

# A Piezoelectrically Tuned RF-Cavity Search for Dark Matter Axions

Christian Boutan

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

Doctor of Philosophy

University of Washington

2017

Reading Committee:

Gray Rybka, Chair

Leslie Rosenberg

Ann Nelson

Program Authorized to Offer Degree:  
Physics

©Copyright 2017

Christian Boutan

University of Washington

## Abstract

A Piezoelectrically Tuned RF-Cavity Search for Dark Matter Axions

Christian Boutan

Chair of the Supervisory Committee:

Dr. Gray Rybka  
Physics

The Axion is a well motivated hypothetical elementary particle that must exist in nature if the strong CP problem of QCD is explained by the spontaneous breaking of a Peccei-Quinn symmetry. Not only would the discovery of the axion solve deep issues in QCD, an axion with a mass of  $\mu\text{eV} - \text{meV}$  could account for most or all of the missing mass in our galaxy and finally reveal the composition of dark matter. The Axion Dark Matter experiment (ADMX) seeks to resolve these two critical problems in physics by looking for the resonant conversion of dark-matter axions to microwave photons in a strong magnetic field. Utilizing state of the art electronics and dilution refrigerator cryogenics, ADMX is the world's leading search for axions - able to discover or rule out even the most pessimistically coupled QCD axions. With multi- $TM_{0n0}$  functionality and with the commissioning of the new high-frequency experiment called the "Sidecar" experiment, ADMX is also sensitive to a wide range of plausible axion masses. Here I motivate axions as ideal dark matter candidates, review techniques for detecting them and give a detailed description of the ADMX experiment. I discuss my contributions to the construction of the ADMX dual-channel receiver, which is the most sensitive microwave receiver on Earth. I describe the data acquisition, data taking and real-time analysis software. The primary focus of this work, however, is the ADMX Sidecar experiment which is a miniature axion haloscope that fits inside of the ADMX insert and has the capability of

searching for axion masses between  $16\mu\text{eV}$  -  $24\mu\text{eV}$  on the  $TM_{010}$  and  $26.4$  -  $30\mu\text{eV}$  on the  $TM_{020}$  mode. I discuss analysis of the Sidecar data and exclude axion-to-two-photon coupling  $g_{a\gamma\gamma} < 6 \times 10^{-12} \text{GeV}^{-1}$  over a mass range of  $3\mu\text{eV}$  ( $\Delta f \sim 708$  MHz) from  $21.05$  -  $23.98\mu\text{eV}$  for axions that compose 100% of dark matter. Over a narrow subsection of this range,  $22.89$  -  $22.95\mu\text{eV}$  ( $\sim 15$  MHz) I set a stricter limit  $g_{a\gamma\gamma} < 10^{-12} \text{GeV}^{-1}$ .

## TABLE OF CONTENTS

	Page
List of Figures . . . . .	v
Chapter 1: Introduction . . . . .	1
Chapter 2: Theory . . . . .	4
2.1 Axions from Quantum Chromodynamics . . . . .	4
2.1.1 Strong CP Problem . . . . .	5
2.1.2 Peccei-Quinn Solution . . . . .	6
2.1.3 Axion Models, Mass and Couplings . . . . .	9
2.2 Dark Matter and Cosmology . . . . .	11
2.2.1 Cosmology before Dark Matter . . . . .	11
2.2.2 Early Evidence . . . . .	13
2.2.3 Rotation Curves of Spiral Galaxies . . . . .	13
2.2.4 Gravitational Lensing and Galaxy Cluster Collisions . . . . .	14
2.2.5 Primordial Nucleosynthesis . . . . .	15
2.2.6 Matter-Radiation Fluctuations . . . . .	17
2.2.7 Cosmic Microwave Background . . . . .	18
2.2.8 Baryon Acoustic Oscillations . . . . .	20
2.2.9 The Temperature of Dark Matter . . . . .	21
2.2.10 The $\Lambda$ CDM Model of our Universe . . . . .	21
2.2.11 Solutions to the Mystery of Dark Matter . . . . .	23
2.3 Dark Matter Axions . . . . .	24
2.3.1 Axion Production . . . . .	24
2.3.2 Galactic Evolution and Signal Linewidth . . . . .	26
2.3.3 Astrophysical Constraints . . . . .	28

2.3.4	Experimental Searches . . . . .	30
2.4	The Axion Haloscope . . . . .	31
2.4.1	Detection Strategy . . . . .	32
2.4.2	Expected Signal to Noise . . . . .	34
2.4.3	Designing an Experiment . . . . .	36
Chapter 3: The Axion Dark Matter eXperiment . . . . .		41
3.1	Resources, Infrastructure and Apparatus . . . . .	41
3.1.1	History . . . . .	41
3.1.2	Collaborators . . . . .	42
3.1.3	Infrastructure . . . . .	43
3.1.4	Insert . . . . .	46
3.2	Dual-Channel Axion Search . . . . .	48
3.2.1	Cavity . . . . .	48
3.2.2	Tuning Rod Motion . . . . .	48
3.2.3	Antenna Coupling . . . . .	50
3.2.4	Cold Amplification . . . . .	51
3.2.5	Room Temperature Receiver . . . . .	54
3.3	Sidecar Axion Search . . . . .	56
3.3.1	Motivation . . . . .	56
3.3.2	Cavity . . . . .	58
3.3.3	Piezoelectric Tuning/Coupling . . . . .	58
3.3.4	Cold Electronics . . . . .	66
3.3.5	Receiver . . . . .	67
3.4	Calibration . . . . .	67
3.4.1	Noise Temperature . . . . .	69
3.4.2	Compression and Clipping . . . . .	71
3.4.3	Synthetic Axion Generator . . . . .	72
Chapter 4: Data-Taking: Navigating Cavity Mode Structure . . . . .		75
4.1	Cavity Basics . . . . .	75
4.1.1	Analytical Solutions . . . . .	76

4.1.2	Swept Measurements . . . . .	79
4.1.3	RF Simulations . . . . .	83
4.2	Mode Following . . . . .	85
4.2.1	Tuning . . . . .	85
4.2.2	Tracking . . . . .	87
4.2.3	Mapping . . . . .	89
4.2.4	Sidecar Structure . . . . .	94
4.3	Making a Run Plan . . . . .	96
4.3.1	Constraints and Goals . . . . .	96
4.3.2	Simultaneous Multi-Mode Tracking . . . . .	100
4.3.3	Candidate Re-scans . . . . .	101
4.4	Data-Taking . . . . .	103
4.4.1	DAQ . . . . .	103
4.4.2	Cadence . . . . .	104
4.4.3	Sidecar Data-Taking Script and Web Interface . . . . .	108
4.4.4	Navigating Mode-Crossings . . . . .	109
4.4.5	Adaptive Scan Rate . . . . .	110
4.4.6	Maintaining Antenna Coupling . . . . .	114
4.4.7	Data-Taking together with Real-Time Analysis . . . . .	116
Chapter 5:	Analysis . . . . .	118
5.1	2016 Sidecar Data Run . . . . .	118
5.1.1	Run Plan . . . . .	118
5.1.2	Timeline . . . . .	119
5.1.3	Performance . . . . .	119
5.1.4	Rescans . . . . .	120
5.2	Analysis Introduction . . . . .	120
5.2.1	Methodology . . . . .	120
5.2.2	Live Analysis vs Offline Analysis . . . . .	121
5.2.3	Pipeline Overview . . . . .	122
5.3	Single Scan Analysis . . . . .	124

5.3.1	Data Packaging . . . . .	125
5.3.2	Calculating Expected Axion Power . . . . .	125
5.3.3	Measured Power . . . . .	131
5.3.4	Weighting . . . . .	134
5.4	Quality Control . . . . .	135
5.4.1	Automated Cuts . . . . .	135
5.4.2	Manual Cuts . . . . .	136
5.5	Grand Analysis . . . . .	138
5.5.1	Co-adding Spectra . . . . .	138
5.5.2	Correlating bins . . . . .	140
5.5.3	Deriving grand spectra from weighted sums . . . . .	140
5.5.4	Calculating Excluded Power . . . . .	142
5.5.5	Setting a limit . . . . .	143
5.5.6	Confirming limit with software injected synthetic axions . . . . .	143
5.6	Results . . . . .	144
Chapter 6:	Conclusions . . . . .	148
Bibliography	. . . . .	151

## LIST OF FIGURES

Figure Number	Page
2.1 Neutron electric dipole moment . . . . .	7
2.2 $a \rightarrow \gamma\gamma$ Feynman diagram . . . . .	10
2.3 Evidence for dark matter: rotation curves . . . . .	14
2.4 Evidence for dark matter: lensing, galaxy collisions and big bang nucleosynthesis	16
2.5 Evidence for dark matter: CMB and BAOs . . . . .	19
2.6 Composition of our universe . . . . .	22
2.7 Romulus25 simulation of axion lineshape . . . . .	28
2.8 Astrophysical constraints on axion mass . . . . .	29
2.9 Axion search window . . . . .	31
2.10 The axion haloscope . . . . .	33
3.1 The ADMX collaboration . . . . .	43
3.2 The ADMX site . . . . .	44
3.3 The ADMX insert . . . . .	47
3.4 Main experiment: frequency tuning . . . . .	49
3.5 Main experiment: antenna coupling . . . . .	52
3.6 Simplified receiver . . . . .	53
3.7 ADMX dual-channel receiver rack . . . . .	55
3.8 ADMX dual-channel receiver schematic . . . . .	57
3.9 The Sidecar cavity . . . . .	59
3.10 Attocube stick-slip system . . . . .	61
3.11 Piezo motors mounted on the Sidecar cavity . . . . .	62
3.12 Sidecar experiment: frequency tuning . . . . .	63
3.13 Sidecar experiment: antenna coupling . . . . .	64
3.14 Sidecar switchbox schematic . . . . .	68

3.15	Sidecar receiver rack . . . . .	69
3.16	Sidecar receiver-chain schematic . . . . .	70
3.17	Synthetic Axion Generator (SAG) . . . . .	72
3.18	SAG calibration . . . . .	73
4.1	Resonant modes in a empty cavity . . . . .	78
4.2	Swept response of a empty cavity . . . . .	81
4.3	TM and TE resonant peaks in a swept response measurement . . . . .	82
4.4	Frequency shifting with tuning rod boundary conditions . . . . .	86
4.5	Sidecar $TM_{010}$ peak tuning . . . . .	87
4.6	Simple peak tracking algorithm . . . . .	88
4.7	Mode-map tutorial . . . . .	90
4.8	Mode mixing . . . . .	92
4.9	mode structure dependence on rod end-gap . . . . .	93
4.10	Mode splitting . . . . .	95
4.11	Analytical predictions, simulations and measurements of the Sidecar cavity .	97
4.12	Sidecar measured mode-map . . . . .	98
4.13	$TM_{010}$ and $TM_{020}$ frequency and form factor simulations . . . . .	101
4.14	ADMX data run plan for 2016/2017 . . . . .	102
4.15	The ADMX DAQ platform . . . . .	105
4.16	Simplified cadence of data-taking . . . . .	107
4.17	Sidecar web-interface: state of the experiment . . . . .	109
4.18	Sidecar web-interface: mode navigation . . . . .	111
4.19	Sidecar web-interface: successful navigation through mode-crossings . . . . .	112
4.20	Sidecar web-interface: tuning throttle . . . . .	113
4.21	Sidecar antenna couple versus antenna depth . . . . .	115
5.1	Conceptual analysis block diagram . . . . .	123
5.2	Raw data scan example . . . . .	126
5.3	Magnetic field in the vicinity of the Sidecar experiment . . . . .	128
5.4	Receiver response removal . . . . .	132
5.5	Histogram of corrected spectra . . . . .	134

5.6	Quality control: identifying mode-crossings . . . . .	137
5.7	Quality control: differentiating real and spurious temperature spikes . . . . .	138
5.8	Coadding spectra . . . . .	139
5.9	Data re-binning schemes . . . . .	141
5.10	Grand spectra . . . . .	145
5.11	Software injected signal calibrations . . . . .	146
5.12	Sidecar limits . . . . .	147
6.1	Future Sidecar scan rate . . . . .	150

## **ACKNOWLEDGMENTS**

I would like to thank my adviser, Gray Rybka for having my back and for making me a more savvy scientist. Thank you to Leslie Rosenberg for your real world outlook and for caring so much about the condition of my corneas and kidneys. They can be quite expensive on the black market. Thanks to both of you for daily fighting the bureaucracy that would otherwise make science impossible for the ADMX team. Thanks to Ann Nelson and John Caspian for reading through this document and providing valuable feedback.

## **DEDICATION**

I dedicate this work to the love of my life, Madeleine Ann Moore Boutan. Thank you for going on this journey with me and for putting up with my grad student salary for so long. Thank you for being my best friend and for being such a wonderful mother to our children.



## Chapter 1

### INTRODUCTION

Despite truly remarkable scientific advancements, achievements and discoveries at the end of the last century, several glaring unsolved mysteries in physics persist. One of the biggest open questions today is what constitutes more than 80% of the mass of the universe. Dubbed “dark matter”, this form of exotic matter is known to be primarily non-relativistic and interact very weakly with normal matter. Astrophysics provides compelling evidence for dark matter and yet its particle nature remains unknown, having thus far evaded detection in the laboratory.

While the concept of dark or “missing matter” is not difficult to communicate to a layperson, there is another, perhaps less appreciated mystery in physics that is equally vexing. It’s called the Strong CP Problem and it arises in quantum chromodynamics (QCD), the branch of physics devoted to explaining the interactions of quarks and gluons within atomic nuclei. Briefly, the puzzle (explained in chapter 2) is why do strong interactions respect CP symmetry, the product of charge conjugation (C) and Parity (P) symmetries, when there is an explicitly CP-violating term in the QCD lagrangian.

$$\mathcal{L}_{CP \text{ violating term}} \propto (\theta + \arg \det M) G_{\mu\nu}^a \tilde{G}^{\mu\nu a} \quad (1.1)$$

where  $G_{\mu\nu}^a$  is the gluon field strength tensor,  $\tilde{G}^{\mu\nu a}$  its dual, the total derivative  $G_{\mu\nu}^a \tilde{G}^{\mu\nu a}$  is the CP violating term,  $\theta$  describes the QCD vacuum,  $M$  is the quark mass matrix and the sum  $\theta + \arg \det M$  controls the amount of CP violation observed in nature. Precision experiments have found  $\theta + \arg \det M < 10^{-10}$  even though these two quantities have no reason or ability to coordinate a near perfect (or perfect) cancelation. As a consequence,

the impact of equation 1.1 on the dynamics of strong interactions is effectively completely suppressed and seemingly by accident. This is the Strong CP Problem and perhaps its most compelling solution, the Peccei-Quinn solution predicts the existence of a new pseudoscalar boson called the axion.

The axion is a very compelling hypothetical particle, not only because its existence would replace a QCD coincidence with a natural mechanism but also because axions, found to have the right mass, are ideal dark matter candidates and could account for some or all of the missing mass in our universe. The world's most sensitive search for these Dark-Matter Axions is being performed by the Axion Dark Matter eXperiment (ADMX) at the Center for Experimental Nuclear Physics and Astrophysics (CENPA) at the University of Washington in Seattle, Washington. Exploiting the inverse primakoff effect where  $a \rightarrow \gamma^* \gamma$  in a magnetic field, the experiment utilizes a high Q, tunable resonant cavity, submerged in a large magnetic field and looks for the resonant conversion of axions into microwave photons.

It has been my privilege to work with the small handful of professors, post-docs, grad students, undergraduates and engineers that make up the ADMX team. From 2011 to the present, my largest contributions to the main experiment were to 1) design, build and test the room temperature portion of the ADMX dual channel receiver, 2) automate the DAQ and write the software responsible for data taking, 3) build a synthetic axion generator used for calibration, that outputs a signal similar to the expected axion line shape, 4) write and implement a real-time analysis program that monitors the data taking, processes incoming spectra and looks for evidence of an axion live and, finally, to 5) determine the data run plans for the experiment based on science goals and mechanical/RF constraints. I was also very involved in the plumbing for the  ${}^4\text{He}$  recovery system and in the construction of the ADMX site infrastructure. In 2014 and 2016 I oversaw data taking and demonstrated simultaneous dual-channel functionality for the experiment. Between September 2014 and June 2016 we upgraded our experiment to accommodate a dilution refrigerator. During this

time I designed and built a new axion experiment called the ADMX-Sidecar. Sidecar is a miniature axion detector that bolts directly to the primary ADMX cavity and is designed to be sensitive to  $16\mu\text{eV} - 30\mu\text{eV}$  axions. It shares the magnetic field and cryogenics with the rest of ADMX but is otherwise a stand alone experiment. In August of 2016, the sidecar and the main experiment took eight weeks of low-field commissioning data at 200mK with a working dilution refrigerator. At the end of these eight weeks, three days of higher field data was taken during a magnet ramping exercise.

In chapter 3, I discuss the hardware and the anatomy of the main ADMX experiment but put a heavier emphasis on the Sidecar experiment. In chapter 4, I then describe how the software brings the hardware to life to take data and track resonant modes. I detail the analysis of the Sidecar data in chapter 5, setting new limits on  $g_{a\gamma\gamma}$  and in chapter 6 I make conclusions and discuss future work.

## Chapter 2

# THEORY

In this chapter I give a brief history of Quantum Chromodynamics (QCD) and lay the minimal mathematical foundation needed to appreciate what appears to be a remarkable coincidence observed in the dynamics of quarks within atomic nuclei. I then introduce a natural mechanism that explains this coincidence and gives birth to a new elementary particle called the axion. I discuss properties of the axion, initial attempts to find it and the reasoning that went into deciding that it must be invisible or very light and weakly coupled. Shifting the focus from the physics of the very small to the physics of the very big, I review the current state of cosmology and address the dark matter problem. Finally I motivate the still undiscovered axion as an attractive dark matter candidate particle and describe methods by which it could be detected.

### **2.1 Axions from Quantum Chromodynamics**

The axion was conceived out of an elegant solution to a fine-tuning problem in Quantum Chromodynamics (QCD). QCD is a gauge theory of symmetry  $SO(3)$  that describes the strong interactions between quarks and gluons within hadrons. Riding on the heels of successes in Quantum Electrodynamics (QED), QCD has survived decades of scrutiny even though its non-abelian nature creates practical difficulties in calculating exact solutions. While it is a remarkable and robust framework, the theory has one source of mystery called the Strong CP Problem.

### 2.1.1 Strong CP Problem

Starting with a minimal form of the QCD Lagrangian which can be written

$$\mathcal{L}_{QCD} = \sum_{f=u,d,s,c,t,b} \bar{q}_f (i\gamma^\mu D_\mu - m_f) q_f - \frac{1}{4} G_{\mu\nu}^a G_a^{\mu\nu} \quad (2.1)$$

where  $f$  runs over quark flavors,  $q_s$  are quark fields,  $\gamma^\mu$  are the gamma matrices,  $D_\mu$  is the covariant derivative,  $m_f$  are the quark masses and  $G_{\mu\nu}^a$  is the gluon field strength tensor.

While this form of the Lagrangian is suitable for obtaining perturbative solutions, it is symmetric under  $U(1)_A$  global axial rotations which would lead to hadron parity doublets such as another proton with opposite parity. Such particles have not been observed in nature so this symmetry must be broken. In the same way that  $SU(2)$  was spontaneously broken to produce the pions  $\pi^-$ ,  $\pi^0$  and  $\pi^+$ , a spontaneously broken  $U(1)_A$  would create a new pseudo-Goldstone boson with the same quantum numbers as  $\eta$  but much lighter, with a mass on the order of the pion masses. Although  $U(1)_A$  must be broken, a particle matching this description has not been found. This is one manifestation of the “ $U(1)_A$  problem”. The  $U(1)_A$  problem is avoided if  $U(1)_A$  is anomalously broken by an axial current. The Adler-Bell-Jackiw (ABJ) anomaly solves the  $U(1)_A$  problem at the cost of introducing a new term to the QCD lagrangian

$$\mathcal{L}_{\bar{\theta}} = \frac{\bar{\theta}}{32\pi^2} G_{\mu\nu}^a \tilde{G}^{\mu\nu a}, \quad (2.2)$$

where the total derivative  $G_{\mu\nu}^a \tilde{G}^{\mu\nu a}$  explicitly violates CP (Charge Conjugation \* Parity) symmetry and  $\bar{\theta}$  is a parameter that controls the level of CP violation in strong interactions.  $\bar{\theta}$  is naively expected to be  $\mathcal{O}(1)$  but a non-zero  $\bar{\theta}$  leads to observable effects, such as the existence of electric dipole moments (EDMs) in neutral elementary particles with non-degenerate ground states like the neutron. Figure 2.1 shows how the separation of charge in a toy model of a neutron can create an electric dipole moment that violates P, CP and T

(time-reversal) symmetries. The most stringent constraint on the value of  $\bar{\theta}$  is achieved with precision measurements of the neutron electric dipole moment. The current bound on the nEDM has been found to be  $d < 2.9 \times 10^{-26}$  which puts an upper limit of  $10^{-10}$  on  $\bar{\theta}$ .

It seems as though our universe was given a special value of  $\bar{\theta}$  to perfectly suppress CP violation in strong interactions but  $\bar{\theta}$  is actually the sum of two unrelated quantities from different sectors in physics

$$\bar{\theta} = \theta + \arg \det M. \quad (2.3)$$

where  $\theta$  describes the QCD vacuum and  $M$  is the quark mass matrix. Both of these contributions are angles between 0 and  $2\pi$  and have no reason to add to zero. The strong CP problem is therefore the mystery of how  $\bar{\theta}$  came to be so tiny.

### 2.1.2 Peccei-Quinn Solution

One way to solve the strong CP problem would be to have a massless quark, such as the up quark that would cause equation 2.3 to be undefined and  $\theta$  unobservable. However, a massless quark seems to be in conflict with lattice QCD. Another approach to solving the strong CP problem is to have a theory where CP is spontaneously broken such that bare  $\theta$  vanishes and careful re-working of the quark mass matrix can potentially achieve  $\arg \det M = 0$  while maintaining a large CKM angle [2, 3]. To me, the most natural and elegant fix to the Strong-CP problem is the Peccei-Quinn solution which effectively turns  $\bar{\theta}$  into a dynamical variable that relaxes to zero.

In 1977, with perhaps the Higgs mechanism in mind, Roberto Peccei and Helen Quinn proposed a new global  $U(1)$  chiral symmetry that was spontaneously broken in the early universe. At temperatures ( $T$ ) above  $\Lambda_{QCD}$ , the  $U(1)_{PQ}$  potential has a classical “wine bottle” form

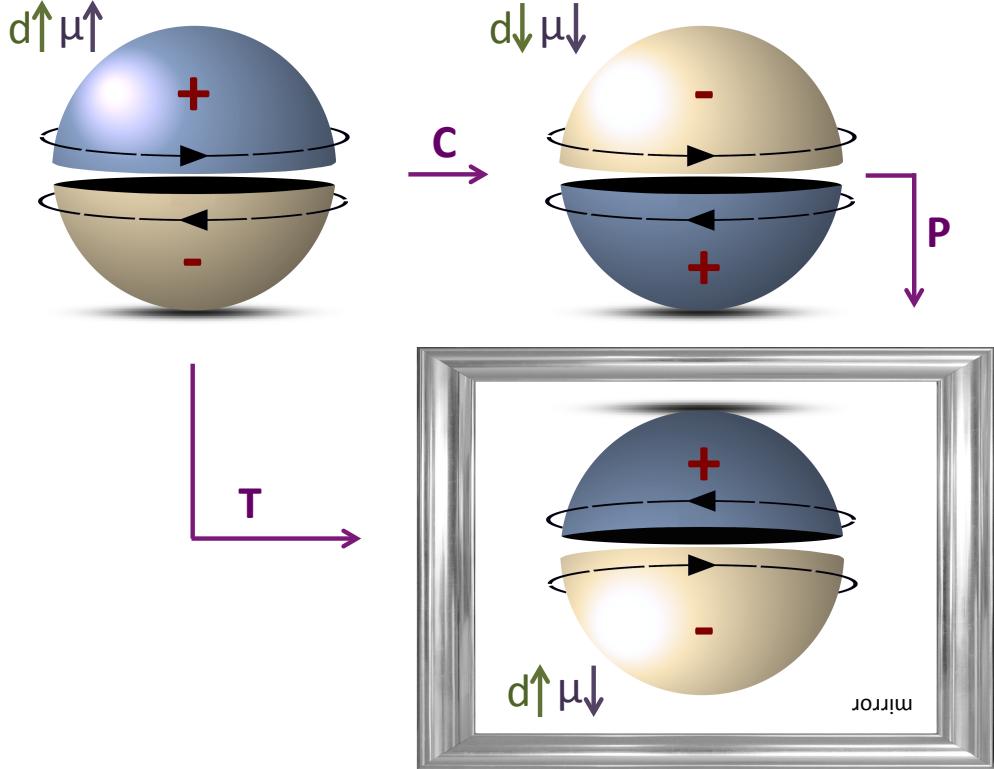


Figure 2.1: A toy model of a neutron with both a magnetic dipole moment ( $\mu$ ) and an electric dipole moment ( $d$ ). Under a charge-conjugation (C) transformation, the baryon number is odd but both  $d$  and  $\mu$  remain aligned. Under parity (P), the two charges switch positions and only the electric dipole moment flips direction. Therefore, the final state of the system is an anti-matter neutron that is physically distinguishable from the initial state. A neutron electric dipole moment (nEDM) violates both P and time-reversal (T) symmetries. By the CPT theorem, CP is violated as well. By measuring the nEDM, experiments have placed limits on the amount of CP violation in strong interactions. While this toy model of a neutron EDM (better termed “the coupling constant between spin and electric field”) provides an intuitive picture, it is an over simplification of the transformation properties of the complete Hamiltonian which govern the time evolution of the system [1].

$$V(\phi) \propto (|\phi|^2 - f_{PQ}^2/2)^2 \quad (2.4)$$

where  $\phi$  is a new a complex pseudo scalar field and  $f_{PQ}$  is the symmetry breaking scale.

Here, the theory promotes  $\bar{\theta}$  from a number to a dynamical variable as the  $\arg(\phi)$ . Above the symmetry breaking scale, the vacuum expectation of  $\phi$  is zero, all values of  $\bar{\theta}$  equally probable and strong CP violation would have gone unchecked.

As the universe cooled below  $f_{PQ}$ , the field was forced to the minima of the potential and developed a non-zero vacuum expectation value at  $\phi = f_{PQ}/\sqrt{2}$ . In the classical potential,  $\bar{\theta}$  could have continued to take on any value, but when  $T \sim \Lambda_{QCD}$ , quantum instanton effects tipped the wine bottle potential and the QCD Lagrangian acquired a new CP violating term

$$\mathcal{L}_{PQ} = C \frac{\phi}{f_{PQ}} \frac{g^2}{32\pi^2} F_{\mu\nu}^a \tilde{F}^{\mu\nu a} + \text{interactions} \quad (2.5)$$

Where  $C$  is a model dependent constant. By grouping like CP violating terms, the complex phase  $\bar{\theta}$  can be redefined to absorb the new term

$$\bar{\theta} = \theta + \arg \det M \rightarrow \theta + \arg \det M - C \frac{\phi}{f_{PQ}} \quad (2.6)$$

such that when the potential is minimized at

$$\langle a \rangle = \frac{f_{PQ}}{C} (\theta + \arg \det M) \quad (2.7)$$

$\bar{\theta}$  vanishes.

Peccei and Quinn had found a way to naturally force  $\bar{\theta}$  to zero and thereby solve the strong CP problem but luckily not without observable consequences. Weinberg [4] and Wilczek [5] realized that the introduction of a new spontaneously broken symmetry implied the existence of a new pseudo-Nambu Goldstone boson that would have acquired a mass when the field had begun to oscillate at the bottom of the potential. Weinberg called this hypothetical particle the “higglet” but Wilczek’s name won out and the “Axion” was born.

### 2.1.3 Axion Models, Mass and Couplings

The Peccei-Quinn-Weinberg-Wilczek (PQWW) Axion was initially associated with the weak scale and would have had a mass on the order of 10 - 100 keV. These heavy axions would have been seen in beam dump experiments ( $a \rightarrow \gamma\gamma$ ,  $a \rightarrow e^-e^+$ ), meson decays ( $J/\Psi \rightarrow \gamma a$ ,  $K^+ \rightarrow \pi^+ a$ ,  $\Upsilon \rightarrow \gamma a$ ) and nuclear de-excitations ( $N^* \rightarrow Na$ ). The failure to detect these heavy axions in accelerator experiments quickly ruled them out [6, 7]. However, the original PQWW model was modified to accommodate a very weakly coupled “invisible axion”. The one free parameter of the theory that must be confirmed experimentally is the symmetry breaking scale which is inversely proportional to the mass and couplings. By setting  $f_{PQ}$  to be very large, the axion can be made to be “invisible” (weakly coupled) with a mass given by

$$m_a = \frac{(m_u m_d)^{1/2}}{(m_u + m_d)} \frac{f_\pi}{(f_{PQ}/N)} m_\pi \sim 0.62 \text{eV} \frac{10^7 \text{GeV}}{(f_{PQ}/N)} \quad (2.8)$$

where  $m_u$  ( $m_d$ ) is the mass of the up (down) quark,  $f_\pi$  is the pion decay constant,  $N$  is the PQ color anomaly and  $m_\pi$  is the pion mass. The interaction relevant to this work is given by

$$\mathcal{L}_{a\gamma\gamma} = -g_{a\gamma\gamma} \phi_a \vec{E} \cdot \vec{B} \quad (2.9)$$

and describes the decay of an axion into two photons. Here  $\phi_a$  is the axion field,  $E$  is the electric field,  $B$  is the magnetic field and the axion-to-two-photon coupling,  $g_{a\gamma\gamma}$  is

$$g_{a\gamma\gamma} = \frac{\alpha g_\gamma}{\pi(f_{PQ}/N)} \quad (2.10)$$

where  $\alpha$  is the fine structure constant and  $g_\gamma$  is a model-dependent constant that can range roughly between .1 and 1. The value of  $g_\gamma$  comes from the fact that the axion, not being electrically charged, cannot couple directly to photons. Instead, the interaction requires an

anomalous fermionic loop shown in figure 2.2. Which fermions are allowed by the model to interact with the axion at tree level determine the strength of the conversion to photons. The two principal invisible axion models are the Kim [8], Shifman, Vainshtein, and Zakharov (KSVZ) [9] and Dine, Fischler, Srednicki [10], and Zhitnitsky [11] (DFSZ) models. The KSVZ “Hadronic” axion only couples to a heavy quark at tree level and not to leptons. In the DFSZ model, the axion is allowed to couple to both quarks and leptons. These two models, with coupling constants  $g_\gamma = -0.97$  (KSVZ) and  $g_\gamma = +0.36$  (DFSZ) essentially form the extremes of a fairly constrained theory. If one could build an experiment sensitive enough to detect the photons in the equation 2.9, the confusion surrounding the strong CP problem could have a definitive answer. Today, the invisible axion remains invisible but a second serious problem in physics makes the potential existence of the axion even more desirable. Before continuing the story of the axion (picked up in section 2.3) I discuss a very different problem in physics that could also be solved with the Peccei-Quinn mechanism.

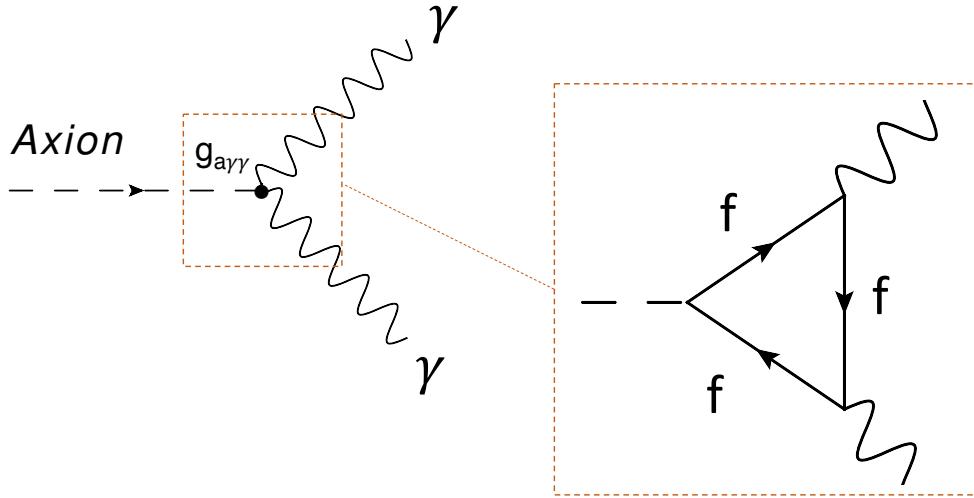


Figure 2.2: Feynman diagram for an axion decaying to two photons via a fermionic loop.

## 2.2 Dark Matter and Cosmology

At the time of writing, one of the most pressing open questions in physics concerns the nature of dark matter. The existence of the axion, thus far introduced in a particle physics context would have serious implications on cosmological scales and may well provide a key to unlocking the dark matter problem. The axion is compelling because it may explain a mystery seen on scales much smaller than the nucleus of an atom while at the same time explain another mystery seen on scales much larger than the size of a galaxy. This section reviews evidence for dark matter, summarizes the current state of cosmology, discusses solutions to the dark matter problem and motivates the axion as a compelling dark matter candidate.

### 2.2.1 Cosmology before Dark Matter

In the early 1900s, it was largely believed that the universe was static, flat and infinite. Einstein laid the foundation for understanding the relationship between energy and the curvature of spacetime in his theory of gravity, general relativity. While this extraordinary theory has not yet been successfully incorporated into the standard model of particle physics, it has been tested time and again and never found to be flawed. His field equations can be written in tensor form

$$R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R + g_{\mu\nu}\Lambda = \frac{8\pi G}{c^4}T_{\mu\nu} \quad (2.11)$$

where  $R_{\mu\nu}$  is the Ricci curvature tensor,  $g_{\mu\nu}$  is the metric tensor,  $R$  is the scalar curvature,  $G$  is the gravitational constant,  $T_{\mu\nu}$  is the stress energy tensor, and finally an added term  $\Lambda$  that Einstein called the cosmological constant. At the time, Einstein believed in the idea of a static universe. Realizing that such a universe with  $\Lambda = 0$  would inherently be unstable and collapse under its own gravity, he inserted this cosmological constant into the theory as a self-repulsive term.

The invention of a cosmological constant appeared at first to be a mistake as new evidence for a non stationary universe emerged. In the late 1910s, Slither and Wirtz determined that most galaxies (they thought were nebulae) were red-shifted and receding from us. In the 1920s, Hubble discovered that the universe is expanding and that two locally stationary galaxies will retreat from each other at a rate equal to the distance between them times the Hubble constant  $H_0$ . Around the same time, a Belgian catholic priest, Georges Lemaître proposed the idea of an expanding universe that started with an “explosion” [12].

The big bang model came about as a natural response to the realization that a growing universe must have been much smaller in its youth and perhaps originated from a single point, when space itself was infinitesimally small. If one assumes that the universe is homogeneous and isotropic on large scales, an assumption which seems naive but not unreasonable, the equations of motion for the scale factor or size of the universe  $a(t)$  is given by the Friedmann equations

$$H^2 = \left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G\rho + \Lambda}{3} - \frac{\kappa}{a^2} \quad (2.12)$$

$$\frac{\ddot{a}}{a} = \frac{\Lambda}{3} - \frac{4\pi G}{3}(\rho + 3p) \quad (2.13)$$

where  $H$  is the Hubble parameter (no longer assumed to be a constant  $H_0$ ),  $\rho$  and  $p$  are the density of and pressure exerted by matter + radiation and  $\kappa$  is the curvature of the universe. With a generic framework for modern cosmology in place, a specific description of our universe would emerge with precision measurements of  $H_0$ , the matter density  $\Omega_{matter}$ ,  $\Lambda$  and other parameters. Over the next hundred years there would be several remarkable discoveries to confirm this early cosmological foundation and a few surprises to keep things interesting.

### 2.2.2 Early Evidence

The first evidence for dark matter started emerging in the 1930s. Jan Oort studied the motions of stars outside the disk of the Milky-Way and noticed that there was a discrepancy between the amount of dynamically-estimated mass and luminous mass in our galaxy [13]. Zwicky and Smith measured the motions of galaxies in the Coma and Virgo clusters and determined that there was not enough visible matter present to keep the systems bound [14]. This early evidence pointed towards the existence of invisible matter dominating the mass of galaxies. However, the significance of these observations went mostly unnoticed at the time.

### 2.2.3 Rotation Curves of Spiral Galaxies

The realization that there was either a problem with missing mass or gravity became apparent in the 1970s when Vera Rubin, Gallagher, Faber et al started measuring the rotation curves of spiral galaxies using the 21 cm wavelength of neutral hydrogen's spin transition. A rotation curve is the measure of the speed of material in a galaxy versus its distance from the center of rotation. Assuming a disk-like distribution of matter and using basic newtonian physics, one expects a Keplerian curve given by

$$\begin{aligned} v &\propto r && \text{inside visible disk} \\ v &\propto 1/\sqrt{r} && \text{outside visible disk} \end{aligned} \tag{2.14}$$

where  $v$  is the velocity of matter some distance  $r$  from the center of the galaxy. Dozens of such rotation curves were cataloged and none of them matched this expected behavior [15]. Instead of falling off like  $1/\sqrt{r}$  (figure 2.3), the curves typically flattened out, indicating that at large enough distances the rate of rotation becomes independent of distance. This evidence suggested that either gravity behaves differently on larger scales or that the visible discs of galaxies are perhaps surrounded by huge diffuse clouds of invisible matter.

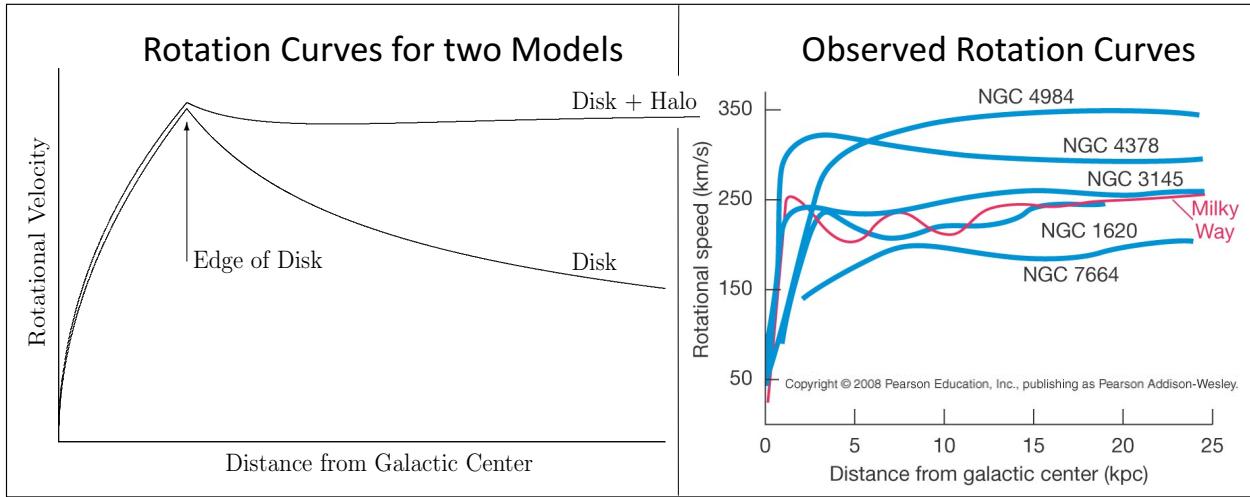


Figure 2.3: Rotation curves of spiral galaxies where the vertical axis is speed of rotation and the horizontal axis is distance from the galaxy center. Left: predicted Keplerian rotation curves with and without a halo of dark matter. Right: observed rotation curves for 6 galaxies (Credit: 2012 Pearson Education, Inc., publishing as Addison Wesley).

#### 2.2.4 Gravitational Lensing and Galaxy Cluster Collisions

Gravitational lensing provides a different probe of dark matter based not upon the dynamics of galaxies but on the fact that spacetime bends around mass. Light from distant sources and passing through massive galaxy clusters creates distorted images. These distortions can be used to infer the masses of galaxy clusters. Figure 2.4 shows light that has been gravitationally bent into a horseshoe shape around a galaxy in between the observer and a distant source.

The evidence presented thus far could possibly be explained by the presence of a lot of non-luminous, baryonic matter in galaxies or by modifying the way in which gravity interacts with matter. A powerful indication that neither of these explanations is sufficient to solve the missing mass problem can be seen in the collisions of galaxy clusters like the MACS or Bullet clusters. A composite image of the bullet cluster is shown in figure 2.4. By combining lensing

techniques and X-ray data, it can be shown that the baryonic matter behaves differently than the bulk of the matter. The figure shows that the dust in each cluster collide and come to a halt, emitting X-rays (pink) but that most of the mass in the clusters (blue) inferred from lensing passes through the baryonic pileup and doesn't interact.

### 2.2.5 *Primordial Nucleosynthesis*

Evidence from the early universe seems to confirm that dark matter must be a new form of exotic, non-interacting and invisible matter. In the first few minutes after the big bang, the universe was a dense plasma of quarks and gluons that cooled to form elementary particles and then the lightest of elements  $^1\text{H}$ ,  $^2\text{H}$ ,  $^3\text{He}$ ,  $^4\text{He}$ , and  $^7\text{Li}$ . The process by which the elements were created is called Primordial Nucleosynthesis or Big Bang Nucleosynthesis (BBN) and is very well understood. The amount of these elements produced depends heavily upon the initial baryon-to-photon ratio. By measuring the ratios of these light element abundances with respect to  $^1\text{H}$ , the initial baryon density can be constrained, thereby constraining the ratio of baryonic matter to total matter.

Deuterium ( $^2\text{H}$ ) provides a particularly useful constraint because it has a low binding energy and is very sensitive to the initial baryon density. By adding more baryonic matter to the early universe, one effectively increases the number of available projectiles to break the deuterium apart. Furthermore, deuterium isn't created in the burning of stars which means that the amount we see today was left over from the start of the universe. The abundance of deuterium is determined by measuring the absorption of Lyman alpha emission from quasars in high redshift clouds of gas.

Although they are not as sensitive as deuterium,  $^3\text{He}$ ,  $^4\text{He}$ , and  $^7\text{Li}$  also provide their own constraints. Barring a slight excess in  $^7\text{Li}$ , the outcome of light element abundance measurements predict that the baryon density is 4% of the critical density (the average density of matter required for a flat universe) or

## Lensing, Galaxy Collision and Nucleosynthesis Constraints

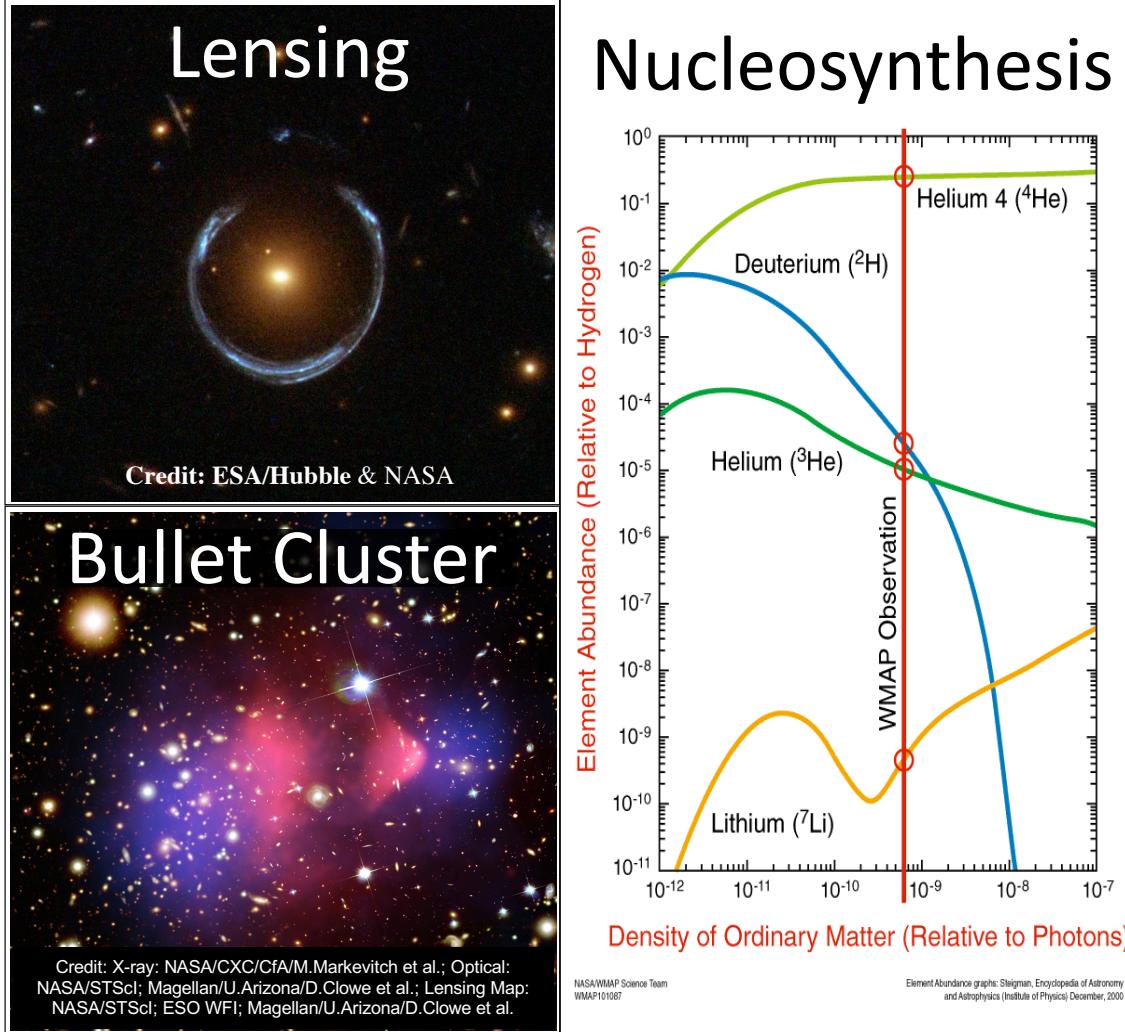


Figure 2.4: Constraints from Lensing, galaxy collisions and big bang nucleosynthesis. Upper left: light has been gravitationally bent around a galaxy in between the observer and a distant source. Lower left: a composite image of the bullet cluster. X-ray emission from the colliding gas is shown in pink, and the regions of highest density inferred from gravitational lensing are shown in blue. Right: light element abundances versus baryon density predicted by big bang nucleosynthesis from NASA/WMAP Science Team [16]. A baryon dominated early universe would evolve into a universe unlike the one observed.

$$\Omega_{baryon} h^2 = 0.0216 \pm 0.0020 \quad (2.15)$$

where  $h \equiv \frac{H}{100} \frac{Mpc}{km/sec}$  is the reduced 100 km/s Hubble constant. All of this to say that evidence left over from minutes after the big bang suggests that normal matter couldn't possibly be responsible for the total amount of missing mass.

### 2.2.6 Matter-Radiation Fluctuations

More evidence for dark matter is found imprinted in the distribution of galaxy formation and in the radiation left over from photons decoupled from matter. Unlike the current vacuum of space that prevents the propagation of pressure waves, for the first three hundred thousand years, the universe was a dense, hot fluid of mostly photons, free electrons, ionized hydrogen and presumably (based on evidence presented thus far) some amount of dark matter. The universe was too hot for neutral hydrogen to form because as soon as an electron would bind to a proton, it would immediately get unbound by a high energy photon. In the same way, photons were not able to freely stream through the medium because they were constantly colliding with and being absorbed and re-emitted by the dense soup of baryons. Much like the interior of the sun, matter was coupled strongly to photons and there was an associated radiation pressure that would push overly dense regions apart. Gravity, on the other hand, was working against the radiation pressure to pull matter together and cause it to clump up.

Feeling both of these forces, the baryons in this hot plasma caused regions in the early universe to oscillate between overly compressed and decompressed states. On all different length scales, baryons would first be driven apart by acoustic, pressure waves, and then in a rarified state gravitationally re-collapse. By contrast, dark matter, not coupling to electromagnetism, would not directly undergo acoustic oscillations but would only communicate gravitationally with itself and with baryonic matter. The outcome of which acoustic harmonics were dominant and which ones were suppressed depended upon the interplay between

both forms of matter and was therefore sensitive to the baryon to dark matter ratio.

When the universe had cooled enough to permit the stable combination of the proton plasma with free electrons to form neutral hydrogen, photons and matter became de-coupled in a process called “recombination” and photons were allowed to stream freely through the universe. As they left, they took with them information about acoustic oscillation temperature fluctuations and left behind matter density fluctuations that would later seed galaxy formation.

#### *2.2.7 Cosmic Microwave Background*

This relic blackbody radiation left over from when the universe was 380,000 years old comes from the “surface of last scattering” from all directions in the sky and is called the Cosmic Microwave Background (CMB). As the universe expanded, it cooled the once 3,000K background down to  $2.72548 \pm 0.00057$ K and has been redshifted to millimeter wavelengths. It was originally predicted by Ralph Alpher and Robert Herman in 1948 and then accidentally discovered by Penzias and Wilson in 1964.

Accurate measurements of the CMB by the WMAP [17] and Planck [18] missions provide some of the most important constraints on dark matter and other cosmological parameters. The upper left hand side of figure 2.5 shows the temperature fluctuations of the CMB after adjusting for the doppler shifts from the earth moving through the galaxy, the galaxy moving with respect to Andromeda, the motion of the local group with respect to the Virgo supercluster and so on. After making these corrections, what is left over is a remarkably isotropic (to one part in 100,000) map of the tiny quantum fluctuations and acoustic oscillations that have been stretched out to cosmological scales. By decomposing the imprint of temperature fluctuations into spherical harmonics to produce a power spectrum in the lower left corner of the figure, the characteristic size and intensity of the temperature irregularities and anisotropies can be used to constrain what kind of universe we live in. The relative size

and locations of each peak put constraints on the curvature of space, the amount of dark energy, the amount of total matter and the amount of baryonic matter. A fit to the CMB power spectrum determines the values for the baryon and matter density to be:

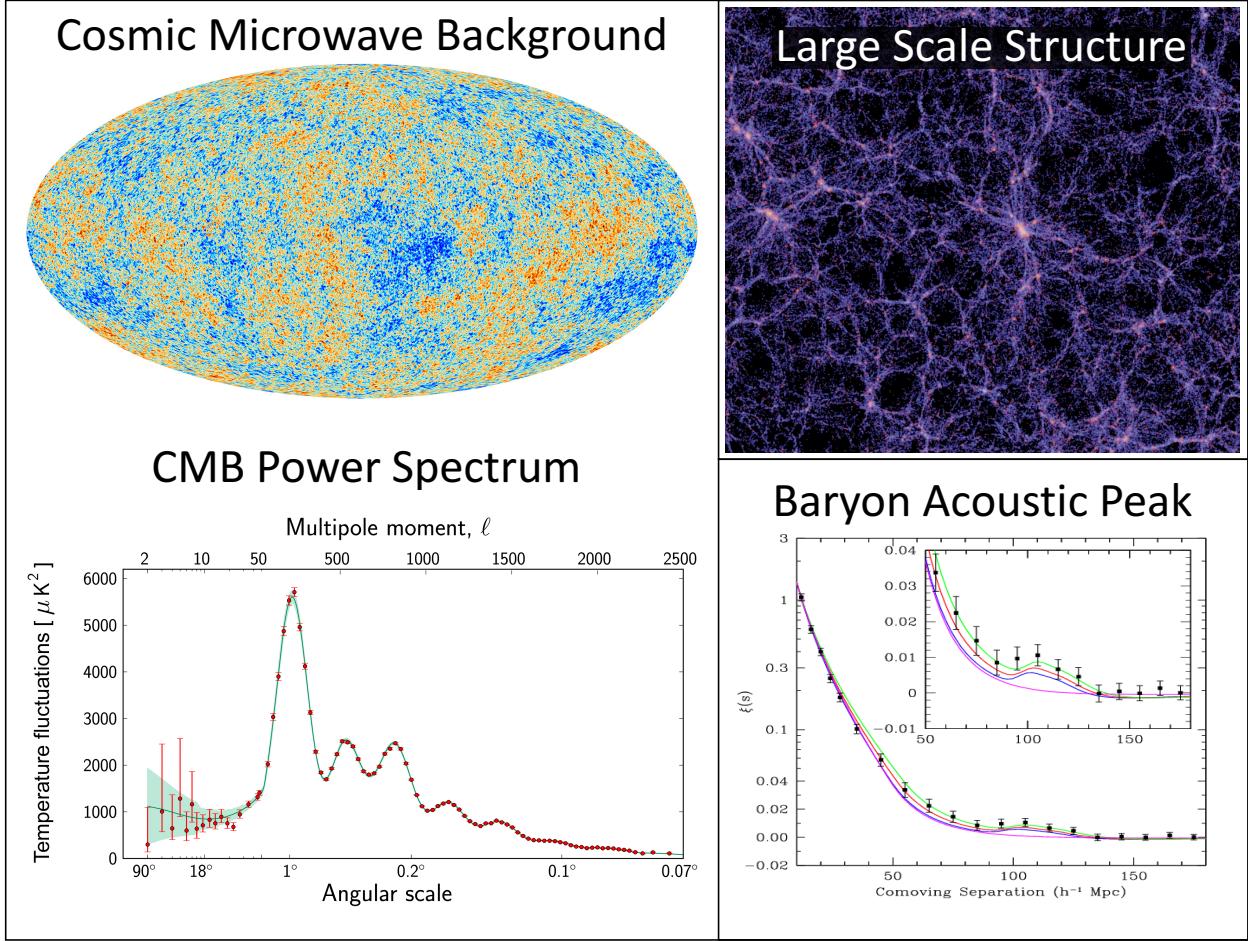


Figure 2.5: Left: the cosmic microwave background measured by Planck and its power spectrum decomposed into a multipole expansion [19]. Measurements (black dots with red error bars) are tied together with their theoretical fit (green) which constrains the amount of dark matter, baryonic matter and radiation in the early universe. Right: a simulation of large scale structure [20] over the baryon acoustic peak seen in the matter-matter correlation function [21].

$$\Omega_{baryon} h^2 = 0.0226 \pm .0006 \quad (2.16)$$

$$\Omega_{matter} h^2 = 0.135 \pm .009 \quad (2.17)$$

These numbers are constant with primordial nucleosynthesis results and show that baryons are only a small fraction of the total matter in the universe.

### 2.2.8 Baryon Acoustic Oscillations

Finally another constraint on the baryon density comes from the matter distribution itself left over from baryon acoustic oscillations. After recombination, there began a period called the “dark ages” when the universe was dominated by neutral hydrogen that neither emitted light or interacted with light from the surface of last scattering. Eventually (approximately 1 billion years after the big bang) slight over-densities from matter density fluctuations seeded galaxy formation by gravitationally pulling in enough mass to light the first stars and re-ionize the universe. This is known as the epoch of re-ionization, the signal from which is currently being searched for by large interferometric radio telescope arrays (such as the Murchison Widefield Array). Today, the characteristic length scales between galaxies has been set by density fluctuations in the early universe corresponding to the first dominant peak in the CMB power spectrum. The power spectrum of the galaxy-galaxy correlation function (lower right hand side of figure 2.5) constrain the matter density to be

$$\Omega_{matter} h^2 = 0.130 \pm 0.010 \quad (2.18)$$

which means that measurements from the Baryon Acoustic Oscillations and the Cosmic Microwave Background and Big Bang Nucleosynthesis all agree upon the amount of total matter in the universe and upon how much of that matter is normal matter described by the standard model of particle physics.

### 2.2.9 The Temperature of Dark Matter

Constraints on whether or not dark matter is “hot” (relativistic) or “cold” (non-relativistic) come from redshift based sky surveys and simulations of large scale structure. It is generally accepted that dark matter can’t be predominately hot because fast moving, relativistic dark matter particles would have washed out small scale density fluctuations needed to prompt the infall of matter to form the first galaxies. Hot dark matter would have therefore led to the “top-down” formation of smaller structures from the eventual fragmentation of superclusters that would have proceeded galaxies [22]. Observations however support a “bottom-up” cold dark matter (CDM) driven structure formation, where small structures formed first and then clumped into larger ones [23]. Simulations confirm that the large scale structure we see today cannot be duplicated with hot dark matter or with baryonic matter but that it must be driven by mostly-cold dark matter. While most theories favor the CDM model, some theories promote some combination of either hot and cold or even warm dark matter to explain certain features of the milky-way.

### 2.2.10 The $\Lambda$ CDM Model of our Universe

The Lambda Cold Dark Matter ( $\Lambda$ CDM) cosmological model is the current consensus for the description of the universe in which we live. It is a simple model that accurately explains the CMB, structure formation and other observations presented here. It assumes that Einstein’s general theory of relativity is the correct theory of gravity, that the universe is isotropic and homogeneous on large scales and that creation started at an infinitely dense point. The theory can also be amended to add a period of rapid inflation [24] that would have started  $10^{-36}$  seconds after the big bang and lasted  $\sim 10^{-32}$  seconds in order to solve the horizon and flatness problems not discussed here.

In particular, the  $\Lambda$ CDM model contains a self attractive substance called cold dark matter and a self repulsive force called dark energy  $\Lambda$ . This  $\Lambda$  is the very same cosmological

## Our Universe

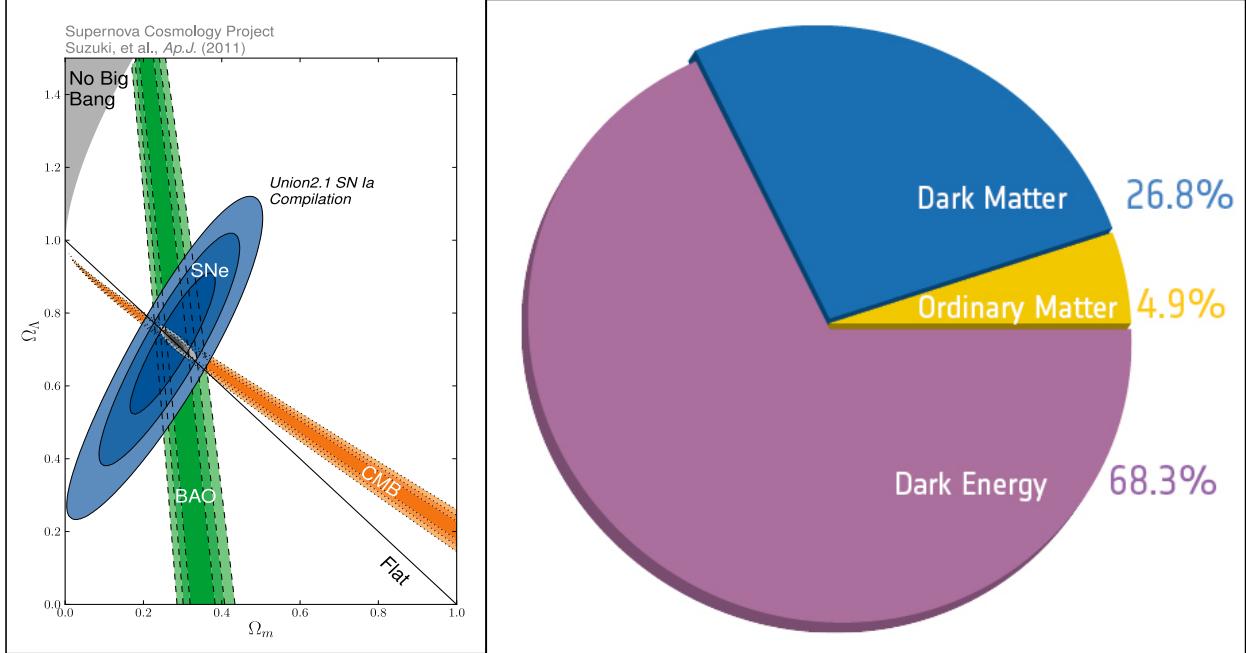


Figure 2.6: Left: is a plot of the cosmological constant vs matter density of the universe. CMB (WMAP and PLANCK), Supernova (SNe) and BAO measurements constrain the universe to be flat and dark energy/dark matter dominated [25]. Right, Planck's most recent constraints show that the composition of of our universe is: Dark Energy = 68.3%, Dark Matter = 26.8%, Baryonic Matter = 4.9% [19].

constant that Einstein added to keep the universe from collapsing. However, with a large enough value, this term would describe a universe being driven apart and accelerating as it does so. In the 1990's, measurements of the Hubble parameter using the recession of type 1a supernova as "standard candles" determined that the universe is in fact accelerating in its expansion [26]. Today, independent CMB, BAO and supernova redshift measurements [25] have been used together to very well constrain the ratio of dark energy to total matter and have found the curvature of space to be flat (figure 2.6). Amazingly, WMAP and Planck constraints from the CMB determine that dark matter and dark energy together account for

more than 95% of the energy density of our universe. In other words, the hugely successful standard model of particle physics only describes about 5% of the known universe. This work will not address the mystery of dark energy but will instead focus on the dark matter problem.

#### *2.2.11 Solutions to the Mystery of Dark Matter*

In some sense, dark matter has already been discovered but it has not been detected in a laboratory and its identity remains unknown. Over the years, tens of candidates have been proposed. The obvious initial response to early evidence was to propose candidates that didn't challenge the standard model: neutrinos, dust, magnetic monopoles, black holes, etc. The other reaction was to suspect that gravity behaves differently on different scales. However, if one accepts the evidence presented here, dark matter must be

- Very stable
- Non-interacting
- Non-baryonic
- Mostly cold (non-relativistic)
- Responsible for 20%-25% of the energy density of universe

It is also worth noting that Einstein's theory of gravity has been tested on micrometer scales [27] and on galactic scales and never found to be wrong. This is especially remarkable in the recent detection of gravitational waves predicted by general relativity. There are few interactions more violent than the merger of two black holes and therefore few remaining opportunities to find the theory to be broken.

With few compelling dark matter candidates left, after imposing the above constraints, I feel the two strongest contenders are Weakly Interacting Massive Particles (WIMPs) and axions. WIMPs were proposed in the 1970’s when it was realized that a particle with a cross-section associated with the weak scale would have about the right abundance today to account for dark matter. While the original coincidence-inspired WIMP has long since been ruled out, WIMPs today have become associated with the lightest stable super-partners as predicted by supersymmetry (typically the Neutralino). With a mass of 10 - 100 GeV, these heavy particles would be detected in the recoil of atomic nuclei. Although WIMPs remain the “go-to” dark matter candidate for many, excitement has grown around the axion as a natural and compelling solution. From here on, this work is devoted solely to the topic of axionic dark matter.

### 2.3 ***Dark Matter Axions***

In section 2.1 I introduced a mechanism by which a vexing problem in QCD is solved, resulting in the creation of a new invisible and non-baryonic particle. In section 2.2 I described a new form of matter that has allowed itself to be discovered without revealing its identity. In this section I start with the assumption that the undiscovered axion and some fraction of the elusive dark matter are one and the same. I then discuss axion evolution in the context of what is already known about the universe in order to better constrain the axion production mechanism, lineshape and symmetry breaking scale  $f_{PQ}$ . The value of  $f_{PQ}$  which determines the axion mass and coupling is not fixed by theory and must be determined experimentally.

#### 2.3.1 *Axion Production*

There are three primary ways in which relic axions could have been produced: thermally, from string decays or by a process called the “misalignment mechanism”. The axions produced from string or domain wall decay don’t give a good prediction for the energy density

contribution of axions and won't be considered here. Thermal axions would have been produced in quark and nucleon interactions when the universe was a hot, dense plasma. The axion mass would have needed to be greater than 1 meV for a cross section large enough to allow the interaction to produce a significant amount of matter but these masses are disallowed by red giant and supernova observations discussed in the following section 2.3.3. Finally, axions produced in this way would be "hot" and if dark matter was mostly comprised hot dark matter, large scale structure formation would have been disrupted.

A third production mechanism comes from the fact that before the PQ symmetry was broken and temperatures were above  $f_{PQ}$ , all values of  $\theta$  were possible. If inflation is ignored, the initial value of  $\theta$  ( $\theta_1$ ) can simply be taken to be the RMS average over all possible initial values between  $-\pi$  and  $\pi$  giving  $(\theta_1)_{RMS} = \pi/\sqrt{3}$ [28]. As the temperature dropped below  $\Lambda_{QCD}$  and instanton effects tipped the potential, a non-zero  $\theta_1$  would roll towards the minimum, overshoot and oscillate back and forth with an equation of motion like that of a damped pendulum

$$\ddot{\theta} + 3H(t)\dot{\theta} + m_a^2 \sin(\bar{\theta}) = 0 \quad (2.19)$$

where  $H(t)$  is the Hubble parameter and  $m_a$  is the mass that the axion acquires from this non-thermal process. For  $H < \theta$ , the oscillations are underdamped and lose energy to "Hubble friction". The axions form a pressure-less zero-momentum condensate with a density

$$\Omega_a h^2 \simeq .7 \left( \frac{10 \mu eV}{m_a} \right)^{7/6} \left( \frac{\bar{\theta}_i}{\pi} \right)^2 \quad (2.20)$$

This result shows that a smaller axion mass counterintuitively gives rise to a higher mass density and can be used to set a suggested lower limit on the axion mass. If  $\theta_1$  is assumed to be  $\mathcal{O}(1)$  the axion mass must be less than  $\mu eV$  to avoid overclosing the universe. This reasoning holds if the PQ symmetry is broken after inflation. However, in the event that PQ symmetry is broken before inflation,  $\theta_1$  would not be an average over all possible values

but a single specific instance of  $\theta$  from a small patch of the universe expanded to fill the whole observable universe. This means that  $\theta_1$  could have any value between 0 at  $2\pi$  and there's no reason it couldn't happen to have a value very near zero. Anthropic arguments can then be invoked to argue that small  $\theta_1$  would have been necessary for life to exist and the constraint on the lower bound of the axion mass is avoided. To summarize this section: axions produced from the misalignment mechanism are cold and have about the right mass and density to account for dark matter but may eliminate the possibility of  $m_a < \mu\text{eV}$ .

### 2.3.2 Galactic Evolution and Signal Linewidth

If axions are responsible for the majority of the dark matter in our galaxy then they are all around us, in large quantities and flowing in and out of our laboratories. The details of their density and velocity distributions throughout the galaxy depend upon how and when the Milky-Way was formed. While rotation curves can be used to measure the amount of mass within a spherical shell of a given radius, it is very difficult to determine a galaxy's full history and composition with only measurements made within that galaxy.

The simplest dark matter model assumes that axions fell into the potential well of our galaxy and became thermalized by gravitational interactions in doing so. Dark matter is treated as a pressure-less gas with a Maxwell-Boltzman like distribution and therefore an easily calculatable velocity dispersion. This model is known as the “Isothermal Sphere”. The axion distribution is expected to have the virial velocity ( $10^{-3}c$ ) and in the lab frame, the isothermal signal is a boosted Maxwellian with a fraction linewidth of  $10^{-6}$ .

While the isothermal sphere is *naïve* and unable to explain all of the features of our galaxy, it has the advantage of being fairly consistent with simulations while remaining a generic model that is easy to modify. For example, simulations have shown that a galaxy that has gone through a merger with another galaxy will sometimes have additional dark matter clumping around the visible disc [29]. The “Dark Disk” model accounts for this over-density

by adding an additional thin disc-like distribution on top of an isothermal distribution. In this case, the net expected axion signal would be the sum of the linewidths from these two contributions. With a lower velocity dispersion, the dark disk contribution would have a narrower signal in frequency space.

In the event that some fraction of axions fell into our halo more recently and haven't had time to come to thermal equilibrium, discrete "cold flows" may have been created from dark matter falling in and around the galactic core. Axions may also form a Bose-Einstein condensate and coalesce into "caustic rings" at the turnaround points where cold flows fall out and back into the galaxy. This would result in a series of high density dark matter rings within the halo that would have a low velocity dispersion and therefore a very narrow signal width [30].

Ultimately, the distribution of dark matter and the validity of these models may be determined by N-body simulations of galaxy formation. The Romulus25 simulation performed by the University of Washington N-body group have found that the expected axion linewidth may be two to three times narrower than the isothermal prediction [31]. In these simulations, the addition of baryons breaks the assumption that a galaxy behaves purely like a collision-less gas. Several billion gas particles, dark matter particles and stars were placed in a cube,  $25 \text{ Mpc } h^{-1}$  on a side, and the system was allowed to evolve. Among the galaxies that were formed, a subset were found to be similar to the milky-way and the axion signals of these case studies are plotted over the (dashed) isothermal signal in figure 2.7.

For the sake of being conservative, my analysis (discussed in section 5) uses the isothermal signal dispersion even though a narrower filter would result in an improved SNR proportional to the change in width.

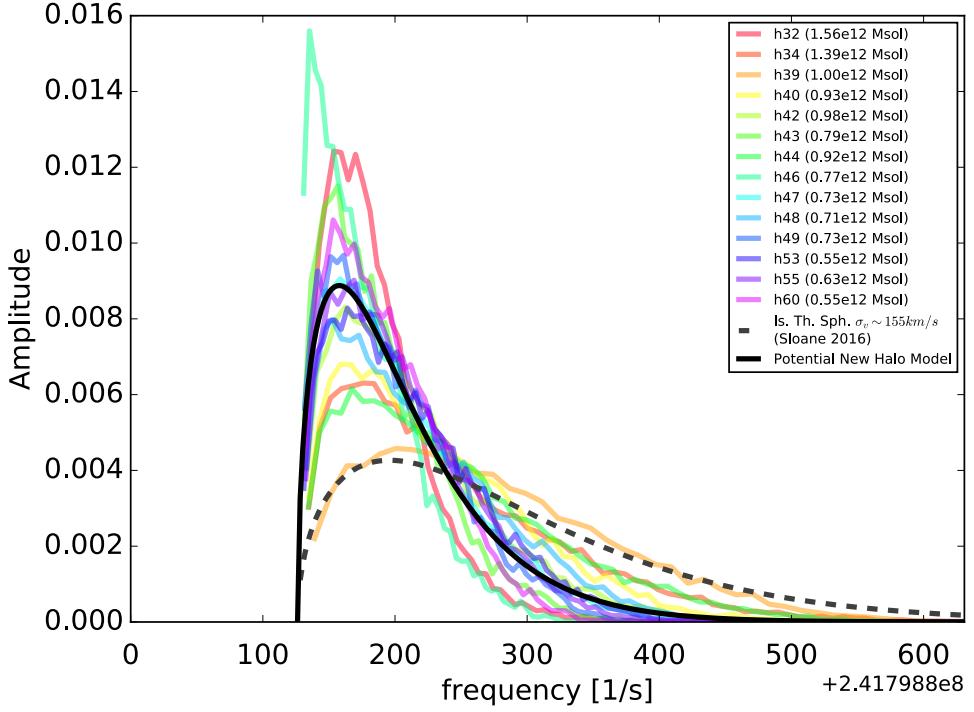


Figure 2.7: Romulus25 N-body simulation of galaxy formation showing an axion linewidth narrower than the isothermal prediction [31].

### 2.3.3 Astrophysical Constraints

Accelerator experiments and overclosure arguments suggest that the axion mass is between  $\mu\text{eV}$  and  $\sim 10$  keV. Astrophysical observations greatly lower this upper limit. Because the axion mass and couplings are both joined by their inverse relationship to  $f_{PQ}$ , a heavier axion implies a stronger coupling. If axions couple too strongly to normal matter, they can easily be created and like neutrinos, provide a method of supercooling stars and supernova. With more ways to lose energy, the speed of stellar evolution in main sequence stars would be increased via  $e\gamma \rightarrow a\gamma$  and noticeably disrupted for masses in the eV - keV range [32]. The strongest constraint on  $g_{a\gamma\gamma}$  comes from horizontal branch (HB) stars which have reached the helium

burning phase. Such stars would have large energy losses due to the primakoff production of axions and would have shortened lifetimes compared to red giants. By comparing the number ratio of HB stars to red giants in globular clusters (GC) a limit  $g_{a\gamma\gamma} < 10^{10} \text{ 1/GeV}$  can be set. Likewise, if axions had a mass between meV and 2eV, their production from nucleons ( $NN \rightarrow aNN$ ) would have effectively diverted a significant fraction of the energy carried by neutrinos leaving the supernova event in 1978. As a result, the duration of the neutrino burst observed on earth would have been dramatically shortened. As far as this constraint is concerned, masses  $< \text{meV}$  and  $> 2\text{eV}$  are allowed because axions don't get appreciably produced in the case of the former and can't free stream out in the latter. Finally, optical searches looking for narrow band photons from the decay of axionic halos have also independently excluded 1-10eV axions [33].

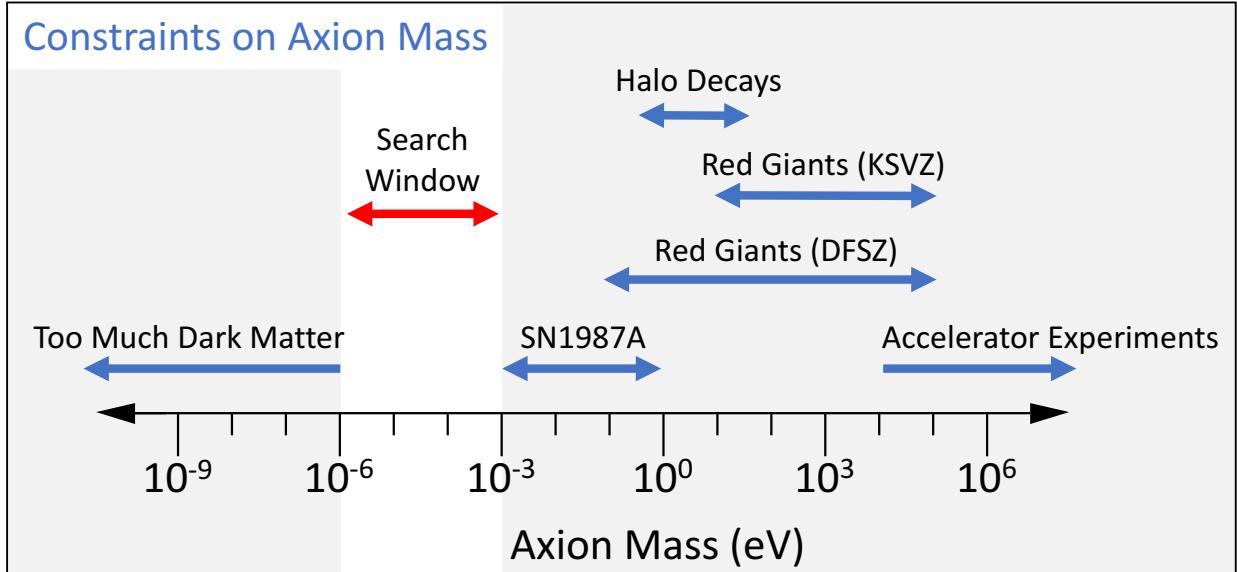


Figure 2.8: A simplified summary of astrophysical constraints on axion mass. This figure is a remake of a figure in [34]. A more detailed summary can be found in [35]

Figure 2.8 summarizes these astrophysical constraints and reveals the “search window”

where plausible axions are waiting to be detected. In the following sections I discuss the distribution of dark matter in our galaxy, the axion signal line shape and the various current experiments that aim to detect this signal or close the search window entirely.

### 2.3.4 Experimental Searches

For the purposes of keeping track of which varieties of axions have been excluded, it is useful to add a  $g_{a\gamma\gamma}$  axis to figure 2.8 and make a plot like figure 2.9, showing  $g_{a\gamma\gamma}$  versus  $m_a$ . Once again, the viable axion search window ( $1\mu eV < m_a < 10meV$ ) sits between the overclosure and SN1978A constraints but now the  $g_{a\gamma\gamma}$  model dependence appears as two diagonal lines. QCD axions are constrained to live between these KSVZ and DFSZ diagonal lines. There is no reason why an “axion-like” particle (ALP) couldn’t be discovered outside of these lines but its existence would not solve the strong CP problem. Most experiments to date are not sensitive to true QCD axions but are able to exclude axion-like particles.

“Light-shining-through-walls” (LSW) attempts to create and then detect ALPs by shining laser light on a wall in a strong magnetic field. The hope is that photons will convert to axions in the magnetic field and pass through the wall. The axions will then convert back to photons on the other side of the wall and be detected. The probability of a detection is proportional to  $g_{a\gamma\gamma}^4$ . Laser experiments have set a limit of  $g_{a\gamma\gamma} < 10^{-7} 1/\text{GeV}$ .

Solar telescopes (“helioscopes”) like the CAST experiment also use a large magnetic field to search for the conversion of broadband axions thermally produced in the sun back into photons [36]. CAST has set a limit of  $g_{a\gamma\gamma} < 8.8 \times 10^{-11} 1/\text{GeV}$  up to a mass  $m_a < 20 \text{ meV}$ . The proposed IAXO (The International Axion Observatory) experiment will be an improved solar axion search with a target sensitivity of  $g_{a\gamma\gamma} < 5 \times 10^{-12} 1/\text{GeV}$ . Notice that the only search technique sensitive to realistic QCD axions is the microwave cavity technique. These detectors are known as Haloscopes and are covered in the following section.

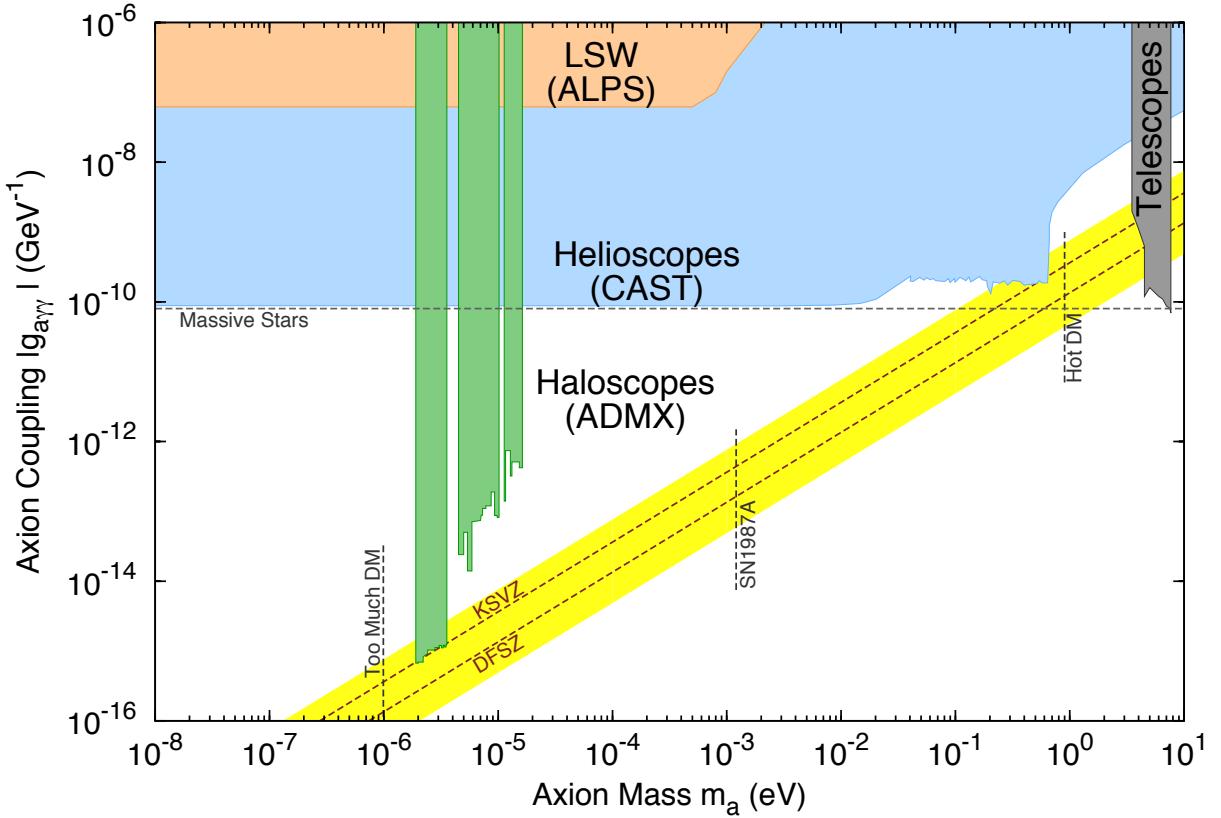


Figure 2.9: Bounds on  $g_{a\gamma\gamma}$  (vertical axis) and axion mass (horizontal axis) which define the axion search window [35]. QCD axions lie between the KSVZ and DFSZ diagonal lines. The lower “too much dark matter” bound applies if the Peccei Quinn symmetry breaking occurs after inflation. The upper SN1978a bound indicates where axions heavier than the bound would have shortened the neutrino burst from supernova 1978a.

## 2.4 The Axion Haloscope

Although it is expected that the number destiny of ueV-meV dark matter axions is  $10^{14}/\text{cc}$  -  $10^{11}/\text{cc}$ , they are incredibly weakly interacting and would be very difficult to detect, with a lifetime

$$\tau_{a\gamma\gamma} \sim \frac{10^{54} \text{sec}}{(m_a/\mu\text{eV})^5} \quad (2.21)$$

While the axion has the potential to be discovered through a number of couplings ( $g_{a\text{ii}}$ ) to normal matter, the axion-to-two-photon coupling ( $g_{a\gamma\gamma}$ ) is particularly promising because it is relatively model independent. The most sensitive technique for searching for axions was invented by Pierre Sikivie in 1983 and named the axion haloscope [37].

#### 2.4.1 Detection Strategy

The technique employs a resonant cavity immersed in a uniform magnetic field and looks for the resonant conversion of axions to photons via an inverse primakoff effect. Equation 2.9 describes the interaction pertaining to this process ( $\mathcal{L}_{a\gamma\gamma} = -g_{a\gamma\gamma}a\mathbf{E} \cdot \mathbf{B}$ ) where  $\mathbf{E} \cdot \mathbf{B}$  is the overlap of the electric and magnetic fields. In this case, the magnetic field is externally applied by a solenoid magnet and the electric field shape is given by the resonant mode in the cavity. Galactic axions interact with the magnetic field, convert to photons and deposit power into a resonant mode in the cavity. By inserting an antenna into the cavity, some of that power can be extracted, amplified and measured in a power spectrum. Thus far, it is assumed that a cavity resonance is coincidentally tuned to the unknown mass of the axion. Not knowing the mass of the axion means that the cavity needs to be tunable. One simple example of a tunable geometry includes a vertically oriented metal or dielectric rod that can be moved around the cavity. The movement of this rod changes the EM boundary conditions in the cavity and alters its frequency (addressed in chapter 4).

The configuration of this experiment thus described is shown in Figure 2.10. This is a simplified cartoon schematic of the Axion Dark Matter eXperiment whose apparatus is described in chapter 3. To the left, axions are shown converting to photons within the cavity. The signal is extracted via an antenna, amplified, mixed down to lower frequencies, digitized, Fourier transformed and squared to create a power spectrum. The most naive

# The Axion Haloscope

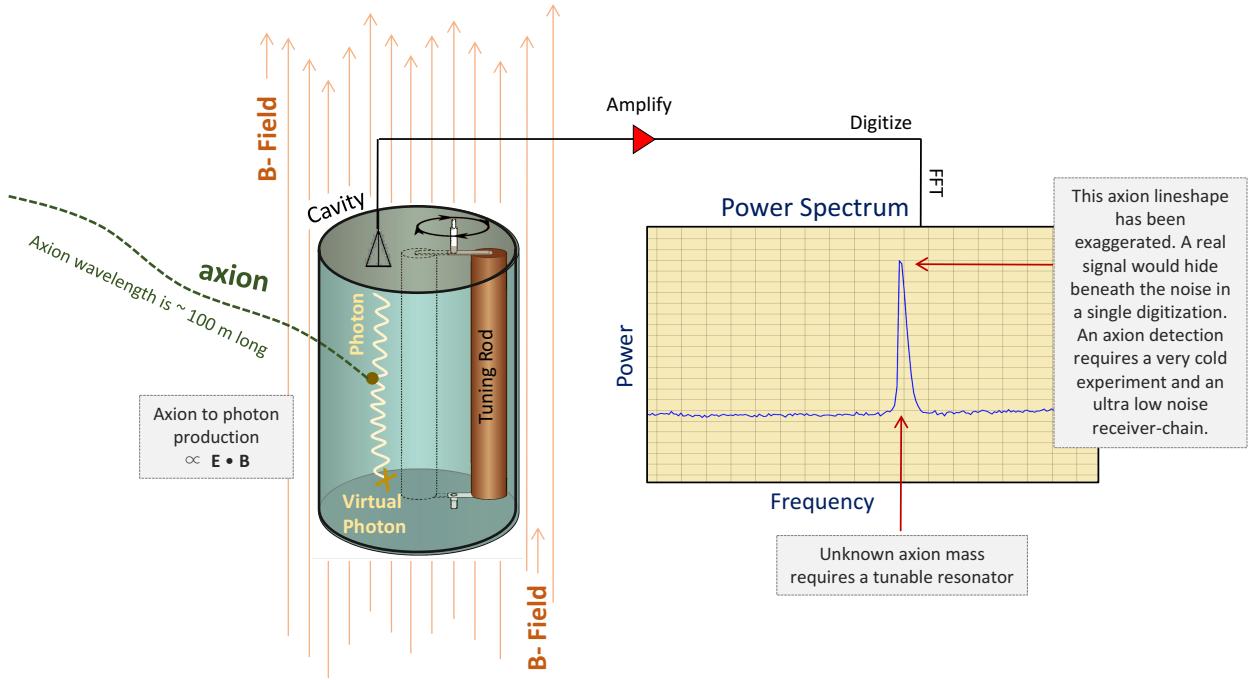


Figure 2.10: Template for an axion haloscope. Axions convert to photons in a large magnetic field and deposit a small amount of power in a resonant mode of a microwave cavity. Axion power from the cavity is extracted with an antenna, amplified, mixed down to lower frequencies and digitized. Evidence for an axion would appear as a surplus of power in a power spectrum.

expectation for the axion signal shape from the isothermal halo model is Maxwellian whose start frequency is given by mass of the axion and the width is of order  $\Delta E/E$ . At 1 GHz the mass of the axion would be approximately 1 kHz wide. The cadence of such an experiment is to digitize the signal coming out of the cavity at a particular frequency, move the tuning rod, digitize again, tune etc. This process repeats over and over again until the axion is found or ruled out for a particular frequency range. However, the expected signal is too small to be seen in a single 1-2 minute digitization. For this reason, the cavity frequency step size

between tunings needs to be small enough to allow adequate integration time per frequency. After taking many digitizations, all spectra are co-added together to form a single wide power spectrum. After this process, assuming enough data has been taken, the spectrum has enough signal to noise to discover the axion or constrain  $g_{a\gamma\gamma}$ .

#### 2.4.2 Expected Signal to Noise

The expected signal from an axion haloscope can be derived from the cavity equivalent circuit [38, 34, 39, 40], by directly solving maxwell's equations that include a new source term [41, 42, 43, 44] or by other means [45, 46]. These derivations show that the axion power developed in the cavity is

$$P_a = \frac{\epsilon_0 c \hbar^2 g_{a\gamma\gamma}^2 \rho_a}{m_a^2} B^2 V f C_{mnp} \quad (2.22)$$

In a form that is more useful to the experimenter, the axion power out of the antenna is

$$P_{out} = 5.7 \times 10^{-27} W \left[ \left( \frac{g_\gamma}{0.97} \right)^2 \left( \frac{\rho_a}{0.45 \text{GeV/cc}} \right) \right] \left[ B^2 V \left( \frac{f}{1 \text{GHz}} \right) \left( \frac{Q_L}{10,000} \right) \left( \frac{C_{mnp}}{0.5} \right) \left( \frac{1 - 2S_{11}}{1 - S_{11}} \right) \right] \quad (2.23)$$

where the values of the terms grouped in the left brackets are fixed by nature and terms in the right brackets are parameters that can be controlled by the experimenter.  $g_\gamma$  is a model dependent constant (-0.97 for KSVZ and +0.36 for DFSZ),  $\rho_a$  is the local density of axions in our halo, B is the magnetic field strength in Tesla, V is the volume of the cavity in liters, f is the resonant frequency of the mode, Q is the quality factor of the resonant mode,  $S_{11}$  is zero if the antenna is critically coupled (discussed in section 5.3.2) and  $C_{mnp}$  is an axion-photon conversion efficiency term given by a number between 0 and 1. It is called the “form factor”, and is defined as the overlap between the axion wave function and the produced photon wave function

$$f_{nlm} \equiv \frac{\left(\int_V dV \mathbf{E}(\mathbf{x}, t) \cdot \mathbf{B}(\mathbf{x})\right)^2}{VB^2 \int_V dV \epsilon_r E^2} \quad (2.24)$$

where  $\mathbf{E}(\mathbf{x}, t)$  is the electric field of the  $nlm$  cavity mode,  $V$  is the cavity volume,  $\epsilon_r$  is the relative permittivity within the cavity, and  $\mathbf{B}$  is the externally applied magnetic field. The ADMX magnetic field is oriented more or less vertically along the cavity's axis. The direction and distribution of  $\mathbf{E}$  within the cavity depends on the cavity mode. Within a tunable right circular cylinder there are transverse electric ( $\text{TE}_{nlm}$ ) modes, transverse magnetic ( $\text{TM}_{nlm}$ ) modes, and transverse electromagnetic ( $\text{TEM}_n$ ) modes. The TE and TEM modes both lack any axial  $E$  field, resulting in negligible form factors. Only the  $\text{TM}_{0n0}$  modes have non-zero form factors and can be used to hunt for axions.

For typical theoretical and ADMX parameter values ( $\rho = .45\text{GeV}/cc$ ,  $B = 7T$ ,  $V = 135$  liters,  $f = 700$  MHz,  $Q = 70,000$  and  $C_{010} = 0.5$ ) the expected power out of the cavity is  $\sim 1.8 \times 10^{-22} W$  ( $\sim 2.5 \times 10^{-23} W$ ) for KSVZ (DFSZ) axions. For the Sidecar experiment discussed in section 3.3, given the smaller cavity and reduced fringe field ( $B = 7/2T$ ,  $V = 0.38$  liters,  $f = 5$  GHz,  $Q = 15,000$  and  $C_{010} = 0.5$ ) the expected signal is  $\sim 2.0 \times 10^{-27} W$  ( $\sim 2.8 \times 10^{-28} W$ ) for KSVZ (DFSZ) axions. To detect this signal, an experimenter must be able to observe it above the noise. The output of the cavity obeys the radiometer equation

$$SNR = \frac{P_{out}}{k_B T_{sys}} \sqrt{\frac{t}{B}} = 4.1 \times 10^{-4} \frac{\left(\frac{g_\gamma}{0.97}\right)^2 \left(\frac{\rho_a}{0.45\text{GeV}/cc}\right) \left(\frac{f}{1\text{GHz}}\right) \left(\frac{Q_L}{10,000}\right) \left(\frac{C_{mnp}}{0.5}\right) (1 - S_{11}) B^2 V t^{1/2}}{\sqrt{B} (T_{physical} + T_{noise})} \quad (2.25)$$

where  $k_B$  is the Boltzmann constant,  $t$  is the integration time,  $B$  is the bandwidth and  $T_{sys}$  is the system temperature which is the sum of the physical cavity temperature and the amplifier noise temperature contributions

$$T_{sys} = T_{phy} + T_N = T_{phy} + T_1 + \frac{T_2}{G_1} + \frac{T_3}{G_1 G_2} + \frac{T_4}{G_1 G_2 G_3} + \dots \quad (2.26)$$

where  $T_i$  and  $G_i$  is the noise temperature and gain of the  $i$ th stage. The challenge of looking for such a weakly coupled particle is overcome by reducing the system temperature and adjusting the integration time enough to boost the SNR.

#### 2.4.3 Designing an Experiment

The motivation behind the design of the ADMX apparatus (Chapter 3) is perhaps more apparent when considering what kind of experiment would optimize each term in equation 2.25 and maximize the SNR. Using this equation as a roadmap, the design of an axion detector similar to the original Pierre Sikivie axion haloscope must be informed by the following

$C_{mnp}$ : The normalized volume integral of  $\vec{E} \cdot \vec{B}$  must be non-zero (and preferably as close to 1 as possible) to achieve axion sensitivity. Traditional haloscopes have used the  $TM_{010}$  of a right circular cylinder because it maximizes the form factor when paired with a uniform magnetic field parallel with the length of the cavity. ADMX has recently demonstrated that simultaneous axion searches can be carried out on both the  $TM_{010}$  and  $TM_{020}$  modes. Other ADMX *R&D* projects such as Orpheus [47] and Electric Tiger [48] make use of different geometries and techniques to create non-zero form factor resonators. Orpheus uses a spatially varying magnetic field to coincide with the electric field of a standing wave between two metal plates. As the electric field changes polarity, the magnetic field also switches direction preserving the sign of  $\vec{E} \cdot \vec{B}$  over the length of the standing wave. Electric Tiger uses a TE mode inside of a rectangular waveguide that is surrounded by a magnetic field perpendicular to the axis of the guide and parallel to the TE electric field direction. Dielectric blocks are placed periodically within the waveguide to localize the electric field of a particular field direction and minimize it's canceling effects on the volume integral of  $\vec{E} \cdot \vec{B}$ . This work focuses on the traditional one or two rod cavity in a uniform magnetic field (discussed further in

chapter 4).

**f:** The mass of the axion determines the frequency at which equation 2.23 is valid. Theoretical axion mass predictions are used to inform the technological development of new resonators able to tune over a wide range of frequencies. The  $TM_{010}$  mode of the current primary ADMX cavity can be tuned from approximately 600 MHz to 950 MHz. The second  $TM_{020}$  channel was added to increased the effective tunability of the experiment. In the event that ADMX rules out axions within the  $\sim 0.5 - 2$  GHz frequency range, new tunable resonators will need to be implemented in order to reach frequencies outside of this range. A traditional axion cavity search below 500 MHz becomes very difficult because lower frequencies are associated with larger geometries and therefore require the bore of the magnet to be very large. Currently, the ADMX group is doing R&D work on an LC circuit design that can be tuned to low frequencies with the variable capacitance of a varactor diode [49]. At higher frequencies, cavity dimensions become smaller and the problem of fitting an experiment into the bore of a magnet is traded for a loss in SNR. The Sidecar cavity (section 3.3) is a pathfinder experiment designed to pave the way for higher mass axion searches.

**V:** The volume of the resonator should be made as large as the magnet bore will allow while still maintaining the target frequency range of the experiment. For a right circular cylinder, the frequencies of the  $TM_{0n0}$  modes are inversely proportional to the cavity radius but are immune to length. Increased sensitivity can be achieved by making the cylinder longer but higher mass axions searches can't avoid a loss in volume from the reduced radius. With an inner diameter of 2.5 inches and a  $TM_{010}$  tuning range of  $\sim 4 - 6$  GHz, the Sidecar experiment is close to the boundary where traditional haloscopes no longer become practical. The volume of the Sidecar cavity is  $\sim 1/200$  the volume of the main ADMX cavity and therefore the SNR of this experiment is reduced by a factor

of 200. Some SNR is made back by the higher operating frequencies but this gain is counteracted by reduced quality factors associated with these higher frequencies. By ganging together multiple co-tuned cavities, the effective volume of the experiment can be increased to regain the lost SNR of a single cavity. The technical challenges that come with this approach are then associated with being able to mechanically tune many cavities in a cramped space and without overheating the cryogenics. The Sidecar piezoelectric tuning/coupling system is a testbed for possible tuning solutions to this problem. At frequencies above 10 GHz it becomes advantageous to use bolometers for single axion to photon counting.

**Q:** Classically, the quality factor is proportional to

$$Q \propto \sqrt{f\sigma} \frac{V}{S} \quad (2.27)$$

where  $\sigma$  is the electrical conductivity and  $S$  is the cavity surface area. For a right circular cylinder:  $S = 2\pi rh$ ,  $V = \pi r^2 h$  and  $f \propto 1/r$  gives

$$Q \propto \sqrt{r\sigma} \quad (2.28)$$

Given that  $r$  is already fixed by some other frequency or geometry constraint, equation 2.28 suggests that improving the bulk properties of the conductor is all that can be done to improve the  $Q$ . However, at high frequencies and high conductances (associated with low temperatures) the classical skin depth becomes shorter than the mean free path of the conducting electrons and the surface finish of the cavity becomes increasingly important. Some have proposed an axion search that uses a superconducting cavity to increase the  $Q$ .

While a higher Q causes more power to be developed in equation 2.23 there are cases where a slight reduction in the Q is actually desirable. A cavity with a reduced Q becomes less sensitive but broader band. By slightly overcoupling the antenna, a larger percentage of power removed by the antenna mostly makes up for the loss in axion power developed and the increase in the bandwidth of the cavity improves the overall scan rate of the experiment.

**B:** While  $f$ ,  $C_{mnp}$ ,  $V$  and  $Q$  are all interconnected,  $B$  is essentially an independent variable limited only by cost. On top of that, unlike the other parameters, the signal power is proportional to  $B^2$  so doubling the field increases the SNR by a factor of 4. Though expensive, higher field magnets are undoubtedly the way forward for small, high mass axion searches.

$T_{phys}$ : The primary background for this kind of experiment is thermal noise where the SNR is  $\propto 1/T_{phys}$  and the scan rate is  $\propto 1/T_{phys}^2$ . Therefore, an axion search requires a huge amount of refrigeration infrastructure to cool the experiment as close to absolute zero as possible. ADMX uses a closed-loop  $^4He$  system to feed a superconducting magnet and to evaporatively cool the experimental apparatus to  $\sim 1K$  by pumping on a pot filled with helium. In addition to the  $^4He$  system, the experiment has been upgraded to use a dilution refrigerator that works together with its own  $^3He/^4He$  closed loop recovery system. The dilution refrigerator allows the experiment to be cooled near 100mK which represents a factor of 100 increase in scan rate.

$T_{noise}$ : Ultimately it is  $T_{sys}$  that appears in equation 2.25 which means that the noise temperature of the amplifiers must be reduced as well as the physical temperature of the experiment. A state of the art solid-state amplifier, with a typical  $T_{noise} \sim 2K$ , would be the dominant source of noise at a physical temperature of 100mK. A sensitive axion search therefore requires the use of a quantum noise limited amplifier on the front end

of the receiver, such as a Superconducting Quantum Interference Device (SQUID) or a Josephson Parametric Amplifier (JPA) detailed in section 3.2.4. The noise temperature of these devices scale with physical temperature and are only noise limited by the uncertainty principle

$$T_{Quantum\ Limit} = \frac{hf}{k_B} \quad (2.29)$$

where  $h$  is Planck's constant and  $k_B$  is the Boltzmann constant. At 1 GHz  $T_{Quantum\ Limit} = 48\text{mK}$ . However, these amplifiers are expensive, finicky and require special care. They are incredibly sensitive to changes in magnetic field (sensitive to a single quanta of flux  $= \frac{h}{2e}$ ) but need to operate in close proximity to a very high field magnet. ADMX overcame this challenge by placing the SQUID/JPA electronics inside their own bucking magnet that cancels the magnetic field in the vicinity of the amplifiers. The bucking magnet is counter-wound for zero mutual inductance between itself and the main magnet which eliminates the huge forces between the two and prevents one magnet from quenching the other.

Having introduced this technique by which galactic axions in our halo may be discovered and outlined the requirements of such a search, I will now describe a specific instance of such an experiment.

## Chapter 3

# THE AXION DARK MATTER EXPERIMENT

The Axion Dark Matter eXperiment (ADMX) is the world’s most sensitive search for dark matter axions. This chapter describes the hardware of ADMX and its particular implementation of the general axion haloscope described in section 2.4. I start with a history of the experiment, then give a description of the infrastructure and hardware that supports three separate axion searches. These three searches are the result of two different cavities (the main and Sidecar cavities) that are able to detect axions converting to photons on two different ( $TM_{010}$  and  $TM_{020}$ ) modes. The emphasis of this chapter is an overview of the project and the anatomy of the apparatus hardware. Future chapters will cover details about how the data acquisition brings the hardware to life by tracking resonant modes, handling incoming information and automating the data taking process.

### **3.1 Resources, Infrastructure and Apparatus**

#### *3.1.1 History*

ADMX has undergone many different technological improvements over the years. Soon after the idea was conceived in 1985, axion haloscope prototypes were tested at the University of Florida [50] and at Brookhaven National Lab [51]. These early experiments were not sensitive to realistic QCD axions but demonstrated the proof of concept. The first phase of ADMX (phase 0), building off of these concepts, developed a full-blown experiment that ruled out optimistically coupled (KSVZ) axions in the mass range  $1.9 \mu\text{eV}$  to  $3.3 \mu\text{eV}$  between 1995 and 2005 [34]. To achieve higher sensitivity, ADMX (in phase 1), between 2005 and 2008,

incorporated quantum-noise limited amplifier technology that had not been available previously. The incorporation of these Superconducting Quantum Interference Devices (SQUIDs) also required a field-free region of the experiment achieved with the addition of a bucking magnet. The experiment took data in 2008 - 2010 and set new limits on KSVZ axions in the mass range of  $3.36 \mu\text{eV}$  to  $3.69 \mu\text{eV}$  [45].

In 2010, the ADMX 9 Ton, 8.5 tesla magnet was moved from Lawrence Livermore National Lab to its new home at the Center for Experimental Physics and Astrophysics (CENPA) at the University of Washington in Seattle, WA. Since then, the experiment has been redesigned and rebuilt. In 2014, dual-channel functionality was demonstrated by simultaneously taking data on the  $TM_{010}$  and  $TM_{020}$  modes. In 2016, the experiment reached temperatures below 200 mK with the addition of a dilution refrigerator. At this time, a new, high frequency (4-7 GHz), piezoelectrically tuned cavity was added to the baseline of the experiment. This “Sidecar” experiment was made to run in parallel with the main experiment and pave the way for higher mass detectors.

### 3.1.2 *Collaborators*

I believe the strength of ADMX has been its light-weight “lean and mean” scientific approach. Also, as a graduate student, there is no better gift than the responsibility to take an active role in all aspects of the experiment just down the hall. At the same time, ADMX has always been a “person-power limited” project and has recently had the opportunity to welcome new expert collaborators. In the last year, the ADMX collaboration between the University of Washington, LLNL, University of Florida, University of Sheffield, NRAO and University of California Berkeley expanded to include the Fermi National Accelerator Laboratory (FNAL), the Pacific Northwest National Laboratory (PNNL) and the Los Alamos National Laboratory (LANL). Filled with fresh blood and new ideas, the people of ADMX (shown in Figure 3.1) are tackling the technological challenges of tomorrow. Today, the current experiment, with

all of its upgrades is poised to hunt down or rule out axions at an unprecedented rate and sensitivity.



Figure 3.1: ADMX collaboration photo from January 2016.

### 3.1.3 *Infrastructure*

CENPA is a Department of Energy “Center of Excellence” and provides ADMX with world class technical staff and support. The experiment resides at the end of an old linear accelerator hall shown in figure 3.2. The heart of the experiment is a resonant cavity assembly called the “insert” (discussed in the following section) that hangs from the ceiling of a clean room during upgrades. In preparation for data-taking, the insert is lifted to the center of the site and lowered into the 60 cm inner diameter bore of a 8.5 T superconducting magnet. Barely visible at the center of the top panel of the figure, the 11' tall magnet hides in a hole

in the ground such that the top plate of the insert sits flush with the recessed floor.



Figure 3.2: Images of the ADMX experimental site. The large superconducting magnet and experimental insert sit in a recess in the floor and are hidden from view in the center of the top of the figure. A  ${}^4He$  recovery system cools the experiment, collects the blown off gas, re-compresses it in a pump shed (lower right hand side of the figure) and turns it back into a liquid to be used again. The lower left had side of the figure shows a separate  ${}^3He/{}^4He$  gas handling system used in conjunction with a dilution refrigerator.

Given that the expected axion signal is many orders of magnitude fainter than the room temperature thermal background, an axion detector is necessarily a cryogenic experiment. ADMX utilizes two separate closed-loop cryogenic systems: a  $^4He$  and a  $^4He/^3He$  dilution system.

The  $^4He$  system keeps the coils of the superconducting magnet submerged and feeds the reservoir in the experimental insert. This reservoir supplies helium to the 1K plots (pumped helium pots) that evaporatively cool the insert. Due to the production nature of a continually scanning experiment, the  $^4He$  system is equipped to recover and reuse the exhausted gas.  $^4He$  blown off from the experiment is stored temporarily in a “gas bag” before being pumped to holding tanks outside the building. The gas is then recompressed (lower right hand panel of figure 3.2) and sent to high pressure storage back inside the building. A Linde helium liquefier purifies and re-liquefies the helium and stores it in the large dewar which feeds the experiment and the cycle continues.

The  $^4He/^3He$  gas handling system is located in a tunnel below the floor (shown in the lower left hand panel of figure 3.2) and works in conjunction with a Janis dilution refrigerator mounted inside the insert. A dilution refrigerator makes use of the fact that when a mixture of  $^3He$  and  $^4He$  is cooled to mK temperatures, the concentration of the solution phase separates into two populations: a fixed ratio mixture of  $^3He$  and  $^4He$  and a pure layer of the lighter  $^3He$  which rises to the top. The ratio of the isotopes in the dilution (6.6%  $^3He$  to 93.4%  $^4He$ ) is fixed by Van der Waals forces and if the concentration is disturbed, the system will try to restore balance. By pumping on the  $^3He/^4He$  dilution, the  $^3He$ , being lighter is preferentially removed from the solution and osmosis will pull some of the pure  $^3He$  into the dilution to maintain the proper concentration. When this happens, entropy is increasing since  $^3He$  from a pure phase is going into a more disorderly phase. This means that the system must absorb heat from the surrounding environment to maintain the  $^3He/^4He$ . The ADMX dilution refrigerator has cooling power of 800  $\mu\text{W}$  and is able to maintain a base

temperature of less than 150 mK.

### 3.1.4 *Insert*

The experimental “insert” is a 59 cm diameter, 3 meter tall cylinder made up of a collection of resonant cavities, cryogenic hardware, gear boxes, cold RF-electronics, radiation baffles and includes a secondary magnet and helium reservoir. Together, these pieces facilitate three low-noise, tunable axion searches that can maintain a base temperature of less than 150 mK. Figure 3.3 shows a cutaway of the insert within the cryostat of the magnet. At the base of the insert sits the main resonant cavity. The region just above the cavity defines the “cold space” where the dilution refrigerator, the sidecar cavity and the quantum electronics are all in thermal contact. Hovering above this space and connected by weak thermal links (stainless steel pillars) is the 1K stage maintained by an evaporatively cooled pot of helium. Gearboxes for the main cavity’s rod and antenna motion are thermally tied to this stage along with various other components. Above the 1K plate there is a stainless steel helium reservoir followed by radiation baffles and finally the top plate of the insert.

A complexity in the design arises from the fact that the cold electronics must remain at mK temperatures but are unable to tolerate a non-zero magnetic field. To accommodate these strict requirements, the helium reservoir was engineered to be hollow, topologically equivalent to a toroid. Within that reservoir, submerged in helium is a secondary magnet. This bucking magnet cancels the field from the main magnet at the center of the reservoir. To prevent quenching and avoid huge forces, the bucking magnet was counter-wound such that there is zero mutual inductance between the two magnets. Finally, tied to the dilution refrigerator, rising above the 1K plate and sitting within the field-free void, without touching, is a component I named the “SQUIDadel”. The SQUIDadel is a copper fortress that sits on top of a long copper pillar and houses a collection of sensitive quantum electronics, circulators and switches. Here, the tiny signal from the cavity receives a large boost in gain without

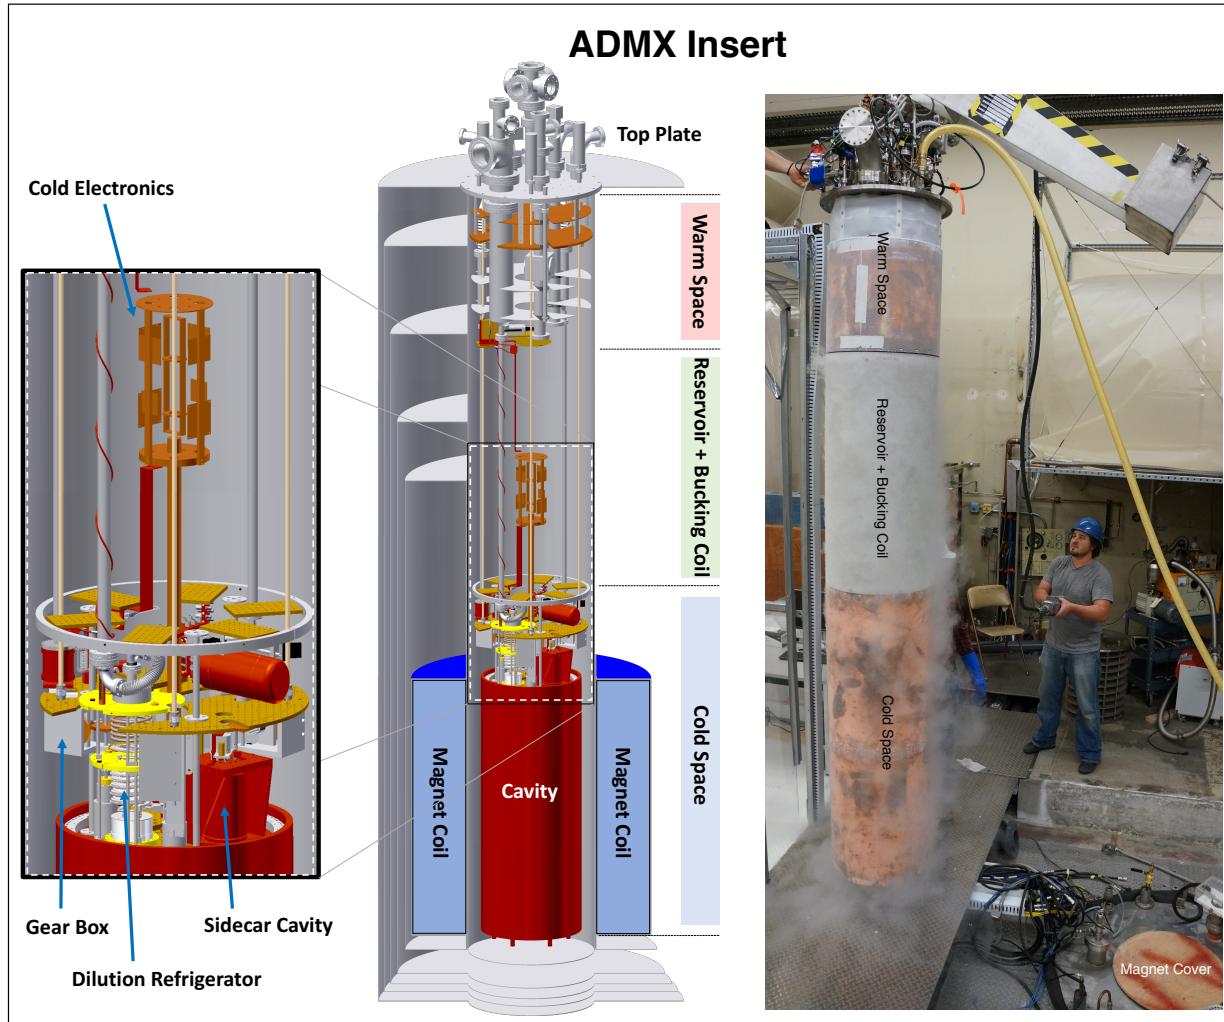


Figure 3.3: Cross-section of the ADMX removable insert sitting inside the superconducting magnet (center) and a picture of me using a crane to remove the insert from the magnet (right). The resonant cavity is at the base of the insert, surrounded by the magnet coils, shown in blue. Directly mounted to the cavity is a dilution refrigerator, the Sidecar cavity (discussed later) and a collection of sensitive electronics also shown in the zoomed in portion of the figure (left). Above this region is a helium reservoir, bucking magnet and radiation baffles.

much added noise, it travels up superconducting RF cables to the top of the insert and makes its way to the room temperature receiver. The details of each of these RF channels are given in the following sections 3.2 and 3.3.

### 3.2 *Dual-Channel Axion Search*

This section describes the  $TM_{010}$  and  $TM_{020}$  searches associated with the main ADMX cavity.

#### 3.2.1 *Cavity*

The interface to a hidden axionic halo and start of the signal path is a high Q resonant cavity. The primary ADMX cavity is a  $\sim$ 140 liter stainless steel cylinder coated with Oxygen-free high thermal conductivity (OFHC). The cylinder is formed with two end caps tightly bolted to both knife-edged ends of a 16.5" inner diameter, 39.9" tall barrel. The knife-edge junction ensures a good electrical connection, lower ohmic losses and therefore a higher Q. The loaded  $Q_L$  at critical coupling (half the weakly coupled  $Q_0$ ) typically ranges between 50,000 and 90,000.

#### 3.2.2 *Tuning Rod Motion*

The resonator is made tunable with two copper-plated tuning rods. These rods are 1" in diameter, 39.4" tall and sweep out  $\sim$ 3.6" radius circular trajectories shown in figure 3.4. This allows the rods to be rotated from the walls of the cavity, where they minimally impact the electromagnetic boundary conditions of the resonator ( $\theta = 0$ ), to the center where the  $TM_{010}$  frequency is at its highest ( $\theta = 180$ ). The armatures that protrude through the end caps and offset the rods are made of alumina. This prevents the rod from acting as an antenna and radiating power out of the system.

Mechanical motion is translated to the rods via room temperature stepper motors mounted to the top plate of the insert. Acting through vacuum feedthroughs, these stepper motors

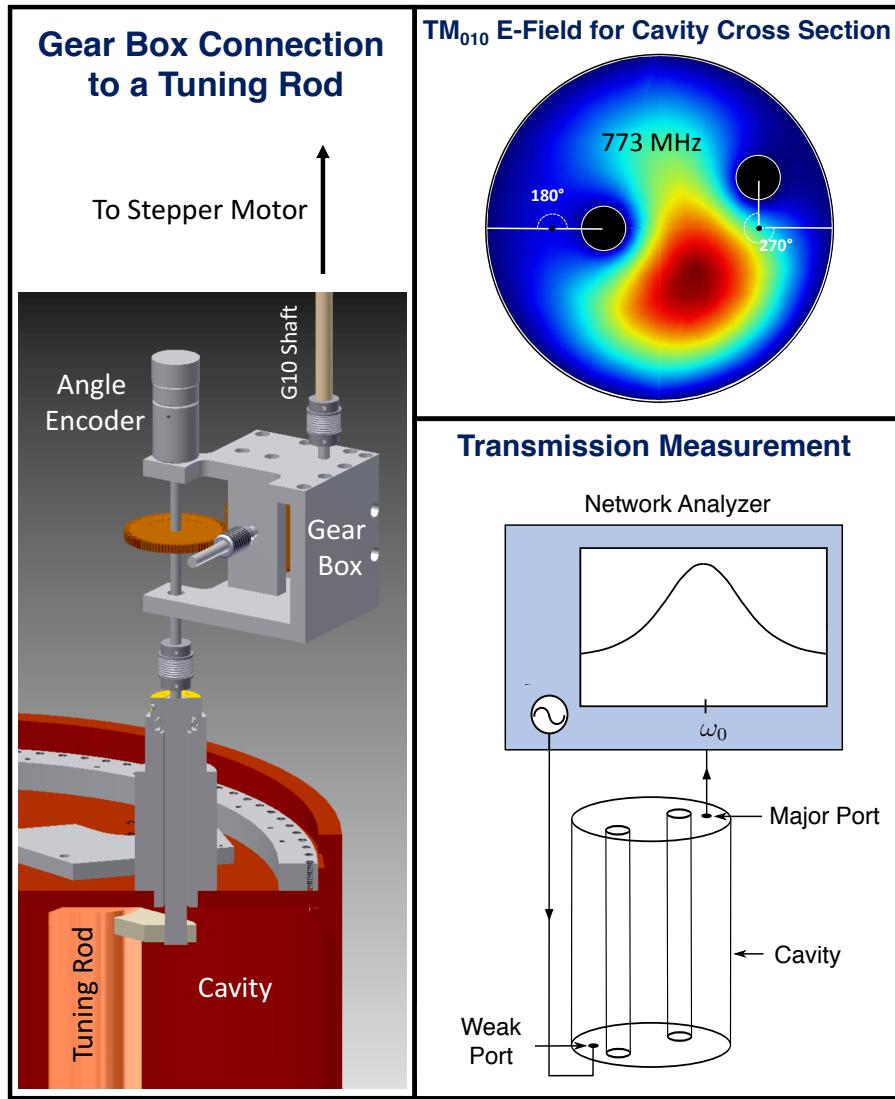


Figure 3.4: Tuning mechanism of the main cavity. Left: a room temperature stepper motor rotates a G10 shaft connected to a gear box. This motion undergoes a 1:19600 gear reduction and rotates the tuning rod in the cavity. Top right: rotating the tuning rods changes the boundary conditions in the cavity and alter the resonant frequency. Bottom right: the resonant frequency of the cavity is determined by injecting swept power from a network analyzer into a weakly coupled antenna. The swept power is picked up by a critically coupled antenna and a transfer function the cavity is measured.

communicate motion to long G-10 shafts which connect to gear boxes like the one shown in the upper right panel of figure 3.4. The gear boxes have minimal backlash a 19,600:1 gear reduction which allows for micro radian positioning. After every time the tuning rods are moved, the new resulting resonant frequency of the cavity is determined from a  $S_{21}$  swept response between two antennas inserted into the cavity. Power is injected from a network analyzer through a weakly coupled antenna on the bottom of the cavity and received by one of the critically coupled antennas used for data taking. The peak of this transmission response is a Lorentzian and is a measurement of the center frequency and Q of a given resonant mode.

### 3.2.3 Antenna Coupling

Apart from the fixed, weakly coupled antenna, used only as an input for calibrations and  $S_{21}$  measurements, each channel (of which there are two) has a variable depth antenna located on the top of the cavity. These two semi-rigid, coaxial antennas attach to gearboxes which turn rotary motion from room temperature stepper motors, into linear depth control. The depth of each antenna is adjusted to maintain a critical coupling to the  $TM_{010}$  and  $TM_{020}$  modes. When an antenna is critically coupled, the power escaping through the antenna equals the power lost in the walls of the cavity. At this coupling, the antenna is impedance matched to the cavity on resonance. On resonance, power incident on the antenna, looking into the cavity is fully absorbed and reflected off resonance. The coupling is determined by reflecting swept power off of the antenna and measuring its response. This is achieved with a network analyzer and a circulator shown in figure 3.5. The swept output of the network analyzer is directed by a circulator towards the antenna. Power reflected off the antenna then travels back to the circulator, up the receiver chain and back to the input of the network analyzer. A good impedance match is marked by a deep trough in the reflected baseline on resonance. The depth of the antenna is then adjusted to minimize the depth of this trough.

Conventionally, when the difference between the minima of the trough and the off-resonance baseline reaches -30dB the antenna is considered critically coupled. This means that only 0.1 percent of the on-resonance incident power is reflected.

It is important to keep the  $TM_{010}$  and  $TM_{020}$  channels (denoted Ch1 and Ch2) from interfering with each other. Low and high pass filters connect directly to the outputs of the antennas. This prevents power from the Ch1 mode from escaping through the Ch2 signal path and visa versa. This precaution doesn't completely isolate the channels. The EM boundary conditions of the cavity are not independent of the positions of the antennas. A localized Ch1 mode can travel up the Ch2 antenna and reflect off the Ch2 filter, such that the propagation length of the Ch1 mode is a function of the Ch2 antenna insertion depth. To ensure the ability to simultaneously couple the antennas, some trial and error is required in the antenna making and installing process.

### 3.2.4 *Cold Amplification*

The central focus of the experiment is to amplify power coming from the cavity and mix it down to frequencies low enough to then be digitized as shown in figure 3.6. Since each element of the receiver chain is an opportunity for added noise and a further attenuation of the already tiny signal, it is necessary that the first stage of amplification be high gain and very low noise. Given that the axion search window must be covered in a finite amount of time and that the scan rate of the experiment goes like  $\propto T_{sys}^{-2}$ , even low noise solid state amplifiers are not suitable for the first stage of amplification. The ADMX receiver is made to be “ultra-low-noise” with the use of Superconducting QUantum Interference Devices (SQUIDs) and Josephson Parametric Amplifier (JPAs).

A SQUID is a flux to voltage transducer where an increasing flux gives rise to a near-sinusoidal voltage output. By inductively coupling the signal current to the SQUID via a striplining resonator, an amplifier can be created if the device is flux biased where the

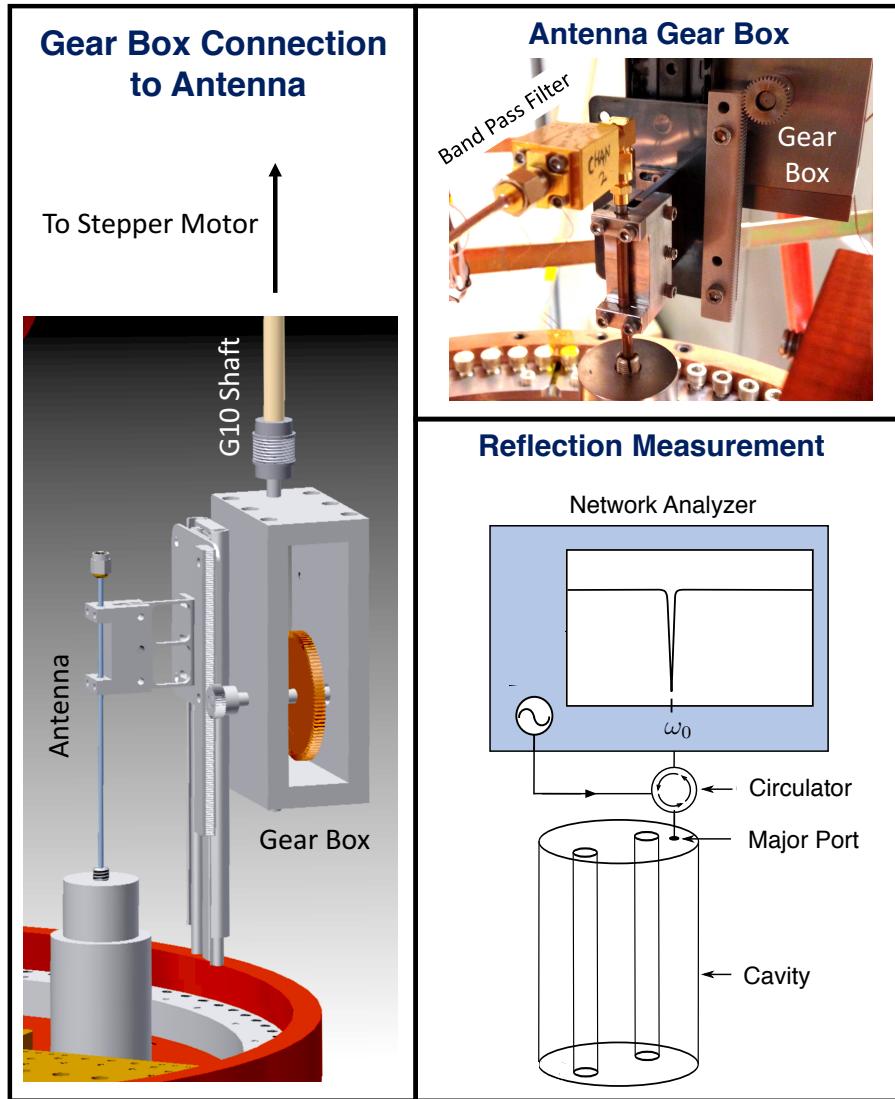


Figure 3.5: Antenna coupling mechanism for the main experiment. Left: a room temperature stepper motor rotates a G10 shaft connected to a gear box. Rotary motion is converted to linear motion which inserts or removes the antenna from the cavity. Top right: is an image of the linear gear box holding the coax antenna. Bottom right: the antenna coupling is determined by reflecting swept power from a network analyzer off of the antenna. If the antenna is critically coupled, all of the incident power on resonance will be absorbed into the cavity, as indicated by a deep trough in the reflection measurement.

# Simplified Receiver

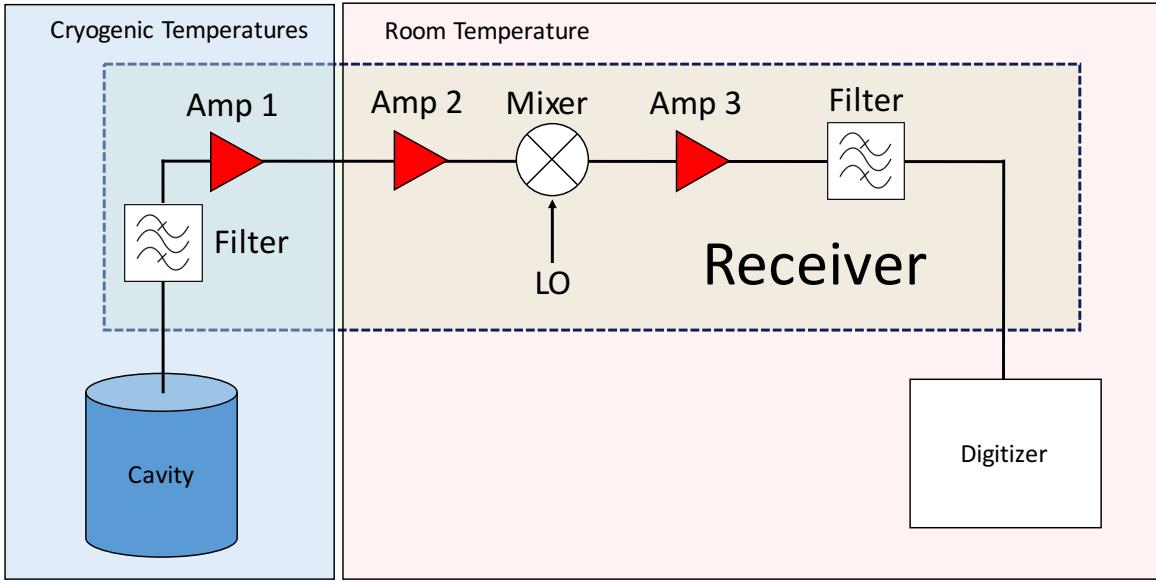


Figure 3.6: A simplified template for a receiver-chain schematic. Power is extracted from the cavity, filtered, amplified and mixed down to lower frequencies so that it can be digitized.

slope of the voltage is at its steepest. This type of amplifier is called a microstrip SQUID amplifier (MSA). An in-depth description of how a SQUID works can be found in [52]. A Josephson Parametric Amplifier is a LC circuit that uses a SQUID as a non-linear inductor. A pumped tone is injected into the circuit and mixes with the response of the SQUID, resulting in a high gain, low noise output. Both of these amplifiers are ultra-low noise amplifiers limited only by quantum mechanics. This means that their noise temperatures follow their physical temperature until running into the quantum limit on the order of a few tens of millikelvin. These cutting-edge devices are also expensive, finicky and require the field free region discussed in section 3.1.4 to operate.

After the first stage of amplification, any injected noise is suppressed by a factor of

$1/G_{1st\ Stage}$  which means that the low noise requirements on subsequent amplifiers in the chain become decreasingly stringent. Depending upon the channel and the data run in question, the second stage of amplification is provided by either High Electron Mobility Transistor (HEMT) amplifiers from the National Radio Astronomy Observatory or the commercial vendor Low Noise Factory. These devices have higher noise temperatures, which are on the order of a few Kelvin but are reliable and ready out of the box. After the second stage of amplification, the signals travel up through the top plate of the insert and encounter room temperature devices for the first time.

### 3.2.5 Room Temperature Receiver

The room temperature receiver is the set of amplifiers, circulators, filters, mixers and switches that define the signal path between the top of the insert and the digitizer. Directly outside of the top plate of the insert, a “post amp box” containing 2, in series, mini-circuits amplifiers give the signal a 40dB boost in gain to counteract 20-25dB of attenuation in the coaxial lines leading to the main receiver rack electronics (figure 3.7 ). When the signal reaches the rack, it gets further amplified, filtered and mixed down to 10.7 MHz so that it can be sampled by the digitizer. In addition to this critical task of digitizing data that may contain proof of an axion, the receiver also performs all of the other operations and measurements that must be in sync with and a part of data-taking. These operations include

- Moving the tuning rods
- Measuring the  $TM_{010}$  frequency
- Measuring the  $TM_{020}$  frequency
- Adjusting the antenna depths

## Main Receiver Rack



Figure 3.7: The ADMX dual-channel receiver rack.

- Measuring the Ch1 antenna coupling
- Measuring the Ch2 antenna coupling
- Performing noise temperature measurements
- Injecting synthetic signals for calibration

No two of these operations can be performed simultaneously without ruining the results of at least one measurement. The hardware that allows the receiver to be re-configured to

allow each of the RF measurements is a custom Arduino-controlled switchbox that I built with an undergraduate student, Skylar Wheaton. The software that decides when and how often each of these operations is performed is covered in chapter 4. Figure 3.8 shows a schematic for the full dual-channel receiver. This figure shows the both the warm and cold sections of the receiver in addition to the RF-lines and components associated with all of the supporting measurements mentioned above.

### 3.3 *Sidecar Axion Search*

This section introduces the experiment that is the focus of this work - the Sidecar experiment. It is a smaller size, higher mass ( $16 - 30\mu eV$ ) axion haloscope that fits inside the existing insert. It shares the benefits of the magnetic field and the cryogenics but is otherwise a stand alone experiment.

#### 3.3.1 *Motivation*

In the event that axions are ruled out in the 500 - 2000 MHz frequency range, ADMX will need to develop new ways of searching for axions outside of this window. While ADMX is investigating low mass axion detection techniques, the hunt for a higher mass axion will likely be the emphasis of future ADMX searches. There are several challenges associated with moving towards smaller cavity, higher frequency axion detectors. The biggest challenge is overcoming the loss in SNR proportional to the loss in cavity volume. The overall volume can mostly be made up with a multi-cavity system but not without paying the price of added complexity. The Sidecar experiment takes advantage of a small patch of available real-estate next to the dilution refrigerator and is an R&D testbed for new techniques, such as piezoelectric tuning, that may be implemented in future, more complicated systems.

### ADMX Dual-Channel Receiver-Chain

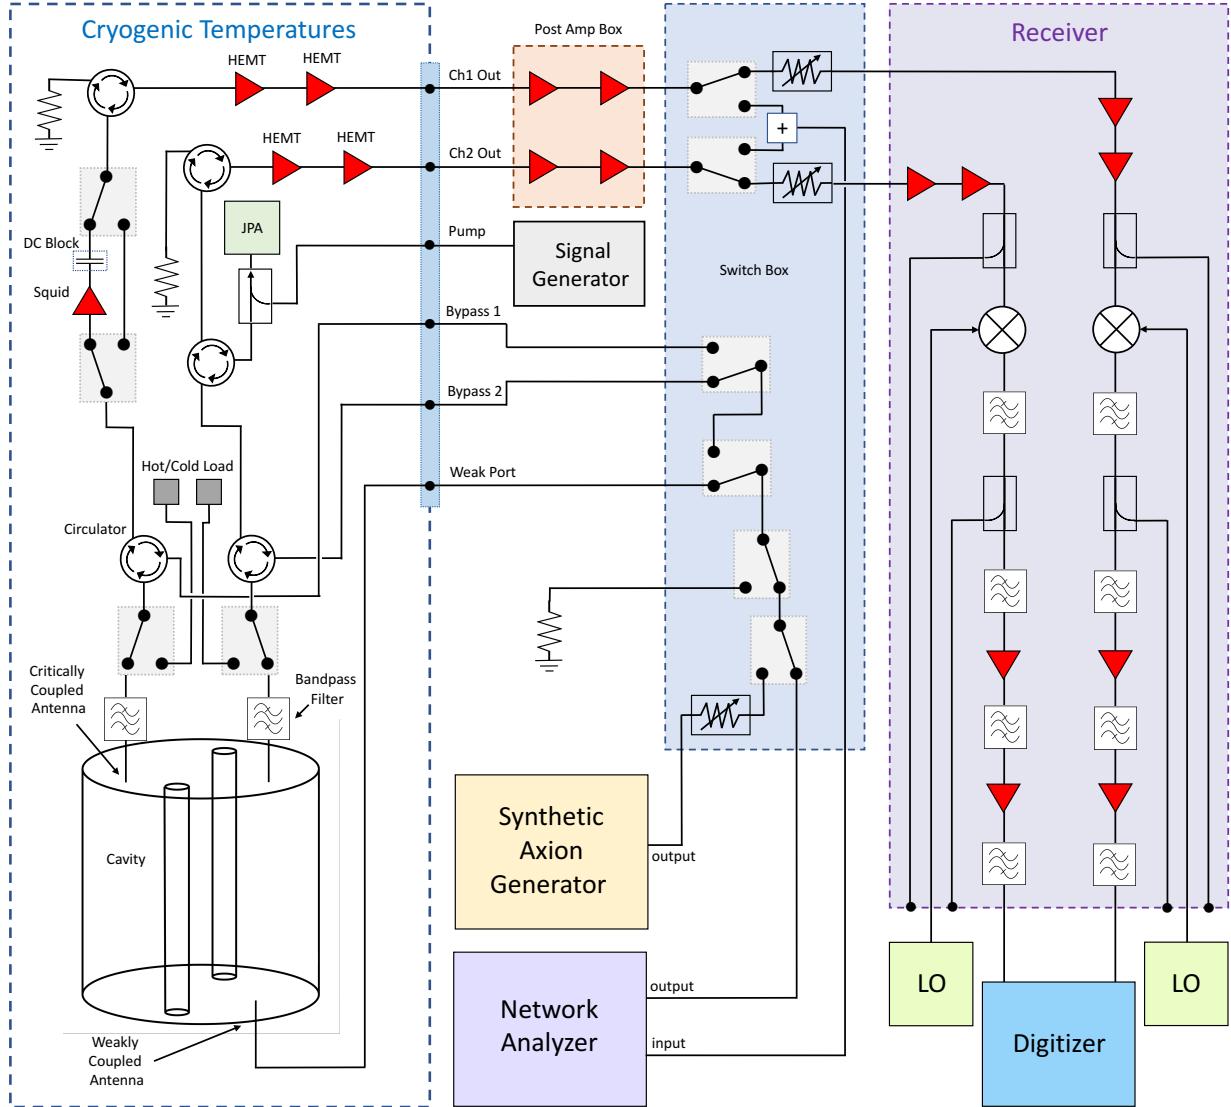


Figure 3.8: Schematic for the full ADMX dual-channel receiver-chain. The figure includes cold RF components and all of the wiring associated with the other functions and measurements performed by the receiver electronics.

### 3.3.2 *Cavity*

The Sidecar cavity was made by Gianpaolo Carosi at Lawrence Livermore National Laboratory (LLNL) and is basically a miniature version of the ADMX main cavity. It is 4.75" tall with a 2.5" inner diameter and a volume of 0.38 liters. Unlike the main cavity, the Sidecar cavity is equipped with only a single 1/2" diameter tuning rod to minimize complexity (see figure 3.9). A cavity of such a small volume has no hope of having a competitive SNR without the help of a quantum noise limited amplifier. With this in mind, the cavity and rod dimensions were specifically chosen to put the  $TM_{010}$  frequency range between  $\sim 4$  and 6 GHz which makes the experiment well suited to work together with a JPA. The end caps are identical and have three holes for the tuning rod, weak port and major port. At room temperature and without the tuning rod, the weakly-coupled Q ( $Q_0$ ) is  $\sim 22,000$ . With the addition of the rod,  $Q_0$  drops to approximately half that value.

### 3.3.3 *Piezoelectric Tuning/Coupling*

Haloscopes like the ones presented in this work require getting mechanical motion into the cryogenic space for the purposes of tuning the experiment and maintaining the antenna coupling. The main cavity achieves this with four room temperature stepper motors, connected to four G10 shafts which connect to four gear boxes. With no more available space on the top plate or in the insert, this system would not have been an option for the sidecar cavity. Besides, the gear boxes are large, they often jam and the worm gears create a heat leak which limits base temperature. For years, ADMX has been looking for a chance to explore the option of using piezo motors for tuning [39]<sup>1</sup>. The sidecar experiment presented the perfect opportunity to finally test this technology in the context of an axion search.

Piezoelectric actuators are electromechanical transducers. They turn a mechanical stress into a voltage developed across the material and likewise, applying a voltage across the

---

<sup>1</sup>Nicole Crisosto also worked on piezo R&D at LLNL

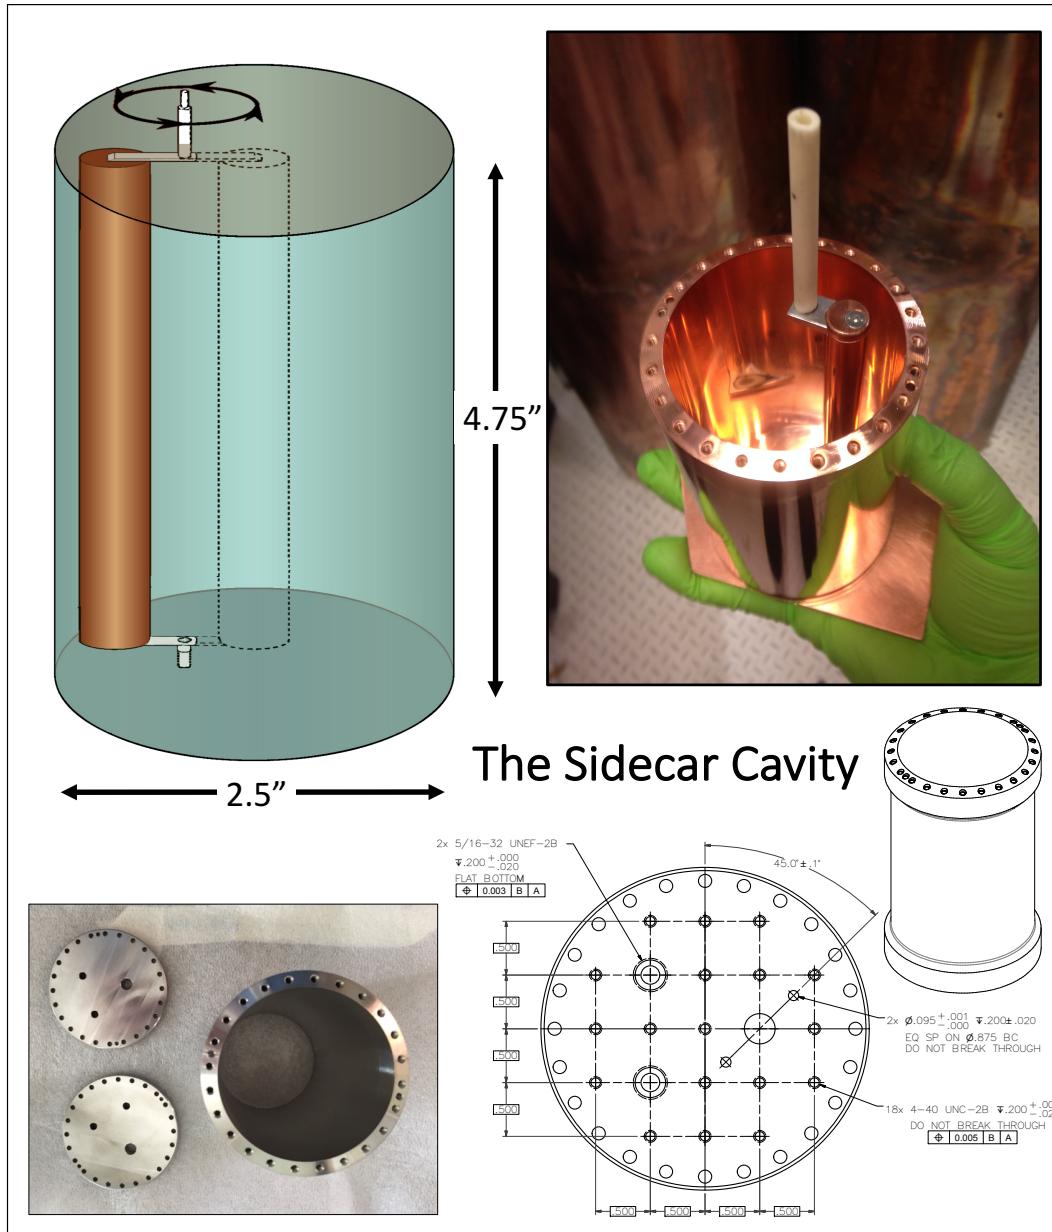


Figure 3.9: The ADMX Sidecar cavity. It is 4.75" inches tall with an inner diameter of 2.5" and a single 1/2" tuning rod. End caps bolt tightly down on a knife edge on the cylinder of the cavity to prevent ohmic losses which would degrade the Q.

material will cause a physical deformation. Depending upon how the actuator is initially cured, this deformation can be made to be tangential or parallel to an applied electric field. Piezoelectric actuators are typically made out of lead-zirconate-titanate (PZT) ceramics. These actuators can be used in conjunction with a stick-slip system to build motors that achieve either linear or rotary piezo motion. Two piezo motors were selected for the Sidecar experiment: the Attocube rotary ANR240/RES for tuning rod control and the Attocube linear ANPz101eXT12/RES for antenna depth adjustment. Both the rotary and linear motors are rated to operate at high vacuum in a 31T field at 10mK and come with their own encoder system.

The Attocube stick-slip system is shown in figure 3.10 and works as follows. The mobile component of the system is a table clamped tightly enough to a guiding rod such that static friction keeps the table fixed unless the rod is accelerated quickly. Attached firmly to the guiding rod is a piezoelectric actuator which is in turn attached to a fixed frame. The voltage on the piezo is then slowly increased such that the table moves with the motion of the rod without slipping. The voltage is then removed, causing the actuator to quickly contract and the rod to retract fast enough such that the inertia of the table keeps it stationary. This action describes a single step of the motor but the motion can be made nearly continuous by supplying the piezo with a voltage sawtooth pattern. The speed and step size of this continuous motion can be altered by changing the peak voltage (V) and driving frequency ( $f_d$ ) input parameters. For fine tuning, a DC voltage can be applied to the actuator, causing it to expand and move the table without slipping. At cryogenic temperatures, the actuator becomes less effective and will not expand and contract as much. This degrades the performance of the motor and requires that the peak voltage and driving frequency be increased. The upside to this consequence is that a higher tuning resolution can be achieved with the decreased piezo step size. Studies on the Sidecar setup have concluded that at temperatures between 150mK and 200mK result in a step size that is  $\sim 1/10$  the room temperature step

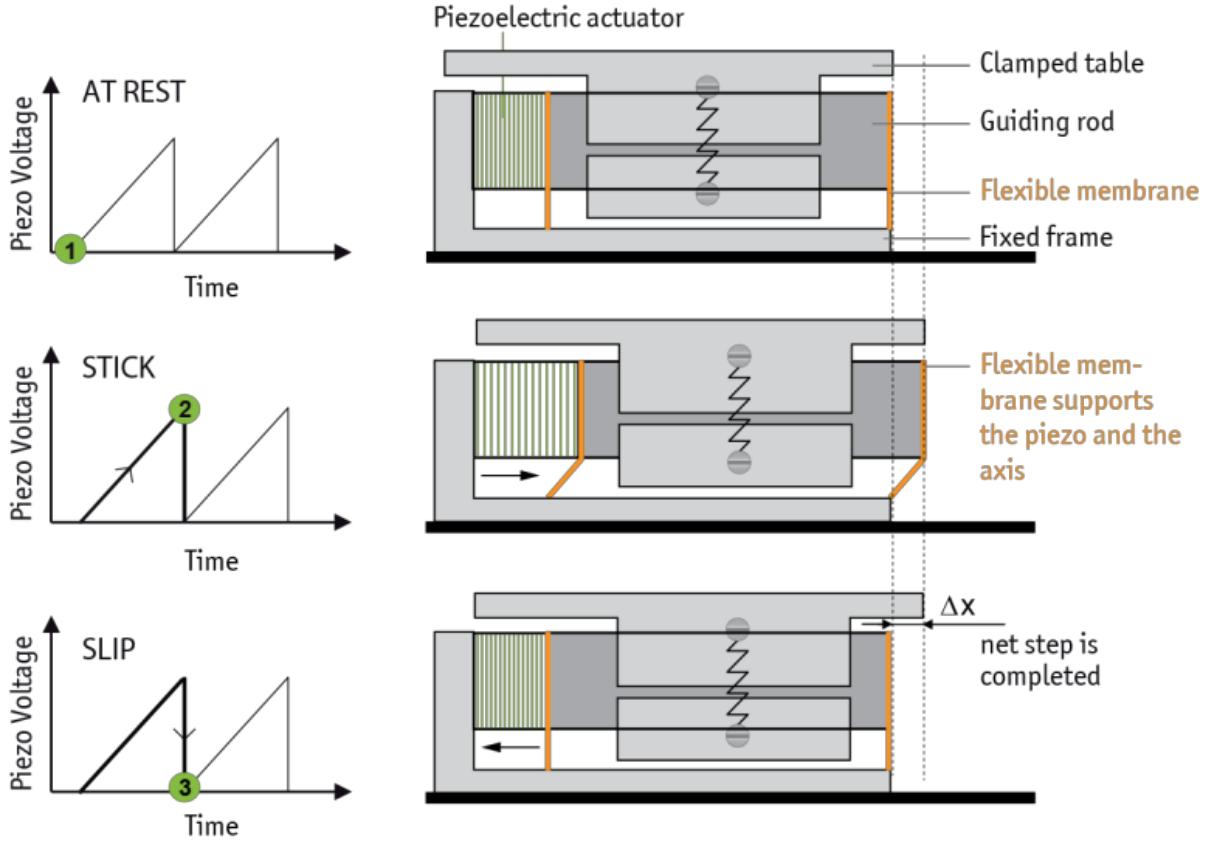


Figure 3.10: The Attocube stick-slip system. A sawtooth voltage is provided to the piezo actuator. The actuator extends and contracts and causes a clamped table to scoot along a guiding rod [53].

size.

Figure 3.11 shows the motors mounted on a 4" x 4" x 1/2" copper block attached to the top of the sidecar cavity. On the left, the rotary piezo is shown screwed down to the copper block and gripping the alumina rod armature with a set screw. Next to the rotary piezo are two antenna ports. One is for the fixed, weakly coupled antenna and the other is for the adjustable, linear piezo driven antenna. The linear motor is mounted upside-down on a copper bracket and holding an antenna/counter-weight system. When the the motor is

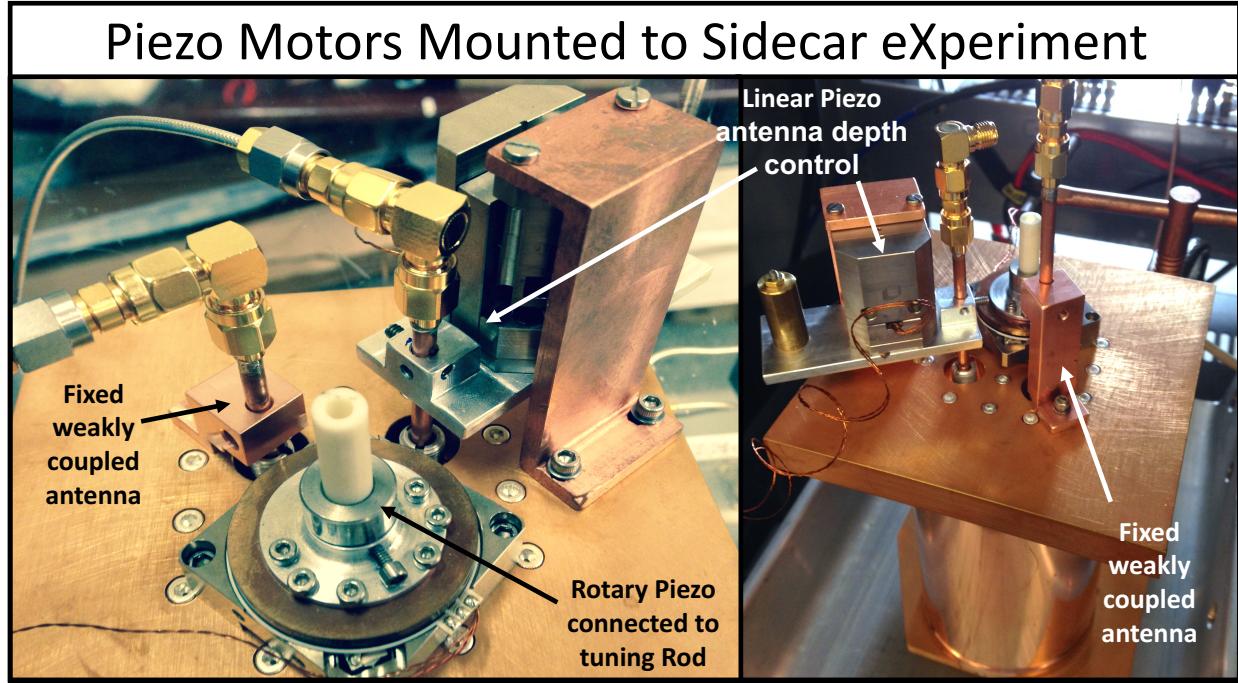


Figure 3.11: Rotary and linear piezo motors mounted on top the the Sidecar cavity.

extended, the antenna is inserted into the cavity. The counterweight was suggested by Attocube engineers to counteract the torque from the antenna.

At room temperature, the rotary piezo has a step size of  $\sim 0.01^\circ$  (with peak voltage = 60V,  $f_d = 1\text{kHz}$ ). This step size falls to  $\sim 0.001^\circ$  at 150mK. Figure 3.12 shows the Sidecar experiment taking single steps over a couple of days while taking data on the  $TM_{020}$  mode. The curve is made up of 1,165 individual steps, each of which shift the frequency of the  $TM_{020}$  by approximately 13.5 kHz. The lower right hand side of the figure shows a test of the DC fine tuning capability. The system started with 0V on the actuator and without tuning, measured the mode frequency for an hour. At that point, 70V (the maximum voltage) was put on the piezo, the system was allowed to continue to take data and 30 minutes later the voltage was set back to 0V. The conclusion is that maximum range of the DC fine-

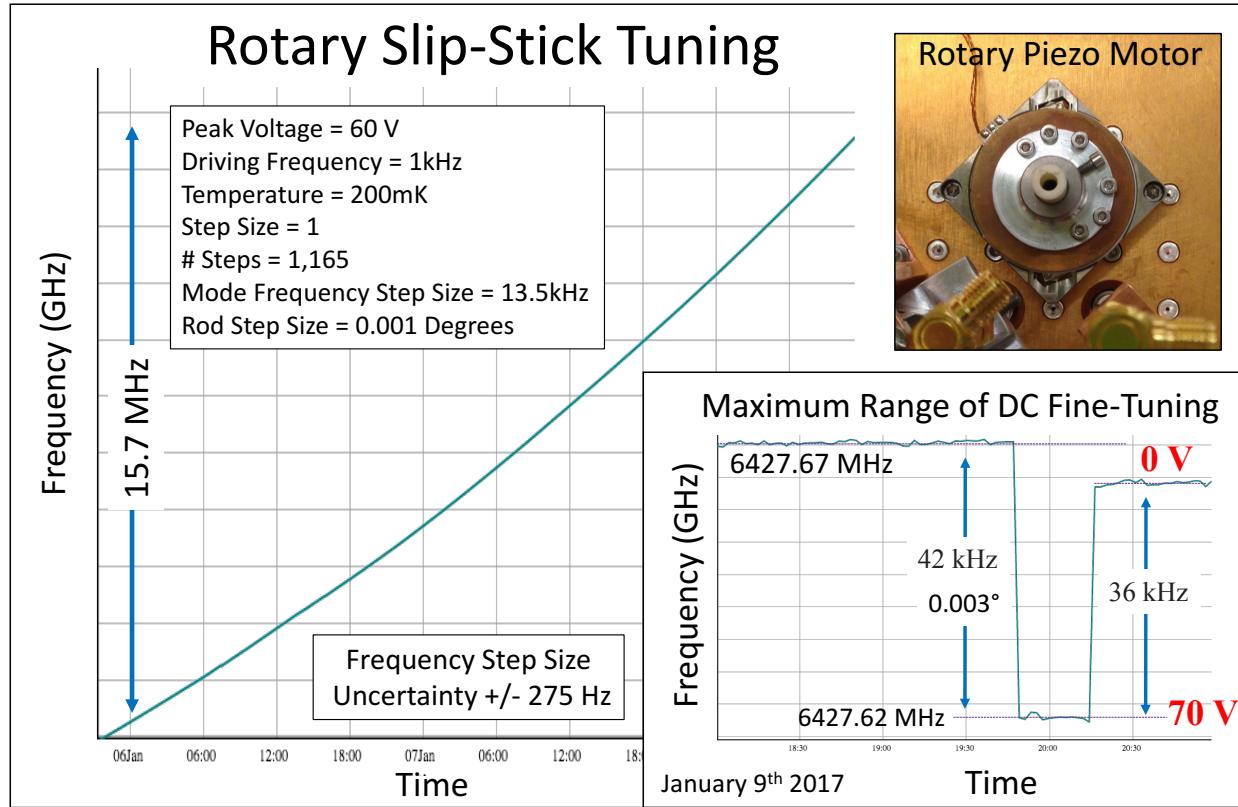


Figure 3.12: A demonstration of the rotary piezo's ability to smoothly move the tuning rod and shift frequency of the cavity. Plots show resonant frequency versus time. Left: the cavity is tuned with stick-slip motion. Right: fine-tuning is achieved by flexing the actuator with a DC offset.

tuning function is  $\sim 0.003^\circ$  which resulted in a mode frequency shift of  $\sim 40\text{kHz}$ . The test also confirmed the claim made by Attocube engineers that clockwise and counter-clockwise rotations are not exactly the same. For this given mode structure, sub kHz positioning could then be achieved by applying a DC voltage  $< 5\text{V}$ . The test also provided a measurement of the noise or frequency jitter from measurement to measurement to be  $\pm 275\text{ Hz}$ . An ultra-high Q cavity search might not be able to tolerate this level of precision but luckily cavity Qs degrade at higher frequencies, the width of the resonant mode washes out jitters at this

level and smooths the SNR.

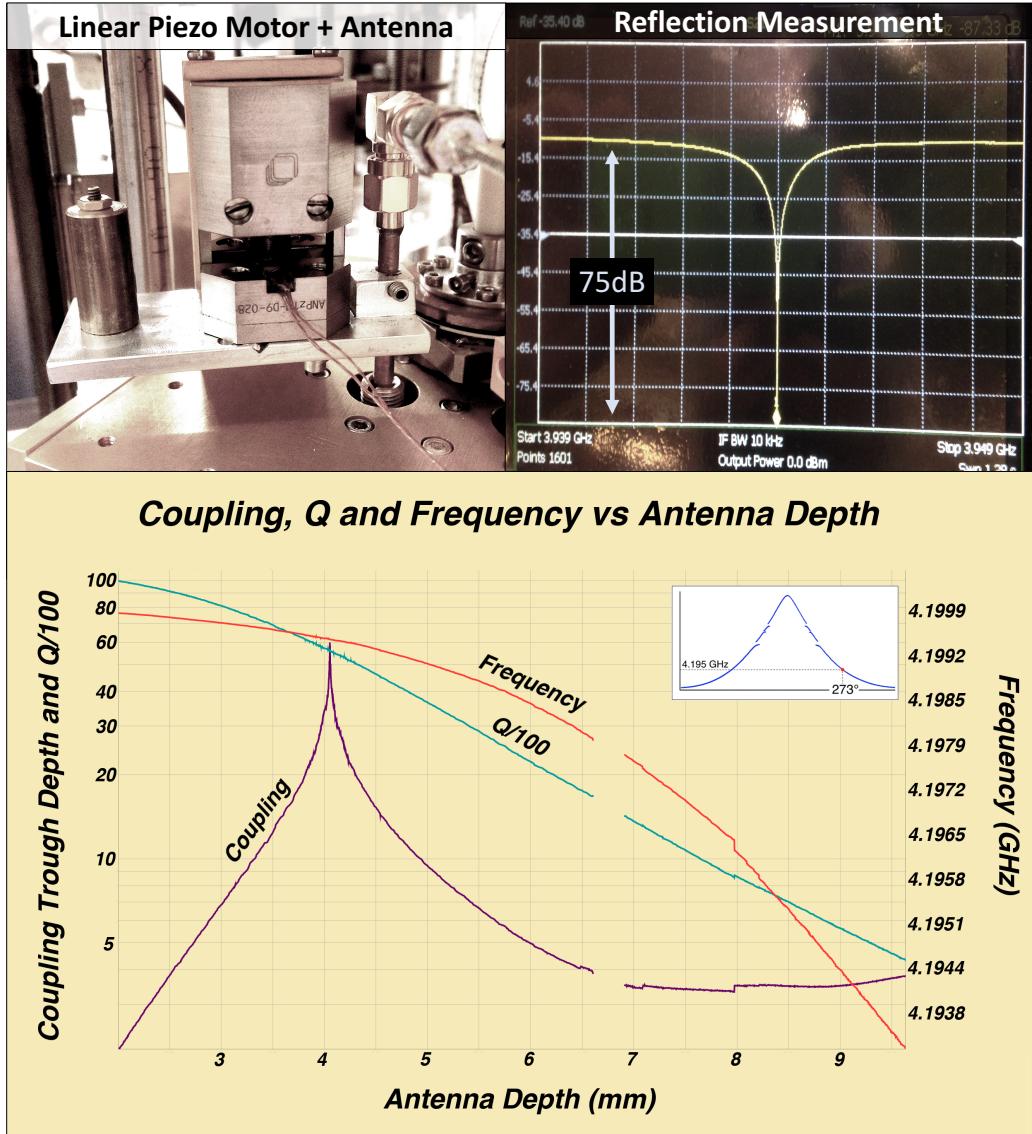


Figure 3.13: Top left: linear piezo motor mounted to the top of the Sidecar cavity and holding the antenna. Top right: reflection measurement where swept power is bounced off the antenna. If the antenna is critically coupled, all of the power on resonance will be absorbed into the cavity, resulting in a deep trough the measurement. Bottom: coupling, Q and cavity frequency versus antenna depth where the coupling trough depth is expressed as a positive number.

The linear piezo, shown in the upper left hand corner of figure 3.13 has a 12 mm travel length and by adjusting the antenna depth, critical coupling can be achieved following the procedure described in section 3.2.3. The upper right hand side of the figure shows a reflection measurement with an absorption trough that is 75dB below the baseline, indicating a remarkably good coupling. The bottom of the figure shows how the mode frequency, loaded Q ( $Q_L$ ) and antenna coupling vary with antenna depth. To display all of these quantities on the same graph,  $Q_L \rightarrow Q_L/100$  and the coupling trough depth (sometimes expressed as a negative number) has been represented by its absolute value. Here, the peak coupling value represents critical coupling and also corresponds to where  $Q_L = Q_0/2$ . The measurement shows that the antenna positioning works well and can be used to dial in the proper coupling. It also gave insights into writing an automated antenna coupling algorithm for the piezo motor discussed in section 4.4.6. The masked data between an antenna coupling of 6 and 7 mm helped identify a kink in the antenna where, at that depth, the antenna was not being properly grounded to the walls of its feedthrough.

There were two options for the heat sinking of the piezo motors. Given that the Sidecar cavity is directly sunk to the mixing chamber of the dilution refrigerator, the first option was to isolate the motors from the cavity and tie them directly to the 1K plate just inches above. However, in this scheme, a huge amount of heat could potentially be dumped through the rod and antenna, into the cavity and cryogenics. The consensus was to opt for the second option which was attach the motors directly to the cavity, heat sink the wiring to each stage of the dilution refrigerator and be careful not to over work the piezos. This choice meant that the piezos could not be run continuously for a fraction of a second without being quickly noticed as a huge temperature spike by the dilution refrigerator sensors. Piezo heating comes from three sources:

- Dissipated Energy =  $CV^2\tan(\delta)$  per step

- Mechanical Work =  $F \cdot x$
- Power from Leakage Currents =  $V^2/R$

where C is the capacitance of the actuator,  $\tan(\delta)$  is the loss tangent (typically 0.02 - 0.04), F is the force on the clamped table, x is the distance the table travels and R is the resistance of the element. The last two sources of heat can be ignored because, typically the mechanical heating is  $\sim 500\text{nJ}$  per step and the resistance of the piezo is  $\sim 10^{10}\Omega$ . Treating only the dominant term, and using typical parameters ( $C = 200\text{nF}$ ,  $\tan(\delta) = 0.03$ ,  $V = 50\text{V}$ ) the energy dumped into the system is  $15\mu\text{J}$  per step. For the rotary piezo, this number is doubled because it contains two actuators. Under normal operating conditions, the rotary piezo takes 1-10 steps per 100 seconds which means the power dissipated is  $.3\text{-}3\mu\text{W}$ . This does not pose a problem for a dilution refrigerator designed to have  $800\mu\text{W}$  cooling power. However, if the motor were run continuously for 3 seconds at  $f_d = 1\text{kHz}$  to move the rod 30 degrees for an axion candidate rescan or to reconfigure the mode structure to accommodate a change in the data run plan, the mixing chamber would receive 100mJ of heating. With good planning and caution, heating was not a major problem for the Sidecar experiment. However, a multi-cavity system needing many micro-adjustments to keep the cavities coupled and mutually tuned to the same frequency will require careful thought.

### 3.3.4 Cold Electronics

While the Sidecar cavity was designed to pair well with a JPA, the goal of the first data run was to demonstrate that the experiment was functional and could be sensitive to axion-like particles. For the data discussed in this work, the signal path from the cavity was routed straight to a cold switch in the SQUIDadel, through a circulator and up to a Low Noise Factory HFET amplifier (model LNF-LNC1-12A *s/n* 455B) which was thermally sunk to the top of the reservoir. These devices are stable, broadband, high gain and low noise (4K).

At this stage, the signal received 40dB of gain, and was then routed to the top of the insert.

### 3.3.5 *Receiver*

Like the main experiment, the Sidecar signal receives a 40dB boost in gain at the top of the insert from two ZX60-83LN+ low noise mini-circuits amplifiers to counteract attenuation in long coax lines leading to the receiver. At the receiver, the weak port, bypass and signal out coax lines from the insert connect to a switch box which directs the flow of RF traffic. The switch box is a mini-circuits USB RF-SPDT Switch Matrix and contains 8 RF single-pole double throw switches. Figure 3.14 shows how the switch box is configured to allow the receiver to digitize incoming data with and without a calibrating input signal and make cavity transmission and antenna reflection measurements. The design also considered lab sociology that might cause the receiver system to be dismantled for a network analyzer measurement. When the switch box is set to its default/digitizing position, manual access to the network analyzer is made available on a custom front panel of the rack.

When the switch box is set to digitize, all inputs to the cavity are killed by a 50 ohm terminator and the signal output from the cavity is routed to the room temperature receiver (figure 3.15). This receiver is a custom box containing mini circuits (ZX60-8008E + and ZFL-500LN+) low noise amplifiers, band pass filters and a Polyphase IRM3570B image-reject mixer which mixes the  $\sim$ 5GHz signal down to 10.7 MHz to be digitized. A full schematic of the full receiver chain is shown in figure 3.16. After the signal is digitized, the power spectrum and all other relevant data is stored in an SQL database. A description of how the signal is processed can be found in chapter 5.

## 3.4 *Calibration*

If a candidate signal were to be detected by the experiment it would be very easy to determine whether or not the signal was an axion or something else. An axion signal would modulate

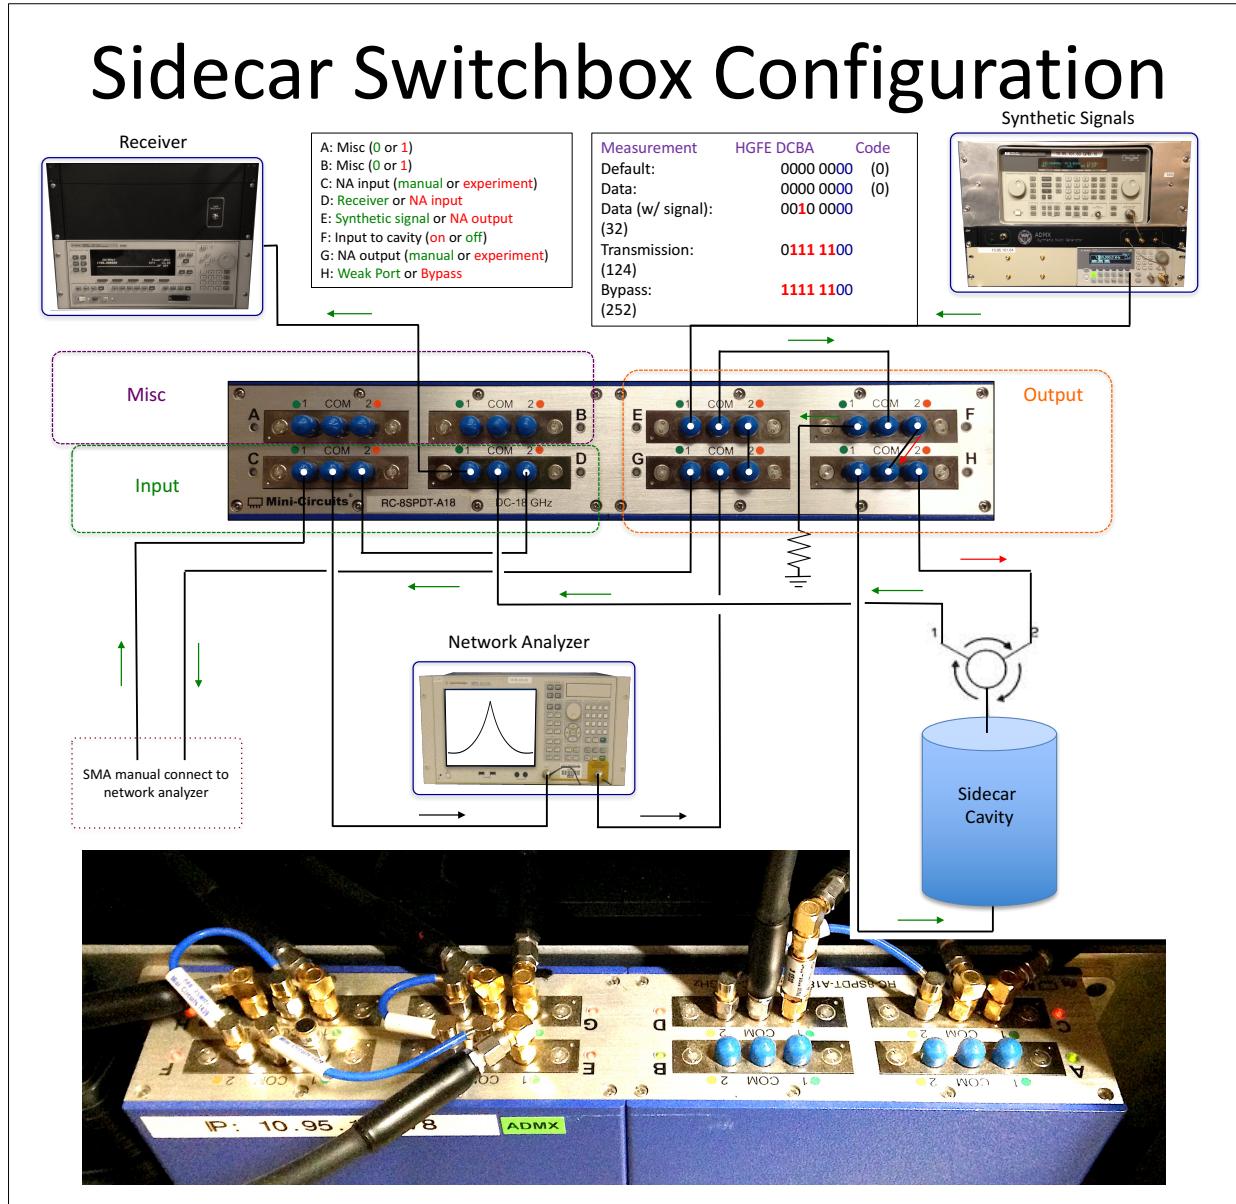


Figure 3.14: Sidecar switchbox schematic showing each of the possible measurements performed by the receiver electronics.

with the rotation of the earth and also vanish when the magnetic field was turned off. In the event that no signal is detected, the task of setting a new limit on  $g_{\alpha\gamma\gamma}$  is much more

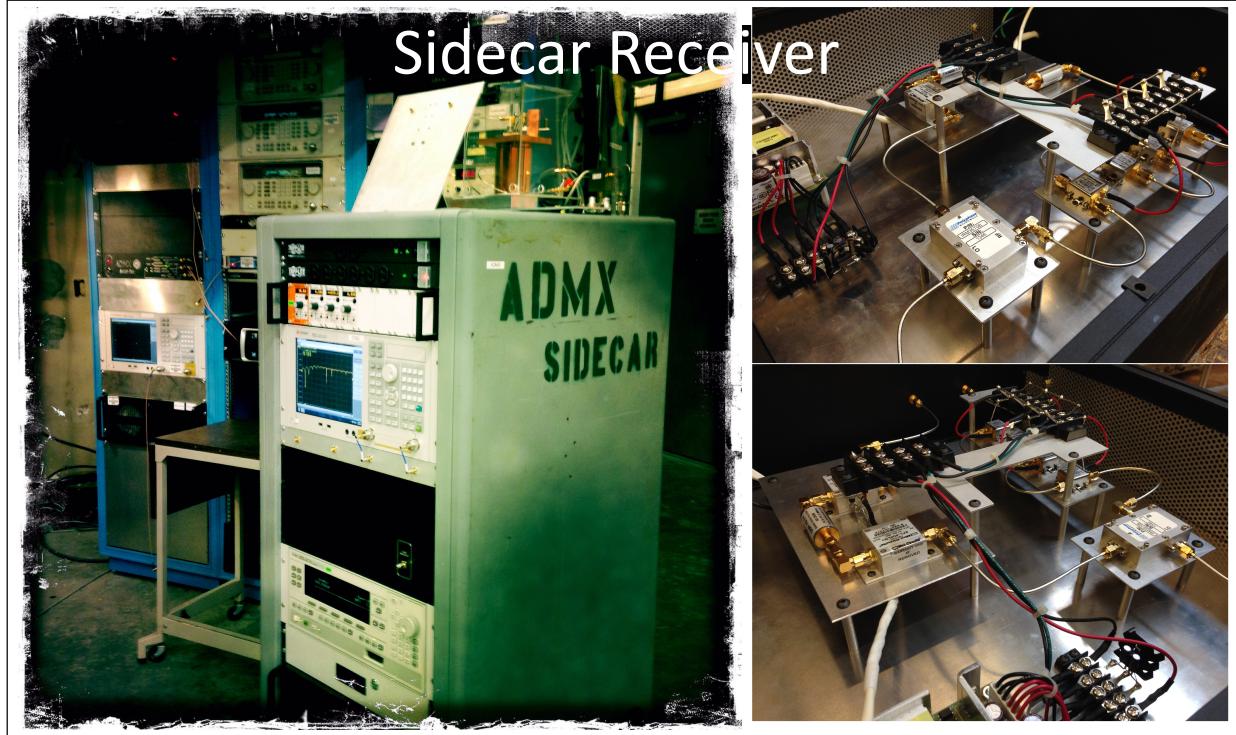


Figure 3.15: Images of the Sidecar receiver. Left: Sidecar receiver rack. Right: Inside of the heterodyne mixing stage of the receiver.

difficult. This section discusses some of the calibration issues that can hide a real signal and yield a null result.

#### 3.4.1 Noise Temperature

ADMX's largest systematic uncertainty is noise temperature. Amplifiers add their own noise into the signals they amplify and this noise power spectral density can be represented as a temperature. If an amplifier turned out to be noisier than believed, the resulting increase in the output power fluctuations would be rescaled by the receiver structure removal processed and go unnoticed. This could lead to a missed opportunity for a discovery or the setting of an unrealistically low limit.

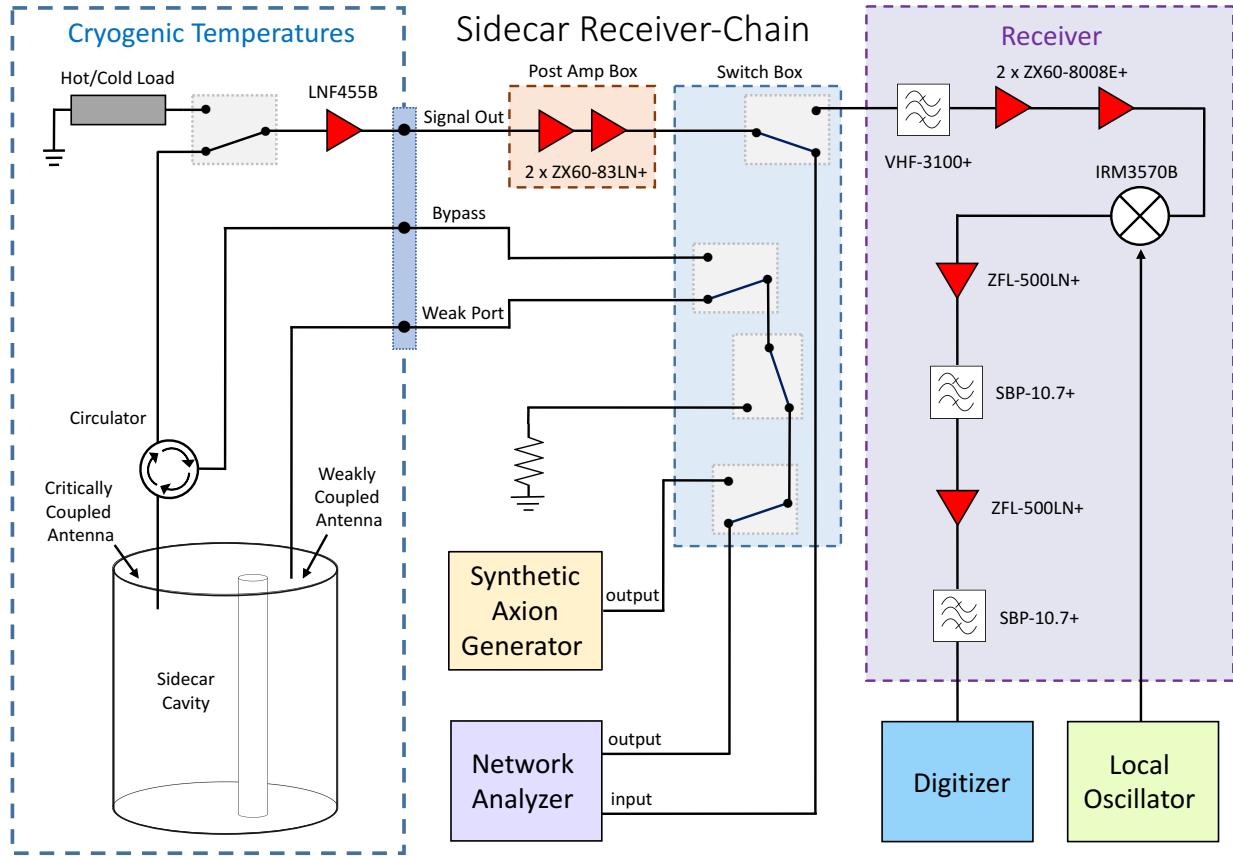


Figure 3.16: Schematic for the full Sidecar receiver-chain. The figure includes cold RF components and all of the wiring associated with the other functions and measurements performed by the receiver electronics.

ADMX measures noise in several ways. One method uses a cryogenic switch and a heatable attenuator in the cold space. By switching the input to the first amplifier away from the cavity and towards the attenuator, the resulting output power should just be  $P = GBkT_{sys}$  where  $P$  is the output power of the receiver,  $G$  is the full gain of the receiver chain,  $B$  is the bandwidth,  $k$  is Boltzmann constant and  $T_{sys} = T_{physical} + T_N$ . By measuring the physical temperature and output power for both cases when the attenuator is “hot” and “cold”, the two equations

$$P_{hot} = GBk(T_{hot} + T_N), \quad (3.1)$$

$$P_{cold} = GBk(T_{cold} + T_N), \quad (3.2)$$

can be solved in terms of known or measurable quantities.

$$T_N = \frac{P_{hot}T_{cold} - P_{cold}T_{hot}}{P_{cold} - P_{hot}} \quad (3.3)$$

### 3.4.2 Compression and Clipping

If the receiver-chain is given too much gain, it could cause amplifiers to compress or the digitizer to clip. This could cause an axion signal to be amplified less than expected and drive power from compressed peaks into other Fourier modes. The fear is that if there is insufficient filtering, broadband power not relevant to the axion search but still propagating through the system may cause amplifiers to saturate. The presence of compression can be tested by injecting both a signal and broadband noise with a known SNR into the input of the isolated receiver and comparing it to the SNR at the output. With the full receiver connected to the rest of the experiment, a similar *in situ* test can be done by injecting a series of tones into the weak port of the cavity. Here the real noise of the experiment takes the place of the noise source in the previous test, but in this case the input SNR is unknown. A plot of output versus input power should be linear over some range and then at higher powers, the line will turn over where compression kicks in. As long as the experiment runs in the linear region, compression should not be an issue. The ultimate test of the system is to recreate a signal that mimics a real axion and see if it propagates through the full receiver-chain and analysis pipeline to trigger a detection.

### 3.4.3 Synthetic Axion Generator

My first contribution to ADMX was a Synthetic Axion Generator that injects axion-like signals into the resonant cavity for the purposes of calibration. The generator is made up of three components: an arbitrary waveform generator, a mixer and a local oscillator (shown in figure 3.17). The resulting signal starts off as low frequency waveform that is then mixed up to the desired frequency. In other words, the waveform generator sets the width and shape of the line width and the mixing stage sets the axion mass.

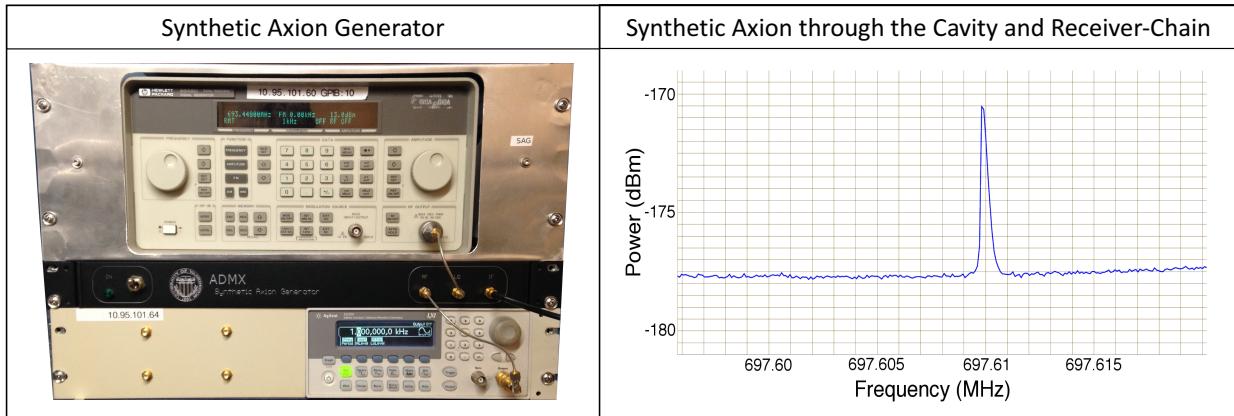


Figure 3.17: Left: synthetic axion generator (SAG) which outputs a signal similar to an axion lineshape. Right: synthetic axion signal detected in a digitized power spectrum.

The output of the arbitrary waveform generator is determined by a list of voltages uploaded to the device. To create that “list”, a frequency representation of the expected axion line shape needs to be Fourier transformed to create a real time-series waveform. Before doing so, the power spectrum is artificially shifted to higher Fourier bins. This causes the start of the waveform to be non-zero such that the resulting signal does not sit directly on top of the blinding local oscillator tone when the signal is mixed up. After the voltage list has been uploaded, the frequency of the waveform generator must be set to the desired frequency resolution of the final signal. This is due to the inverse relationship between fre-

quency resolution and time-series intervals. If the frequency of the generator is set too high, the resulting axion signal will look like a discrete set of delta functions. If the frequency is set too low, the local oscillator tone will no longer be out of band. With the right input parameters, a signal like the one in figure 3.17 can be achieved.

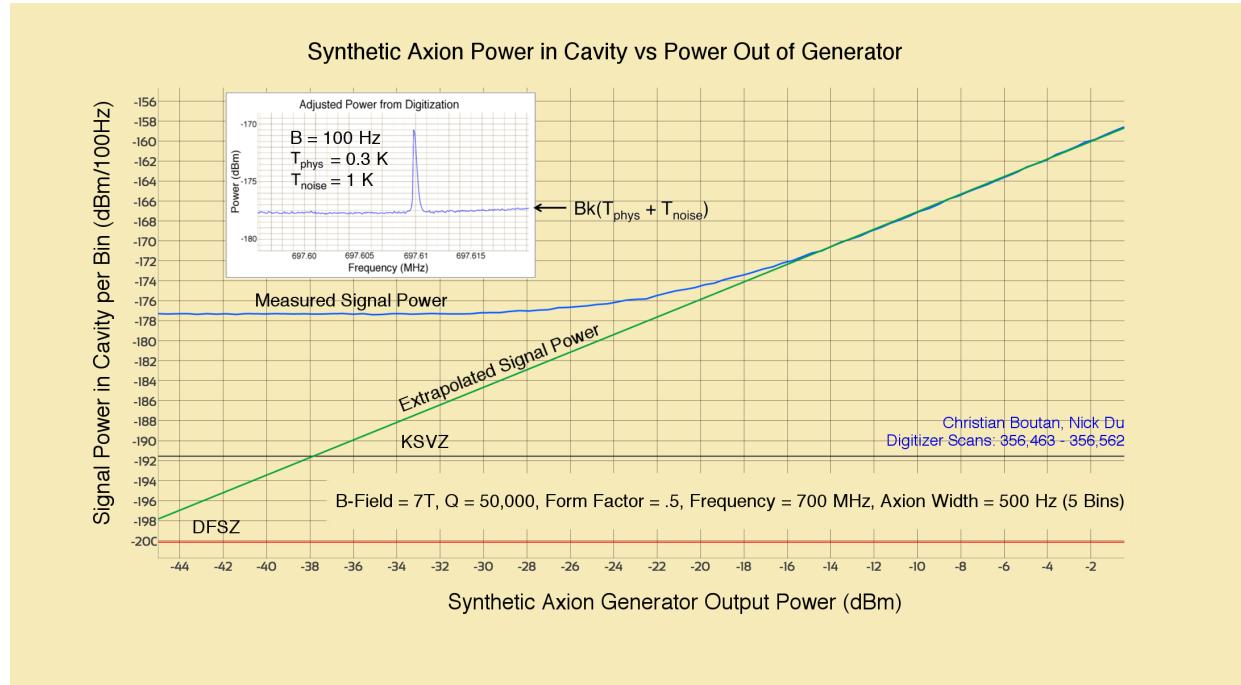


Figure 3.18: Signal power detected in a power spectrum versus the power output from the SAG (blue). At low powers, the signal is beneath the noise corresponding to the flat region of the blue line. As the power is increased, the linear regime of the blue line indicates that the signal is above the noise but not is not causing the amplifier chain to compress. An extrapolation from this linear regime (green) maps the relationship between power output by the generator and power that makes it into the cavity.

Due to the change in attenuation in the cables and reflection off the weakly coupled port after the insert is cooled, it is impossible to know precisely how much power makes it into the cavity. In order to map the relationship between the power out of the generator and the power that makes it into the cavity, non-realistic, high power axions must be injected. If

$T_{sys}$  is known, fluctuations in the power around the signal can be used to calibrate the peak. This overpowered signal can then be attenuated to the desired power below the noise. In the style of the compression test discussed in section 3.4.2, figure 3.18 shows this measurement for 200 different input powers. The blue line is the value of the peak power measured in the cavity versus the power output by the generator. Where the signal is below the noise, only noise is being measured and the line is flat. As the signal creeps above the noise the relationship between the input and output power is linear. An extrapolation from this linear region provides a conversion between desired axion power and required generator output power. The hardware and software of the synthetic axion generator is already in place and has been used successfully for calibration purposes and to test the analysis pipeline. Future blind injections are planned.

## Chapter 4

### DATA-TAKING: NAVIGATING CAVITY MODE STRUCTURE

In chapter 3 I outlined the hardware and anatomy of the ADMX experiment. This section describes how the data-taking software brings the experiment to life. The experiment relies on the mechanism discussed in section 2.4 where local axions in the milky-way halo interact with a magnetic field and convert to photons, thereby depositing a small amount of power into a resonant mode in the cavity. This conversion only happens for non-zero form factor modes where form factor, defined in equation 2.24 is the overlap the electric and magnetic fields integrated over the volume of the cavity. The objective of data taking is to digitize the signal from one or more of these resonant modes in hopes of detecting a small surplus in power from converted axions. This inherently narrow-band search which requires the high Q mode be tuned to the axion mass can be made wideband with a tunable cavity. A discussion about how resonant modes tune, mix and interact is integral to an explanation of the data taking process. This section begins with an introduction to resonant cavities and builds to a description of the ADMX data taking software which automates the axion search through a complicated mode structure.

#### 4.1 *Cavity Basics*

A typical microwave cavity is a simple stainless steel cylinder plated in OFHC copper and sealed with two end caps. I begin the resonant cavity discussion with a bare, untunable cavity.

#### 4.1.1 Analytical Solutions

Starting with a bare cavity that is axially oriented in the  $z$  direction, resonant modes in the cavity can be obtained analytically by solving the Helmholtz equation

$$\nabla^2\Phi + k^2\Phi = \nabla^2\Phi + (\mu\epsilon\omega^2 - k_\perp^2)\Phi = 0 \quad (4.1)$$

where  $\mu$  is the permeability,  $\epsilon$  the permittivity,  $\omega$  the frequency and  $k_\perp$  the eigenvalue for components transverse to the  $z$  direction. Depending upon the boundary conditions that are enforced, bare cavities yield two different kinds of solutions: there are transverse electric (TE) modes where the electric field is always in the plane that is orthogonal to the  $z$  direction and transverse magnetic (TM) modes where the magnetic field is always in the plane that is orthogonal to the  $z$  direction.

TE modes are associated with electrical currents around the azimuthal direction that give rise to a magnetic field in the  $z$  direction. Since the electric field is orthogonal to  $\hat{z}$ , and therefore the static magnetic field from the solenoid magnet, the dot product  $\vec{E} \cdot \vec{B}$  in the form factor is zero. With no coupling to the axion field, these modes are harmless bystanders when out of the band where data is to be taken but an annoyance when they are near enough to interfere with high form factor modes.

TM modes are associated with electrical currents flowing along the length of the cavity and have the form

$$\vec{E} = E_0 J_m \left( \frac{x_{mn} r}{R} \right) e^{-i\omega t} \hat{z} \quad (4.2)$$

$$\vec{B} = -i\sqrt{\epsilon_r} E_0 J_n \left( \frac{x_{mn} r}{R} \right) e^{-i\omega t} \hat{\phi} \quad (4.3)$$

where  $R$  is the cavity radius and  $x_{mn}$  is the  $n$ th root of the Bessel function  $J_m(x) = 0$ . The frequencies of these modes can be solved analytically [54]

$$f_{mnp}^{TM(TE)} \equiv \frac{1}{2\pi\sqrt{\mu\epsilon}} \sqrt{\frac{x_{mn}^{(')2}}{R^2} + \frac{p^2\pi^2}{d^2}} \quad (4.4)$$

where  $f$  is the frequency of the TM (or TE) mode,  $d$  is the cavity length,  $p$  is an integer and  $x_{mn}^{(')}$  stands for either  $x_{mn}$ , the  $n$ th root of the Bessel function  $J_m(x) = 0$  or  $x'_{mn}$ , the  $n$ th root of the Bessel function  $J'_m(x) = 0$ . Figure 4.1 shows the electric field for 6 example resonant modes in a cross section of the bare Sidecar cavity.

Keeping in mind that the field from the ADMX superconducting magnet is oriented vertically, it is easy to work out which of these modes give rise to a zero or non-zero form factor. It is clear the  $TE_{111}$  (and all TE modes) in (panel a) has a zero form factor because  $\mathbf{E} \cdot \mathbf{B} = 0$  at all locations in the cavity. Panels b, c and d show TM modes where most locations in the cavity have  $\mathbf{E} \cdot \mathbf{B} \neq 0$  but still have no coupling to axions because the form factor dies in the integral over the volume. Finally, panels e and f show the  $TM_{010}$  and  $TM_{020}$  modes that have non zero form factors, and if tuned to the mass of the axion, may contain a tiny surplus of energy from the conversion of axions to photons. By studying Figure 4.1, it is easy to convince yourself that only  $TM_{0n0}$  modes have a sensitivity to axions.

The expected power deposited by axions is proportional to the quality factor of the cavity. This depends upon the stored energy in the cavity mode divided by the amount of loss per cycle

$$Q = \omega_0 \frac{U}{P_{lost}} \quad (4.5)$$

where  $\omega_0$  is the frequency of the mode,  $U$  is the stored energy and  $P_{lost}$  is the time averaged power lost. For this reason, a good electrical connection between the cavity and end caps is

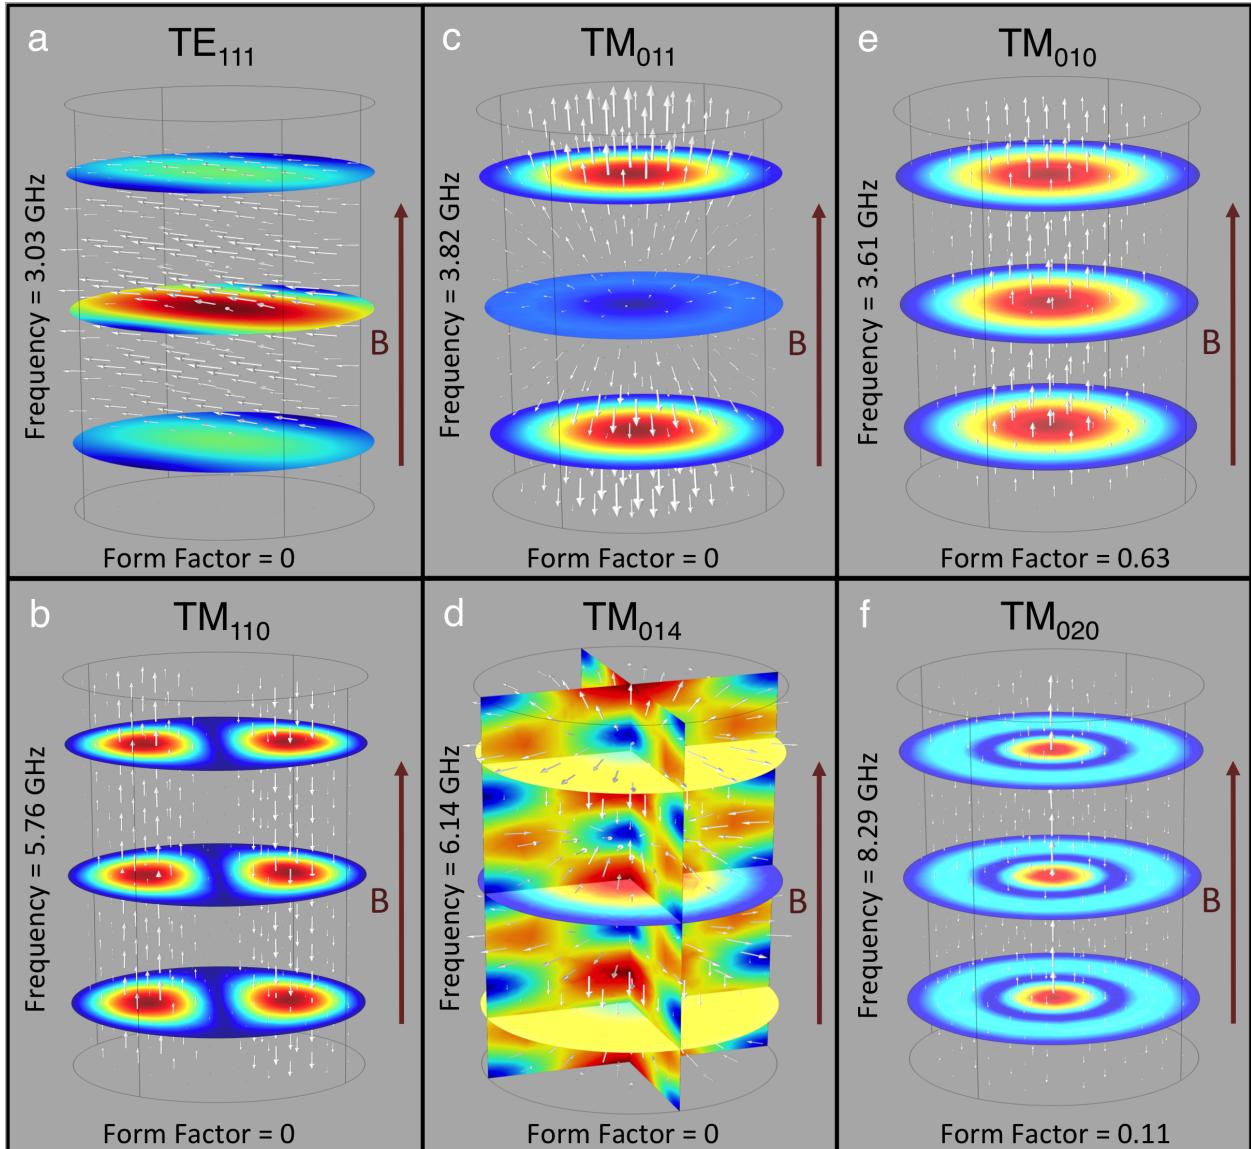


Figure 4.1: Electric field for different TM and TE modes in a bare cavity. Red indicates regions of higher field for cross-sections of the cavity.

important because currents need to flow across the junction to create the modes that couple to axions. To avoid Ohmic losses that would degrade the Q, the cavity is given a knife-edge and the end caps are bolted tightly down. The Q can be written in terms of the geometry

and skin depth ( $\delta$ ) of the cavity

$$Q = \frac{\lambda}{\delta} \frac{(x_{mn}^2 + (p\pi r/d)^2)^{1/2}}{2\pi(1 + 2R/d)} \quad (4.6)$$

where  $\lambda$  is the wavelength of the cavity mode in free space. In the classical regime, at moderate temperatures and frequencies, the skin depth is given by

$$\delta_{classical} = \left( \frac{2}{\omega \mu \sigma} \right)^{1/2} \quad (4.7)$$

where  $\mu$  is the relative permeability and  $\sigma$  is the bulk conductivity. However, at higher frequencies and lower temperatures,  $\delta$  shrinks and the conductivity increases. When the mean free path of electrons in the metal exceeds the skin depth, the surface finish of the resonator plays a bigger role and equation 4.7 is no longer valid. In this new regime, the “anomalous skin depth” is given by

$$\delta_{anomalous} = \left( \frac{3^{1/2} c^2 m_e v_F}{8\pi^2 \omega n e^2} \right)^{1/3} \quad (4.8)$$

where  $m_e$  is the electron mass,  $v_F$  is the Fermi velocity,  $n$  is the electron number density and  $e^2$  is the electron charge.

#### 4.1.2 Swept Measurements

The frequency and  $Q$  of a cavity mode is determined by injecting a swept signal into the apparatus with a weakly coupled antenna and measuring the transmitted power through a second antenna that is typically critically coupled (antenna coupling is addressed later in this section and in sections 3.3.3 and 4.4.6). A device that performs such a measurement is called a network analyzer. The top two panels of Figure 4.2 show a test cavity sandwiched

between two supporting end caps. Inserted into two holes in the top end cap are weakly and critically coupled antennas that also connect to the input and output ports of a network analyzer.

On resonance, the swept power from the network analyzer is absorbed through the “weak port” into the cavity and is picked up by the critically coupled antenna at the “major port” and received by the network analyzer. Off resonance, most of the power from the network analyzer is reflected back from the weak port and is not transmitted through the cavity. The shape of the transfer function is a Lorentzian

$$h(f) = \frac{1}{1 + 4(f - f_0)^2/\Gamma^2} \quad (4.9)$$

where  $f$  is the swept frequency of the network analyzer,  $f_0$  is the center frequency of the resonant mode and  $\Gamma$  is the full width half maximum (FWHM) of the response. The middle panel of Figure 4.2 shows the transmitted power vs frequency response of the network analyzer over-plotted with a lorentzian fit. This fit provides the center frequency and  $Q$  of the cavity. Of coarse this is just a single mode in the cavity and equation 4.6 claims that there are many such modes. The bottom panel of Figure 4.2 shows a wide sweep of the network analyzer as a purple line. Every mode picked up by the antenna appears as a peak. By entering the dimensions of the cavity into equation [54], it is easy to calculate what these frequencies aught to be. These calculated frequencies, independent from the measurements, are shown as blue vertical bars. Labels appear above the peaks and form factors obtained from a simulation are printed sideways and just to the left of expected frequencies. Shown in red are the  $TM_{010}$  and the  $TM_{020}$  modes with very large form factors by comparison with the other TM mode form factors. These other non  $TM_{0n0}$  mode form factors are technically equal to exactly zero but likely due to uneven meshing in the simulation, they are calculated to be of order  $10^{-4}$ .

Present in the mode structure but hidden in this measurement are also a host of TE

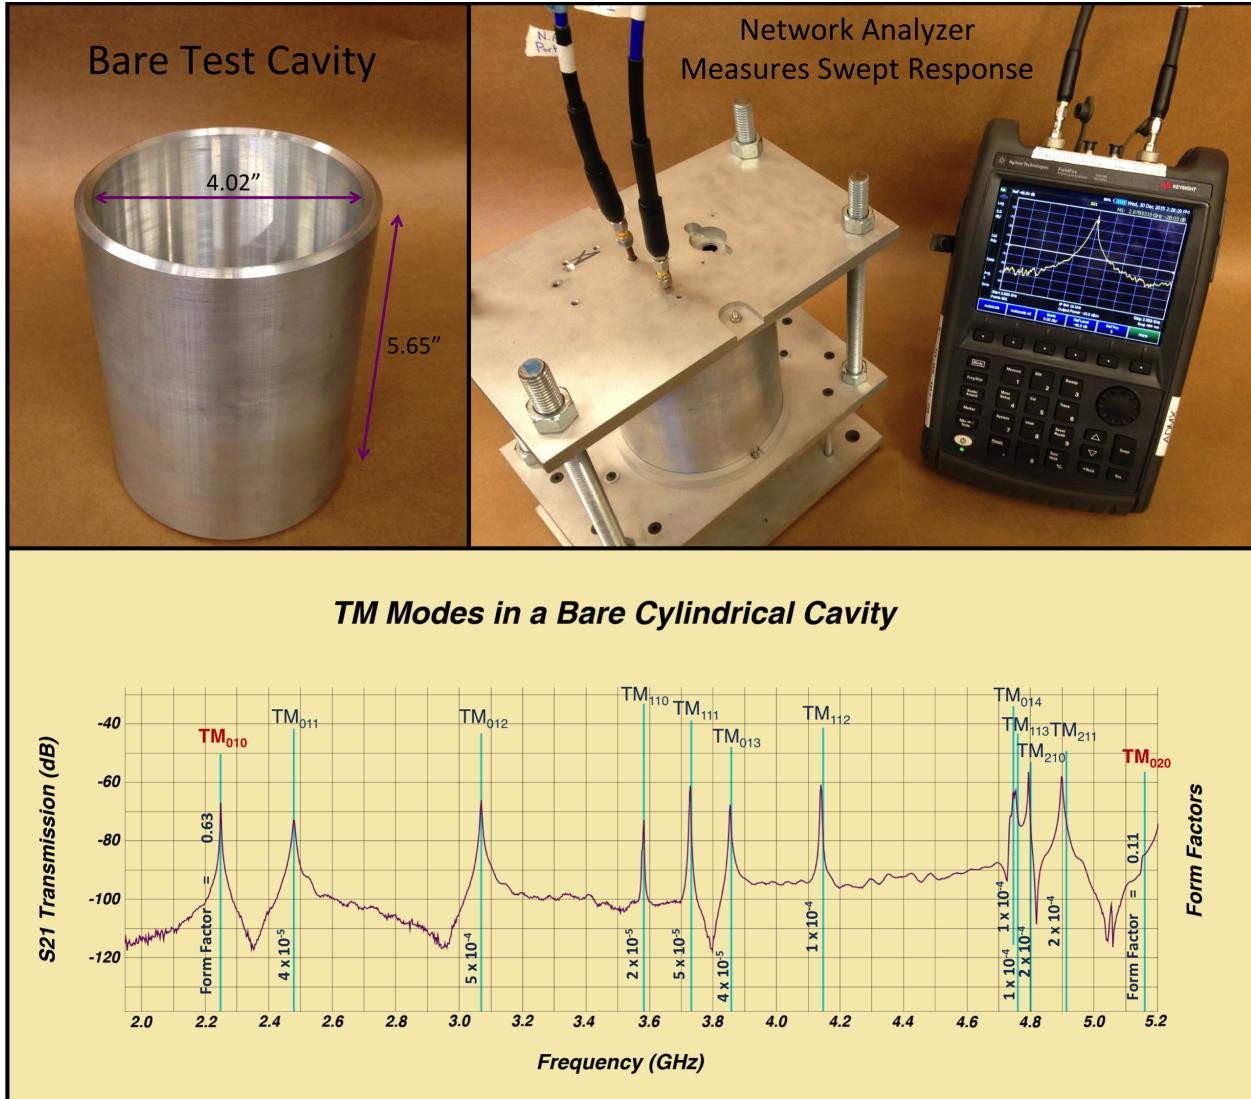


Figure 4.2: Resonant peaks in a bare cavity. Modes are measured by injecting swept power from a network analyzer into a weakly coupled antenna and detecting the power transmitted through the cavity from a critically coupled antenna. The peaks in the measurement are in good agreement with the vertical blue lines representing the analytical predictions. Only TM modes are visible because the electric field of TE modes is perpendicular to the orientation of the antenna.

modes. However, because these TE modes have no vertical electric field component, they don't couple well to the vertically positioned antenna. To illustrate that they are indeed there and add complexity to the mode structure, the straight coax antenna was bent to pick up radial and azimuthal electric field components and the measurement was performed again. Figure 4.3 shows another network analyzer transmission measurement plotted in purple. The TM labels and form factors from Figure 4.2 are plotted above the purple line and TE information is plotted below.

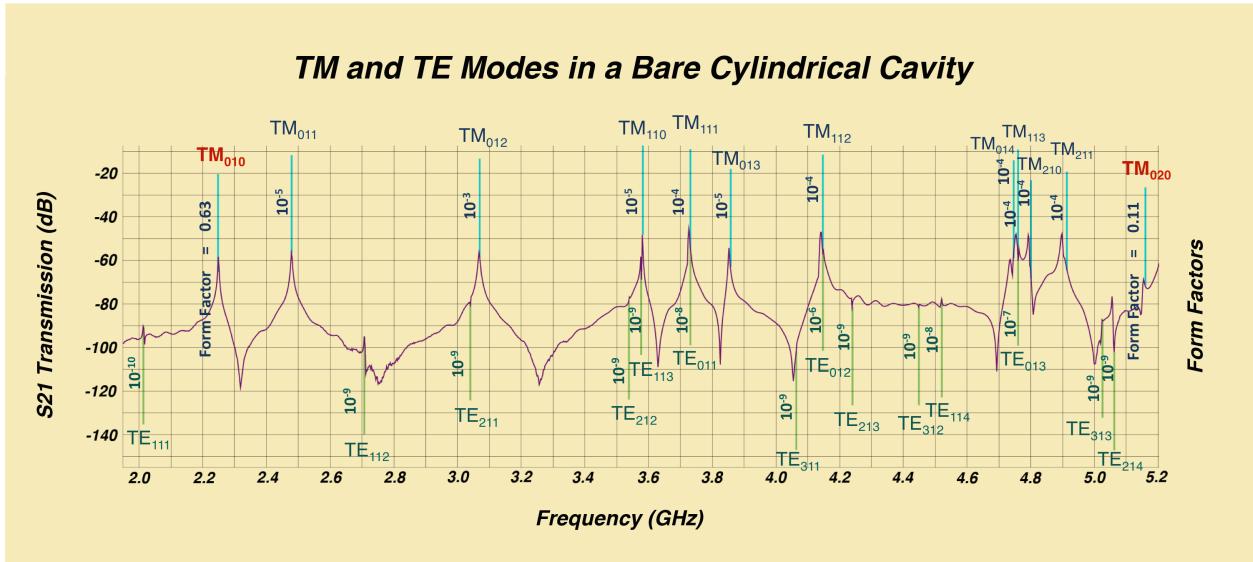


Figure 4.3: Resonant peaks in a bare cavity from a network analyzer transmission measurement. In addition to TM mode, TE modes were also detected by bending the antenna. TM (blue) and TE (green) analytical predictions are shown as vertical lines.

As described above, making a measurement of the transfer function requires that two antennas be inserted into two ports on the cavity. The port through which power is injected is called the “weak port”. Inserted into the weak port is a fixed antenna that is made to be very weakly coupled to the cavity so as not to disturb the geometry or remove power. The port through which power is received clutches an antenna with a variable depth and is called

the “major port”. This antenna can be made to be weakly coupled by almost fully removing it and the unloaded  $Q$  of the cavity can be measured. As the antenna is inserted into the cavity, the coupling increases and the  $Q$  decreases. When the loaded  $Q$  reaches half the unloaded  $Q$ , the antenna is said to be critically coupled. At this coupling, half of the power deposited from axions is being absorbed by the walls of the cavity and half of it escapes through the antenna. It is also at this coupling that the impedance looking towards the cavity is equal to the probe characteristic impedance. During data taking, the antenna(s) must be periodically adjusted to maintain a critical coupling for optimal signal propagation.

Antenna coupling can be measured by looking at a signal reflected off of the antenna. This is done using a network analyzer and a circulator (Figure 3.5). Power is swept from the network analyzer into a port on the circulator that directs the signal towards the probe. Off resonance, the impedance is predominantly imaginary, the probe is not impedance matched to the cavity and the signal is reflected back to the network analyzer (shown as a flat baseline in the figure). When critically coupled, a signal that is on resonance sees a real  $50\Omega$  looking in, is completely absorbed by the cavity and is represented by the large dip. Critical coupling is maintained by adjusting the antenna depth to minimize this absorption trough. When the trough depth is 30dB below the baseline it is considered critically coupled.

#### 4.1.3 RF Simulations

A swept response cannot measure the form factor of a mode. Also, for any case other than the minimally perturbed bare cavity, the form factor cannot be computed analytically. Without this knowledge, it is impossible to know the SNR of the experiment and in some cases even the correct mode. Past versions of ADMX used 2D relaxation techniques until it was discovered that the small deviations in axial symmetry had a dramatic effect on the mode structure. It was determined that full 3D electromagnetic simulations were necessary [46]. Microwave Studio, HFSS and COMSOL are all examples of software packages that

perform finite - element analysis and are capable of solving 3D problems. The 3D electric field plots shown in Figure 4.1 were generated by J Sloan using COMSOL.

In COMSOL, the user first builds a virtual model of the experimental apparatus by inputing the geometry, boundary conditions and materials used. The user also picks a particular mesh and a starting guess as to the frequency of interest. COMSOL will then solve for eigenfrequencies, electric field components and any other desired output parameters. In the case of a tunable cavity, tens to thousands of these solutions will be performed to characterize a changing geometry. Resolution, number of output solutions other parameters are optimized to avoid the case where a simulation takes weeks to finish.

While such simulations are time consuming, very sensitive to slight deviations in geometry and difficult to integrate into an automated DAQ, they remain essential for computing form factor. Eventually, the ability to apply intuition from the bare cavity scenario breaks with the addition of tuning rods at higher frequencies, and a simulation is the only practical way to differentiate between modes that couple to axions and modes that do not. A simulation, however, is only useful to the degree to which it agrees with real frequency measurements of the cavity.

Here, starting with a simplified, bare cavity, I have presented three different approaches to understanding the mode structure within a cavity. All techniques provide frequency information necessary for direct comparison and have their own advantages: analytical solutions provide mode labels, simulations provide form factors and transmission measurements provide the only actual real world information about the cavity itself without making assumptions. Figure 4.3 contains information from all three of these techniques. With the addition of tuning rods, analytical solutions which assume radial, azimuthal and translational symmetry become no longer valid and the mode structure becomes much more complicated.

## 4.2 Mode Following

### 4.2.1 Tuning

If one knew the mass of the axion, the discussion about understanding mode structure would be very short. One would simply fabricate a cavity with a radius

$$r \equiv \frac{2.405h}{2\pi\sqrt{\mu\epsilon}m_a c^2(1 + \beta^2/2)} \quad (4.10)$$

and a length such that no TE modes appear at the same frequency as the  $TM_{010}$ . Here  $h$  is planck's constant,  $m_a$  is the mass of the axion,  $c$  is the speed of light  $\beta$  is the Lorentz boost from the expected velocity of dark matter near the earth. However, not knowing the mass of the axion, requires the use of a tunable cavity (Subsection 4.2.1) which complicates the mode structure and has no analytical solutions. One can imagine many schemes for a tunable geometry. The primary ADMX cavity is equipped with two vertically oriented copper rods the can rotate from the wall of the cavity to the center and back again. Figure 4.4 shows how the addition of tuning rods alter the boundary conditions of the cavity and perturb the electric field for the  $TM_{010}$ . With the addition of tuning rod(s), equation 2.24 is no longer valid, and TEM solutions are also possible. At this point, mode labels have technically lost their meaning but needing a vocabulary to distinguish modes, I will continue to use the language of the bare cavity. From now on, I define the  $TM_{010}$  to be the lowest frequency tunable mode with the highest form factor, and a  $TM_{020}$  mode as a mode with a non-zero form factor that is roughly twice the frequency of the  $TM_{010}$  (more about mode splitting in Subsection 4.2.1).

The Sidecar experiment takes the same approach but with only a single rod. The movement of a tuning rod shifts the frequencies of the Lorentzian peaks shown in a swept response measurement. Figure 4.5 shows a collection of swept responses of the Sidecar cavity for different tuning rod positions. Here, the  $TM_{010}$  is shown for 70 different rod angles between  $9^\circ$

## Magnitude of the Electric Field for a Cross-Section of the Cavity

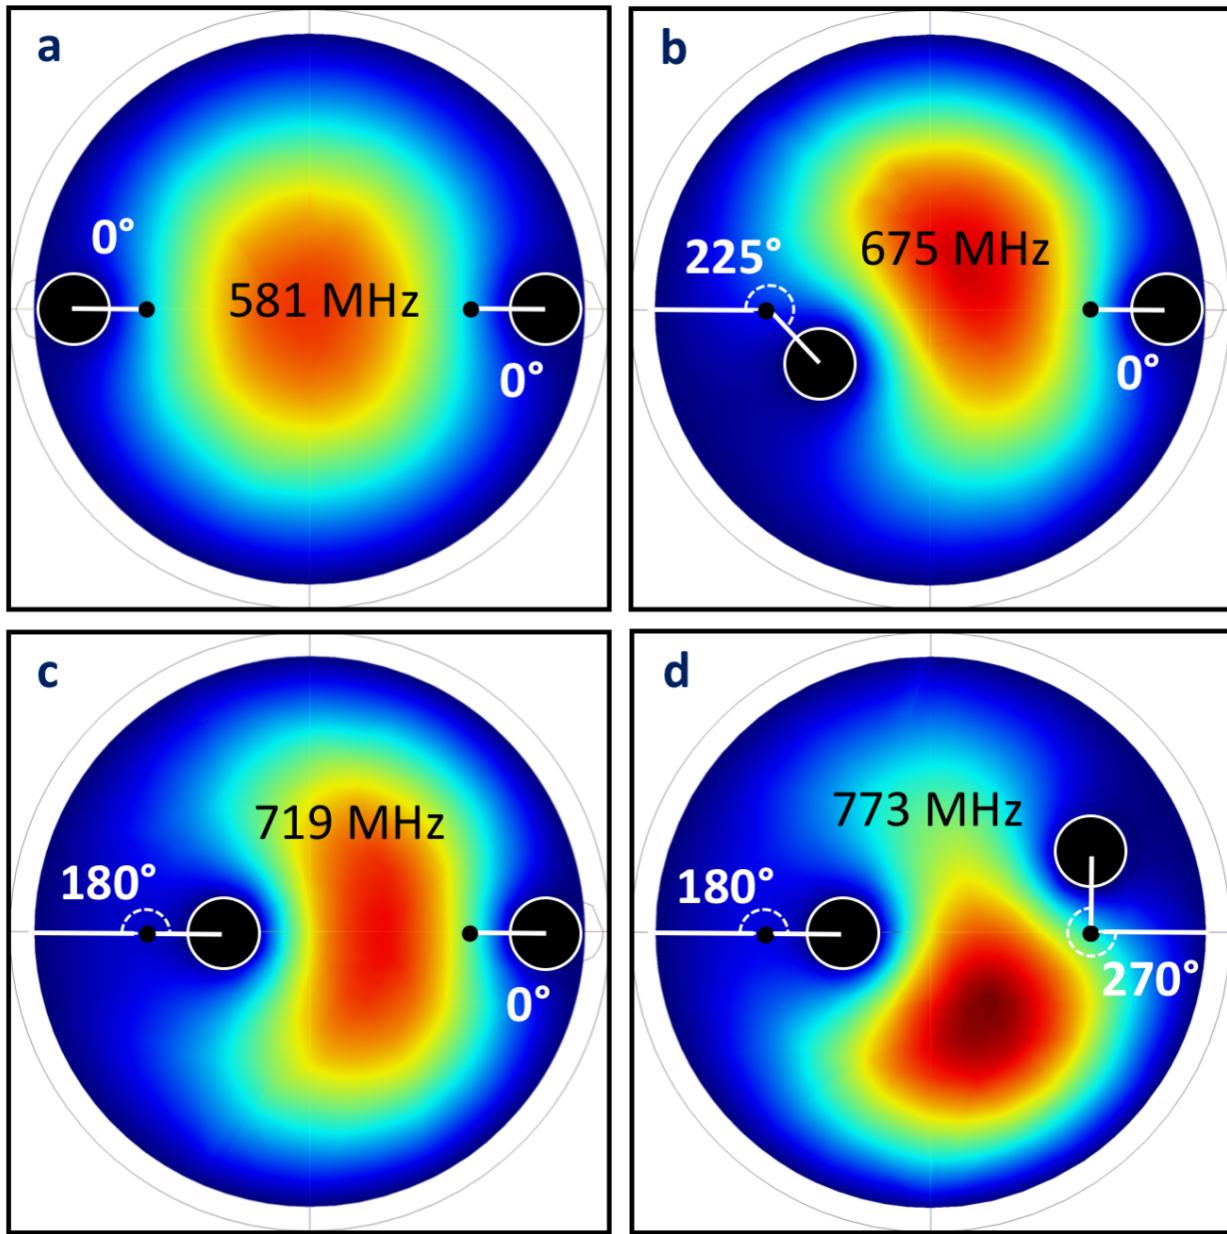


Figure 4.4: Tuning rods alter the boundary conditions of the cavity and perturb the electric field for the  $TM_{010}$ . Four cross-sections of the main ADMX cavity are show for different tuning rod configurations. The magnitude of the  $TM_{010}$  electric field has been mapped to a color: red indicating a strong field and blue in regions where the field is weak. Tuning rod angles are shown in white and the frequency of the resulting  $TM_{010}$  appears in black.

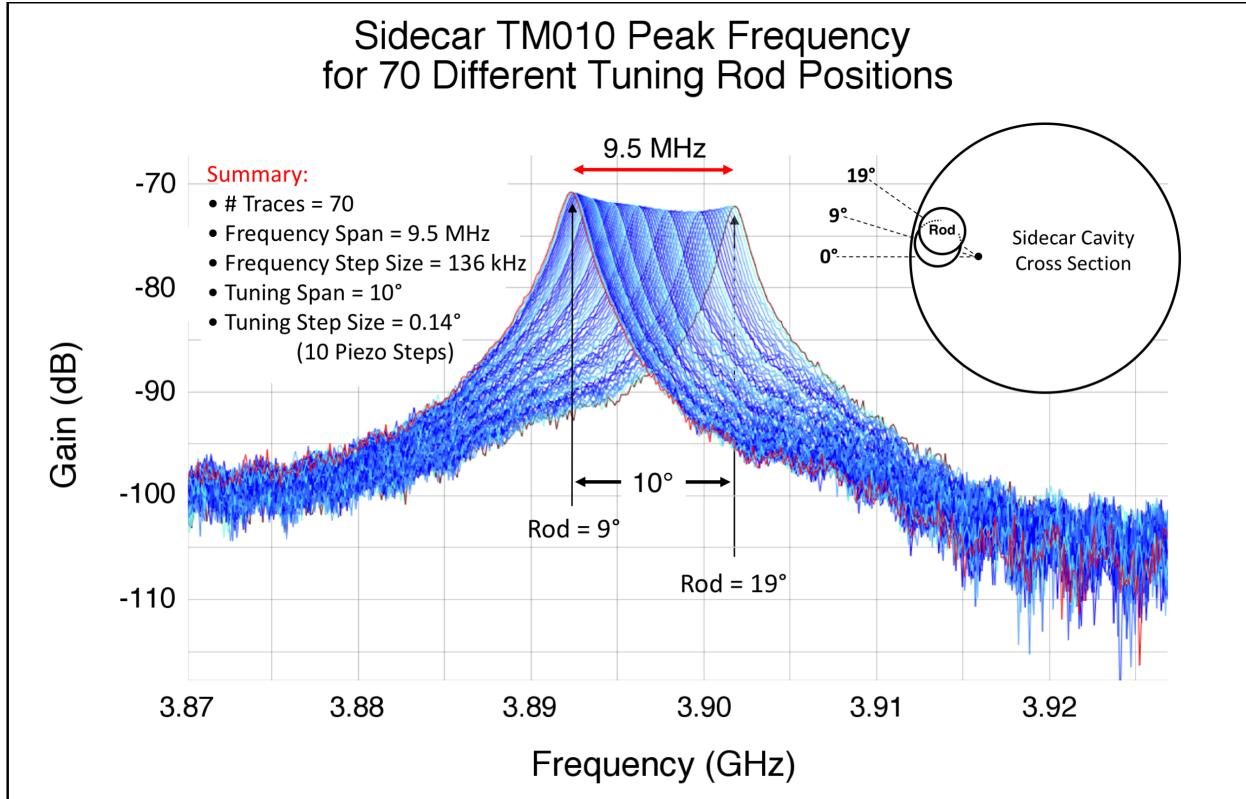


Figure 4.5: Lorentzian peak of the  $TM_{010}$  mode of the Sidecar cavity for 70 different tuning rod positions. The  $TM_{010}$  is smoothly tuned nearly 10MHz by moving the tuning rod 10 degrees.

and  $19^\circ$ . Here, the mode frequency is seen to tune nearly 1 MHz per tuning angle degree in 136 kHz/0.14° steps.

#### 4.2.2 Tracking

During data taking, the experiment needs to be able to track the mode frequency as it tunes. If the tuning rod step size between network analyzer scans is small, this can be done by fitting each  $S_{21}$  response with a Lorentzian and using the resulting measured center frequency and Q to set the center and span frequencies for the next network analyzer scan. This simple

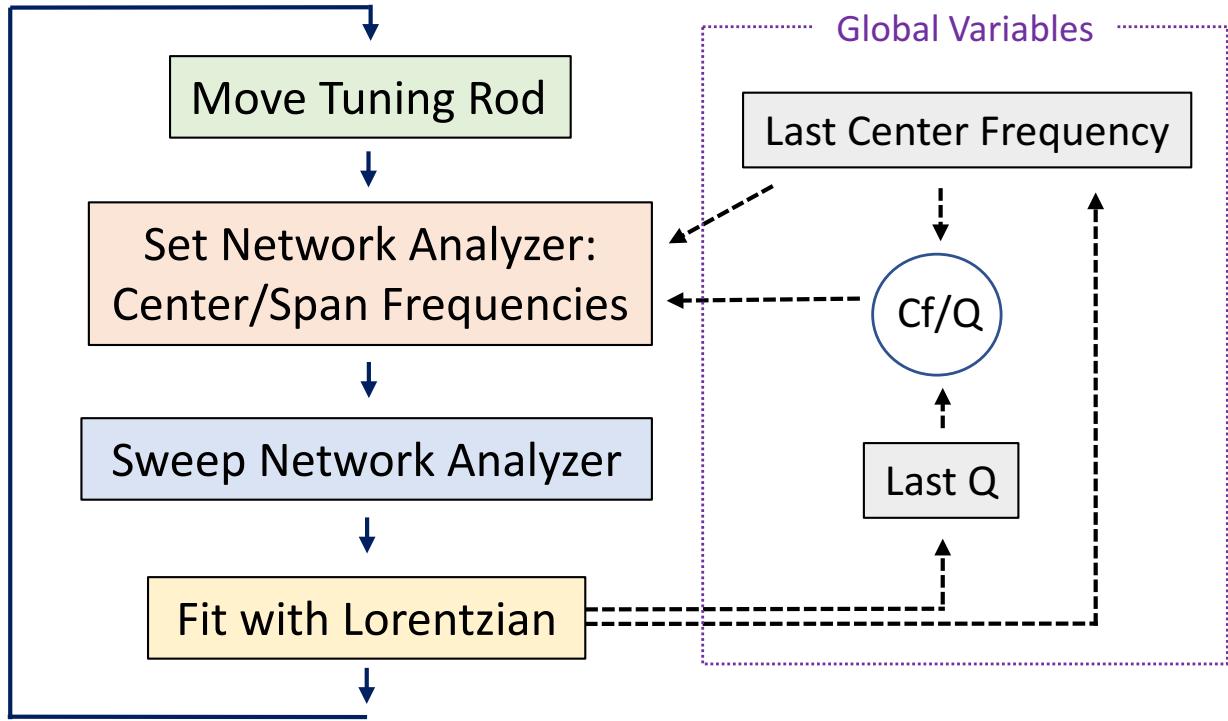


Figure 4.6: Simple algorithm for tracking a resonant mode. Move tuning rod, set center/span frequencies using their last measured values, sweep cavity with the network analyzer, fit peak with a Lorentzian, save the center frequency/Q and repeat.

algorithm is shown in figure 4.6

By dynamically changing the span width and center frequency of the network analyzer, the mode of interest is measured with adequate resolution and in most cases, the only mode in the network analyzer window. However, the mode structure is complicated. At certain frequencies, modes will mix and this simple algorithm may start to follow the wrong mode. Tuning too fast, vibrations of the insert and the failure to identify a Lorentzian peak in the fitting process may also cause the system to find a new mode that does not couple to the axion field. Finding the correct mode and following that mode without getting lost requires a map of the cavity frequencies for different tuning configurations. This is known

as a mode-map.

#### 4.2.3 *Mapping*

A mode-map is a chart of cavity mode frequencies versus tuning rod positions. A mode-map is made by making a series of wide  $S_{21}$  swept response measurements for slightly different tuning positions such that each measurement contains several resonant modes. A peak identifying function then iteratively decomposes the swept response into its constituent Lorentzians and saves the values for each rod configuration. For a single rod cavity, a mode-map is a 2D plot of cavity frequencies versus rod angle. For a double rod cavity, a full mode-map is intrinsically a 3D entity projected onto 2 dimensions. For that reason, it is always necessary to specify the location of the rod not shown. Figure 4.7 gives an example of how a mode-map is made. A section of a mode-map is shown at the center of the figure. Each vertical slice of the map is made from swept response measurements - two examples of which are shown, rotated sideways to match the axis of the mode-map.

A mode map contains the information needed to associate a tuning rod configuration with a frequency, the slope of which can be used to determine how much a rod needs to be rotated to tune a mode by a given amount. A mode map reveals which modes tune well (TM modes), which modes barely tune at all (TE/TEM modes) and where their paths cross. Mode mixing occurs when two modes are tuned to nearly the same frequency and start to interact with each other. For some modes, this interaction is very subtle and modes pass right through each other. For others, the mutual interaction is quite large, creating voids where no data can be taken. Keeping track of which mode is the correct mode is difficult without form factor information. A simulation with numerically determined form factor values must be performed and compared directly with a measured mode map. Figure 4.8 shows a measured mode crossing potted above a simulation of the same crossing. In the simulation, form factor is indicated by color; gray indicates no sensitivity to axions. The measured map shows two

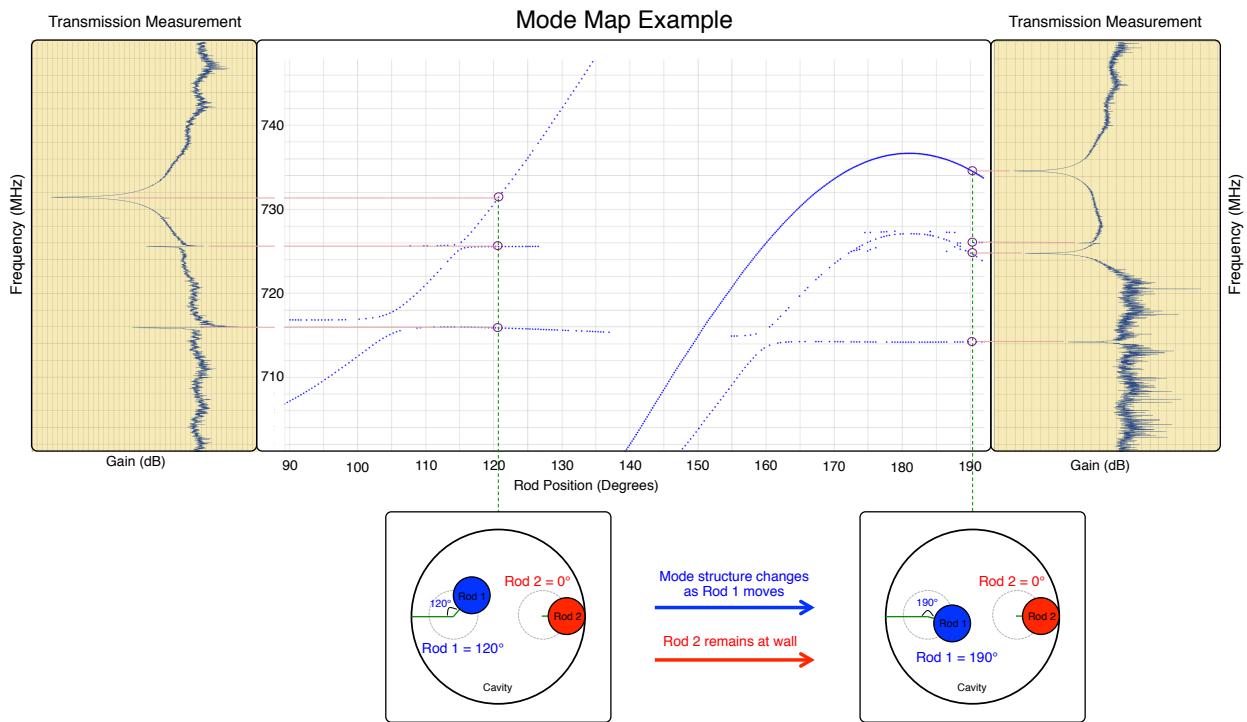


Figure 4.7: An explanation of how a mode-mate. The frequencies of peaks from transmission measurements are recorded plotted versus tuning rod positions. A cavity with two rods requires a 3D mode-map to characterize all possible tuning configurations and therefore requires specifying the location of the second rod in a 2D map. A mode-map reveals how a given mode tunes and what tuning rod configuration is required to hit a target frequency.

modes that appear to repel each other when in fact, the simulation shows that the modes have swapped places.

The figure illustrates why it is necessary for measurements and simulations to go hand in hand. However, achieving good agreement between the two can be difficult. Tiny milling imperfections and details about the geometry that are not properly accounted for can have a very large impact on certain modes. This is especially true for higher frequency modes. For example, [46] found that the tiny gaps between the tuning rod and the end cap of the cavity greatly affect the mode structure. A more recent study performed a set of simulations measuring the main cavity  $TM_{010}$  mode structure for different rod endgaps [55]. Figure 4.9 summarizes these results and show that increasing the gap size increases the frequency and impact of interfering modes (highlighted in green and purple). The study also shows that by shorting the capacitive coupling between the rod and endcap with a thin filament greatly increases the possible frequency coverage of the  $TM_{010}$  and nearly resembles a cavity without endgaps. Such RF simulations are necessary not only for identifying the correct mode but also for future cavity designs.

As a final example of how simulations must always go together with measurements, I consider this simple perturbation of a bare cavity: the creation of a tuning rod at the wall. Until recently, it was assumed that a rod at the wall of the cavity would minimally affect the mode structure and in this configuration, the  $TM_{010}$  could be inferred from the bare cavity prediction. To test this intuition, a simulation of the bare Sidecar cavity was gradually perturbed with the growth of a tuning rod. This simulation is shown in figure 4.10. The left side of the plot is the bare cavity with analytical predictions for the TM and TE modes. The green points in the middle of the plot show how the mode structure of the cavity changes as the tuning rod diameter is increased. At the right of the plot, the rod has reached full size at a diameter of 0.5 inches. If one were attempting to locate and identify a particular mode without the ability to perform such a simulation, the naive assumption would be, for

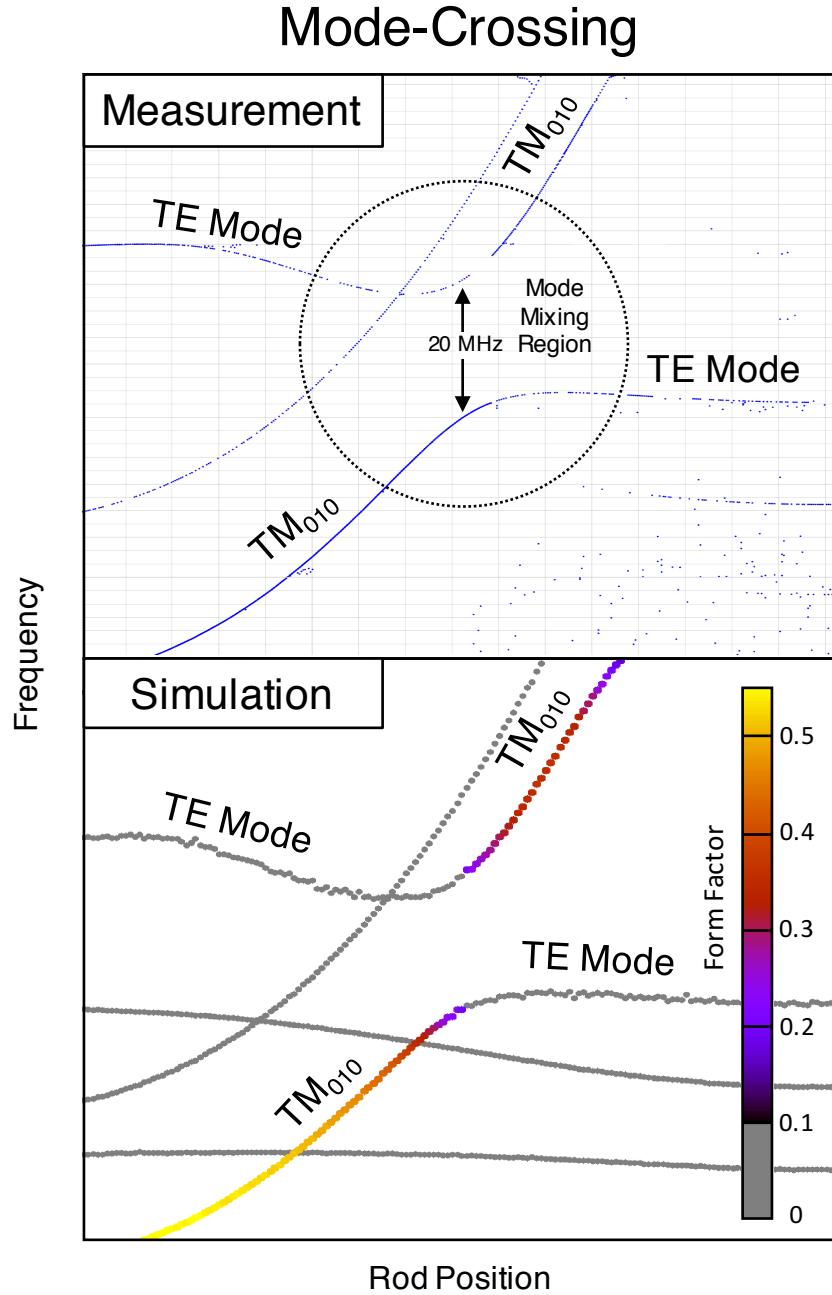


Figure 4.8: An example of a mode-crossing. When two modes are tuned to similar frequencies, the modes mix and often leave a gap where no data can be taken. From the perspective of a set of network analyzer measurements, two modes appear to move towards each other, slow down and then move apart. A measured mode-map representation of this crossing is shown at the top of the figure. A simulation, shown at the bottom of the figure is required to determine which mode couples to the axion field. Color indicates the value of the form factor and shows that the  $TM_{010}$  mode mixes with a  $TE$  mode and swaps frequencies with it.

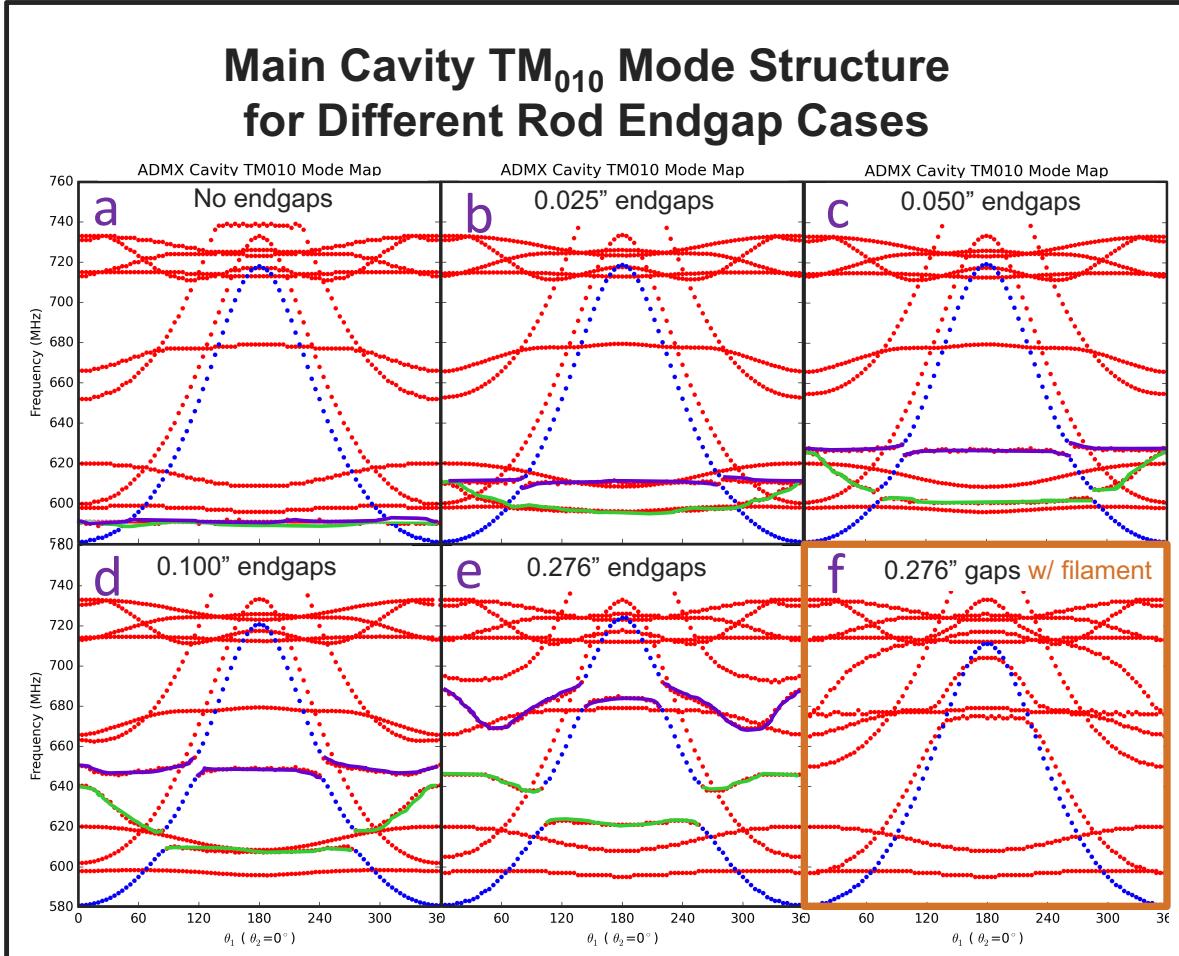


Figure 4.9: Simulation of the main cavity mode structure for different gap sizes between the end of the tuning rods and the end cap of the cavity. As the gap is increased, two intruding modes (shown in purple and green) increase in frequency and create large voids where no data can be taken. By adding a filament which shorts the end-gap capacitive coupling (simulation f), the mode structure becomes comparable to the minimally perturbed no end-gap scenario (simulation a). Simulations were performed by Mark Jones at PNNL.

example that the  $TM_{012}$  mode should be at about 4.4 GHz and several hundred MHz below the  $TE_{113}$  and  $TE_{211}$ . Instead, it turns out that the  $TM_{012}$  has tuned up over 200 MHz and that the breaking of the degeneracy in the cavity has split both of the TE modes on

either side of the  $TM_{012}$ . The fact that there are so many mode crossings in this simple perturbation suggests that bare cavity intuition is not even valid in the case where the rod is at the wall of the cavity.

The conclusion is that the mode structure of a rod perturbed cavity is complicated and that measurements and simulations must be used together to characterize the system. In some cases it may be possible to tune around a crossing or perturb the geometry somehow (with for example a second antenna) to move the frequency at which the crossing occurs. It is the author's opinion that time spent attempting to navigate around a crossing will result in diminishing returns. Instead, clean data should be taken where ever possible while a second cavity is machined to have crossing that don't coincide with the blind spots of the first cavity.

#### 4.2.4 Sidecar Structure

The Sidecar cavity (discussed in section 3.3.2) is a 4.75" long, 2.5" inner diameter OFHC plated cylinder designed to have a frequency range between 4-6 GHz. At room temperature, the cavity typically has a  $Q_{unloaded}$  of 11,000 ( $Q_{loaded} \sim 5,500$ ). Figure 4.11 gives a detailed description of the Sidecar, single rod mode structure using bare cavity analytical predictions, swept response measurements and simulations of tuning rod growth and movement. On the far left, the TM and TE analytical predictions for the bare cavity are shown. To the right of these predictions, is the simulated mode structure of the cavity vs rod diameter (green). This COMSOL simulation of a growing rod, performed by James Sloan at the University of Washington provides the missing information that maps the mode structure of the bare cavity to the tunable cavity on the right of the figure. Plotted over the measurement (blue) is a HFSS simulation (red) that was performed by Mark Jones at PNNL. The agreement is not perfect and not all modes can be identified and traced back to their bare cavity origins. There are also TE modes that appear in the simulation but not in the measurements because

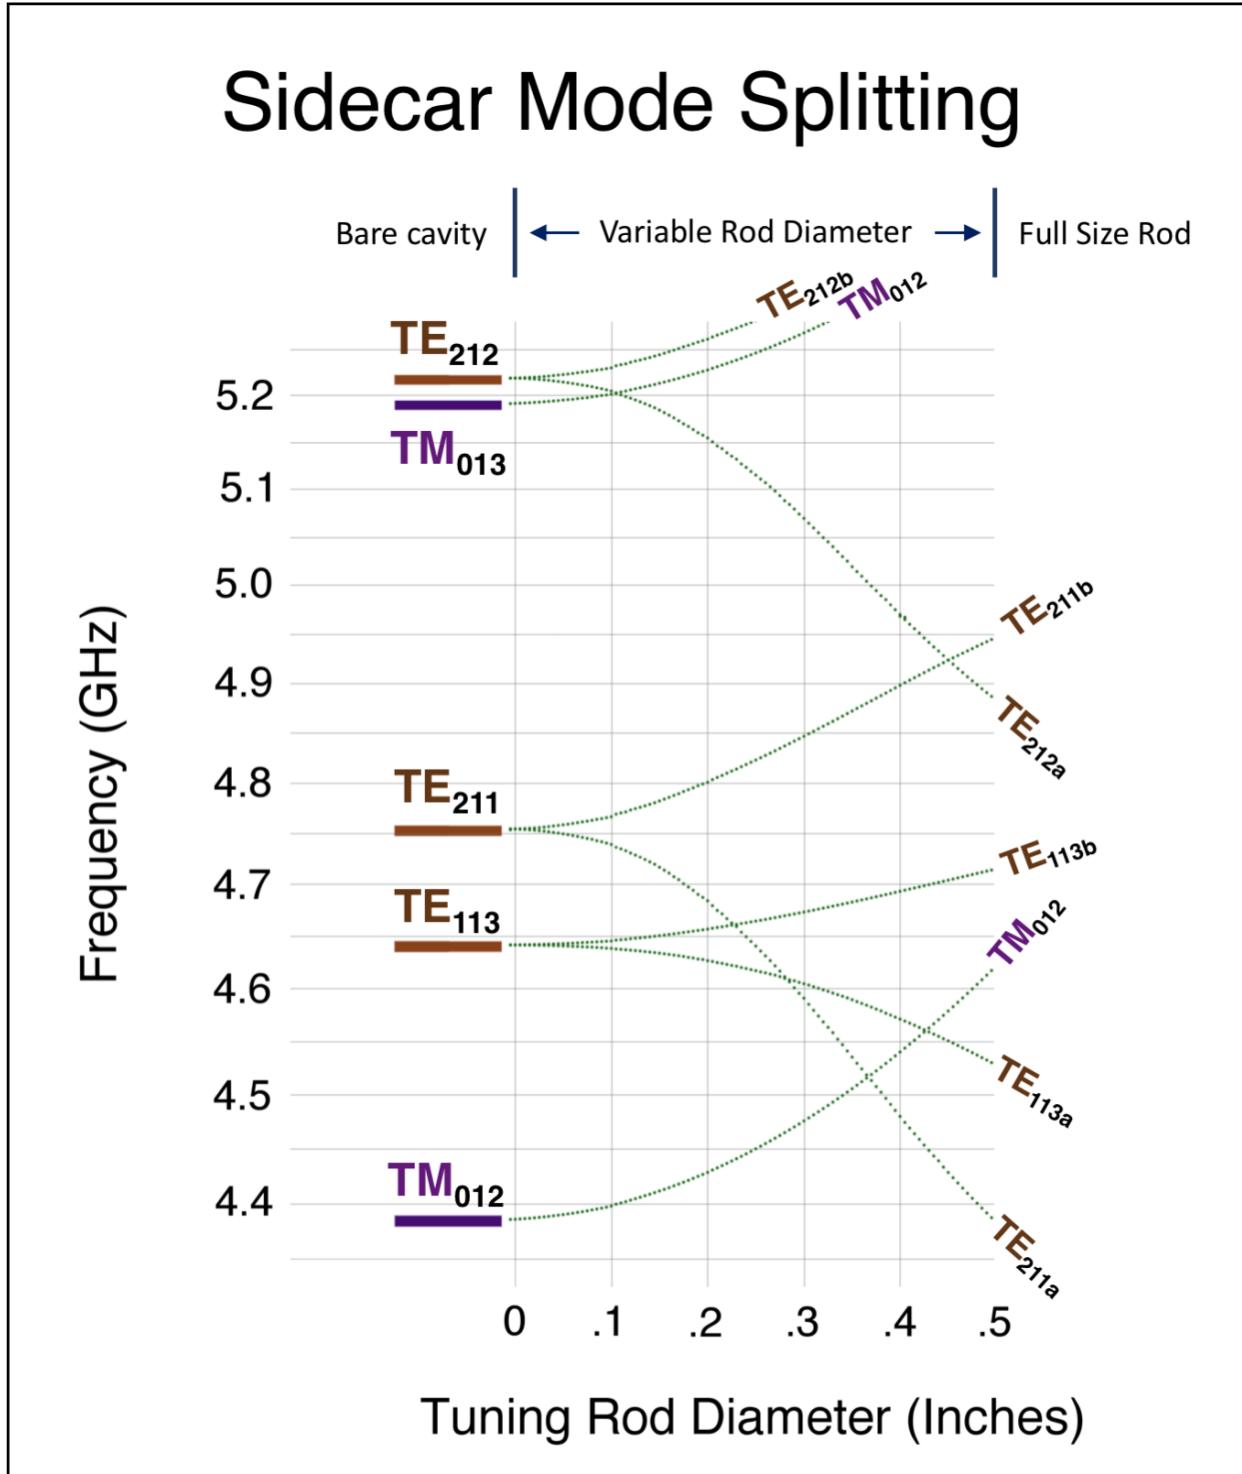


Figure 4.10: Analytical frequency predictions of the bare Sidecar cavity are compared to simulations the rod perturbed scenario. The simulations (green) show changes in the mode structure as the Sidecar tuning rod is grown from a diameter of zero to 1/2". The creation of a tuning rod breaks the symmetry in the cavity and causes TE modes to split.

of the vertical orientation of the antenna. With the addition of a tuning rod, TEM modes are also possible and depend only on the length of the cavity. The frequency spacing between these modes occurs every  $\Delta f = c/2L$  where  $c$  is the speed of light and  $L$  is the length of the cavity. For Sidecar, this spacing happens every 1.2424 GHz and are visible as noisy green dots in the rod growth COMSOL simulation. These modes were identified because they had the correct spacing and because they don't connect to any analytical TE or TM predictions.

Figure 4.12 is the measured mode map by itself with form factor values plotted over the  $TM_{010}$  orange highlighted line. This plot represents three weeks of piezo rotations and network analyzer measurements. Red points represent a mode map that was made with the antenna weakly coupled and blue points represent measurements made with the antenna critically coupled. The blue map (plotted over and mostly hiding the weakly coupled measurements) is the first full critically coupled mode map produced by ADMX. This was made possible by a the Sidecar auto-couple algorithm (section 4.4.6). A careful inspection of the blue and red maps shows that the measurements agree everywhere except where they diverge near 4.6 GHz between 130 - 230 degrees. The form factor for the  $TM_{010}$  starts at 0.585 when the rod is at the wall, reaches its lowest value of 0.464 when the rod is near 145 degrees and rises to it's highest value of 0.65 at the peak. The 2016 data run, discussed in chapter 5, took data in the upper section of the mode-map between 5.090 and 5.799 GHz.

### 4.3 Making a Run Plan

This section discusses the planning that must go into a data run before the insert is lowered into the magnet.

#### 4.3.1 Constraints and Goals

Formulating a “run plan” is a multi-disciplinary exercise that shouldn't be underestimated. Before the finalization of a run plan, the following questions should have been asked and

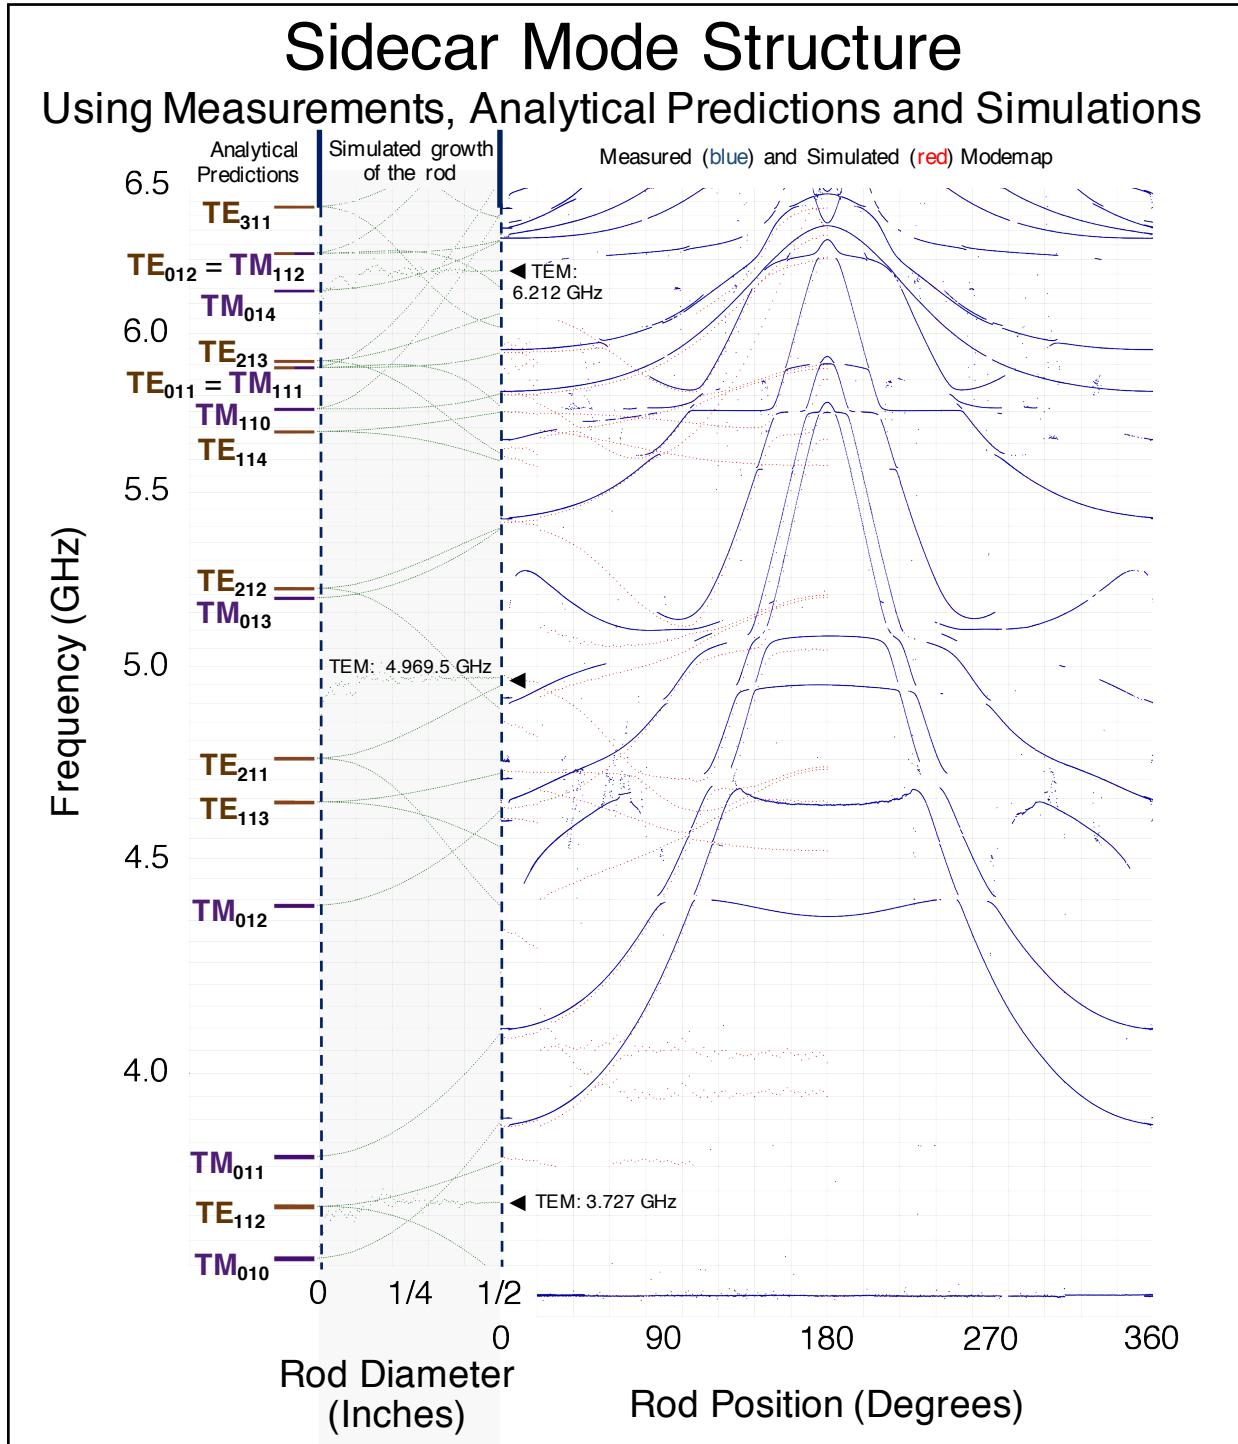


Figure 4.11: Analytical predictions, simulations and measurements of the Sidecar cavity. Simulations of the growth of a tuning rod (green) link the bare cavity predictions on the left of the figure to the measured (blue) and simulated (red) sidecar mode map on the right.

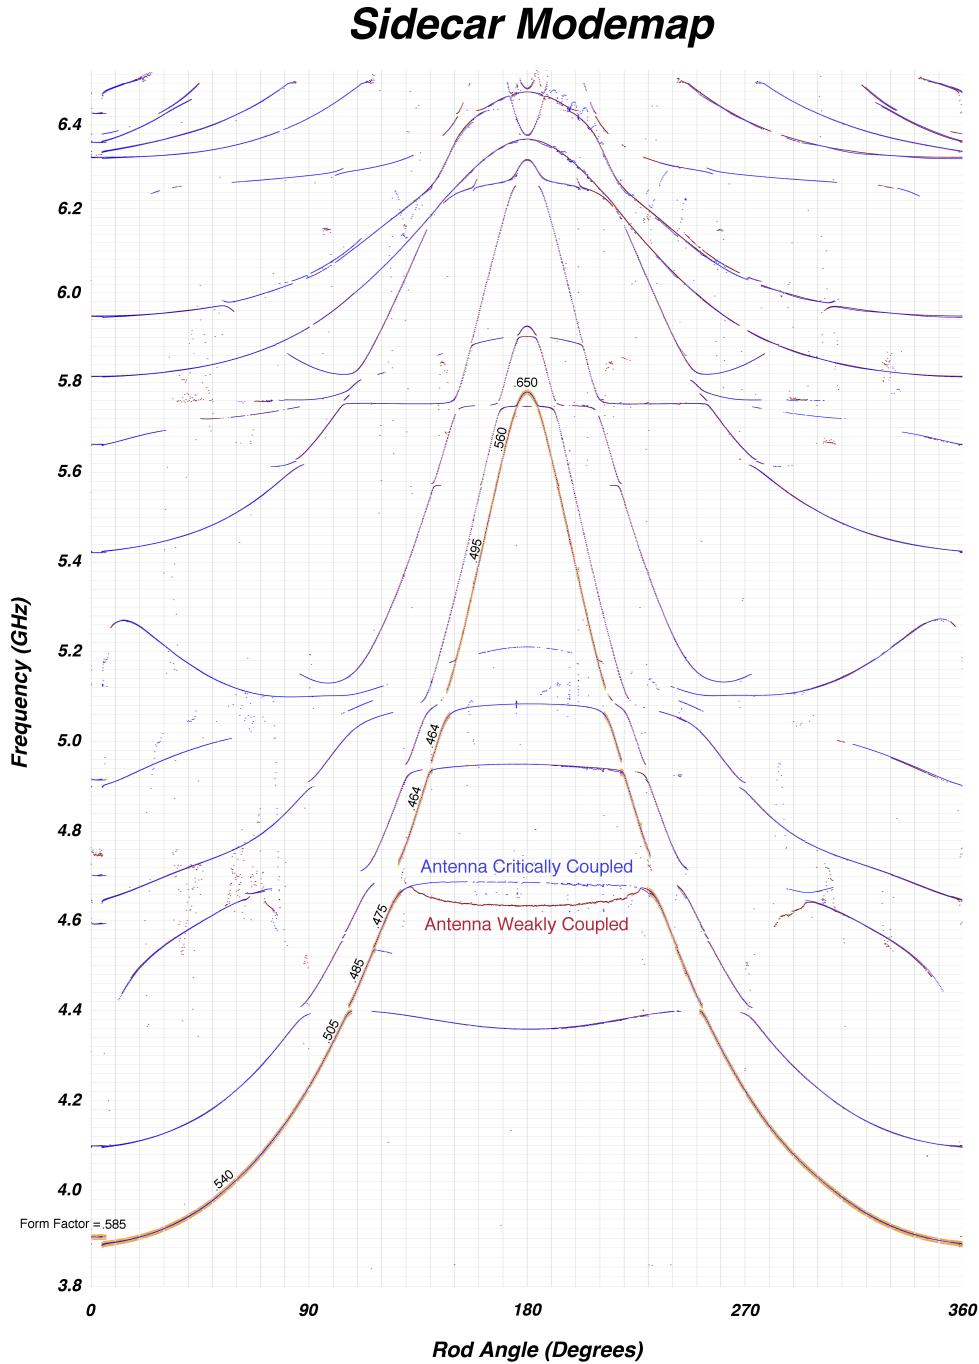


Figure 4.12: Weakly and critically coupled measured Sidecar mode-maps plotted on top of each other. Red weakly coupled measurements are hidden by and coincide with the blue critically coupled measurements except for a region at 4.6 GHz at the center of the figure where they break. The  $TM_{010}$  mode is highlighted in orange and is accompanied by simulated form factor values.

answered

- What are the science goals of the experiment?

Example: Is it more desirable to set great limits over a narrow mass range or set good limits over wide range?

- What are the milestones of the funding agency?

Example: dilution refrigerator must achieve below 150mK base temperature.

- What are the hardware constraints of the experiment?

Example: The MSA has a frequency range 644MHz - 715MHz

- What tuning trajectory satisfies the objectives of the experiment and results in several months of good, continuous data?

- Are there contingency plans to account for broken hardware?

Example: if the gearbox of the active tuning rod jams, does the mode structure allow for good data by tuning the other rod?

- Should blind, synthetic axion signals be injected

- How often should candidate peaks be rescanned?

- How many weeks/months have been allocated for the run?

When the run plan is finalized, the following should be specified: timeline, target sensitivity, start/stop rod positions, start/stop frequencies, intended scan rate, protocol for synthetic axion injection and protocol for axion candidate rescans. With a single rod and a single data channel, it is relatively simple to make such a plan for the Sidecar experiment. However, the main dual-channel, multi tuning rod experiment requires careful planning.

### 4.3.2 *Simultaneous Multi-Mode Tracking*

In 2014, a dual-channel axion search was demonstrated. These searches were performed simultaneously on the  $TM_{010}$  and  $TM_{020}$  using two separate receiver chains (section 3.2). Bandpass filters connected directly to the antennas separated the high frequency from the low frequency data and prevented ch1 data from being lost through the ch2 port. Simultaneously operating with two modes meant finding tuning rod configurations that gave rise to good data on both channels. This involved avoiding mode crossings, low form factor regions and frequencies where the hardware was out of specification while at the same time maximizing the science goals of the experiment for both channels. By identifying, the frequency and form factor regions that were not allowed, a set of tuning configuration constraints was combined to reveal where data could be taken.

COMSOL was used to solve for the electric field for multiple modes at all possible rod configurations, the form factors were computed and used to sort the modes. By picking out the mode with the highest form factor for all solutions, the  $TM_{010}$  was identified. At higher frequencies, the situation was a little more complicated. Three different modes in distinct frequency regimes were found with form factors similar to what would be expected for the  $TM_{020}$ . One of these modes was picked out as having a frequency range that complimented the Ch2 RF-constraints. Figure 4.13 shows the frequencies and form factors for both the Ch1 and Ch2 modes. The lower left hand corner of each plot represents the configuration where the rods are near the wall of the cavity. The upper right hand corners represent the configuration where the rods are at the center of the cavity. The figure shows the entire tuning range and sensitivity to axions for both modes at every tuning configuration. Taking good data on these modes involved looking at the bottom plots and picking regions where the form factors were high. Hardware frequency constraints were also considered. While the rich mode structure was very complicated, by ignoring the unallowed and low form factor regions, tuning trajectories were found that provided good data on both channels.

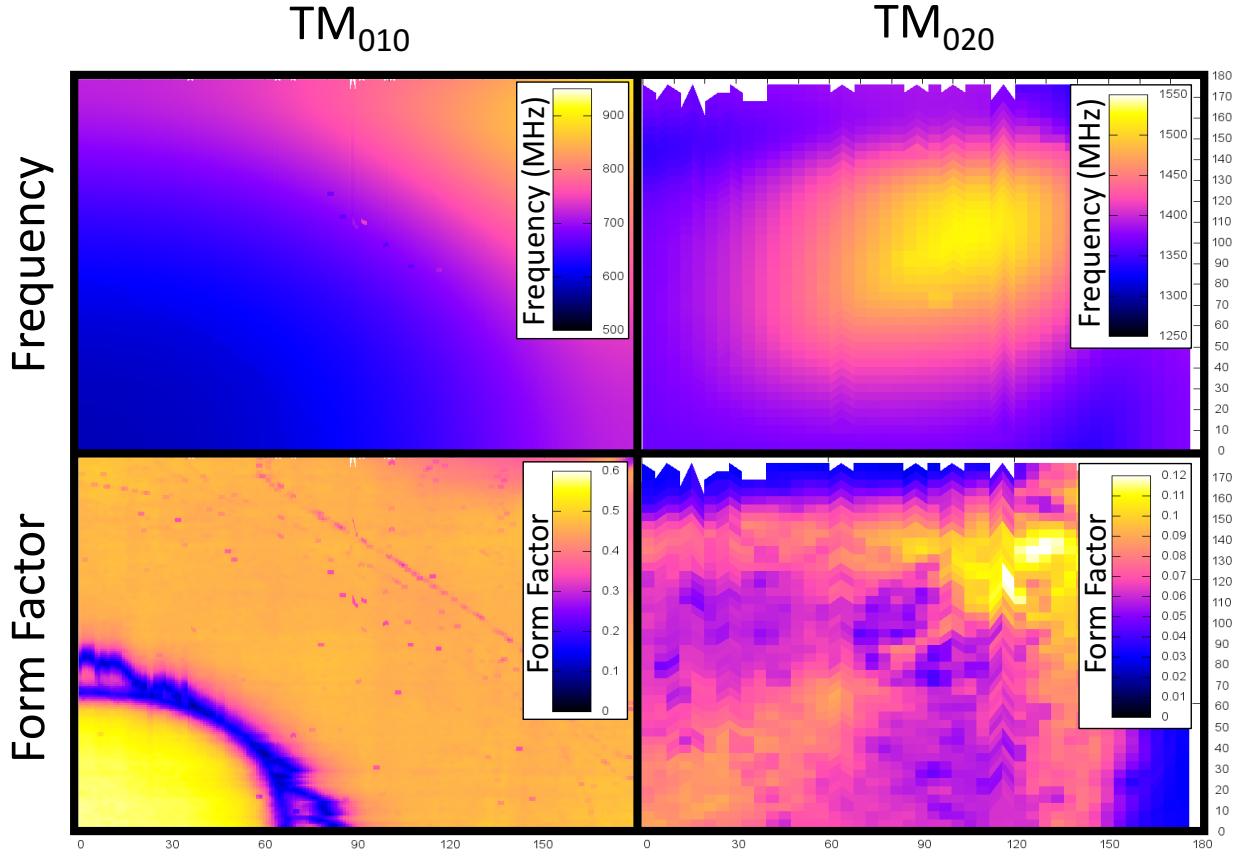


Figure 4.13: Frequency and form factor simulations of the  $TM_{010}$  and  $TM_{020}$  modes for all tuning rod configurations.

Once allowable tuning trajectories have been identified, simulated mode map representations of specific trajectories can be compared directly to measurements. Figure 4.14 shows the allowable data regions and rod configurations that I compiled in planning several ADMX data runs. This information was used to guide data taking in 2016 and 2017.

#### 4.3.3 Candidate Re-scans

A run plan must also specify the protocol for candidate rescans. After many spectra have been co-added, persistent peaks and potential axions are periodically observed. A certain

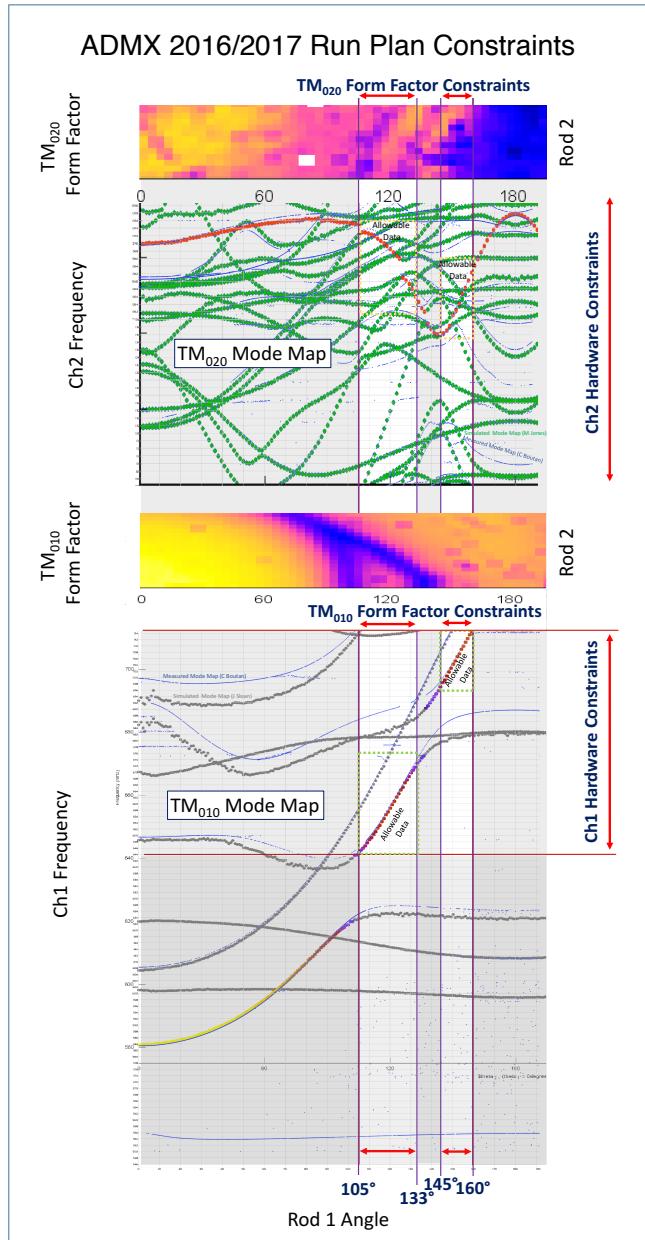


Figure 4.14: Simulated (gray/green) and measured (blue) mode maps for the  $TM_{010}$  and  $TM_{020}$  modes. Hardware constraints and low form factor regions for both channels define clear allowable tuning trajectories where good data was taken in 2016 and 2017.

amount of time must be allocated to the rescan process when the cavity is tuned back to those peaks and more spectra are taken. If peaks are statistical, the signal will not persist. If a peak is persistent, an axion can be differentiated from RFI by ramping the magnet down. Depending on how often rescaning occurs, this operation may involve tuning back over hundreds of MHz. Rushing this process would result in a lot of heating that could overpower the dilution refrigerator. The optimal protocol for candidate rescanning is being considered.

#### **4.4 Data-Taking**

This section shifts the focus from the cavity mode structure to the cadence of data taking and software behind the automated axion search.

##### **4.4.1 DAQ**

All of the hardware is incorporated into the DAQ via an Experimental Physics and Industrial Control System (EPICS) platform. A Lua interpreter runs custom Lua scripts which can “SET” or “GET” EPICS variables that correspond to measurements or physical hardware states. All data is stored in a Structured Query Language (SQL) database. Web interfaces tailored for monitoring and/or controlling certain aspects of the experiment query the database and display the information. These interfaces can also be made to send requests to the database to execute specific scripts. These scripts are written directly into a web interface and can be executed via a command entry from anywhere in the world. Typically, scripts pull data out of the database or directly from the experimental hardware via EPICS, custom functions are then be used to manipulate the data and then save it back to the database. The SQL database stores the working memory of all parts of the experiment and serves as a mediator between EPICS, the Lua interpreter and web interfaces. An example of this interaction can be seen in the obtaining of a network analyzer measurement in figure

## 4.15.

1. A button on the bottom left hand corner of Figure 4.15 can be pressed that requests that the database run the Lua script on the right.
2. This Lua script then talks to EPICS and is able to control the network analyzer by setting the output power, resolution, center frequency, span frequency etc. EPICS is then told to sweep the network analyzer and return the data. The Lua script then converts real and imaginary values to powers, creates an array of frequencies, fits the data with a Lorentzian, identifies a peak and saves all of this information into the database.
3. By querying the datatable in which the network analyzer data was saved, the website is then able to parse the information and display the  $S_{21}$  measurement.

Different types of data is saved in different SQL datatables. Network analyzer scans are saved in a table called “na\_log”, digitizations are stored in “digitizer\_log” and the actual axion search data is saved in “axion\_scan\_log”. Each new entry into these tables is auto assigned a scan number. This allows “axion\_scan\_log” to save network analyzer and digitizer scan numbers as pointers rather than creating data duplicates. “DAQ variables” are variables that can be stored in the database and used like global variables. These variables are used extensively to allow parallel scripts to listen to each other and enable the ability to steer the progress of a script after it has been executed.

4.4.2 *Cadence*

The end goal of the experiment and central purpose behind the process of taking data is to digitize the signal from the cavity, make a grand power spectrum and look for evidence of an axion signal in the noise. Given that at any instant, a high Q cavity experiment is

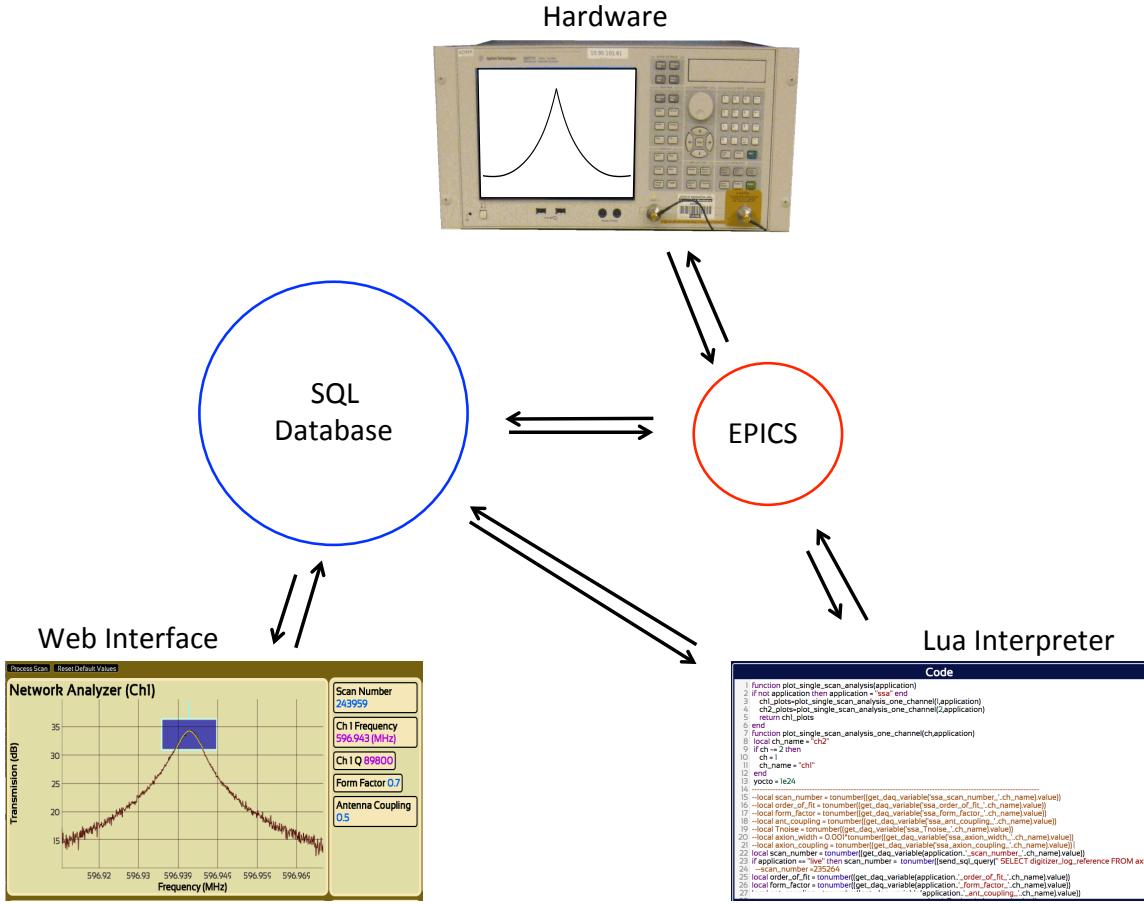


Figure 4.15: The ADMX DAQ is built on an EPICS platform. Lua scripts control the experiment hardware via EPICS and data is stored in an SQL database. Web interfaces query the database and run lua scripts to monitor and control the experiment.

only sensitive to axions over a narrow ( $f/Q$ ) frequency range, many digitizations over many different tuning configurations are needed to cover a wide frequency range. By iteratively tuning the cavity and digitizing the signal at that frequency, the experiment is able to co-add many spectra together to form a grand spectrum and thereby achieve a wide frequency range axion search and obtain enough integration time to reach a given target sensitivity. When the data is analyzed, all individual spectra of like frequencies are co-added to obtain a grand

spectrum. By adjusting the tuning frequency step size, frequency overlap of spectra can be regulated to hit a given target sensitivity.

Practically, the process of taking data involves performing a list of measurements and operations over and over again. The rhythm at which the experiment carries out its list of operations (like, tuning, digitizing etc) defines the “cadence” of data taking. A simplified algorithm for the cadence of data taking requires one more step; a swept response measurement (described in section 4.1.2) after tuning and before digitizing. The cavity mode center frequency is then used to determine the frequency at which the local oscillator must be set such that the axion signal is mixed down to 10.7 MHz.

The center of figure 4.16, shows the cadence for a toy model of an axion search at the top of the figure. The program is initialized by setting essential initial parameters of the network analyzer, digitizer, etc. The program then enters the body of the while loop and begins the list of operations to be repeated over and over. The cavity is tuned, the  $TM_{010}$  mode frequency is measured with a network analyzer, signal coming out of the cavity is digitized, all relevant data is saved, and the whole process repeats again. Each of these high-level operations is made up of a set of many low-level subroutines. For example, the high level task of sweeping the cavity involves, setting the switch box, getting the expected  $TM_{010}$  center/span frequency, setting various network analyzer parameters, making an  $S_{21}$  measurement, fitting the data with a Lorentzian and saving parameters from the fit. Notice that the central frequency and  $Q$  are saved as “DAQ variables” introduced in section 4.4.1 and made globally readable to all other scripts. The Digitize script, when run will use this information to set the frequency of the local oscillator. Also, the next time the Sweep script is run, it will use these values as expected initial guess parameters and set the network analyzer to sweep where the mode was last. By making the network analyzer span frequency a function of the  $Q$ , the span ( $\propto \frac{f}{Q}$ ) will dynamically change to hug the resonance, ensuring that the measurement doesn’t find another mode and that there is enough resolution for

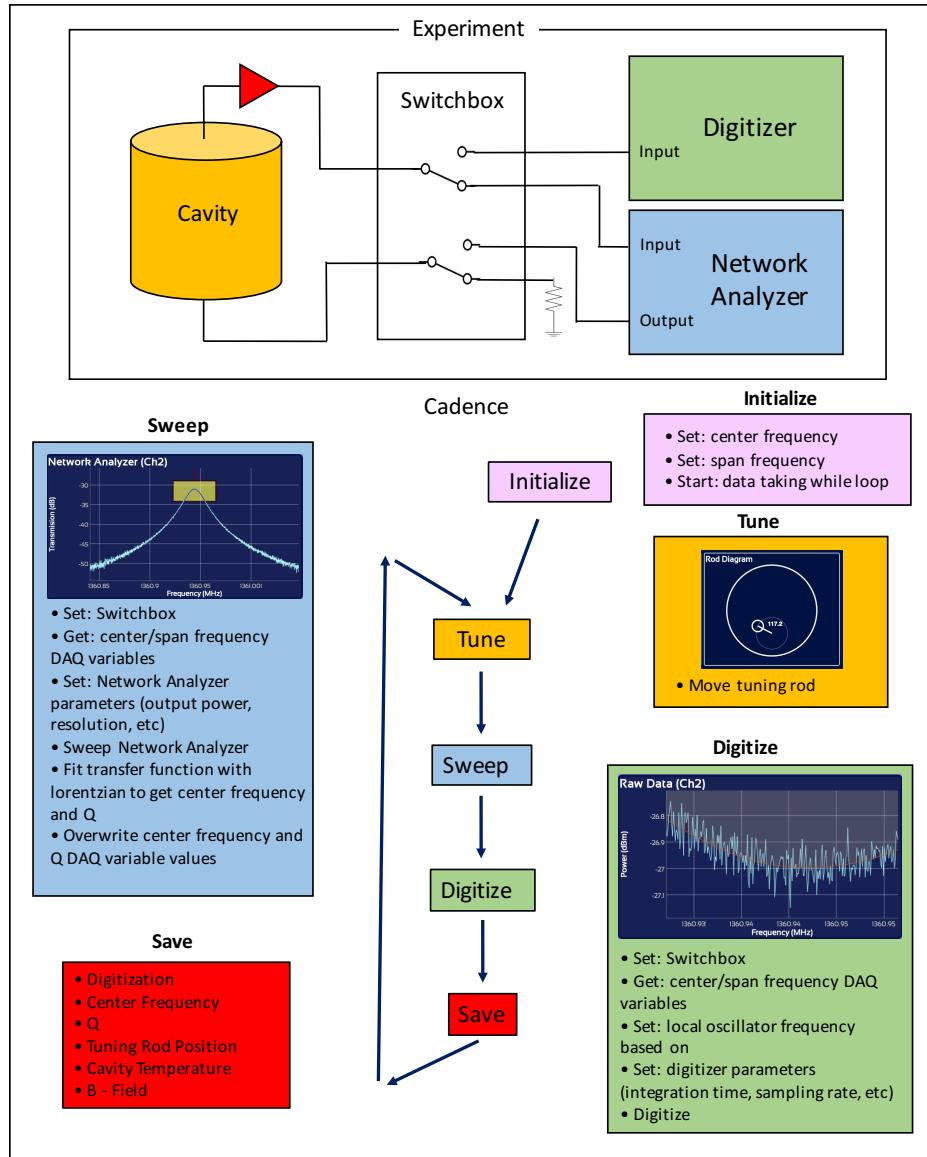


Figure 4.16: The cadence of data taking for this simplified experiment. 1) hardware is initialized, the cavity is tuned, the center frequency of the cavity is measured with a network analyzer, signal out of the cavity is digitized, the data is saved and this process is repeated over and over again. A switchbox allows the experiment to toggle between taking digitized spectra and making network analyzer measurements.

accurate fits. Provided that tuning rod step size remains small and that no mode crossings are encountered, this simplified experiment and algorithm above will reliably tune and track a mode for months on end while at the same time taking digitized spectra.

For a full description of the cadence of the ADMX experiment, this toy model needs to be expanded to include the following operations

- Antenna reflection measurements
- Depth adjustments of the antenna
- Occasional wide transmission sweeps for navigational purposes.
- Occasional injections of axion-like signals

For the dual-channel search,  $TM_{020}$  measurements must be added to this list. Data taking is also made more complicated by the fact that mode crossings will be encountered, the conversion between rod step size and frequency is not constant and the task of automating antenna coupling is difficult. These additional complexities are covered in the following subsections.

#### 4.4.3 *Sidecar Data-Taking Script and Web Interface*

The Sidecar data-taking script is a while loop that calls sub-routine functions for each of the mentioned cadence operations over and over again. Each subroutine communicates with a RF-switchbox which allows for the switching between digitizing, calibrating and various network analyzer measurements. This central data-taking script communicates with a web-interface via the daq variables mentioned earlier. For example, at the start of each subroutine, the operation being performed and the time it takes to complete are stored as two separate daq variables. The web-interface then queries these variables and displays

them. Figure 4.17 shows all of the options for the current state of the experiment. The lit “Sweeping” indicator reveals that the experiment is currently performing a transmission measurement and that the last time it did so, the script took 3 seconds to complete. The green timer indicators point out which sub routines are taking too long to execute and help diagnose which processes are broken.

Current State								
sweeping	3	bypassing	1	digitizing	119	tuning	3	stalled
widescanning	12	coupling	0	saving	0	mapping	27	stopped

Figure 4.17: The current state of the Sidecar experiment. The web interface indicates that the cavity is currently being swept with the network analyzer. The number of seconds the experiment recently took to perform a given operation is shown in green.

Buttons on the web-interface, can be pushed to run Lua scripts or alter the values of daq variables which will, in turn affect the way certain scripts are executed. For example, by clicking a button called “Digitize” the value of a particular daq variable will be toggled. Every time the data-taking script prepares to digitize, it first queries the database for this variable and the script will know to either begin digitizing or ignore the digitizing sub-routine.

#### 4.4.4 Navigating Mode-Crossings

At a given tuning, with only wide swept measurement of the cavity, keeping track of the correct mode is not trivial. Several features of the Sidecar software aid mode-navigation. Figure 4.18 shows the piece of the web-interface associated with tuning. The bottom of the figure shows a swept response of the  $TM_{010}$  next to a collection of tuning options. The center frequency and Q from the Lorentzian fit are also displayed. One tier up in the figure reveals more context about the current tuning of the cavity. On the right hand side, the position of the tuning rod is displayed and on the left, a wider swept measurement showing a second mode on the horizon. A peak finding algorithm has identified and fit both peaks. Finally,

at the top of the figure, the recent fits to the wide swept measurements is plotted over a static mode-map which provides a context for the current state of the mode structure. The yellow points are from a simulation of the  $TM_{010}$  mode plotted over a measured mode-map in blue. The small red collection of points represent measurements from the last 30 minutes for the mode on which data is being taken. The fact that the red dots are in agreement with the yellow simulation indicates that the experiment is following the correct mode. The ability to see new measurements over an old mode-map also allows the piezo encoders to be re-calibrated after reaching base temperature. The calibration is reset by tuning the rod through the center of the cavity and measuring the difference in the red and blue points at the top of the mode-map. Once the rod encoder is calibrated, the yellow simulation can be used as a lookup table to guide the experiment through mode-crossings. By clicking the “Simulation Aided Tuning” button, the sub-routine that performs transmission measurements will start comparing the fit frequency with the frequency predicted by simulation and current piezo rod coder value. If the difference between the measured and expected frequencies is too great (a parameter also set by the website), the cavity is swept again but this time with a center frequency equal to the predicted frequency. Figure 4.19 shows, once again the narrow swept response of the  $TM_{010}$  mode, the wider swept response of the local mode structure along with a vertical green bar indicating the expected  $TM_{010}$  frequency and the mode map. The red points represent the mode on which the data was taken over the last day. In light blue are shown other non -  $TM_{010}$  peaks identified in the wide sweeps. The fact that the red measurements follow the yellow simulation in the figure demonstrates the experiment’s ability to tune through complicated crossings without getting lost on the wrong mode.

#### 4.4.5 Adaptive Scan Rate

The default tuning option is to take fixed piezo steps because it is safer than making the step size a variable. However, the goal is to achieve a smooth and flat SNR which ultimately

## Mode Tuning and Navigation

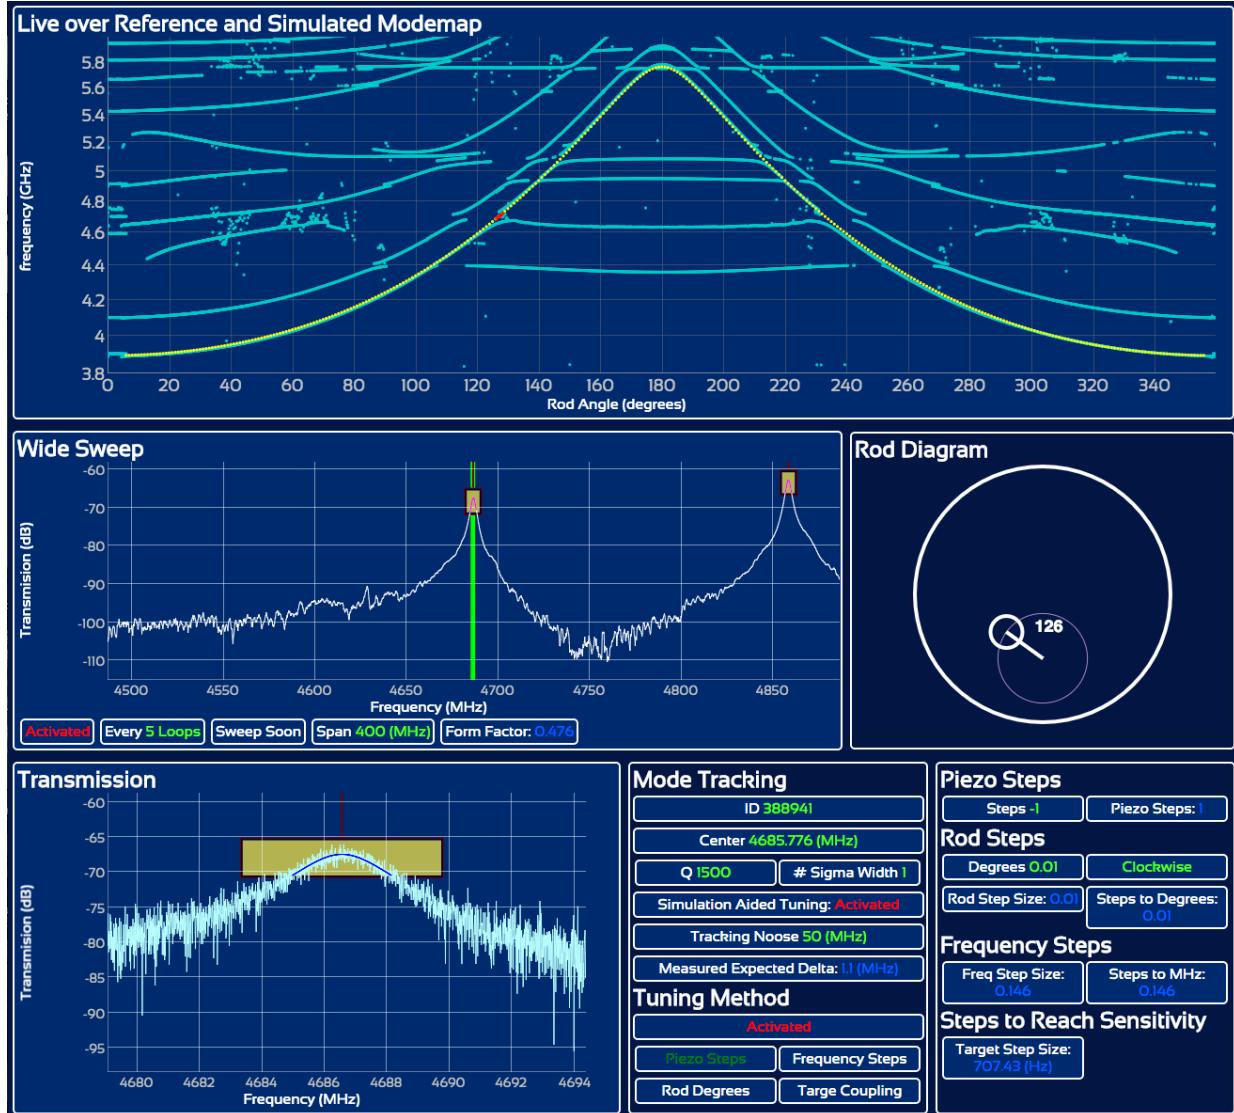


Figure 4.18: Mode structure navigation portion of the web interface. Swept response measurements show the  $TM_{010}$  mode in the context of the greater mode structure. Recent measurements are plotted over a static mode-map. A diagram of the cavity shows the current position of the tuning rod. Buttons allow different tuning options to be selected.

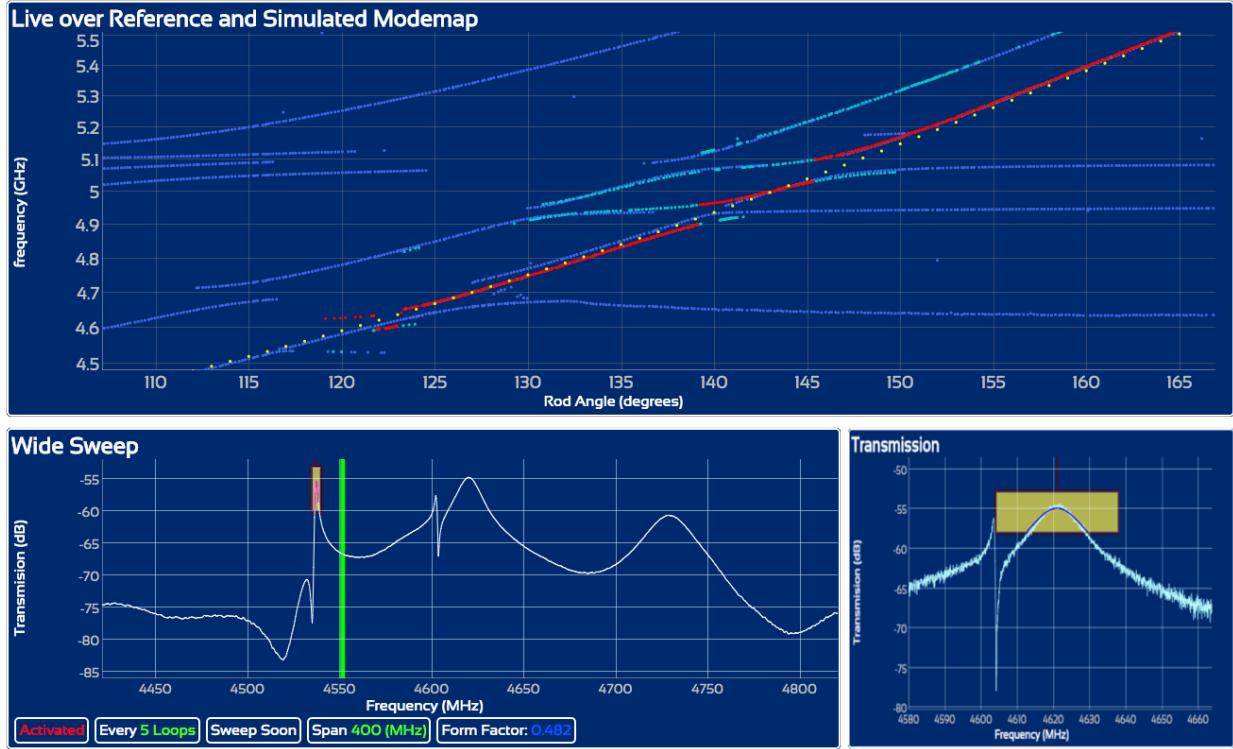


Figure 4.19: Mode structure navigation portion of the web interface showing the recently taken data (red) over a static (blue) mode map. The recent history of the experiment shows that the mode-tracking algorithm has successfully tracked the  $TM_{010}$  through several mode crossings without getting lost.

depends on the magnetic field,  $Q$ , system temperature, form factor etc. Also, as the rod steps along at fixed angle intervals, the slope in the mode-map may change, causing the frequency step size to increase or decrease. To aid the human driving the experiment, the past history of the experiment, the expected SNR and ideal step size is computed for every iteration of the data-taking loop. Figure 4.20 shows the tuning rate required to hit a given axion coupling for a given SNR set by the user at the center of the figure. On the left, the required number of spectra needed to achieve the KSVZ and target coupling versus SNR. Plotted on either side of the the vertical target SNR is the required and current number of

overlapping spectra. Under perfect operating conditions, these two cross-hairs would line up. At the right of the figure is the same information put in terms of frequency step size. The plot shows the current and ideal step size to reach the KSVZ and target couplings. The displayed information can be used passively by the user, to adjust the tuning speed of the experiment or the “Target Coupling” tuning option can be selected in figure 4.18. The software will then look at the recent history of the experiment, throw away outliers, fit frequency step size versus piezo steps and adjust the throttle of the experiment accordingly. This means that if the magnet is ramped down or if the temperature of the cavity is raised, the tuning rod step size will be reduced accordingly. This functionality was demonstrated at room temperature but not used during the cold data run. This was because the software does not, at present consider the heating restrictions on the dilution refrigerator and in the unlikely event of a runaway rod, the cryogenics could be damaged.

### Required Number of Integrations and Frequency Step Size to hit Target



Figure 4.20: Scan rate and tuning throttle portion of the web interface. Periodically, the experiment queries recently taken data, predicts trends and calculates an ideal frequency step size to hit a target sensitivity for a given SNR. Left: number of spectra needed to achieve a target sensitivity versus SNR. Right: frequency step size required to achieve a target sensitivity versus SNR.

#### 4.4.6 Maintaining Antenna Coupling

Periodically, as the cavity is tuned, the antenna will slowly drift from critical coupling and become either under or over coupled. If this happens overnight or when no human is around, the quality of the data will suffer. However, slight vibrations, intermittent antenna grounding issues and other unknown anomalies make it very difficult to automate the coupling of an antenna in a stable way. In order to better understand this process, the antenna coupling,  $Q$  and frequency of the cavity were measured for many antenna depths. Figure 4.21 shows the results of this test which are identical to the measurements in figure 3.13 from the previous chapter except that, in this case, the rod is at the wall. As the major antenna is inserted, the frequency and the  $Q$  (replaced by  $Q/100$  for plotting reasons) are reduced. In the figure, the absolute value of the coupling is plotted such that the peak value represents the optimal coupling. It is important to notice that near critical coupling, the line is relatively symmetric about the peak. Because of the symmetry, it is difficult to know if the antenna is under or over coupled based purely on the single measurement of a reflection trough. By making several measurements at slightly different antenna depths and following the slope of the curve, the location of the peak may be inferred. The down side to this technique is that there are sometimes local minima and glitches cause by antenna grounding errors which can throw off the peak location prediction. Another method is to utilize the  $Q$ , which is clearly not symmetric about the optimal coupling and can be used to indicate whether or not the antenna should be removed or inserted into the cavity. If  $2Q_L/Q_0$ , is  $> 1$  the antenna should be inserted; if it is  $< 1$  it should be removed. However, by comparing the two figures 4.21 and 3.13, representing two tests that occurred on the same day, with the same electronics, it becomes clear the  $Q_0$  is different for different rod positions. An in situ measurement of  $Q_0$  could be done periodically during data taking but not without excessive heating. The uncertainty in  $Q_0$  makes this technique ill-suited for situations where the antenna is already very close to being critically coupled.

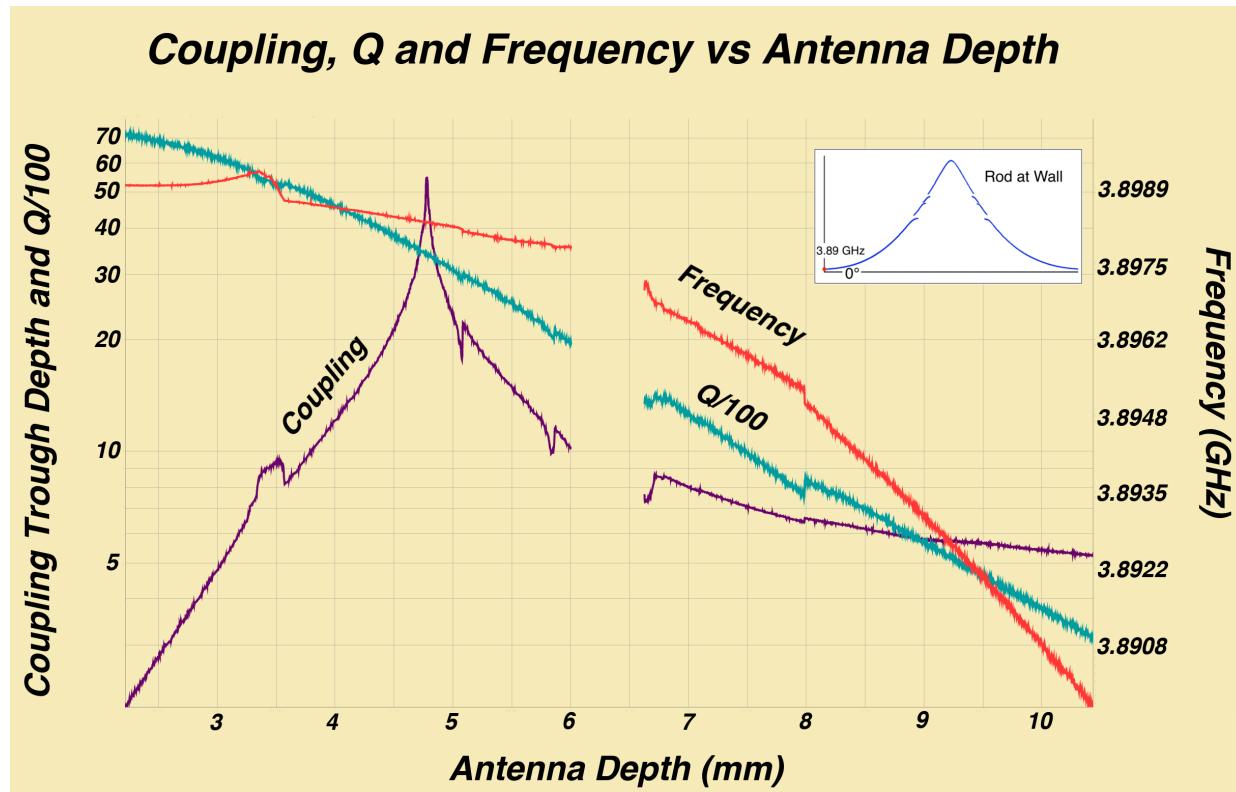


Figure 4.21: Sidecar antenna coupling, Q and cavity frequency versus antenna depth where the coupling trough depth is expressed as a positive number.

The sidecar auto-coupling algorithm works as follows. If the coupling is very far from the critical coupling, a measurement of  $Q_L$  with some assumption about  $Q_0$  indicates whether or not the coupling is under or over coupled. The difference between critical and current coupling is then used together with the shape of the measurements made in figures 4.21 and 3.13 to determine the approximate number of steps to critical coupling. The number of steps is made to undershoot the peak in order to see a trend in the increasing slope. As the coupling approaches the optimal coupling, the depth size adjustments are reduced with increasing slope.

The algorithm for maintaining an already good coupling is different than the one used in

scenarios where the coupling is far from ideal. After much experimentation, it was discovered that the optimal method for maintaining a good coupling was to have the software learn from past experiences. If the last time the antenna tried to couple itself resulted in a success and, for example, required removing the antenna 20 piezo steps, the antenna would perform the same operation for the next attempt. By doing what it had just done last time, the antenna chose the correct direction towards critical coupling more than 90% of the time. If after this procedure, the antenna had not reached critical coupling, it would repeat the operation until critical coupling was achieved. If a trend of worsening coupling was observed, the direction of the piezo motion was reversed until either the desired coupling was achieved or the script ran out of the time it had been allotted to couple the antenna. Because of local minima in the coupling versus antenna depth it was determined that the algorithm needed to be fairly insensitive to individual measurements such that a clear trend was noticed before the piezo was given orders to alter its behavior.

While the auto-couple software was demonstrated to work at room temperature and was used to produce a complete map of the  $TM_{010}$  while maintaining a critical coupling, it was not used in the 2016 data run. After the insert was put into the bore of the magnet, a set screw came loose and caused the antenna to remain stuck in an over-coupled state. Since then, the set screw was fixed but the auto-couple option, like the adaptive scan rate option has not been used because of the need for further optimizations that consider the heat load on the dilution refrigerator.

#### 4.4.7 Data-Taking together with Real-Time Analysis

In past incarnations of ADMX, the act of taking data and the job of analyzing that data were two separate processes. Data would be taken for a few months to a year, time-series digitizations would be stored on a hard drive and eventually analyzed. However, since it was my task to write the data taking routines in Lua and having access to the SQL database, I

had the opportunity to write a program on the same platform that could be run in parallel with data taking and used in a feedback loop to guide the experiment. This has the benefit of informing the data-taking routine when the measured SNR is low or when a persistent peak has been detected. Given this information, the experiment, in real-time has access to where where axion-like signals have occurred and where more data needs to be taken to flatten the SNR. The “Real-Time Analysis” software was designed to work in both live and after the fact situations, such as in the final analysis discussed in the following chapter.

## Chapter 5

# ANALYSIS

In this chapter I present my analysis of the 2016 Sidecar data. The analysis was performed in Lua, the scripting language of the ADMX data acquisition. First I describe the taking of the raw data and give an overview of the whole data set. I then discuss how individual spectra are conditioned and weighted in preparation for the stage at which all of the data is co-added in quadrature. Just before the co-adding stage, the data associated with individual spectra go through a quality control in which each contributing variable is examined. Spectra associated with clearly bad or even slightly suspicious measurements were cut. Finally, I discuss the grand analysis where many individual spectra became a single spectrum and present new limits on  $g_{a\gamma\gamma}$ .

### 5.1 2016 Sidecar Data Run

This section describes the data run, the performance of the experiment, and gives an overview of the resulting data set.

#### 5.1.1 Run Plan

Given the tiny volume of the cavity and the lack of a quantum-limited front-end amplifier, this pathfinder experiment was not equipped to set limits on QCD (KSVZ or DFSZ) axions. In order to demonstrate the rotary piezo's ability to tune the experiment over a wide frequency range, a target sensitivity of  $10^{-13}/GeV$  was picked. The higher frequency end of the cavity tuning range, between 5.1 and 5.8 GHz was chosen for the data run. This corresponded to the tuning rod angle being rotated between  $145^\circ$  and  $180^\circ$ . This region was selected because

of the few number of mode crossings and because it was the goal of the experiment to pave the way for higher mass searches. Also, this frequency range corresponded to the steepest section of the mode-map which meant less heating per frequency step size. Given several months of high field data, the Sidecar experiment was projected to exceed the sensitivity of the CAST [36] and future IAXO experiment [56] over a mass range greater than the mass range excluded by past ADMX searches.

### 5.1.2 *Timeline*

The experimental apparatus was inserted into the magnet on the 10th of June and extracted on the 6th of October, resulting in 17 weeks in which the experiment sat in the bore. Soon after its insertion, it was discovered that there was a leak in the helium reservoir which required further studies. Given that the highlight of this run was the commissioning of the new ADMX dilution refrigerator, priority was given to cryogenic tests which made the data run inefficient as an axion search. On the 9th of August, the magnet was ramped to 1.7 T and data taking began. This low field was maintained for just under 8 weeks so as not to push the system during this commissioning run. Finally, to achieve some high field data, at the end of the run, the magnet was ramped to 5.5 T over a weekend. The magnet was then ramped down, the experiment was allowed to warm up and the insert was removed. Due to the location of the Sidecar cavity being in the fringe field of the magnet (discussed in section 5.3.2), this meant that the Sidecar search had 8 weeks of 0.78 T data and three days of 2.55 T data.

### 5.1.3 *Performance*

Apart from the loosening of a set screw, which caused the antenna to remain over-coupled throughout the run (discussed in section 5.3.2) the performance of the experiment was excellent. In frequency step sizes ranging between several kHz and tens of kHz, the rotary piezo

smoothly tuned the axion search back and forth between 5.090 GHz ( $21.050\mu\text{eV}$ ) and 5.799 GHz ( $23.983\mu\text{eV}$ ), covering a 709 GHz ( $2.93\mu\text{eV}$ ) range. During this time, 33,500 100-second digitized spectra were taken resulting in 930 hours of integration time. Over the course of the data run, the mechanical systems (barring the over-coupled antenna), RF-systems and data acquisition software worked well.

#### 5.1.4 Rescans

1 week was allotted to the rescanning of frequencies with power excesses that could represent the detection of an axion-like signature. Post data-taking analysis revealed these candidates were actually artifacts of the live analysis (fixed after the end of the data run), so the rescan data was simply combined with the rest of the data set to produce an upper bound on axion-photon coupling.

### 5.2 Analysis Introduction

The objective of the analysis is very simple: discover the axion or set new limits on  $g_{a\gamma\gamma}$ . To do this, many raw power spectra need to be conditioned and co-added to form a single power spectrum in which the axion lineshape may or may not be hiding. This section is an overview of the analysis process and sets the tone for the “single-scan”, “quality control” and “grand” analysis stages that follow.

#### 5.2.1 Methodology

This analysis was originally designed to mainly run in parallel with data taking and be a light-weight version of a time tested ADMX process. It builds on Ed Daw’s “gold standard” analysis from 1998 [34]. It was made to work in a live environment as well as a more intensive offline environment and has many visualization features designed to prevent poor data from propagating through the pipeline unnoticed. Written in the same language that controls

the data taking, and visible to everyone on the project, this analysis was designed to be transparent to the collaboration and one day improved upon. In each stage of the analysis, I have attempted to make few assumptions, and be as meticulous and conservative as possible.

### 5.2.2 *Live Analysis vs Offline Analysis*

Any challenge associated with performing a live analysis is inherently due to the fact that data must be processed with incomplete information about the data set as a whole. The following is a list of some of the issues that made this live analysis unique and affected the design of the pipeline.

**Some signs of poor quality data are only visible as a trend over time:** While, automated cuts can be made on low-field, low  $Q$ , high temperature data immediately, some evidence for low quality data isn't apparent without the context of more data. The prime example of this is data taken when the experiment has tuned through a mode crossing and has started to track the wrong mode. The transition is a gradual change in tuning speed and can't be identified without seeing a trend in the frequency trajectory in the context of a mode map. To combat this issue, the analysis was made reversible. As data is taken, if it passes the automated cuts, it is presumed to be good data. When the data is co-added to form a single grand spectrum, a separate list keeps track of which data has contributed and which data is believed to be good. If at some point, data that was believed to be good is determined to be bad, that data can be flagged and then subtracted from the set. The software also prevents data from being added twice or subtracted if it were never added in the first place.

**The frequency span of the data changes as it's being processed:** In the co-adding stage of the analysis, many spectra with their own bin frequencies and resolutions get “re-gridded” to a much larger array with its own native resolution and quantized

set of bin frequencies. If a data set is being processed after the fact, this step can be done all at once and with full knowledge of how wide the frequency span of the “grand array” needs to be. While a run plan gives a guideline for the start and stop frequencies, this plan is often subject to change. The Sidecar live analysis was made to grow organically as new data is taken. When a new spectrum is saved to the data base, the analysis queries this data along with data that has already been co-added. Based on the frequency spans of both spectra, the analysis creates a new empty grand array, re-grids both spectra to it, re-defines the resulting array as the “grand spectrum” and saves it back to the database. By continually saving the growing data back to the database, the analysis doesn’t need to start from scratch in the event of a power outage or a script runner crash. If the data-taking is ahead of the live analysis or if the analysis is being re-run after all data has been taken, the software can be set to process all data at once without intermediate saves.

**Spectral structure can’t be corrected with an averaged receiver response:** The process of removing the receiver structure from digitized spectra often involves dividing the data by an averaged receiver response. This response can be created in advance of the data run by injecting white noise through the system but it is easier to get the response by averaging a lot of spectra from the data set itself. If this procedure were used in the live analysis, the corrected spectra at the end of the data set would have been treated on unequal footing with the data taken at the beginning. A high pass filter routine was developed to help with the structure removal process.

### 5.2.3 Pipeline Overview

As the experiment takes data, digitized power spectra are saved to the SQL database and assigned a “scan number”. The data taking script then saves a list of pointers that point to all of the data associated with that scan number and links it to the name of the data run.

The name of the data run paired with the data in this work is ‘*first\_data*’. The analysis script is then passed the run name as an input parameter, a database query identifies all of the scan numbers linked to that name and the script enters a while loop where, one at a time, each scan is individually processed, poor quality scans are discarded and finally, high quality scans are co-added to the grand spectrum. A conceptual block diagram of the analysis is shown in Figure 5.1. Though an accurate block diagram of the analysis is too complicated to show here, each stage of the analysis tends to fit into one of three categories: “Single Scan Analysis”, “Quality Control” and “Grand Analysis”. The following is a short description of each of these stages.

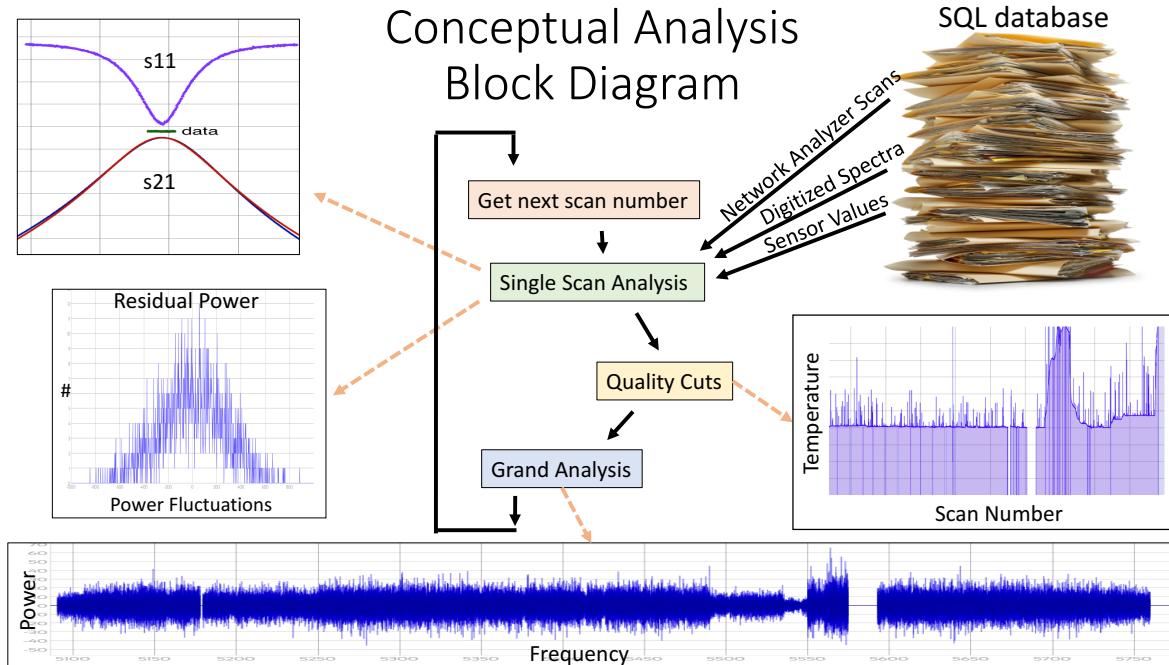


Figure 5.1: A conceptual block diagram for the analysis. Data is repackaged, spectra are conditioned, poor quality data is cut and spectra are co-added to form a grand power spectrum. Limits on then set on  $g_{\alpha\gamma\gamma}$ .

**Single Scan Analysis:** When a scan number is passed to the single scan analysis, relevant data from the database is queried and packaged into scans. In this stage, expected power is calculated, raw spectra are conditioned and data relevant to the co-adding stage is weighted such that low and high quality data can be summed in quadrature.

**Quality Control:** The output of the single scan analysis is passed to quality control. This stage keeps a log of every relevant data parameter versus scan number. This stage makes automated cuts on poor quality data but also has a web interface such that a human can easily look through the data and make cuts by hand. This stage also keeps a log of what data has contributed to the grand spectrum.

**Grand Analysis:** The output of the quality control is passed to the grand analysis. At this stage, one at a time, individual spectra are added to the grand spectrum. “the grand spectrum” is actually an over simplification. Several grand spectra are created: number of contributing spectra, measured power, expected SNR and RMS-noise. From these grand spectra a limit on  $g_{a\gamma\gamma}$  is set.

The following sections describe each of these stages in more detail and conclude with new limits set on the axion-to-two-photon coupling.

### 5.3 Single Scan Analysis

The job of the single scan analysis is to take raw data out of the database, calculate the expected axion power, condition raw spectra and turn them into a measured power in the cavity. Finally these quantities are weighted such that less sensitive data contributes less to future stages in the analysis.

### 5.3.1 Data Packaging

All of the data associated with a given digitizer scan number is removed from the database and packaged into a Lua structure. These data packages are referred to as “scans”. Figure 5.2 is an example of all of the raw data associated with a scan. The single scan analysis web-interface allows a user to view the cavity transmission, antenna reflection and digitized spectrum all at once. This allows a human to check to see that the local oscillator hasn’t stopped responding or that the digitizer span is set appropriately to match the Q of the cavity. If there happens to be a dip or an anomaly in the digitized power spectrum, the explanation is sometimes visible in one of the network analyzer measurements. A scan also contains relevant resonant mode parameters (such as  $Q_L$ ), digitizer parameters (such as integration time) and meta data (such as rod angle).

### 5.3.2 Calculating Expected Axion Power

By modifying equation 2.23 to include the sensitivity off resonance, plugging in  $\rho_a = 0.45\text{GeV}/\text{cc}$  and  $V = V_{cavity} - V_{rod} = 0.367$  liters, adjusting the normalization to reflect normal operating parameters for this run and including the full expression for fraction of power out of the antenna, the expected axion power per bin is

$$P_{per\ bin} = 5.7 \times 10^{-27} W \left[ \left( \frac{f}{5\text{GHz}} \right) \left( \frac{Q_L}{2,000} \right) \right] \left[ B^2 \left( \frac{C_{mnp}}{0.5} \right) \left( \frac{1 - 2S_{11}}{1 - S_{11}} \right) \left( \frac{1}{1 + 4Q^2(f/f_0 - 1)^2} \right) \right] \dots \times \left[ \left( \frac{g_\gamma}{0.97} \right)^2 \frac{Axion\ Width}{Digitizer\ Resolution} \right] \eta \quad (5.1)$$

where the pre-factor is a typical total expected axion power out of the a Sidecar cavity, values in the first set of brackets has already been measured and quantities in the second set of brackets must be calculated from the raw data. Contained in the final set of brackets is

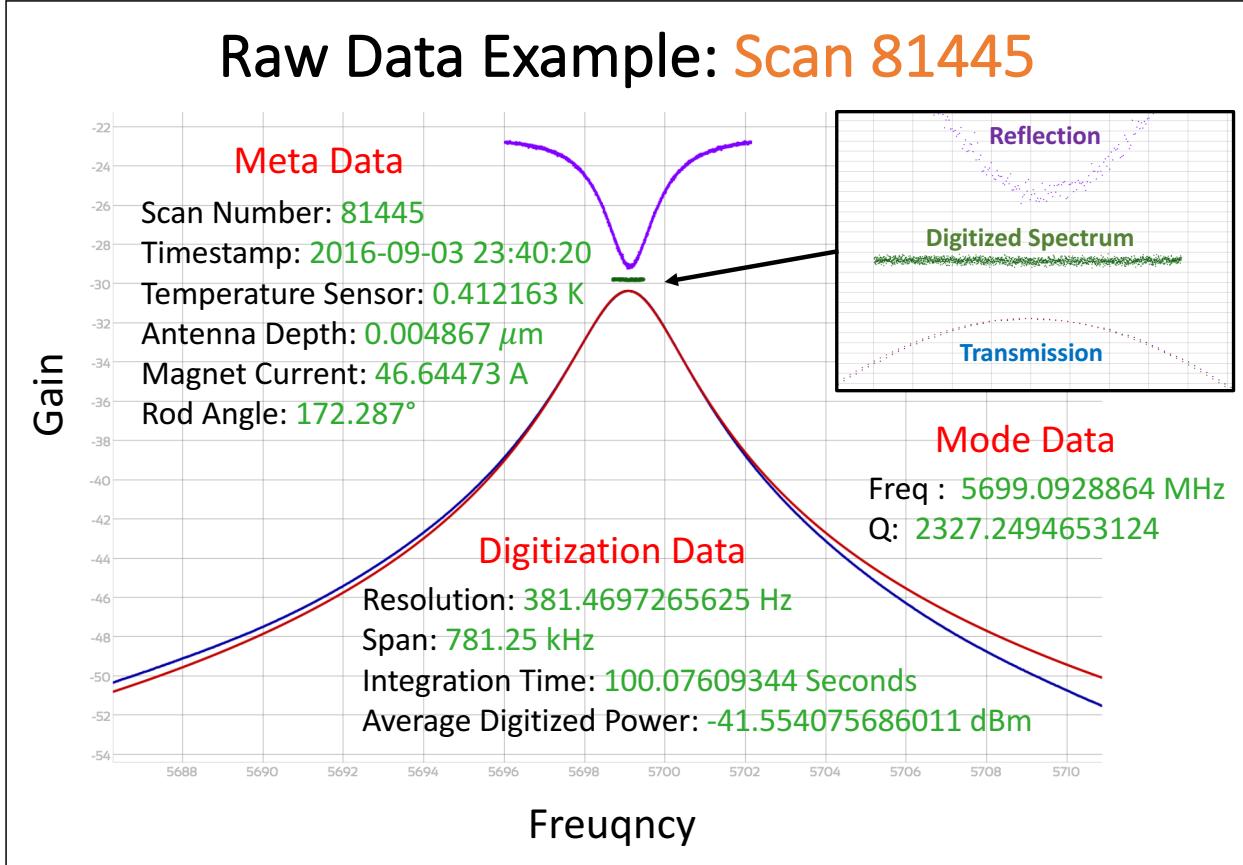


Figure 5.2: Example of a raw data scan. A raw data scan includes a digitized spectrum, cavity transmission measurement, antenna reflection measurement and meta data related to the digitizer, mode structure, physical temperature, piezo motors magnetic field.

information ultimately fixed by nature, but is practically chosen by the experimenter during the analysis process. The (axion width/digitizer resolution) term is a loss term that takes into account the fact that this analysis assumes that the axion width is  $\sim 5$ kHz but the single scan spectral resolution of the search is 381.5 Hz. Therefore only  $\sim 1/13$ th of the axion line shape would be seen in a single bin. While this naive term assumes that axion lineshape is a box filter, the  $\eta$  term takes into account additional losses associated with the actual lineshape and from receiver structure removal.  $\eta$  is initially set to 1 and must be determined

later on with the injection of software synthetic axions. Each of the values calculated from the raw data in equation 5.1 are discussed below.

**B:** The magnetic field in the vicinity of the Sidecar experiment was determined with a hall probe located inches away from the cavity and with a simulation performed by an undergraduate student, Kunal Patel. The value of the simulation at the location of the hall probe, and the measurement of the the hall probe agreed to 5%. The value of the magnetic field in the Sidecar cavity was determined by believing the hall probe measurement and then using the simulation to extrapolated the value of the field in the cavity from the position of the hall probe. Figure 5.3 shows the intensity of the magnetic field in the vertical direction. The value of the field has been put in terms of the fraction of hall probe and main cavity fields. The RMS value of the field in terms of the current through the magnet coils is  $B = (0.0156T/A)I\hat{z}$ . Therefore, if the coils were ramped to 250 A, the field at the Sidecar cavity would be 3.9 T. During this run, for 8 weeks, the value of the Sidecar field was 0.78T. At the end of the run, the field was increased to 2.55T.

$C_{010}$ : Form factors were determined by a simulation performed by James V Sloan. As the experiment tuned, a lookup table provided unique  $TM_{010}$  form factor values for every rod position in the cavity. Erring on the side of caution, these values where reduced by 5%. During this run, form factors ranged from 0.44 to 0.61.

$S_{11}$ : In past ADMX data sets, the antenna was either critically coupled or assumed to be critically coupled. In this case, exactly half of the axion power escapes through the antenna while the other half dies in the walls of the cavity. However, in this maiden Sidecar run, a set screw holding the antenna loosened, causing it to remain over-coupled for the duration of the run. For this data set, it was therefore unacceptable to assume

# Sidecar Magnetic Field

from Computer Simulations and Hall Probe Measurements

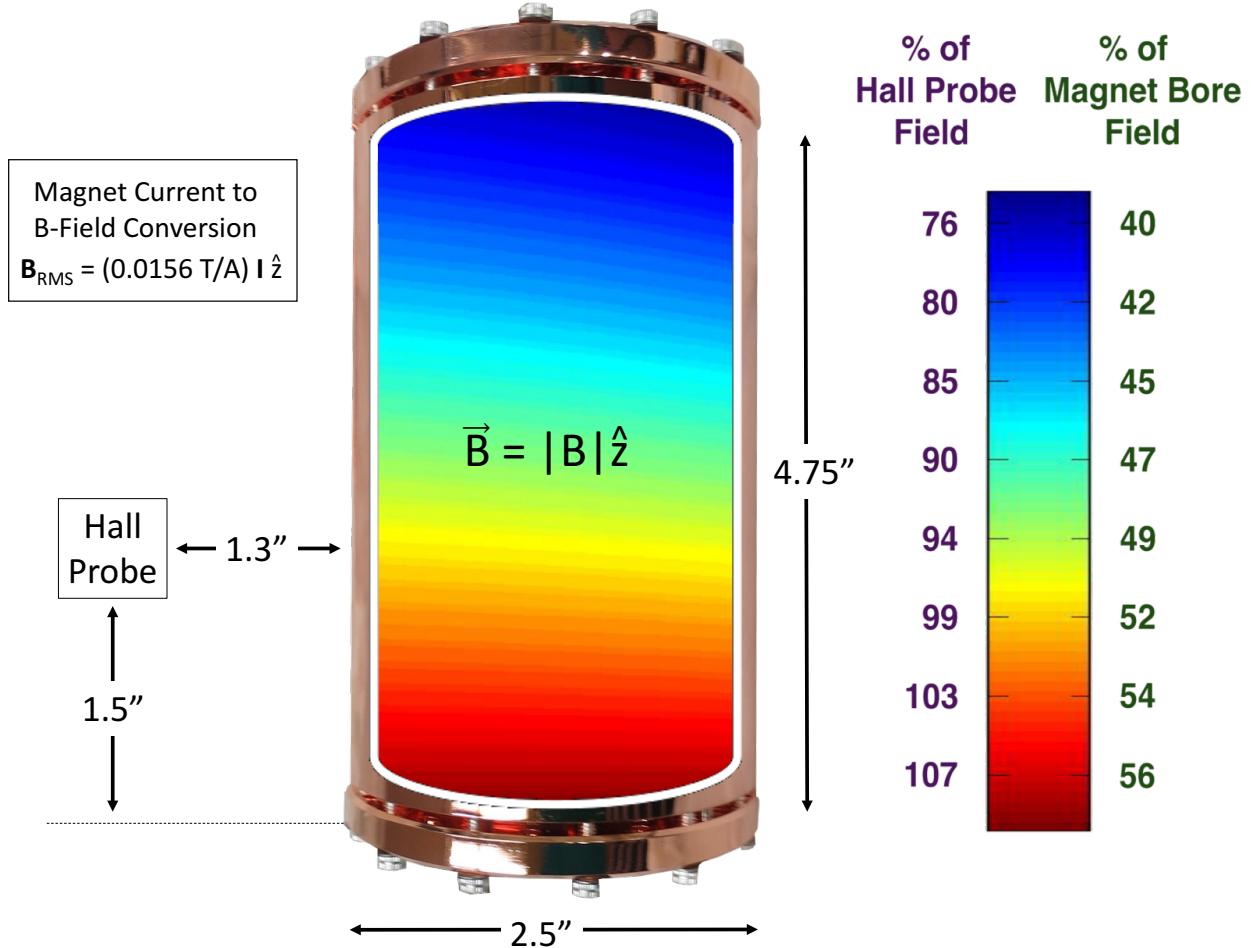


Figure 5.3: Computer simulated  $B\hat{z}$  in the vicinity of the Sidecar cavity.  $B\hat{z}$  is given in terms of a percentage of the main cavity field and of a measurement made by the hall probe located inches away from the cavity.

a critical coupling. The full expression for the fraction of axion power out of the cavity [57] can be obtained from

$$S_{11} = \frac{Q_{antenna} - Q_0}{Q_{antenna} + Q_0} = \pm \sqrt{power\ reflected\ off\ antenna} \quad (5.2)$$

where  $S_{11}$ , the reflection coefficient is zero if the antenna is critically coupled, positive if under-coupled, negative if over-coupled and  $Q_0$  and  $Q_{antenna}$  are functions of the mode frequency ( $\omega$ ), total stored energy (U) and average power lost ( $P_i$ )

$$Q_0 = \omega U / P_{walls} \quad (5.3) \quad Q_{antenna} = \omega U / P_{antenna} \quad (5.4)$$

where  $P_{walls}$  and  $P_{antenna}$  are the time-average powers lost in the walls of the cavity and out of the antenna. Together, these losses define the loaded Q ( $Q_L$ ) of the cavity

$$\frac{1}{Q_L} = \frac{1}{Q_{antenna}} + \frac{1}{Q_0} \quad (5.5)$$

where  $\frac{1}{Q_L}$  is the total fractional power lost per cycle. The fraction of axion power out of the cavity, denoted by  $\kappa$ , is therefore a function of  $S_{11}$

$$\kappa = \frac{1/Q_{antenna}}{1/Q_L} = \frac{1}{2} \frac{1 - 2S_{11}}{1 - S_{11}} \quad (5.6)$$

where the factor of 1/2 was already included in the equation 5.1 pre-factor. Calculating the fraction of axion power out of the cavity can then be achieved with an  $S_{11}$  reflection measurement. If the dip in the reflection trough is -20dB and the antenna is known to be under-coupled, 1% of the power incident on the antenna on resonance is reflected and according to equation 5.2  $S_{11} = + 0.1$ . Then by equation 5.6, the fraction of developed power out of the cavity 0.44. Notice that at critical coupling ( $S_{11} = 0$ ),  $\kappa$  reduces to 1/2.

During this run,  $Q_L$  was degraded by the over-coupled antenna. While the low Q would have caused less axion power to develop, the fraction of axion power out of the

antenna would have been  $> 1/2$ . Because the antenna was over-coupled, resulting  $S_{11}$  measurements were often asymmetric and not easily fit by an inverse Lorentzian. For this analysis, the scan by scan antenna coupling was not assigned by direct reflection measurements but instead by using the relation

$$S_{11} = \frac{2Q_L}{Q_0} - 1 \quad (5.7)$$

easily obtained from equations 5.2 and 5.5. Together  $Q_L$  and  $Q_0$  were used to calculate the value of  $\kappa$  in equation 5.6. During the data run,  $\kappa$  ranged from 0.6 to 0.8.

$T_{sys}$ : The system temperature is the sum of the physical and noise temperatures and is necessary in a SNR calculation. The physical temperature of the cavity was measured by two sensors thermally sunk to the top and the bottom of the cavity. One of these sensors behaved erratically and was ignored in the analysis. The other is discussed in the quality control section 5.4. Ideally, the noise temperature of the receiver-chain is easily measured with the procedure discussed in section 3.4.1. Unfortunately, the RF-connection to the hot/cold load appeared to have been broken at some point in the cool-down, which meant that noise-figures from amplifier data sheets needed to be believed at face value and not measured in situ. The system temperature was computed by adding up each of the terms in equation 2.26. This calculation is shown in table 5.3.2.

At the end of this stage in the analysis, every bin in a digitized power spectrum can be assigned a “before gain” expected axion power. The next step in the single scan analysis is to turn the high gain power coming out of the digitizer back into a power that can be compared directly with the power calculated above.

Component	Gain (dB)	Total Gain (dB)	$T_N$ (K)	$T_{sys}$ (K)	Notes
Cavity	0	0	0.4	0.4	notes
Cable	-0.5	-0.5	0.0435	0.4435	notes
Circulator	-0.4	-0.9	0.0308	0.4781	notes
Cable	-0.5	-1.4	0.4350	1.0132	notes
Amp Input Reflection	-0.14	-1.5396	0.1265	1.1879	notes
LNF Amp	35	33.4604	4	6.8898	notes
Cable	-1	32.4604	15.837	6.8969	notes
Post Amp 1	21	53.4604	125.67	6.9682	notes
Post Amp 2	21	74.4604	125.67	6.9688	notes
Cable	-20	54.4604	76.230	6.9688	notes

Table 5.1: Gain and noise temperature values for each component in the receiver-chain relevant to the overall system temperature. The  $T_{sys}$  column shows that components after the 2<sup>nd</sup> post amplifier don't contribute to the final system temperature.

### 5.3.3 Measured Power

If one could see a power spectrum of the blackbody power in the cavity (without time averaging), the spectrum would have fluctuations about some baseline where the mean of the baseline and the standard deviation of the fluctuations are both equal to  $BkT_{physical}$ . The power out of the cavity then propagates through the entire receiver chain and receives gain from many stages of amplifiers. At this point the baseline power increases and the fluctuations in the power are exaggerated. In addition to a huge boost in gain, each component of the chain can add some shape or distortion to the signal. The goal of this stage is to turn the output of the digitizer back into a “power in the cavity”. The first step in this process involves dividing the raw spectra by the response of the receiver. Figure 5.4 shows

this process in three stages

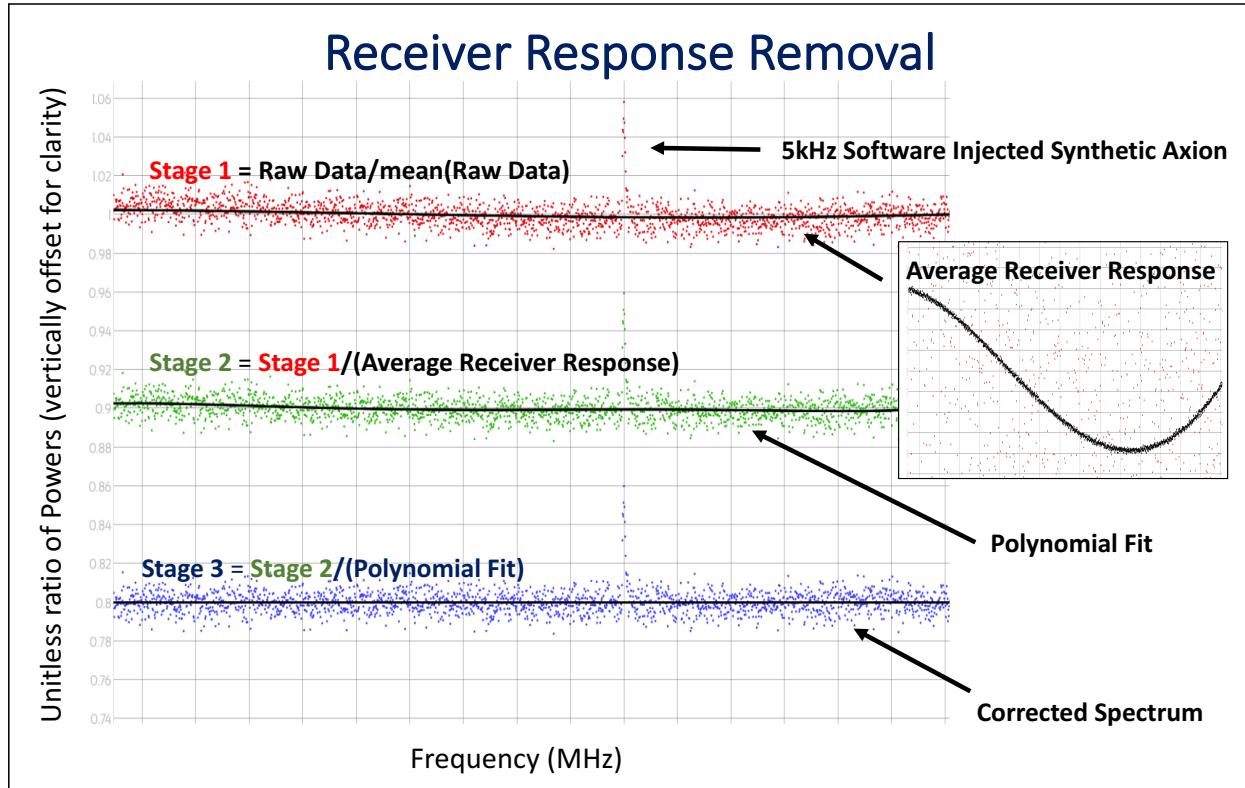


Figure 5.4: Receiver response removal process. Stage 1: raw data is divided by the mean such that the resulting fluctuations are unitless with a mean of 1. Stage 2: stage 1 data is divided by an average of the entire data set. This average characterizes a time independent response for the receiver. Stage 3: stage 2 data is divided by a low order polynomial. The resulting conditioned spectrum should be flat, noise-like and still contain axion-like signals.

**Stage 1: Divide by mean:** This step is not necessary to the process but allows for comparison between the raw data and the corrected data. After this stage, fluctuations are unitless and have a mean of 1.

**Stage 2: Divide by Receiver Response:** In this step, the goal is to remove any time independent structure from the spectrum. If one has not been careful and has only

used data over a very narrow frequency range, the average response could characterize a true axion signal and divide it out. If the cavity has tuned over a very wide range, this step mostly removes post-mixer structure since pre-mixer structure would be washed out in the average. The response used on this data set was the average of all of the data that passed the quality control stage.

**Stage 3: Divide By Low-Order Polynomial Fit:** This step aims to remove any time dependent (or any parameter dependent) change in the receiver shape. Some have used a more complicated multi-parameter fit that is meant to characterize the physics of the receiver chain [34] while others have opted for a more generic polynomial fit motivated by a similar number of degrees of freedom [45, 46]. This analysis used a 6-order polynomial fit because I found it to be effective while still being very insensitive to 5kHz wide axion-like injected signals. For large amplitude injected signals, approximately 10% of the original signal was found to be removed from the resulting corrected spectrum. Realistic, axion-like signals below the noise are discussed in section 5.5.6.

After the receiver response removal, the spectrum should be flat with a mean of one. The distribution will be Gaussian if the assumption of thermal noise being the only noise source is true. Figure 5.5 shows a histogram of the entire data set after having undergone these corrections. A gaussian fit shows that the resulting conditioned data is noise-like. At this point, the spectrum is multiplied by  $BkT_{sys}$  to convert it back into a power in the cavity except that now, the temperature in the cavity is effectively larger because of the added  $T_N$  contribution. The fluctuations in the power are now less than the baseline power by a factor of  $1/\sqrt{Bt}$ . In the co-adding stage, when it comes time to add many spectra together, the fluctuations in the noise will fall like  $\sim 1/\sqrt{N}$  but an overall offset will not. For this reason, the baseline is subtracted from the spectra to give power fluctuations about a mean of zero. The resulting noise of each spectrum is  $\sigma BkT_{sys}$  where  $\sigma$  is the standard deviation of the

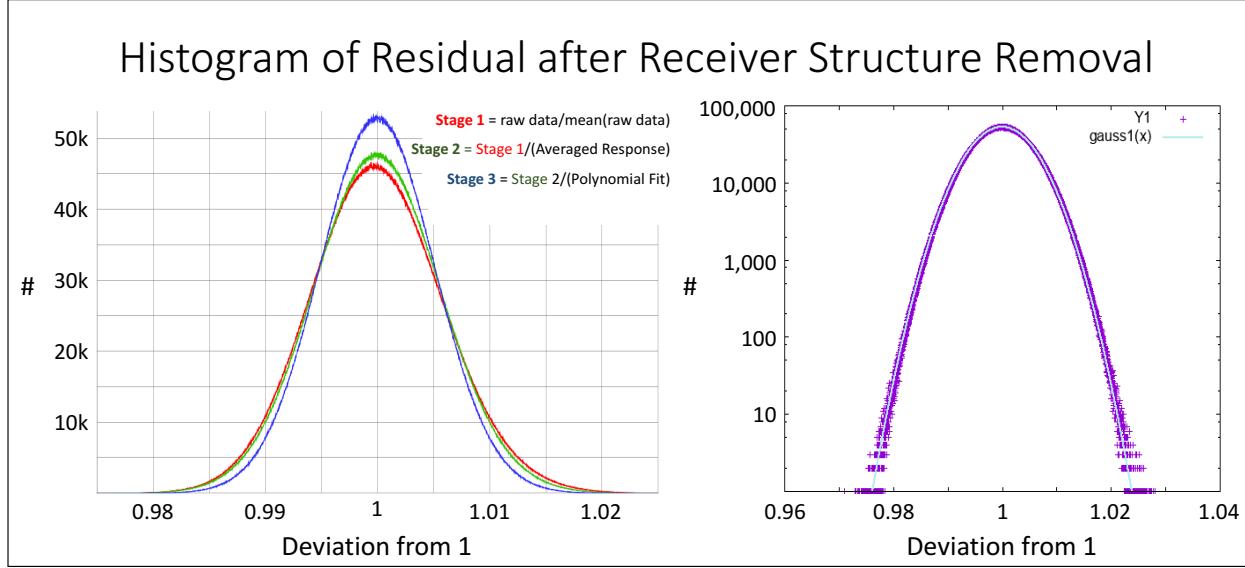


Figure 5.5: Left: a histogram of the entire data set for each of the stages of correction in figure 5.4. Right: histogram of the stage 3 residual after receiver structure removal on a log plot along with a Gaussian fit. Non-thermal noise, such as RF interference would be visible as deviations from Gaussian shape.

fluctuations and plays the role of the  $\sqrt{N}$  noise decrease after integrating for 100 seconds. At the end of this stage in the single scan analysis, corrected spectra are called deltas ( $\delta$ ) because they represent real power deviations about zero.

#### 5.3.4 Weighting

In preparation for the co-adding stage, the expected axion power ( $p$ ) and measured power fluctuations ( $\delta$ ) must be weighted such that high field, low noise data is not summed on equal footing with low field, high noise data. A weight ( $w$ ) is defined as

$$w = \frac{p}{\sigma^2} \quad (5.8)$$

High quality scans therefore produce high weights. The measured and expected powers

are scaled by the weight and denoted

$$p^w = wp = \frac{p^2}{\sigma^2} \quad (5.9)$$

$$\delta^w = w\delta = \frac{p\delta}{\sigma^2} \quad (5.10)$$

A simple array of numbers ( $n$ ) keeps track of the number of spectra that contribute to a given frequency bin. Since no data has been summed in the single scan analysis and the resolution has not been altered,  $n$  is just an array of ones at this point. At the end of the single scan analysis,  $n$ ,  $w$ ,  $p^w$  and  $\delta^w$  are passed to the next stage in the analysis.

## 5.4 Quality Control

In addition to receiving the quantities prepared for the co-adding stage, the quality control function is passed the values of many parameters associated with a given scan such as:  $Q_L$ , frequency,  $B$ ,  $C_{010}$ ,  $S_{11}$ ,  $T_{phys}$ ,  $T_{noise}$ , average measured digitizer power etc. The goal of this stage is to monitor the health of the data and cut poor quality scans.

### 5.4.1 Automated Cuts

Automated cuts were made on malformed data or where data parameters were outside of their typical values. Table 5.4.1 summarizes these cuts.

Parameter	Lower Threshold	Upper Threshold
$B$ Field	0.1T	8T
Frequency	5GHz	6GHz
$Q_L$	1,000	20,000
$T_{phys}$	10mK	2K

The goal of this stage is to remove nonsensical measurements that are clearly bad, not remove lower quality data. Lower quality data has already been weighted such that its contribution to the final spectrum will be negligible.

#### 5.4.2 *Manual Cuts*

Not all signs of bad data can be detected in an automated way. Two examples of this are instances when the experiment has gradually started following the wrong mode and when temperature spikes are observed. Both of these cases require the context of more data. As the analysis slowly churns through the data, scan by scan, a log of every parameter is stored and displayed on the web-interface so that a human eye can look for trends or even correlations between parameters. Each parameter is plotted versus scan number and if bad data is identified, their scan numbers can be read directly off the plot and flagged. Figure 5.6 shows mode frequency versus scan number for the entire data set and identifies the mode crossings that were flagged.

A human eye was also required to examine temperature spikes reported by the Sidecar temperature sensor. The lower section of figure 5.7 shows the temperature of the Sidecar experiment as reported by the cavity sensor versus scan number. The upper left hand side of the figure shows a zoomed in portion of the plot. Displayed are clearly spurious temperature spikes that die instantly as well as spikes from real heating that take time to decay. In either case, the data should not be cut. However, spurious temperature spikes make the data appear less sensitive than it actually is. Automating the differentiation between the two is difficult to do in the context of a single scan.

A thermal analogue of a voltage divider circuit was used to model and then predict the Sidecar temperature as a function of the 1K plate and dilution refrigerator mixing chamber temperatures. The results of this exercise is shown in the upper right hand portion of the figure. The model (barely visible right under the Sidecar temperature sensor) tended to

## Identifying Mode-Crossings

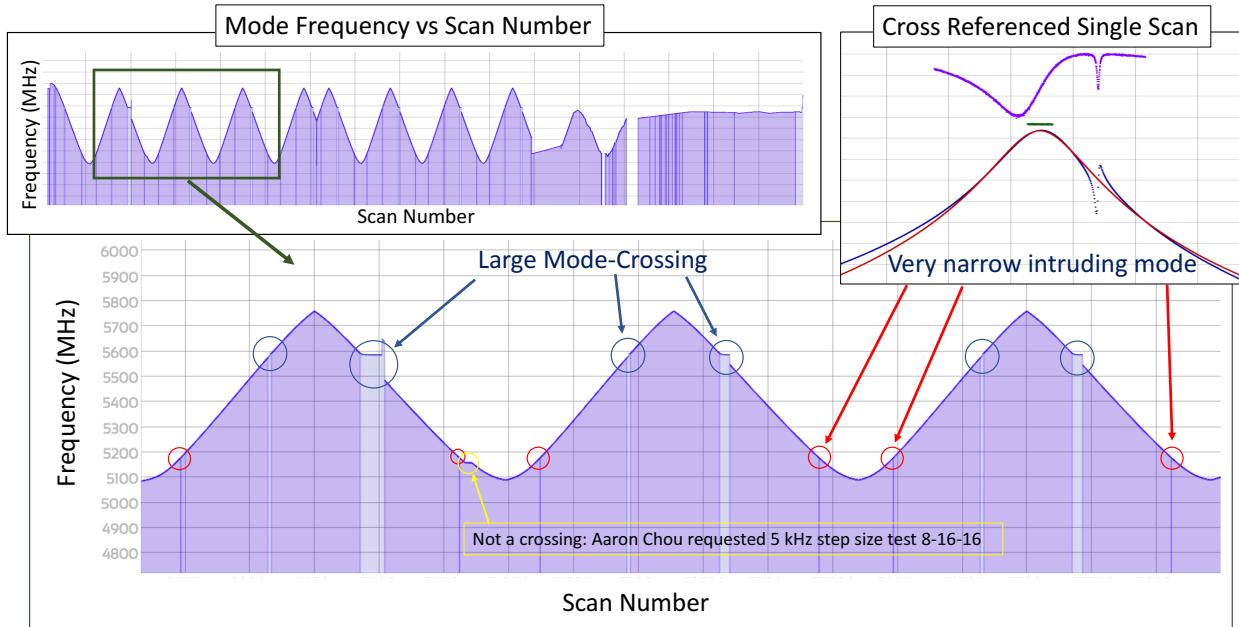


Figure 5.6: Cavity frequency versus scan number for the entire data set. Mode crossings are identified by eye and the data associated with those scan numbers is manually cut.

slightly overshoot real heating while of course, being completely uncorrelated with spurious peaks. The voltage divider model was originally proposed as a test to determine if the 200mK offset from the mixing chamber was caused by real heating or by a systematic miscalibration. The test was inconclusive but the prediction did allow spurious peaks to be identified and set equal to the model.

The data set, in its entirety, is comprised of 33,333 spectra. 25,084 survived the quality control stage and 8,249 (24.7%) were cut.

## Differentiating Real and Spurious Temperature Spikes

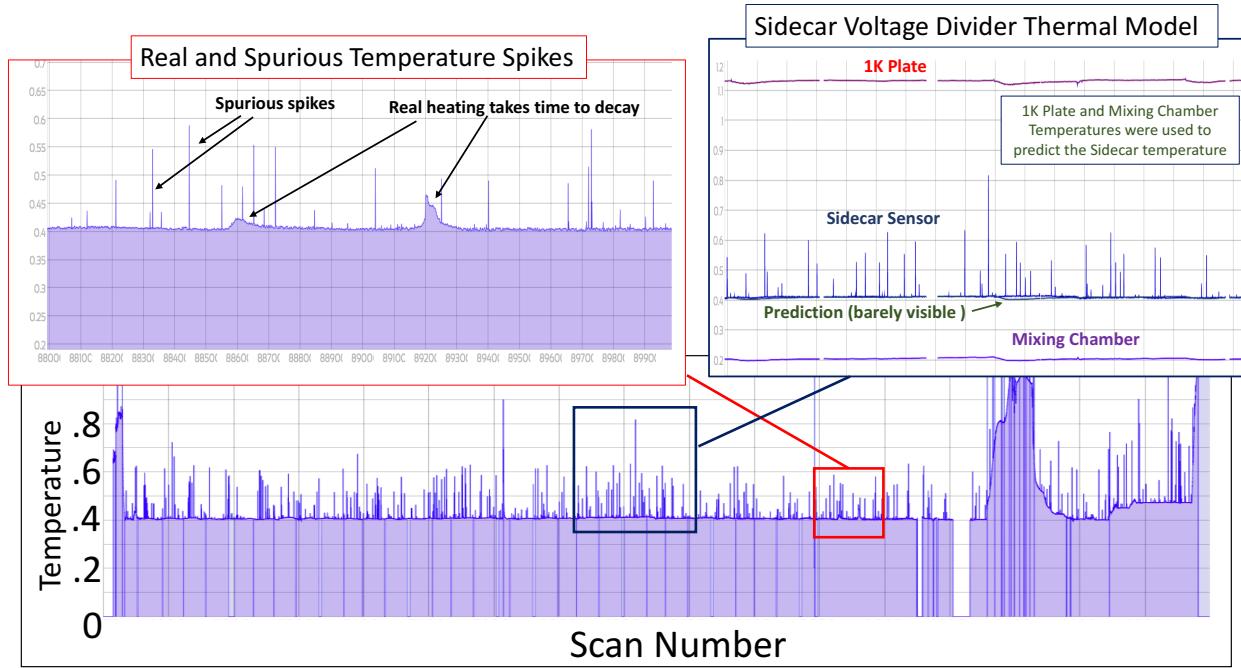


Figure 5.7: Cavity temperature sensor values versus scan number for the entire data set. Mixing chamber and 1K plate temperatures were used to predict the temperature of the Sidecar cavity. This model was used to differentiate between spurious and real temperature spikes. Spurious values were set equal to the temperature model prediction.

## 5.5 Grand Analysis

### 5.5.1 Co-adding Spectra

Finally quantities  $n_{si}$ ,  $w_{si}$ ,  $p_{si}^w$  and  $\delta_{si}^w$  of like frequency indices (i) were summed over scan number indices (s)

$$N_i = \sum_s n_{si} \quad (5.11) \quad W_i = \sum_s w_{si} \quad (5.12) \quad P_i^w = \sum_s p_{si}^w \quad (5.13) \quad \Delta_i^w = \sum_s \delta_{si}^w \quad (5.14)$$

to produce 4 kinds of grand spectra representing the entire data set. Figure 5.8 illustrates the

co-adding process that takes place separately for  $N$ ,  $W$ ,  $P^w$  and  $\Delta^w$ . Here,  $N$  is the number of spectra, defined by the resolution of the analysis that fit into a given  $i$ th frequency bin. If the frequency resolution is set to be very coarse, the number of  $i$  indices in the  $N$  array is small but a given value of  $N_i$  will be large. For reasons explained in the following section, the original 381.5 Hz resolution of the raw digitized spectra was initially preserved in this stage. At this resolution, over a tuning range of  $\sim 700$  MHz, the resulting grand arrays would nearly be  $2 \times 10^6$  bins long, which would slow the analysis and prevent the data from loading on the web-interface. For this final analysis, scan numbers were re-sorted by mode frequency and the full dataset was cut into 30 chunks that were analyzed separately.

## Co-Adding Scheme

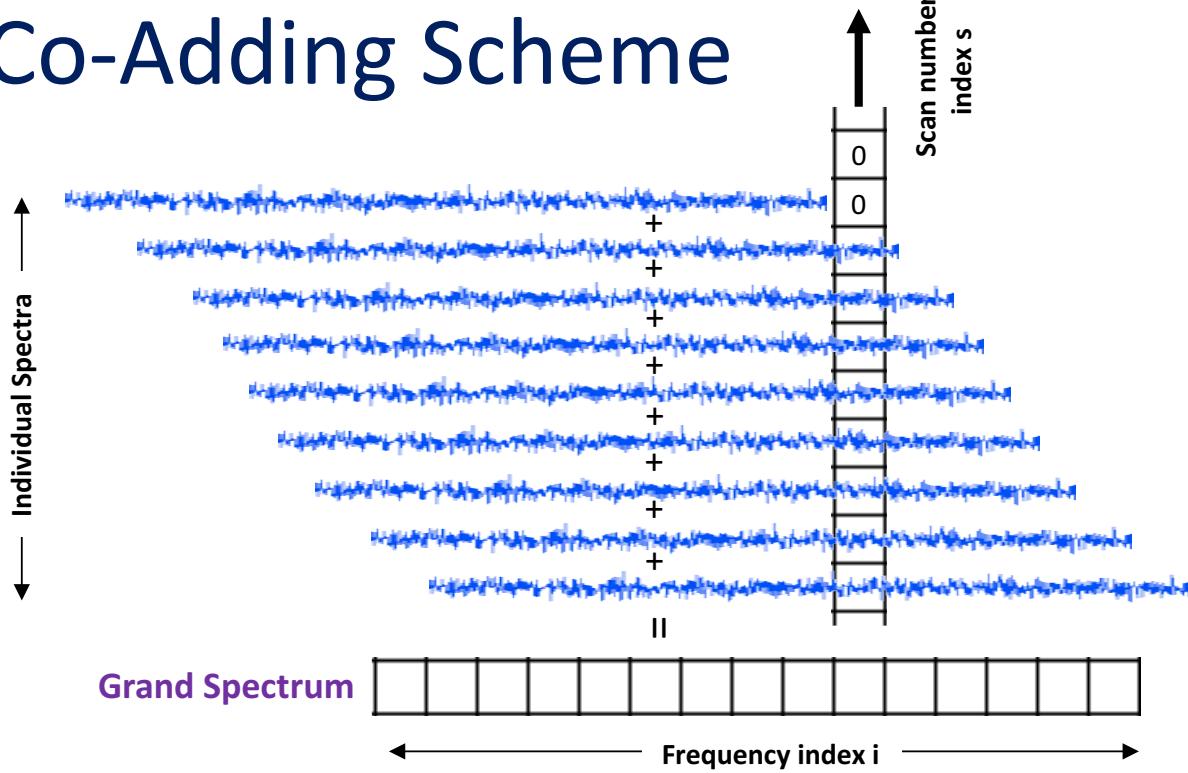


Figure 5.8: All of the spectra of like frequencies were added in quadrature, resulting in a single grand spectrum.

### 5.5.2 Correlating bins

This analysis was a search for 5kHz wide axions and was done by looking for a surplus in power in each bin of the grand power spectrum. However, the expected axion power  $p$  calculated in the single scan analysis was reduced by a factor of  $\sim 13$  since an axion of this width would approximately span 13 of the 381 Hz bins. To regain this lost sensitivity, the data was re-binned to a courser resolution. The straight forward way to do this is to co-add the data in every 13 bins to a single bin that is 13 times as wide, resulting in a course grand array that has  $1/13$  the number of bins as the higher resolution array. This naive approach is shown at the top of figure 5.9. The downside to this scheme is that the axion signal may land between two of the course resolution bins causing it to evade detection.

This undesirable side effect is avoided by correlating bins (shown at the bottom of figure 5.9). Here, the contents of the  $i$ th bin in the correlated array is the sum of the  $i-6$  through the  $i + 6$  bins of the original array. This scheme guarantees that at least one of the resulting bins has most of the axion signal. The downside to this choice is that the resulting array has the same number of bins as the original array and is therefore computationally intensive. The other downside is that the correlated bin search is non-statistical and the final sensitivity of the limit must be calibrated with software injected synthetic axion signals. This method of correlating neighboring bins is effectively equivalent to convolving the data set with a box filter and could yield better limits with the use of a more realistic axion filter shape. In this analysis, the generic box filter was used. At some point in the future, this data may be re-analyzed to produce slightly better limits.

### 5.5.3 Deriving grand spectra from weighted sums

The data was originally weighted such that it could be summed and convolved in quadrature. At this stage the data can be unweighted and used to derive several different grand spectra such as: RMS noise, measured power, SNR etc. Table 5.5.3 summarizes the unweighting

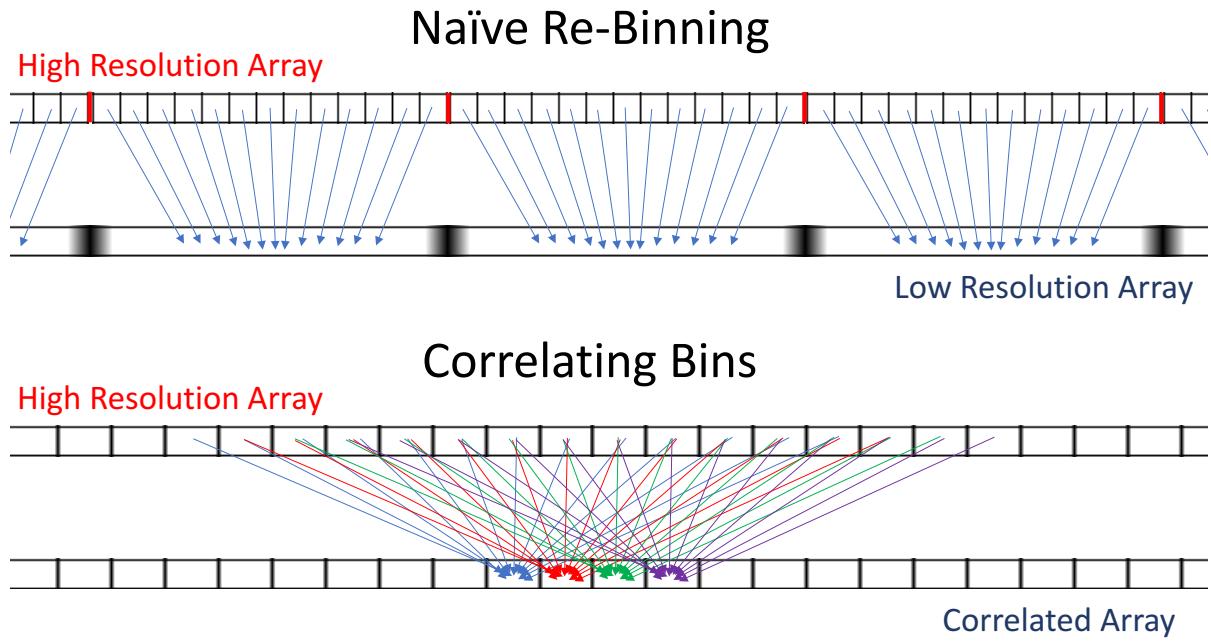


Figure 5.9: Top: data is re-binned to a courser resolution. This naive procedure could cause an axion signal to be split between two bins, resulting in a missed opportunity for a detection. Bottom: the final re-binned array is essentially a convolution of the initial array such that at least one of the bins contains almost all of the axion power.

process.

Grand Spectra Calculated	Symbol	Derivation	Notes
Number of Spectra	N	N	Number of spectra contributing to ith bin
RMS Noise	$\sigma$	$\frac{\sqrt{P^w}}{W}$	Watts/5kHz
SNR	$P/\sigma$	$\sqrt{P^w}$	Computed for $g_{a\gamma\gamma} = 5 \times 10^{-12} \text{GeV}^{-1}$
Expected Power	P	$\frac{P^w}{W}$	Computed for $g_{a\gamma\gamma} = 5 \times 10^{-12} \text{GeV}^{-1}$
Candidate Search	$\Delta/\sigma$	$\frac{\Delta^w}{\sqrt{P^w}}$	Measured power in units of noise
Measured Power	$\Delta$	$\frac{\Delta^w}{W}$	Deviation from mean power in watts

Figure 5.10 shows each of these grand spectra. Here the expected power and SNR and calculated for an axion-to-two-phonton coupling of  $g_{a\gamma\gamma} = 5 \times 10^{-12} \text{GeV}^{-1}$ .

#### 5.5.4 Calculating Excluded Power

Having defined the “Candidate Search” spectrum as the fluctuations in the measured power in units of the noise ( $\Delta/\sigma$ ), an upper limit must be set on the power sensitivity of the data. The confidence level was determined by finding the probability of a signal ( $p_{excluded}/\sigma$ ) passing the measured fluctuations  $\Delta/\sigma$  defined by

$$c.l.(p_{excluded}) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{p_{excluded}/\sigma - \Delta/\sigma} e^{-x^2/2} dx = \frac{1}{2} (1 + \text{erf}(\frac{p_{excluded}/\sigma - \Delta/\sigma}{\sqrt{2}})) \quad (5.15)$$

so if  $\Delta/\sigma = 3$ , a signal with power 4.6 would be detected with 95% confidence.  $p_{excluded}$  was solved numerically for all measured values of  $\Delta/\sigma$  at a 95% confidence. The excluded power is plotted (green) above the measured fluctuations (blue) in figure 5.10. In other words, powers above this line have been excluded with a 95% confidence.

### 5.5.5 Setting a limit

Finally a limit can be set with

$$\frac{g_{a\gamma\gamma}^2(\text{target})}{g_{a\gamma\gamma}^2(\text{C.L.95\%})} = \frac{P_{\text{target}}}{P_{\text{C.L.95\%}}} \rightarrow g_{a\gamma\gamma}(\text{C.L.95\%}) = 5 \times 10^{-12} \text{GeV}^{-1} \sqrt{\frac{P_{\text{C.L.95\%}}}{P}} \quad (5.16)$$

where  $g_{a\gamma\gamma}(\text{target})$  is the target coupling used in the single scan analysis ( $5 \times 10^{-12} \text{GeV}^{-1}$ ),  $P_{\text{target}}$  is the power (P) calculated from that coupling,  $P_{\text{C.L.95\%}}$  is the power excluded at 95% confidence and  $g_{a\gamma\gamma}(\text{C.L.95\%})$  is the new limit to be set.

### 5.5.6 Confirming limit with software injected synthetic axions

Up until this point,  $\eta$ , the loss term which accounts for axion power lost in the receiver structure removal and from the tail in the axion line shape evading the 5kHz bin width has been ignored. The analysis was intentionally run with  $\eta = 1$  to produce limits which were overly optimistic. Using the output of the analysis in a feedback loop, the analysis was then repeated but this time, software synthetic axions were injected according to the  $\eta = 1$  limit on  $g_{a\gamma\gamma}$ . Given that the data had already been cut into 30 frequency chunks (for computational reasons), within a given chunk, a random number generator picked the frequency locations of 20 synthetic axions. Knowing that the first prediction on  $\eta$  would create overly optimistic limits and therefore unrealistically low powered axions, most of the signals were not detected. The value of  $\eta$  was then reduced and the analysis was repeated over and over again until 95% of the synthetic axions were detected.  $\eta$  was measured to be 0.649, meaning that 35% of the power was removed in the tail of the peak and in the fitting stage. Figure 5.11 shows the software signal shape cloned from a real hardware source, as well as several examples of detected and undetected signal injections. Using this value of

eta, the 95% confidence limit correctly accounted for systematics introduced by background removal. During this stage of the analysis, several hundred synthetic signals were injected.

## 5.6 Results

Using the method presented above, taking the width of the axion to be 5kHz and the local dark matter density to be .45 GeV/cc, new limits on  $g_{a\gamma\gamma}$  were set. The results of the analysis are shown in Figure 5.12. The Sidecar limits are shown next to the limits set by the CAST experiment, the Rochester-Brookhaven-Fermilab (RBF) collaboration [51], the University of Florida [50] and by the narrow bandwidth Yale axion search [58]. On this low magnetic field commissioning run, the results of the data are not competitive with the limits set by Yale. However, with a Josephson Parametric Amplifier and full field, this wide  $3\mu\text{eV}$  span could be ruled out at  $10^{-14}\text{GeV}^{-1}$  and would be competitive, indeed, being both very sensitive and covering a broad mass range.

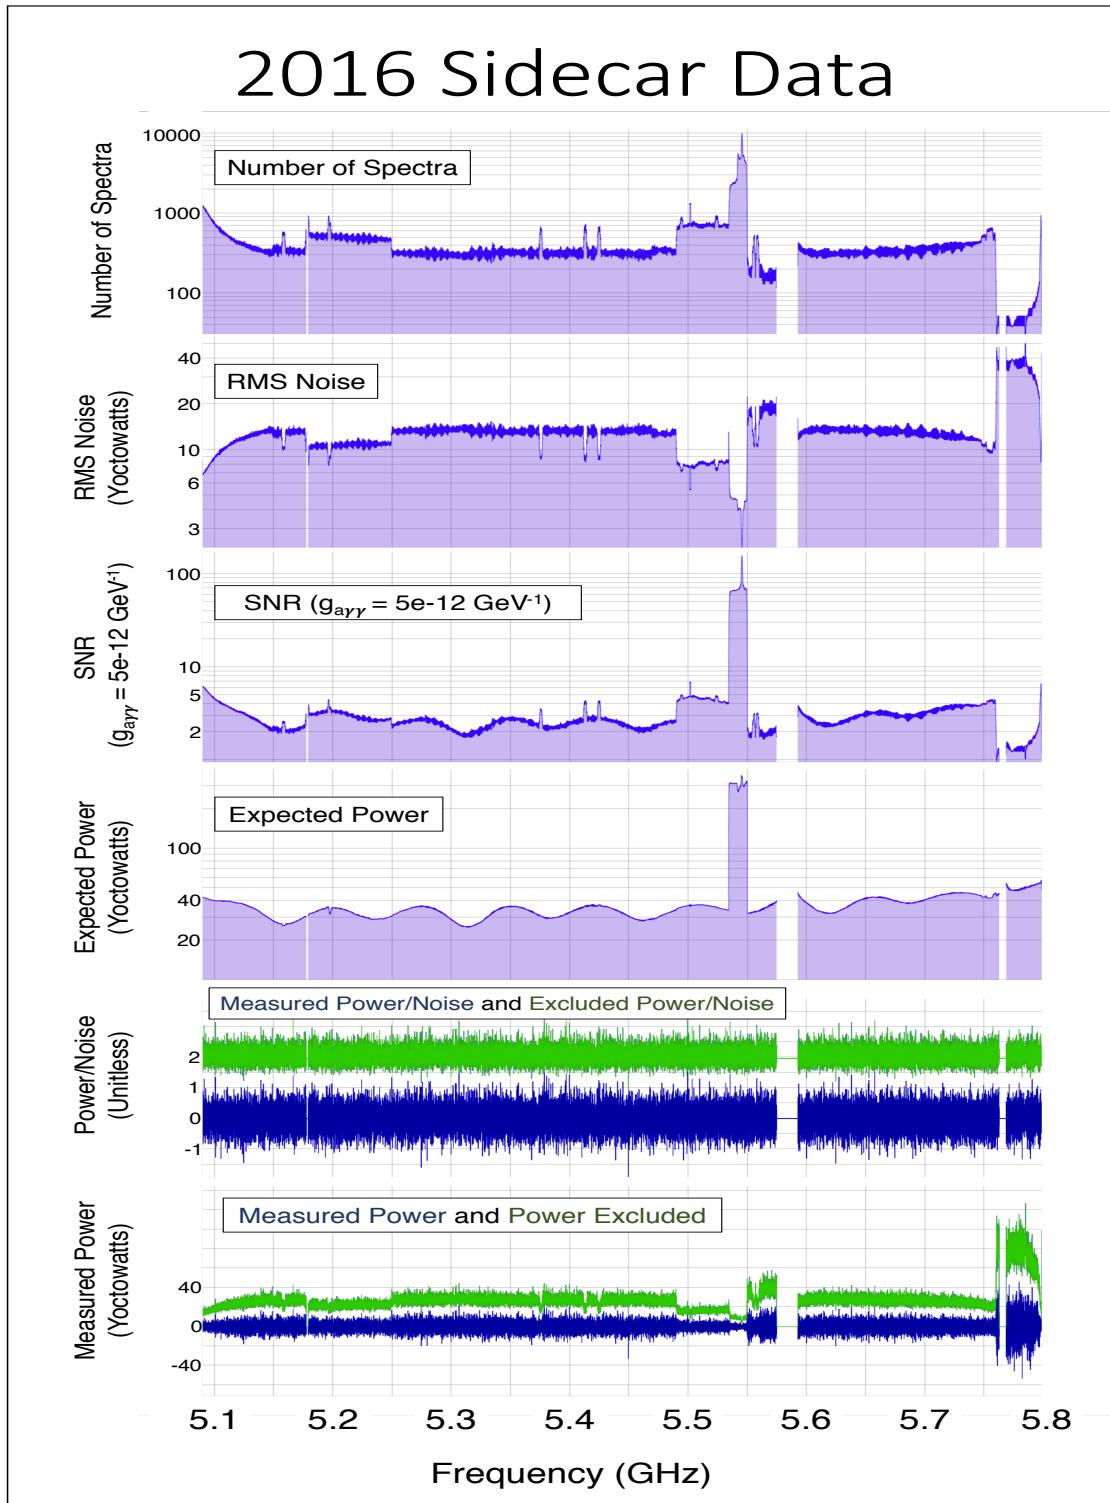


Figure 5.10: 6 final grand spectra for the entire data set: number of spectra taken, rms noise, SNR, expected axions signal, an axion candidate search and measured power fluctuations.

# Software Injected Signal Calibrations

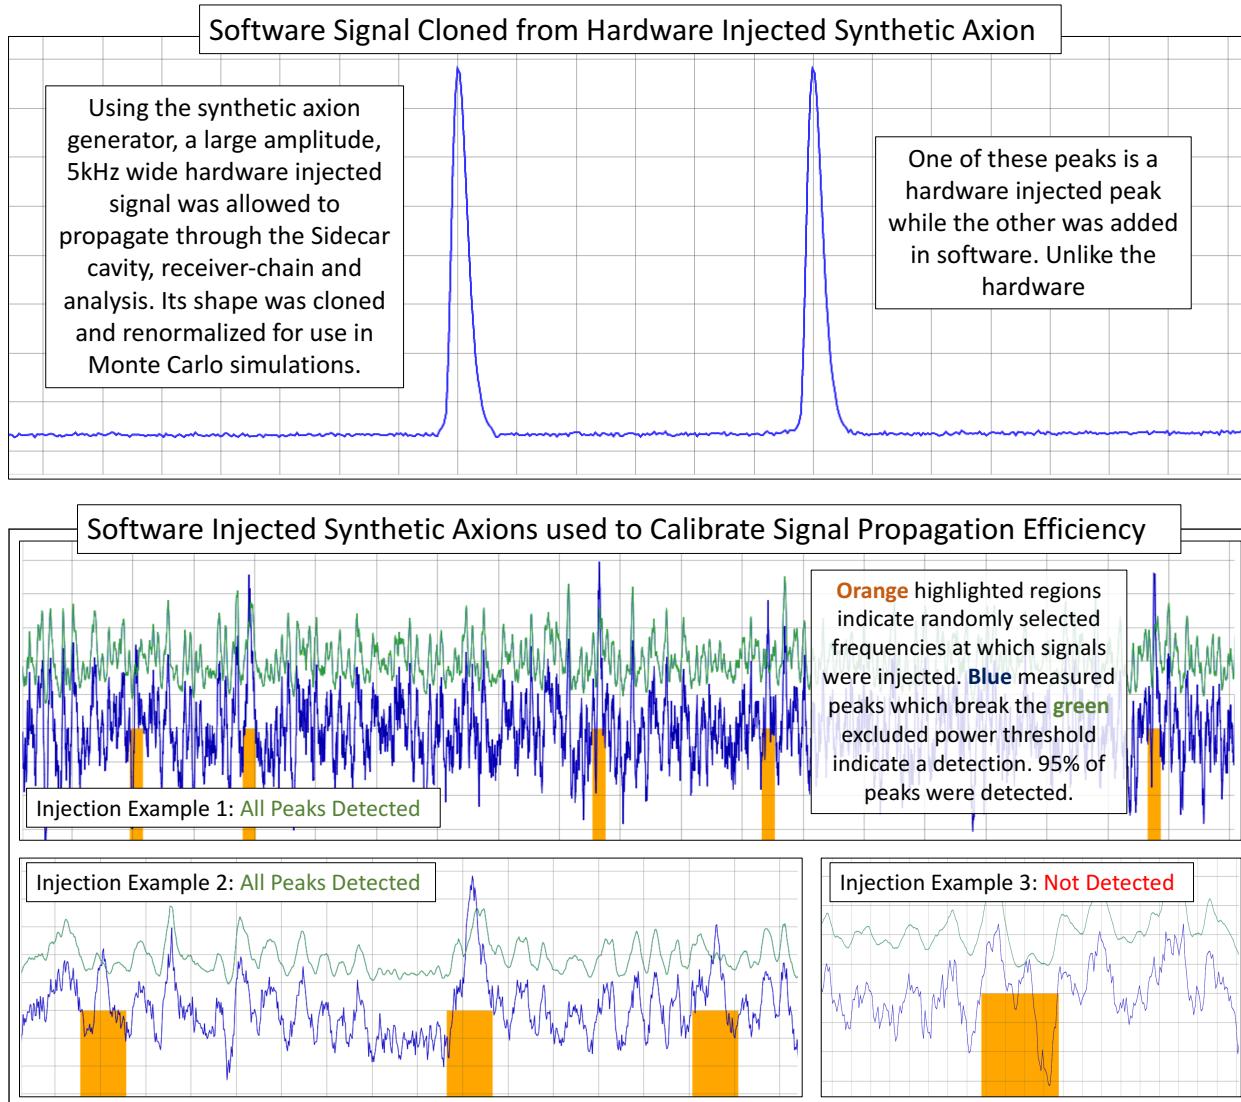


Figure 5.11: A software axion signal shape was cloned from hardware injected synthetic axions. Software signal injections were then used to measure the signal propagation efficiency  $\eta$ . For  $\eta = 0.649$ , 95% of the software injected axions were detected.

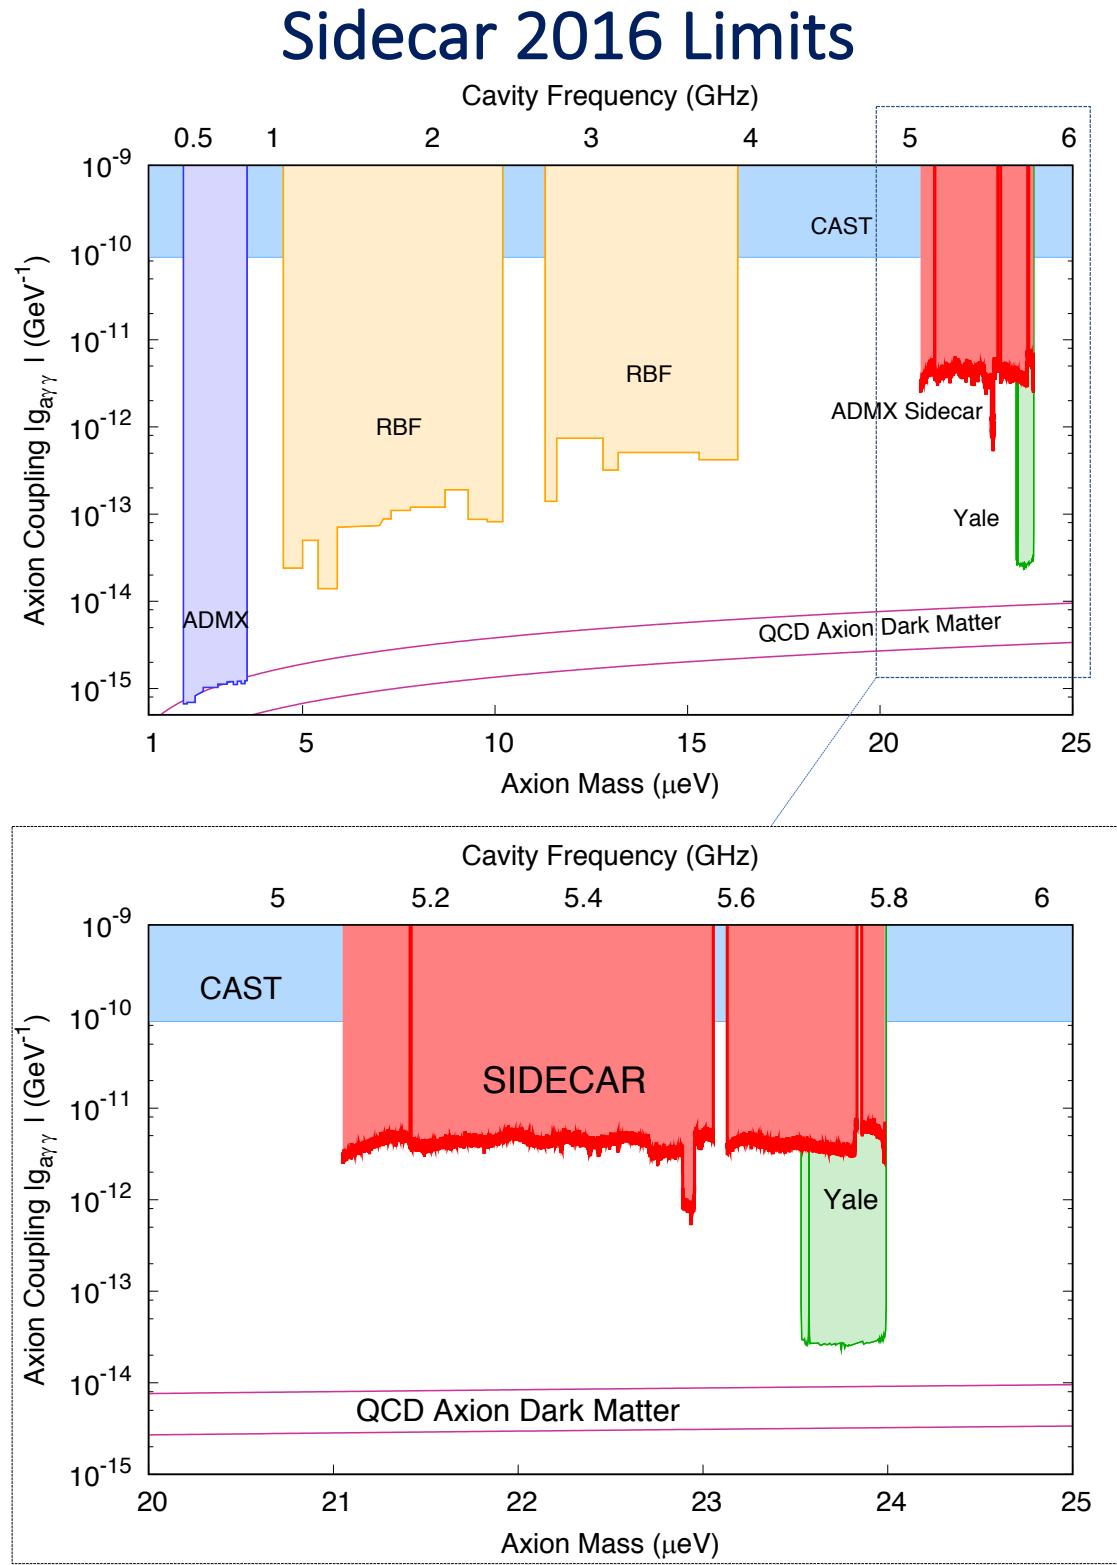


Figure 5.12: New limits on  $g_{a\gamma\gamma}$  set by the 2016 Sidecar data run over nearly a  $3\mu\text{eV}$  mass range.

## Chapter 6

### CONCLUSIONS

The axion is a well motivated hypothetical particle that could simultaneously solve the strong CP and dark matter problems. The tunable microwave cavity technique discussed in section 2.4 and throughout this work provides the most sensitive method of searching for weakly coupled axions in the Milky-Way halo. The Axion Dark Matter eXperiment is the most sensitive experiment searching for  $\mu\text{eV}$  axions using this technique. Armed with a working dilution refrigerator, ADMX is now able to scan orders of magnitude faster in frequency compared to earlier versions of the experiment. In the event that  $1\text{-}10\mu\text{eV}$  are ruled out, ADMX will develop new technology and techniques to preform axion searches at higher and lower masses. The Sidecar experiment is a pathfinder experiment created to test new piezoelectric tuning and coupling techniques which will help ADMX evolve as it encounters new challenges associated with higher frequencies. In 2016, a low magnetic field run, aimed towards commissioning a Janis dilution refrigerator took data for 8 weeks with the Sidecar experiment. Smooth piezoelectric tuning was successfully demonstrated by the experiment over a very wide 709 MHz ( $3\mu\text{eV}$ ) range. The experiment set new limits on the axion-to-two-photon coupling  $g_{a\gamma\gamma} < 6 \times 10^{-12} \text{GeV}^{-1}$  over a mass range of  $3\mu\text{eV}$  ( $\Delta f \sim 708$  MHz) from 21.05 - 23.98  $\mu\text{eV}$  for axions that compose 100% of dark matter. Over a narrower range of 22.89 - 22.95  $\mu\text{eV}$  ( $\sim 15$  MHz), a stricter limit of  $g_{a\gamma\gamma} < 10^{-12} \text{GeV}^{-1}$  was set.

The conclusion was that this piezoelectrically tuned/coupled experiment was a huge success though the heating from the motors will need to be considered carefully for future, more complicated multi-cavity systems. Also, the ability to piezoelectrically adjust antenna depth could be implemented now on the main experiment. With future data runs, limits

from the Sidecar experiment using this cavity could be taken down to  $16\mu\text{eV}$ , with the rod near the wall of the cavity, or extended all the way from  $26.4 - 30\mu\text{eV}$  by taking data on the  $TM_{020}$  mode. Masses higher than  $30\mu\text{eV}$  would be accessible with a new cavity. The modest sensitivity demonstrated in the 2016 run could be improved dramatically and with better amplifiers and a fully ramped magnet. Finally, if the main cavity were removed and replaced by a quantum limited Sidecar-like cavity that is 10 times longer, a competitive scan rate and sensitivity could be achieved. The scan rate for this hypothetical “Future Sidecar” is shown in figure 6.1. With a  $T_{sys}$  of  $150\text{mK}$ , this next-generation experiment could achieve DFSZ with a scan rate of 10 MHz a month. This scan rate could be further improved by filling the bore of the magnet with many such cavities. At the time of writing, the future multi-cavity system is in its planning stage.

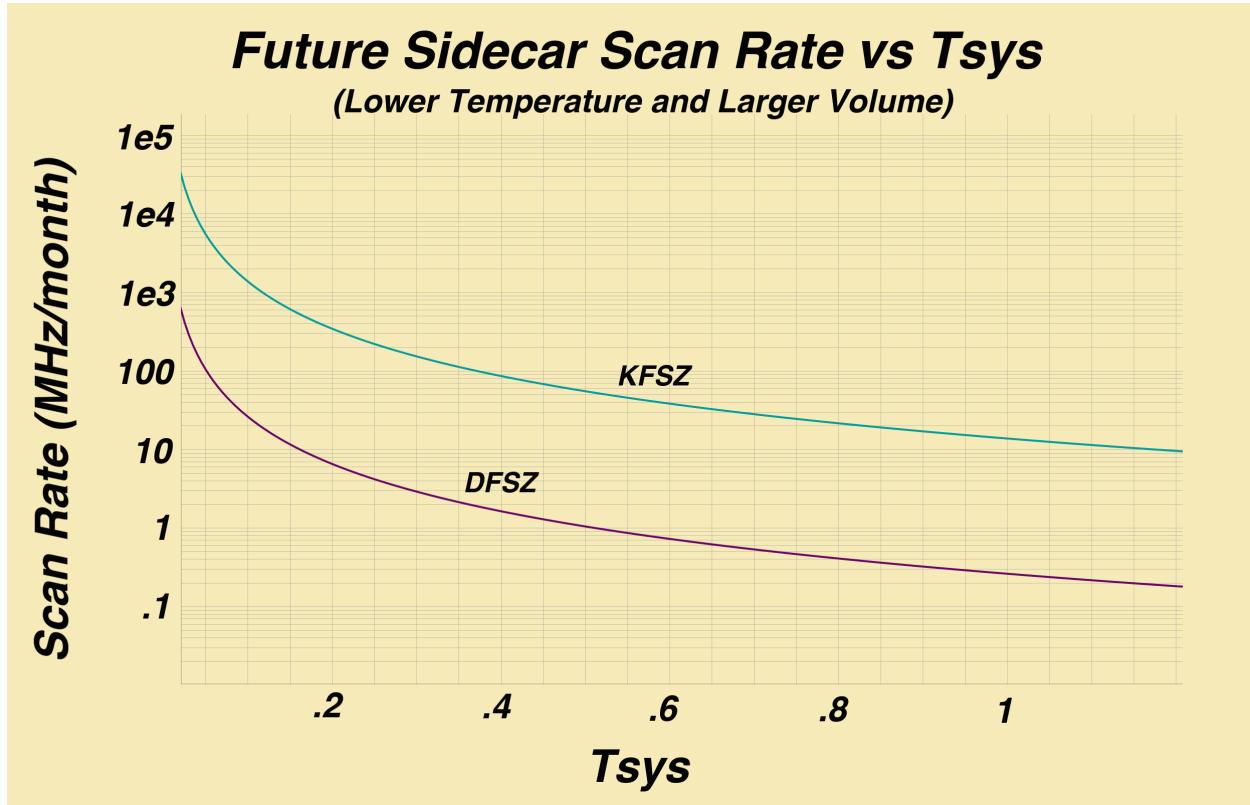


Figure 6.1: Future scan rate versus  $T_{sys}$  for a Sidecar-like experiment, 10 times as long with a JPA at full field.

## BIBLIOGRAPHY

- [1] B.L. Roberts and W.J. Marciano. *Lepton Dipole Moments*. Advanced series on directions in high energy physics. World Scientific Publishing Company Pte Limited, 2010.
- [2] Ann Nelson. Naturally weak CP violation. *Physics Letters B*, 136(5):387–391, 1984.
- [3] Stephen M Barr. Solving the strong CP problem without the Peccei-Quinn symmetry. *Physical Review Letters*, 53(4):329, 1984.
- [4] S. Weinberg. A new light boson. *Phys. Rev. Lett.*, 40(4), 1978.
- [5] F. Wilczek. Problem of strong P and T invariance in the presence of instantons. *Phys. Rev. Lett.*, 40(5), 1978.
- [6] C. Edwards et al. Upper limit for  $\text{J}/\Psi \rightarrow \gamma + \text{Axion}$ . *Phys. Rev. Lett.*, 48(14), 1982.
- [7] M. Sievertz et al. Upper limit for axion production in radiative  $\Upsilon$  decay. *Phys. Rev. D*, 26(3), 1982.
- [8] J. Kim. Weak-interaction singlet and Strong CP invariance. *Phys. Rev. Lett.*, 43(2), 1979.
- [9] M. Shifman, A. Vainshtein, and V. Zakharov. Can confinement ensure natural CP invariance of strong interactions? *Nuclear Physics B*, 166(3):493–506, 1980.
- [10] M. Dine, W. Fischler, and M. Srednicki. A simple solution to the Strong CP problem with a harmless axion. *Phys. Lett.*, 104B(3), 1981.
- [11] A. Zhitnitsky. On possible suppression of the axion hadron interactions. *Sov. J. Nucl. Phys.*, 31(260), 1980.
- [12] Jean-Pierre Luminet. Lemaitre’s big bang. *arXiv preprint arXiv:1503.08304*, 2015.
- [13] J. H. Oort. The force exerted by the stellar system in the direction perpendicular to the galactic plane and some related problems. *BAIN*, 6:249, August 1932.

- [14] F. Zwicky. On the Masses of Nebulae and of Clusters of Nebulae. *ApJ*, 86:217, October 1937.
- [15] V. C. Rubin, W. K. J. Ford, and N. . Thonnard. Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605 /R = 4kpc/ to UGC 2885 /R = 122 kpc/. *ApJ*, 238:471–487, June 1980.
- [16] NASA/WMAP Science Team. Nucleosynthesis in the Early Universe. [https://map.gsfc.nasa.gov/universe/bb\\_tests\\_ele.html](https://map.gsfc.nasa.gov/universe/bb_tests_ele.html). Page Updated: Friday, 12-10-2010.
- [17] G. Hinshaw et al. Five-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Data Processing, Sky Maps, and Basic Results. *Astrophys.J.Suppl.*, 180:225–245, 2009.
- [18] P. Ade et al. Planck Collaboration. Planck 2013 results. xvi. cosmological parameters. *AA*, 571:A16, 2014.
- [19] European Space Agency. Planck and the cosmic microwave background. [http://www.esa.int/Our\\_Activities/Space\\_Science/Planck/Planck\\_and\\_the\\_cosmic\\_microwave\\_background](http://www.esa.int/Our_Activities/Space_Science/Planck/Planck_and_the_cosmic_microwave_background), 2017.
- [20] Croton et al. The many lives of active galactic nuclei: cooling flows, black holes and the luminosities and colours of galaxies. *MNRAS*, 365:11–28, jan 2006.
- [21] Eisenstein et al. Detection of the baryon acoustic peak in the large-scale correlation function of sdss luminous red galaxies. *The Astrophysical Journal*, 633:560–574, November 1, 2005 2005. n/a.
- [22] C. S. Frenk and Simon D. M. White. Dark matter and cosmic structure. *Annalen Phys.*, 524:507–534, 2012.
- [23] Superclusters. *Annual Review of Astronomy and Astrophysics*, 21(1):373–428, 1983.
- [24] Alan H. Guth. Inflationary universe: A possible solution to the horizon and flatness problems. *Phys. Rev. D*, 23:347–356, Jan 1981.
- [25] N. Suzuki et al. The hubble space telescope cluster supernova survey. v. improving the dark-energy constraints above  $z > 1$  and building an early-type-hosted supernova sample. *The Astrophysical Journal*, 746(1):85, 2012.

- [26] Adam G. Riess et al. Observational evidence from supernovae for an accelerating universe and a cosmological constant. *Astron. J.*, 116:1009–1038, 1998.
- [27] C. D. Hoyle, D. J. Kapner, Blayne R. Heckel, E. G. Adelberger, J. H. Gundlach, U. Schmidt, and H. E. Swanson. Sub-millimeter tests of the gravitational inverse-square law. *Phys. Rev.*, D70:042004, 2004.
- [28] Edward W. Kolb and Michael S. Turner. *The Early Universe*. Frontiers in Physics. Addison-Wesley Publishing Company, 1990.
- [29] C. W. Purcell, J. S. Bullock, and M. Kaplinghat. The Dark disk of the Milky Way. *ApJ*, 703(2):2275, 2009.
- [30] L.D. Duffy and P. Sikivie. The Caustic Ring Model of the Milky Way Halo. *Phys. Rev.*, D78:063508, 2008.
- [31] Erik W Lentz, Thomas R Quinn, Leslie J Rosenberg, and Michael J Tremmel. A New Signal Model for Axion Cavity Searches from N-Body Simulations. *submitted to Ap. J.*, 2017.
- [32] Dicus et al. Astrophysical bounds on the masses of axions and Higgs particles. *Phys. Rev. D*, 18(6), 1978.
- [33] Daniel Grin, Giovanni Covone, Jean-Paul Kneib, Marc Kamionkowski, Andrew Blain, and Eric Jullo. Telescope search for decaying relic axions. *Phys. Rev. D*, 75:105018, May 2007.
- [34] Edward Daw. *A Search for Halo Axions*. PhD thesis, Massachusetts Institute of Technology, 1998.
- [35] C. Patrignani et al. Review of Particle Physics. *Chin. Phys.*, C40(10):100001, 2016.
- [36] S. Aune et al. CAST search for sub-eV mass solar axions with  ${}^3\text{He}$  buffer gas. *Phys. Rev. Lett.*, 107:261302, 2011.
- [37] P. Sikivie. Experimental tests of the invisible axion. *Phys. Rev. Lett.*, 51(16).
- [38] Joseph Thomas Rogers. *Limits on the Electromagnetic Coupling and Density of Galactic Axions*. PhD thesis, University of Rochester, 1987.

- [39] Darin Kinion. *First Results from a Multiple-Microwave-Cavity Search for Dark-Matter Axions*. PhD thesis, Massachusetts Institute of Technology, 2001.
- [40] D. B. Yu. *An Improved RF Cavity Search for halo Axions*. PhD thesis, Massachusetts Institute of Technology, 2004.
- [41] Lawrence Krauss, John Moody, Frank Wilczek, and Donald E. Morris. Calculations for cosmic axion detection. *Phys. Rev. Lett.*, 55:1797–1800, Oct 1985.
- [42] Christian A Hagmann. *A Search for Cosmic Axions*. PhD thesis, University of Florida, 1990.
- [43] Ana T. Malagon. *Search for 140 micro-eV Pseudoscalar and Vector Dark Matter Using Microwave*. PhD thesis, Yale University, 2014.
- [44] Ben Brubaker. Derivation of admx-hf sensitivity, 2015. Yale University.
- [45] Michael Hotz. *A SQUID-Based RF Cavity Search for Dark Matter Axions*. PhD thesis, University of Washington, 2013.
- [46] Dmitry Lyapustin. *An improved low-temperature RF-cavity search for dark-matter axions*. PhD thesis, University of Washington, 2015.
- [47] Rybka et al. Search for dark matter axions with the Orpheus experiment. *Phys. Rev. D*, 91, 2015.
- [48] K. Patel D. Garrett J. Sloan, G. Rybka. Electric tiger, 2014. Open Discussion & Presentation of New Ideas.
- [49] P. Sikivie, N. Sullivan, and D. B. Tanner. Proposal for axion dark matter detection using an *lc* circuit. *Phys. Rev. Lett.*, 112:131301, Mar 2014.
- [50] C. Hagmann, P. Sikivie, N. S. Sullivan, and D. B. Tanner. Results from a search for cosmic axions. *Phys. Rev. D*, 42:1297–1300, Aug 1990.
- [51] W. U. Wuensch, S. De Panfilis-Wuensch, Y. K. Semertzidis, J. T. Rogers, A. C. Melissinos, H. J. Halama, B. E. Moskowitz, A. G. Prodell, W. B. Fowler, and F. A. Nezrick. Results of a laboratory search for cosmic axions and other weakly coupled light particles. *Phys. Rev. D*, 40:3153–3167, Nov 1989.

- [52] J. Clarke and A. I. Braginski. *The SQUID Handbook: Volume 2 Applications of SQUIDS and SQUID Systems*. Wiley, third edition, 2006.
- [53] Attocube. Premium line introduction, 2017. <http://www.attocube.com/attomotion/>.
- [54] John David Jackson. *Classical Electrodynamics*. Wiley, third edition, 1998.
- [55] Mark Jones. Main cavity  $tm_{010}$  mode-map for different rod end gaps, 2016. Pacific Northwest National Laboratory.
- [56] Igor G Irastorza. The International Axion Observatory IAXO. Letter of Intent to the CERN SPS committee. Technical Report CERN-SPSC-2013-022. SPSC-I-242, CERN, Geneva, Aug 2013.
- [57] Aaron S Chou and Mark Jones. private communication, 2016.
- [58] B. M. Brubaker, L. Zhong, Y. V. Gurevich, S. B. Cahn, S. K. Lamoreaux, M. Simanovskaia, J. R. Root, S. M. Lewis, S. Al Kenany, K. M. Backes, I. Urdinaran, N. M. Rapidis, T. M. Shokair, K. A. van Bibber, D. A. Palken, M. Malnou, W. F. Kindel, M. A. Anil, K. W. Lehnert, and G. Carosi. First results from a microwave cavity axion search at 24  $\mu$ eV. *Phys. Rev. Lett.*, 118:061302, Feb 2017.