind that the cavity bandwidth is $\Delta f = \frac{f_0}{Q_L}$, the instantaneous scan rate is

$$\frac{df}{dt} = \frac{\Delta f}{\Delta t} = \frac{f_0 Q_L}{b} \left(\frac{g_{a\gamma\gamma}^2 \rho_a B_0^2 V_{eff} \beta}{\text{SNR} m_a T_n (\beta + 1)} \right)^2 \quad (2.12)$$

ADMX has successfully applied this technique to search for DFSZ axions with masses between $2.7 \mu\text{eV}$ (650 MHz) and $3.3 \mu\text{eV}$ (800 MHz) and continues to push the search to higher masses [32, 33, 34]. For Run 1B, ADMX employed a 136 L cylindrical cavity. The cavity was immersed in a 7.6 T magnetic field and operated at the TM_{010} mode such that the axion's electric field maximally aligned with the external magnetic field (the axion's electric field looks like that of a parallel plate capacitor). For Run 1B, $Q_L = 30000$ and $\frac{df}{dt} = 543 \text{ MHz/yr}$.

2.2 Haloscopes for Dark Photons

The dark photon story is similar but with some key differences.

The dark photon interacts with the SM photon through kinetic mixing. The Lagrangian for this interaction is shown in Equation 1.3. The consequence is that the dark photon and SM photon will oscillate into each other. One can derive the electromagnetic field produced by the dark photon to be [19]

$$|\mathbf{E}_0| = \left| \frac{\chi m_X}{\epsilon} \mathbf{X}_0 \right| \quad (2.13)$$

where \mathbf{X} is the dark photon field. One should note that the SM photon polarization is determined by the dark photon polarization rather than some external magnetic field. This polarization can be misaligned with the probing mode of the cavity.

The resulting dark photon signal power in a cavity is [18] (in natural units)

$$P_S = \eta \chi^2 m_{A'} \rho_{A'} V_{eff} Q_L \frac{\beta}{\beta + 1} L(f, f_0, Q_L) \quad (2.14)$$

$$V_{eff} = \frac{(\int dV \vec{E}(\vec{x}) \cdot \vec{X}(\vec{x}))^2}{\int dV \epsilon_r |\vec{E}(\vec{x})|^2 |\vec{X}(\vec{x})|^2} \quad (2.15)$$

Let θ be the angle between the cavity field and the dark photon field, such that $\vec{X} \cdot \vec{E} = |\vec{X}| |\vec{E}| \cos \theta$. The V_{eff} can be rewritten as

$$V_{eff} = \frac{(\int dV |\vec{E}(\vec{x})| |\vec{X}(\vec{x})|)^2}{\int dV \epsilon_r |\vec{E}(\vec{x})|^2 |\vec{X}(\vec{x})|^2} \langle \cos^2 \theta \rangle_T \quad (2.16)$$

where $\langle \cos^2 \theta \rangle_T$ is the time-averaged $\cos^2 \theta$ value. The actual value of $\langle \cos^2 \theta \rangle_T$ depends on the dark photon cosmology and the detector design and orientation. Often I will write $V_{eff} = V_{eff,max} \langle \cos^2 \theta \rangle_T$ because it allows me to rescale my dark photon limits to any cosmology.

2.3 Scaling the Haloscope Concept to Higher Frequencies

ADMX is currently using this haloscope method to look for axions around a few μeV with great success. Unfortunately, this haloscope design becomes increasingly difficult to implement at higher frequencies. Increasing mass corresponds to higher frequency photons. Operating at the TM_{010} mode would require smaller-diameter cavities, and a smaller cavity volume reduces the signal strength. The volume would scale by $V_{eff} \propto f^{-2}$, or $V_{eff} \propto f^{-3}$ if one wanted to keep the same aspect ratio. Furthermore, the decreased volume-to-surface ratio decreases Q , further decreasing the signal. Q also reduces as a function of frequency because of the anomalous skin effect, so $Q_u \propto f^{-2/3}$. The smaller Q_u also makes it more difficult for the receiver to couple critically. The quantum noise limit also increases linearly with f . So putting all these effects together, for a single closed cavity operating at the lowest order mode with a quantum noise limited amplifier, the axion signal power scales optimisti-

cally as $P_s \propto f^{-2.66}$ and the scan rate scales optimistically as $\frac{df}{dt} \propto f^{-7.66}$. This unfavorable frequency scaling motivates the design of more sophisticated resonators².

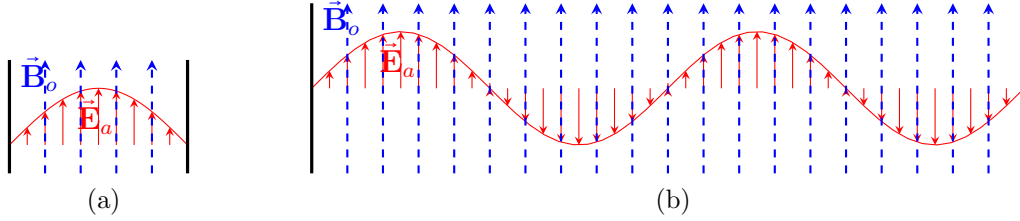


Figure 2.2: Traditional haloscope signal power scale poorly with frequency. (a) Small cavity operating at the fundamental mode, e.g., TM_{001} mode. The volume of such a cavity scales as $\sim f^{-3}$ for the same aspect ratio. (b) Operating a large cavity at higher frequencies requires operating at higher order modes. The effective volume $V_{eff} \propto \int dV \vec{E}_a \cdot \vec{B}_o \approx 0$.

One can think about keeping a large volume and operating at a mode higher than the TM_{010} . But then portions of \vec{E}_a are anti-aligned with \vec{B}_o , and $|\int dV \vec{B}_o \cdot \vec{E}_a| \approx 0$. The effective volume approaches zero even though the physical volume is large. Thus, there is little benefit to operating an empty cylindrical cavity at a higher-order mode.

One can also combine many cylindrical cavities. This is the plan for future ADMX runs. However, once the frequency approaches 15 GHz, the wavelength is about 1 cm^3 . To have $V_{eff} = 1 \text{ L}$, one would need to coherently power combine about 2000 cylindrical cavities. It would be difficult to instrument a 1 cm^3 cavity, and it would be unfeasibly complex to instrument 2000 cavities in a coordinated way.

2.4 Multimode Dielectric Haloscopes and Orpheus Conceptual Design

High-order modes can couple to the axion when dielectrics are placed inside of the resonator. Dielectrics suppress electric fields. If dielectrics can be placed where the electric field is anti-aligned with the magnetic field, then the overlap between the axion's electric field and

²Someone at a workshop called axions above 10 GHz axions from hell. Orpheus came out of hell but lost someone he loved in the process.

the external magnetic field is greater than zero (see Figure 2.3). Thus the effective volume can become arbitrarily large, and the axion signal power is greater than what it would have been for a cylindrical cavity operating at the TM_{010} mode. Overall, dielectric resonators can be operated at higher-order modes while maintaining coupling to the axion and can be made arbitrarily large, making them suited for higher-frequency searches.

Orpheus will implement this dielectric haloscope concept to search for axions around $70 \mu\text{eV}$. Orpheus³ is a dielectrically loaded Fabry-Perot open cavity placed inside of a dipole magnet. Dielectrics are placed every fourth⁴ of a half-wavelength to suppress the electric field where it is anti-aligned with the dipole field. Orpheus is designed to search for axions around 16 GHz with over 1 GHz of tuning range. The cavity tunes by changing its length, and the dielectrics are automatically adjusted to maintain even spacing throughout the cavity.

There are several benefits to the open resonator design. Less metallic walls lead to less ohmic losses and a higher Q. Less metallic walls also mean fewer resonating modes. This leads to a sparse spectrum and fewer mode crossings, making it easier to maintain the mode of interest.

However, this experiment has many challenges. First, the optics must be designed to maintain good axion coupling for over 1 GHz of tuning range. This includes choosing the right radius of curvature for the Fabry-Perot mirrors and appropriate dielectric thicknesses. Another challenge is that a dipole magnet is required to have B_0 aligned with the TEM electric fields. Dipoles are about ten times more expensive than solenoids. In addition, the mechanical design for such a cavity is complicated because there are many moving parts that

³Orpheus was originally designed to have a spatially alternating magnetic field rather than a periodic dielectric structure[35]. However, this alternating magnetic field design is difficult to scale to many Tesla.

⁴I could have also designed Orpheus so that the dielectrics were at every other half-wavelength. Intuitively, this would have resulted in a much greater V_{eff} . But that didn't happen for various practical reasons that will be addressed in the next few sections. To summarize, it's because it would have been harder to design a cryo-compatible mechanical structure with the dielectric plates that close together, and I was having a hard time simulating the mode of interest⁵ for the entire tuning range. I managed to make Orpheus work when the dielectrics were spaced every fourth of a half-wavelength and kept forging ahead with what worked. Orpheus could have worked if I placed dielectrics at every other half-wavelength, but it would have been harder to implement, and the tuning range would have likely been more limited.

have to work in a cryogenic environment. Finally, both diffraction and dielectric losses will decrease the resonator Q.

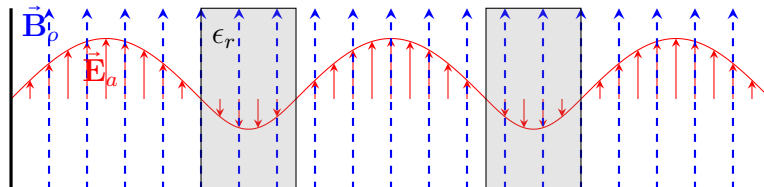


Figure 2.3: A multimode dielectrically-loaded cavity. Dielectrics are placed where the electric field is anti-aligned with the external magnetic field. Dielectrics suppress electric fields, so $V_{eff} \propto \int dV \vec{E} \cdot \vec{B} > 0$.

2.5 Other related experiments

Dielectric haloscopes are a growing field. Without going into detail, there is MADMAX that uses a very similar periodic dielectric structure [36]. Baryakhtar, et. al., has developed a similar concept for the infrared and optical range [37]. Quiskamp, et. al., has also developed a cylindrical cavity with a periodic dielectric structure [38]. This design has the advantage of working in a solenoid rather than a dipole magnet. RADES is not a dielectric haloscope but is a multimode cavity designed to work in a dipole magnet [39].

Chapter 3

ORPHEUS RF DESIGN AND CHARACTERIZATION

This chapter describes the RF characterization of Fabry-Perot cavity with dielectrics evenly-spaced at every fourth half-wavelength¹. Measurements are taken to extract the resonant frequency, quality factor, and cavity coupling coefficient of the mode of interest. These measurements are compared to Finite Element Analysis simulations. Simulations are also used to determine V_{eff} and tolerance for misplaced dielectrics. The chapter ends with describing different cavity coupling schemes and why the current coupling configuration was chosen.

3.1 The Empty Fabry-Perot Cavity

Intuition for the dielectrically-loaded Fabry-Perot cavity builds on the well-established intuition for the empty Fabry-Perot cavity. The empty cavity modes have analytical solutions. But an analytical solution to the dielectrically-loaded case is infeasible², and simulations are required to understand the electric field. One way to corroborate the simulations is to study the empty cavity in detail. If the mode frequencies from the measurement, simulation, and analytical calculation match within experimental/simulation uncertainty, then the simulation and measurement techniques are validated and can be trusted for the dielectrically-loaded case.

¹Section 4.1 describes how the cavity was built.

²Intuition from geometric optics suggest that the solutions from the empty cavity if the optical length is modified from the dielectrics. That's a good way to guess the mode frequencies. But it turns out that the resonant frequencies and electric field deviate quite strongly from what one would suppose from geometric optics. Therefore, simulations are needed to get the electric field.

3.1.1 Theoretical Basis

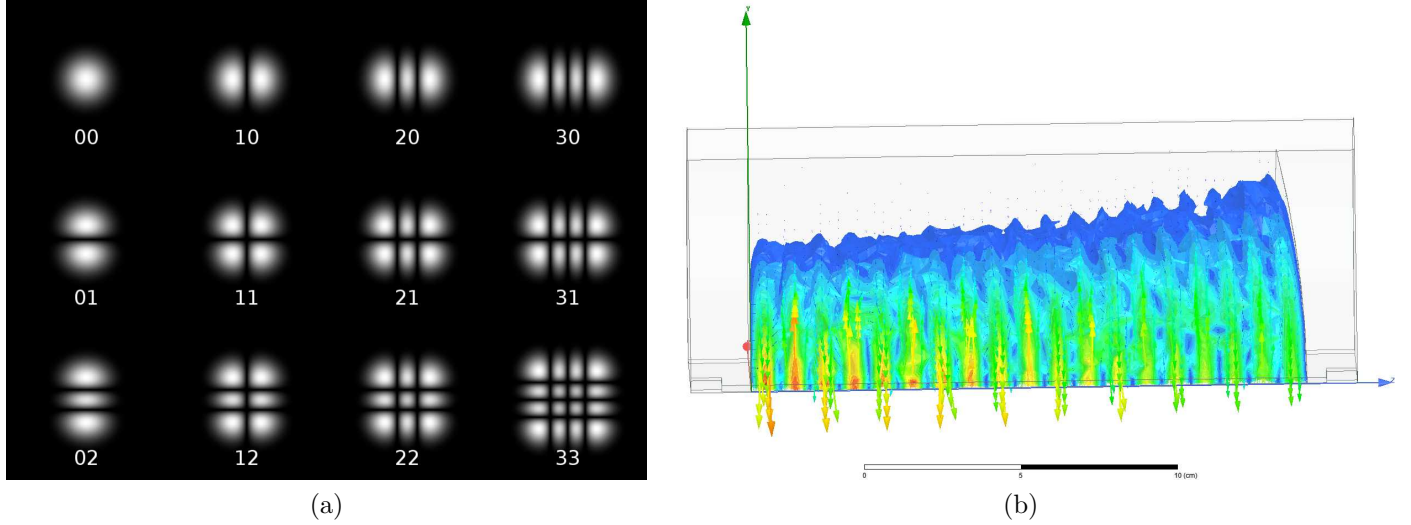


Figure 3.1: (a) The Hermite-Gaussian modes of a Fabry-Perot cavity. Figure from Kogelnik[40]. (b) FEA simulation of the TEM₀₀₋₁₈mode in an empty Fabry-Perot cavity.

The Fabry-Perot cavity modes are described by TEM Gaussian modes, as shown in Figure 3.1. The details and derivations are in references [40, 41]. I will outline the solutions and main ideas relevant to Orpheus.

The traveling wave Gaussian beam electromagnetic wave is described by

$$u_{mn} = \frac{w_0}{w} \left[H_m \left(\sqrt{2} \frac{x}{w} \right) H_n \left(\sqrt{2} \frac{y}{w} \right) \exp \left(\frac{-\rho^2}{w^2} \right) \right] \exp \left(-i(kz - \varphi) - \frac{ik\rho^2}{2r} \right) \quad (3.1)$$

where k is the wavenumber, z is the position coordinate along the beam axis, ρ is the transverse distance from the beam axis, $w(z)$ characterizes the width of the beam, w_0 is the beam waist (the smallest width of the beam), r is the radius of curvature of the wavefront, H_m and H_n are the Hermite polynomials that characterize modes in the transverse direction. Finally, φ is the Gouy phase. It's a phase change additional to the free space propagation that results from the wavefront radius of curvature making the phase velocity $v_p > c$. The

Gouy phase is derived to be

$$\varphi = (1 + m + n) \tan^{-1}(z/z_0) \quad (3.2)$$

Note that for the fundamental Gaussian beam, $H_0 = 1$. The beam width and wavefront radius of curvature are hyperbolic functions that depend on z .

$$w^2 = w_0^2[1 + (z^2/z_0^2)] \quad (3.3)$$

$$r = z[1 + (z_0^2/z^2)] \quad (3.4)$$

$$z_0 = kw_0^2/2 \quad (3.5)$$

Inside the Fabry-Perot cavity, the resonant condition for the Gaussian beam is

$$\frac{f}{f_0} = (q + 1) + \frac{1}{\pi}(m + n + 1) \arccos \sqrt{(1 - L/R_1)(1 - L/R_2)} \quad (3.6)$$

$$f_0 = \frac{c}{2D} \quad (3.7)$$

where L is the length of the cavity, R_1 and R_2 are the radii of curvature of the mirrors, and q is the mode number along the axis of the cavity. If $q = 0$, then there is one half-wavelength along the axis and no nodes. If $q = 18$, then 18 half-wavelengths along the axis and 18 nodes. For Orpheus, one of the mirrors is flat, so $R_2 = \infty$ and $\frac{f}{f_0} = (q + 1) + \frac{1}{\pi}(m + n + 1) \arccos \sqrt{1 - L/R_1}$.

The stability condition for the Fabry-Perot cavity is

$$0 \leq (1 - L/R_1)(1 - L/R_2) \leq 1. \quad (3.8)$$

The condition is derived in [40] and can be intuited from ray tracing multiple bounces between

the mirrors.

A stable beam with a good Q can be achieved if the cavity operates in the near-confocal configuration. For a symmetric cavity where $R_1 = R_2 = R_0$, the near confocal geometry describes a configuration where the mirrors are placed at each others focus, i.e. $L = R_0/2$.

For good Q in the near-confocal regime, the Fresnel number should be[41]

$$N = \frac{a_1 a_2}{\lambda L} > 1 \quad (3.9)$$

where a_1, a_2 are the diameters of the mirrors.

The quality factor of the cavity is[42]

$$Q = 2\pi f_n \frac{2L}{c} \frac{1}{\beta} \quad (3.10)$$

where f_n is the resonant frequency and β is the fraction of power lost from the cavity in each round trip³. β is comprised of the ohmic losses in the mirror, diffraction, and dielectric losses if dielectrics are present. The diffraction loss caused by the spillover of the mode around the edges of the mirrors is $\exp\left(-\frac{a^2}{2w_m^2}\right)$, where w_m is the beam size at the mirror.

The cavity modes are excited through a rectangular waveguide and aperture. The magnetic field \mathbf{H} in the waveguide will fringe through the aperture and appear like a magnetic

³Not to be confused with the cavity coupling coefficient. The term only appears in this section, so I will keep it as is to stay consistent with the literature.

dipole moment[30, 42, 43]⁴. The effective magnetic moment of the hole is [42]

$$\mathbf{m}_h = \alpha_M \mathbf{H}_h \quad (3.11)$$

$$\alpha_M = 4/3 f_t r_h^3$$

$$f_t = \exp\left(\frac{-2\pi A_m t}{\lambda_c} \sqrt{1 - (\lambda_c/\lambda)^2}\right)$$

$$\lambda_c = 3.412 r_h$$

$$A_m t = 1.0064t + 0.0818 r_h \quad \text{for } t > 0.2 r_h$$

where \mathbf{H}_h is the magnetic field at the hole, t is the thickness of the hole, and f_t is a factor that accounts for the thickness of the hole.

For achieve critical coupling, the radius of the hole r_h should be [42]

$$r_h = \left(\frac{9\pi w_m^2 L a^2}{256 Q_0 k_g f_t^2}\right)^{1/6} \quad (3.12)$$

$$(3.13)$$

where k_g is the wavenumber in the waveguide.

The cavity can also be understood using network theory. The Orpheus cavity will have a strongly-coupled port and a weakly-coupled port⁵. So it can be thought of as a two-port network. Networks can be fully described by the scattering matrix[30]. The scattering matrix element is defined as

$$S_{ij} = \left. \frac{V_i^-}{V_j^+} \right|_{V_k^+ = 0 \text{ for } k \neq j} \quad (3.14)$$

where V_j^+ is the amplitude of the voltage wave incident on port j and V_i^- is the amplitude of the voltage wave reflected from port i . S_{11} is the reflection coefficient Γ of port 1, and S_{21}

⁴Pozar 4th ed. Section 4.8 has a good illustration.

⁵Sometimes colloquially called the strong port and the weak port.

is the transmission coefficient T from port 1 to port 2.

Near resonance of the cavity, i.e., near the mode frequency, the transmitted and reflected power can be derived to be Lorentzian (Section A.2):

$$\begin{aligned} |T|^2 &= \frac{\delta y}{1 + 4\Delta^2} + C_1 \\ |\Gamma|^2 &= -\frac{\delta y}{1 + 4\Delta^2} + C_2 \\ \Delta &\equiv Q \frac{f - f_0}{f_0} \end{aligned} \tag{3.15}$$

where T is the transmission coefficient, R is the reflection coefficient, δy is the depth of the Lorentzian, and C_1 and C_2 are constant offsets.

3.1.2 Simulation of the empty cavity

The mode structure of the cavity can be simulated using Finite Element Analysis. The Orpheus mode structure was simulated using ANSYS HFSS. These simulations are resource-intensive because the cavity is electrically large, and there aren't any simplifying assumptions other than two symmetry boundaries. It took me over a year to learn to simulate Orpheus for the entire tuning range. I will outline my simulation setup so that it may help others.

The simulation for the empty cavity is shown in Figure 3.1. I model only a quarter of the cavity (in the +X,+Y quadrant). I place a Perfect H symmetry plane where the electric field is parallel to the boundary and a Perfect E symmetry plane where the electric field is perpendicular to the boundary. The cavity is surrounded by vacuum, and there exists a cylindrical radiation boundary 1 cm away from the mirrors.

I apply curvilinear meshing to all curved surfaces. For the curvilinear meshing, I use a resolution of 7 (maximum is 9). An artificial cylindrical volume surrounds the beam axis that helps the simulation converge faster. The cylindrical volume connects the weakly-coupled port aperture and the strongly-coupled port aperture. Because of the high-resolution

curvilinear meshing, the initial mesh starts with more mesh elements near the beam axis where they are needed.

To simulate the electric field, I use a driven modal solution. The eigenmode solver is untenable because of the large density of modes, particularly due to the subresonances of the dielectrics. The majority of these modes are irrelevant. Using the driven modal solution filters out modes that don't couple well to the aperture. I used the iterative solver instead of the direct solver to save RAM. I ask for a Mixed Order Basis and a Maximum Delta S of 0.001 for the convergence condition.

For simulating a mode, I use the Fast Sweep, which seems to be optimized for narrowband simulations. The frequency sweep is centered around the mode frequencies predicted by Equation 3.7 and will span ~ 30 MHz. After the frequency sweep is completed, $|S_{21}|^2$ is plotted to find the resonant frequency. The electric field is plotted at the simulated resonant frequency to check that it is the expected Gaussian TEM mode.

To find all the resonances between 15 GHz and 18 GHz, I will use an interpolating sweep. The interpolating sweep is better suited for simulating multiple resonances over a broader range. However, the interpolating sweep doesn't simulate any field data. Therefore, for plotting the electric field, I recommend using the fast sweep and focusing the frequency range around the mode of interest.

The simulated TEM₀₀₋₁₈mode is shown in Figure 3.1.

Simulations can help determine the dominant loss mechanism in the cavity. For this study, two simulations were performed, one where the mirrors are perfect conductors, and another simulation where mirrors are made of aluminum with a conductivity of 38×10^6 siemens. Figure 3.2 shows that the simulated Q are nearly identical, suggesting that the ohmic losses are negligible compared to diffraction losses. One can also calculate the beam size at the mirrors for the TEM₀₀₋₁₈mode, calculate the diffraction of the mode spilling around the cavity, and find that it overwhelmingly dominates the ohmic losses.

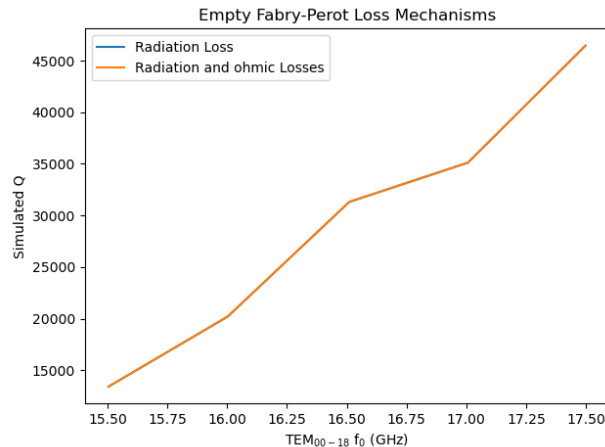


Figure 3.2: A comparison of simulated quality factor of the TEM_{00-18} mode with different sets of loss mechanisms. The blue curve corresponds to a cavity with perfectly-conducting mirrors. The orange curve corresponds to a cavity with resistive aluminum mirrors. They are nearly identical, suggesting that the dominant loss is diffraction.

3.1.3 Measurement of the empty cavity

Now I measure the mode structure of a tabletop empty Fabry-Perot cavity, shown in Figure 3.3⁶. The mechanical implementation of the cavity is described in Section 4.1. Assuming the cavity is built, the modes can be characterized using a Vector Network Analyzer (VNA), which can measure S_{21} and S_{11} . The Orpheus test stand uses a Keysight E5063A.

The reflection coefficient is needed to extract the cavity coupling coefficient of the strongly-coupled port. However, the strongly-coupled port is connected to an amplifier, which is a one-directional device. Thus the reflection measurement needs to bypass the amplifier to reach the strongly-coupled port. The amplifier is bypassed using the directional coupler, as shown in Figure 3.4. To measure the transmission and reflection coefficient of the cavity, the VNA makes an S_{21} measurement, where the port numbers correspond to the VNA. For a transmission measurement, the VNA injects a signal into the weakly-coupled port. For a reflection measurement, the VNA injects a signal into the directional coupler.

⁶This picture is out of date. But I can't seem to find a better photo of the empty cavity.

I may still refer to the reflection measurement as an S_{11} measurement. If so, port 1 corresponds to the cavity's strongly-coupled port instead of the port on the VNA.



Figure 3.3: The empty Fabry-Perot cavity consists of a flat mirror and a curved mirror. The cavity coupling scheme shown in this picture is out of date. This picture is just meant to show the cavity without dielectric plates.

For these measurements, the aperture diameter is 0.2130 in. Equation 3.13 suggests a hole larger than the waveguide dimension. This is untenable. So I started with a hole size of 0.18 in and increased the size until I got close to critical coupling for some frequency range.

The transmission and reflection measurements are used to characterize the mode structure of the cavity. First, “wide scan” measurements are taken, where the transmission and reflection coefficients are measured from 15 GHz to 18 GHz. The wide scan measurements are repeated for cavity lengths between 16 cm and 19.2 cm so that the TEM_{00-18} mode is tuned from 18 GHz to 15 GHz. Figure 3.5 shows the wide scan for different cavity lengths.

The wide scan measurements can be combined into a 2D heat map to form a mode map, as shown in Figure 3.6. The mode map illustrates the mode structure of the cavity for the entire tuning range and demonstrates the mode frequencies decrease with increasing cavity length, as is expected from Equation 3.7. The simulated modes and analytical modes are then overlaid on the mode map. All three solutions agree within a few MHz, and the

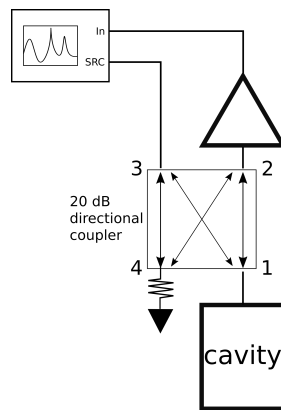


Figure 3.4: The reflection coefficient of the cavity’s strongly-coupled port is measured using a directional coupler. S_{11} can’t be measured directly because the amplifier is a non-reciprocal device. The directional coupler allows an injected signal to bypass the amplifier to reach the cavity’s strongly-coupled port. The signal from the VNA travels from port 3 of the directional coupler to port 1 (with a 20 dB attenuation), bounces off the cavity’s strongly-coupled port, then reaches port 2 of the directional coupler. The signal is amplified and then reaches the VNA’s port 2.

disagreement can be attributed to a systematic uncertainty in the absolute cavity length. The fact that the measured and simulated modes agree with the analytical modes lends great confidence to both the simulation and measurement techniques. This confidence is important because the dielectrically-loaded case doesn’t have analytical solutions for comparison. So the simulations need to be trusted to correctly identify the TEM_{00-18} mode. V_{eff} calculated from simulations also need to be trusted because it’s infeasible to measure the electric field in the cavity.

Other important parameters for the cavity modes are the loaded quality factor and the cavity coupling coefficient. To extract these parameters, “narrow scans” are taken, where the VNA measures the S_{21} and S_{11} parameters within a few Q -widths of the TEM_{00-18} mode, as shown in Figure 3.7. These narrow scans are fitted to the Lorentzian functions in Equation 3.15. From the Lorentzian fits, one can extract f_0 and Q_L . The cavity coupling coefficient β is extracted from the value of the reflection coefficient on resonance and the phase change on resonance (See A.2 for determining whether the cavity is undercoupled or overcoupled

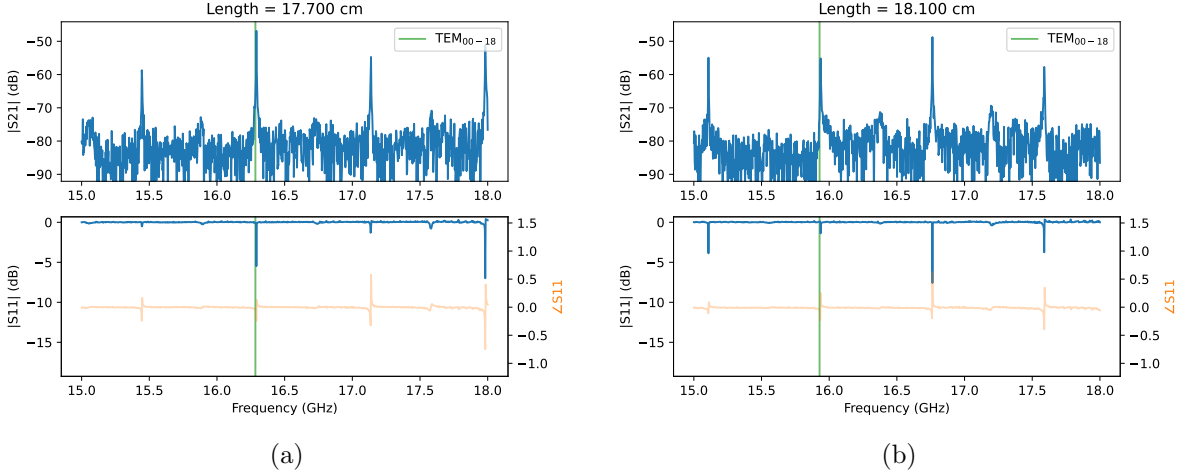


Figure 3.5: The measured scattering parameters S_{21} and S_{11} of an empty Fabry-Perot cavity as a function of frequency for two different cavity lengths. The prominent peaks are Lorentzian and correspond to different Gaussian modes.

based on the phase change.)

$$\beta = \begin{cases} \frac{1-|\Gamma_c(f_0)|}{1+|\Gamma_c(f_0)|} & \text{if } \beta \leq 1 \\ \frac{1+|\Gamma_c(f_0)|}{1-|\Gamma_c(f_0)|} & \text{if } \beta \geq 1 \end{cases} \quad (3.16)$$

For a chi-square minimization fit, the uncertainty in the reflected and transmitted power is derived to be (Section A.1).

$$\sigma_P^2 = 4\sigma_V^2 P \quad (3.17)$$

The extracted Q_L and β are shown in Figure 3.8. Q_L should be the same whether it's extracted from the transmission measurement or reflection measurement. The measured discrepancy between the quality factors is caused because of interference effects inside the directional coupler. Section A.3 derives a skewed Lorentzian function that more fully captures

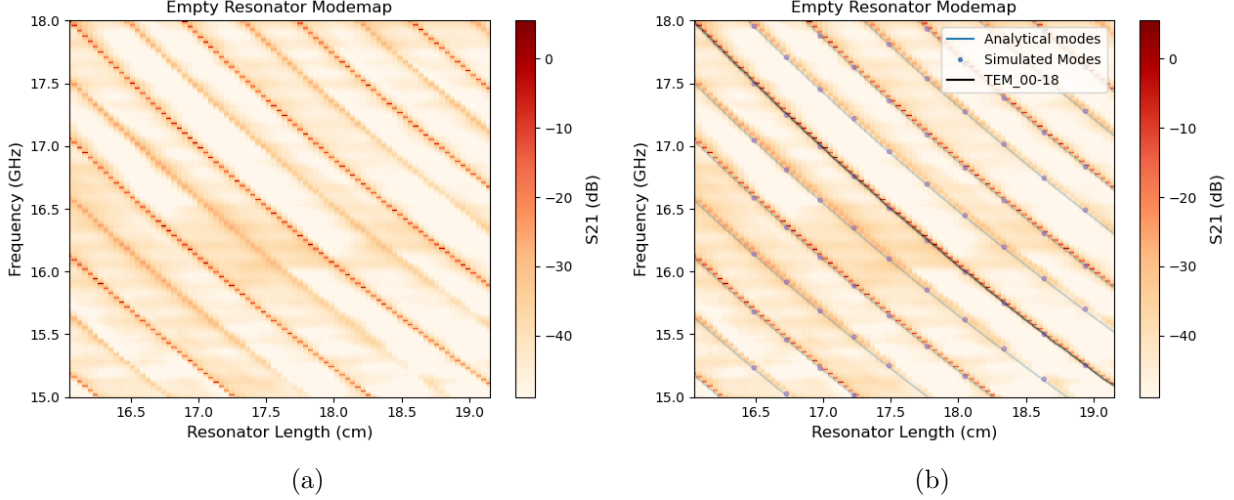


Figure 3.6: (a) Measured mode map for an empty Fabry-Perot Cavity. This is the measured scattering parameter S_{21} as a function of frequency and cavity length. The dark lines correspond to the modes of the cavity where transmission is highest. (b) The analytical TEM modes and the simulated modes are overlaid on the measurement. The agreement between the analytical formula, FEA simulation resonances, and measured resonances are within a few MHz.

the physics of the directional coupler. Figures A.5 and A.6 show that the more sophisticated reflection model fits the data better, which results in the reflection Q_L matching the transmission Q_L . The cost of implementing this model is that there is significant degeneracy in the rest of the parameters and the cavity coupling coefficient is difficult to interpret. So for the rest of the analysis, the straightforward Lorentzian model from Equation 3.15 is used to extract β .

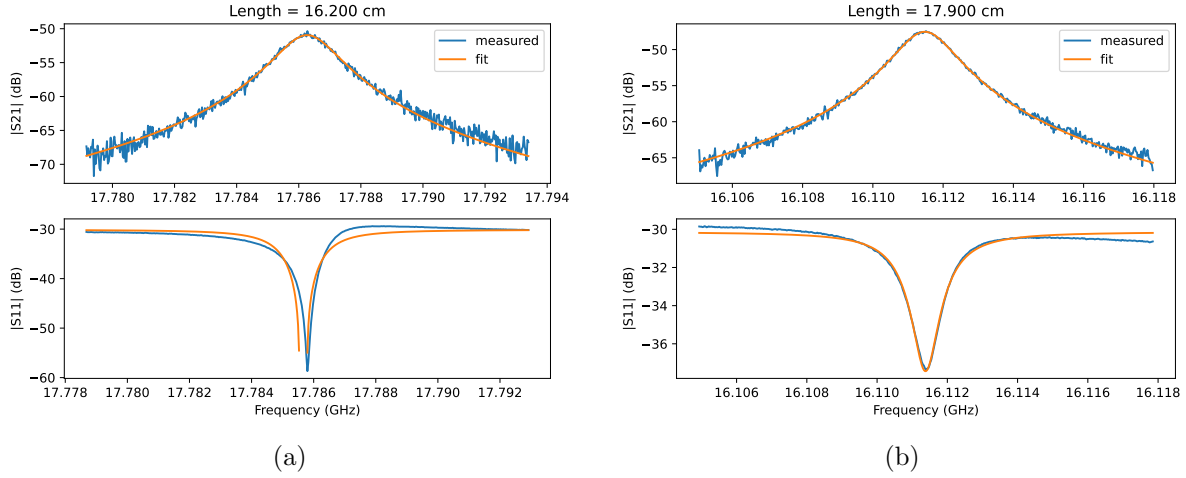


Figure 3.7: S_{21} and S_{11} of an empty Fabry-Perot cavity near the TEM_{00-18} mode for two different cavity lengths. The transmitted power is a Lorentzian. The reflected power is approximately Lorentzian, but effects discussed in Section A.3 distort the Lorentzian. Because of the skew in the data, the fitted Lorentzian in (a) predicts a small negative reflected power at resonance, which is unphysical

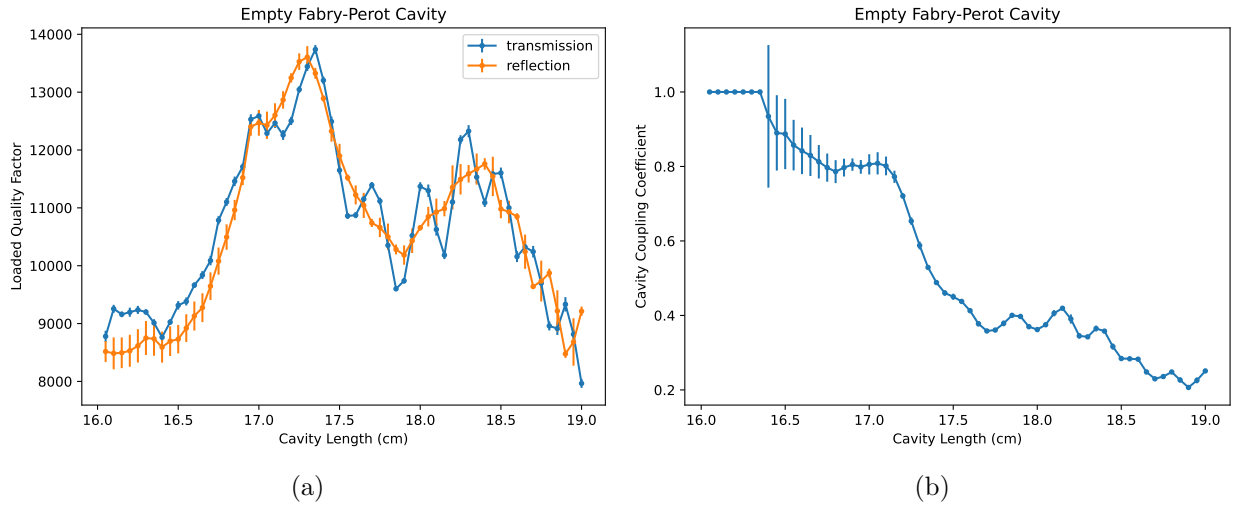


Figure 3.8: The loaded quality factor and cavity coupling coefficient of the TEM_{00-18} mode as a function of cavity length.

3.2 Adding dielectrics to a Fabry-Perot Cavity

Dielectrics are added at every fourth half-wavelength to increase V_{eff} . The TEM_{00-18} mode is engineered to be the mode of interest. Much of the formalism and techniques built up for the empty cavity apply to the dielectrically-loaded case.

3.2.1 Choice of Dielectrics

99.5% alumina sheets were purchased from Superior Technical Ceramic because they offered high-purity alumina at custom dimensions. The alumina plates are about 5.97 in \times 5.97 in \times 3 mm. The plates are octagonal to approximate a circular shape, but the straight lines are easier to machine. The dielectrics are 3 mm thick because that is approximately half a wavelength inside a $\epsilon_r = 9.8$ dielectric at 16.5 GHz.

3.2.2 Simulated TEM_{00-18} Mode

The same simulation setup used in Section 3.1.2 applies to simulating the dielectrically-loaded case. A major difference is that there aren't analytical solutions for the cavity mode frequencies. An excellent first guess is to use the free space solutions in Equation 3.7 and replace L with the optical length. For some cavity lengths, this initial guess was correct within 10 MHz. But for other frequencies, this guess was off by over 100 MHz, and I had to spend a lot of time searching for the mode.

The TEM_{00-18} mode fields are simulated and shown in Figures 3.9 and 3.10. The fields resemble their free space Gaussian counterparts in the empty cavity case with a few notable exceptions. At higher frequencies, the beam waist is at the curved mirror instead of the flat mirror. This is likely a result of the properties of the dielectric that's not considered by physical optics. The other notable difference is that the fields deviate from the Gaussian profile for frequencies near 16.35 GHz and 15.43 GHz. This is the result of a mode crossing seen in Figure 3.19. In a mode crossing, two modes become degenerate, and the field shape is distorted. The new field shape is a linear combination of the two modes, only one of which

couples to the axion or dark photon.

The intruding mode (or set of modes) is shown in Figure 3.11. This is the mode at 16.86 GHz when the cavity is $L = 14.4219$ cm. The measured mode map (Figure 3.19) shows that this mode tunes more slowly than the TEM_{00-18} mode. The vector field suggests that this mode is some combination of a TM mode (from the arrows that point into the mirror) and a TEM mode (from the arrows that point perpendicular to the beam axis). TEM modes tune with the cavity length, but TM modes do not. Hence, the intruding mode tunes more slowly than the more pure TEM modes.

The simulated resonant frequency and quality factor are plotted in Figure 3.12. The simulated cavity is very undercoupled, so $Q_L \approx Q_u$. The simulated V_{eff} and form factor are plotted in Figure 3.13. All plots have a notch at 16.4 GHz at the location of the mode crossing. The discontinuity in the resonant frequency is actually caused by the simulation not accurately modeling the mode crossing. The mode map in Figure 3.19 shows that the TEM_{00-18} mode tunes very smoothly, even through the mode crossing.

The form factor in Figure 3.13 shows that only about 2% of the cavity volume is used to detect axions. That's not good, but the cavity properties haven't been optimized for maximal V_{eff} and Q_L . Changing mirror curvatures, dielectric thicknesses, dielectric positioning, and adding more dielectrics will increase the form factor.

The width of the notch in the simulated Q matches the width of the notch of the simulated V_{eff} . This suggests that the measured Q can be used as a proxy for the width of the mode crossing.

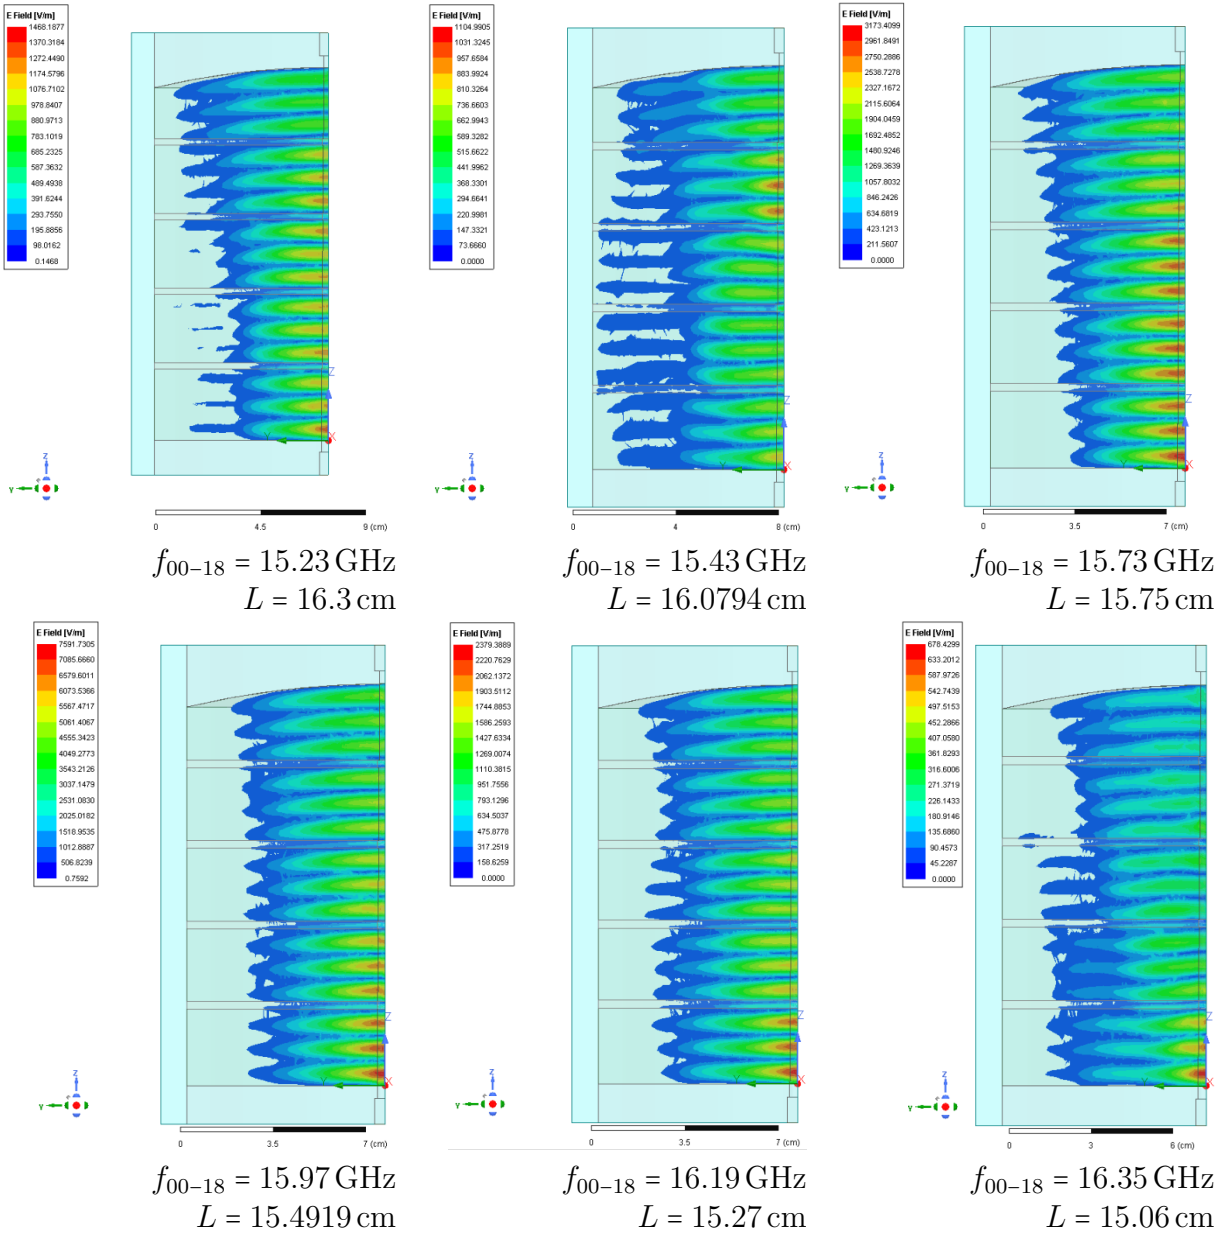


Figure 3.9: The simulated electric field magnitude of the TEM_{00-18} mode with different cavity lengths. There are mode crossings around 15.4 GHz and 16.3 GHz.

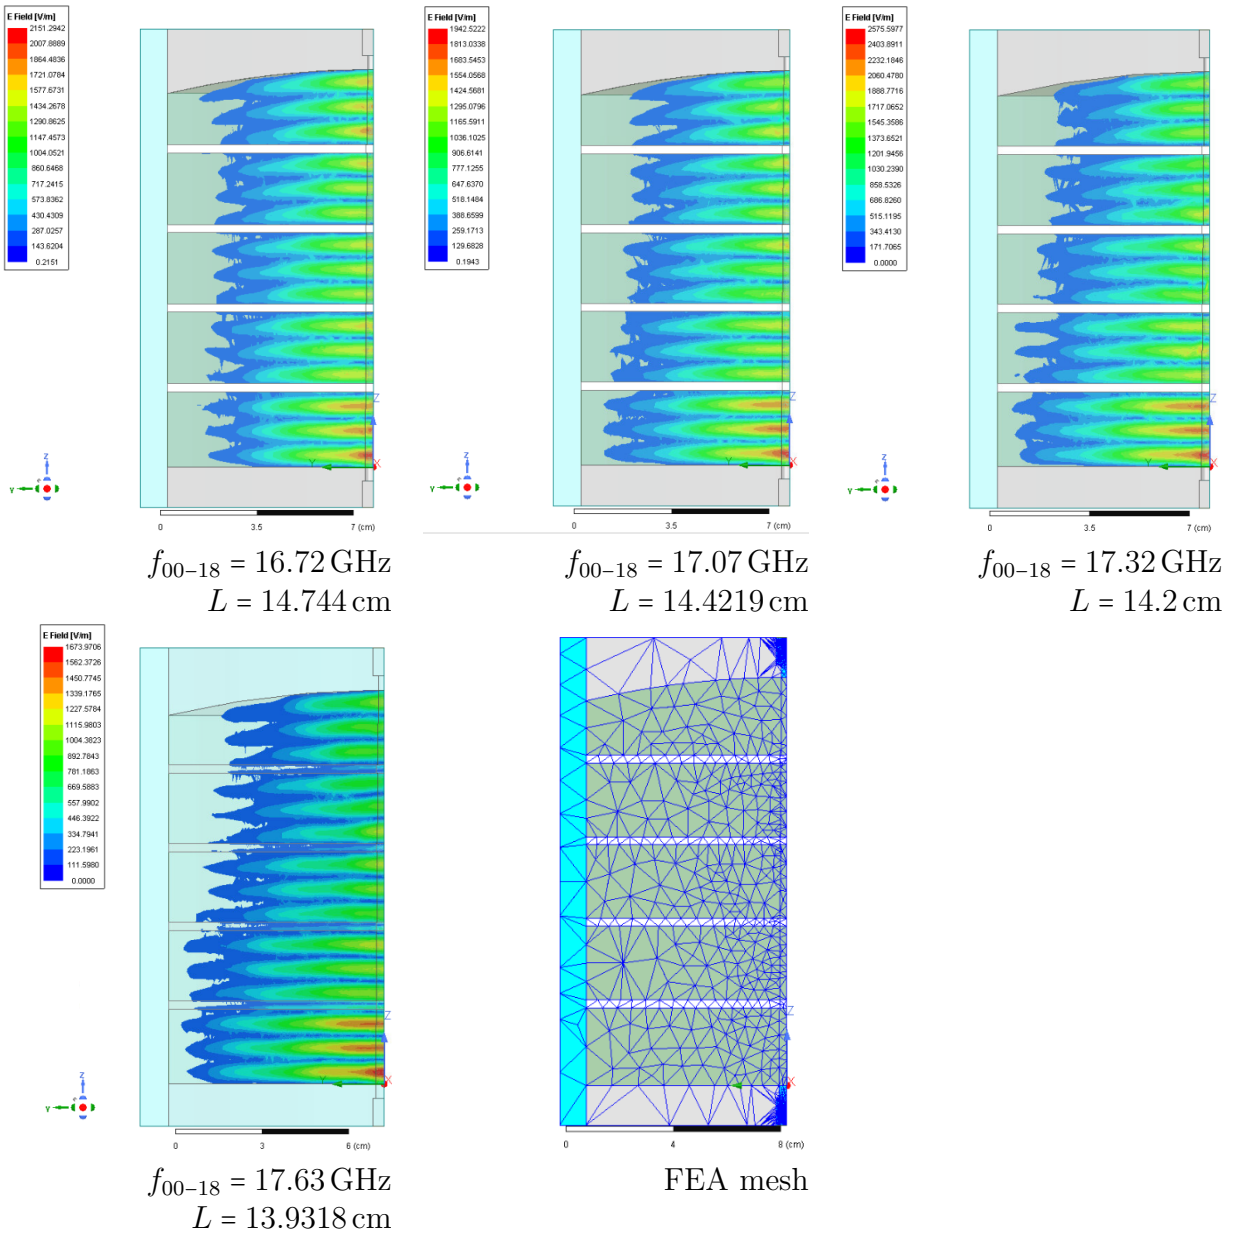


Figure 3.10: The simulated electric field magnitude of the TEM_{00-18} mode with different cavity lengths. The field profile starts to deviate from what is expected from physical optics. One intuit from physical optics that the waist of the beam would be on the flat mirror instead of the curved mirror. But that is not the case.

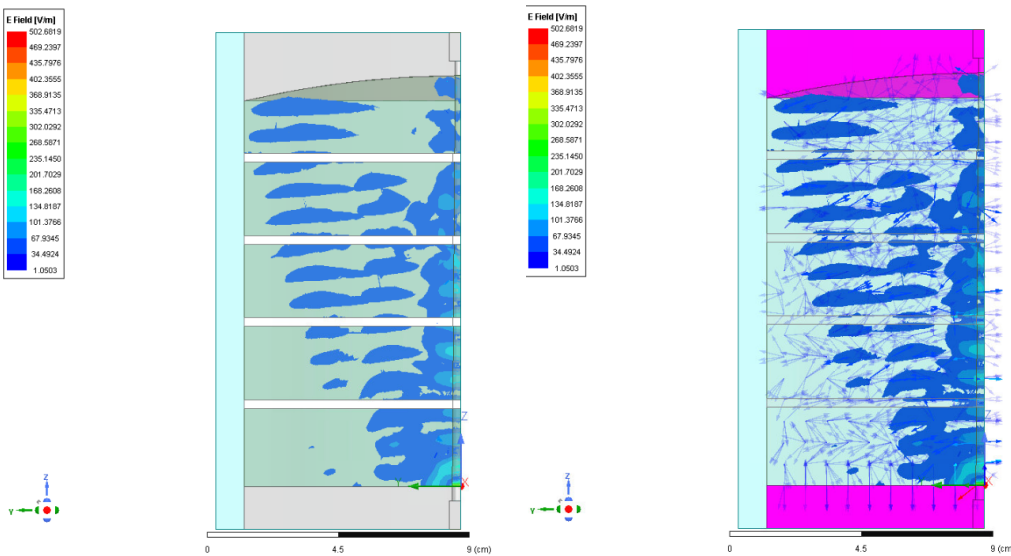
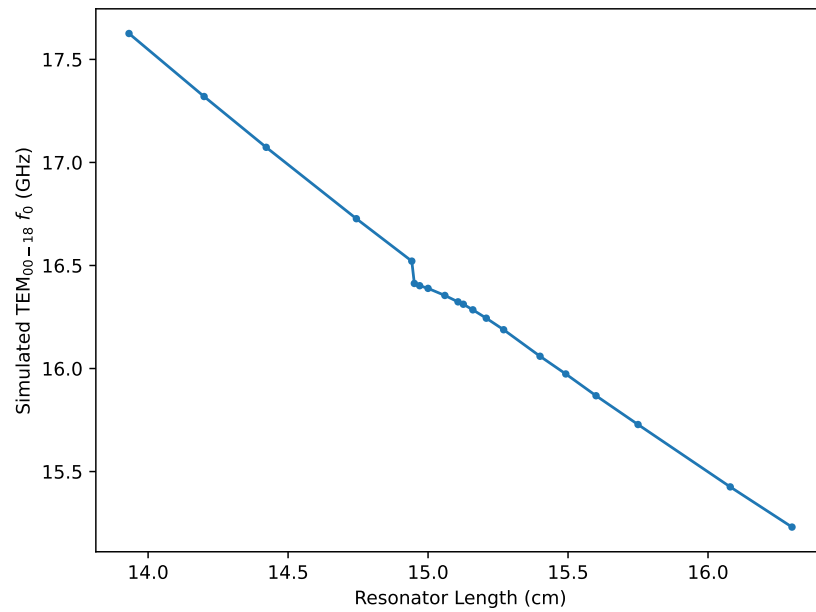
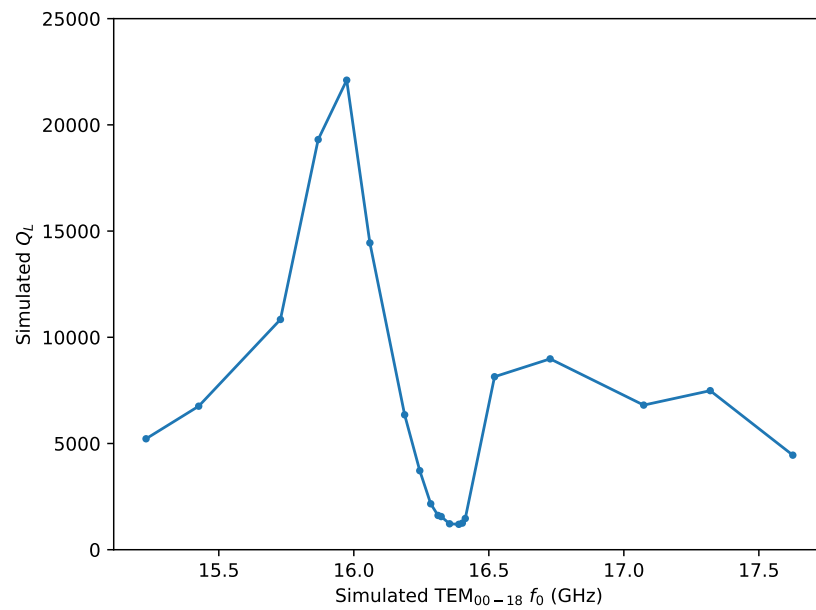


Figure 3.11: Intruder mode while it's separated from the TEM_{00-18} mode. Here, the cavity length is $L = 14.4219$ cm. The TEM_{00-18} resonance is at $f_{00-18} = 17.07$ GHz while the intruder mode frequency is at 16.86 GHz. They will cross around $L \approx 14.95$ cm, $f_{00-18} \approx 16.4$ GHz. Left: Magnitude of the intruder mode electric field. Right: Complicated vector field superimposed on the magnitude of the electric field. The field is complicated, but it looks like some mixture of a TM mode (based on the field pointing into the mirror) and a TEM mode (based on the vector field being transverse to the axis for much of the interior of the cavity). TM modes don't tune with the changing cavity length, which may explain why the intruder mode tunes slower than the TEM_{00-18} mode.

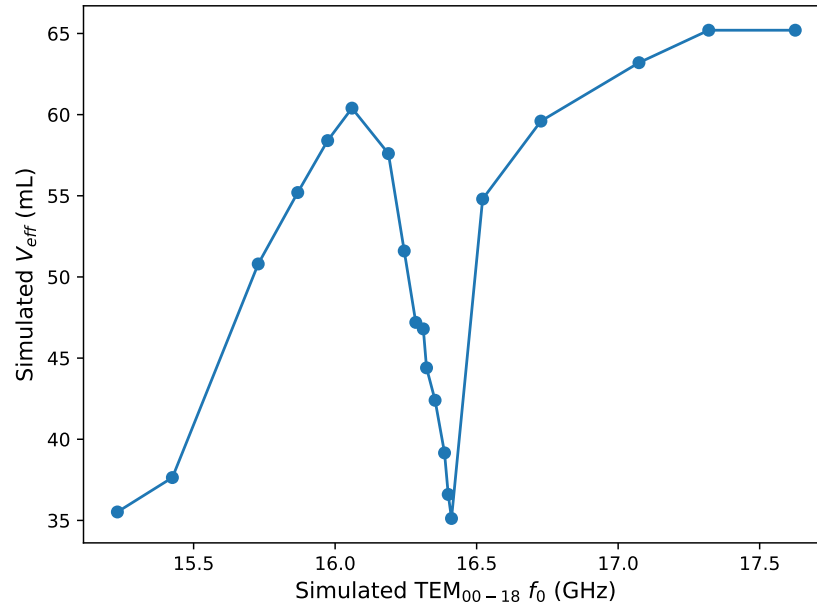


(a)

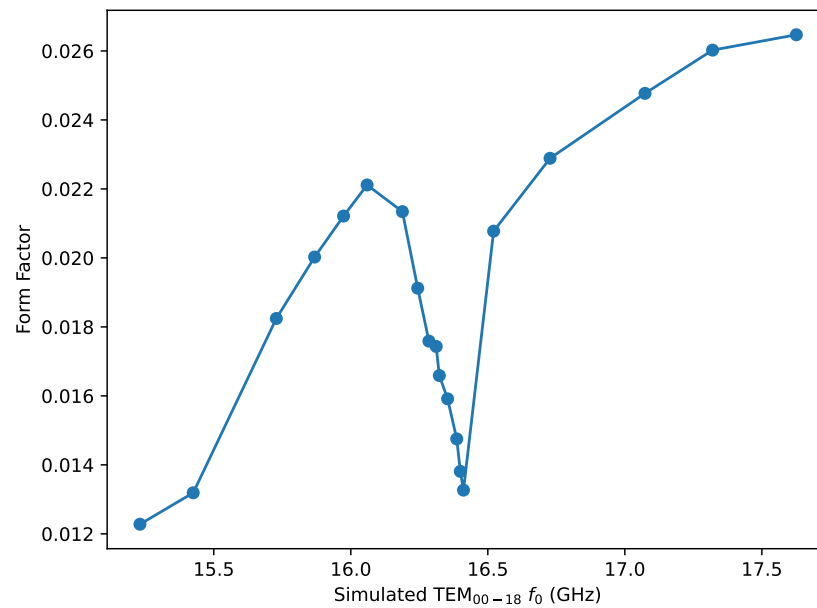


(b)

Figure 3.12: The simulated frequency and loaded quality factor for the TEM₀₀₋₁₈ mode. (a) The simulations struggles to simulate the mode crossing, hence the discontinuity in resonant frequency at around 16.4 GHz. Perhaps we need higher accuracy requirements. (b) The simulated loaded quality factor. In our simulations, the cavity is severely undercoupled, so this is close to the unloaded quality factor.



(a)



(b)

Figure 3.13: The effective volume and form factor for the TEM_{00-18} mode. The sharp dip at 16.4 GHz corresponds to the mode crossing.

3.2.3 V_{eff} Tolerance to Misplaced Dielectrics.

Previous simulations assumed the dielectrics are evenly spaced inside the cavity. But V_{eff} is an important parameter that can only be obtained through simulation. Thus, it's essential to understand the uncertainty in V_{eff} and how much it tolerates positioning errors of the dielectrics. Understanding this tolerance will be crucial for the data analysis (Section 6.3).

Let the position error δ be defined as the deviation from the evenly spaced configuration (actual position - intended position). Orpheus is oriented vertically, such that the flat mirror is on the top and the curved mirror is on the bottom, as shown in Figure 4.1. A positive position error means the dielectric plate is closer to the curved mirror than what would be intended by the evenly spaced configuration. This sign convention is consistent with the direction of increasing or decreasing cavity length (the flat mirror is fixed, and the curved mirror moves to adjust the cavity length). The middle two dielectric plates are constrained by the scissors jacks and stay evenly spaced between the top and bottom dielectric plates.

The dielectric plate closest to the curved mirror is the bottom dielectric plate, and the plate closest to the flat mirror is the top dielectric plate. This is reverse from the simulation shown in Figure 3.9, where the flat mirror is on the bottom, and the curved mirror is on top. However, the physical cavity (Figures 3.17 and 4.4) has the curved mirror on the bottom, and I will stick with this convention.

Figure 3.14 compares the electric field with and without any error position. The mode is robust against error positions as much as 2 mm.

Figure 3.15 shows how the error position affects V_{eff} at a fixed cavity length of $L = 15.66$ cm. Surprisingly, V_{eff} improves with negative position error. Comparing the unperturbed and perturbed cases in Figure 3.14 provides intuition for why V_{eff} increases with negative position error. For the unperturbed case, the portion of the wavefront radially farther away from the beam axis gets pushed towards the dielectrics, reducing V_{eff} . Pulling the dielectrics away from the curved mirror also pulls it away from the curved wavefront. As I think about it, the negative position errors give the mode more room to breathe. For

positive position errors, the curved wavefront gets pushed more into the dielectrics, and V_{eff} decreases.

It appears that an position error of ≈ -1 mm optimizes V_{eff} . So not only are perturbations tolerated, they are actually beneficial.

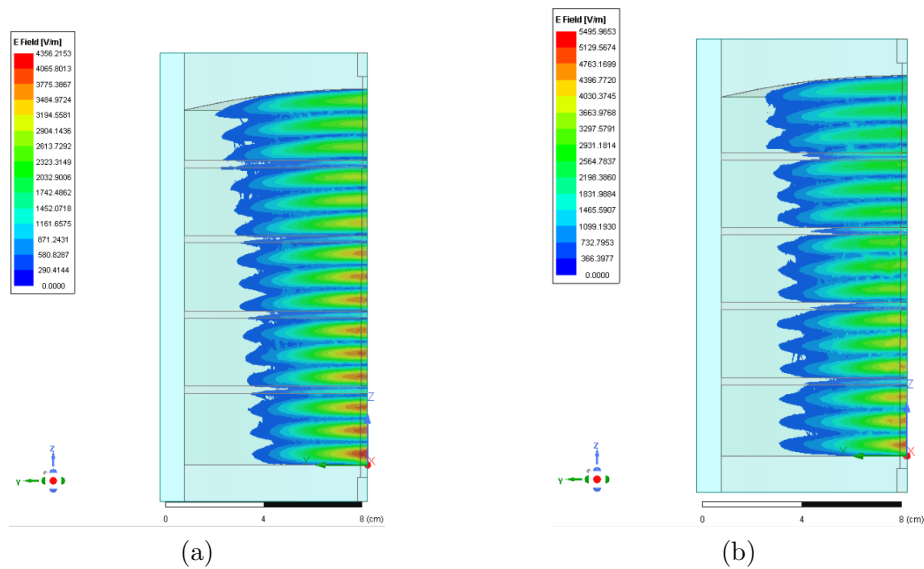


Figure 3.14: The effect of dielectric plate position error on the electric field, simulated at about 15.82 GHz. (a) No position error. (b) $\delta_{top} = -0.52$ mm, $\delta_{bottom} = -2$ mm. Note the δ_{bottom} corresponds to the bottom dielectric plate closest to the curved mirror, so it appears as the uppermost plate in this figure.

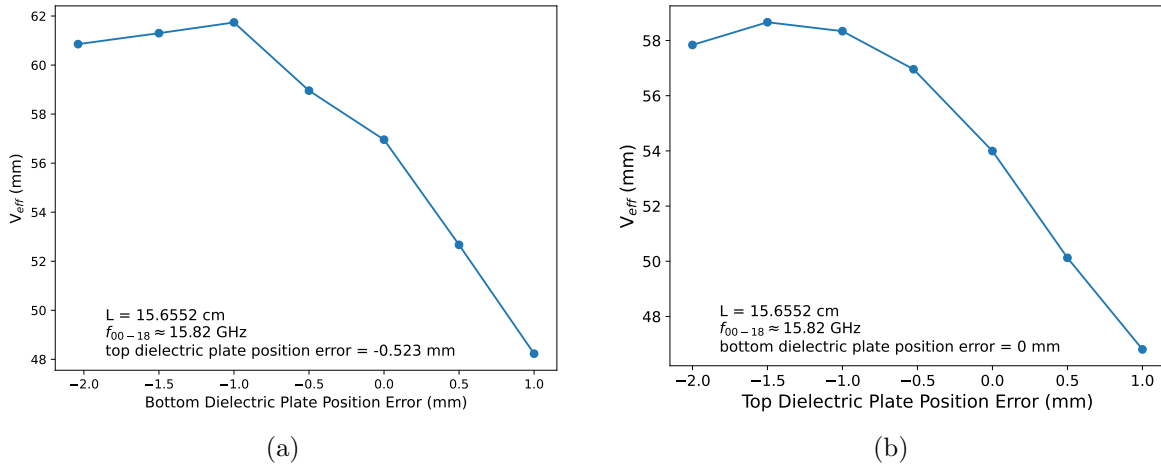


Figure 3.15: The effect of position error on the effective volume. This study suggests that it would actually be more optimal to deviate from the evenly-spaced configuration. At 15.82 GHz, it appears that the effective volume can increase by more than 8% by moving the dielectrics further away from the curved mirror.

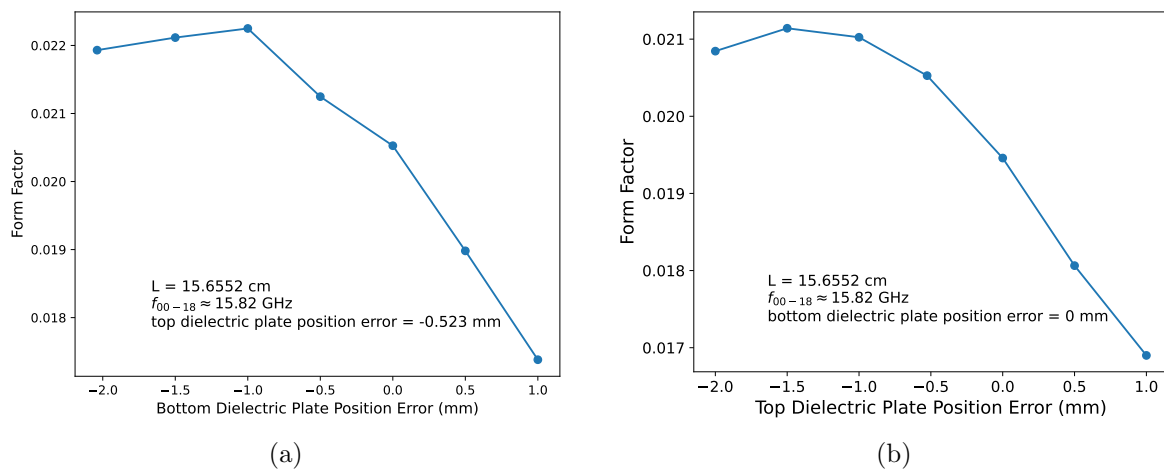


Figure 3.16: The effect of position error on the form factor.

3.2.4 Characterization of TEM_{00-18} Mode

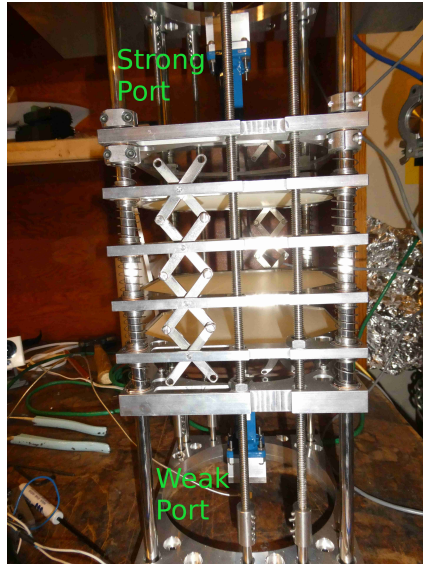


Figure 3.17: The tabletop Orpheus cavity.

Figure 3.17 shows a tabletop dielectric-loaded Fabry-Perot cavity. The TEM_{00-18} mode is characterized in the same way as the empty cavity (Section:3.1.3). The transmission and reflection coefficients are measured, and Q_L and β are extracted from those measurements.

The wide scan measurements are plotted in Figure 3.18. They are more feature-rich than the empty case because there are more resonating substructures in the cavity⁷. The wide scan measurements can be combined to form the mode map shown in Figure 3.19.

The measured wide scans can be compared to the simulated wide scans, as shown in Figure 3.20. The simulated modes are then overlaid on the mode map. With some discrepancies, the simulated modes match the measured mode structure well. There is some deviation at higher frequencies and smaller cavity lengths. This discrepancy has several possible explanations:

1. There is a fixed systematic uncertainty in the cavity length, which results in a larger

⁷The dielectrically-loaded cavity can be thought of as a series of resonators.

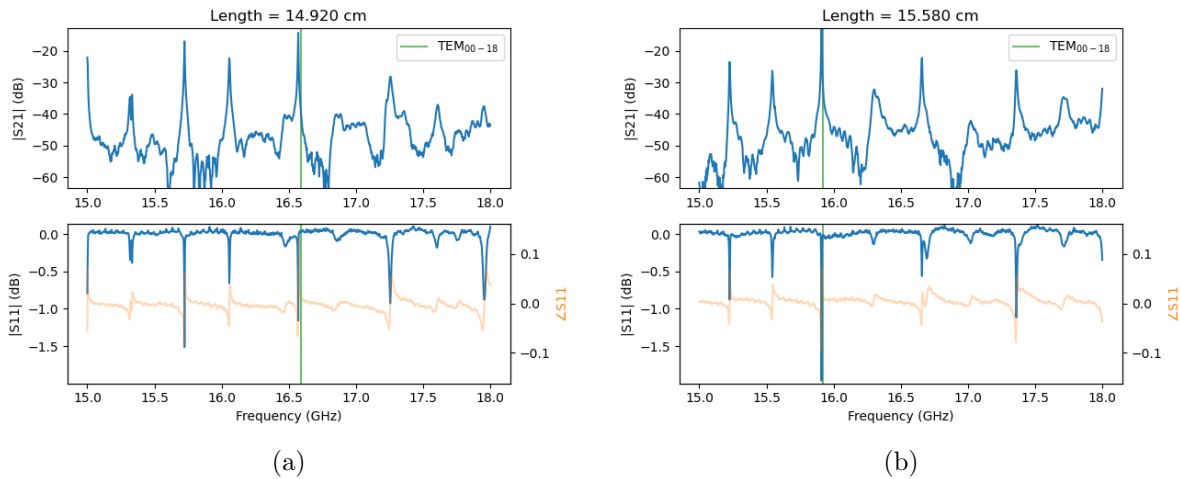


Figure 3.18: The measured scattering parameters S_{21} and S_{11} of an dielectrically-loaded Fabry-Perot cavity as a function of frequency for a fixed cavity length. The prominent peaks are Lorentzian and correspond to different Gaussian modes.

relative uncertainty for smaller cavity lengths.

2. The simulation assumes a fixed dielectric constant. But the dielectric constant may decrease notably at higher frequencies.
3. I am using a single interpolating sweep to simulate the cavity modes from 15 GHz to 18 GHz. The simulation mesh may not have enough resolution for the interpolating sweep to accurately locate the cavity modes at higher frequencies.
4. The simulated geometry doesn't capture the entire cavity geometry. For example, I don't simulate the metallic dielectric plate holders, scissor jacks, or guide rails. All these components may shift the resonant frequencies. But I believe this effect to be subdominant because the empty cavity simulations agreed so well with measurement and analytical predictions.

Regardless, it is clear that the simulated mode structure follows the same patterns as the measured mode structure.

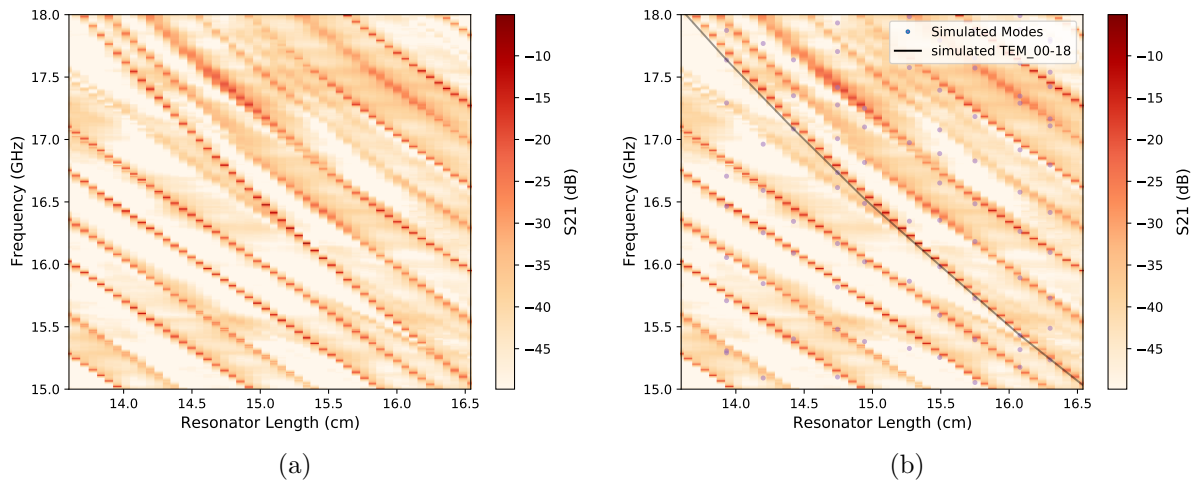


Figure 3.19: Measured mode map for dielectrically-loaded Fabry-Perot cavity. This is the measured scattering parameter S_{21} as a function of frequency and cavity length. The dark lines correspond to the modes of the cavity where transmission is highest. (b) The simulated modes are overlaid on the measurement.

The quality factor and cavity coupling coefficient for the TEM_{00-18} mode are shown in Figure 3.21. At room temperature, $Q_L \sim 5000$ and the $\beta \sim 0.4$. There is a notch in the middle of the tuning range, at about 16.4 GHz. This notch corresponds to the mode crossing seen on the mode map in Figure 3.19 and is corroborated by simulations discussed in Section 3.2.2.

This measurement also suggests Orpheus has a natural bandwidth. The loaded Q drops off below 16 GHz and above 17 GHz. This makes sense because the dielectric thickness is chosen to be about $\lambda/2$ thick at 16.5 GHz. The more the dielectric thickness deviates from $\lambda/2$, the more destructively interfering the dielectrics become. If the dielectric thickness is $\lambda/4$, then the wave destructively interferes in the dielectric and would not transmit through the cavity. My intuition tells me that the more dielectrics there are in the cavity and the higher the dielectric constant, the smaller the cavity bandwidth. This system is reminiscent of a cavity filter.

At cryogenic temperatures, the dielectric losses are expected to drop considerably. I expect the cryogenic quality factor to be similar to that of the empty cavity in Figure 3.8.

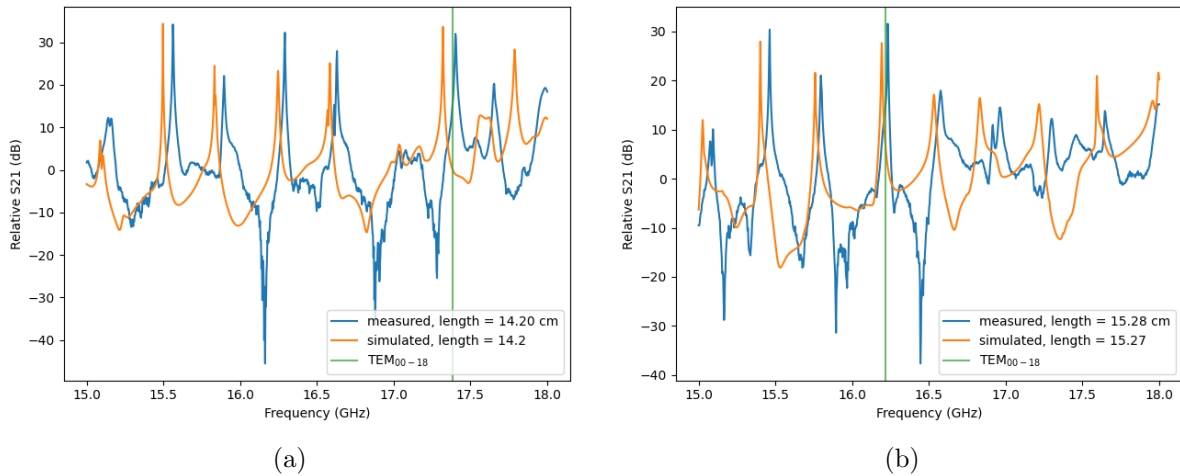


Figure 3.20: The measured transmission wide scans are compared to the simulated wide scans. The general shape between the two matches. There is some discrepancy in the location of the resonances. This may be improved by demanding more accuracy from the wideband simulation, but this improving the agreement doesn't have much bearing on the science results.

I predict that the quality will double. Since $\beta \propto Q_0$, I expect the cavity coupling to also double. With many caveats, that is roughly what is observed in the cryogenic measurement (Figures 6.7 and 6.8).

The simulated Q_u can be compared to the measured Q_u , as shown in Figure 3.22. While the values between the two don't match, the general shape matches⁸. The notch width of the measured Q matches that of the simulated Q . And the notch in simulated Q matches the notch in V_{eff} . This suggests that notches in the measured quality factor can be used as a proxy for notches in the V_{eff} . This is important because V_{eff} can't be measured directly.

⁸The mismatch is to be expected because simulating quality factors is notoriously difficult. One has to make a lot of assumptions about material properties and alignment. To have a more accurately simulated Q_u , one may have to include small features like screws and gaps into the simulation. This makes the mesh size untenably large for a desktop computer.

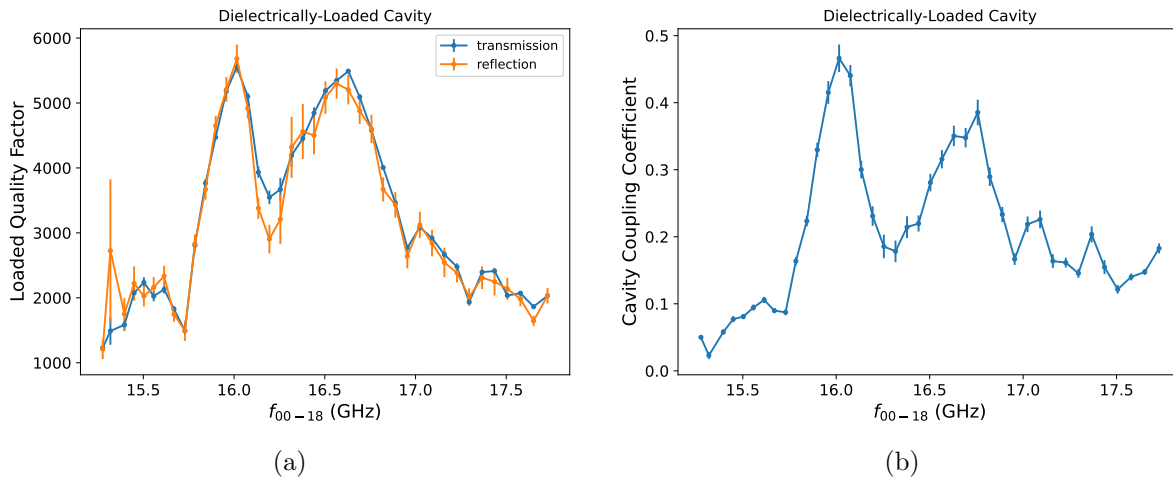


Figure 3.21: The measured loaded quality factor and cavity coupling coefficient of the TEM_{00-18} mode as a function of cavity length. Measurements are done in air at room temperature.

3.3 Coupling to Orpheus

Ultimately, power is extracted from the Orpheus cavity with aperture coupling and with a WR62 waveguide. The aperture is about 0.213 in in diameter and 0.15 in thick. There are several alternatives and modifications to this scheme that were explored but not pursued.

3.3.1 Coaxial antenna vs. waveguide aperture coupling

The downside to aperture coupling is that the hole size is fixed for the entire tuning range, and the coupling coefficient can't be adjusted in-situ.

Using a coaxial antenna to couple to the cavity is attractive because there are well-established ways to adjust the coupling; a motor could adjust the insertion depth of the antenna. An antenna setup is shown in Figure 3.23. Two holes are drilled on the flat mirror for the antenna ports. The weakly-coupled antenna is placed 3 in away from the strongly-coupled antenna at the center.

Unfortunately, the antennas performed worse than aperture coupling. The quality factor

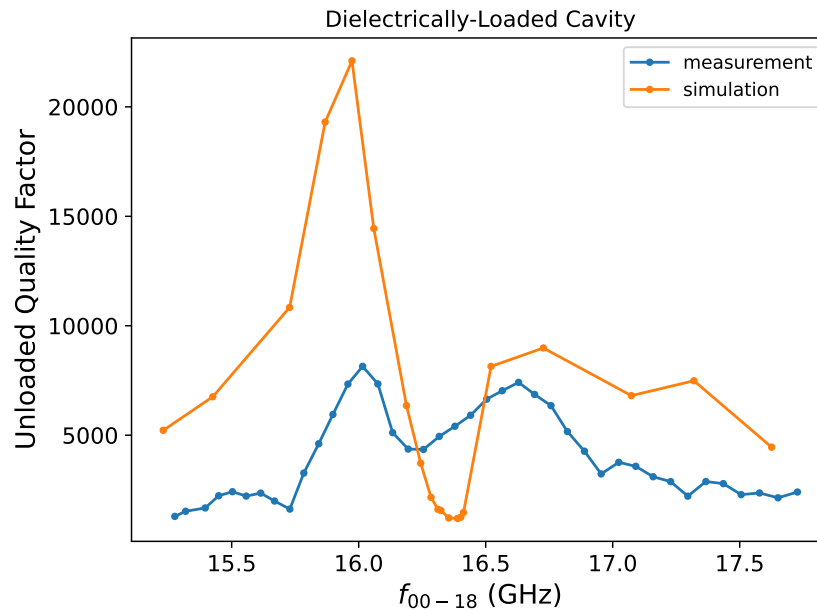


Figure 3.22: The simulated Q_u is compared to the measured Q_u . The quality factor is notoriously difficult to simulate, so the values don't match. But the general shape between the two curves are similar.

was worse by a factor of two or three. Even more unfortunate, the degenerate modes split and were shifted by about 100 MHz from where they are predicted to be. In contrast, the modes measured by aperture coupling are consistent with what is expected. It is evident that this antenna coupling mechanism strongly perturbs the cavity mode structure.

Antenna coupling is also harder to simulate than aperture coupling. This is because the aperture coupling affords me two symmetry boundaries to reduce the needed computing power. However, an antenna would only allow me at most one symmetry boundary.

It's possible that antenna coupling may have worked better if the antennas were inserted from the side of the cavity instead of through the mirror. But it would have been challenging to design a cavity coupling adjustment mechanism.

For all these complications, coaxial antenna coupling was not pursued any further.

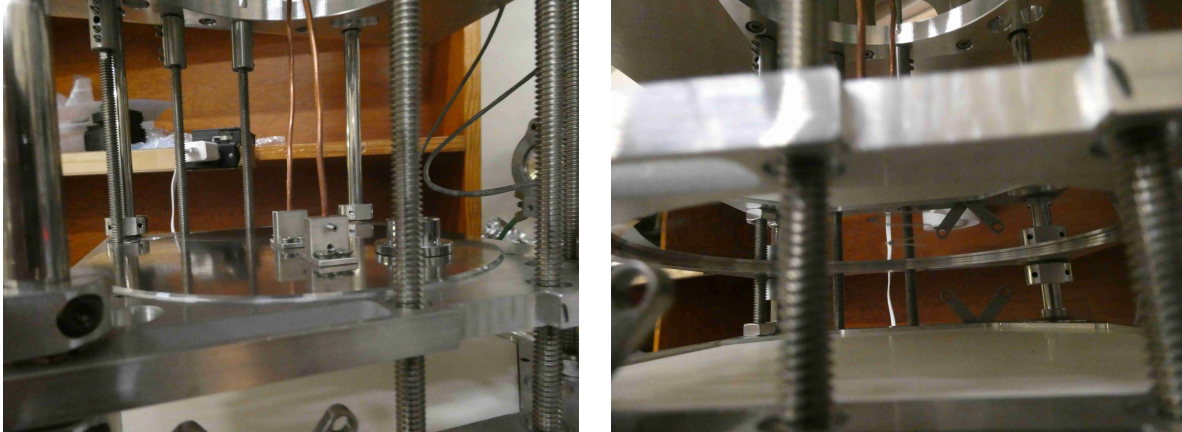


Figure 3.23: Two dipole antennas were inserted into the cavity through the flat mirror and bent so they are transverse to the axis of the cavity. The antenna in the center of the mirror corresponds to the strongly-coupled port; the antenna further from the center corresponds to the weakly-coupled port.

3.3.2 *How thin should the hole be?*

The cavity coupling can be increased by making the hole as thin as possible while maintaining a sturdy mirror surface. To understand why thinner apertures increase coupling, one can examine Equation 3.13. One can also think of the exponential decay of evanescent modes. The aperture thickness is reduced by machining a pocket for a rectangular waveguide into the back of the mirror, as shown in Figure 3.25. The mirror thickness at the pocket is about 0.15 in. The thin surface improves the coupling substantially. Unfortunately, the thin mirror surface is warped in a subtle way around the beam axis, likely leading to a reduced quality factor.

3.3.3 *Diagnosing and avoiding crosstalk between ports*

Crosstalk occurs when photons from one port reach the other port without interacting with the cavity. Crosstalk is problematic because it can interfere with the cavity signals and distort the Lorentzian peaks.

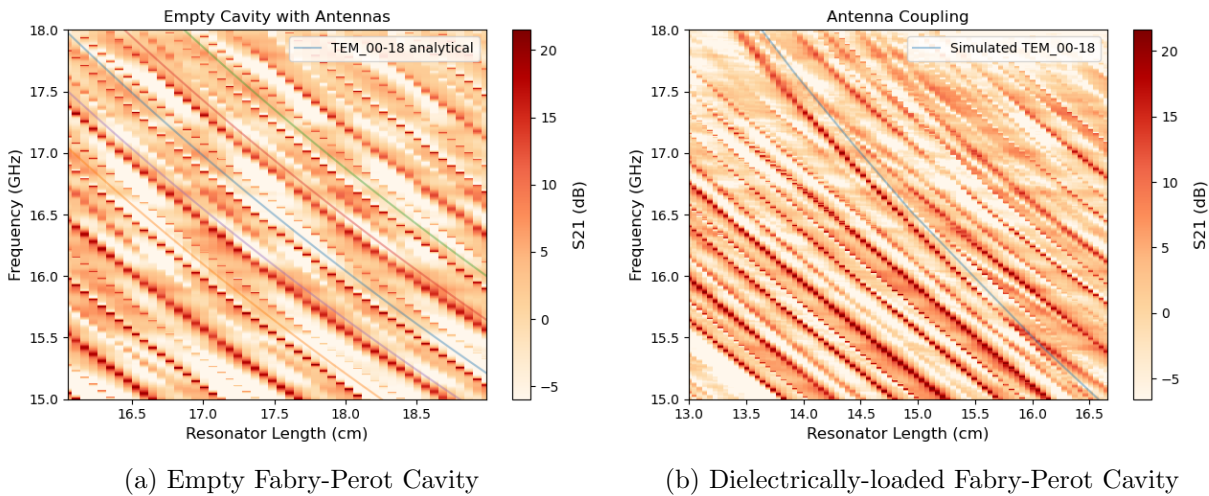


Figure 3.24: The mode map of an empty Fabry-Perot cavity and dielectrically-loaded Fabry-Perot cavity that’s coupled using coaxial antennas. (a) The doubly-degenerate TEM modes are split into two. The modes with the polarization that most closely lines up with the orientation of the dipole antennas are pushed lower in frequency by about 100 MHz compared to where they are expected to be from the analytical formula. The other polarization is pushed to higher frequencies than the analytical expectation and couples more weakly to the antennas. (b) The TEM_{00-18} mode is lower than the simulated prediction by about 100 MHz.

Crosstalk makes it challenging to measure f_{00-18} and Q_L because the transmission and reflections no longer look like Lorentzians around the resonances. In fact, sometimes, a resonance in a transmission looks like a dip instead of a peak.

The first Orpheus design had the weakly-coupled port and strongly-coupled port on the flat mirror 3 in apart from each other. The corresponding transmission measurements and mode map are shown in Figure 3.26. The effects of crosstalk are evident when narrow scan measurement and mode map of Figure 3.26 are compared to Figures 3.7 and 3.19, where the aperture ports are on opposite ends of the cavity and crosstalk is negligible. Crosstalk between the strongly-coupled port and weakly-coupled must be mitigated for proper characterization of the cavity.

Reference [44] details how to understand and diagnose crosstalk. One can plot S_{21} near resonance on the complex plane, as shown in Figure 3.27. This is known as a Q circle.

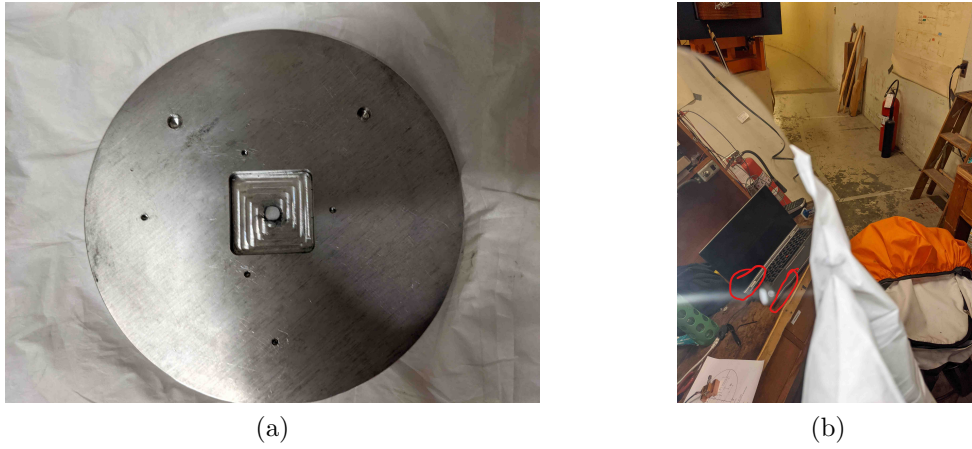


Figure 3.25: The waveguide fits into a mirror pocket shown in (a). The mirror thickness at the pocket is 0.15 in. (b) The thin mirror surface improves the cavity coupling, but resulted in warping of the mirror surface around the beam axis. This likely reduced the Q .

Crosstalk will shift the Q circle away from its canonical position. This crosstalk shift was observed for all transmission measurements when the ports were next to each other.

Moving the weakly-coupled port to the curved mirror mitigated crosstalk, as shown by the Q circle on Figure 3.28.

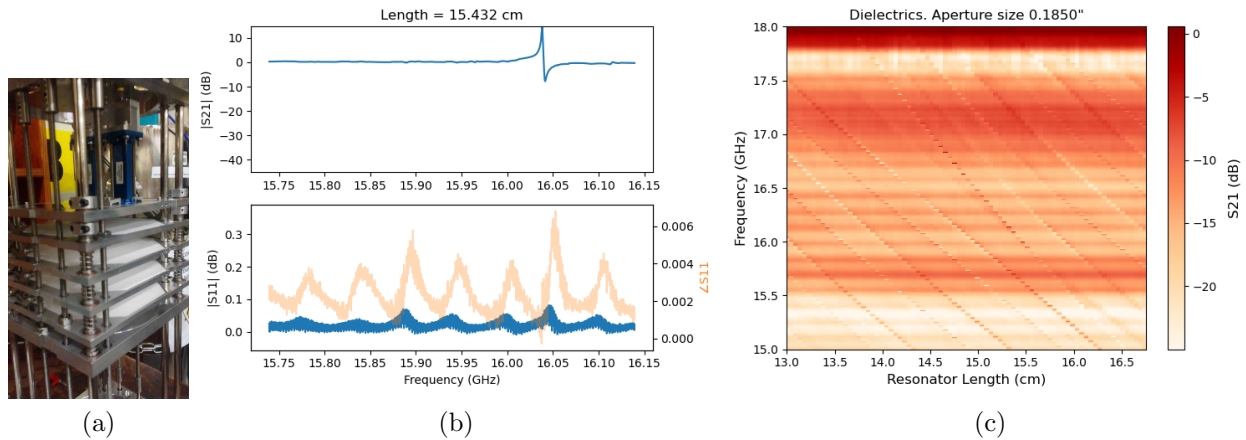


Figure 3.26: Having the strongly-coupled port and weakly-coupled port right next to each other resulted in significant crosstalk. The crosstalk interferes with the cavity signal, which causes the Lorentzian peaks to distort. In the most perverse case, the crosstalk and cavity signal interfere perfectly destructively, causing a dip in the S_{21} measurement where a Lorentzian peak is expected.

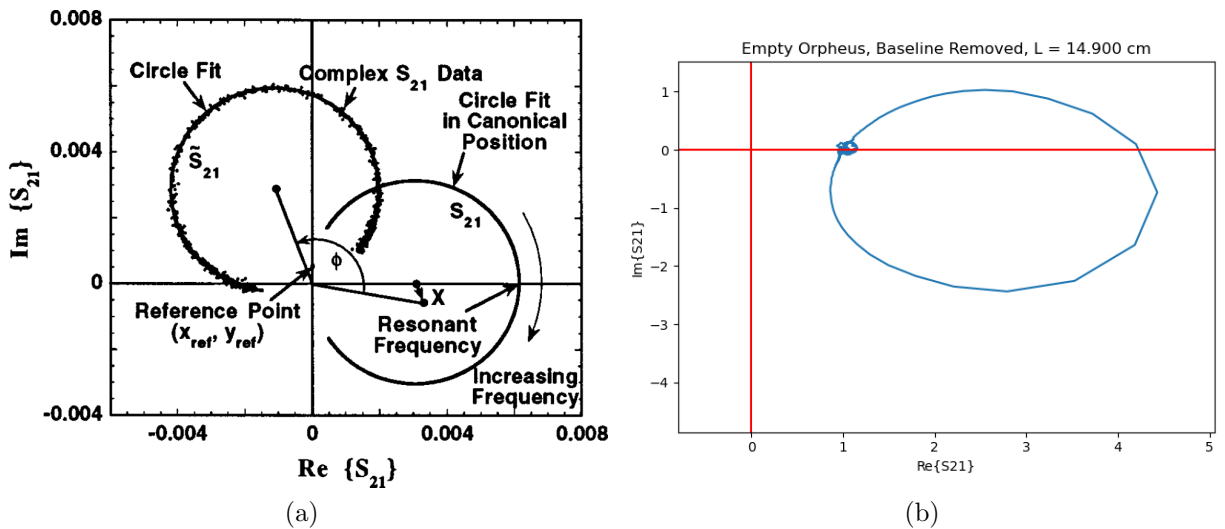
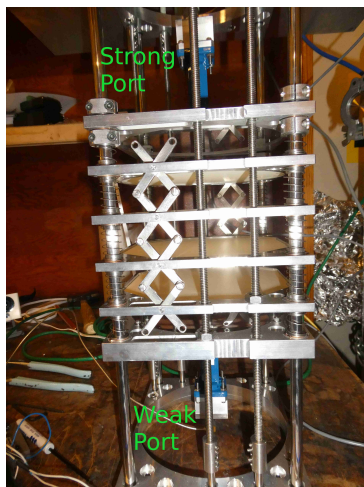
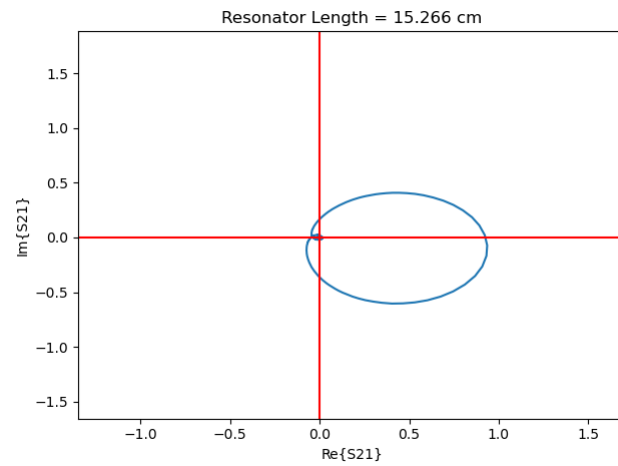


Figure 3.27: (a) S_{21} can be plotted in the complex plane. A cavity without any crosstalk or phase shift will have the S_{21} in the canonical position. Crosstalk will cause the q-circle to be shifted from this canonical position[44]. (b) This shift is observed for the measured S_{21} for the configuration when the strongly-coupled and weakly-coupled ports were next to each other.



(a)



(b)

Figure 3.28: Moving the weakly-coupled port to the other side of the cavity mitigated crosstalk.

3.3.4 Adjustable cavity coupling

Theoretically, impedance matching networks can be designed to improve coupling to the Orpheus cavity without relying solely on the size of the aperture. Pozar 4th ed. Ch.5 describes how to use single-stub tuning and double-stub tuning to match impedances. There exists waveguide tuners (shown in Figure 3.29) that are analogous to double-stub tuners[45]. The waveguide tuner has a series of stubs spaced 0.35in apart, and these stubs can be inserted into the waveguide to change its impedance. The stub makes the waveguide either more inductive or more capacitive, depending on how far it's inserted into the waveguide. One could use a Smith chart and the double-stub tuning recipe outlined in Pozar to achieve an impedance matched network. Unfortunately, I just never got this to work. Adjusting the stubs seemed to just change the susceptibility displayed on the Smith chart without ever moving the impedance closer to the center of the chart. I tried to make this work for more than two weeks and moved on.

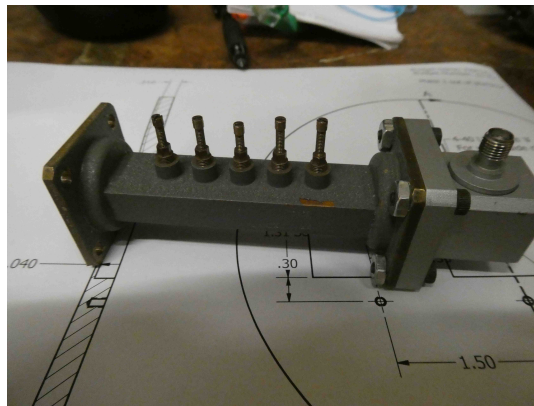


Figure 3.29: Waveguide tuner used to attempt to create an impedance matching network. I was unsuccessful in using it in that manner.

Chapter 4

ORPHEUS CRYOGENIC AND MECHANICAL DESIGN

Previous chapters discussed the conceptual design and electrostatics of an open cavity with periodically-placed dielectrics. It is now time to talk about the practical implementation of such a cavity. There must be a mechanical structure that moves the mirrors and dielectric plates, and the mechanical structure must work in a vacuum and cryogenic environment. The cavity should also work as a room-temperature, tabletop setup for easy RF characterization and rapid prototyping. The machining tolerances need to be tight enough to maintain good alignment and achieve good Q , but there needs to be enough slop¹ in the system so that the tuning mechanism works even with small amounts of misalignment.

4.1 *Orpheus Cavity Mechanical Design*

Most of the Orpheus mechanics were designed by Richard Ottens. I prototyped some of the mechanical concepts and made adjustments, and afterward, I commissioned the UW Physics machine shop to fabricate the cavity components in November 2019. I received the completed parts around February 2020². The Orpheus cavity CAD model is shown in Figure 4.1. I will describe the main design choices.

The mirror and alumina plates are held in place by aluminum holders (Figure 4.2). They rest on a lip and rely on gravity for mechanical stability³. The lip can be seen in the dielectric plate assembly in Figure 4.3. The dielectric plates sit inside the aluminum holder with lots of slack to accommodate the differential thermal expansion coefficients during cooldown. If the

¹Also known as backlash, lash, or play. Mechanical systems need it to prevent jamming.

²For those reading this in the far future, COVID-19 lockdowns started around March 2020.

³This is not a viable choice if the cavity needs to rest horizontally.

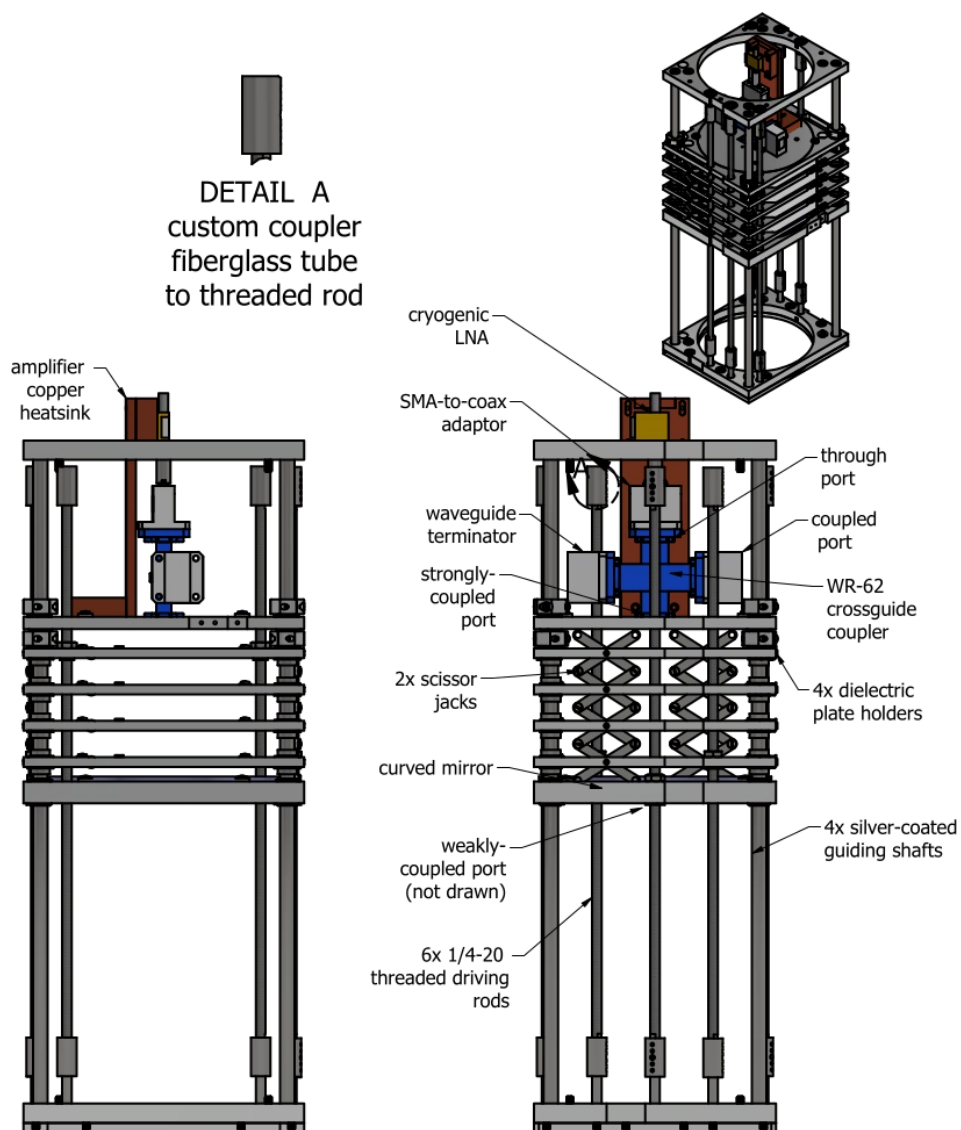


Figure 4.1: CAD model for the Orpheus cavity.

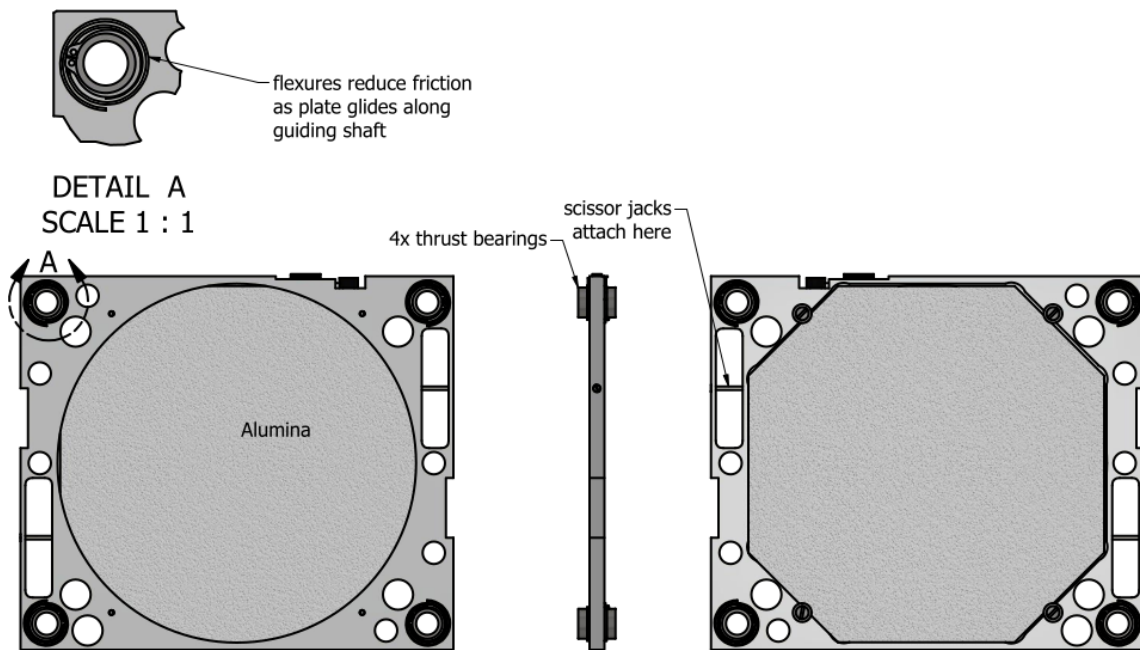


Figure 4.2: CAD model for one of the dielectric plates.

tolerances were tight at room temperature, the aluminum would contract faster and crush the alumina⁴. Nothing is clamping down the alumina because the clamping mechanism would contract more quickly than the alumina and crush it. The mirrors are made out of aluminum, so there is no differential expansion between the mirror holders. Tighter tolerances are viable and even desired to keep the coupling holes in the axis of the cavity.

The aluminum holders containing the curved mirror, top dielectric plate, and bottom dielectric plate are moved vertically by a pair of 1/4-20" stainless steel threaded rods. This results in a tuning mechanism with three degrees of freedom. The middle two dielectric plates are constrained to be spaced evenly between the top and bottom dielectric plate because they are attached to a pair of scissor jacks. The scissor jacks are connected to all dielectric plates, but not the mirrors. All the plates slide along four guide rails. To compensate for any possible misalignment that would cause the plates to jam, each aluminum holder has spiral

⁴Having this much slop isn't a viable option if the dielectrics were curved and needed to be centered on the axis of the cavity.

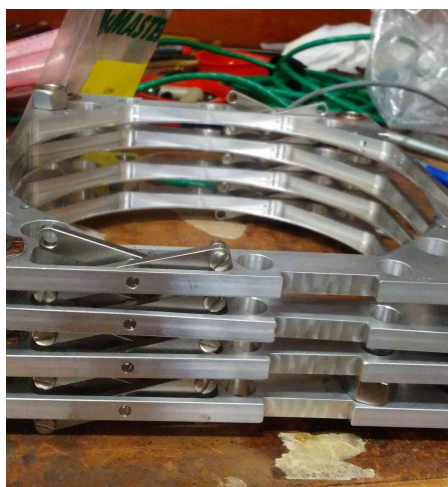


Figure 4.3: The dielectric holder assembly without the dielectrics. There is a lip inside the octagon where the dielectrics freely rest. Note that the thrust bearings in this photo are 0.25 in long are shorter than the final bearing length of 0.615 in.

cut flexures around the bearing surface.

Cavity materials were chosen to accommodate thermal contractions. Alignment is maintained after cooldown because all vertical structures are made from stainless steel, and all horizontal structures are made from aluminum. Bearings are made out of stainless steel so that the fitting tolerance between the shaft and bearing remains the same after cooldown⁵. However, having stainless steel thrust bearings rub against the stainless steel guide rails would lead to galling⁶. All bearing surfaces are coated with silver to prevent galling and reduce friction.

Without further measures, the friction between the bearing and the guide rail would cause the plates to tilt as the cavity tunes. Two measures are taken to mitigate this tilting. First, the thrust bearings (0.615 in) are much longer than the dielectric plate thickness (0.25 in).

⁵It would be bad if you had a stainless steel shaft and you made your sleeve bearing out of aluminum. Even if your sleeve bearing slides freely along your shaft at room temperature, the bearing will shrink more than the stainless steel after cooldown. Now your aluminum bearing is gripping the stainless steel shaft, and it's completely jammed.

⁶In general, having similar materials slide against each other causes both sliding surfaces to scratch each other, leading to galling and possibly cold welding



Figure 4.4: Orpheus cavity attached to the cryogenic insert.

The bearing length gives the plates less leeway to tilt. Second, a weak compression spring is placed around the guide rails in between the dielectric plates and mirror plates. For this experiment, LP 024L 03 S316 from Lee Spring was used. It has a spring constant of 0.37 lb/in. If the springs are too stiff, then they deform the spiral flexures and cause large position errors.

The real cavity is shown in Figure 4.4. The tuning mechanism performed to satisfaction in the tabletop setup. Measurements with a circular bubble level showed that plates tuned with negligible tilting. The position error was formally studied with a cavity that didn't have the springs. The measurements showed that the position error was often around 0.2 mm and rarely exceeded 0.4 mm. Simulations in Section 3.2.3 show that V_{eff} is tolerant to this magnitude of position error.

Overall, the cavity tunes accurately and maintains good alignment while tuning.

4.2 *Insert*

The cavity is the core of the experiment, but the cavity needs to be cooled down to 4 K under vacuum. The cavity is kept in a vacuum instead of directly under cryogen because the cavity

isn't designed to remain mechanically stable under a boiling fluid⁷. The cavity is to remain cold for about a week to allow enough time to scan the tuning range. Room temperature stepper motors tune the cavity, and the stepper motors need to remain far away from the magnet. The insert in Figure 4.5 and Figure 4.6 were designed to meet these requirements. I will describe the main components from the top down.

1. Motor stage shown in Figure 4.7. Each stepper motor (Applied Motion Products STM23S-2EE) drives a pair of miter gears that transmit power to the vertical shafts that are connected to the vacuum rotary feedthroughs on the top flange.
2. Top flange that sits on top of the dewar (Figure 4.8). Has many of the vacuum and cryogenic ports listed below.
 - Six MDC Precision 670000 rotary motion feedthroughs.
 - AMI vapor-cooled magnet leads with superconducting busbars. They have a 8 L/day boil off.
 - Kurt-Lesker IFDRG197018B 19 pin electrical feedthrough.
 - Isocross for vacuum port, 51 pin micro type-D port (MDC 9163004), and RF coaxial port.
 - cryogen fill port
 - cryogen vent port
3. 300 L Dewar with a boil 18 L/day boil off.
4. 50" long, 4" OD stainless steel tube.
5. Six 50" long, 0.75" ID stainless steel tubes to house motor shafts in vacuum space.

⁷The cavity modes were not stable when I tried to dunk it in liquid nitrogen. I've since been told that liquid helium can be very quiet. So that might be a viable option, but was not pursued in this experiment.

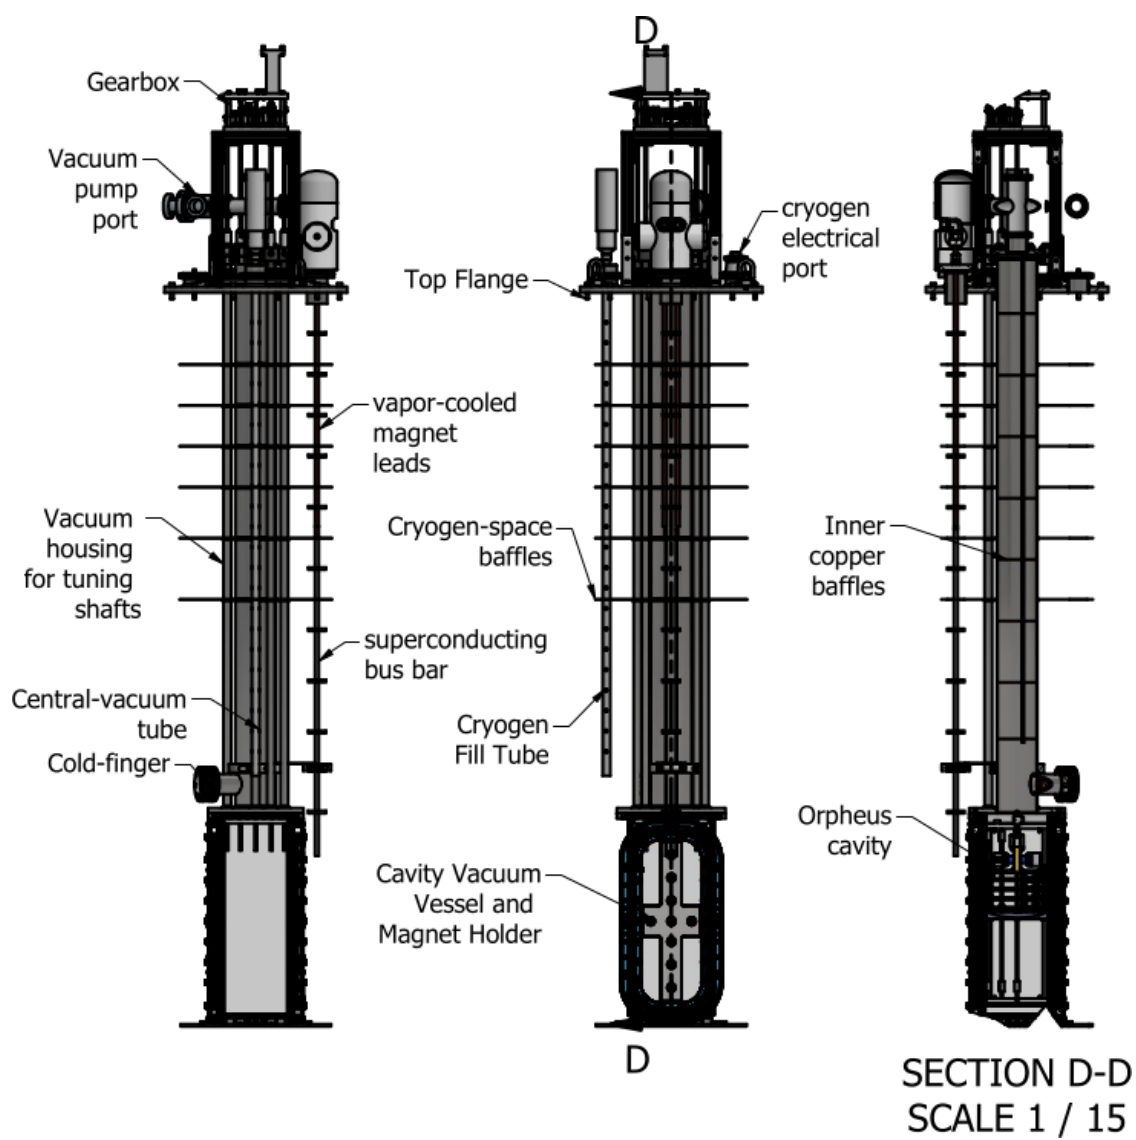


Figure 4.5: Orpheus cryogenic insert.



Figure 4.6: Cryogenic insert in real life.

6. Six motor shafts. Each motor shaft consists of a stainless steel tube coupled to a G10 fiberglass tube. The fiberglass tube reduces thermal contact to the top flange. The stainless steel tube is not as thermally isolating as the fiberglass but has less elasticity and would lead to less mechanical backlash.
7. Copper plug to act as cold finger. Used to thermally sink cryogen space to cavity top.
8. Rectangular experimental support plate welded to the bottom of the tube. The cavity attaches to the experimental support plate.
9. Cavity vacuum vessel. The dipole magnet attaches to this vessel.

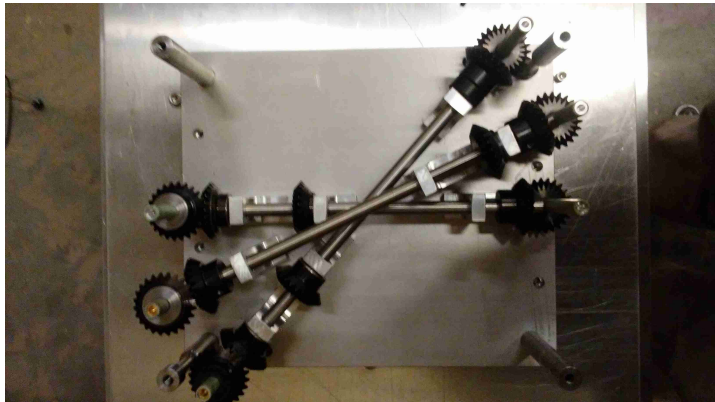


Figure 4.7: Orpheus gearbox

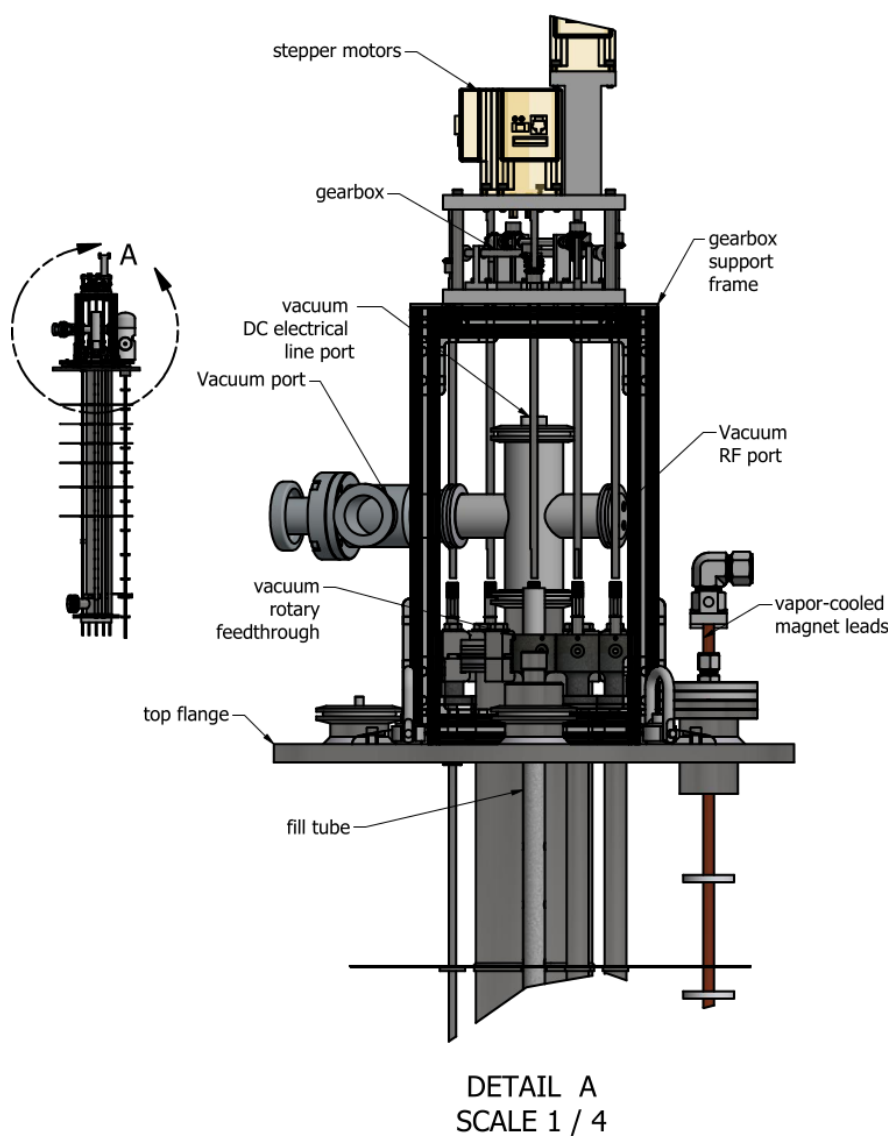


Figure 4.8: Insert top flange with all the vacuum ports, cryogen ports, and motor stage.

4.3 *Issues with the Tuning Mechanism*

There were several shortcomings of this tuning mechanism, as implemented on Sept. 2021. The most severe is in the friction caused by shaft misalignments. One of the most prominent sources of misalignment is in the vacuum port where the vacuum rotary feedthroughs attach (Figure 4.9). The vacuum ports were unfortunately welded with a lot of tilt. This meant that the gearbox shafts were very misaligned from the rotary feedthroughs. Double universal joints were used to compensate for the misalignment. However, these double u-joints would apply a radial force on the gearbox shaft during rotation, which caused the miter gear to move closer or farther away from the mating gear. The changing mounting distance caused the mating gears to mesh poorly, leading to friction, inefficient power transfer, and motor stalls. A flexible coupler with a spring joint may have mitigated this issue since it wouldn't have provided as much radial force, but vendors were out of stock at the time.

The vacuum port misalignment had propagating effects into the vacuum space. The tilted vacuum rotary feedthrough caused the vacuum shafts to also be tilted. The vacuum shafts consist of a 1/4" OD stainless steel tube connected to a 1/4" OD fiberglass tube connected by a 5/8" OD spiral flexible shaft coupler. The stainless steel vacuum housing tube has a 3/4" inner diameter, so the misalignment caused the flexible coupler to rub against the inner walls of the stainless steel tube. This touch might have been especially problematic during cooldown if residual moisture was frozen onto the coupler. The touch was evident when the insert was pulled out of the dewar, as seen by the frost marks were localized around the location of the flexible coupler (Figure 4.9).

The effects of the misalignment can propagate even further. If the vacuum shafts are misaligned, some shafts may need to be longer than others to reach the resonator. Because the vacuum shafts may have different lengths, they would have contracted by different amounts during cooldown. That could lead to the curved mirror and dielectric plate becoming tilted after cooldown, and the tilted plates would experience more friction.

Besides the myriad misalignment issues, there is no mechanism that mitigates backlash.

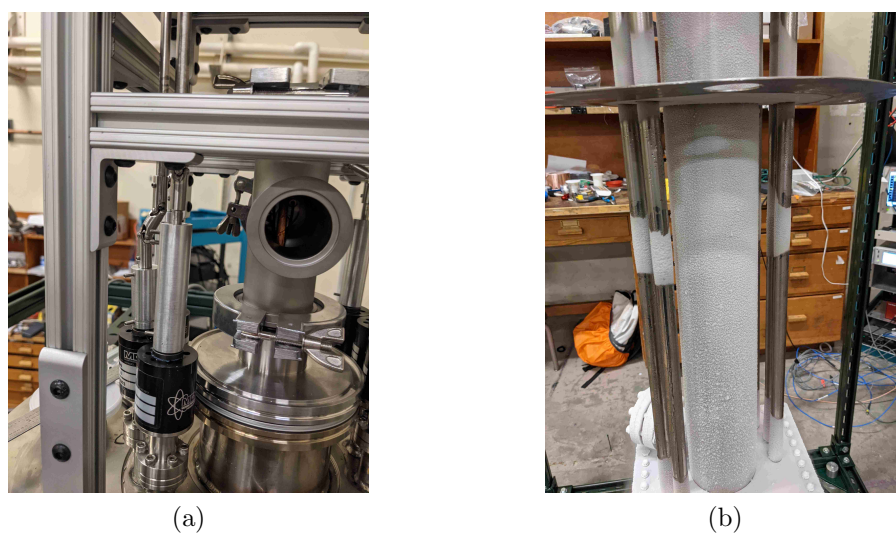


Figure 4.9: (a) The vacuum ports for the rotary feedthroughs are misaligned. This contributed to a cascade of effects that caused the motors to stall frequently after the cooldown. (b) The vacuum shaft consists of a stainless steel tube connected to a G10 fiberglass tube connected by a spiral shaft coupling with a $5/8$ " outer diameter. The vacuum housing tube has a $3/4$ " inner diameter, and the misalignment from the top flange caused the shaft coupler to touch the vacuum walls. The point of contact is made obvious by the localized frost marks.

The gear ratios in the gearbox are all one-to-one. The only increase in gear ratio is in the transfer of power to the threaded rods, and the resulting gear ratio is about 16. The backlash is approximately a quarter of a turn. This amounts to sub-millimeter systematic uncertainty in the absolute length of the cavity and relative positions of the dielectric plates. This systematic uncertainty caused by the backlash, while undesirable, can be mitigated by always tuning continuously in one direction.

4.4 Thermal Gradient Issues

There are also multiple issues with the thermal design. The resonator is cooled because the flat mirror and copper amplifier heatsink is connected to a cold finger. The mirror is then supposed to cool down the rest of the cavity. However, the flat mirror can lose thermal contact with the rest of the cavity, and this thermal connection is lost. The contact is lost if the flat mirror is lifted by the contracting copper braids or copper RF coaxial cables.

The flat mirror also has a higher heat leak than the rest of the cavity because it is connected to two coaxial cables that connect straight to the room-temperature top flange. In contrast, the curved mirror at the bottom only has one coaxial cable attached. The flat mirror also has a higher heat load since it is cooling down the cryogenic amplifier.

The dielectric plates and mirrors are also somewhat thermally isolated from each other as only the stainless steel provides thermal connections.

The combination of these effects causes the temperature of the flat mirror to diverge significantly from the curved mirror during the data run, as will be shown in Figure 5.7. This divergence makes it harder to characterize a cavity temperature and consequently makes the extracted noise temperature have more uncertainty.

Chapter 5

ELECTRONICS, DATA ACQUISITION SYSTEM, AND OPERATIONS

This chapter focuses on how measurements in the Orpheus experiment are taken. The chapter talks about the radio frequency (RF) and intermediate frequency (IF) electronics, the steps taken for each data-taking cycle, and the software that controls and monitors the experiment. The chapter ends with a narrative of the real-time operation of the first science run.

5.1 *Cold Electronics*

The diagram for the cold electronics is shown in Figure 5.1. The cavity has a weakly-coupled port and a strongly-coupled port. The strongly-coupled port is connected to a WR-62 20 dB crossguide coupler (PEWCP1047). The crossguide coupler is attached to a waveguide to coax adapter (PE9803). The coax adaptor connects directly to the cryogenic low noise amplifier (LNF-LNC6_20C). The cryogenic amplifier output is connected to an RG405 coaxial cable that connects directly to the room-temperature SMA bulkhead¹.

During a science run, the strongly-coupled port transmission and reflection coefficients need to be measured for each data-taking cadence. For a transmission measurement, the VNA port 1 is connected to the weakly-coupled port. For a reflection measurement, the VNA port 1 is connected to the crossguide coupler coupled port, as shown in Figure 5.1. The signal then travels to the strongly-coupled port, gets reflected, and reaches the input of the cryogenic amplifier. The coupler² is needed because the VNA source signal needs to

¹This direct connection causes a large heat leak but was easy to implement. Future runs will take steps to mitigate this heat leak.

²A broadwall coupler would have been better than a crossguide coupler because it has a coupling value

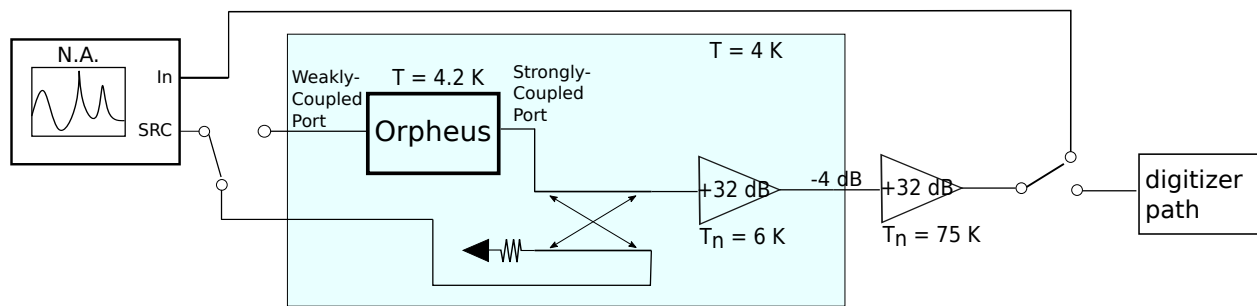


Figure 5.1: A diagram of how transmission, reflection, and power measurements are taken. Room temperature Teledyne switches are used to switch between transmission and reflection measurements, and from VNA measurements to power measurements.

bypass the cryogenic amplifier to reach the strongly-coupled port. The remaining port in the crossguide coupler is terminated with a waveguide terminator (PE6804³). A Teledyne SMA switch is used to switch between a reflection measurement and transmission measurement.

5.2 Warm Electronics

The schematic for warm electronics is shown in Figure 5.2, a photograph is shown in Figure 5.3, and the list of components with the associated cascade analysis is shown in Table 5.1.

After the signal makes it out of the vacuum space, it is directed to a room-temperature low noise amplifier (LNF-LNR6.20A SV) to ensure that the gain is high enough so the noise from the rest of the noisy high-power electronics doesn't reduce the SNR. The signal is amplified again by the ZX60-183-S. The 18 GHz signal needs to be downsampled so that it can be digitized by a 180 MSPS digitizer card. Downsampling is achieved by sending the RF signal to an Image-Reject Mixer (IRM140180B). The mixer uses the non-linear semiconductor properties to multiply the RF signal with angular frequency ω_{RF} by

frequency response is much flatter, and its directivity is far higher. Unfortunately, a broadwall coupler doesn't fit inside the 4 in insert tube. A circulator may have also been more attractive. But they are far more expensive and require magnetic shielding. Also, I didn't find an appropriate cryogenic 18 GHz circulator from Pasternack or QuinStar.

³PEWTR1004 has a much better VSWR. Unfortunately, there was no room for it inside the vacuum vessel.

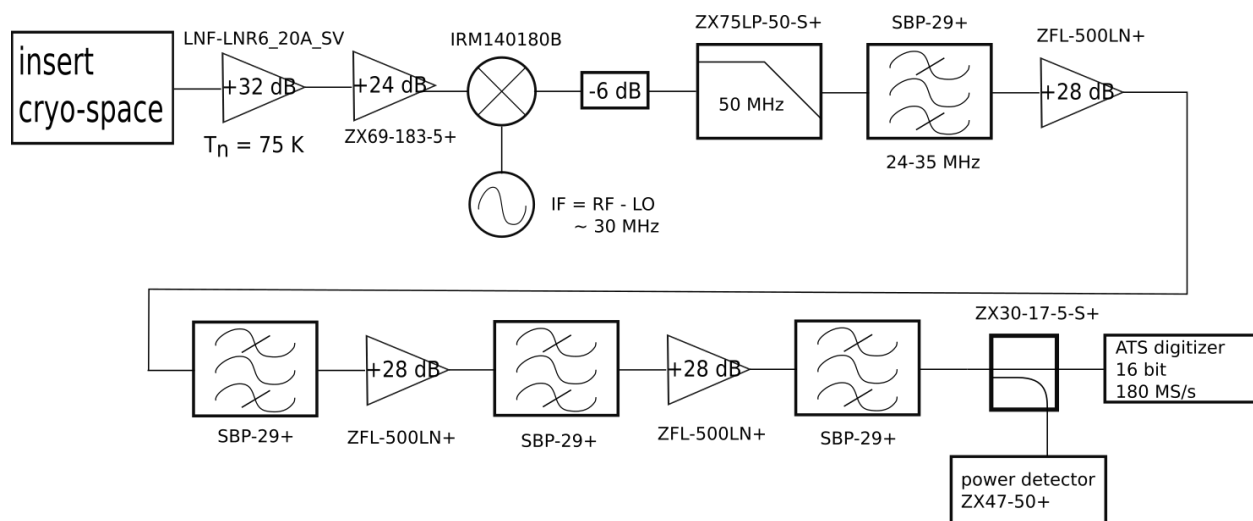


Figure 5.2: Orpheus measurement scheme. There are 3 paths: transmission, reflection, and digitization.

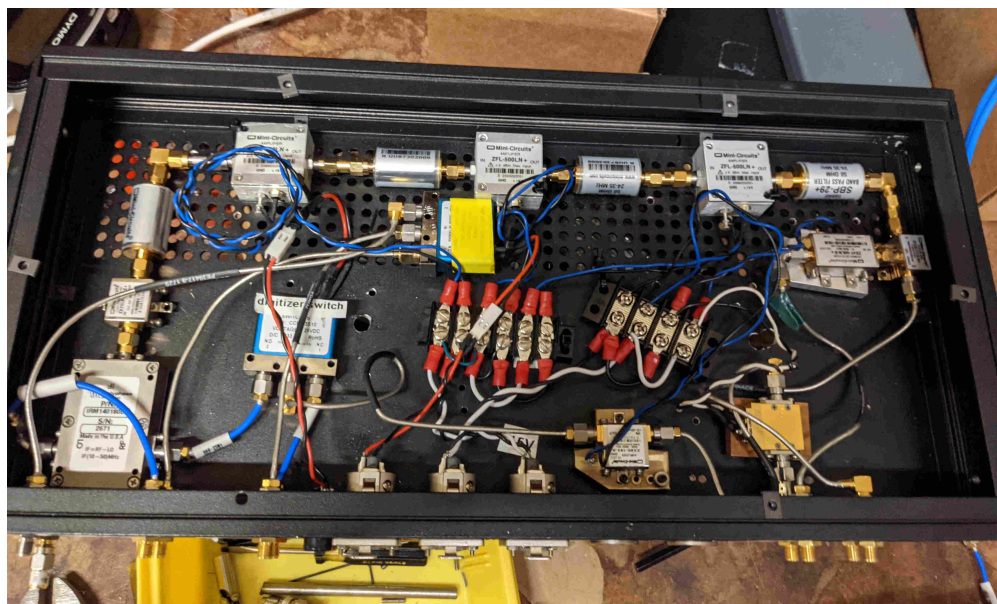


Figure 5.3: RF/IF electronics box. Contains warm RF stage, IF stage, and coaxial switches to switch between a VNA transmission measurement, VNA reflection measurement, and a power spectrum measurement.

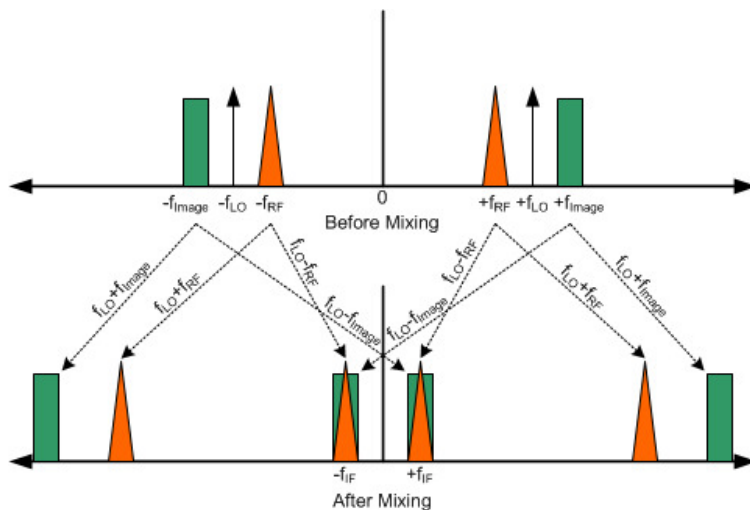


Figure 5.4: A demonstration of image noise after mixing. The negative image frequency folds into the positive IF frequency, and the positive image frequency folds into the negative IF frequency. Figure from [46].

a local oscillator signal with angular frequency ω_{LO} . The result is a signal with spectral components with both the sum and difference of the two frequencies (and higher-order harmonics), i.e. $\sin(\omega_{RF}t) \sin(\omega_{LO}t) \propto \sin((\omega_{RF} - \omega_{LO})t) + \sin((\omega_{RF} + \omega_{LO})t)$. The signal component with $\omega = \omega_{RF} + \omega_{LO}$ is removed using a low pass filter, and the signal component with $\omega = \omega_{RF} - \omega_{LO}$ is digitized. The downsampled frequency $2\pi(\omega_{RF} - \omega_{LO})$ is known as the Intermediate Frequency (IF). The IF is chosen to be about 30 MHz to be compatible with the IRM140180B.

Image rejection is important because otherwise, the noise would unnecessarily double. To see why this is true, one can examine Figure 5.4. One can keep track of all the combinations of beat frequencies to see that the negative image frequency folds into the positive IF frequency and the positive image frequency folds into the negative IF frequency. The image rejection works using phasing techniques, which can be read about in Edward Daw's thesis[12].

After the RF signal is multiplied by the local oscillator signal, the resulting signal passes through a low pass filter, and only the IF frequency survives. There is an attenuator between

the mixer and low pass filter to eliminate standing waves between the two⁴. After this, the IF signal passes through a series of IF amplifiers and bandpass filters. The IF amplifiers amplify the IF signal to the dynamic range of the digitizer card, and the bandpass filters attenuate noise outside the frequency of interest. The amplified IF signal then goes to a directional coupler. The coupled port is connected to a power detector (ZX47-50+), allowing for the real-time measurement of the IF signal power. The directional coupler's through port is connected to the Alazartech digitizer card (ATS9462).

The Alazartech digitizer is set to a sampling rate of 125 MSPS. This rate is chosen so that the Nyquist frequency is well above the IF band. Each subspectrum contained 50,000 time samples, resulting in a 2.5 kHz bin width. Depending on the run settings, each spectrum is either the average of 75,000 (30 seconds) subspectra or 250,000 (100 seconds) subspectra. No windowing is applied to each subspectrum. It is a 16-bit digitizer with a dynamic range of ± 200 mV to ± 16 V

The cascade analysis in Table 5.1 shows that the gain from the RF and IF electronics is about 140 dB. For $T_n \approx 11$ K and bin width $b = 2.5$ kHz, the noise power without the system gain is $P_n = k_B T_n b = -118$ dBm. Thus the noise power seen by the digitizer is $P_n = 21.8$ dB = 8 V_{p-p}. That's well within the dynamic range of the digitizer.

⁴From Christian Boutan: Mixers are sometimes temperamental with impedance matching and reflections off the mixer can actually cause filters to detune. The fix is having a few dB of attenuation between mixers and filters to kill standing waves.

| Part Description | vendor | part number | Relative power gain | absolute gain | device noise temperature | noise | cascaded noise temperature |
|----------------------------------------------------------|---------------------|----------------|---------------------|---------------|--------------------------|-------|----------------------------|
| LNF-LNC6-20C s/n 1556Z | Low Noise Factory | | 33 | 33 | 6 | | 10.5 |
| SMA Male to SMA Male Cable Using RG405 Coax, RoHS | Pasternack | PE3818LF-72 | -7.6 | 25.4 | 0 | | 10.5 |
| C3146 SMA Hermetic Bulkhead Adapter 18GHz | Centric RF | C3146 | -0.2 | 25.2 | 0 | | 10.5 |
| Low Loss Test Cable 12 Inch Length, PE-P142LL Coax | Pasternack | PE341-12 | -0.82 | 24.38 | 0 | | 10.5 |
| Hand-Flex Interconnect, 0.086" center diameter, 18.0 GHz | Minicircuits | 086-SBSMR+ | -0.3 | 24.08 | 0 | | 10.5 |
| LNF-LNR6_20A_SV s/n 1257Z | Low Noise Factory | | 32 | 56.08 | 100 | | 10.8908408957924 |
| Right Angle Semi-Flexible Cable | Pasternack | PE39417-6 | -2 | 54.08 | | | |
| Wideband Microwave Amplifier 6 to 18 GHz | Minicircuits | ZX60-183-S+ | 24 | 78.08 | 1378.8 | | 10.8942410708306 |
| Semi-Flexible Cable | Pasternack | PE39417-9 | -3 | 75.08 | | | |
| Hand-Flex Interconnect, 0.086" center diameter, 18.0 GHz | Minicircuits | 086-4SM+ | -0.38 | 74.7 | 0 | | 10.8942410708306 |
| Hand-Flex Interconnect, 0.086" center diameter, 18.0 GHz | Minicircuits | 086-4SM+ | -0.38 | 74.32 | 0 | | 10.8942410708306 |
| DC-18 GHz/DC-22 GHz SPDT Coaxial Switch | Teledyne | CCR-33S/CR-33S | -0.4 | 73.92 | 0 | | 10.8942410708306 |
| Hand-Flex Interconnect, 0.086" center diameter, 18.0 GHz | Minicircuits | 086-2SM+ | -0.32 | 73.6 | 0 | | 10.8942410708306 |
| IMAGE-REJECT MIXER 14.0 - 18.0 GHz | Polyphase Microwave | IRM140180B | -8.5 | 65.1 | 3360.9 | | 10.8943877794366 |
| 3 dB Fixed Attenuator | Pasternack | PE7005-3 | -3 | 62.1 | 0 | | 10.8943877794366 |
| Low Pass Filter | Minicircuits | ZX75LP-50-S+ | -1.38 | 60.72 | 0 | | 10.8943877794366 |
| Hand-Flex Interconnect, 0.086" center diameter, 18.0 GHz | Minicircuits | 086-2SM+ | -0.32 | 60.4 | 0 | | 10.8943877794366 |
| Lumped LC Band Pass Filter, 24 - 35 MHz, 50Ω | Minicircuits | SBP-29+ | -0.88 | 59.52 | 0 | | 10.8943877794366 |
| Low Noise Amplifier, 0.1 - 500 MHz, 50Ω | Minicircuits | ZFL-500LN+ | 28.21 | 87.73 | 275 | | 10.8946949168298 |
| Lumped LC Band Pass Filter, 24 - 35 MHz, 50Ω | Minicircuits | SBP-29+ | -0.88 | 86.85 | 0 | | 10.8946949168298 |
| Low Noise Amplifier, 0.1 - 500 MHz, 50Ω | Minicircuits | ZFL-500LN+ | 28.21 | 115.06 | 275 | | 10.8946954848093 |
| Lumped LC Band Pass Filter, 24 - 35 MHz, 50Ω | Minicircuits | SBP-29+ | -0.88 | 114.18 | 0 | | 10.8946954848093 |
| Low Noise Amplifier, 0.1 - 500 MHz, 50Ω | Minicircuits | ZFL-500LN+ | 28.21 | 142.39 | 275 | | 10.8946954858597 |
| Lumped LC Band Pass Filter, 24 - 35 MHz, 50Ω | Minicircuits | SBP-29+ | -0.88 | 141.51 | 0 | | 10.8946954858597 |
| 17.5 dB Directional Coupler, 5 - 2000 MHz, 50Ω | Minicircuits | ZX30-17-5-S+ | -0.7 | 140.81 | 0 | | 10.8946954858597 |

Table 5.1: List of electronics and cascade analysis.

5.3 System Noise Temperature

The system noise power can be described as the sum of the blackbody radiation of the cavity and the added Johnson noise of each electrical component, and is written as $P_n = Gk_b b T_{sys}$, where G is the system gain that depends on both frequency and physical temperature, b is the frequency bandwidth, and T_{sys} is the system noise temperature referenced to the cavity.

From examining Figure 3.4, one can determine that

$$T_{sys} = T_{cav}(1 - |\Gamma|^2) + T_{amp,input}|\Gamma|^2 + T_{rec} \quad (5.1)$$

where T_{cav} is the physical temperature of the cav, $T_{amp,input}$ is the noise temperature coming from the input of the cryogenic amplifier, T_{rec} is the noise temperature of the receiver chain from the output of the cryogenic amplifier outward (discussed in Section 5.4), and Γ is the reflection coefficient and $|\Gamma|^2$ is the fraction of power reflected. As a reminder of Equation 3.15, $|\Gamma|^2$ depends on both the cavity coupling coefficient β and the detuning factor Δ .

To understand why Equation 5.1 is plausible, consider several limiting cases. When the receiver is critically coupled to the cavity, the cavity on resonance looks like a blackbody.

Thermal photons are in thermal equilibrium with the cavity, and $T_{sys} = T_{cav} + T_{rec}$. If the receiver is poorly coupled or if the RF frequency is far off resonance, the cavity looks more like a mirror. In this scenario, thermal photons are emitted from the input of the amplifier. These thermal photons reach the flat mirror and are reflected back into the amplifier. Thus, $T_{sys} = T_{amp,input} + T_{rec}$. Generally, the cavity is only partially reflecting, and the noise temperature is described by Equation 5.1.

5.4 Receiver Noise Temperature

The noise temperature of the receiver is derived from datasheets from Low Noise Factory and the Friis cascade equation [47]

$$T_{rec} = T_1 + \frac{T_2}{G_1} + \frac{T_3}{G_1 G_2} + \dots \quad (5.2)$$

T_1 and G_1 are the noise temperature and gain of the 1st stage cryogenic amplifier. T_2 and G_2 are the noise temperature and gain of the 2nd stage room-temperature LNA. T_3 is the noise temperature of the 3rd stage amplifier. The 3rd term is negligible because of the gain from the first two amplifiers. The derived receiver noise temperature is shown in Figure 5.5.

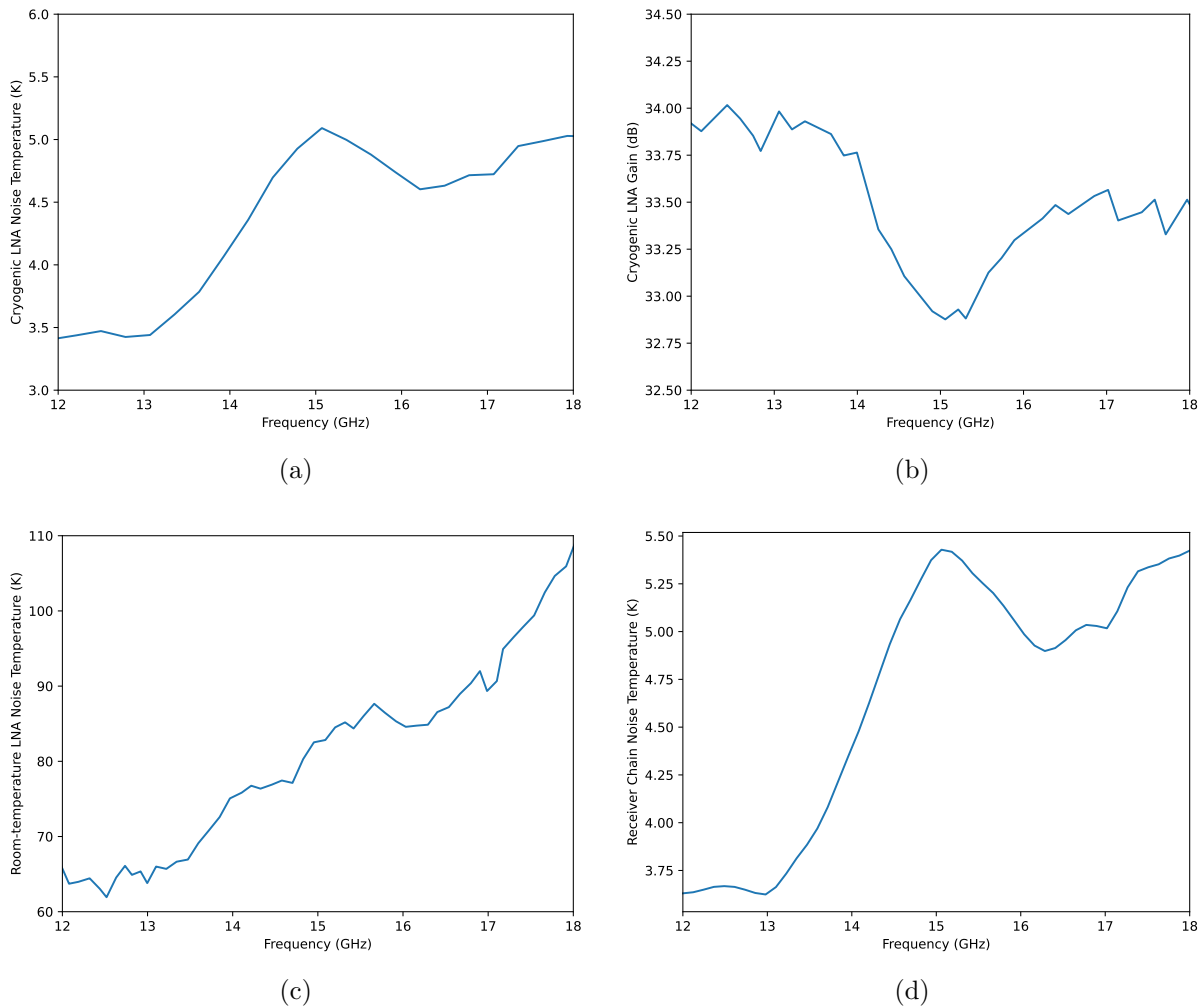


Figure 5.5: The derived receiver noise temperature. The Friis equation is used to derive the receiver noise temperature (d) from the noise temperature (a) and gain (b) of the 1st stage cryogenic LNA and the noise temperature of the 2nd stage room-temperature amplifier (c). These values are obtained from data sheets from Low Noise Factory. The cavity temperature can be added to the receiver noise temperature to obtain the system noise temperature, as referenced to the cavity.

5.5 Software Stack

The control software stack consists of Python, Postgresql, RabbitMQ, Dripline, Docker, Kubernetes, Helm, and Grafana. Except for the digitizer driver, all software is open source.

The software consists of modular, self-healing, loosely-coupled services with a standardized messaging protocol for all serial communication between devices. The software stack consists of:

- **Dripline:** The standardized messaging protocol developed by Project 8 and maintained by Pacific Northwest National Laboratory. Dripline allows for a standardized way to talk to different devices running on different servers, as long as they accept socket commands. One can use the same messaging protocol to talk to stepper motors, digitizer cards, and databases. A fundamental design philosophy of Dripline is RESTfulness and that software services should run without making assumptions of what other services are doing. This philosophy allows for a continuously evolving DAQ system. More information can be found on their website [48]. A schematic of the Dripline architecture is shown in Figure 5.6.
- **Python:** The scripting language of choice. Python scripts control the data-taking cadences and online data analysis.
- **RabbitMQ:** A messaging broker. It's like the post office for Dripline messages.
- **Postgresql:** An object-relational database.
- **Docker containers:** Provides OS-level virtualization. Packages an application and its dependencies in a virtual container.
- **Kubernetes:** Manages the lifecycle of Docker containers. For example, Kubernetes can be configured to always have a certain Docker container at all times. If that container

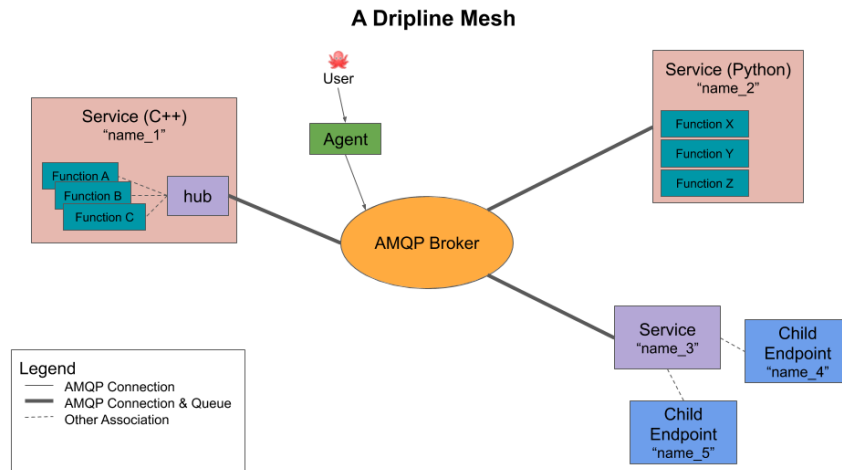


Figure 5.6: Dripline software architecture. Figure from [48].

crashes, Kubernetes will detect that crash and create another container to replace the crashed container.

- Helm: A package manager for Kubernetes configuration files.
- Grafana: Real-time data visualization. Integrates really well with Postgresql.

Since Orpheus took data only for a few days, the overhead of implementing this software stack was likely not worth it. But Dripline is used for the ADMX flagship experiment, where data is collected continuously for six months, and a fault-tolerant DAQ system is required. The flagship experiment uses Dripline v2. Orpheus is the first implementation of Dripline v3 for science data-taking. Orpheus uses much of the same equipment as the flagship experiment, so implementing the newly developed Dripline v3 allowed for stress testing, debugging, and development of new services in a smaller test stand.

5.6 Data-taking cadence

The basic strategy is to tune the cavity to scan for dark photons with different masses. For each cavity length, a series of ancillary measurements are taken to extract a noise power

calibration and expected dark photon signal. The power is then measured out of the cavity in search of non-statistical peaks that may correspond to a dark matter signal. The data-taking cadence is as follows:

1. The state of the system is recorded. The temperature sensors and motor encoders are logged into the database. These measurements allow the extraction of the noise power, cavity length, and dielectric plate positions.
2. The transmission and reflection coefficient is measured 20 MHz around the TEM₀₀₋₁₈ mode. Parameters f_0 , Q_L and β are extracted from these measurements. The extracted parameters are logged to the database.
3. The VNA output is disabled, and the power spectrum is measured. The power spectrum is integrated for either 30 s or 100 s.
4. The motors are tuned in a coordinated way. This means the curved mirror is moved by an amount specified in a configuration file. Simultaneously, the top dielectric plate is moved by 1/5 the number of steps of the curved mirror, and the bottom dielectric plate is moved by 4/5 the number of steps as the curved mirror.
5. Repeat data-taking cadence for new cavity length.
6. Every 20 data-taking cycles, the transmission coefficient is measured with a frequency range of 15 GHz to 18 GHz. This is known as a wide scan.

5.7 September 1st Commissioning Narrative

I will end this chapter with a real-time narrative of the September science run, as it may explain the quirks in the data.

On 9/1/2021, I tuned the cavity so that the TEM₀₀₋₁₈ mode was at 16.15 GHz. I knew this to be a frequency region with good Q and coupling. We had 350 L of liquid helium at our

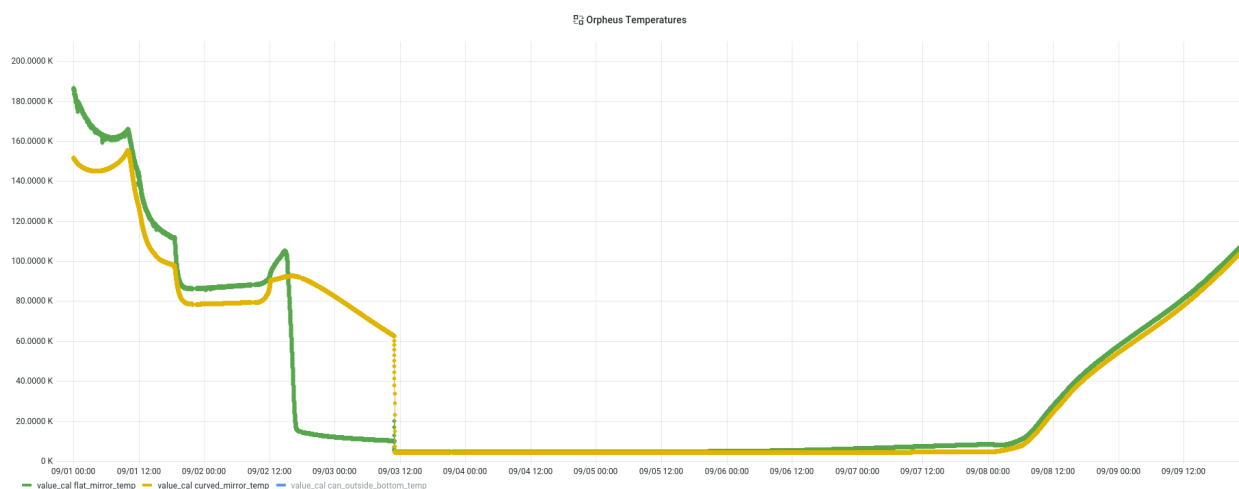


Figure 5.7: Flat mirror and curved mirror temperatures throughout the entire commissioning of the experiment, including warmup and cooldown.

disposal, and I wanted to make the best use of it. Charles, Grant, and I began pre-cooling Orpheus with liquid nitrogen in the reservoir space and introduced a nitrogen exchange gas in the vacuum space. The cooling and temperature history is shown in Figure 5.7.

On 9/2/2021, we boiled off the LN₂ and pumped out the nitrogen exchange gas from the vacuum space. This was so that we didn't waste LHe freezing LN₂. Then we filled the reservoir with liquid helium.

By the morning of 9/3/2021, the top mirror cooled down to 10 K, but the curved mirror was cooling very slowly and had only reached about 60 K. Furthermore, most of the liquid helium in the reservoir had boiled off, and we didn't have any liquid helium left to refill the reservoir. It seemed like the bottom of the cavity was thermally detached from the flat mirror, and we couldn't wait for the curved mirror to cool to 4 K. Around noon, we introduced helium exchange gas into the vacuum space, and things cooled to below 5 K within minutes. Once the cavity was cooled, we pumped out some of the exchange gas. No one knows how much exchange gas we introduced or how much we pumped out, but the procedure worked well.

Once the cavity was cold, I tried tuning, but the top dielectric plate motor was stalled.

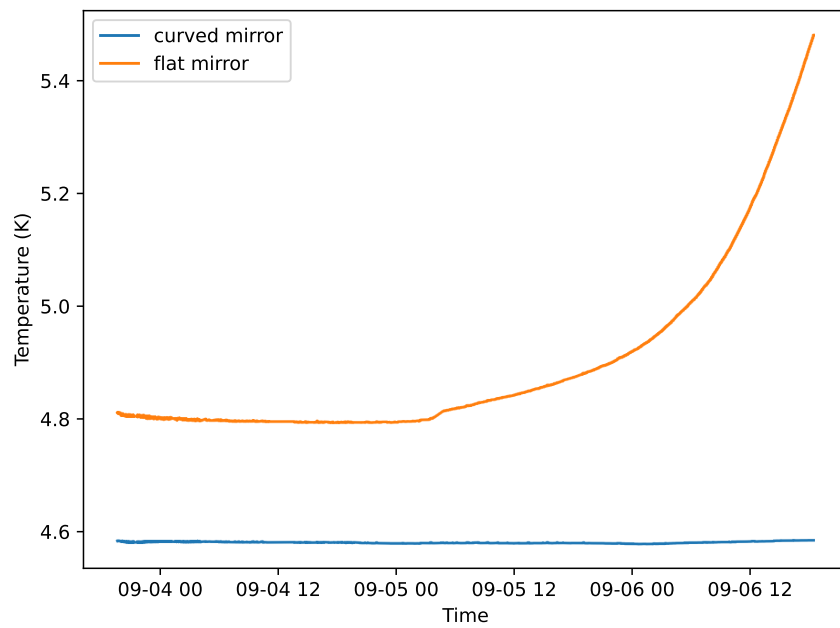


Figure 5.8: Flat mirror and curved mirror temperatures throughout the data-taking run. Physics data-taking stopped because the flat mirror temperature was deviating significantly from the curved mirror temperature, causing a significant increase in the uncertainty in the cavity temperature.

I found that with great force, I could tune the top dielectric plate by hand. So I decided to skip the top dielectric plate tuning during a data-taking cadence and to tune by hand every few hours to correct for the top dielectric plate lagging behind. The top dielectric plate needed to be tuned with $1/5$ of the steps of the curved mirror. I was able to tune the TEM_{00-18} mode without tuning the top plate, so I decided that this top plate position error would just amount to a perturbation in V_{eff} . Luckily, the motor encoders kept count of the steps, even if the motor was turned manually. The recorded motor positions are shown in Figure 5.10. The actual motor positions vs intended motor positions are shown in Figure 5.11.

I set the digitization time to 30 seconds and tuned downwards until I reached about 15.8 GHz (Figure 5.9). I stopped here because the Q and cavity coupling were decreasing. I decided to tune the cavity in the opposite direction so that $f_{00-18} \approx 16.8$ GHz. Unfortunately, all three motors stalled while tuning in the reverse direction. So I had to turn each motor by hand. Ironically, the Orpheus experiment couldn't look back without great effort. Once the cavity was set to 16.8 GHz, I programmed the cavity to tune downwards continuously.

At 9/4/21 4:20 AM, I decided to take a two-hour nap. Coincidentally, at 4:26 AM, the power supply that controlled the RF switches stopped working, and so the digitizer was not connected to the cavity while measuring power spectra. So the resulting data was useless. I went back to the lab at around 7:30 AM to manually reverse the motors so that I could rescan the affected frequency range.

At 12PM, I noticed that Orpheus was closer to becoming critically coupled. So I decided the cavity had returned to a more optimal frequency range, and I increased the digitization time to 100 s.

At 10PM, I noticed that the flat mirror's temperature was increasing while the curved mirror temperature stayed constant. I feared I was running out of liquid helium. Because I wanted to sleep and I was unsure how long the experiment could keep running, I decided to stop tuning. I would remain at 15.96 GHz until the end of the run. At the time, I didn't know the dark photon coupling sensitivity would be, so I decided that I could have world-

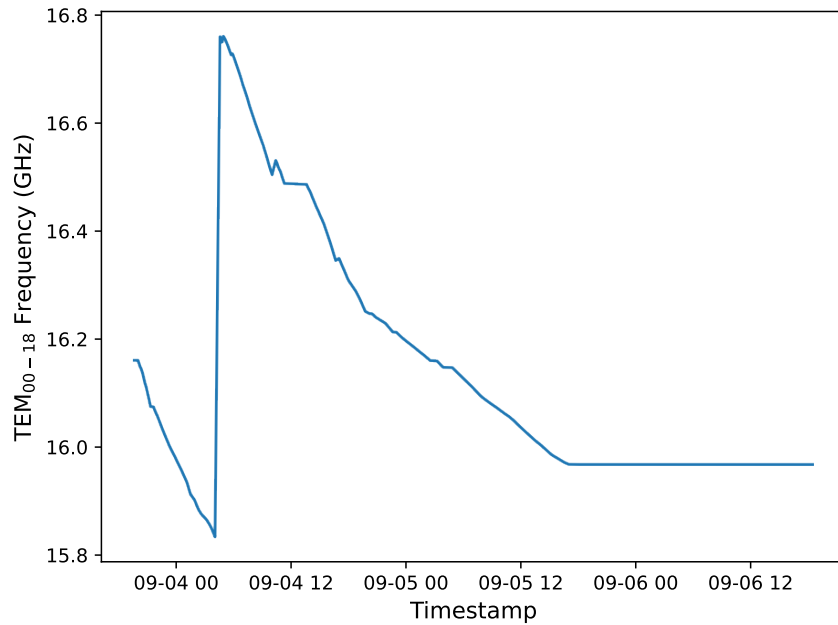


Figure 5.9: The TEM₀₀₋₁₈ resonant frequency vs time.

leading sensitivity for at least one frequency. This was a bad decision since dark photon sensitivity scales poorly with integration time ($\chi_{exc} \propto t^{-1/4}$). This extended integration time at 15.96 GHz is shown in the histogram of the number of IF bins corresponding to an RF bin, shown in Figure 5.12.

At 9/5/2021 6PM, I noticed that the flat mirror was 0.5 K hotter than the curved mirror. I suspected that it would be difficult to interpret a cavity temperature from this data. So that's why I ended the run.



Figure 5.10: The motor steps as recorded by the encoder. From the motor steps, one derives the distance traveled by the curved mirror, bottom dielectric plate, and top dielectric plate. The top dielectric plate motor stalled after cooldown, preventing the top dielectric plate from being tuned remotely and in a continuous manner. I needed to turn the motor feedthroughs by hand periodically (explaining the discrete tuning steps for the top dielectric plate).

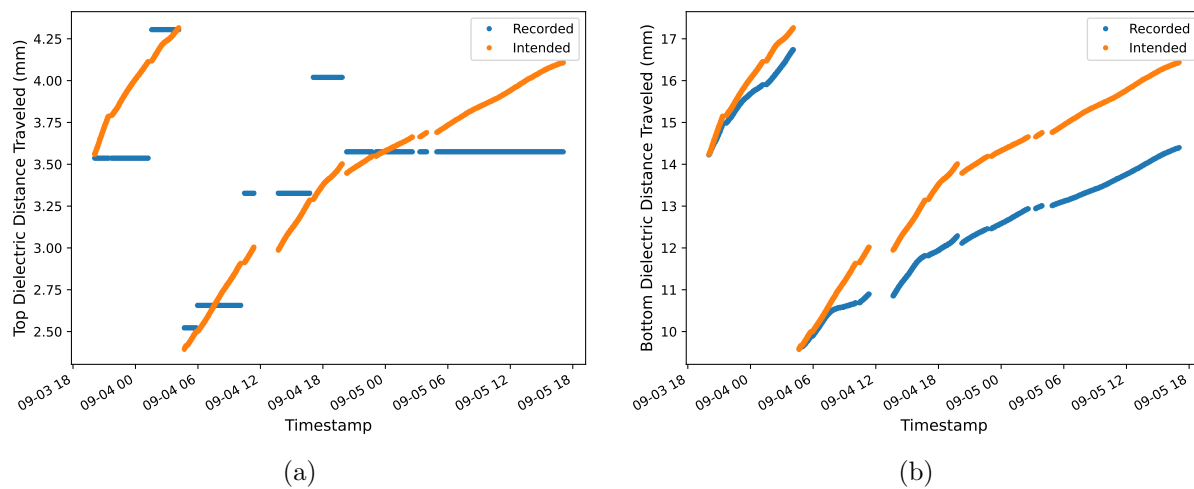


Figure 5.11: Due to a cascade of mechanical failures, software design flaws, and operation errors, the top dielectric plate and bottom dielectric plate deviated from their intended positions throughout the data-taking run.

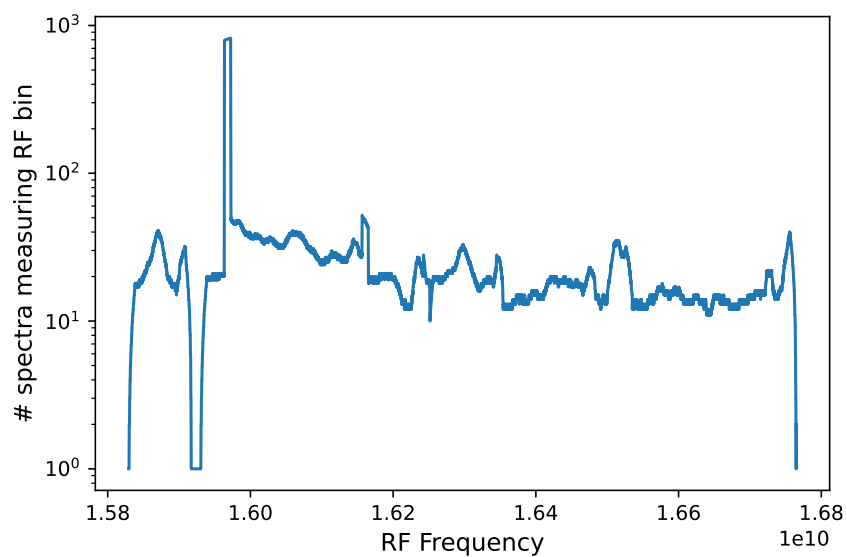


Figure 5.12: Number of spectrum containing a particular RF bin.

Chapter 6

DARK PHOTON SEARCH ANALYSIS

The data collected between 9/3/2021 and 9/7/2021 was used to search for dark photons between $65.5 \mu\text{eV}$ (15.8 GHz) and $69.5 \mu\text{eV}$ (16.8 GHz). No candidate for a dark photon was found, so a 90% confidence level exclusion limit was placed on the kinetic mixing angle χ in this mass range. The procedure for deriving the exclusion limits in this chapter follows the procedure developed by ADMX and HAYSTAC [49, 50, 51], and is adapted for dark photon searches [19, 18]. Perhaps one of the more innovative aspects of this analysis is producing dark photon limits directly from haloscope data rather than rescaling dark photon limits from axion limits.

The dark photon signal would show up as spectrally narrow power excess over a noise background (Section 2.1). The dark photon signal power is

$$P_S = \begin{cases} \eta\chi^2 m_{A'} \rho_{A'} V_{eff} Q_L \frac{\beta}{\beta+1} L(f, f_0, Q_L) & \text{natural units} \\ \frac{1}{\hbar} \eta\chi^2 (m_{A'} c^2) (\rho_{A'} c^2) V_{eff} Q_L \frac{\beta}{\beta+1} L(f, f_0, Q_L) & \text{with physical constants} \end{cases} \quad (6.1)$$

$$L(f, f_0, Q_L) = \frac{1}{1 + 4\Delta^2} \quad \Delta \equiv Q_L \frac{f - f_0}{f_0} \quad (6.2)$$

where χ is the kinetic mixing angle between the dark photon and Standard Model (SM) photon, η is a signal attenuation factor, $m_{A'}$ is the dark photon mass, $\rho_{A'}$ is the local density of dark matter, V_{eff} is the detector's effective volume, Q_L is the loaded quality factor of the relevant mode, and β is the cavity coupling coefficient. $L(f, f_0, Q_L)$ is the Lorentzian term and depends on the SM photon frequency f , cavity resonant frequency f_0 , and Q_L . The effective volume is a focal point of this thesis and is discussed in various other sections,

including Section 3.2.2. It is calculated by

$$V_{eff} = \frac{\left(\int dV \vec{E}(\vec{x}) \cdot \vec{X}(\vec{x})\right)^2}{\int dV |\vec{E}(\vec{x})|^2 |\vec{X}(\vec{x})|^2} \quad (6.3)$$

where \vec{E} is the relevant mode's electric field and \vec{X} is the dark photon field. Let θ be the angle between the cavity field and the dark photon field, such that $\vec{X} \cdot \vec{E} = |\vec{X}| |\vec{E}| \cos \theta$. Orpheus is only sensitive to dark photons polarized in the \hat{y} direction (along the short side of the rectangular waveguides). If the dark photon is polarized, then $\langle \cos^2 \theta \rangle_T$ can be as low as 0.0025. If the dark photon is unpolarized, Orpheus can detect a third of the dark photons and $\langle \cos^2 \theta \rangle_T = 1/3$ [19].

The dark photon search strategy is to look for a spectrally narrow power excess over the noise floor. In broad strokes, the strategy is to first remove the low-frequency structure from each power spectrum, such that the population mean of each bin is zero and deviation from zero is either from statistical fluctuation due to the noise temperature of the detector or from a coherent RF signal. This results in a unitless power excess normalized to system noise power. In searching for potential dark photon candidates, power excesses with high SNR are sought. Thus the power excess is rescaled such that the power excess is in units of single-bin dark photon power. That leads to the population mean of a bin of 1 in the presence of a single-bin dark photon signal. To account for the dark photon power being spread across many bins, a lineshape filter, i.e., matched filter with dark photon kinetic energy distribution, is applied to the rescaled power excess to increase the SNR of the potential signal. The different, partially overlapping spectra are then combined using a maximum likelihood weighting procedure. In the absence of any dark photon signal candidates, the sample mean and sample standard deviation of each bin in the grand spectrum can be used to place a 90% confidence exclusion limit on the scanned dark photon mass ranges.

6.1 Confirmation that the DAQ Followed the Correct Mode

An ideal way to confirm whether digitization was centered on the dark photon-coupling mode is to measure the mode map, similar to the tabletop measurement Figure 3.19. From the mode map, one determines if the measured cavity mode structure matches expectation, leading to greater confidence that the DAQ followed the correct mode. Unfortunately, the motor mechanism could not tune smoothly throughout the entire range of motion. Manual intervention was needed to handle portions of the tuning range where the motor would stall, so an automated mode map was infeasible. Instead, VNA wide scans were taken once every 20 data-taking cycles. The wide scans measured the transmission coefficient S_{21} from 15 GHz to 18 GHz, as shown in Figure 6.1. These wide scan measurements for various cavity lengths can be stitched together to form the mode map shown in Figure 6.2. The blue line in Figure 6.3 is the frequency the DAQ determined to be the TEM_{00-18} mode in real-time, and digitization was always centered around this frequency. The simulated TEM_{00-18} mode is overlaid in both Figure 6.2 and Figure 6.3 and is shown to be offset from the measured mode directly above by about 0.7 mm. This can be attributed to a systematic error in measuring the absolute length of the cavity, likely caused by the motor system's backlash. Overall, Figures 6.2 and 6.3 show that the TEM_{00-18} mode is easy to follow and that the DAQ digitized around the correct frequency.

The motor encoder values are used to calculate the cavity length. The motor encoders are first set to 0 when the cavity length is 13.6 cm long. The steps counted by the motor encoders are then used to calculate relative changes in distance. One full revolution corresponds to 20,000 motor steps, and twenty revolutions correspond to a plate moving one inch along the 1/4"-20 threaded rod.

There are some features and oddities in Figures 6.1, 6.2, and 6.3 that must be explained. The wide scans in Figure 6.1 are more feature-rich than the wide scans of the tabletop setup shown in Figure 3.18. That is because the cavity sits inside a stainless steel rectangular vessel, and the non-tuning modes of the vessel pollute the mode structure. These non-tuning modes

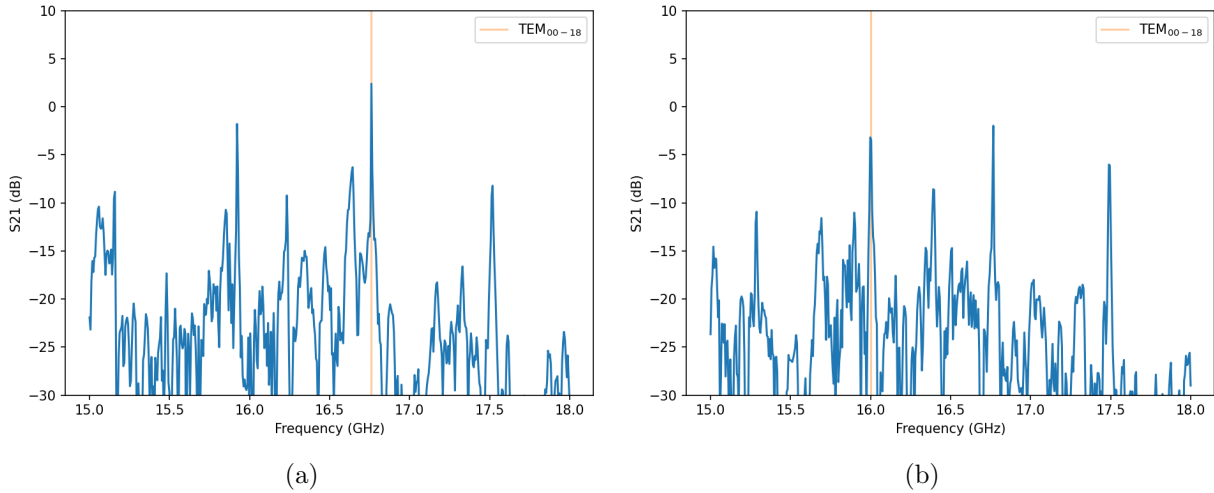


Figure 6.1: Wide scan transmission measurements with the simulated mode overlaid for different cavity lengths. Because the resonator is encased in a stainless steel can, there are low-Q non-tuning modes that make the transmission measurement messier than in Figure 3.18.

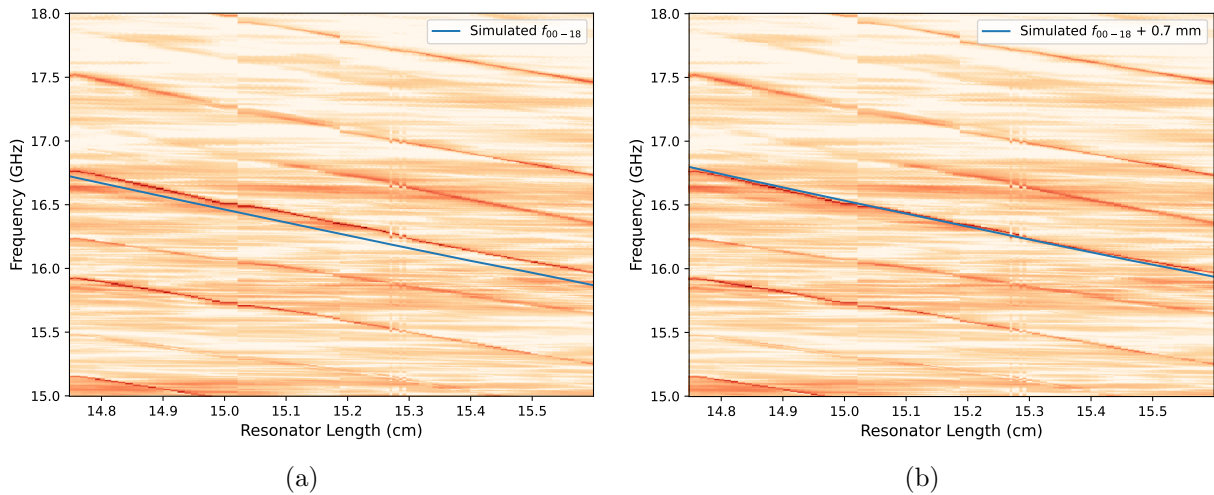


Figure 6.2: Piecing together wide scan measurements to form a mode map. The discontinuities near 15 cm and 15.3 cm are from adjusting the top dielectric plate every once in a while. (b) is the same as (a), but the simulated frequency is shifted over by 0.7 mm to show the systematic error in measuring cavity length.

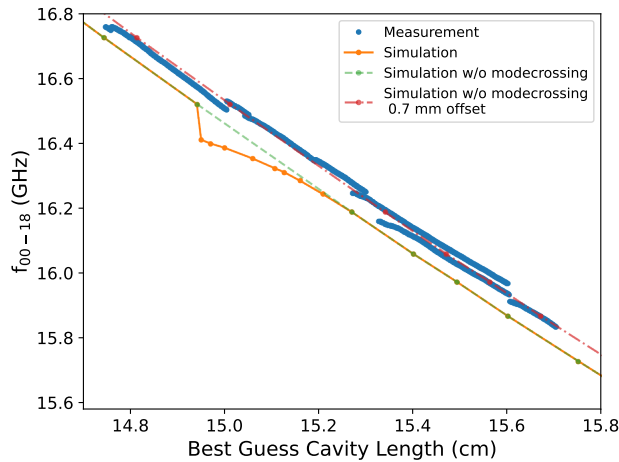


Figure 6.3: The TEM_{00-18} mode frequency vs the cavity length. The cavity length is determined from the stepper motor encoder values.

are also evident as horizontal lines in the mode map on Figure 6.2. The rectangular cavity was not included in the simulations because it would have been computationally expensive. Although it's possible these rectangular cavity modes can mix with the TEM_{00-18} mode and affect V_{eff} , I believe this to be a negligible effect. The quality factors of these rectangular cavity modes are around 200, which is much lower than $TEM_{00-18}Q$ of 5000-10,000. Because of the difference in Q , these modes will hardly mix, and I will use the V_{eff} from the simulations.

Another oddity is that the f_{00-18} in Figure 6.3 seems double-valued and discontinuous. The discontinuities were created when the top dielectric plate was tuned by hand every few hours to compensate for the fact that the motor would stall under normal operation. Each time the top dielectric plate was tuned, f_{00-18} shifted by more than 10 MHz. The discrete steps are shown in Figure 5.10. The discontinuities are evident in the mode map in Figure 6.2 and the measured mode frequency in Figure 6.3.

The resonant frequencies in Figure 6.3 look double-valued because the cavity lengths are swept through twice. The same VNA measurements taken with the same motor encoder values at different times don't result in the same resonant frequency because the motor

backlash would lead to different cavity lengths. The mode map avoids the double-valued resonant frequencies because timestamps below 2021-09-03 21:30:00 UTC are filtered out, so the mode map only shows the second sweep.

The last oddity is mentioned in Section 3.2.2. The simulated frequencies near the simulated mode crossing deviate from the straight line. But tabletop mode map measurements (Figure 3.19) shows that the TEM₀₀₋₁₈ mode tunes smoothly and continuously. This deviation in simulation is likely because the simulations don't have enough precision to deal with the mode crossing. To deal with this, I interpolate the resonant frequency and V_{eff} through the simulated mode crossing. Section 6.3 will describe how V_{eff} is determined for data analysis.

6.2 Confirmation of a Healthy Receiver

Two in-situ measurements were performed to make sure that the receiver was healthy, i.e., no compression and every component is operating within the desired dynamic range.

One measurement is to record the power spectrum as one injects a sine wave, i.e., tone injection, into the cavity's weak port. If the receiver is healthy, the injected RF tone should show up as a tone in the digitized spectrum. A signal generator wasn't available for this injection, so the output of the VNA was used as a tone generator. The VNA behaves like a tone generator if the frequency span and IF bandwidth are set to zero. The center of the VNA sweep was set to the resonant frequency of the cavity. For this measurement (Figure 6.4), the resonant frequency was $f_{00-18} = 16.0742$ GHz, the loaded Q is $Q_L = 10300$, and the cavity was critically coupled.

Another measurement is to record the system noise power as a function of cavity temperature. This measurement can be performed during a cool down or warm up. The system noise power can be described as the sum of the blackbody radiation of the cavity and the added Johnson noise of each electrical component, and is written as $P_n = Gk_b b T_{sys}$, where G is the system gain that depends on both frequency and physical temperature, b is the frequency bandwidth, and T_{sys} is the system noise temperature referenced to the cavity. The

thermal model is described in Section 5.3.

$$T_{sys} = T_{cav}(1 - |\Gamma|^2) + T_{amp,input}|\Gamma|^2 + T_{rec} \quad (6.4)$$

Figure 6.5 shows that the integrated digitized power (the sum of the bins in a power spectrum) tracked linearly with the cavity temperature. For this measurement, the resonant frequency was about 16 GHz and the cavity was critically coupled. The digitized spectrum had a frequency span of six cavity widths and was centered on a frequency two cavity widths away from resonance¹. This linear dependence demonstrates a healthy, non-compressed receiver.

T_{rec} may be determined from the Y-factor measurement at 16 GHz from the relationship between the on-resonance power and cavity temperature. The thermal model for this Y-factor measurement is $T_{sys} = T_{cav} + T_{rec}$. Unfortunately, $\sigma_{T_{cav}} \sim 1$ K during the warm up because of the temperature gradient. So T_{rec} is determined from the measurements performed by Low Noise Factor (Section 5.4).

I also confirmed the amplifiers were not being compressed during operation by adjusting the VNA output power and seeing that the transmission measurement remained the same. I did not check if the IF amplifiers were compressed, but the cascade analysis suggests this to be highly unlikely². All this suggests a healthy, happy, low-noise receiver chain.

¹I was trying to do a Y-factor measurement off-resonance because I was trying to avoid the uncertainty in the cavity temperature. But there is also a lot of uncertainty in the noise temperature coming from the input of the cryogenic amplifier.

²In future work, I recommend checking compression of the IF amplifiers by varying the tone power during the tone injection studies.

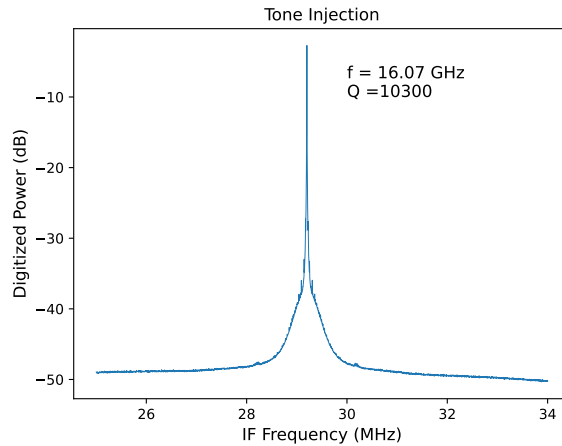


Figure 6.4: Use the VNA to inject a tone into the cavity. Both the tone frequency and the resonant frequency was at 16.07 GHz. The loaded Q was about 10,000.



Figure 6.5: Plotted integrated digitizer power as the cavity is warming up. For reasons outlined in the text, this measurement is difficult to interpret as a Y-factor measurement to obtain the receiver noise temperature. But the digitized power rises more-or-less linearly with the cavity temperature. This suggests a healthy, low-noise receiver chain.

6.3 Parameter extraction

For each data-taking cadence, ancillary measurements were taken to extract the parameters needed to determine the noise temperature and the dark photon power. The necessary parameters are the cavity temperature, resonant frequency, loaded Q, cavity coupling, and effective volume.

The cavity temperature T_{cav} is taken to be the mean of the flat mirror and curved mirror temperature shown in Figure 6.6, and the uncertainty is taken from a continuous uniform distribution, $\sigma_{T_{cav}} = \frac{T_{flat} - T_{curved}}{\sqrt{12}}$. This temperature T_{cav} is thought to be an overestimate of the mean temperature of the thermal photons coming from the cavity and is therefore a conservative choice. The flat mirror was hotter than the rest of the system because it is thermally sunk to the amplifier, which deposits heat. The flat mirror is also subject to a greater heat leak than the curved mirror because the flat mirror is connected to two copper coaxial cables that connect directly to a room temperature SMA vacuum feedthrough, whereas the curved mirror is only connected to one such copper coaxial cable. Due to issues with cooling down described in Section 5.7, I believe the flat mirror is in poor thermal contact with the rest of the cavity. I also believe that the temperature of the dielectrics is closer to the temperature of the curved mirror than to the flat mirror, and most of the cavity's thermal photons come from the curved mirror and dielectric plates.

The resonant frequency, loaded Q, and cavity coupling coefficient were derived using the same type of transmission and reflection measurements used in Chapter 3. The Lorentzian fits were performed online and offline³ while building the combined spectrum. For this analysis, both online and offline fitted parameters were used to overcome different data quirks. The online fitted resonant frequencies were used because the starting and ending frequency for each VNA measurement were improperly logged into the database⁴. The DAQ also logged the online fitted parameters but didn't log the fitting uncertainties. So the fit

³I will refer to online fits as fits that were performed during the data-taking run and offline as fitting performed after the data-taking run.

⁴But I know that the frequency span is 20 MHz. This helped with the offline Q and cavity coupling fits.

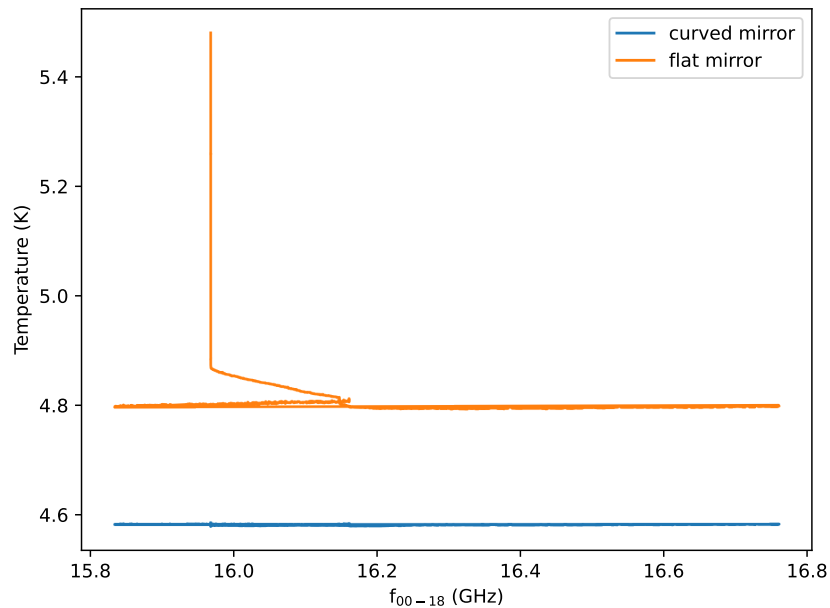


Figure 6.6: Cavity temperature as a function of mode frequency.

was reapplied offline to transmitted and reflected power to obtain these fitting uncertainties ($\frac{\sigma_{f_0}}{f_0}$ is small compared to $\frac{\sigma_{Q_L}}{Q_L}$, so I deemed it ok to just use the online fitted f_0). Q_L and σ_{Q_L} is determined from the Lorentzian fit of the transmission measurement, and β and σ_β is determined from the Lorentzian dip of the reflection measurement (Section: 3.1.3). It's possible to determine Q_L from the reflection measurement, but crosstalk-like effects in the waveguide coupler distort the Lorentzian shape of the reflection measurement⁵, making the reflection Q_L fit unreliable. It's possible to add extra parameters to the reflection fitting measurement to better fit the distorted Lorentzian shape and extract a more accurate Q_L (as is done in Appendix A.3), but the extra fitting parameters have a degeneracy, and extracting a coupling coefficient seems infeasible. The offline fitted Q_L and β are shown in Figures 6.7 and 6.8.

Because the Lorentzian fit doesn't capture the full physics of a reflection measurement,

⁵An example of a distorted Lorentzian is found in Figure A.5.

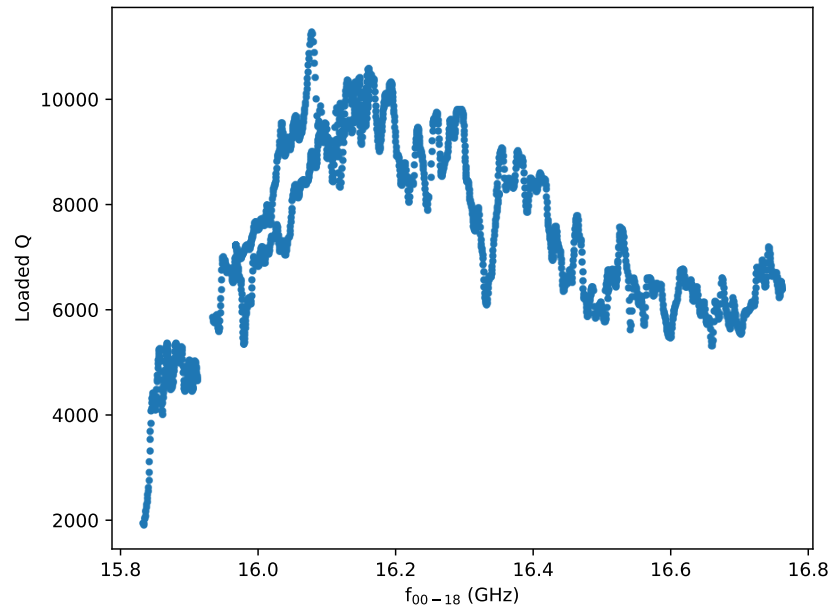


Figure 6.7: Loaded quality factor as a function of mode frequency.

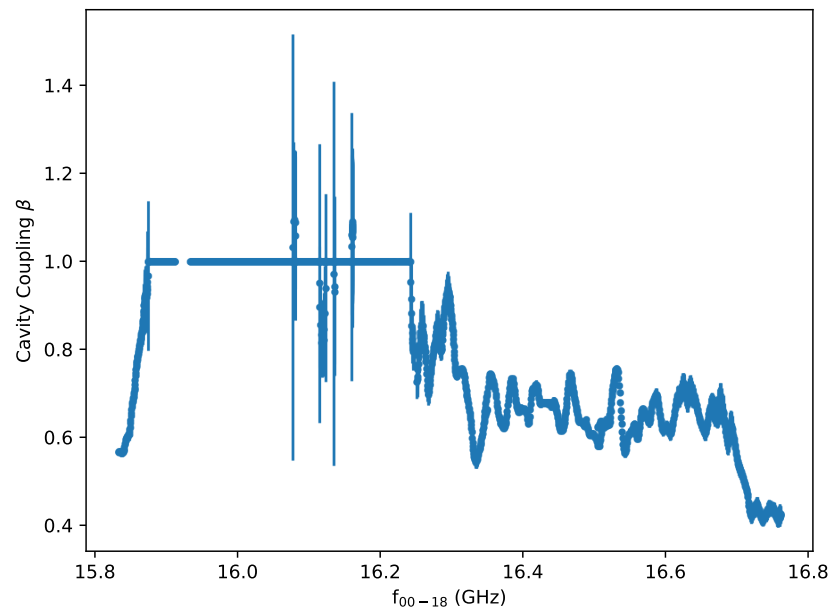


Figure 6.8: Cavity coupling as a function of mode frequency.

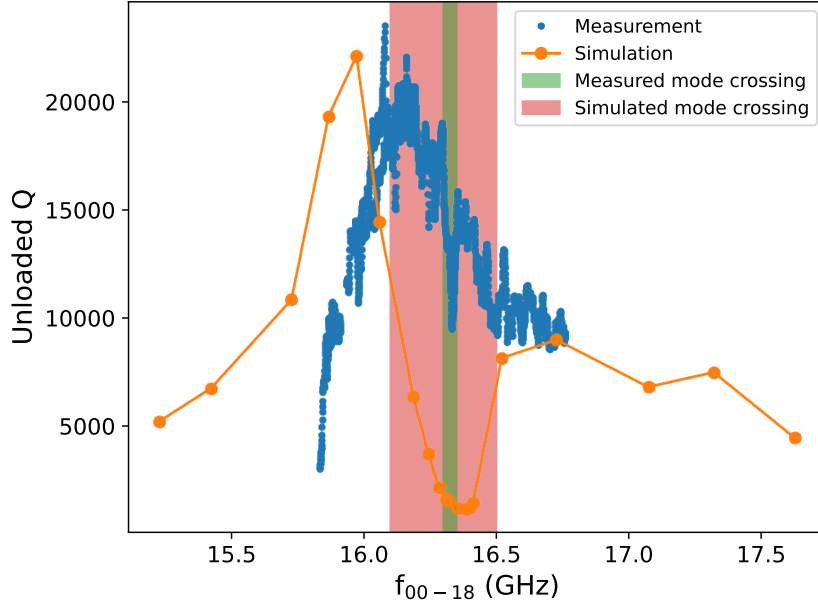


Figure 6.9: Unloaded quality factor as a function of frequency, as derived from the measured loaded quality factor and cavity coupling. The results from the simulation are overlaid for comparison and to identify the mode crossing in the data. The notch in the simulated Q is much wider and steeper than the measured Q . This discrepancy suggests the mode crossing in the simulation is much larger than the mode crossing of the physical cavity.

there are often fits, particularly when the cavity is critically-coupled, when the fitted reflected power is negative on resonance. When this occurs, the Lorentzian dip is often greater than 15 dB, which is very close to critical coupling. In this case, the coupling coefficient is taken to be 1 with no uncertainty. This choice seems safe because near critical coupling, the dark photon power is insensitive to uncertainty in β , as shown in Figure A.7.

V_{eff} is determined from simulation because directly measuring the electric field is infeasible. The simulations are described in detail in Section 3.2.2. This section will focus on dealing with the mode crossing near 16.35 GHz. As has been discussed in Sections 3.2.2 and 3.2.4, a measured dip in quality factor can be used as a proxy for a dip in V_{eff} . Figure 6.9 compares the measured unloaded quality factor $Q_u = Q_L(\beta + 1)$ and the simulated Q_u . The dip in Q_u corresponds to the mode crossing, but the comparison shows that the measured

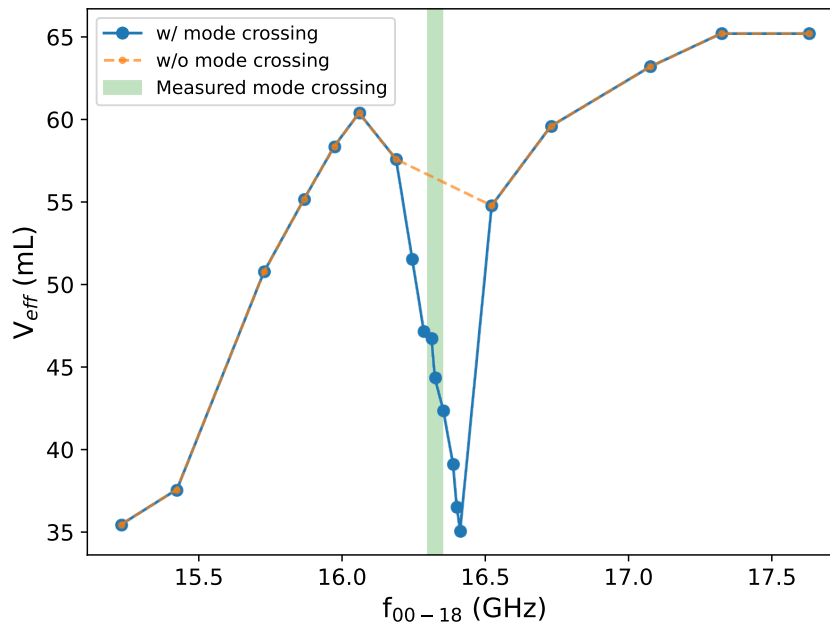


Figure 6.10: The simulated effective volume a function of mode frequency. V_{eff} is plotted with and without the mode crossing because I will have to take care of how I handle the mode crossing in the data analysis.

mode crossing is much narrower than the simulated mode crossing. This suggests that the actual mode crossing does not affect V_{eff} as dramatically as simulations suggest. To deal with this discrepancy, V_{eff} is interpolated through the mode crossing, as shown by the orange curve in Figure 6.10. Within the measured mode crossing region between 16.3 GHz and 16.35 GHz, the effective volume is set to $V_{eff} = 35$ mL, the smallest V_{eff} predicted by simulation. Mode crossings are often omitted from haloscope science limits, but I don't think that's necessary for Orpheus. The intruder mode has a much smaller Q than the TEM_{00-18} mode, and the mode crossing in the measured mode map lacks the avoided crossing structure that's ubiquitous to mode crossings in other haloscopes. $V_{eff} = 35$ mL is a conservative estimate considering that the measured mode crossing is milder than the simulated mode crossing. Outside the measured mode crossing, the orange curve in Figure 6.10 is used.

V_{eff} is simulated for when the dielectrics are evenly spaced throughout the cavity. Un-

fortunately, neither the top nor bottom dielectric plate was in this configuration the entire run. A mechanical failure caused the top dielectric plate motor to stall, so automated tuning was not possible. A cascade of DAQ design flaws and operating errors caused the bottom dielectric plate motor to undershoot the intended position. The position error is shown in Figure 6.11. Fortunately, the simulations in Section 3.2.3 show that at 15.8 GHz, a negative position error actually increases the form factor. The position error is always negative for the bottom dielectric plate and almost always negative for the top dielectric plate. It would take more time than I have to simulate V_{eff} for the position errors of the entire run. But there is a compelling physical intuition for why a negative position error increases the form factor, and this ought to hold true for all frequency ranges scanned.

The position error is used to place an uncertainty on V_{eff} . The position error for the bottom plate is always negative. If the position error of the top dielectric plate is negative, then the analysis uses V_{eff} for the evenly configured case with no uncertainties. Since negative position errors increase V_{eff} , this is a conservative estimate. However, positive position errors decrease V_{eff} . The top dielectric plate is rarely above 5 mm, and Figure 3.15 suggests that a 5 mm error corresponds to a 7.4% decrease in V_{eff} . In the absence of a full simulation campaign, I use the heuristic that any positive position error corresponds to a relative uncertainty of $\frac{\sigma_{V_{eff}}}{V_{eff}} = 0.074$.

For this analysis, the attenuation factor $\eta = 1$. The attenuation factor is comprised of both the insertion loss of the waveguide and electronic components and analysis attenuation from removing the spectral baseline. The insertion loss is largely ohmic, and the resistivity of the waveguide components diminishes at 4 K. As a concrete example, the waveguide-to-coaxial adaptor has an insertion loss of 0.3 dB at room temperature and is made out of aluminum. The residual resistance ratio of aluminum is around 32 [52], so the insertion loss would be about 0.01 dB, subdominant to other sources of uncertainty. The signal attenuation induced by the baseline removal method (discussed in Section 6.5.1) has not yet been studied but should be before publication.

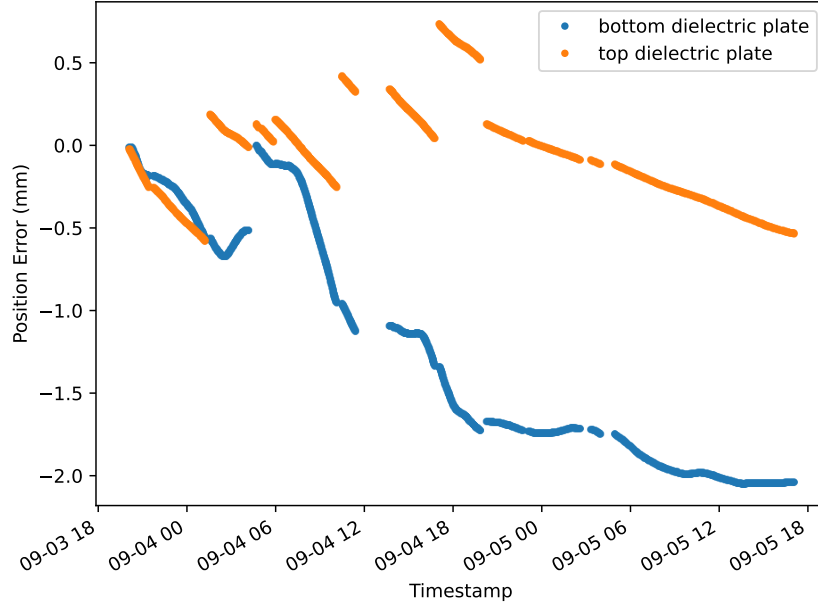


Figure 6.11: The deviation of the plate positions from their targeted positions. This position error was used to estimate the uncertainty in V_{eff} .

6.4 Estimating expected sensitivity

Before assembling the combined spectrum to derive an exclusion limit for χ , I first use the operating parameters to estimate the detector sensitivity. I will use natural units for this calculation. The operating parameters are shown can be read in the plots from the previous section and are shown in Table 6.1.

$$P_{A'} = \chi^2 m_{A'} \rho_{A'} V_{eff} Q_L \frac{\beta}{\beta + 1}$$

$$P_n = bT_n; \quad \sigma_{P_n} = \frac{P_n}{\sqrt{N}} = \frac{P_n}{\sqrt{tb}}$$

$$\text{SNR} = \frac{P_s}{\sigma_{P_n}} = \frac{P_s}{P_n} \sqrt{tb}$$

$$\chi = \sqrt{\frac{\beta + 1}{\beta} \frac{\text{SNR} \times bT_n}{m_\rho V_{eff} Q_L} \left(\frac{1}{bt}\right)^{1/4}}$$

| parameter | value | natural units conversion if needed |
|---------------|--------------------------|----------------------------------------|
| β | 1 | |
| $V_{eff,max}$ | 60 cm ³ | |
| $\rho_{A'}$ | 0.45 GeV/cm ³ | |
| $m_{A'}$ | 65×10^{-15} GeV | |
| t | 100 s | 1.5×10^{26} GeV ⁻¹ |
| Q | 10000 | |
| T_n | 10 K | 8.617×10^{-13} GeV |
| SNR | 3 | |

Table 6.1: Operating parameters for Orpheus used to estimate Orpheus science reach.

So

$$\chi_{max} = 4.94 \times 10^{-14}; \quad \chi = \frac{\chi_{max}}{\sqrt{\langle \cos^2 \theta \rangle_T}}$$

$$\chi_{unpolarized} = \frac{\chi_{max}}{\sqrt{1/3}} = 8.56 \times 10^{-14}$$

$$\chi_{polarized} = \frac{\chi_{max}}{\sqrt{0.0025}} = 9.88 \times 10^{-13}$$

6.5 Processing Individual Raw Spectrum

6.5.1 Removing Spectral Baseline

The raw spectrum consists (Figure 6.12) of the system noise power $P_n = Gbk_b T_n$, plus fluctuations about this noise power. T_n affects the SNR of the dark photon data. The system gain, i.e., transfer function, does not affect the SNR of the dark photon data, but causes the large-scale gain variation seen in Figure 6.12. The first step of processing the raw spectra is to remove this gain variation. The gain variation is caused by both the IF electronics and RF electronics. The IF gain and RF gain are thought of as separate because the IF gain applies equally to all spectra, whereas the RF gain depends on the frequency of the cavity. So to remove the gain variation, we first divide the spectrum by G_{IF} , and then by G_{RF} .

Figure 6.13 shows the same raw spectrum but with a hypothetical dark photon signal of

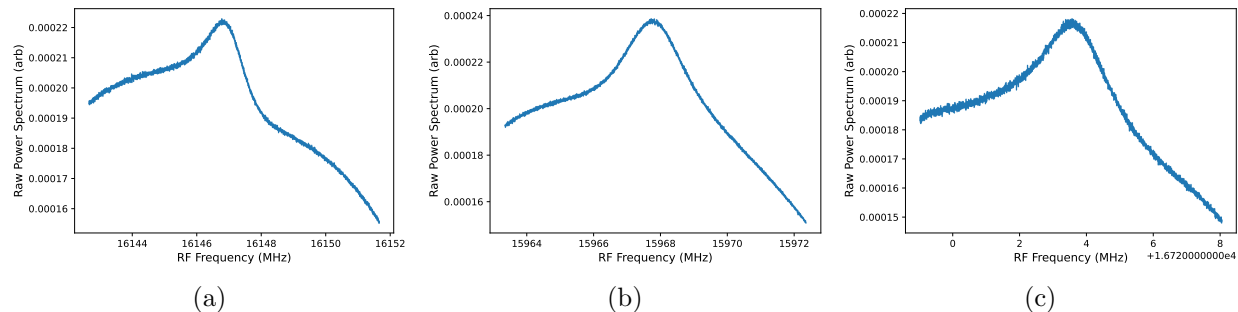


Figure 6.12: Raw spectrum as measured by the digitizer. Assuming the spectrum is just noise, the raw spectrum is the noise power of the system.

arbitrary power superimposed. The dark photon signal is much narrower than the large-scale gain variation caused by the electronics. Thus it is safe to apply filters to remove this gain variation.

The IF baseline is estimated by averaging all the power spectrum while the cavity was actively tuning. Power spectra with timestamps after 2021-09-05 17:00:00 UTC were excluded because the cavity rested at 15.8 GHz, and including those power spectrum in this baseline would have heavily weighted that RF frequency⁶. The IF baseline is then smoothed using a fourth-order Savitzky-Golay software filter with a window length of 501 bins. Note that the Savitzky-Golay window is much larger than the hypothetical dark photon signal in Figure 6.13, which is about eight bins wide. The resulting IF baseline is shown in Figure 6.14. Every spectrum is divided by this IF baseline, and the result is a unitless power spectrum that is still affected by the RF gain variation, as shown in Figure 6.15.

To remove the RF gain variation, a second-order Savitzky-Golay software filter with a window length of 101 bins is applied to the individual spectrum[50], as shown in the orange curves in Figure 6.15. The spectra are then divided by the filtered curve to remove the RF gain variation. The processed spectra are now normalized such that the mean of the bins is one. The dark photon signal would show up as a power excess, so the process spectra

⁶In practice, neglecting this cut didn't seem to make much difference.

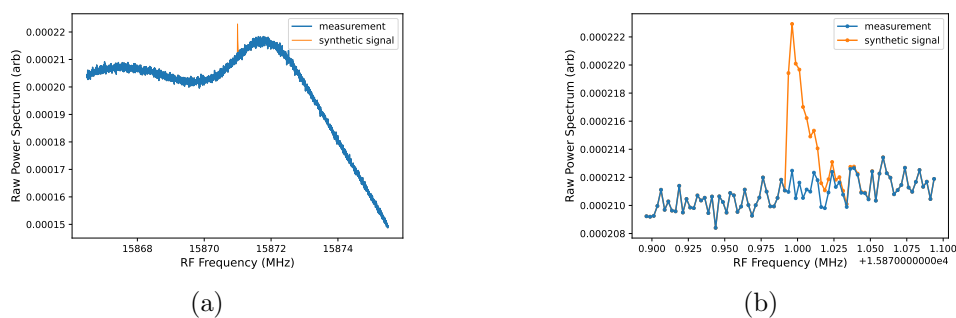


Figure 6.13: Raw spectrum with a hypothetical dark photon signal of arbitrary power superimposed. (a) The dark photon signal is much narrower than the large-scale structure of the noise background. Thus a low-pass filter can be safely applied to each power spectrum. (b) Same as (a), but zoomed into the dark photon signal. The hypothetical dark photon signal covers about 8 RF bin widths.

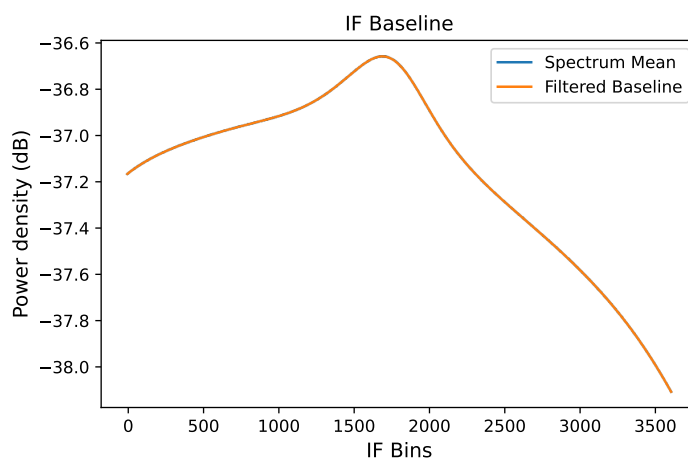


Figure 6.14: The average of all the recorded spectrum. This is thought of as a good estimation of the gain variation caused by the IF electronics. The estimation is less good if you don't have even and continuous tuning through your frequency range.

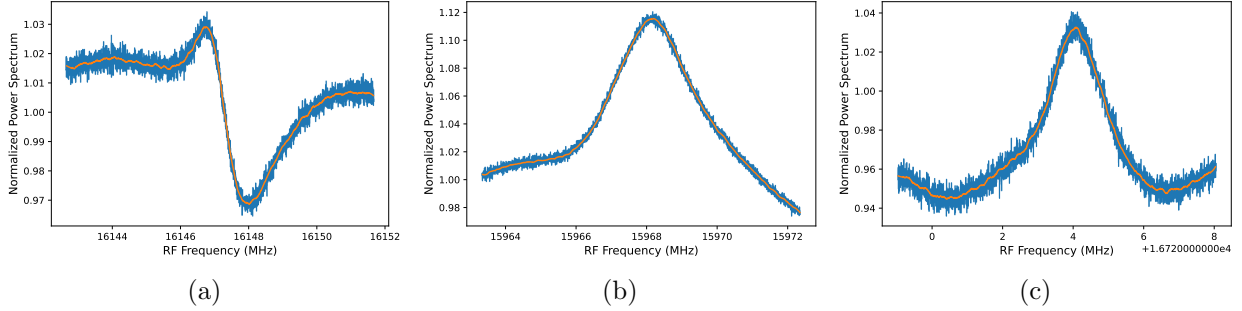


Figure 6.15: The raw spectra are all divided by the IF baseline to form some normalized spectrum. A Savitzky-Golay filter is then applied (orange) to obtain gain variation from the RF electronics.

are subtracted by one to yield the flat spectra fluctuating about zero, shown in Figure 6.16. Each bin is either an average of 75,000 spectra or 250,000 spectra, and so the fluctuations about zero are Gaussian by the Central Limit Theorem. The power excess for each bin in the processed spectrum is denoted as δ_p . All bins in a spectrum are pulled from the same Gaussian distribution, and so the uncertainty of each bin σ_p is taken as the standard deviation of all the bins in a spectrum. δ_p is the power excess normalized to units of system noise power P_n .

6.5.2 Rescaling Spectra to be in units of dark photon power

Next, the data is rescaled to be in units of dark photon power. Different operating conditions, such as different temperatures or Q_L , will change the noise temperature or expected signal power. The data is rescaled so that a single-bin dark photon signal would have one unit of excess power, regardless of the operating conditions. This rescaling allows for different spectra to be more directly comparable. This rescaling also makes the SNR (the signal being that from the dark photon) the true figure of merit for which potential dark photon candidates are sought. Currently, the spectra are in units of noise power. To rescale, multiply

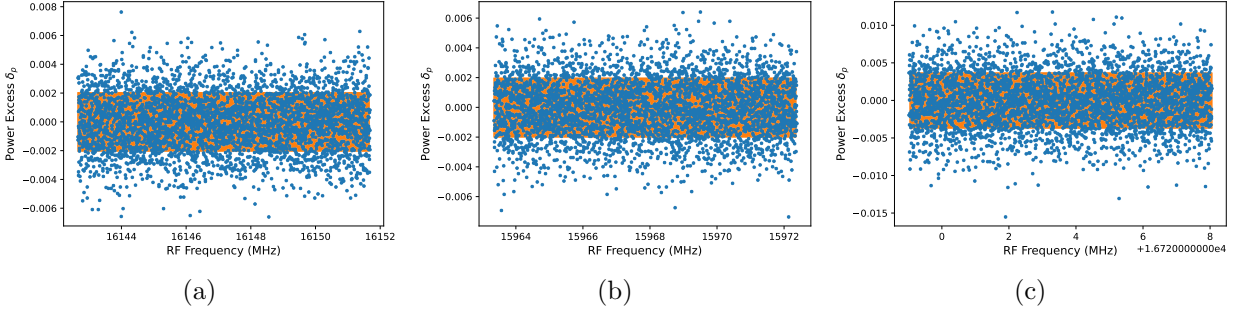


Figure 6.16: The power excess in units of system noise power.

each bin by the noise power and divide by the hypothetical single-bin dark photon power,

$$\delta_s = \delta_p \frac{k_b b T_n}{P_s} \quad (6.5)$$

There is an ambiguity in Equation 6.2 because χ and θ parameters that are not determined by experimental conditions. In particular, there is no natural benchmark model for dark photon (like there are with axion models) to use as a baseline. So for spectrum processing, $\chi = 1$ and $\theta = 0$. The analysis procedure will self-correct for the arbitrarily chosen $\chi = 1$, and the limits can be rescaled at the very end to match any cosmological model that claims a value for θ .

The uncertainty is determined using standard uncertainty propagation techniques.

$$\frac{\sigma_s^2}{\delta_s^2} = \frac{\sigma_p^2}{\delta_p^2} + \frac{\sigma_{P_N}^2}{\delta_{P_N}^2} + \frac{\sigma_{P_S}^2}{\delta_{P_S}^2} \quad (6.6)$$

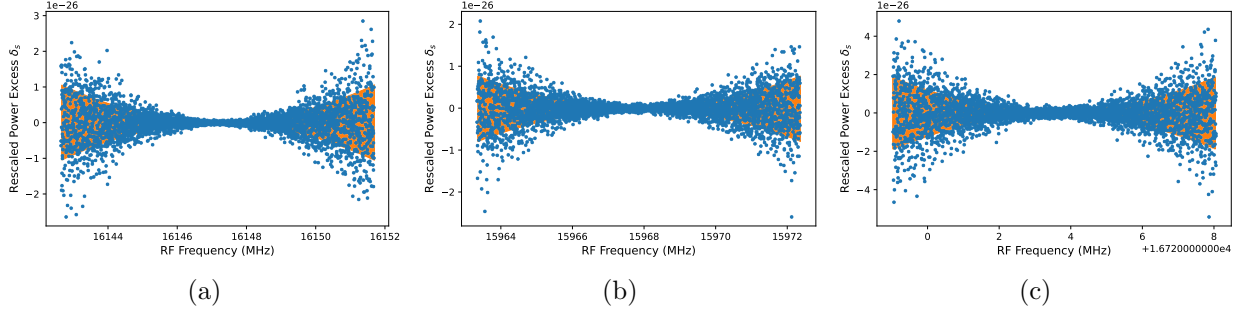


Figure 6.17: The power excess rescaled to units of single-bin dark photon power.

6.5.3 Taking into account the dark photon lineshape

The spectra were rescaled so that a dark photon signal would have unit height if the dark photon signal was confined to a single bin. However, Figure 6.13 shows that the expected hidden photon lineshape is spread over about eight bin widths, and this reduces the SNR of each bin.

The spectral shape of the dark photon signal is proportional to the dark photon kinetic energy distribution. The most conservative energy distribution assumes a virialized, isothermal halo that obeys a Maxwell-Boltzmann distribution,

$$f(f) = \frac{2}{\sqrt{\pi}} \sqrt{f - f_a} \left(\frac{3}{f_a \langle \beta^2 \rangle} \right)^{3/2} \exp \frac{-3(f - f_a)}{f_a \langle \beta^2 \rangle} \quad (6.7)$$

where f is the measured frequency, f_a is the frequency of the associated SM photon, $\langle \beta^2 \rangle = \frac{\langle v^2 \rangle}{c^2}$, and $\langle v^2 \rangle$ is the rms velocity of dark matter halo. The “quality factor” of the lineshape is $\mathcal{O}(10^6)$. In the lab frame, the lineshape takes a more complicated form, but is well approximated by the Equation 6.7 if $\langle \beta^2 \rangle \rightarrow 1.7 \langle \beta^2 \rangle$ [50].

One naive approach is to rebin the spectrum so that it has lower frequency resolution, as shown in Figure 6.18. However, the dark photon power can still be divided into two neighboring bins. At worst, the power will be divided equally into two bins.

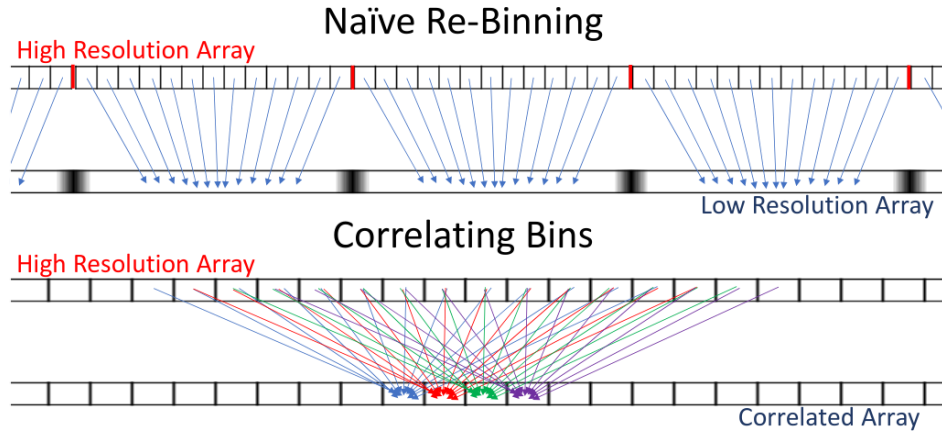


Figure 6.18: A diagram of how convolution can enhance the signal of a multi-bin dark photon signal. Figure courtesy of Christian Boutan [53].

A more elegant approach that increases the SNR is to convolve the spectrum, such that each bin is a weighted sum of the neighboring bins. The weight is determined by the line-shape. This has the effect of enhancing bumps that span multiple neighboring bins, increasing the SNR of a potential dark photon signal. The convolution is shown at the bottom of Figure 6.18. In digital signal processing, this is known as applying a matched filter.

There is a subtlety. How does one perform a convolution when neighboring bins have different uncertainties associated with them, and what is the resulting uncertainty after the convolution? The issue is resolved by treating it like a χ^2 minimization problem. Let

$$\chi^2 = \sum \frac{(Ak_i - \delta_{si})^2}{\sigma_{si}} \quad (6.8)$$

where δ_{si} (σ_{si}) is δ_s (σ_s) corresponding to bin i , k_i is the element of the convolution kernel corresponding to bin i , and A is a amplitude of the signal shape to which we are fitting our data. The more the bins δ_{si} match the shape of k_i , the larger the amplitude A . χ^2 is minimized by setting the derivative of χ^2 to zero and solving for A .

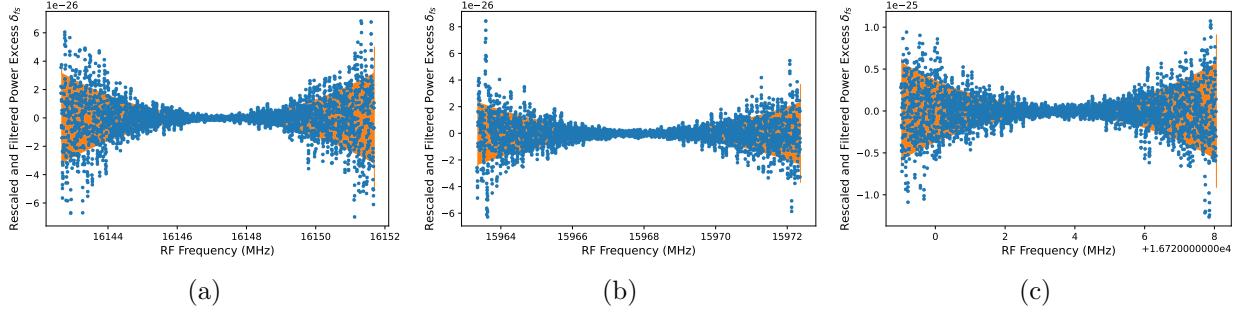


Figure 6.19: The rescaled power excess after the dark matter halo lineshape filter has been applied.

$$A = \frac{\sum \frac{\delta_{si} k_i}{\sigma_{si}^2}}{\sum \frac{k_i^2}{\sigma_{si}^2}} \quad (6.9)$$

A is the value for each bin in the optimally convolved spectrum. The uncertainty for parameters derived from χ^2 minimization fits is known to be $\sigma_A^2 = \left(\frac{d^2 \chi^2}{dA^2} \Big|_{\chi_{min}^2} \right)^{-1}$, resulting in

$$\sigma_A^2 = \frac{1}{\sum \frac{2k_i^2}{\sigma_{si}^2}} \quad (6.10)$$

In other words, a moving chi-square minimization fit is performed on the data. This is reminiscent of pulse shape fitting done in other analyses.

Applying the matched filter with the lineshape kernel leads to the spectra in Figure 6.19.

Once all the spectra are processed, one can obtain a histogram of the power fluctuations normalized to the standard deviation, as shown in Figure 6.20. The normalized processed power fluctuations $\frac{\delta_p}{\sigma_p}$ and rescaled process spectrum $\frac{\delta_s}{\sigma_s}$ follow the normal distribution beautifully, with zero mean and unit standard deviation. The normalized filtered spectrum $\frac{\delta_{fs}}{\sigma_{fs}}$,

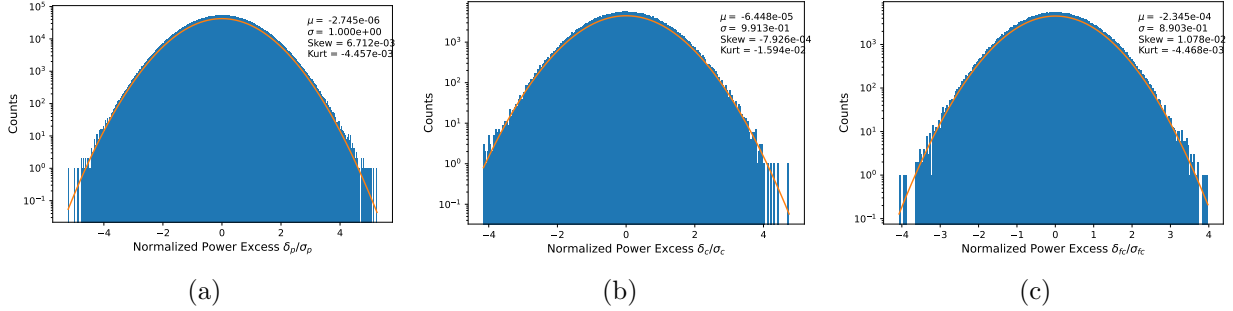


Figure 6.20: The (a) power excess δ_p and (b) rescaled power excess δ_s has a Gaussian distribution. The filtered power excess has a smaller standard deviation than what is expected of Gaussian noise. This can be attributed to correlations induced by the convolution.

follows the shape of a standard normal distribution, but has a standard deviation of 0.89. This is because the convolution correlates neighboring bins.

6.6 Combined Spectrum

The individual spectra corresponding to different RF frequencies have been processed and are now ready to be combined into a single combined spectrum. The Maximum-Likelihood estimate of the mean and uncertainty of the combined spectrum is obtained by a weighted average of all IF bins corresponding to a particular RF bin. The weights of each contributing bin are the inverse variance.

$$\delta_c = \frac{\sum_i \frac{\delta_{si}}{\sigma_{si}^2}}{\sum_i \frac{1}{\sigma_{si}^2}} \quad (6.11)$$

The standard deviation for each bin in the combined spectrum is

$$\sigma_c = \sqrt{\frac{1}{\sum_i \frac{1}{\sigma_{si}^2}}} \quad (6.12)$$

This method of combining spectrum is well-established in haloscope data analysis [50,

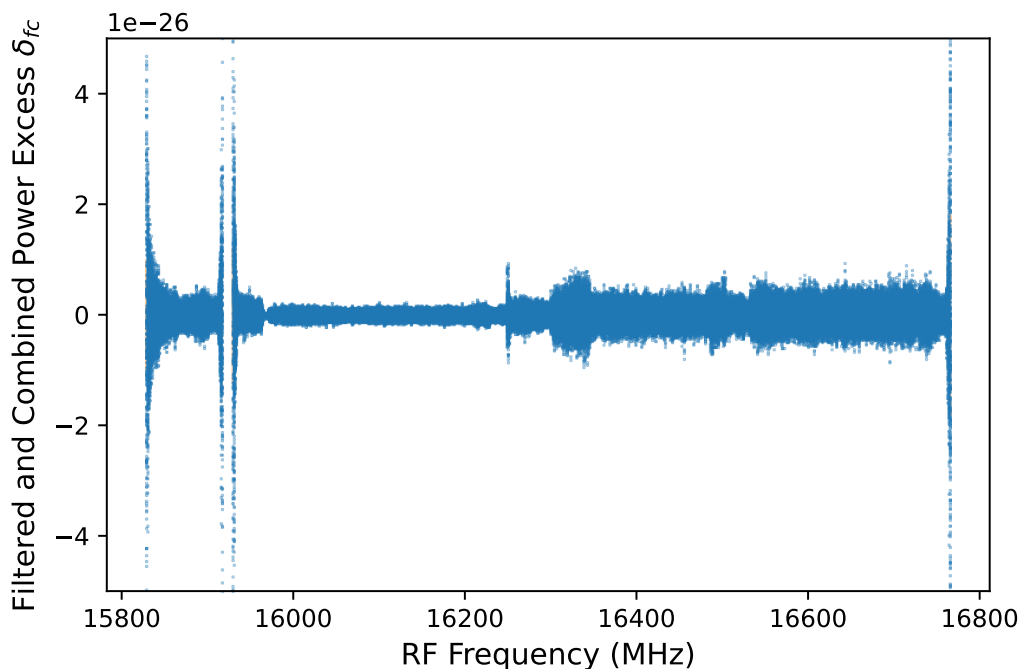


Figure 6.21: The combined spectrum.

51, 49]. Finding the Maximum-Likelihood estimations of means and standard deviations is a general and ubiquitous problem in experimental data analysis.

The combined power excess is shown in Figure 6.21. The gap near 15.95 GHz is a result of manually tuning the top dielectric plate.

6.7 Placing 90% Exclusion Limit

Figure 6.20 suggests that there are no signals that are astatistical. In the absence of any obvious dark photon candidates, I choose to set an upper limit on what dark photon signals could exist based on the statistics of the combined spectrum.

The probability distribution of measuring a power excess δ_c given a haloscope signal power is a normal distribution centered around the dark photon signal power (the signal

power can be 0).

$$P(\delta_c|S) = \frac{1}{\sqrt{2\pi}\sigma_c} \exp\left(-\frac{(\delta_c - S)^2}{2\sigma_c^2}\right) \quad (6.13)$$

where δ_c , σ_c are the mean and standard deviation of the combined power excess, and S is the dark photon signal power. I've measured power excess and want to know the probability distribution of the dark photon signal. In other words, I want to know probability distribution of S given that I know the power excess δ_c , $P(S|\delta_c)$. Using Bayes Theorem

$$P(S|\delta_c) = P(\delta_c|S) \frac{P(S)}{P(\delta_c)} \quad (6.14)$$

where $P(S)$ is the probability of measuring dark photon signal power S , and $P(\delta_c)$ is the probability of measuring power excess δ_c . δ_c in Equation 6.14 is a parameter and not a continuous variable, so $P(\delta_c)$ is just a constant. Our only prior is that the signal power can not be negative⁷. In the absence of any other information, any value for S above zero is as good as any other. Consequently,

$$\frac{P(S)}{P(\delta_c)} \propto H(x) \quad (6.15)$$

where $H(x)$ is the unit step function. That means $P(\delta_s|S)$ is just a normal distribution truncated at zero⁸.

$$P(S|\delta_c) \propto \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(\delta_c - S)^2}{2\sigma^2}\right) H(x) \quad (6.16)$$

The 90% confidence limit for the dark photon signal power is the value of δ_c for which 90% of the probability distribution lies below. The inverted, but equivalent, statement is

⁷ S also has an upper limit based on previous exclusions, but this limit has a negligible effect on the subsequent analysis.

⁸[54, 55] talk about similar subjects.

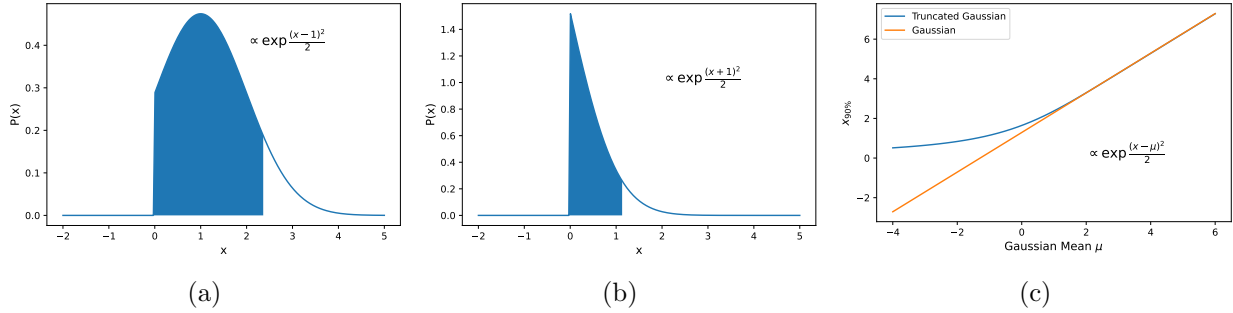


Figure 6.22: The probability distribution of finding a dark photon signal, given a measured power excess, is represented as a truncated Gaussians. The truncated Gaussians are plotted for (a) $\delta_{fc}/\sigma_{\delta_{fc}} = 1$ and (b) $\delta_{fc}/\sigma_{\delta_{fc}} = -1$. The 90% confidence interval is shaded in the plots. The value representing the upper limit of the confidence interval is plotted for different values of $\delta_{fc}/\sigma_{\delta_{fc}}$. The confidence interval for a full normal distribution is plotted in orange for comparison.

that signal powers above this δ_c are excluded with 90% confidence. The confidence interval means that if the same experiment were repeated, there is a 90% chance the dark photon signal (even if that signal is 0) would lie in this interval. The confidence interval for truncated normal distributions with different means is shown in Figure 6.22. The 90% confidence limits for dark photon signal power is determined by applying the percent point function to the truncated normal distribution⁹. The result is shown in Figure 6.23.

By examining Equation 6.2, one sees that

$$\frac{\chi_{90\%}}{\chi = 1} = \sqrt{\frac{P_{90\%}}{P(\chi = 1)}} \quad (6.17)$$

Finding $\chi_{90\%}$ amounts to taking the taking the square root of Figure 6.23. However, this $\chi_{90\%}$ corresponds to the maximum V_{eff} . This is the maximally benevolent model where the photon polarization always aligns with the polarization of the TEM₀₀₋₁₈ mode, even as the Earth spins and orbits through the cosmos. The realistic mixing angle is $\chi_{90\%} = \frac{\chi_{90\%,max}}{(\cos^2 \theta)_T}$.

⁹Python's SciPy module has a native function for this.

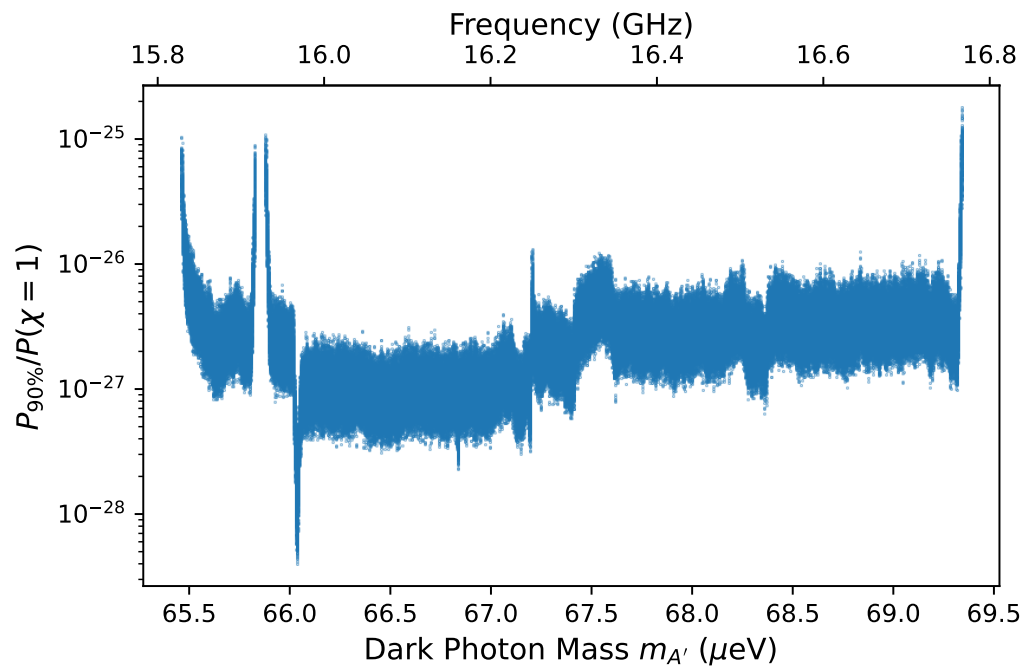


Figure 6.23: The 90% confidence limit on the dark photon power, normalized to dark photons signal power of mixing angle $\chi = 1$.

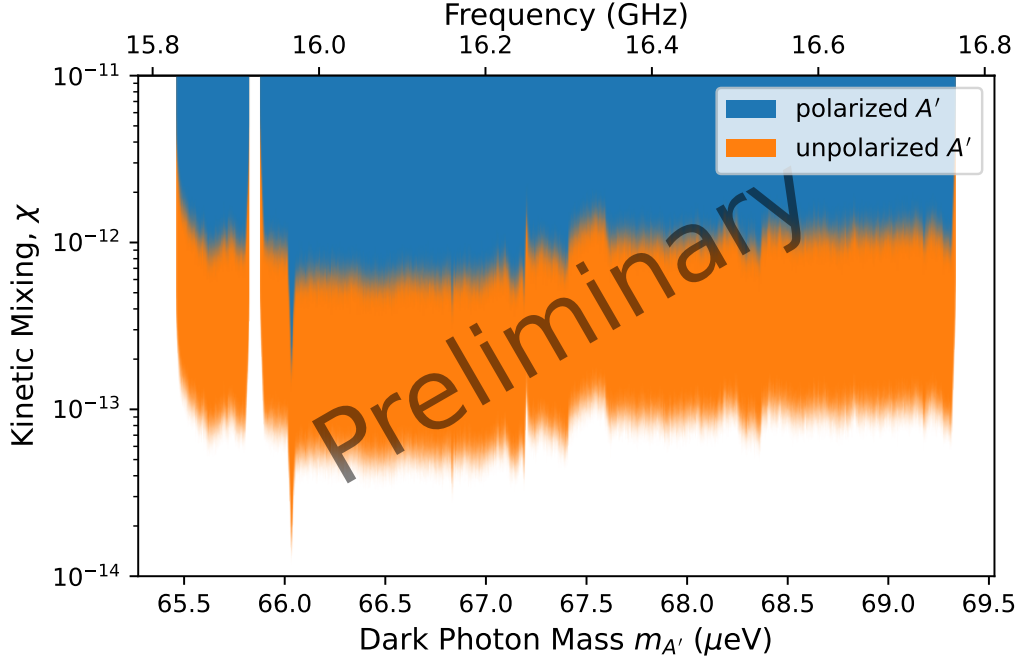


Figure 6.24: A 90% exclusion on the mixing angle parameter space.

$\langle \cos^2 \theta \rangle_T = 1/3$ if the photons are unpolarized, and $\langle \cos^2 \theta \rangle_T \geq 0.0025$ if they are polarized. Between $65.5 \mu\text{eV}$ and $69.3 \mu\text{eV}$, the excluded dark photon mixing angle is $\chi_{90\%} \sim 1 \times 10^{-12}$ for the polarized dark photon case and $\chi_{90\%} \sim 1 \times 10^{-13}$ for the unpolarized dark photon case. The limits are shown in Figure 6.24 and are consistent with expectations laid out in Section 6.4. The limits are compared to other haloscope experiments in Figure 6.25. Orpheus achieves similar sensitivities at substantially higher frequencies than other larger, more well-funded haloscope experiments while maintaining a respectable tuning range.

The limits in the polarized case could be improved by perhaps an order of magnitude for selected frequencies by using the fact that the same frequencies were rescanned at different times of the day[19]. This may be accounted for in a future analysis.

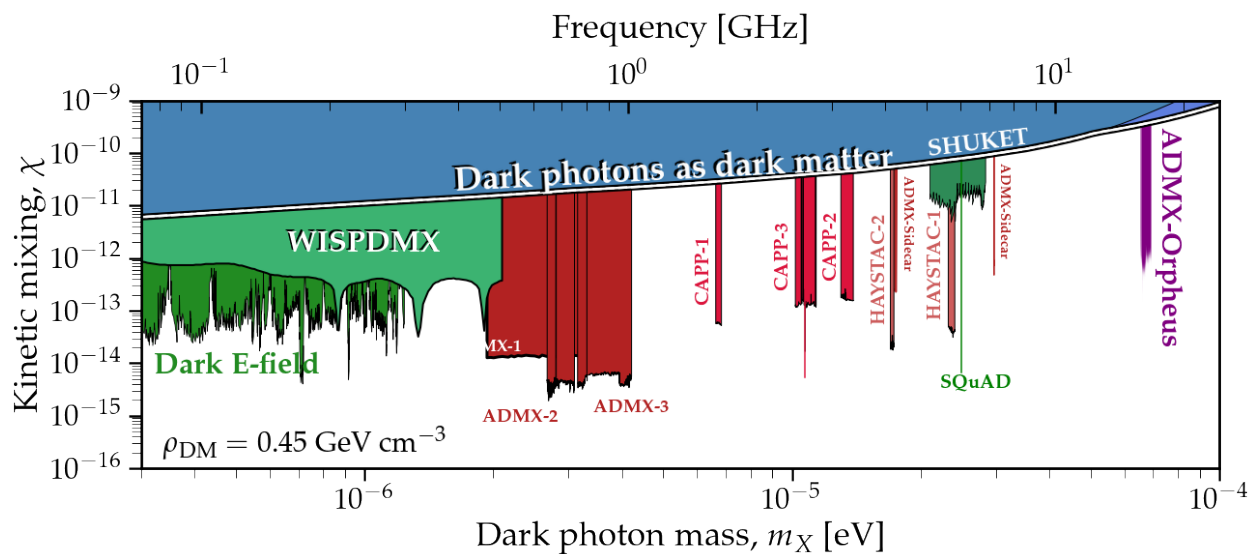


Figure 6.25: Orpheus science reach in the context of other microwave cavity results. The limits are for the uniformly polarized case. Figure adapted from [16].

Chapter 7

FUTURE DIRECTION

I spent three years simulating, prototyping, designing, characterizing, debugging, and commissioning Orpheus. The result was two very fruitful days of collecting dark photon data. The Orpheus test stand remains and is ready to gather more science data as needed. This chapter describes how to get more science out of Orpheus.

7.1 Upcoming Axion Run

The last run was without a magnetic field, and so Orpheus was only sensitive to dark photons. But adding an imaginary 1.5 T magnet to the very real haloscope data demonstrates that Orpheus would have easily been sensitive to $g_{a\gamma\gamma} \sim 3 \times 10^{-12}$ from 15.8 GHz to 16.8 GHz (Figure 7.1), over an order of magnitude more sensitive than CAST. That would be the first demonstration of a tunable haloscope beyond 7 GHz, as shown in Figure 7.2.

Orpheus is currently preparing for an axion data-taking run, with James Sinnis at the helm. James Sinnis is winding a 1.5 T superconducting dipole magnet and has spent the Summer of 2021 developing and practicing the winding technique. See Appendix A.5 to learn more about the magnet design and winding technique.

There are near-term hardware upgrades to improve the automated tuning mechanism so the motors don't stall. Mainly, the coupled shafts need to be better aligned, the motor vacuum ports need to be straightened, and more appropriate shaft couplers are needed to handle any remaining misalignments. Section 4.3 explains the myriad of motor problems in need of fixing. Thermal issues also limited the science reach of this past run. The thermal issues described in Section 4.4 should be addressed by clamping the mirror to its holder so that it won't lose thermal contact during cool down. The cold finger will also be thermally

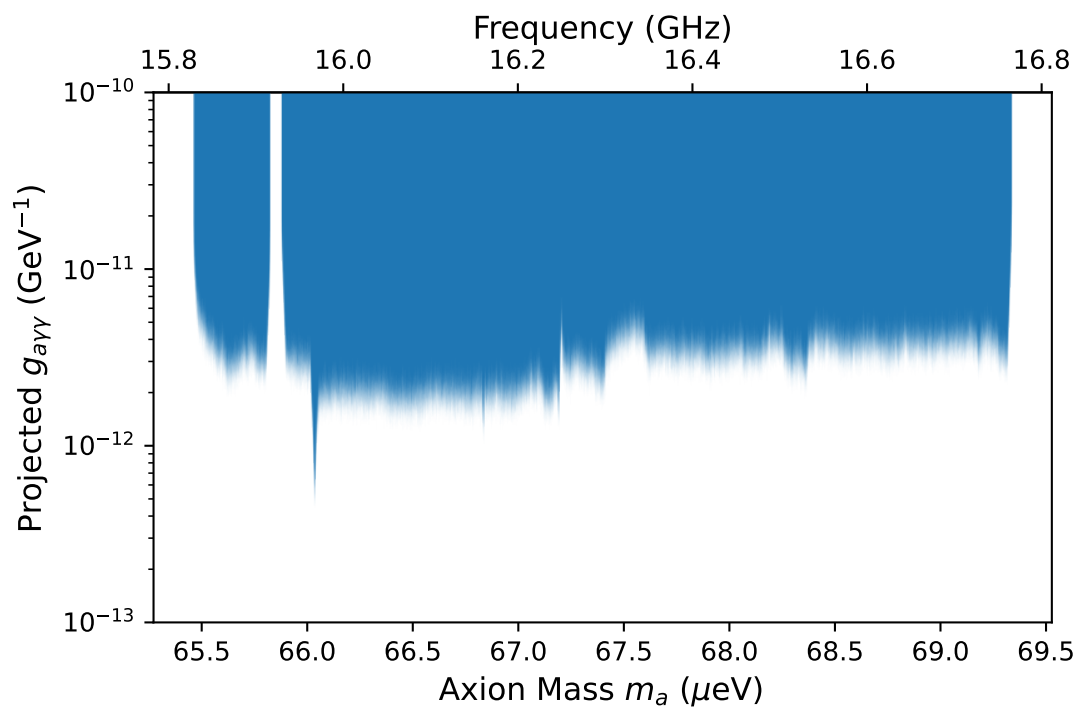


Figure 7.1: A 95% exclusion on the axion-photon coupling $g_{a\gamma\gamma}$ had there been a 1.5 T dipole magnet.

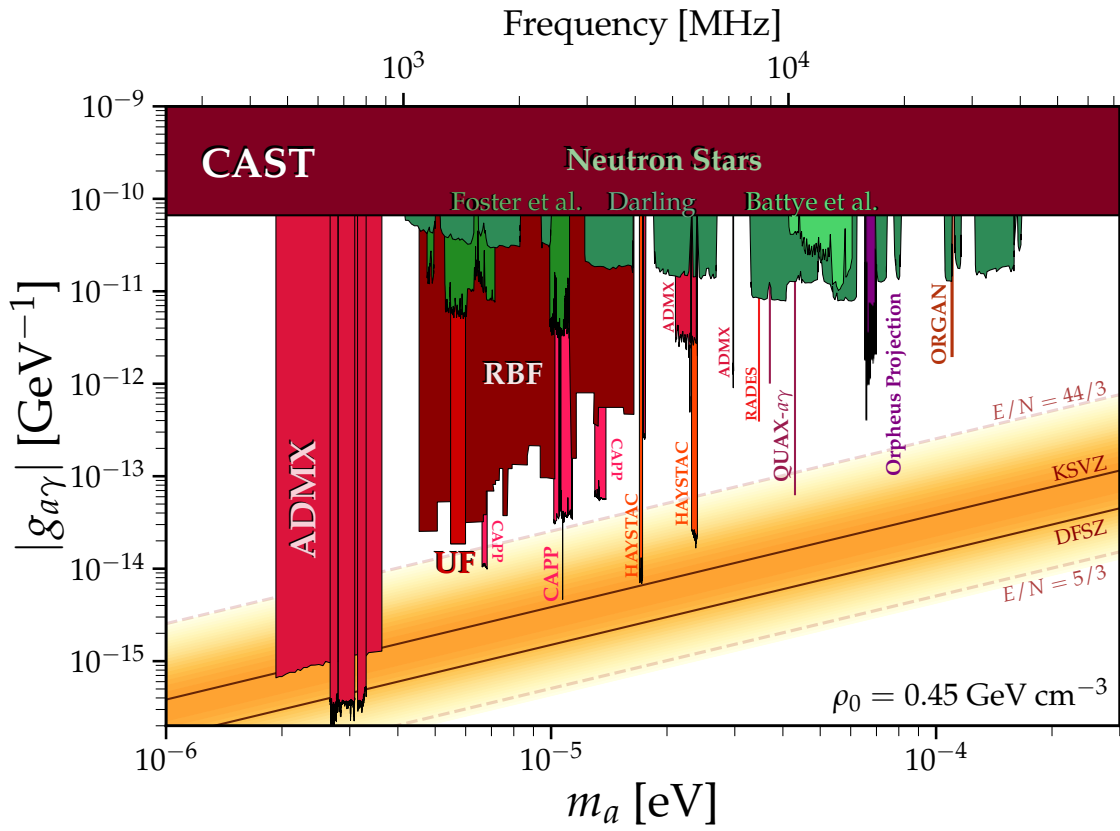


Figure 7.2: The projected axion limits in the context of other microwave cavity searches. Figure adapted from [16].

connected to other portions of the cavity for more robust cooling in case the flat mirror loses thermal contact during cooldown. Temperature sensors will be added to the dielectric plates to better understand the thermal gradient along the cavity. The cavity must also be better thermally anchored to itself for better temperature uniformity.

A series of DAQ bugs mentioned in Chapters 5 and 6 made data analysis awkward and will be fixed for the next run.

7.2 *More Science with Same Apparatus*

Beyond this coming axion run, there is more science one can do with the same apparatus.

7.2.1 *Electrodynamic Optimizations*

Orpheus has lots of degrees of freedom in its design, and one can spend time optimizing the cavity properties. Many design choices were just reasonable first guesses from what I knew as a third-year grad student. I chose the dielectric thickness to be about $\lambda/2$ at 16.5 GHz and mirror radius of curvature to be about twice the cavity's optical length¹ near 18 GHz. Simulations were difficult, and I had little computing power to simulate different geometries. The simulations disagreed with what I expected from physical optics, and I didn't know if the simulations were accurate. I built the Orpheus cavity as soon as I could to validate the simulations and to just have something working as quickly as I could². Before building the cavity, no one knew if the cavity design would work, and many around me were skeptical. It was not until the Orpheus cavity was assembled and the modes matched the simulation that I felt comfortable trusting my simulation.

Now that the cavity concept has been verified and simulations have been validated, there are a lot of degrees of freedom to optimize to improve the quality factor and effective volume. As currently configured, the Orpheus V_{eff} is about 2% of the physical volume, and that's not

¹This is the confocal configuration for a Fabry-Perot cavity.

²Software developers have a maxim, "Make it work. Make it right. Make it fast." In my context, I wanted to get something working before I thought about optimizations.

good. Surprisingly, moving the dielectrics away from the curved mirror by 1 mm improved V_{eff} by more 7% (Section 3.2.3). Now that the simulations can be trusted, one can perform a parametric sweep of the dielectric positions, mirror radius of curvature, and dielectric thickness to try and optimize the Q_u and V_{eff} . One could also increase V_{eff} by placing dielectrics at every other half-wavelength instead of every fourth half-wavelength. But I suspect that reduces the bandwidth of the cavity.

For scaling the Orpheus concept to larger sizes, it's important to manage the diffraction losses. Diffraction is a dominant source of loss³ and will only become worse as the cavity length increases relative to the mirror size. Diffraction losses can be mitigated by optimizing the mirror radius of curvature (likely by decreasing it). One can also reduce diffraction by increasing the mirror size, but that increases the physical volume of your cavity, which would necessitate a larger magnet, and large dipole magnets are difficult to fabricate⁴. It may be possible to mitigate diffraction losses by curving the surface of the dielectric plate so that they act as lenses that collimate the field. However, given that the dielectrics are only a few mm thick, its radius of curvature would have to be about 100 in, and may be difficult to machine. One workaround is to have dielectrics with a smaller diameter than the mirrors, but that's another parameter to optimize. It's also likely that curving the dielectrics reduces the cavity's bandwidth.

The dielectric losses can be reduced by using sapphire instead of alumina. Sapphire has a dielectric loss tangent $\tan_d \sim 0.00002$ compared to alumina's $\tan_d \sim 0.0001$. But one would have to deal with sapphire's birefringence.

Developing a reliable and adjustable cavity coupling mechanism is important for scaling the concept. For a fixed coupling hole size, the cavity coupling coefficient reduces for longer cavities. Impedance matching is a theoretically solved problem (see Pozar 4th ed. Ch. 5 [30]). However, my prototyping with a waveguide tuner proved unfruitful (Section 3.3).

³I haven't studied whether diffraction losses are greater than losses from dielectric dissipation. But it's more dominant than ohmic losses.

⁴Orpheus is an open resonator, so it *may* be possible to design mirrors that extend outside the volume of the magnet. But then the dipole magnet may distort the electric field.

Throughout this thesis, I've focused on the TEM_{00-18} mode. It's possible that other modes have substantial coupling to the axion or dark photon. The V_{eff} of these modes should be simulated. The electronics in the Orpheus DAQ have a pretty broad frequency range and can digitize around various modes. If multiple modes couple to dark matter, then collecting data around these modes would be an easy way to increase the scanned mass range of a data-taking run.

7.2.2 Mechanical Optimizations

One can also think about how to redesign the cavity's mechanical design if one were to start from scratch. I would redesign it so that the tuning is less dependent on the flexures. The flexures reduce friction while the plates move vertically, but they also allow for misalignments. The combination of springs with the correct stiffness constant and correct length of bearings mitigates tilting while tuning. But that solution isn't viable when the dielectrics are placed at every other half-wavelength instead of every quarter of a wavelength. A trained mechanical engineer with cryogenic experience would be beneficial in redesigning the cavity. But I believe each plate should have three points of contact. Perhaps one could modify the current design so that it uses three guiding rods and three scissor jacks. Each of the thrust bearings should also have 2° misalignment capability.

The cavity has to be better thermally sunk to itself, such that each mirror and dielectric plate are close to the same temperature.

The gear ratio can also be increased to reduce the backlash.

7.2.3 Scanning More Dark Matter Parameter Space

Even without the mentioned optimizations, the same detector can be used to scan through more parameter space. The limiting factors are that the VNA's upper limit is 18 GHz, the WR62 waveguides have a frequency range between 12.4 GHz and 18 GHz, and the Image Reject Mixer has a range between 14 GHz and 18 GHz. Without changing anything, the electronics are capable of searching for axions between 14 GHz and 18 GHz at possibly similar

sensitivities. We can then scan between 12.4 GHz and 14 GHz if we swap out the Image Reject Mixer.

Scanning from 12.4 GHz to 18 GHz with the same detector would require several runs with different dielectric thicknesses and mirror curvatures. It's possible that going to lower frequencies would increase sensitivity to the axion since lower frequencies need longer cavities, leading to an increased V_{eff} . However, it's not clear to me how decreasing the frequency would change the quality factor or cavity coupling coefficient.

If we swap out dielectric thicknesses and mirror curvatures at the appropriate frequency ranges, the Orpheus detector may be used to scan axion masses at photon couplings of $g_{a\gamma\gamma} \sim 10^{-12}$ for axion masses between $45 \mu\text{eV}$ and $75 \mu\text{eV}$, as shown in Figure 7.3. All that is required is additional funding and a graduate student lifecycle.

7.3 How to reach KSVZ and DFSZ

In natural units

$$P_a = \frac{g_{a\gamma\gamma}^2}{m_a} \rho_a B_0^2 V_{eff} Q_L \frac{\beta}{\beta + 1}$$

$$g_{a\gamma\gamma} = \frac{1}{B_0} \sqrt{\frac{\beta + 1}{\beta} \frac{\text{SNR} \times \Delta f T_n m_a}{\rho_a V_{eff} Q_L}} \left(\frac{1}{\Delta f t} \right)^{1/4}$$

I will use the projected results in Figure 7.1 as a baseline. To reach KSVZ sensitivity, V_{eff} and Q_L need to double and the noise temperature needs to reduce ten-fold ($T_n \sim 1 \text{ K}$). That would require the optimizations I mentioned in this chapter, cooling the cavity with a dilution refrigerator, quantum noise limited amplifiers, and technological advances in winding superconducting dipole magnets. Except for the dipole magnet, all this seems tractable with current technology.

From the KSVZ estimates, one could reach DFSZ by increasing V_{eff} more. So increase V_{eff} ten-fold instead of two-fold. Sub quantum-noise-limited amplification with photon counters would also be viable in this regime.

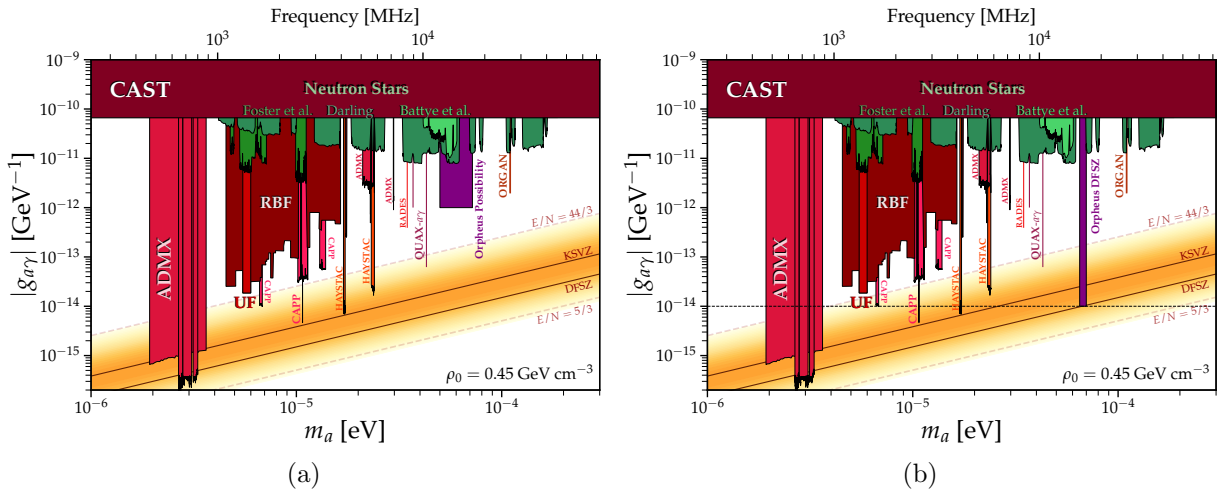


Figure 7.3: (a) It is likely that we can achieve sensitivity with the same insert and electronics by swapping out dielectrics of different thicknesses and mirrors with different radius of curvature. (b) DFSZ can be reached with a similar cavity design at similar frequency ranges if the magnetic field was increased to 10 T, the effective volume increased by a factor of 10, the quality factor doubled, and the noise temperature reduced by a factor of 10. Figures adapted from [16].

Chapter 8

CONCLUSION

Orpheus is a pathfinder experiment that demonstrates that a dielectrically loaded structure can be implemented and tuned in a cryogenic environment to search for axions and dark photons.

In the past three years, I spearheaded this experiment from the design¹ phase through construction, commissioning, and data analysis. I excluded parameter space that is comparable to that of larger, more well-funded experiments. I figured out how to simulate the dielectrically-loaded cavity reliably on a desktop, which was difficult because the cavity is electrically large and has a high mode density. I demonstrated that the cavity had promising sensitivity to wavelike dark matter, and that the mode of interest had good tuning capabilities and forgiving tolerances to mechanical perturbations. I commissioned this experiment at 4.2 K with a low noise HFET receiver.

Between $65.5 \mu\text{eV}$ and $69.3 \mu\text{eV}$, the excluded dark photon kinetic mixing angle is $\chi_{90\%} \sim 1 \times 10^{-12}$ for the polarized dark photon case and $\chi_{90\%} \sim 1 \times 10^{-13}$ for the unpolarized dark photon case. With modest alterations and several iterations, the same apparatus may be used to exclude larger parameter space from $45 \mu\text{eV}$ to $80 \mu\text{eV}$ with similar sensitivities.

This is a pathfinder experiment with a limited scope. It's my hope that the hard-earned knowledge gained from Orpheus proves useful to other experiments that implement dielectric haloscopes, such as MADMAX and DBAS.

This work was supported by the U.S. Department of Energy through Grants No. DE-SC0011665 and by the Heising-Simons Foundation.

¹Rich Ottens designed the vast majority of the mechanical structure and cryogenic infrastructure before I joined ADMX.

BIBLIOGRAPHY

- [1] P.A. Zyla et al. “Review of Particle Physics”. In: *PTEP* 2020.8 (2020), p. 083C01. DOI: [10.1093/ptep/ptaa104](https://doi.org/10.1093/ptep/ptaa104).
- [2] Vera C. Rubin et al. “Rotational properties of 23 SB galaxies”. In: *Astrophys. J.* 261 (1982), p. 439. DOI: [10.1086/160355](https://doi.org/10.1086/160355).
- [3] K. G. Begeman, A. H. Broeils, and R. H. Sanders. “Extended rotation curves of spiral galaxies: dark haloes and modified dynamics”. In: *Monthly Notices of the Royal Astronomical Society* 249.3 (Apr. 1991), pp. 523–537. ISSN: 0035-8711. DOI: [10.1093/mnras/249.3.523](https://doi.org/10.1093/mnras/249.3.523). eprint: <https://academic.oup.com/mnras/article-pdf/249/3/523/18160929/mnras249-0523.pdf>. URL: <https://doi.org/10.1093/mnras/249.3.523>.
- [4] A. N. Taylor et al. “Gravitational Lens Magnification and the Mass of Abell 1689”. In: *The Astrophysical Journal* 501.2 (July 1998), pp. 539–553. DOI: [10.1086/305827](https://doi.org/10.1086/305827). URL: <http://dx.doi.org/10.1086/305827>.
- [5] Priyamvada Natarajan et al. “Mapping substructure in the HST Frontier Fields cluster lenses and in cosmological simulations”. In: *Monthly Notices of the Royal Astronomical Society* 468.2 (Feb. 2017), pp. 1962–1980. ISSN: 0035-8711. DOI: [10.1093/mnras/stw3385](https://doi.org/10.1093/mnras/stw3385). eprint: <https://academic.oup.com/mnras/article-pdf/468/2/1962/11210742/stw3385.pdf>. URL: <https://doi.org/10.1093/mnras/stw3385>.
- [6] M. Markevitch et al. “Direct Constraints on the Dark Matter Self-Interaction Cross Section from the Merging Galaxy Cluster 1E 0657-56”. In: 606.2 (May 2004), pp. 819–824. DOI: [10.1086/383178](https://doi.org/10.1086/383178). URL: <https://doi.org/10.1086/383178>.

- [7] N. Aghanim et al. “Planck 2018 results”. In: *Astronomy & Astrophysics* 641 (Sept. 2020), A6. ISSN: 1432-0746. DOI: [10.1051/0004-6361/201833910](https://doi.org/10.1051/0004-6361/201833910). URL: <http://dx.doi.org/10.1051/0004-6361/201833910>.
- [8] Lam Hui. *Wave Dark Matter*. 2021. arXiv: [2101.11735](https://arxiv.org/abs/2101.11735) [[astro-ph.CO](#)].
- [9] C. Abel et al. “Measurement of the Permanent Electric Dipole Moment of the Neutron”. In: *Physical Review Letters* 124.8 (Feb. 2020). ISSN: 1079-7114. DOI: [10.1103/PhysRevLett.124.081803](https://doi.org/10.1103/PhysRevLett.124.081803). URL: <http://dx.doi.org/10.1103/PhysRevLett.124.081803>.
- [10] Wikimedia Commons. *File:NEDM P&T violation.png* — *Wikimedia Commons, the free media repository*. [Online; accessed 7-June-2020]. 2018. URL: https://commons.wikimedia.org/w/index.php?title=File:NEDM_P%5C%26T_violation.png%5C&oldid=299337593.
- [11] Michael Dine. *Solutions of the Strong CP Problem: A Scorecard*. 2015.
- [12] Edward J. Daw. *A Search for Halo Axions*. 2018. arXiv: [1807.09369](https://arxiv.org/abs/1807.09369) [[hep-ex](#)].
- [13] Stefan Knirck. *MADMAX: A new way of probing QCD AxionDark Matter with a Dielectric Haloscope*.
- [14] Yvonne Y.Y. Wong. “Axions”. In: *COSMO2011*. 2011.
- [15] Elisa G. M. Ferreira. *Ultra-Light Dark Matter*. 2020. arXiv: [2005.03254](https://arxiv.org/abs/2005.03254) [[astro-ph.CO](#)].
- [16] Ciaran O’Hare. *cajohare/AxionLimits: AxionLimits*. Version v1.0. July 2020. DOI: [10.5281/zenodo.3932430](https://doi.org/10.5281/zenodo.3932430). URL: <https://doi.org/10.5281/zenodo.3932430>.
- [17] R. Essig et al. *Dark Sectors and New, Light, Weakly-Coupled Particles*. 2013. arXiv: [1311.0029](https://arxiv.org/abs/1311.0029) [[hep-ph](#)].
- [18] Sumita Ghosh et al. “Searching for Dark Photons with Existing Haloscope Data”. In: (Apr. 2021). arXiv: [2104.09334](https://arxiv.org/abs/2104.09334) [[hep-ph](#)].

- [19] Andrea Caputo et al. *Dark photon limits: a cookbook*. 2021. arXiv: [2105.04565](https://arxiv.org/abs/2105.04565) [[hep-ph](#)].
- [20] Bob Holdom. “Searching for ϵ charges and a new U(1)”. In: *Physics Letters B* 178.1 (1986), pp. 65–70. ISSN: 0370-2693. DOI: [https://doi.org/10.1016/0370-2693\(86\)90470-3](https://doi.org/10.1016/0370-2693(86)90470-3). URL: <https://www.sciencedirect.com/science/article/pii/0370269386904703>.
- [21] Bob Holdom. “Two U(1)’s and ϵ charge shifts”. In: *Physics Letters B* 166.2 (1986), pp. 196–198. ISSN: 0370-2693. DOI: [https://doi.org/10.1016/0370-2693\(86\)91377-8](https://doi.org/10.1016/0370-2693(86)91377-8). URL: <https://www.sciencedirect.com/science/article/pii/0370269386913778>.
- [22] Maxim Pospelov, Adam Ritz, and Mikhail Voloshin. “Bosonic super-WIMPs as keV-scale dark matter”. In: *Phys. Rev. D* 78 (11 Dec. 2008), p. 115012. DOI: [10.1103/PhysRevD.78.115012](https://doi.org/10.1103/PhysRevD.78.115012). URL: <https://link.aps.org/doi/10.1103/PhysRevD.78.115012>.
- [23] Peter W. Graham, Jeremy Mardon, and Surjeet Rajendran. “Vector dark matter from inflationary fluctuations”. In: *Phys. Rev. D* 93 (10 May 2016), p. 103520. DOI: [10.1103/PhysRevD.93.103520](https://doi.org/10.1103/PhysRevD.93.103520). URL: <https://link.aps.org/doi/10.1103/PhysRevD.93.103520>.
- [24] P. A. R. Ade et al. “Planck2015 results”. In: *Astronomy & Astrophysics* 594 (Sept. 2016), A20. ISSN: 1432-0746. DOI: [10.1051/0004-6361/201525898](https://doi.org/10.1051/0004-6361/201525898). URL: <http://dx.doi.org/10.1051/0004-6361/201525898>.
- [25] Samuel D. McDermott and Samuel J. Witte. “Cosmological evolution of light dark photon dark matter”. In: *Phys. Rev. D* 101 (6 Mar. 2020), p. 063030. DOI: [10.1103/PhysRevD.101.063030](https://doi.org/10.1103/PhysRevD.101.063030). URL: <https://link.aps.org/doi/10.1103/PhysRevD.101.063030>.

- [26] P. Sikivie. “Experimental Tests of the ”Invisible” Axion”. In: *Phys. Rev. Lett.* 51 (16 Oct. 1983), pp. 1415–1417. DOI: [10.1103/PhysRevLett.51.1415](https://doi.org/10.1103/PhysRevLett.51.1415). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.51.1415>.
- [27] Alexander J. Millar et al. “Dielectric haloscopes to search for axion dark matter: theoretical foundations”. In: 2017.01 (Jan. 2017), pp. 061–061. DOI: [10.1088/1475-7516/2017/01/061](https://doi.org/10.1088/1475-7516/2017/01/061). URL: <https://doi.org/10.1088/1475-7516/2017/01/061>.
- [28] Jonathan Ouellet and Zachary Bogorad. “Solutions to axion electrodynamics in various geometries”. In: *Phys. Rev. D* 99 (5 Mar. 2019), p. 055010. DOI: [10.1103/PhysRevD.99.055010](https://doi.org/10.1103/PhysRevD.99.055010). URL: <https://link.aps.org/doi/10.1103/PhysRevD.99.055010>.
- [29] John David Jackson. *Classical electrodynamics*. 3rd ed. New York, NY: Wiley, 1999. ISBN: 9780471309321. URL: <http://cdsweb.cern.ch/record/490457>.
- [30] David M. Pozar. *Microwave Engineering*. Hoboken, New Jersey: John Wiley & Sons, Inc, 2012.
- [31] Benjamin M. Brubaker. *First results from the HAYSTAC axion search*. 2018. arXiv: [1801.00835](https://arxiv.org/abs/1801.00835) [astro-ph.CO].
- [32] N. Du et al. “Search for Invisible Axion Dark Matter with the Axion Dark Matter Experiment”. In: *Phys. Rev. Lett.* 120 (15 Apr. 2018), p. 151301. DOI: [10.1103/PhysRevLett.120.151301](https://doi.org/10.1103/PhysRevLett.120.151301). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.120.151301>.
- [33] T. Braine et al. “Extended Search for the Invisible Axion with the Axion Dark Matter Experiment”. In: *Phys. Rev. Lett.* 124 (10 Mar. 2020), p. 101303. DOI: [10.1103/PhysRevLett.124.101303](https://doi.org/10.1103/PhysRevLett.124.101303). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.124.101303>.
- [34] ADMX Collaboration et al. *Search for ”Invisible” Axion Dark Matter in the 3.3 – 4.2 μeV Mass Range*. 2021. arXiv: [2110.06096](https://arxiv.org/abs/2110.06096) [hep-ex].

- [35] Gray Rybka et al. “Search for dark matter axions with the Orpheus experiment”. In: *Phys. Rev. D* 91 (1 Jan. 2015), p. 011701. DOI: [10.1103/PhysRevD.91.011701](https://doi.org/10.1103/PhysRevD.91.011701). URL: <https://link.aps.org/doi/10.1103/PhysRevD.91.011701>.
- [36] P. Brun et al. “A new experimental approach to probe QCD axion dark matter in the mass range above $40 \mu\text{eV}$ ”. In: *The European Physical Journal C* 79.3 (Mar. 2019), p. 186. ISSN: 1434-6052. DOI: [10.1140/epjc/s10052-019-6683-x](https://doi.org/10.1140/epjc/s10052-019-6683-x). URL: <https://doi.org/10.1140/epjc/s10052-019-6683-x>.
- [37] Masha Baryakhtar, Junwu Huang, and Robert Lasenby. “Axion and hidden photon dark matter detection with multilayer optical haloscopes”. In: *Phys. Rev. D* 98 (3 Aug. 2018), p. 035006. DOI: [10.1103/PhysRevD.98.035006](https://doi.org/10.1103/PhysRevD.98.035006). URL: <https://link.aps.org/doi/10.1103/PhysRevD.98.035006>.
- [38] Aaron P. Quiskamp et al. “Dielectric-Boosted Sensitivity to Cylindrical Azimuthally Varying Transverse-Magnetic Resonant Modes in an Axion Haloscope”. In: *Phys. Rev. Applied* 14 (4 Oct. 2020), p. 044051. DOI: [10.1103/PhysRevApplied.14.044051](https://doi.org/10.1103/PhysRevApplied.14.044051). URL: <https://link.aps.org/doi/10.1103/PhysRevApplied.14.044051>.
- [39] A. Álvarez Melcón et al. *First results of the CAST-RADES haloscope search for axions at $34.67 \mu\text{eV}$* . 2021. arXiv: [2104.13798](https://arxiv.org/abs/2104.13798) [hep-ex].
- [40] H. Kogelnik and T. Li. “Laser beams and resonators”. In: *Proceedings of the IEEE* 54.10 (1966), pp. 1312–1329. DOI: [10.1109/PROC.1966.5119](https://doi.org/10.1109/PROC.1966.5119).
- [41] R N Clarke and C B Rosenberg. “Fabry-Perot and open resonators at microwave and millimetre wave frequencies, 2-300 GHz”. In: *Journal of Physics E: Scientific Instruments* 15.1 (Jan. 1982), pp. 9–24. DOI: [10.1088/0022-3735/15/1/002](https://doi.org/10.1088/0022-3735/15/1/002). URL: <https://doi.org/10.1088/0022-3735/15/1/002>.
- [42] D P Dunseith et al. “A high quality, efficiently coupled microwave cavity for trapping cold molecules”. In: *Journal of Physics B: Atomic, Molecular and Optical Physics*

- 48.4 (Jan. 2015), p. 045001. DOI: [10.1088/0953-4075/48/4/045001](https://doi.org/10.1088/0953-4075/48/4/045001). URL: <https://doi.org/10.1088/0953-4075/48/4/045001>.
- [43] R.K. Mongia and R.K. Arora. “Equivalent circuit parameters of an aperture coupled open resonator cavity”. In: *IEEE Transactions on Microwave Theory and Techniques* 41.8 (1993), pp. 1245–1250. DOI: [10.1109/22.241661](https://doi.org/10.1109/22.241661).
- [44] Paul J. Petersan and Steven M. Anlage. “Measurement of resonant frequency and quality factor of microwave resonators: Comparison of methods”. In: *Journal of Applied Physics* 84.6 (1998), pp. 3392–3402. DOI: [10.1063/1.368498](https://doi.org/10.1063/1.368498). eprint: <https://doi.org/10.1063/1.368498>. URL: <https://doi.org/10.1063/1.368498>.
- [45] M.P. Weidman. *A Method for Designing Multi-Screw Waveguide Tuners*. NBS Technical Publications, 1970.
- [46] Kirt Blattenberger. *Mixer Image Frequency*. URL: <https://www.rfcafe.com/references/electrical/image-frequency.htm>.
- [47] Kirt Blattenberger. *Cascaded Noise Figure & Noise Temperature*. URL: <http://www.rfcafe.com/references/electrical/noise-figure.htm>.
- [48] Dripline Maintainers. *Welcome to Dripline!* 2020. URL: <https://driplineorg.github.io/develop/index.html> (visited on 10/22/2021).
- [49] S. Asztalos et al. “Large-scale microwave cavity search for dark-matter axions”. In: *Phys. Rev. D* 64 (9 Oct. 2001), p. 092003. DOI: [10.1103/PhysRevD.64.092003](https://link.aps.org/doi/10.1103/PhysRevD.64.092003). URL: <https://link.aps.org/doi/10.1103/PhysRevD.64.092003>.
- [50] B. M. Brubaker et al. “HAYSTAC axion search analysis procedure”. In: *Phys. Rev. D* 96 (12 Dec. 2017), p. 123008. DOI: [10.1103/PhysRevD.96.123008](https://link.aps.org/doi/10.1103/PhysRevD.96.123008). URL: <https://link.aps.org/doi/10.1103/PhysRevD.96.123008>.
- [51] C. Bartram et al. “Axion dark matter experiment: Run 1B analysis details”. In: *Phys. Rev. D* 103 (3 Feb. 2021), p. 032002. DOI: [10.1103/PhysRevD.103.032002](https://link.aps.org/doi/10.1103/PhysRevD.103.032002). URL: <https://link.aps.org/doi/10.1103/PhysRevD.103.032002>.

- [52] Frederick Roland Fickett. “Electrical Properties of Materials and Their Measurement at Low Temperatures”. In: 2019.
- [53] Christian Boutan. “A Piezoelectrically Tuned RF-Cavity Search for Dark Matter Axions”. PhD thesis. University of Washington, Jan. 2017.
- [54] D. A. Palken et al. “Improved analysis framework for axion dark matter searches”. In: *Phys. Rev. D* 101 (12 June 2020), p. 123011. DOI: [10.1103/PhysRevD.101.123011](https://doi.org/10.1103/PhysRevD.101.123011). URL: <https://link.aps.org/doi/10.1103/PhysRevD.101.123011>.
- [55] Gary J. Feldman and Robert D. Cousins. “Unified approach to the classical statistical analysis of small signals”. In: *Phys. Rev. D* 57 (7 Apr. 1998), pp. 3873–3889. DOI: [10.1103/PhysRevD.57.3873](https://doi.org/10.1103/PhysRevD.57.3873). URL: <https://link.aps.org/doi/10.1103/PhysRevD.57.3873>.

Appendix A

REFLECTION FITTING

A.1 First pass at a reflection fit

The measured reflected power is modeled as

$$|\Gamma|^2 = C - \frac{\delta y}{1 + 4\Delta^2}; \quad \Delta \equiv Q \frac{f - f_0}{f_0} \quad (\text{A.1})$$

where Γ reflection coefficient, f_0 (Q) is the resonant frequency (quality factor) of the cavity, C is some constant offset, and δy is the depth of the Lorentzian. One can check that $\Gamma(f_0) = C - \delta y$ and $\Gamma(f_0 - \delta f/2) = C - \delta y/2$, where the bandwidth $\delta f = f_0/Q$. To obtain the resonant frequency and Q of the cavity, we fit Equation A.1 to the measured reflected power, as shown in Figure A.1. Non-linear fits are often sensitive to the initial guess of fit parameters, so the initial guesses should be close to the correct answer (aim for an order of magnitude). One can come up with reasonable guesses from abstracting features of the reflected power. Below is a list of methods to obtain the initial guesses.

1. $f_{0,\text{guess}}$: Find the frequency that corresponds to Γ_{\min}^2
2. C_{guess} : Take the median of the upper 2/3 of the $|\Gamma|^2$ values. The lower 1/3 probably corresponds to the reflection on resonance.
3. $\delta y_{\text{guess}} = C_{\text{guess}} - |\Gamma|_{\min}^2$
4. Q_{guess} : Find bandwidth δf by determining frequencies where $\Gamma = C_{\text{guess}} - \delta y_{\text{guess}}/2$, i.e., find the full-width-half-max. $Q_{\text{guess}} = f_{0,\text{guess}}/\delta f$.

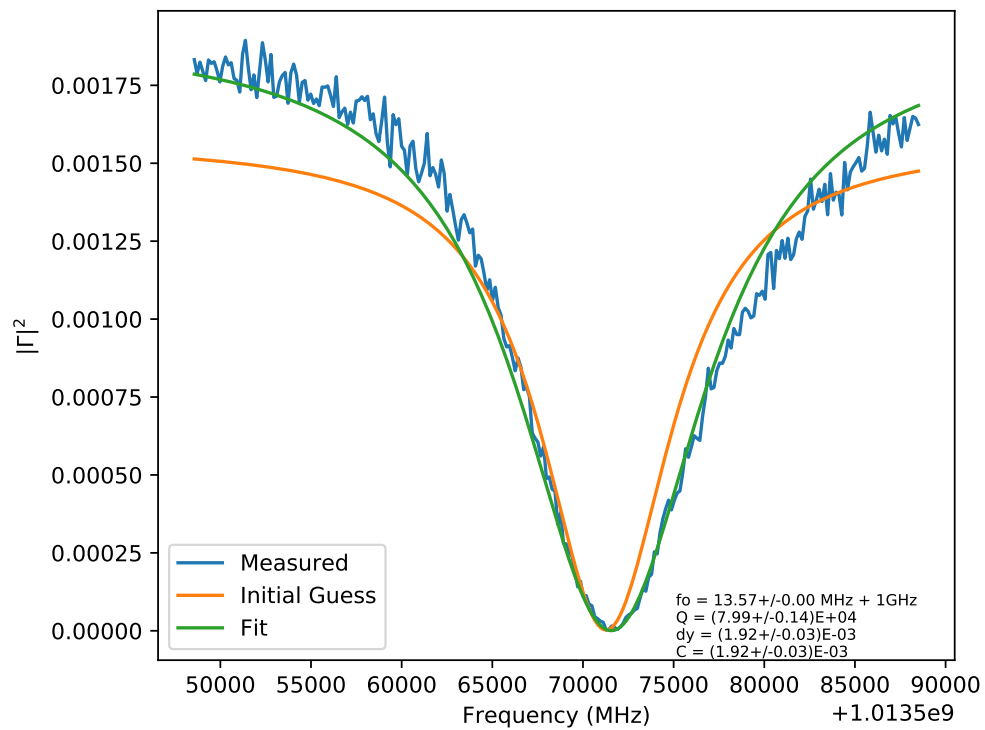


Figure A.1: The process of fitting a Lorentzian to the reflection coefficient to extract the resonant frequency and quality factor. The blue curve is the square of the measured reflected coefficient. The orange line is the initial guess of the Lorentzian. The green line is the curve fit, which tries to minimize the χ^2 .

Of course, if you already roughly know the Q of the system, then you can manually put in that guess.

Now let us extract the antenna coupling coefficient β . In later sections, I derive that the reflection coefficient of the cavity is $\Gamma_c = \frac{1-\beta}{1+\beta} + \frac{2i\Delta}{1+2i\Delta}$. Figure A.2 shows the phase response for an undercoupled cavity ($\beta < 1$), overcoupled cavity ($\beta > 1$), and critically-coupled cavity ($\beta = 1$). We can check the phase of the measured IQ data around resonance to determine if we are undercoupled or overcoupled. As one option, one can check if $\angle\Gamma(f_0^+) > \angle\Gamma(f_0^-)$ to check if the cavity is undercoupled.

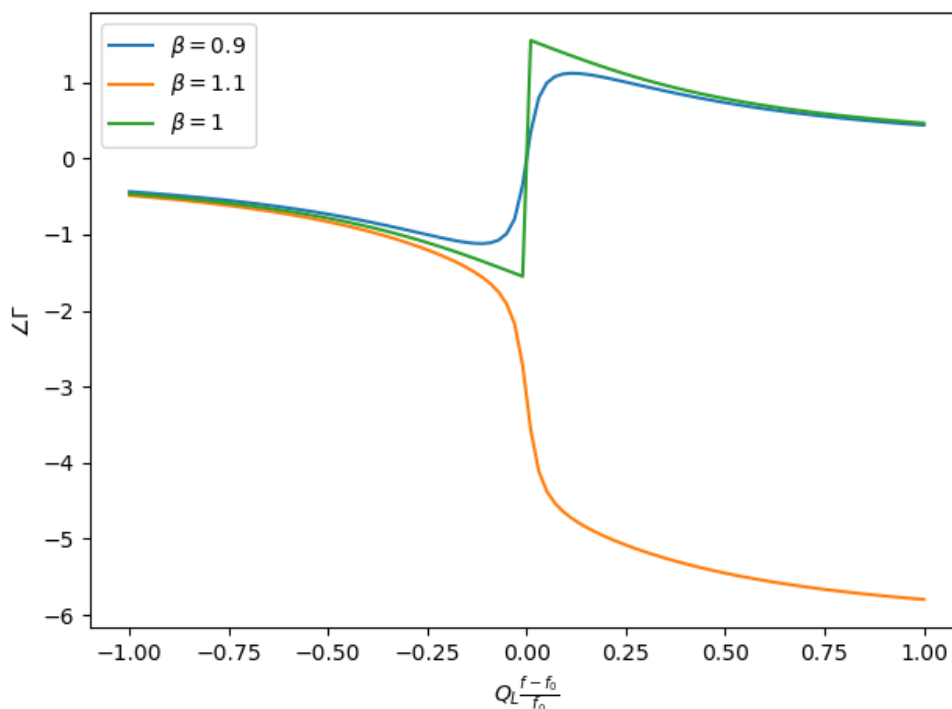


Figure A.2: The phase change of the reflection coefficient for the undercoupled case (blue), critically-coupled case (green), and overcoupled case (orange). The derivative on resonance is positive for the undercoupled case and negative for the overcoupled case.

We can calculate the coupling coefficient from the reflected power. It's best to use the value from the fitted function since the measured values at low power tend to be noisy.

$$\beta = \begin{cases} \frac{1-|\Gamma_c(f_0)|}{1+|\Gamma_c(f_0)|} & \text{if } \beta \leq 1 \\ \frac{1+|\Gamma_c(f_0)|}{1-|\Gamma_c(f_0)|} & \text{if } \beta \geq 1 \end{cases} \quad (\text{A.2})$$

But we do not measure the cavity reflection coefficient directly. The signal is amplified and then attenuated by cables. To estimate $|\Gamma_c(f_0)|$, we note that the cavity acts like a reflector off resonance, so $|\Gamma_c(\Delta)| \rightarrow 1$ for large Δ . After normalizing the reflected power to estimate $|\Gamma_c(\Delta)|^2$, we can calculate $|\Gamma_c(f_0)|$ using the fitted parameters $|\Gamma_c(f_0)| = \sqrt{1 - \frac{\delta y}{C}}$.

A proper chi-square minimization requires inputting accurate relative uncertainties of each datapoint in the reflected power. The uncertainty of each data point can be derived as $\sigma_P^2 = 4\sigma_V^2 P$, where σ_V is the uncertainty of the IQ voltages. σ_V depends on the VNA settings such as the number of averages and the IF bandwidth. σ_V can be estimated by measuring the standard deviation of the recorded IQ voltages from the VNA noise floor.

The derivation for the uncertainty in reflected power follows:

Proof.

$$P \propto V_I^2 + V_R^2$$

Let $Re \equiv V_R^2$ and $Im \equiv V_I^2$.

$$\begin{aligned} P &\propto R + I \\ \sigma_P^2 &= \sigma_R^2 + \sigma_I^2 \\ \left(\frac{\sigma_R^2}{R}\right)^2 &= \left(\frac{2\sigma_{V_R}}{V_R}\right)^2 \\ \sigma_R^2 &= (2V_R\sigma_{V_R})^2 \end{aligned}$$

Similarly

$$\begin{aligned}\sigma_I^2 &= (2V_I\sigma_{V_I})^2 \\ \sigma_P^2 &= (2V_R\sigma_{V_R})^2 + (2V_I\sigma_{V_I})^2 \\ \sigma_{V_R} &= \sigma_{V_I} \equiv \sigma_V \\ \sigma_P^2 &= (2\sigma_V)^2(V_R^2 + V_I^2) \\ \sigma_P^2 &= 4\sigma_V^2 P\end{aligned}$$

□

A.2 Deriving the reflection off of the cavity

In this section, we derive why the reflected power off a cavity follows a Lorentzian profile.

A cavity can be modeled as a series RLC circuit (Pozar 4th ed. Ch. 6), and the cavity impedance Z_c is

$$\begin{aligned}Z_c &= R + \frac{1}{i\omega C} + i\omega L \\ &= R + \frac{1 - \omega^2 LC}{i\omega C}\end{aligned}$$

For an RLC circuit, $\omega_o = \frac{1}{\sqrt{LC}}$. Define $\omega \equiv \omega_o - \delta\omega$. Then $\omega_o^2 \approx \omega_o^2 - 2\omega_o\delta\omega$.

$$\begin{aligned}Z_c - R &\approx \frac{1 - \frac{1}{\omega_o^2}(\omega_o^2 - 2\omega_o\delta\omega)}{i\omega_o C} \\ &= \frac{-2\delta\omega}{i\omega_o^2 C}\end{aligned}$$

Define $Q_o \equiv \frac{1}{\omega_o RC} = \frac{1}{R}\sqrt{\frac{L}{C}} = \frac{\omega_o L}{R}$.

$$Z_c = R \left(1 + 2i \frac{\delta\omega Q_o}{\omega_o} \right)$$

Let's connect the cavity to a load with impedance $Z_o = R_o$. Assume that the transmission line connecting the cavity to the load is impedance-matched to the load. Then the reflection coefficient, from Pozar 4th ed. Equation 2.35, is

$$\Gamma_c = \frac{Z_c - Z_o}{Z_c + Z_o}$$

Define an antenna coupling coefficient $\beta = \frac{Q_o}{Q_e} = \frac{Z_o}{R}$.

$$\begin{aligned} \frac{1}{Q_L} &= \frac{1}{Q_o} + \frac{1}{Q_e} \\ \frac{1}{Q_L} &= \frac{1}{Q_o} \left(1 + \frac{Q_o}{Q_e} \right) \\ Q_L &= \frac{Q_o}{1 + \beta} \end{aligned}$$

On resonance, $Z_c = R$ and the magnitude of the reflection coefficient depends on the coupling coefficient.

$$\begin{aligned} \Gamma_c(f_o) &= \frac{1 - \frac{Z_o}{R}}{1 + \frac{Z_o}{R}} \\ \Gamma_c(f_o) &= \frac{1 - \beta}{1 + \beta} \end{aligned} \tag{A.3}$$

Let's now solve for the reflection coefficient around the resonance.

$$\begin{aligned}
\Gamma_c &= \frac{Z_c - Z_o}{Z_c + Z_o} \\
&= \frac{1 + 2i \frac{\delta\omega Q_o}{\omega_o} - \frac{Z_o}{R}}{1 + 2i \frac{\delta\omega Q_o}{\omega_o} + \frac{Z_o}{R}} \\
&= \frac{1 + 2i \frac{\delta\omega Q_o}{\omega_o} - \beta}{1 + 2i \frac{\delta\omega Q_o}{\omega_o} + \beta} \\
&= \frac{1 + 2i \frac{\delta\omega Q_o}{\omega_o} - \beta}{1 + 2i \frac{\delta\omega Q_o}{\omega_o} + \beta} \times \left(\frac{1 + \beta}{1 + \beta} \right) \\
\Gamma_c &= \frac{\frac{1-\beta}{1+\beta} + 2i \frac{\delta\omega Q_L}{\omega_o}}{1 + 2i \frac{\delta\omega Q_L}{\omega_o}}
\end{aligned}$$

Define $\Delta \equiv Q_L \frac{\delta\omega}{\omega_o}$

$$\Gamma_c = \frac{\frac{1-\beta}{1+\beta} + 2i\Delta}{1 + 2i\Delta} \quad (\text{A.4})$$

The reflected power is the magnitude-squared of the reflection coefficient.

$$\begin{aligned}
|\Gamma_c|^2 &= \frac{\frac{1-\beta^2}{1+\beta} + 4\Delta^2}{1 + 4\Delta^2} \\
\frac{4\Delta^2}{1 + 4\Delta^2} &= 1 - \frac{1}{1 + 4\Delta^2} \\
|\Gamma_c|^2 &= 1 - \frac{4\beta}{(1+\beta)^2} \frac{1}{1 + 4\Delta^2}
\end{aligned}$$

Let $L(f) \equiv \frac{1}{1+4\Delta^2}$. This is the functional form of the Lorentzian. Then

$$|\Gamma_c|^2 = 1 - \frac{4\beta}{(\beta + 1)^2} L(f) \quad (\text{A.5})$$

A.3 Model of a resonator connected to directional coupler and amplifier

In the previous section, we modeled a cavity connected to a load, and the transmission line was impedance-matched to the load. More realistically, the transmission line is composed of many components with various sources of impedance mismatches.

Figure A.3 demonstrates the reflection measurement off of the Orpheus cavity. The VNA sends a signal to port 3 of the directional coupler. The signal travels to port 1 and reflects off the cavity. Then the signal travels from port 1 to port 2, and then through the amplifier and back to the VNA. However, the amplifier could be mismatched, allowing for standing waves to form in the transmission path. Furthermore, there might be additional crosstalk-like effects when one considers the directivity and isolation of the coupler.

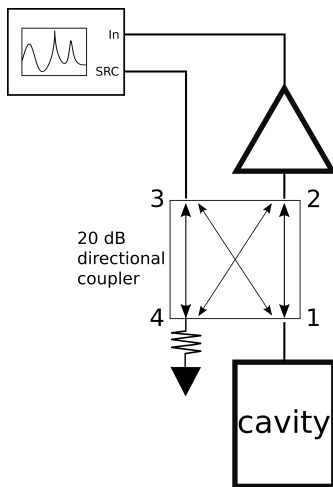


Figure A.3: The reflection measurement off of the Orpheus cavity. The VNA sends a signal to port 3 of the directional coupler. The signal travels to port 1 and reflects off the cavity. Then the signal travels from port 1 to port 2, and then through the amplifier and back to the VNA.

To simplify the analysis, I've simplified the reflection model in Figure A.3 to Figure A.4. Our VNA measures $\frac{V_{out}}{V_{in}}$. Let V_1 be the voltage reflected off of the cavity heading towards the amplifier, and let Γ_a be the reflection coefficient of the amplifier. Furthermore, let V_x represent the crosstalk-like effects of the directional coupler that are abstracted away.

$V_{out} \propto V_{in}$, so let's solve for V_{in}

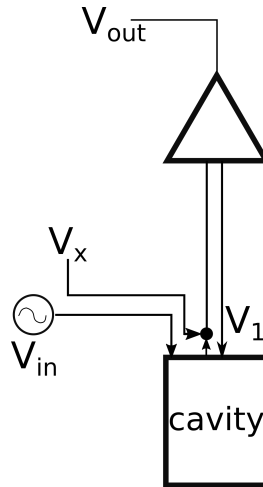


Figure A.4: A simplified mathematical model for a reflection measurement. A signal is injected into the cavity. V_1 is the signal reflected off the cavity heading towards the direction of the amplifier. V_1 reaches the amplifier and is partially reflected back to the cavity. That signal is then reflected again off the cavity and is added to V_1 . There is an additional V_x term that represents the complications from the injected signal interacting with the four different ports in the directional coupler.

$$V_1 = V_{in}\Gamma_c + V_1e^{i\theta}\Gamma_c\Gamma_a + V_x$$

$$\frac{V_1}{V_{in}} = \frac{\Gamma_c + \frac{V_x}{V_{in}}}{1 - e^{i\theta}\Gamma_c\Gamma_a}$$

The denominator is reminiscent of the mismatched generator problem in Pozar 4th ed. 2.6. Let's say $\frac{V_x}{V_{in}} = e^{i\phi}\chi$ and assume that χ is constant across the resonance of the cavity.

Now, let's assume that the amplifier is matched well enough that reflections are minimal and set $\Gamma_c = 0$

$$\frac{V_1}{V_{in}} \approx \Gamma_c + e^{i\phi} \chi$$

We want to see how this additional term perturbs the reflected power away from the Lorentzian.

$$\left| \frac{V_1}{V_{in}} \right|^2 = |\Gamma_c|^2 + |\chi|^2 + \chi(\Gamma_c e^{-i\phi} + \Gamma_c^* e^{i\phi})$$

By using Equation A.4 and expanding, one can show

$$\Gamma_c e^{-i\phi} + \Gamma_c^* e^{i\phi} = 2 \cos \phi - \frac{4\beta}{1+\beta} \cos \phi L(f) + \frac{8\beta}{1+\beta} \sin \phi \Delta L(f) \quad (\text{A.6})$$

Using Equation A.6, we get

$$\begin{aligned} \left| \frac{V_1}{V_{in}} \right|^2 &= 1 - \frac{4\beta}{(\beta+1)^2} L(f) + |\chi|^2 + 2\chi \cos \phi - \frac{4\chi\beta}{1+\beta} \cos \phi L(f) + 8\chi \sin \phi \frac{\beta}{1+\beta} \Delta L(f) \\ &= 1 + |\chi|^2 + 2\chi \cos \phi - 4 \left(\frac{\beta}{(\beta+1)^2} + \frac{\chi\beta}{1+\beta} \cos \phi \right) L(f) + 8\chi \sin \phi \frac{\beta}{1+\beta} \Delta L(f) \end{aligned} \quad (\text{A.7})$$

Let $C \equiv 1 - 2\chi^2 \cos \phi + \chi^2$, $D \equiv \chi \cos \phi$, $F \equiv 8\chi \sin \phi$.

$$\left| \frac{V_1}{V_{in}} \right|^2 = C - 4 \left(\frac{\beta}{(\beta+1)^2} + D \frac{\beta}{1+\beta} \right) L(f) + F \frac{\beta}{1+\beta} \Delta L(f)$$

Finally,

$$\left| \frac{V_{out}}{V_{in}} \right|^2 = A \left(C - 4 \left(\frac{\beta}{(\beta+1)^2} + D \frac{\beta}{1+\beta} \right) L(f) + F \frac{\beta}{1+\beta} \Delta L(f) \right) \quad (\text{A.8})$$

where A represents the amplification of the amplifier and the attenuation from the amplifier to the VNA.

If we had not set $\Gamma_a = 0$, then Equation A.8 becomes

$$\left| \frac{V_{out}}{V_{in}} \right|^2 = A \left(\frac{C - 4 \left(\frac{\beta}{(\beta+1)^2} + D \frac{\beta}{1+\beta} \right) L(f) + F \frac{\beta}{1+\beta} \Delta L(f)}{1 + \Gamma_a (2 \cos \theta - 4 \frac{\beta}{\beta+1} (\cos \theta + \Delta \sin \theta)) + \Gamma_a^2 \left(1 - \frac{4\beta}{(\beta+1)^2} L(f) \right)} \right)$$

Equation A.8 is a promising candidate to fit reflected powers that deviate from the standard Lorentzian shape. This model fits Orpheus data better than a Lorentzian (Figure A.5) and more accurately extracts the Q parameter, as evident by the fact that the reflection Q_L matches the transmission Q_L (Figure A.6). Unfortunately, there is a large degeneracy between β , χ , and ϕ . The resulting elements in the covariance matrix are huge. So Equation A.8 hasn't proved useful for extracting β . One may be able to make progress by constraining χ to be appropriately small.

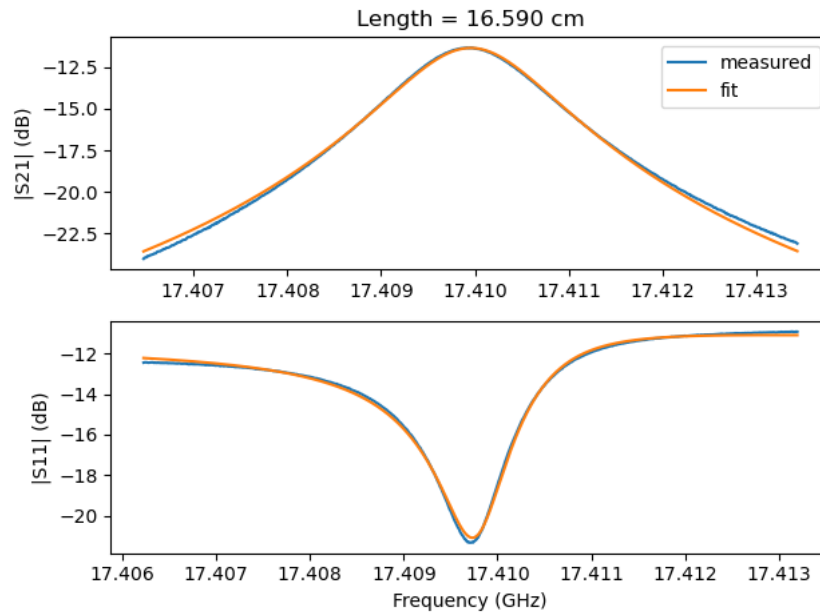


Figure A.5: The fitted transmitted (standard Lorentzian fit) and reflected power (perturbed Lorentzian fit, Equation A.8).

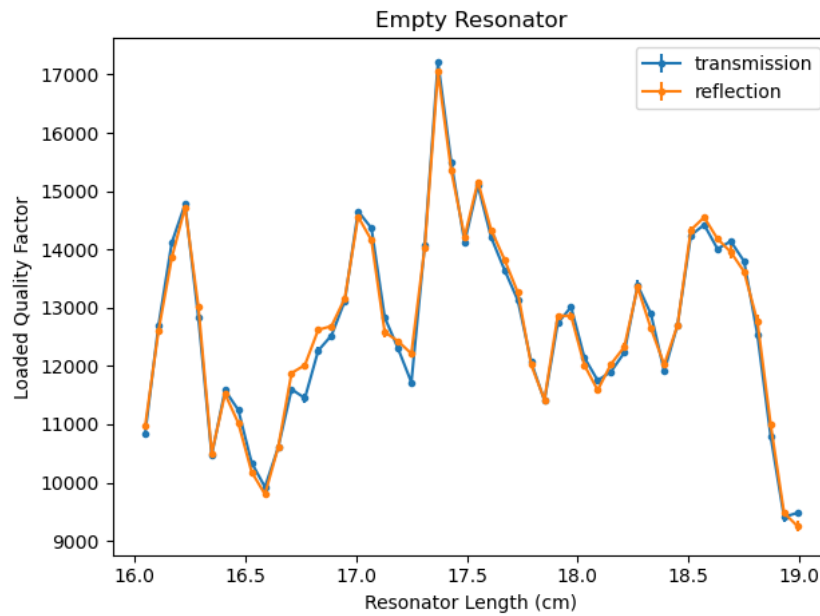


Figure A.6: The fitted loaded quality factors for transmission (standard Lorentzian fit) and reflection (perturbed Lorentzian fit, Equation A.8).

A.4 Axion coupling sensitivity to uncertainty in antenna coupling coefficient

A resonators sensitivity to the axion is related to the antenna coupling coefficient

$$g_{a\gamma\gamma} \propto \left(\frac{\beta + 1}{\beta} \right)^{1/2}$$

The uncertainty in $g_{a\gamma\gamma}$ is related to the uncertainty in β by

$$\Delta g_{a\gamma\gamma} = \left| \frac{\partial g_{a\gamma\gamma}}{\partial \beta} \right| \Delta \beta$$

$$\frac{\partial g_{a\gamma\gamma}}{\partial \beta} = -\frac{1}{2\beta^2 \sqrt{\frac{\beta+1}{\beta}}}$$

$\frac{\partial g_{a\gamma\gamma}}{\partial \beta}$ is plotted in Figure A.7. This suggests that near critical-coupling, the uncertainty in axion sensitivity is insensitive to the uncertainty in antenna coupling. Then it becomes very sensitive when we are weakly coupled.

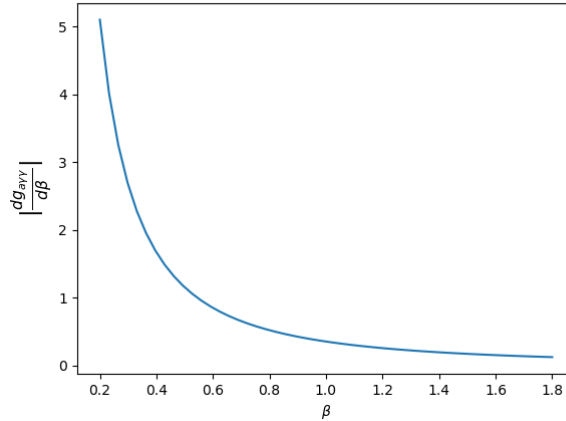
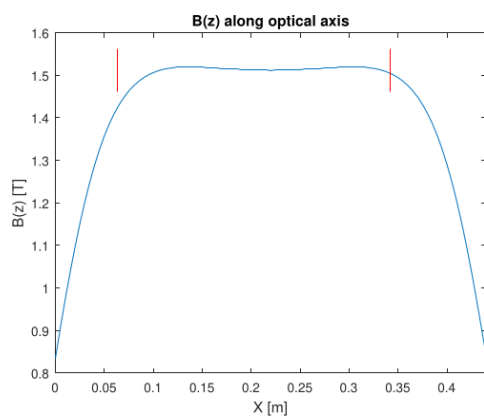


Figure A.7: The derivative of the axion-photon coupling constant vs antenna coupling coefficient. This relationship determines how the uncertainty in β affects the uncertainty in $g_{a\gamma\gamma}$.

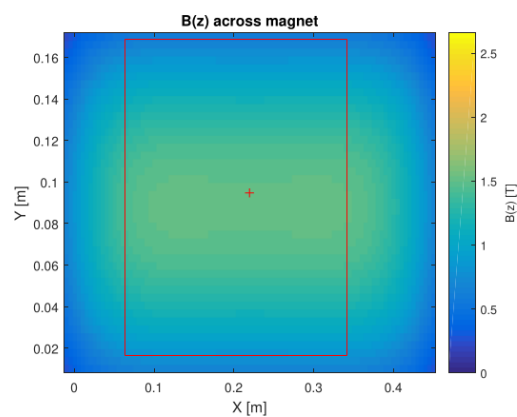
A.5 Magnet Design

Seth Kimes and Richard Ottens designed the magnet. The planned dipole magnet has dimensions $44 \times 18.5 \times 18$ cm³. The NbTi superconducting is from Supercon, Inc. The model is 56S53, the diameter is 0.30 mm bare and 0.33 mm with the insulation. The critical current is 125 A at 3 T. Each dipole side will have 57 layers of winding, and each layer will have 57 turns. The dipole magnet will run 103 A to produce a 1.5 T field inside the detector volume. The magnetic field profile is shown in Figure A.8. The force between the wires of the dipole magnet, depending on the direction, is $\mathcal{O}(2000$ lbs) of force. Seth Kimes designed the vacuum vessel to withstand that force with a certain safety factor.

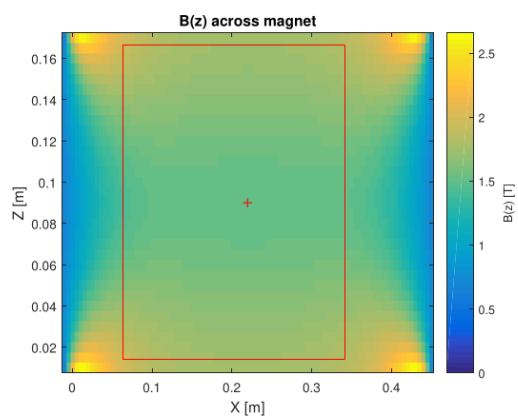
We are currently in the process of winding the superconducting dipole magnet. James Sinnis spent the Summer of 2021 practicing the winding and potting technique, and the results in Figure A.9 look satisfactory. We are now waiting for cryogenic epoxy to arrive, but the supply chains are currently in chaos.



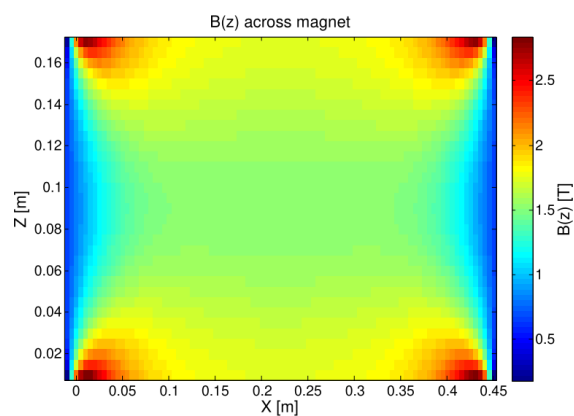
(a) Along the axis of the cavity



(b) XY Plane

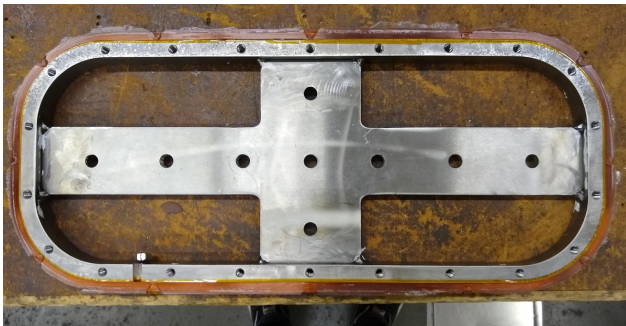


(c) XZ Plane

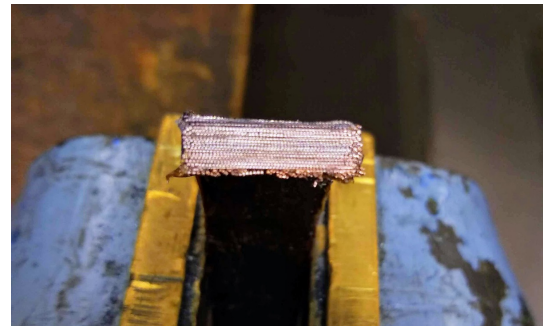


(d) YZ Plane

Figure A.8: Calculated magnetic field of the planned Orpheus magnet. For this calculation, Z is the axis of the magnet, X is along the long side of the dipole magnet, and Y is along the short side of the dipole magnet.



(a)



(b)

Figure A.9: Practicing Orpheus magnet winding with copper wire.