

An improved low-temperature RF-cavity search for dark-matter
axions

Dmitry Lyapustin

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

Doctor of Philosophy

University of Washington

2015

Reading Committee:

Leslie Rosenberg, Chair

Ann Nelson

Alejandro Garcia

Program Authorized to Offer Degree:
UW Physics

©Copyright 2015

Dmitry Lyapustin

University of Washington

Abstract

An improved low-temperature RF-cavity search for dark-matter axions

Dmitry Lyapustin

Chair of the Supervisory Committee:
Dr. Leslie Rosenberg
Physics

Axions are hypothetical elementary particles that may help provide the answer as to why QCD preserves the discrete symmetries P and CP. Light axions also have properties that make them ideal dark-matter candidates. The Axion Dark Matter eXperiment (ADMX), has been at the forefront of the search for dark-matter axions for over a decade, and over the past few years has undergone upgrades to dramatically improve its sensitivity. I give a brief motivation for dark matter axions, discuss the hardware and software aspects of ADMX, review the 3D cavity simulations that are an integral part of the experiment, analyze the 2014 data set and present those results, and highlight some of the ongoing R&D efforts within ADMX. New limits on the axion-to-photon coupling constant $g_{a\gamma\gamma}$ are set on select frequency ranges between 600 - 800 MHz and 1050 - 1400 MHz.

TABLE OF CONTENTS

	Page
List of Figures	iv
Chapter 1: Introduction	1
Chapter 2: Motivation for dark matter axions	2
2.1 Introduction	2
2.2 The Strong CP problem	2
2.2.1 Possible solutions to the Strong CP problem	4
2.3 Axion theory and phenomenology	4
2.3.1 Axion mass and couplings	5
2.3.2 The PQWW axion	6
2.3.3 The invisible axion	8
2.4 Axions as dark matter	8
2.4.1 Axion cosmology	9
2.4.2 Limits on axion mass	11
Chapter 3: The ADMX experiment	13
3.1 Introduction	13
3.2 The axion haloscope	13
3.2.1 Power deposited in a haloscope	14
3.2.2 Signal temperature dependence	17
3.3 ADMX hardware	17
3.3.1 Resonant cavity	18
3.3.2 SQUID amplifier	24
3.3.3 HEMT amplifiers	29

3.3.4	Cold space	32
3.3.5	Reservoir	36
3.3.6	Warm space	39
3.3.7	Receiver chain	39
3.4	ADMX DAQ software	42
3.5	Data taking cadence	44
3.5.1	2014 operations	45
Chapter 4:	3D cavity simulations	47
4.1	Introduction	47
4.2	Need for cavity studies	47
4.2.1	Motivation for 3D modeling	48
4.3	Methodology	51
4.3.1	3D model	52
4.4	Simulation results	53
4.4.1	Guide for data-taking	53
4.4.2	Capacitive effect of tuning rods	63
Chapter 5:	Data analysis	69
5.1	Introduction	69
5.2	Power spectrum conditioning	69
5.2.1	Quality control	70
5.2.2	Structure removal	70
5.2.3	Power scaling	73
5.3	Grand spectrum	77
5.3.1	Cavity response	78
5.3.2	Mini grand spectrum	79
5.3.3	2014 data set grand spectrum	82
5.4	Axion candidate search	86
5.5	Limit generation	86
5.5.1	Dilution factor	86
5.5.2	Binning a Maxwellian	88

5.5.3	Upper limit on power per bin	89
5.5.4	Axion-to-photon limit	91
Chapter 6:	Conclusions and future work	96
6.1	Future work	96
6.1.1	Dilution refrigerator	96
6.1.2	Use of JPAs	98
6.1.3	Improved thermal design	99
6.1.4	ADMX Sidecar	100
6.1.5	Ouroboros	100
6.1.6	New cavity designs	101
6.1.7	Final words	109
Bibliography	112
Appendix A:	Ouroboros	115

LIST OF FIGURES

Figure Number	Page
2.1 Spinning neutron under CP transformation	3
2.2 Feynman diagram for the process $a \rightarrow \gamma\gamma$	6
2.3 Cartoon of axion-mediated energy loss in red giant stars	7
2.4 QCD potential of the PQ field Φ	10
2.5 Plot of $g_{a\gamma\gamma}$ vs. m_a	12
3.1 Axion detection in a Sikivie haloscope	14
3.2 Circuit diagram of a haloscope cavity	15
3.3 Cutaway model of ADMX Phase II	19
3.4 ADMX cavity interior and model with tuning rods	21
3.5 Tuning rod gear box	22
3.6 Schematic of transmission and reflection measurements	23
3.7 dc SQUID circuit diagram	25
3.8 dc SQUID V - Φ_B curve	26
3.9 ADMX SQUID amplifier cartoon	27
3.10 MSA noise temperature versus physical temperature	28
3.11 Squidadel	30
3.12 SQUID amplifier noise temperature circuit schematic	31
3.13 Schematic of a balanced HEMT amplifier	32
3.14 Insert Cold space	34
3.15 Cold space temperatures	35
3.16 Reservoir with bucking coil	37
3.17 Bucking coil field cancellation during a ramp.	38
3.18 Insert Warm space	40
3.19 Insert being lowered into cryostat	41

3.20	Receiver chain	43
3.21	Electronics rack during data-taking	44
3.22	Illustration of two possible data-taking schemes	45
4.1	Cavity mode map around TM_{010}	49
4.2	Cavity mode map around TM_{020}	50
4.3	COMSOL cavity model	52
4.4	Simulated TM_{010} mode	54
4.5	TM_{010} mode map with simulation	55
4.6	TM_{010} form-factor map	56
4.7	$TM_{020(A)}$ form-factor map	57
4.8	$TM_{020(B)}$ form-factor map	57
4.9	$TM_{020(C)}$ form-factor map	58
4.10	f_{020} over 900 MHz - 1500 MHz	59
4.11	Contour map of $f_{010} \times f_{020(C)}$ in the 2104 data taking region	60
4.12	Mode map used to guide data taking	61
4.13	Mode map overlaid on simulation data to guide data-taking	62
4.14	TM_{010} form-factor map from 2D simulation	64
4.15	f_{010} over TM_{010} tuning range with 0" tuning rod gap	64
4.16	f_{010} over TM_{010} tuning range with 0.025" tuning rod gap	65
4.17	f_{010} over TM_{010} tuning range with 0.25" tuning rod gap	65
4.18	f_{010} over TM_{010} tuning range with 0.75" tuning rod gap	66
4.19	Capacitive tuning rod effect	68
5.1	Raw power spectrum	72
5.2	Power spectrum with gross structure fit	73
5.3	Power spectrum less gross structure	74
5.4	Power spectrum less gross structure with fit	74
5.5	Power spectrum less gross and fine structures	75
5.6	Histogram of deviations overlaid with a Gaussian fit	76
5.7	Corrected spectrum	77
5.8	Lorentzian for typical cavity parameters	78

5.9	Mini grand spectrum	80
5.10	Grand spectrum uncertainty vs number of points	81
5.11	Axion lineshape	81
5.12	Mini grand spectrum with synthetic axions	82
5.13	Grand spectrum for 590 - 630 MHz	83
5.14	Grand spectrum for 720 - 760 MHz	83
5.15	Grand spectrum for 1050 - 1080 MHz	84
5.16	Grand spectrum for 1210 - 1280 MHz	84
5.17	Grand spectrum for 1350 - 1390 MHz	85
5.18	Mini grand spectrum with synthetic axions and 5σ curve	87
5.19	Maximum 1 kHz axion power spread over two bins	89
5.20	Minimum 1 kHz axion power spread over two bins	90
5.21	Exclusion plot for 590 - 630 MHz	91
5.22	Exclusion plot for 720 - 760 MHz	92
5.23	Exclusion plot for 1050 - 1080 MHz	92
5.24	Exclusion plot for 1210 - 1280 MHz	93
5.25	Exclusion plot for 1350 - 1390 MHz	93
5.26	Exclusion limits on $g_{a\gamma\gamma}$ for 1050 - 1400 MHz compared to previous limits . .	94
5.27	Exclusion limits on $g_{a\gamma\gamma}$ for 590 - 760 MHz compared to previous limits . . .	95
6.1	Picture of the ADMX dilution refrigerator	97
6.2	JPA in open container and installed in squidadel	99
6.3	Picture of improved Cold space	102
6.4	Prototype Sidecar cavity.	103
6.5	Model of Sidecar cavity installed on main cavity	104
6.6	Ouroboros recent results	105
6.7	Miniature cavity prototype	106
6.8	UF vaned cavity	107
6.9	Electric tiger prototype	108
6.10	Orpheus experiment schematic	109
6.11	Orpheus prototype setup	110
6.12	Orpheus results	111

A.1	Ouroboros schematic	115
A.2	Ouroboros setup	116

ACKNOWLEDGMENTS

First and foremost, I would like to thank my advisor Leslie Rosenberg for his years of good advice and astute social commentary. I thank my reading committee, my colleagues at ADMX, and the many physicists, engineers, and machinists at the University of Washington who have been ever receptive to my many questions over the years. Last but certainly not least, I thank my wife for thinking of the long term and supporting a struggling scientist for so many years.

DEDICATION

This thesis is dedicated to my two cats, without whom daily life would be a lot less full of cat hair.

Chapter 1

INTRODUCTION

My thesis describes the third incarnation of the Axion Dark Matter eXperiment (ADMX), known as ADMX Phase II. This experiment is located at the Center for Experimental Physics and Astrophysics (CENPA) at the University of Washington in Seattle, WA. Over the past five years ADMX has been completely redesigned and rebuilt, and as a result it has been able to place new bounds on the axion-to-photon coupling constant $g_{a\gamma\gamma}$ in the 1050 - 1080 MHz and 1350 - 1390 MHz regions from data collected in 2014. ADMX Phase II represents the start of a larger-timescale project known as ADMX Gen II. The Gen II experiment will operate at a temperature more than an order of magnitude lower than any previous RF cavity axion experiment, and it will be able to find or exclude even the most pessimistically-coupled axions over a wide range of axion masses.

I have been fortunate to be involved in all aspects of ADMX. After joining the group in early 2010, I was tasked with designing the new experimental insert. For the next three years I was kept busy designing components, machining them or having them professionally made, and integrating them into the ever-growing insert. At the same time, I designed and helped assemble many ancillary pieces of the experiment, such as ADMX clean room and the plumbing for ADMX's ^4He liquefaction system. From the end of 2013 through 2014 my tasks additionally included running the experiment and analyzing data, as well as performing 3D cavity simulations.

Chapter 2

MOTIVATION FOR DARK MATTER AXIONS

2.1 Introduction

In this chapter I argue that the axion is a very well-motivated particle that also happens to be an attractive dark matter candidate. I first present the historical basis for axions and show how their existence would provide an explanation for a long-standing problem in physics. I then discuss the evolution of axion theory as experiments began to constrain heavy-axion models. Next, I show that the light “invisible” axion is one of the favored dark matter candidates. I finish this chapter by looking at the current limits on the axion mass.

2.2 The Strong CP problem

The story of axions begins with Quantum Chromodynamics (QCD). QCD is the part of Standard Model of Physics governing the strong force and describes the interactions of quarks and gluons. It has been rigorously tested by a myriad of experiments since the late-20th century. For a detailed study of this theory, see, for example, reference [1].

Though QCD has been an incredibly successful theory, it exhibits one glaring fault, known as the Strong CP problem. This problem arises from a CP¹-violating term in the QCD lagrangian of the form

$$\mathcal{L}_{\bar{\theta}} = \frac{\bar{\theta}}{32\pi^2} G_{\mu\nu}^a \tilde{G}^{\mu\nu a}, \quad (2.1)$$

where $G_{\mu\nu}^a$ is the gluon field strength tensor and $\tilde{G}^{\mu\nu a}$ its dual. $\bar{\theta}$ is a number which characterizes the strength of the CP violation, and has contributions from both the strong and

¹CP is the product of charge and parity symmetries. Under a charge (C) transformation all particles are turned into their antiparticles. Under a parity (P) transformation the coordinate system becomes inverted.

weak sectors:

$$\bar{\theta} = \theta + \text{arg det } M. \quad (2.2)$$

Here M is the quark mass matrix, which plays an integral role in weak physics, and θ is a value which characterizes the QCD vacuum. Since $G_{\mu\nu}^a \tilde{G}^{\mu\nu a}$ is a total derivative it does not contribute to the equations of motions; nevertheless, the CP violation from $\mathcal{L}_{\bar{\theta}}$ leads to observable effects. Both of the terms contributing to $\bar{\theta}$ can take on values from 0 to 2π , and naively one would expect them to be $\mathcal{O}(1)$. Further, as these values arise from different sectors of physics there is no reason to expect they should relate to one another. Thus $\bar{\theta}$ is also expected to be $\mathcal{O}(1)$.

The most sensitive experiment one can do to obtain a value for $\bar{\theta}$ is to look for CP violation in QCD through observation of a neutron electric dipole moment (nEDM). Figure 2.1 illustrates the idea of why a nonzero nEDM would necessarily violate CP symmetry. Recent results give an upper limit for the nEDM of $|d_n| < 2.9 \times 10^{-26} \text{ e cm}$ [2], which




Figure 2.1: A Spinning neutron with an EDM undergoing a CP transformation. The spin of the neutron gives it a magnetic moment μ , and if I assume that the quarks in the neutron are arranged such that it has net positive and negative lobes, then the neutron also has an electric dipole moment d . If μ and d are initially parallel, it is easy to see that under a CP transformation d is inverted while μ is unchanged, so the two vectors are now antiparallel. A neutron with parallel electric and magnetic moments is physically distinguishable from one with antiparallel moments, thus CP symmetry is broken in this system.

corresponds to $\bar{\theta} < 10^{-10}$ [3, 4]. This 10 orders of magnitude discrepancy between the upper limit on $\bar{\theta}$ and its expected value is known as the Strong CP problem. It should be noted that this is a fine-tuning problem: in principle there is no reason we cannot have $\bar{\theta} = 0$ by making $\theta = \arg \det M$. However, simply imposing this fact requires one to choose the values for these numbers precisely, which is neither an elegant nor intuitive solution.

2.2.1 Possible solutions to the Strong CP problem

There have been many attempts over the years at a solution to the Strong CP problem. The most common theories are based either on “naturalness,” where $\bar{\theta}$ is forced to zero at tree level in Feynman diagrams, or “relaxation,” where $\bar{\theta}$ becomes a variable that relaxes to 0 [5]. Natural theories tend to have the issue that higher order loops tend to contribute to $\bar{\theta}$, thus it becomes difficult to realize $\bar{\theta} = 0$ once these higher-order effects are taken into account. Relaxation theories are based on the Peccei-Quinn mechanism, which is arguably the most popular solution to the Strong CP problem.

In 1977, Roberto Peccei and Helen Quinn demonstrated that with the addition of a global $U(1)$ symmetry, designated $U(1)_{PQ}$, to the QCD lagrangian, $\bar{\theta}$ becomes a dynamic variable that affects the value of the QCD potential [6]. The potential is minimized for $\bar{\theta} = 0$, so there is now a elegant explanation for why CP violation is suppressed in the strong force.

2.3 Axion theory and phenomenology

Less than a year after Peccei and Quinn proposed their brilliant solution to the Strong CP problem, Steven Weinberg [7] and Frank Wilczek [8] realized that this solution would have physical consequences. The proposed $U(1)_{PQ}$ symmetry is explicitly broken by instanton² effects, which results in the creation of a neutral pseudoscalar boson. Weinberg and Wilczek

²Instantons are pseudoparticles that connect the degenerate but inequivalent vacuum ground states of QCD. They arise from the non-Abelian nature of QCD and its multiple degenerate vacua.

named this boson the “axion.”

2.3.1 Axion mass and couplings

The Peccei-Quinn axion model has one main parameter, the axion mass m_a , or, alternatively, the Peccei-Quinn symmetry breaking scale f_{PQ} . The two are related through

$$m_a = \frac{\sqrt{z}}{1+z} \frac{f_\pi m_\pi}{f_{PQ}/N}, \quad (2.3)$$

where $z = m_u/m_d$, m_u and m_d are the masses of the up and down quarks, m_π and f_π are the mass and decay constant of the pion, and N is the PQ color anomaly [9]. It is important to note that the mass is inversely proportional to the coupling constant, so a larger coupling constant will lead to a smaller mass.

The axion coupling strength to normal matter and radiation is characterized by the coupling constants $g_{a\gamma\gamma}$, g_{aee} , etc... Though the axion does couple to photons, electrons, and nucleons, for $m_a \leq 2m_e \approx 1 \text{ MeV}$, where m_e is the mass of the electron, the $a \rightarrow \gamma\gamma$ interaction dominates. The lagrangian term describing this interaction is given by

$$\mathcal{L}_{a\gamma\gamma} = -g_{a\gamma\gamma} a \vec{E} \cdot \vec{B}, \quad (2.4)$$

where $g_{a\gamma\gamma}$ is the axion-to-two-photon coupling constant, a is the axion field, and \vec{E} and \vec{B} are the usual electric and magnetic fields. Figure 2.2 shows the Feynman diagram for this interaction.

The specific $a \rightarrow \gamma\gamma$ effective coupling is given by

$$g_{a\gamma\gamma} = \frac{\alpha/2\pi}{f_{PQ}/N} \left(\frac{E}{N} - \frac{2(4+z)}{3(1+z)} \right), \quad (2.5)$$

where E is the PQ electromagnetic anomaly. Current values of m_u and m_d [10] give $2(4+z)/(3(1+z)) = 1.95$. The value E/N is model-dependent, and in principle there is no reason it cannot be ~ 1.95 , thus strongly suppressing the axion-to-two-photon interaction;




Figure 2.2: Feynman diagram for the process $a \rightarrow \gamma\gamma$. Since the axion is a neutral particle it cannot couple to photons at tree level. Instead this interaction must proceed through an electromagnetic or color anomaly.

however, DFSZ, the simplest Grand Unified Theory (GUT) model incorporating axions, uses $E/N = 8/3 \approx 2.67$.

Model dependence is also inherent in the axion couplings to electrons and nucleons. Different models make different predictions on the axion mass, thereby giving experimentalists possible axion search windows.

2.3.2 The PQWW axion

The original axion that Weinberg and Wilczek posited, dubbed the Peccei-Quinn-Weinberg-Wilczek (PQWW) axion, had a symmetry breaking scale around the weak scale $\sqrt{2G_F} = 246 \text{ GeV}$. With f_a of this order the axion is expected to have mass $10 \text{ keV} \lesssim m_a \lesssim 1000 \text{ keV}$.

Astrophysical observations can provide limits on the PQ axion mass, though model dependence cannot be completely circumvented. The most sensitive observational evidence

that helped kill the PQWW axion came from red giant evolution. The basic idea is that if axions exist, they would provide red giant stars an additional energy-loss mechanism. For example, the Primakoff ($Z, e^- + \gamma \rightarrow Z, e^- + a$) or Compton ($e^- + \gamma \rightarrow e^- + a$) processes would allow energy to escape stars through the axion intermediary. This would in turn speed up stellar evolution, and if the calculated evolution rate disagrees with observed values a limit can be placed. Utilizing this method, Dicus et al. helped show that the PQWW axion could not exist [11]. Figure 2.3 illustrates the idea behind this observation.




Figure 2.3: Cartoon of axion-mediated energy loss in red giant stars. Photons have a very short mean free path, typically on order of 1 cm, within red giant stars. This strongly limits the rate at which these stars can radiate energy to the vacuum of space, which in turn elongates their lifetimes. If axions of the proper mass existed, they would effectively increase photons' mean free path. A star is far more transparent to axions than to photons, so axions would have no problem escaping the star and carrying away energy. If $m_a \approx 0.2 \text{ meV}$, for example, red giant stars' lifetimes would be reduced to just 1 year! [11]

Among the most sensitive laboratory experiments searching for PQWW axions involved looking for rare meson decays. Specifically, $J/\Psi \rightarrow a + \gamma$ and $\Upsilon \rightarrow a + \gamma$ were attractive to search for since the product of the branching ratios for these processes is model-independent. For instance, the CUSB collaboration looked for $\Upsilon \rightarrow a + \gamma$, and combining their results

with data on $J/\Psi \rightarrow a + \gamma$ from the Crystal Ball detector [12], they were able to rule out the PQWW axion with high confidence [13].

2.3.3 *The invisible axion*

With the demise of the PQWW axion, new models were required to maintain the PQ solution to the Strong CP problem. Thus was the “invisible” axion introduced, so called because it was thought this kind of particle had such a small mass it could not be detected. Two very prominent invisible-axion models emerged, one known as KSVZ, and the other as DFSZ.

The KSVZ model, pioneered by Kim [5] and Shifman, Vainshtein, and Zakharov [14], is the first of the invisible axion models. Under KSVZ, the axion couples neither to the standard quarks nor to electrons at tree level (so the electromagnetic anomaly $E = 0$). Instead, axions have a tree level coupling to a new kind of heavy quark. KSVZ is called a hadronic model, since the axions have bare couplings only to quarks and not leptons. While KSVZ does allow for an axion which would have evaded early searches, it is problematic in its positing a new heavy quark, which is not embedded in most extensions of the Standard Model.

The DFSZ model, proposed by Dine, Fischler and Srednicki [15] and Zhitnitsky [16] was developed several years after KSVZ. DFSZ is a leptonic model, meaning that axions in this model have bare couplings to leptons as well as quarks. It is important to note that the axion-to-matter couplings under DFSZ are suppressed by a factor $r = \sqrt{f_u^2 + f_d^2}/f_{PQ}$, where f_u and f_d are the decay constants of the two scalar doublets in the original PQ theory. This results in a decreased value for the coupling constant $g_{a\gamma\gamma}$ under DFSZ than under KSVZ for a given f_{PQ} .

2.4 *Axions as dark matter*

The invisible axion makes an attractive dark-matter particle candidate. Since $g_{a\gamma\gamma} \propto 1/f_{PQ}$ and $m_a \propto 1/f_{PQ}$, it is clear that $g_{a\gamma\gamma} \propto m_a$; therefore, a very light axion has a very weak

coupling to matter. This ensures not only that axions cannot easily be detected ("dark"), but further it means axions are remarkably stable. The process $a \rightarrow \gamma\gamma$ proceeds with a lifetime of

$$\tau_{a\gamma\gamma} = 6.8 \times 10^{54} \text{sec} \frac{(m_a/\mu\text{eV})^{-5}}{[(E/N - 1.95)/0.72]^2}, \quad (2.6)$$

so for m_a of order μeV and $E/N \approx 8/3$, the axion decays to two photons after about 10^{37} lifetimes of the universe.

Since heavy PQ axions have long since been ruled out and the light invisible axions are the only plausible choice, it is clear that axions are indeed very promising dark-matter candidates. The more interesting question to ask is whether or not they constitute a large portion of the total dark matter we see in the universe.

2.4.1 Axion cosmology

To ascertain how much of the dark matter is comprised of axions, it is necessary to consider how they are created. First, though, it is important to note that any significant axion contribution to dark matter likely came from relic axions. These are axions which were created very early in the history of the universe.

One of the main sources of relic axions is a process known as the "misalignment mechanism." In this mechanism axions are created from oscillations of the axion field about the minimum of the QCD potential for the PQ field Φ , as shown in figure 2.4. The basic idea is that the axion field in some region initially has a random value for $\bar{\Theta}$ when instanton effects "turn on" and tilt the potential. The field then rolls towards the new minimum of the potential at $\bar{\Theta} = 0$. It overshoots, and thereafter oscillates about the minimum with frequency m_a . This is clearly a non-thermal process, and axions produced in this manner are in a Bose condensate [9]. They are quasi thermalized through gravitational interactions, so the local axion velocity is expected to be that of the local virial velocity, namely $10^{-3}c$. Their fractional linewidth near Earth is therefore 10^{-6} .




Figure 2.4: QCD potential of the PQ field Φ . Axions become relevant when the PQ symmetry is spontaneously broken around a temperature $T \approx f_{PQ}$. While $f_{PQ} \geq T \geq \Lambda_{QCD}$, where the QCD scale parameter $\Lambda_{QCD} \approx 200 \text{ MeV}$, axions are massless. Instanton effects turn on around $T \approx \Lambda_{QCD}$ and cause the potential to tip, so that a single minimum forms. This explicit symmetry breaking gives the axion a small mass.

Axion production by the misalignment mechanism gives a contribution to the energy density of the universe of the form

$$\Omega_a = \frac{0.85 \times 10^{\pm 0.4}}{h^2} \left(\frac{\Lambda_{QCD}}{200 \text{ MeV}} \right)^{0.7} \left(\frac{m_a}{10^{-5} \text{ eV}} \right)^{-1.18}, \quad (2.7)$$

where the reduced Hubble factor $h \approx 0.7$. Recent results give $\Omega_{dark\ matter} \approx 0.27$ [17], so for $10^{-6} \text{ eV} \lesssim m_a \lesssim 10^{-4} \text{ eV}$, axions would make up a considerable fraction of the dark matter

we see in the universe.

Axions can also be created in the very early universe when temperatures were extremely high, though such “thermal” axions do not contribute significantly to Ω_a . There could exist other possible axion production mechanisms, too, such as through cosmic strings or domain walls. However, there is no agreement in theory on the contribution to Ω_a by these mechanisms.

2.4.2 Limits on axion mass

The dark matter axion region of $m_a \approx 1 - 100 \mu eV$ occupies a significant portion of the axion search window, shown in figure 2.5. The lower bound on m_a is provided by theoretical arguments, and the upper bound comes from observations and experiments.

The axion contribution to the energy density of the universe, given by (2.7), provides the basis for $m_a \gtrsim 1 \mu eV$. Clearly $m_a < 1 \mu eV \Rightarrow \Omega_a / \Omega_{dark\ matter} > 1$, meaning a too-light axion produces more dark matter than we observe in the universe today.

Observations from Supernova 1987A provide one of the most sensitive upper bounds for the axion mass. One of the signatures of SN1987A was a 12 second neutrino pulse observed on Earth. If $m_a \gtrsim 1000 \mu eV$, axions would have carried away enough energy from the supernova that the neutrino pulse would have been considerably shorter.³

It is important to note that attempting to measure such an extremely-weakly coupled particle as the invisible axion is very difficult, which is why only one experiment, ADMX, has been able to set bounds on the axion mass within the plausible model range. The next chapter describes ADMX and shows how this experiment has been able to achieve such good sensitivity to axions.

³This argument does not apply for $m_a > 0.02 eV$, since axions that are too massive and thus too strongly coupled would interact with normal matter before escaping the nascent neutron star’s core. Other observations and experiments exist which exclude $m_a > 0.02 eV$.




Figure 2.5: The axion search window, shown as $g_{a\gamma\gamma}$ vs. m_a . The yellow region around the KSVZ and DFSZ lines represent uncertainties in the models. Realistic QCD axions are expected to lie somewhere within this yellow region. An upper bound of $\approx 1000 \mu\text{eV}$ and a lower bound of $\approx 1 \mu\text{eV}$ tightly constrain the possible axion mass.

Chapter 3

THE ADMX EXPERIMENT

3.1 Introduction

This chapter goes into the details of ADMX. I first introduce the haloscope idea that ADMX is based on, and then discuss how an axion signal could be resolved in a haloscope experiment. Next, I describe the hardware of ADMX and detail its various components. Included here is a discussion of the amplifier chain that allows ADMX to extract signals from the experiment. Then, I give a brief overview of the software of the ADMX data acquisition system (DAQ), which takes the raw data coming from ADMX and processes it for analysis. Finally, I discuss the data-taking cadence of the 2014 data runs.

3.2 The axion haloscope

The idea at the heart of ADMX is that of Pierre Sikivie's axion haloscope [18]. A haloscope detects relic axions in our local dark matter halo. This is done by immersing a resonant cavity in a very strong static magnetic field and looking for the axion signal coming out of the cavity. Axions from our galactic halo would resonantly convert inside the cavity to microwave photons, and the interaction lagrangian for this process is given by $\mathcal{L}_{a\gamma\gamma} = -g_{a\gamma\gamma}a\vec{E}\cdot\vec{B}$. An antenna, critically coupled to the cavity, would then extract the microwave signal. Figure 3.1 illustrates this idea.

With the axion mass known, one could simply construct a cavity with the proper physical dimensions and resonant frequency which would be able to quickly detect the axion. As the mass is not known, the cavity must be made such that its resonant frequency is tunable. One method to achieve this is to deploy within the cavity one or multiple rods that can rotate off




Figure 3.1: Axion detection in a Sikivie haloscope. Axions from our galactic halo scatter off of the virtual photons of a magnetic field and convert to microwave photons inside a resonant cavity. This process is enhanced by the Q factor of the cavity when $f_{cavity} \approx m_a$.

center with respect to the cavity axis. This technique is further described in the “Resonant cavity” section below.

3.2.1 Power deposited in a haloscope

To estimate the power axions deposit in a haloscope, I will treat the cavity as a resonant RLC circuit. R is a dissipative term, and axions are represented in this circuit by a voltage term V_a . The equivalent circuit is shown in figure 2.2. For the following discussion, charge will be denoted by q .




Figure 3.2: Circuit diagram of a haloscope cavity.

I first equate the energy deposited by charges moving through the circuit, qV_a , to the energy deposited by axions in a cavity of volume V . This is just $\mathcal{L}_{a\gamma\gamma}$ integrated over the volume of the cavity:

$$qV_a = -g_{a\gamma\gamma}a \int_V dV \vec{E} \cdot \vec{B}. \quad (3.1)$$

For a static axial magnetic field, $\vec{E} \cdot \vec{B} = B \times \vec{E} \cdot \vec{z}$. I now define a quantity f_{nlm} known as the “form factor,” which is a measure of how well the electric and magnetic fields overlay within the cavity:

$$f_{mnp} \equiv \frac{\left(\int_V dV \vec{E} \cdot \vec{z} \right)^2}{V \int_V dV E^2}. \quad (3.2)$$

The subscripts on f_{mnp} differentiate individual cavity modes. Substituting (3.2) into (3.1) gives me

$$qV_a = -g_{a\gamma\gamma}aB \sqrt{f_{mnp}V \int_V dV E^2}. \quad (3.3)$$

I then relate the total energy in a cavity mode to the total energy stored in the circuit,

$$\frac{1}{2} \int_V dV E^2 = \frac{1}{2} \frac{q}{C^2}. \quad (3.4)$$

Combining (3.4) and (3.3) results in

$$V_a = -g_{a\gamma\gamma} a B \sqrt{\frac{f_{mnp} V}{C}}. \quad (3.5)$$

Now I write the power absorbed by the circuit as

$$P_a = \frac{V_a^2}{R} = g_{a\gamma\gamma}^2 a^2 B^2 V f_{mnp} \frac{1}{RC}. \quad (3.6)$$

Finally, I equate the circuit quality factor $Q = (1/R)\sqrt{L/C}$ and resonant frequency $\omega_0 = 1/\sqrt{LC}$ to the same values for the resonant cavity to obtain

$$P_a = g_{a\gamma\gamma}^2 a^2 B^2 V f_{mnp} \omega_0 Q. \quad (3.7)$$

To get a numerical value for this power I rewrite the dimensionless quantity $ag_{a\gamma\gamma}$ in terms of the local axion density ρ_a , axion mass $m_a \approx \omega_0$, and a model-dependent constant g_γ , and give (3.7) in MKS units:

$$P_a = 1.01 \times 10^{-41} c^2 \epsilon_0 g_\gamma^2 \rho_a^2 B^2 V f_{mnp} \omega_0 Q. \quad (3.8)$$

Here c is the speed of light in m/s, B is measured in Tesla, V in m^3 , ρ_a in GeV/cc, and ω_0 in 1/s. The dimensionless constant g_γ equates to 0.36 under DFSZ and -0.97 under KSVZ. For parameter values typical of ADMX, the power comes to

$$P_a = 1.70 \times 10^{-21} W \left(\frac{V}{140 L} \right) \left(\frac{B}{7 T} \right)^2 \left(\frac{f_{mnp}}{0.5} \right) \left(\frac{\rho_a}{0.3 \frac{\text{GeV}}{\text{cc}}} \right) \left(\frac{g_{a\gamma\gamma}}{0.97} \right)^2 \left(\frac{Q}{10^5} \right) \left(\frac{\omega_0}{750 \text{ MHz}} \right). \quad (3.9)$$

An antenna can maximally extract only half of the total power within a cavity mode, while the other half is dissipated within the cavity walls. An antenna in such a configuration is ‘‘critically coupled’’ to the cavity mode. Assuming critical coupling, I therefore expect axions

to produce of order $10^{-21} W$ in the ADMX cavity. This minuscule power strongly suggests that thermal noise will contribute significantly to the ADMX data, and it also highlights the fact that lowering the temperature of the experiment can make it far more efficient at searching for axions.

3.2.2 Signal temperature dependence

I can quantify the temperature dependence of a signal observed by a haloscope by considering the noise power generated alongside the axion signature. I assume that blackbody radiation as well as electronic voltage fluctuations in the receiver chain make up most of the noise in this system. Assigning these two contributions an equivalent temperature $T_{system} = T_{blackbody} + T_{electronics}$, I write down the noise power P_N as

$$P_N = k_B T_{system} B. \quad (3.10)$$

This is just Johnson noise where k_B is the Boltzmann constant and B is the bandwidth of the signal. I now look at the signal to noise ratio S/N , given by the Dickie radiometer equation as

$$\frac{S}{N} = \frac{P_S}{P_N} \sqrt{Bt} = \frac{P_a}{k_B T_{system}} \sqrt{\frac{t}{B}}, \quad (3.11)$$

where t is the signal integration time and $P_S = P_a$ [19]. Thus $S/N \propto 1/T_{system}$, or, in terms of the integration time, $t \propto T_{system}^2$. Any axion haloscope experiment that does not intend to take data for hundreds of years must therefore work towards minimizing temperature.

3.3 ADMX hardware

ADMX has been around since the early 1990s, with major upgrades ushering in new “phases” of the experiment. Phase 0 demonstrated that an axion haloscope could indeed achieve good enough sensitivity to detect or rule out plausible QCD axions within a given mass range. Phase I incorporated new amplifier technology that allowed even better sensitivity to

axions than was previously achieved. For the details of those experiments see, for example, references [20] or [21].

This thesis came out of the third incarnation of the experiment, ADMX Phase II. One of the major features of Phase II is the addition of a dilution refrigerator, which will allow ADMX to run far cooler than ever before and decrease T_{system} by a factor of 40 relative to the previous version of the experiment. This temperature reduction has no effect on the magnet cryostat, which has not changed since Phase 0. However, the “insert,” which sits in the bore of the cryostat and is the object that actually looks for axion signals, required a complete redesign and rebuild. Figure 3.3 presents a cutaway model of the upgraded ADMX Phase II experiment. The insert is organized by temperature into three general areas: the “Cold space,” “Reservoir,” and “Warm space.” The cold space contains the most crucial component of the experiment, the resonant cavity.

3.3.1 Resonant cavity

The ADMX cavity consists of a stainless steel barrel and two stainless steel endcaps. The endcaps are secured to the barrel with 96 bolts each. They form a low-resistance seal through tight contact with a knife-edged surface on the barrel. This is essential to prevent ohmic loss which would degrade its quality factor Q , and thus diminish axion signal power. The interior surfaces of the barrel and endcaps have a very smooth OFHC copper plating, which is vital to achieving a high Q . Typical Q values of the ADMX cavity are of order 10^5 . The cavity has an inner diameter of 16.5”, and a length of 39.9”.

The cavity is made tunable by the inclusion of two 1” diameter, 39.4” long rods that are loosely affixed to the cavity endcaps. These rods can rotate about axes offset from the cavity axis, which allows them to be radially closer or farther from the center of the cavity. The rods, just as the rest of the cavity interior, are made of stainless steel and plated with OFHC copper. They attach to the cavity endcaps via alumina armatures. The total cavity




Figure 3.3: Cutaway model of ADMX Phase II. The cryostat, shown in dark gray, is roughly 1.5 m in diameter and 3.5 m tall. Its superconducting magnet, shown in blue, can reach fields above 8 T. The insert is presented in a light gray color and sits in the bore of the cryostat. It can be broken up into three general sections: the Cold space, the Reservoir, and the Warm space. The Cold space includes the cavity, the 1K plate, and 1K pots, along with other components. The 4K space is comprised of a helium reservoir and various components that heat sink to it. The Warm space extends above the reservoir, and includes baffles, supporting structures, and the top plate, which is the only visible part of the experiment when ADMX is taking data.

volume (less the tuning rod volume) is approximately 135 L. Figure 3.4 shows the inside of the empty cavity and a model with the tuning rods included. Stepper motors at the very top of the insert provide the torque that rotates the tuning rods. Each of the two motors drives a G10 rod that connects to a gear box which has a net 19,600:1 gear reduction. The gear box output motion rotates another G10 rod which is directly connected to one of the tuning rod armatures. The extreme gear reduction allows for control of the tuning rod rotation to better than 1 micro radian and reduces backlash on the relatively long G10 rod. The motion of the rods within the cavity is fully characterized by one parameter, θ_n , which measures the position of the rod along its rotational path. The subscript n serves only to differentiate the rods and is set to either 1 or 2. An angle of 0° corresponds to a rod near the wall of the cavity, while at 180° the rod is near the center of the cavity.

Signals from the cavity are detected by antennas going through “major” antenna ports in top cavity endcap. The antennas are coupled to the two lowest axial transverse magnetic (TM_{0n0}) modes of the cavity, TM_{010} and TM_{020} . These are used because they maximize the form factor f_{mnp} . In fact, any modes other than those of the form TM_{0n0} give $f_{mnp} = 0$ for an ideal right-circular cylinder. On the other hand, $f_{010} \approx 0.5$ and $f_{020} \approx 0.1$. A detailed discussion of cavity modes is given in Chapter 4: 3D cavity simulations. The antennas themselves are coaxial lines terminated in stubs that can be moved vertically into or out of the cavity to maximize their coupling to cavity modes. Motion of these antennas is achieved in much the same way as with the tuning rods. There are again stepper motors at the insert top that drive G10 rods which themselves drive gears in a gear box. The output of these gear boxes is linear vertical motion, with a resolution of better than $1 \mu m$. Figure 3.5 shows a picture of a tuning rod gear box, as well as its CAD model.

The degree to which the antennas are coupled to the TM_{010} and TM_{020} modes is determined by a reflection measurement. This is obtained by sending power down one of the antennas and measuring how much of the incident power is reflected back. When the antenna




Figure 3.4: The left picture shows the ADMX cavity interior without tuning rods. The walls and endcaps of the cavity interior are plated with 5 mil OFHC copper. The endcaps have through holes to allow for antennas and tuning rod attachments. The right picture is a model of the cavity with rods installed. Rotating the tuning rods from a position where both are near the cavity wall to where both are near the cavity center increases the cavity resonant frequency by about 80%.




Figure 3.5: Tuning rod gear box. At left is a picture of the gear box installed in the insert, while a CAD model of the box is shown at right. Two circular gears are used along with a worm gear to provide a high gear reduction of 19,600:1, as well as strongly damp backlash of the G10 rods driving the gears. A potentiometer³ at the top of the gear box measures the location of the tuning rod driven by that box. A circulator⁴ is used to ensure that no unwanted power travels back into the cavity.

is uncoupled, all of the power should be reflected. We define it to be critically coupled when the reflected power at the $TM_{01(2)0}$ frequency is at -30 dB relative to the power sent in.

The frequency of a certain mode is predicted by simulations, the details of which are again given in Chapter 4. Modes are experimentally verified by performing a transmission measurement. This is done by injecting power into the cavity through the “weak” port antenna, another stub-terminated coaxial cable located on the bottom endcap, and measuring the cavity response through one of the major-port antennas. The resonant frequency of the cavity for a given mode will show up as a frequency response in the transmission. A schematic of how the reflection and transmission measurements are made is shown in figure 3.6.

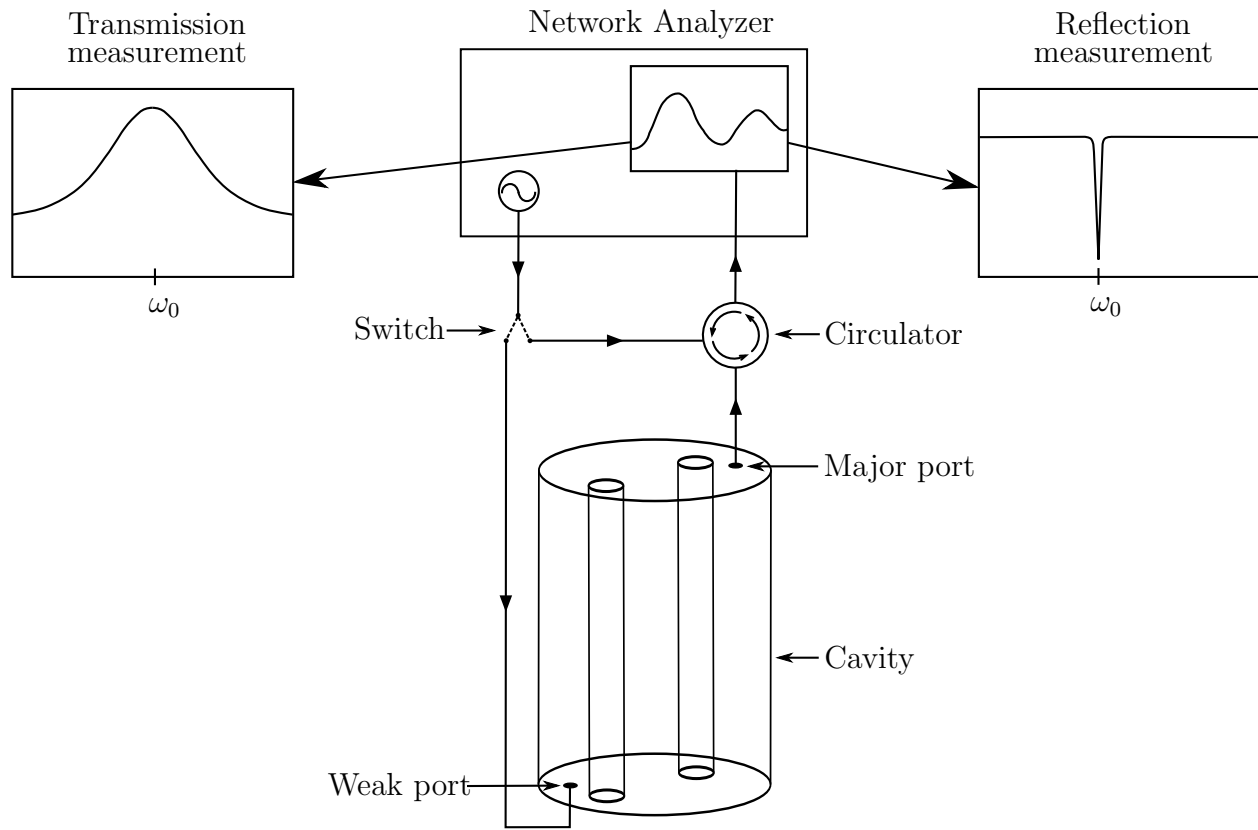


Figure 3.6: Schematic of transmission and reflection measurements. The switch connects the network analyzer (NA) to either the weak port or to one of the major ports. Only one major port is shown to make the diagram more clear. For a transmission measurement, the NA injects power into the cavity through the weak port and the outgoing power is picked up by the major port antenna. The expected response for such a damped, driven oscillator is a Lorentzian function. To make a reflection measurement the NA sends power through the major port antenna. When the antenna is decoupled from the cavity there is an impedance mismatch between the antenna and the cavity, resulting in all of the power being reflected back up the line. When the antenna is critically coupled, a -30 dB trough is observed in the reflection spectrum due to power being absorbed within the cavity.

The cavity blackbody radiation is one of the two dominant noise sources in ADMX, with amplifier noise being the other major contributor. Since minimizing this noise is essential to achieving a high sensitivity to axions, ADMX uses an ultra-low-noise “SQUID” amplifier as well as low-noise HEMT amplifiers.

3.3.2 SQUID amplifier

The most important amplifier in a serial chain of such devices is the first one. Subsequent amplifiers’ noise contributions are reduced by the product of all of the previous amplifiers’ gains. For example, the second amplifier in an amplifier chain adds noise $P_2 \approx P_1/G_1$ if both amplifiers are at the same physical temperature, where $P_{1(2)}$ is the noise power of amplifier 1(2), and G_1 is the first amplifier’s gain. In ADMX a Superconducting QUantum Interference Device (SQUID) amplifier is the first in the amplifier chain for the TM_{010} channel. Below I discuss the basic properties of SQUIDS, and I make no attempt to describe the devices in detail. Those interested in such details may consult, for example, reference [22].

A dc SQUID, the kind used in ADMX’s SQUID amplifier, is essentially a loop of superconducting material broken in two places with weak links known as “Josephson junctions.” Figure 3.7 shows a simple circuit diagram for this kind of device, where “X” is a Josephson junction.

A biasing current I is applied to the SQUID so that the Josephson junctions act as resistive elements, and therefore a voltage develops across the device. If an external magnetic field threads the SQUID, a counter current will develop in the loop such that the induced field cancels the applied field. The net magnetic flux through the SQUID loop must be an integer number of flux quanta Φ_0 , so in general

$$\Phi_S = \min(n\Phi_0 - \Phi_B), \quad (3.12)$$

where Φ_S is the flux developed by the counter current, n is an integer, Φ_B is the flux




Figure 3.7: dc SQUID circuit diagram. With a proper biasing current I applied, the Josephson junctions in the loop, represented by the two Xs, become resistive. The voltage V across the superconducting loop halves then varies as a function of magnetic flux Φ through the loop.

resulting from the applied magnetic field, and by $\min(n\Phi_0 - \Phi_B)$ I mean the minimum possible magnitude of $n\Phi_0 - \Phi_B$. The solution for Φ_S as a function of Φ_b is a sinusoid, as shown in figure 3.8. With the SQUID properly flux biased to ensure

$$\Phi_S = \Phi_0(n/2 + 1/4), \quad (3.13)$$

such that Φ_S is at the steepest part in its response, the dc SQUID is made a very sensitive magnetometer.

To utilize the dc SQUID as an amplifier, one needs only to inductively couple it to a signal current. The simplest way this can be realized is simply by inductively coupling the input coil to the SQUID; however, such a method is not effective above roughly 100 MHz, for at these higher frequencies, parasitic capacitance between the coil and the SQUID severely degrades the gain [23]. Therefore, the ADMX amplifier is made by connecting the SQUID to




Figure 3.8: dc SQUID V - Φ_B curve. The voltage induced across the SQUID loop varies with the applied magnetic flux as a sine function. If the SQUID is flux biased such that $V \propto \Phi_0(n/2 + 1/4)$ where n is an integer, any further change in Φ will produce a large change in V .

the input coil ground and floating one of the coil ends, as shown in figure 3.9. This turns the coil into a stripline resonator, which is much more adept at handling the ~ 500 MHz signals coming out of the cavity than a simple coil of wire would be. The physical dimensions of the resonator determine the center frequency and bandwidth of the amplifier. This device is known as microstrip SQUID amplifier (MSA).

All amplifiers add some electrical noise to a system. ADMX uses an MSA because it is able to provide good gain while adding very little noise power. Figure 3.11 shows equivalent noise temperature versus physical temperature for a typical MSA. The noise temperature decreases as half the physical temperature until the MSA cools to around 100 mK, at which point the noise approaches to the quantum limit of roughly 30 mK at ~ 600 MHz. Accordingly, ADMX keeps this amplifier as cold as possible by thermally sinking it to the coldest part of




Figure 3.9: ADMX SQUID amplifier cartoon. A grounded niobium washer serves as the base for the device. Current flowing through the stripline resonator induces a counter current through the SQUID and generates a large voltage at the output. Resistive shunts are used to prevent hysteresis.

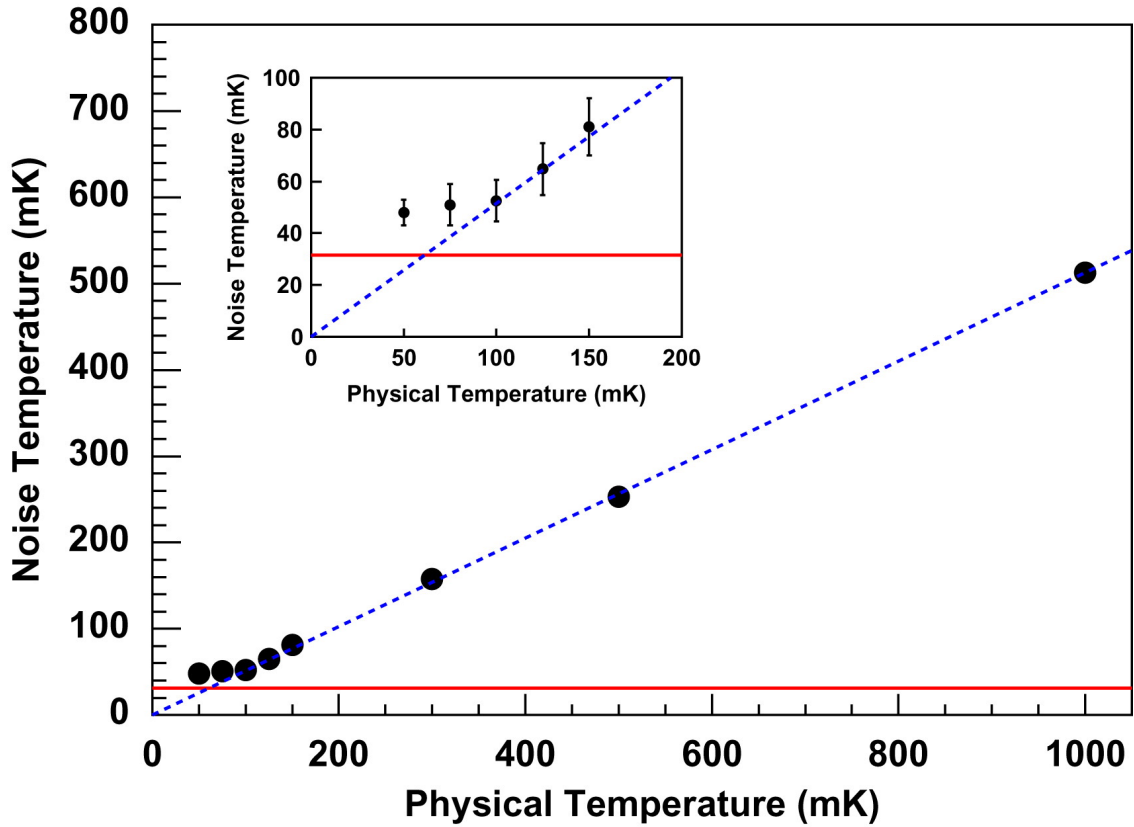


Figure 3.10: MSA noise temperature versus physical temperature. The center frequency of this device is approximately 600 MHz. Black dots are data, the blue dashed line is a linear fit, and the red line is the minimum possible noise as dictated by quantum mechanics. The inset shows a close-up view of the 0 - 200 mK physical temperature part of the plot. Noise temperature decreases linearly as half the physical temperature until about 100 mK physical, at which point the noise begins to asymptote to the quantum-limited noise floor. [24]

the experiment, the resonant cavity. The thermal bridge that links the cavity and the MSA is known as the “squidadel,” shown in figure 3.10. The squidadel is made of OFHC copper and serves not only to thermally sink the MSA and ancillary electronic components (circulators, isolators⁵, and switches⁶) to the cavity, but it also provides structural support for these elements. Figure 3.12 shows a schematic of the circuit used to measure the SQUID amplifier’s noise temperature, which demonstrates why these electronic components are necessary in the experiment.

Within the squidadel, the MSA has two principle enclosures. The first is a lead cylindrical box that the MSA is packaged in, which both supports the printed circuit board (PCB) on which the SQUID is mounted and helps keep stray magnetic fields from reaching the MSA. The device so packaged comes to ADMX from John Clarke’s SQUID group at UC Berkeley. Immediately exterior to the lead box is a cryoperm shield that provides further passive magnetic shielding for the SQUID, with factory-specified theoretical field attenuation of better than 10^6 .

3.3.3 HEMT amplifiers

An MSA can achieve good gain while adding very little noise, though it becomes harder to deploy for frequencies above ~ 1 GHz. Therefore at higher frequencies and temperatures ADMX uses High Electron Mobility Transistor (HEMT) amplifiers, which are built by from Richard Bradley at the National Radio Astronomy Observatory (NRAO).

In general, a single HEMT amplifier is not going to easily have a 50Ω input impedance, which means it would be mismatched from all of the 50Ω RF cables carrying the axion signal. For this reason AMDX employs HEMTs in a balanced design, where two HEMTs

⁵An isolator is a two-port device which transmits (microwave) power in only one direction.

⁶An electronic switch is a multi-port device which takes power through an input port and delivers it to one of several output ports.




Figure 3.11: At left is the 3D rendering of the squidadel. The central structure has two circular pads which can each hold one packaged MSA. Once the MSA is installed, a cryoperm shield is fit around the central region, as can be seen in the middle picture. The square pads, two of which are attached to each of the 4 arms, hold circulators, isolators, and switches. The picture to the right shows the squidadel with all components installed. The entire structure is made of OFHC copper, which allows for good heat transfer from every electrical component down the central copper coldfinger. Another cylindrical coldfinger, not shown here, connects the squidadel with the cavity.

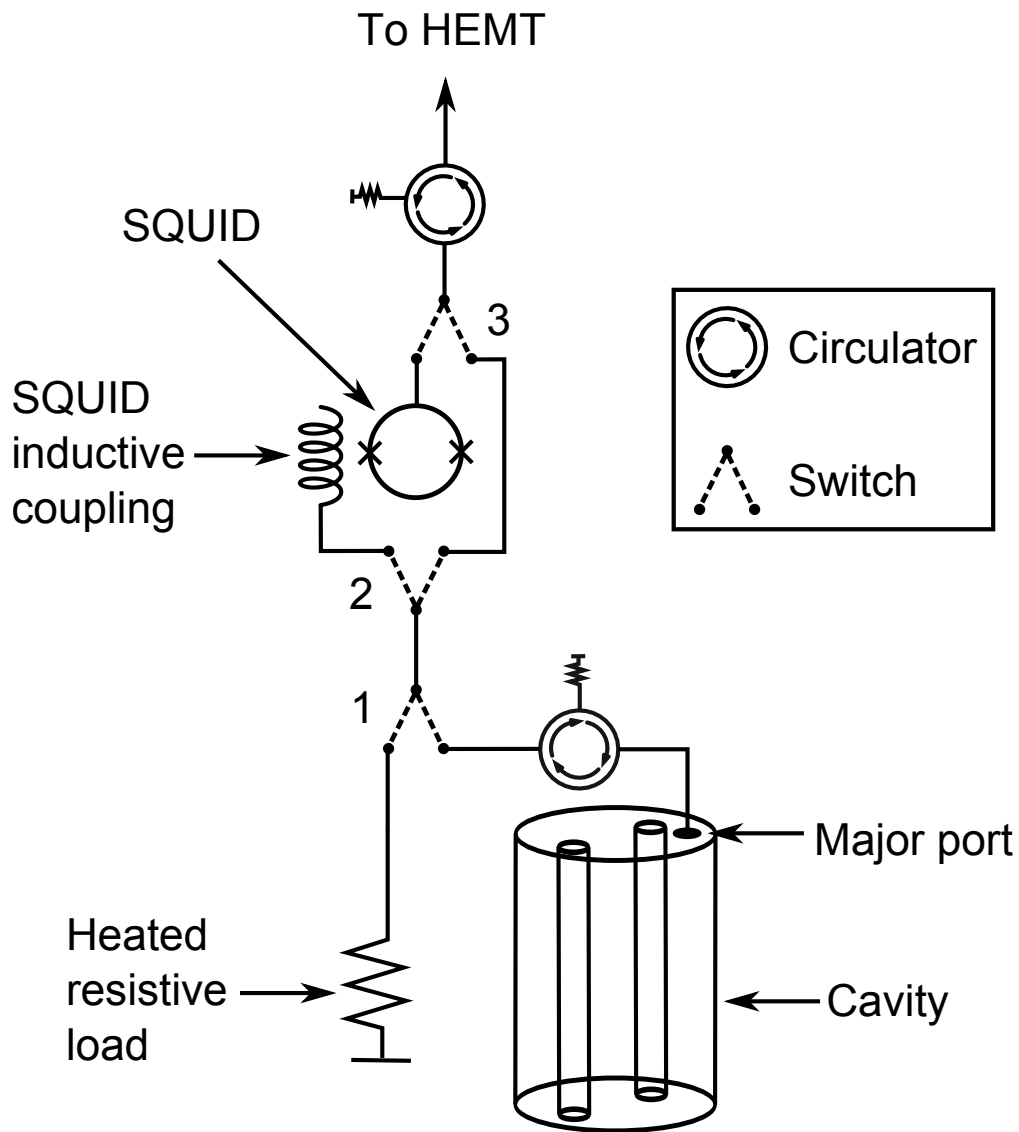


Figure 3.12: SQUID amplifier noise temperature circuit schematic. To measure the noise temperature that the SQUID amplifier produces, power spectra must be taken at different system temperatures. The noise temperature is then solved for through the equation $P = G(T_{\text{physical}} + T_{\text{noise}})$, where P is the average power read out, G is the gain of the system, and T_{noise} is the noise of the SQUID amplifier. To do this in situ, first switches 1 and 2 connect the resistive load to the SQUID inductive coupling. The load is heated to a known temperature, and the resulting power drives the SQUID amplifier. Switch 3 connects the SQUID amplifier to the system, which sends the amplified signal up through the rest of the receiver chain. Circulators are used with one port terminated in a $50\ \Omega$ resistor to prevent signals from reflecting back into the system. Switch 1 is only ever used for the in-situ noise measurement test, while switches 2 and 3 are used to connect the SQUID amplifier to the receiver chain or to bypass the amplifier.

work in tandem with two 90° hybrids⁷, as depicted in figure 3.13. The resulting amplifier is then well matched to $50\ \Omega$ and thus experiences almost no signal loss.

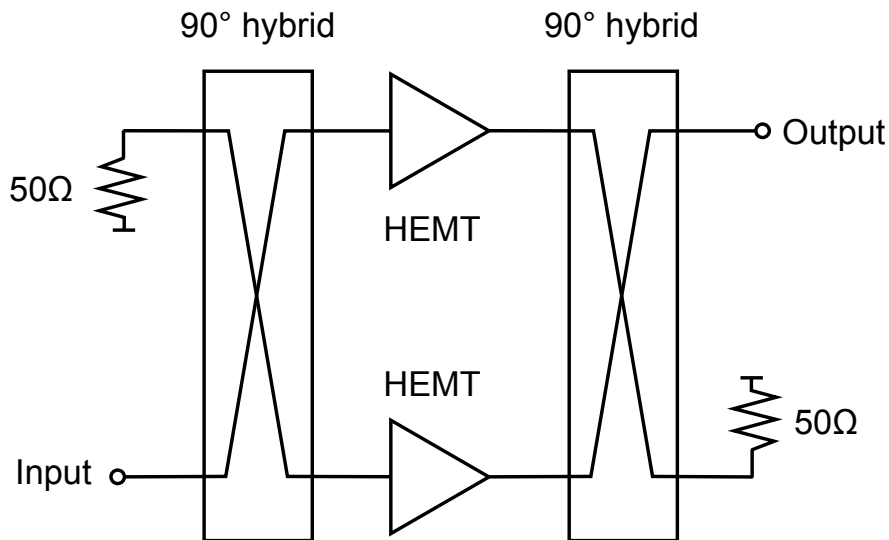


Figure 3.13: Schematic of a balanced HEMT amplifier. A single HEMT amplifier is not well-matched to incoming and outgoing RF lines over its full frequency band, but such a balanced design will be. Balanced HEMT amplifiers are used in both frequency channels in ADMX.

The noise temperature of HEMT amplifiers does not depend strongly on the physical temperature of the devices below some threshold, which is $\sim 30\text{K}$ for ADMX's HEMTs. As with all amplifiers, however, noise temperature does scale with frequency. In the TM_{020} range of $\sim 1.0 - 1.5\ \text{GHz}$ where ADMX takes data, the HEMTs contribute 4K noise temperature.

3.3.4 Cold space

The cavity makes up the largest part of the Cold space, which itself comprises roughly the bottom third of the insert. The other main components in this region are the 1K plate, the

⁷A 90° hybrid coupler evenly splits an incoming signal into two output signals that are 90° phase shifted with respect to one another.

1K pot, and the 4K and 1K radiation shields. The 1K plate is a large oxygen-free high-purity (OFHC) copper baffle that provides support and thermal sinking for all of the components in the Cold space. The 1K plate itself is cooled by a large 1K pot, which is sunk to the plate with a solid OFHC copper block. The cavity is in turn sunk to the 1K plate with 10 OFHC copper braids, each of approximately 0.625 cm^2 cross sectional area. The cavity hangs from three thin-walled stainless steel supports bolted to the 1K plate. Figure 3.14 presents a picture of the Cold space.

In understanding the design of the Cold space, it is important to realize that it is made with the dilution refrigerator in mind. So, for example, the stainless steel rods connecting the 1K plate and the cavity transfer very little heat between these two regions, yet a slew of thermal braids then sink them together. With a dilution refrigerator in place those braids would not be part of the experiment, and only a very little amount of heat would be transferred from the 1K plate down to the cavity through the support rods.

Radiation shields are employed between regions that should not be radiatively connected. Physically these shields can be any opaque material that does not allow the passage of photons from one area to the other. There are two radiation shields in the Cold space, one thermally sunk to 4K and another to the 1K plate. The 4K shield blocks 30K blackbody photons from the cryostat bore from reaching the 1K area, and the 1K shield blocks the 4K photons from warming up the 100 mK region. As with many other components in the Cold space, these shields are made of OFHC copper.

For the data in this thesis, the 1K pot has a typical operating temperature of 1.4K which keeps various parts the Cold space around 1.5K - 2.0K. The graph in figure 3.15 shows typical temperatures for various components in the Cold space.




Figure 3.14: Insert Cold space. ADMX’s resonant cavity forms the bottom of the cold space. It is attached to the 1K plate by way of three thin-walled stainless steel support tubes. The two large stainless steel boxes seen hanging from the bottom of the 1K plate are gear boxes. The left box control the rotation of one of the tuning rods in the cavity, whereas the right box is used to vertically translate the TM_{010} antenna. Above the 1K plate is a mass of wires, breakout boxes, sensors, and copper braid. The stainless steel barrel of the reservoir is seen at the very top of the picture.




Figure 3.15: Cold space temperatures. The 1K pot cools the entire Cold space, so it governs the temperatures of all components in the region. With the 1K pot at 1.4K, temperatures in the rest of the Cold space stay between about 1.5K and 1.7K.

3.3.5 Reservoir

Immediately above the Cold space is a large stainless steel cylindrical tank known as the “reservoir.” The reservoir has a volume of 165 L, and during data taking it is kept 35% to 100% filled with ^4He . The reservoir serves many purposes, not least of which is supporting the 1K plate, which hangs from the reservoir by three thin-walled stainless steel support rods. It also provides a steady flow of liquid helium to the 1K pot below, and will in the future feed a second 1K pot with the integration of the dilution refrigerator. In addition, many of the HEMT amplifiers ADMX uses are thermally sunk to this stage.

Perhaps most importantly, the reservoir provides cooling for the superconducting “bucking coil” that serves as the active magnetic shielding for ADMX’s MSA. The MSA is sensitive to single quanta of magnetic flux, and the passive shielding provided by the aforementioned cryoperm shield does not attenuate the 6 - 8 T field threading the SQUID to below one quantum of magnetic flux. Without the bucking coil’s active shielding the MSA would at best stop functioning as an amplifier, and at worst be destroyed. The bucking coil is meter-tall bobbin that holds two loops of superconducting wire. The lower “field cancellation” loop serves to counteract the magnetic field produced by the main cryostat magnet at the geometric center of the bucking coil. The second “counterwound” loop generates a force on the bucking coil opposite of that which the field cancellation loop produces, thus resulting in no net force on the bucking coil spool piece. This active field-canceling component is critical for the SQUID amplifier, as passive shielding alone would get saturated and become useless at the multi-Tesla field generated by the main magnet. Figure 3.16 shows the reservoir and bucking coil in various stages of production, and the graph in figure 3.17 shows data of the bucking coil successfully canceling the magnetic field in the region of the MSA.

The top of the reservoir varies in temperature from 4 K when the reservoir is fully filled with liquid helium, to around 10 K when it is filled 35% or less. This is an ideal temperature range in which to use HEMT amplifiers since in this range they operate at their minimum



Figure 3.16: At left, a picture of the bucking coil being installed within the reservoir. The bucking coil fits around the reservoir bore tube and bolts to the reservoir bottom. At right, a picture of the reservoir as it is being machined. The large central bore is to accommodate the squidadel with the MSA and other electronics. The surrounding through tubes run various plumbing pipes and wires that go from warmer stages down the Cold space.

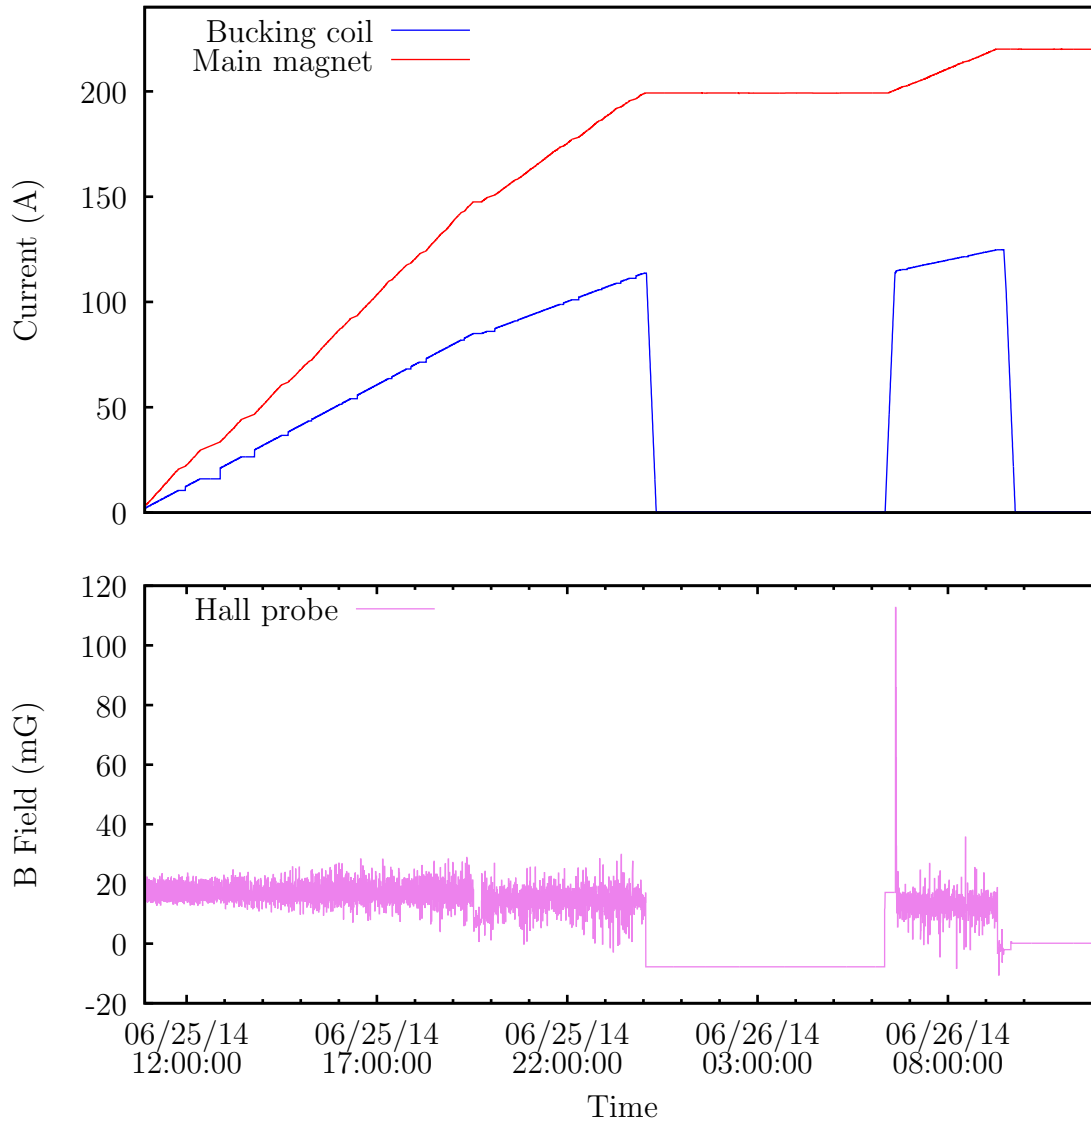


Figure 3.17: Bucking coil field cancellation during a magnet ramp⁹. The upper plot shows the currents in both the bucking coil and the main magnet during a ramp, while the lower plot gives the magnetic field in the SQUID mace region. While the main magnet ramps at a constant voltage, the magnetic field in the SQUID mace region remains at tens of miliGauss. The hall probe is located just outside the cryoperm shielding in the SQUID mace, so the MSA itself has two further passive shields protecting it from magnetic field intrusion. The bucking coil is automated to get short-lived voltage injections when the hall probe reads too-high a magnetic field. This ramp was briefly stopped during the night, and the bucking coil current is only read while the ramp is ongoing.

noise temperature. Therefore all cold (typically kept at 4 K - 10 K) amplifiers, save for the MSA, in the TM_{010} channel affix to the top of and are thermally sunk to the reservoir.

3.3.6 Warm space

The Warm space represents the top third of the insert, and is comprised of all of the components above the reservoir. The “top plate” is the uppermost part of the Warm space, and is the only visible part of the insert during data taking. The top plate can attach to a crane which allows the entire insert to be moved into and out of the cryostat. Figure 3.18 shows the Warm space while the insert is secured in the clean room, and figure 3.19 shows the insert completely assembled as it is lowered into the cryostat. Four radiation baffles layered with aluminized mylar separate the 300K top plate from the reservoir. The second baffle from the top plate is actively cooled with a dedicated pulse tube cryocooler¹⁰, and as such it stays at a constant 40 K. A radiation shield hangs from and is thermally sunk to the 40K plate, and a thin nylon shield blocks off the rest of the warm space.

3.3.7 Receiver chain

The receiver chain is the network of analog and electronic devices that takes signals from the cavity and ultimately saves digital data to disk. The starting point of this chain for both frequency channels is the coaxial stub antenna that connects to the 1st stage amplifiers.

As mentioned earlier, the first amplifier in the TM_{010} channel is an MSA while for TM_{020} a pair HEMTs in a balanced configuration is used. The gain of each of these devices is 15 dB - 20 dB, meaning that the 1.5 K noise of the MSA is amplified to ~ 100 K, while the 4 K noise of the HEMT is amplified to ~ 300 K. The room-temperature electronics have noise temperatures of ~ 100 K, which would add significantly to the overall noise of the receiver

¹⁰A pulse tube cryocooler is a cryogen-free refrigerator that works by oscillating the pressure of a gas within a fixed volume, where the gas can exchange heat with both a warm and cold reservoir. For details on how these refrigerators work, see, for example, reference [25]




Figure 3.18: Insert Warm space. Two copper and two aluminum baffles separate the 300 K top plate at the top of the Warm space from the reservoir at its bottom. All of the baffles are covered with aluminized mylar to improve their reflectivity. Other than baffles, this stage of the insert is replete with structural support tubes, plumbing for cryogenic gases, and various data cables and wires. When the insert is ready to be inserted into the bore, an OFHC copper radiation shield and a nylon shield cover the Warm space.




Figure 3.19: The insert being lowered into the cryostat. The 4K radiation shield, reservoir, 40K radiation shield, and nylon shield are visible below the top plate. Each of these elements has a 1/8" radial clearance to the inner diameter of the bore.

chain. Therefore between the 1st stage amplifiers and the ones at room temperature there are a number more HEMTs that serve to take the 1st stage amplifier noise up to over a billion Kelvin, so as to completely dwarf the added noise of the room temperature amplifiers. After the various amplification stages, RF cables carry the signal out of the insert to the “warm” room-temperature electronics.

Once at room temperature, the microwave-frequency signals from the cavity are further amplified, and then mixed down to an intermediate frequency of 10.7 MHz¹¹. Signals are then again amplified, and finally go through two cascaded bandpass filters, with bandwidths of 2 MHz for the first filter and 150 kHz for the second. At this point, the raw data is digitized with a Signatec PX1500-2 board operating at 200 Megasamples per second (MSPS), and hereafter software takes over the data processing. Figure 3.20 shows a diagram of the receiver chain from the cavity up to the digitizer.

All of the warm-electronics machines are bolted into racks which are located just a few feet from the cryostat. Figure 3.21 shows a picture of the computers and racks housing the various receiver-chain machines.

3.4 ADMX DAQ software

The first step on the software side of the ADMX DAQ is to downsample the data. Incoming voltages are digitized at 200 MSPS, which is far more than needed. The Nyquist frequency for a 10.7 MHz signal is 21.4 MHz, or roughly a tenth of the frequency of data coming from the digitizer. As extra samples would only use up more computer resources but otherwise provide no benefit, every 8 consecutive voltages are averaged, which downsamples the 200 MHz signal to 25 MHz. The voltage time series then undergoes a fast Fourier transform (FFT) to produce a voltage spectrum. The FFT code is written in C and utilizes the

¹¹Though the axion mass is ~ 1 GHz, I expect its linewidth to be only ~ 1 kHz. It is therefore advantageous to mix the axion signal down to a lower frequency, which is far easier to transmit than a high-frequency signal. 10.7 MHz is a common RF mixing frequency.

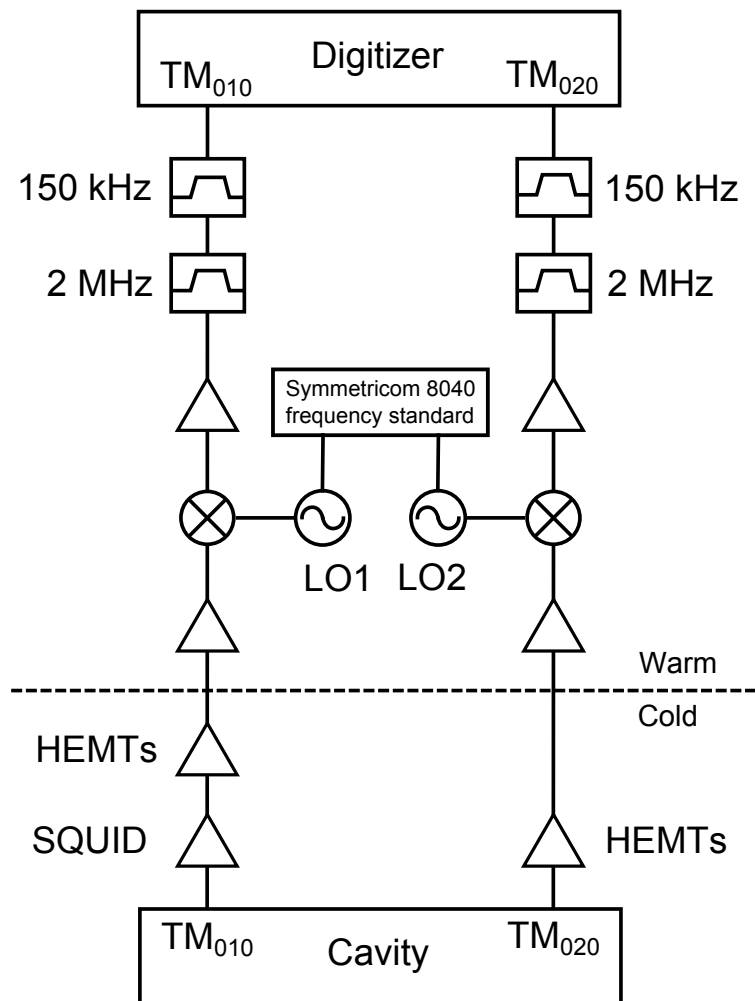


Figure 3.20: Receiver chain. The TM_{010} and TM_{020} channels use almost identical hardware in their receiver chains. The only differences are that TM_{010} channel has an additional SQUID amplifier in its chain, and the channels employ different mixers as appropriate for their particular frequency range.

FFTW3 library [26]. The final steps are to digitally filter and mix the voltage spectrum. First, a 25 kHz window centered at 10.7 MHz is cropped out from the full voltage spectrum, and then all of the voltages are squared. The resulting (arbitrarily scaled) power spectrum is saved to a hard disk as well as to an SQL database. Chapter 5 discusses analysis of these




Figure 3.21: Electronics rack during data-taking. All of the machines and components that constitute the receiver chain are housed in two racks which are located several feet from the cryostat. Two desktop computers display the various temperatures and pressure from the insert and cryostat, as well as instantaneous power spectra and other information from cavity sweeps. A large TV is likewise set to display this information.

data.

3.5 Data taking cadence

The main practical limitation to the ADMX experiment is time: in principle, given enough time, ADMX could scan the full axion mass range and either detect these particles or rule them out. In reality, with current technologies any one version of the experiment can only

cover some subset of the full range. Therefore in data taking there are two possible approaches, illustrated in figure 3.22, which establish a data-taking cadence of the experiment. The first of these is to choose some coupling strength and take enough data to achieve that




Figure 3.22: Illustration of two possible data-taking schemes. In the first scheme, a frequency range is specified and the tuning rods are moved in such a way as to cover the full frequency range. The scan speed can be variable, and typically depends on how much data-taking time is available. That time, then, determines the sensitivity to axions that can be achieved. In the second scheme, an SNR is chosen and little bits of frequency space are scanned over repeatedly until that SNR is reached. Using this method, the available time determines how large a frequency range can be covered.

coupling for all scanned frequencies. In this case the time spent taking data determines frequency range covered. The second approach, opted for during the 2014 data runs, is to assume a specific frequency range and have available time determine the axion-to-photon coupling.

3.5.1 2014 operations

The decision to focus on a particular frequency range means that ADMX sweeps back and forth across that range for however long data taking goes on. Digitization for each scan takes 100 seconds and between successive scans the tuning rods are moved to change the

resonant frequency by 20 - 200 Hz. The frequency change between scans depends on the rate of rotation of the tuning rods, which varies between data runs. Between digitizations the cavity is swept to measure its Q , and the antennas are recoupled (inserted further into the cavity or slightly pulled out of it) if they are no longer critically coupled to the desired cavity modes. Measuring Q and recoupling the antenna takes roughly 25 seconds, so the downtime between scans is approximately 20% of total data-taking time.

Chapter 4

3D CAVITY SIMULATIONS

4.1 *Introduction*

In this section I talk about the 3D cavity simulations that are an integral part of ADMX Phase II. First, I explain why simulations are necessary for the experiment. Then, I briefly describe the software package I use, after which I discuss the model of the cavity. I finish this chapter by discussing the simulation results, focusing on how they guide data taking and also how they give insight into a previously-unconsidered capacitive effect of the tuning rods.

4.2 *Need for cavity studies*

Previous versions of ADMX have made extensive use of simulation studies of the cavity, and this incarnation of the experiment is no different. Such studies are crucial to ADMX for two main reasons. First, although the form factor as defined by equation 3.2 can be simply calculated for an ideal right-circular cylinder, the addition of tuning rods in the cavity makes an analytical approach largely intractable. The form factor is integral to the data analysis presented in the next chapter, so it must be known to relatively high precision. The second reason for the necessity of simulations is that the cavity mode structure is in general quite complicated, so simulations are needed to establish the resonant frequency of the TM_{0n0} modes for given tuning rod positions. Furthermore, the resonant frequency of cavity modes moves in response to motion of the tuning rods, so it is necessary to track modes as they change in frequency. This is accomplished by making “mode maps,” which measure the resonant frequency of different modes at given rod positions. Simulations are necessary to

pick out which of these modes correspond to the TM_{0n0} set. Figures 4.1 and 4.2 show mode maps of actual data for frequencies around the TM_{010} and TM_{020} modes. These mode maps are made by moving the tuning rods to a certain position and sweeping the cavity within a wide frequency range. If the signal exhibits a peak above the noise, indicating a cavity mode, the center frequency of that peak is recorded, and subsequently appears as a dot in the mode map. Scanning over all tuning rod positions of interest generates the full mode map.

4.2.1 Motivation for 3D modeling

ADMX phases 0 and 1 used custom-written code to make 2D models of the cavity with tuning rods, from which the form factor and TM_{010} resonant frequency could be extracted. These models worked well because pure TM_{0n0} modes are independent of the axial dimension in a right-circular cylinder. Physically, of course, the tuning rods do not run the length of the cavity, but instead have an approximately 1/4" gap at each end. These small gaps were initially considered inconsequential to both the frequency and the form factor.

In principle the same modeling could be done in this version of the experiment for TM_{020} as was done for TM_{010} in the past. However, the mode maps of the two different frequency regions make it abundantly clear that there are far more mode crossings that need to be dealt with in TM_{020} than in TM_{010} . The slight differences to resonant frequency and form factor that might be introduced by the gaps between the tuning rods and cavity endplates need to be studied in detail so that the TM_{020} mode can be scanned with any confidence. Furthermore, if a high-accuracy 3D mode map can be made for TM_{020} , it might reveal a path for the tuning rods which avoids mode crossings¹ yet still covers the full TM_{020} tuning range.

¹A mode crossing in the cavity occurs when two modes have approximately equal resonant frequencies for a given tuning rod position. Even if the correct TM_{0n0} mode can be picked out, the form factor in these regions tends to be very poor.




Figure 4.1: Cavity mode map around TM_{010} . The TM_{010} mode is present, and is identified in the “3D model” section below. Also present are other TM, TE, and TEM modes. Simulations are necessary to differentiate TM_{0n0} modes from the rest. Mode map courtesy of C. Boutan.




Figure 4.2: Cavity mode map around TM_{020} . The mode structure at these higher frequencies is considerably more complicated than it is at the lower TM_{010} frequencies. Therefore simulations are even more critical to the second channel than they are to the first. Mode map courtesy of C. Boutan.

4.3 Methodology

The software package I use to do 3D cavity analysis is called COMSOL. This is a finite-element analysis program that is well-suited for a wide array of physics simulation problems. COMSOL has a number of different physics modules, each with their own particular set of variables that can be solved for. I primarily use the RF module, which allows solving for variables like eigenfrequencies, electric field components, and so on. Conducting a cavity simulation involves building a model, applying proper boundary conditions, specifying the desired outputs, and letting the program run.

Most often I run COMSOL in parametric mode, where I provide some set of parameters and the program solves for my desired output variables, each parameter set specifying a new model. I use three parameters, two for specifying the angles of the tuning rods and another for the center frequency around which the calculation should start. To get decent resolution I use 4° steps and traverse a full 360° with each rod², which gives me a total of 8100 parameter sets. For each set COMSOL builds a new model and does a new calculation. The results of these calculations are output to a text file, which I then analyze using C++.

For all its robustness, COMSOL does not feature an easy way to track modes. I cannot specify that I want to look only at TM modes, for example. Instead, I save the output of a set number of modes around a center frequency, and in my analysis of the COMSOL data I calculate the form factor for each of those frequencies. I consider TM_{0n0} modes to be those which have the highest form factor.³

²There is a degeneracy in the data in doing this, and I could indeed get the full 3D mode map by traversing 180° with one rod and 360° with the other. I do a full revolution of both rods to have another check on the data - if my data does not feature a plane of symmetry I know I've made a mistake either in performing the simulation or analyzing the results.

³Due to the tuning rods in the cavity, my algorithm will not actually ever find a "pure" TM_{0n0} mode. Instead the mode my algorithm selects will be a mix of mostly TM_{0n0} plus other nearby modes; the form factor is the most important quantity, however, so it does not matter that I have a mixed mode versus a pure mode.

4.3.1 3D model

The model I use in COMSOL is shown in figure 4.3. It is a much-simplified design where



Figure 4.3: COMSOL cavity model. The top endplate and one quarter of the barrel are hidden from view. The cavity itself is represented as a right circular cylinder, as are the two tuning rods. The rods have a 1/2" gap to the top and bottom cavity endplates.

the cavity and the two tuning rods are appropriately-sized cylinders. Each of these cylinder have perfect electrical conductor surface boundary conditions. I use this simplified model instead of the exact cavity model because it takes COMSOL only 20 seconds to complete one study with the three cylinders, whereas that number increases to several minutes using

the full cavity model. With 8100 parametric steps, this translates to a simulation taking several days versus one running for a few weeks.

This simple model is still robust enough to generate results that agree well with predictions and with experimental data. The TM_{010} form factor and frequency that I get from running a simulation of a bare cavity with no rods nearly exactly match the theoretical values of $f_{010} = 0.698$ and $\omega_0 = 546$ Mhz. Furthermore, the mode visibly looks precisely as it should, as is evident from figure 4.4. Also, with COMSOL, I get form factors and frequencies for the cavity plus tuning rods that mirror the same variables obtained through Microwave Studio by Gianpaolo Carosi, one of ADMX’s collaborators.

Agreement with data, of course, is even more convincing than agreement with other software. In figure 4.5 I overlay a simulated mode map onto the cavity mode map near TM_{010} frequencies from figure 4.1.

4.4 *Simulation results*

My simulations provide a mix of both expected and unexpected results. The former category includes maps of form factors and frequencies that are used to guide data taking in both the TM_{010} and TM_{020} channels. The latter category deals with the fine details of the small gap between the tuning rods and the cavity endplates.

4.4.1 *Guide for data-taking*

Before data-taking can commence, the cavity resonant frequency at given rod positions must be known. In principle, then, a list of rod positions corresponding to some frequency range can be fed to an automated program which moves the tuning rods through each set of positions over some amount of time. For the 2014 data run, the process is a little more complicated since there are two channels to consider. The first thing I do is make mode maps and form-factor maps for both the TM_{010} and TM_{020} modes. As it turns out, TM_{020}

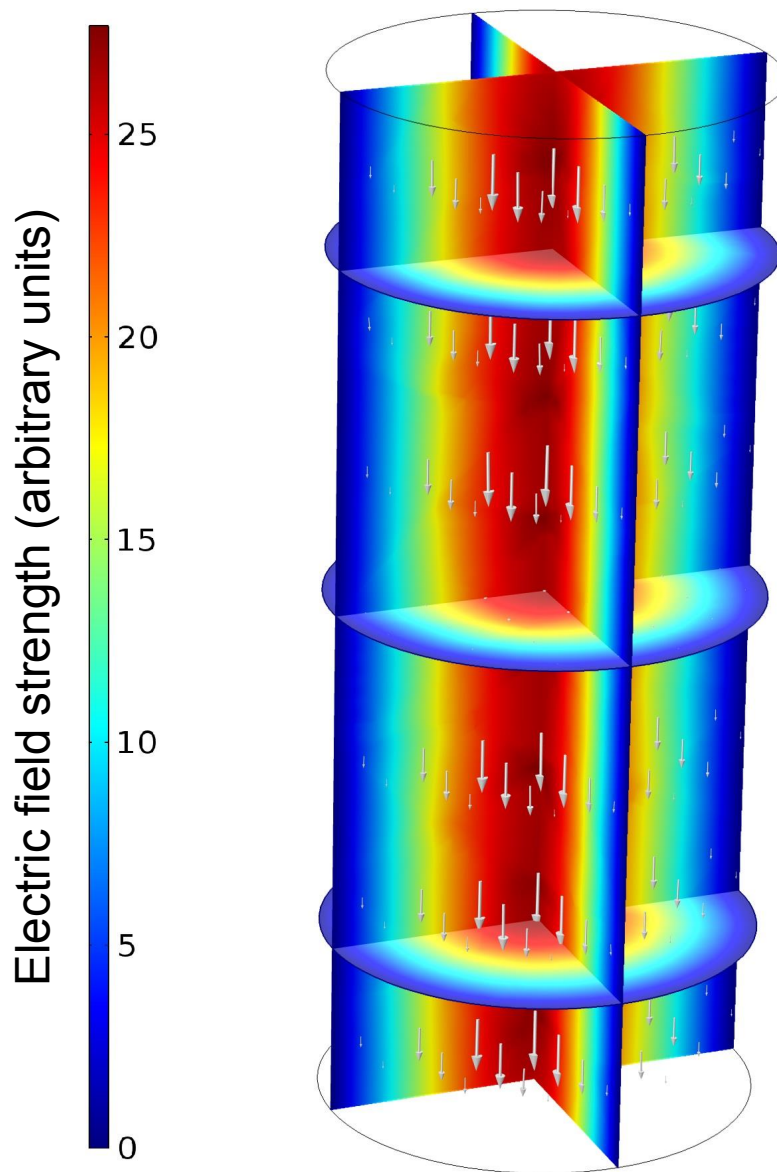


Figure 4.4: Simulated TM_{010} mode. The coloring represents electric field strength. The arrows represent electric field line vectors at various points within the cavity.




Figure 4.5: TM_{010} mode map with simulation. The difference between the data and the simulation is less than 0.5%. Given these results it is easy to identify the TM_{010} mode within the cavity. The simulation is consistently slightly lower in frequency than the data, which suggests that some of the cavity features I did not include in my simple model likely reduce the effective cavity size. Mode map courtesy of C. Boutan.

is not continuous over its full frequency range, but is instead split into distinct regions. Over the 900 - 1500 MHz range I identify and label three such regions: $TM_{020(A)}$ at 950 - 1150 MHz, $TM_{020(B)}$ at 1200 - 1250 MHz, and $TM_{020(C)}$ between 1325 - 1475 MHz. The 3D form factors maps, which are considerably easier to interpret than the 3D mode maps, are given in figures 4.6 - 4.9 for TM_{010} and the three TM_{020} regions. The form-factor maps give the form factor of a given cavity mode as a function of the tuning rod positions, which are measured by θ_1 and θ_2 . A color scale for the form factor is used with each map to make it easier to interpret.




Figure 4.6: TM_{010} form-factor map. θ_1 and θ_2 measure tuning rod positions, and the form factor is measured in the z axis as well as by the color scale.




Figure 4.7: $TM_{020(A)}$ form-factor map. θ_1 and θ_2 measure tuning rod positions, and the form factor is measured in the z axis as well as by the color scale.




Figure 4.8: $TM_{020(B)}$ form-factor map. θ_1 and θ_2 measure tuning rod positions, and the form factor is measured in the z axis as well as by the color scale.




Figure 4.9: $TM_{020(C)}$ form-factor map. θ_1 and θ_2 measure tuning rod positions, and the form factor is measured in the z axis as well as by the color scale.

To figure out a whether or not I can step continuously between one TM_{020} region and the next, I plot form factor vs frequency for all modes in the range 900 - 1500 MHz, taking only the highest form factor at any given frequency. This results in the plot in figure 4.10, which confirms that the three regions of good form factor are distinct from one another. I




Figure 4.10: f_{020} over 900 MHz - 1500 MHz.

should note that the “jagged” nature of this plot is a result of the 4° rod resolution. At finer resolutions, the excess structure of sharp peaks and valleys is not present and a smooth curve of form factor versus frequency emerges. The same applies to subsequent form factor versus frequency plots I will present.

For the 2014 run, the TM_{010} data channel frequency range is limited by the bandwidth of various electronics, primarily the SQUID. As a result, only certain frequency ranges corresponding to particular rod positions are allowed, and this limits the accessible channel two frequency range. $TM_{020(C)}$ has the best form factor for the given rod positions, so the

majority of the 2014 data set is taken around 1360 MHz for the TM_{020} data channel. The decision of where ADMX ultimately collects data takes into account simultaneously maximizing the form factors of both channels. This is visually represented in figure 4.11, which gives a contour plot of $f_{010} \times f_{020(C)}$ vs rod angles.




Figure 4.11: Contour map of $f_{010} \times f_{020(C)}$ in the 2014 data taking region. The tuning rods are directed to move in such a way that the product of the modes' form factors is maximized.

Finally, when running the experiment it is necessary to track modes as they change in frequency as a result of tuning rod motion. The cartoon in figure 3.5 suggests that there is one obvious peak in a transmission measurement, but in reality this is rarely the case. To pick out the correct mode I use data from plots like the one in figure 4.12, which overlays results from two simulations. Figure 4.13 shows this same simulation data with an overlaid mode map.

In the first simulation I look at all cavity modes around the expected TM_{0n0} frequency for particular rod positions, and in the second simulation I pick out only the specific TM_{0n0}




Figure 4.12: Mode map used to guide data taking. This kind of plot is generated for each frequency range in which ADMX takes data. Knowing ω_0 and its slope versus rod angle for all frequencies in over some bandwidth allows TM_{010} to be accurately picked out.




Figure 4.13: Mode map overlaid on simulation data to guide data-taking. With the combination of the simulation data and the mode map, it is relatively straightforward to pick out the correct TM_{010} even as it changes frequency with rod motion.

mode. I can then physically track the mode of interest by moving the tuning rods and ensuring that the resonant frequency, rate of change, and relative position of the mode I am tracking all behave as expected. In fact, the process could be completely automated with the help of high-precision simulation data. This is not necessary for the 2014 data run since the scan rate is relatively low; however, the addition of a dilution refrigerator will increase scan speed by several orders of magnitude, so an automated system of mode checking will likely become necessary.

4.4.2 Capacitive effect of tuning rods

In addition to providing critical information on how to tune the cavity, the TM_{010} form-factor map in figure 4.6 suggests something is amiss. There are areas with very clear and sudden drop offs in form factor, though according to studies done with earlier versions of ADMX, TM_{010} should have a slowly-varying form factor across the full range of the rods' motion. These earlier studies were done using a 2D model for the cavity, so to verify that I am not overlooking something simple, I run a 2D simulation and generate the form-factor map shown in figure 4.14. This agrees very strongly with the mode maps created and used by earlier versions of ADMX, and it exhibits all of the same characteristics as the plot in figure 4.6, save only for those chasms. This leads me to conclude that my COMSOL analysis is working as intended, and suggests the drop off in form factor is due to the gap between the tuning rods and the cavity endcaps.

To figure out if this affects the tuning range of the cavity, I make a plot of form factor over the TM_{010} tuning range, just as I did earlier for TM_{020} . Further, to ensure the effect works as I think it does, I run three more simulations with varying gap sizes. The resulting data are shown in figures 4.15 - 4.18.




Figure 4.14: TM_{010} form-factor map from 2D simulation. θ_1 and θ_2 measure tuning rod positions, and the form factor is measured in the z axis as well as by the color scale.




Figure 4.15: f_{010} over TM_{010} tuning range with 0" tuning rod gap.




Figure 4.16: f_{010} over TM_{010} tuning range with 0.025" tuning rod gap.




Figure 4.17: f_{010} over TM_{010} tuning range with 0.25" tuning rod gap.




Figure 4.18: f_{010} over TM_{010} tuning range with 0.75'' tuning rod gap.

It is indeed the case that a larger gap size leads to a wider inaccessible frequency region in the tuning range, though interestingly enough, the center frequency of this range increases, as well. It is also encouraging that with the tuning rods set to the length of the cavity and a resulting gap size of 0, the 3D simulation produces a result identical to that of the 2D simulation.

I can clearly see that this is a capacitive effect when I look at the electric field profiles of the different models. Figure 4.19 shows the differences between a models with gap sizes of 0" and 0.75".

These results shine light on a previously overlooked effect that has great ramifications for future cavity design, yet no doubt they only scratch the surface of this topic. I am sure that 3D modeling will play an even more active role in future incarnations of ADMX than it does now.




Figure 4.19: Capacitive tuning rod effect. The model on the left has tuning rods at the same length as the cavity, leaving no gap between the rods and the cavity endplates. The model at right uses rods that are 1.5" shorter than the cavity, leading to a 0.75" gap. For both models $\theta_1 = 70^\circ$ and $\theta_2 = 80^\circ$. The left model has $\omega_0 = 633$ MHz and $f_{010} \approx 0.55$, while these values are lowered to $\omega_0 = 625$ MHz and $f_{010} \approx 0.35$ for the right model. It is easy to see that in a certain frequency range, introducing a gap between the tuning rods and cavity endplates causes electric field to concentrate around the endplates and thus leads to a loss of form factor.

Chapter 5

DATA ANALYSIS

5.1 Introduction

In this chapter I discuss the data analysis that takes power series from the ADMX Structured Query Language (SQL) database and turns them into a limit on axion-to-photon coupling. The analysis software consists of thousands of lines of custom-written C++ code, and I will not go into the details of how I implement various analytic methods. Instead, I will focus on the net results of these methods.

I first describe power spectrum “conditioning,” which is necessary to remove the shape imparted by the ADMX receiver chain on data coming from the insert. I then describe how many individual power series spectra are added together to form a single “grand spectrum.” Finally, I show how the grand spectrum is used to set a limit on axion-to-photon coupling.

5.2 Power spectrum conditioning

As discussed in chapter 3, data from ADMX is saved as many individual power series in an SQL database. This power, however, is very far removed from that which axions deposit in the cavity. Running the course of the receiver chain, axion power has been amplified and distorted by the various analog and electronic components in the system. Therefore, before I start hunting for axions I must first condition the power spectrum into something sensible. The very first part of the conditioning involves discarding spectra which are relatively insensitive to axions.

5.2.1 *Quality control*

Each spectrum is parametrized by variables, such as magnetic field strength, which describe the state of the experiment at the time the spectrum was taken. Since data taking occurs whenever reasonably possible, the variables for some spectra are either too high or too low, which corresponds to a relative insensitivity to axions. In other words, such spectra would only add noise to a result. Therefore they are discarded, and not used in the rest of the analysis. Table 5.1 summarizes the cuts made on spectra. In total, 49,575 spectra were used from TM_{010} channel while 22,491 were cut, and 36,715 spectra were used from the TM_{020} channel while 35,489 were cut. Most of these cuts were due to overly-high temperatures, as often during and between runs thermal tests would be performed to determine the specifics of heat flow within the experiment. The TM_{020} channel has more cuts than the TM_{010} channel due to the greater difficulty of finding the appropriate higher-frequency modes. After these cuts are made, I move to processing the remaining spectra.

If I assume that noise in the spectra is Gaussian, I might expect to see a normal distribution of power around some value which depends on the net amplification. This naive assumption can be compared to figure 5.1, which shows a raw spectrum pulled off of the SQL database. This spectrum is clearly not normally distributed, so the next step in analysis is to remove the excess structure in the power spectra.

5.2.2 *Structure removal*

The shape added to the power spectrum by the receiver chain is corrected in two stages. First, the gross, large-scale structure is removed. This is accomplished by averaging the shape of many consecutive spectra and dividing each individual spectrum by that average. In principle the receiver chain should add a constant shape to all spectra within a some frequency range. However, I assume that this shape can vary slightly from day to day, possibly from thermal effects or as a result of other unknown processes. Therefore, averaging of spectra is done

Variable	Upper value	Lower value	Note
B field	8 T	1 T	Too high a magnetic field means there is an error in the data. The magnet is never taken above 7.6 T.
Cavity Q	200,000	10,000	Same as B field regarding the upper value: above 200,000 means there is an error in data collection.
Cavity temperature	4.5 K	0.5 K	Operating with only 1K pots during 2014, it is impossible for the cavity to get much below 1K, so too low a reading is an error. On the other hand, a temperature above that of the liquid Helium transition likely means something is wrong in the insert.
SQUID temperature	10 K	0.5 K	Same as cavity temperature regarding the lower value.
HEMT temperature	20 K	1 K	The HEMTs are thermally sunk to the reservoir, so they should never vary too much from 4 K.
Notes	* *		Anything other than “none” triggers a cut. Common cuts in this category are on ”SAG,” which represents a spectrum with a synthetically generated axion, or something similar to ”wrong mode,” which marks data that was taken on a non- TM_{0n0} mode.

Table 5.1: Cuts on spectra




Figure 5.1: Raw power spectrum from the from the TM_{020} channel, obtained directly from the ADMX SQL database.

on a day-to-day basis, with each new day generating a unique average shape. In figure 5.2 I overlay the daily power average over the raw data for the spectrum in figure 5.1, and in figure 5.3 I show that spectrum divided by the average.

There is residual structure left over in the power spectrum after taking out the gross shape, presumably due to the interactions of the cavity, probe, transmission line, and amplifier input, so the next step is to remove the small-scale fine structure from the spectrum. To achieve this I fit a 5th order polynomial¹ to the spectrum and again divide the data by the fitting function. Each single spectrum is fit individually since it should have unique fine structure. For the fit I employ the ASA047 Nelder-Mead minimization algorithm library [27], which is distributed under the GNU LGPL license. Figure 5.4 overlays the best-fit function to the

¹The order of the polynomial is based on wanting to fit out excess structure, but not fitting out any axion signal. Historically, ADMX has used a 5th order function because it was found to provide the best structure removal while still being insensitive to axions. Further, the cavity, probe, transmission line, and amplifier input interactions can be described by a 5-parameter circuit model[20].

data, and figure 5.5 presents the spectrum with this fine structure removed.




Figure 5.2: Power spectrum with gross structure fit. The shapes of all of the power spectra for a single day are averaged, producing a daily average function.

5.2.3 Power scaling

With both gross and fine structures removed, I now convert the power in these “structureless” spectra from arbitrary units to Watts. In order to do this, I first assume that the noise in the data is Gaussian. To test this assumption I calculate the goodness of a Gaussian fit, based on having a reduced χ^2 close to 1, to a histogram of the rms noise of many structureless spectra. Figure 5.6 shows the Gaussian fit overlaid on the histogram of noise data, and as expected, the agreement is very good. Given, then, that I am dealing with normally-distributed Johnson noise with an unknown scale factor, I have

$$C \times \sigma = k_B T_{sys} B, \quad (5.1)$$

where C is the unknown scale factor, σ is the rms noise power, and k_B , T_{sys} , and B are the usual Boltzmann constant, system temperature, and bandwidth. The scale factor depends




Figure 5.3: Power spectrum less gross structure. Each power spectrum from a given day is divided by the daily average function for that day, resulting in a roughly-flat spectrum.

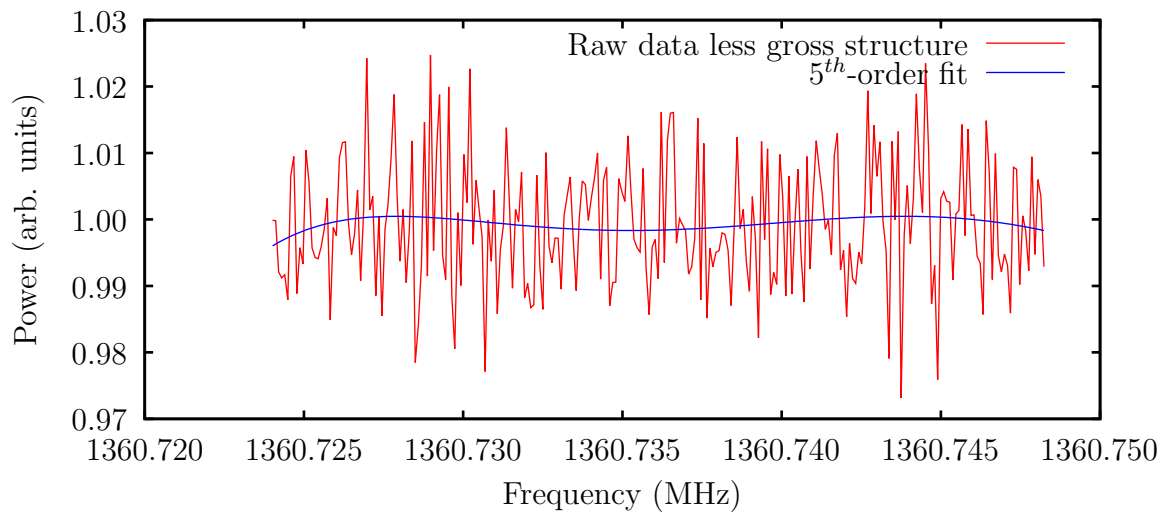


Figure 5.4: Power spectrum less gross structure with fit. There is residual small-scale structure left over in each spectrum, as evidenced by the 5^{th} -order fit to the data. Each spectrum gets a unique fitting function.




Figure 5.5: Power spectrum less gross and fine structures. Each spectrum is divided by its unique fit function, which generates a flat, structureless data set.

only on my analysis algorithm, so I expect it to be constant across all spectra. I calculate σ as the average of the standard deviations of 1200 consecutive structureless spectra, and for each of these spectra I ensure that T_{sys} is nearly identical. The bandwidth here is just the bin size for data in the power spectrum, which is 95.4 Hz. Substituting all appropriate values, I obtain $C = 9.2 \times 10^{-19}$. This constant converts the arbitrary power units of structureless spectra into Watts. Figure 5.7 gives the “corrected” spectrum, which is a measure of the power coming out of the cavity.

It is important to note that the standard deviation of the corrected spectrum is of order 10^4 yoctowatts, whereas axion conversion is expected to deposit only 10^3 yoctowatts in the cavity; therefore, a single scan will not be able to resolve an axion signal. Instead, many scans are averaged together to produce a grand spectrum which has much lower relative noise than a single spectrum, and in which axion signals can be identified.




Figure 5.6: Histogram of deviations overlaid with a Gaussian fit. The top plot uses a linear y-scale, while the bottom plot presents the same data on a log y-scale. The spectra used for this plot come from two days of data taking. There are 499 degrees of freedom in the fit and $\chi^2 = 445$, giving just a 4% chance that the noise is not normally distributed.




Figure 5.7: Corrected spectrum.

5.3 Grand spectrum

The idea behind forming a grand spectrum is that when averaging values, the uncertainty in the average decreases as the square root of the number of samples. Each raw spectrum has a bandwidth of approximately 25 kHz, and the average run speed for the 2014 data set shifts the raw spectrum center frequency by only 30 Hz from scan to scan. Furthermore, the bin size of raw spectra is 95.4 Hz, whereas I use a 1 kHz² bin size for the grand spectrum. Thus I expect to see a noise decrease of $\sqrt{(25,000/30)(1000/95.4)} \approx 100$ in the most sensitive bins. Since an axion deposits about 1/10 the noise power in a single scan, with the noise decreased by two orders of magnitude I can reasonably expect to detect axion signals present in the data.

When adding many spectra together, I have to take into account the fact that, due to

²To achieve the maximum possible signal to noise ratio the bin size of the grand spectrum should be equal to the axion width [21]. As I expect an axion with approximately 1 kHz width, I set the bin size to 1 kHz for both channels.

the cavity's frequency-dependent response, each 95.4-Hz bin in a given 25-kHz spectrum has a different sensitivity to axions.

5.3.1 Cavity response

A resonant cavity driven by axions converting into photons will behave in the same general way as any other damped, driven oscillator. The response of all such oscillators takes the form of a Lorentzian function, given by

$$L(\omega) = \frac{1}{1 + 4Q^2\left(\frac{\omega}{\omega_0} - 1\right)^2}. \quad (5.2)$$

$L(\omega)$ is plotted in figure 5.8 for Q and ω_0 typical of the TM_{020} frequency.




Figure 5.8: Lorentzian for typical cavity parameters. Here $\omega_0 = 1360$ MHz and $Q = 80,000$.

I account for this frequency-dependent cavity response when I add many individual spectra together. When adding spectra, the power and uncertainty associated with each bin

in a given spectrum are weighted by the cavity’s relative sensitivity to axions at that bin’s frequency for that spectrum. The weighing parameter is simply $L(\omega)$, which has an average over the entire 2014 data set of approximately 0.65, with a range of about 0.35 - 0.99.

With the cavity response accounted for, I there is only one more step to perform before assembling the full 2014 grand spectrum. Namely, I must ensure that my raw-spectrum processing has worked properly. Specifically, there are two properties of the processing output that I want to test: first, noise should decrease with increasing number of data points per bin, and second, axions should not be fit out. To ensure these properties I run some tests on a “mini” grand spectrum, which covers only about 250 kHz. This subset of the data is used because it is large enough to be representative, yet small enough such that I can still distinguish the important features I am looking for.

5.3.2 *Mini grand spectrum*

To average many corrected spectra, I simply bin the grand spectrum frequency range into 1 kHz chunks and add power from individual spectra to those bins. After all of the data is collected, I average the power in each bin by the number of points that were added to that bin. When building the grand spectrum I want to consider only the excess power in the cavity, since axions can only ever add to the power being read out, never detract from it. Thus after averaging many corrected spectra, I subtract the average power from the data and focus only on the positive range of values. The resulting mini grand spectrum is shown in figure 5.9.

To ascertain that $P_N \propto 1/\sqrt{N}$ for N data points in a bin, I compare the standard deviation of each bin to the number of points in it. The graph of those two curves is shown in figure 5.10, and it agrees with the expected result.

In verifying that I can detect axions in the grand spectrum, I will first describe what kinds of signals will be picked out by my analysis algorithm. The axion lineshape is that




Figure 5.9: Mini grand spectrum. The data set used for this spectrum was acquired over approximately three weeks. The apparent structure in this mini grand spectrum is a result of large variations in the number of individual power spectra co-added in each bin. The central frequencies have far fewer added spectra, resulting in considerably more noise, than the frequencies at the ends of this mini grand spectrum.

of from a Maxwellian velocity distribution, as shown in figure 5.11 [28]. Thus an axion signal would appear as a peak above the background in the grand spectrum. So I need only to inject an axion lineshape with the proper amount of power into the cavity, and make sure that my data analysis finds it with high confidence. One of the ways this is achieved is with the help of the ADMX Synthetic Axion Generator (SAG). This device takes an axion lineshape, Fourier transforms it to a power time series, amplifies it to some amount of power (usually on order of 10^{-21} W), and finally injects that signal into the cavity through the weak port. Another way to accomplish axion injection is digitally, by adding an axion's worth of power at some frequency to the raw spectra. Figure 5.12 shows the mini grand spectrum with one SAG axion and one digital axion. Both of these are clearly visible, so I can be confident that my analysis will not completely fit out real axions.




Figure 5.10: Grand spectrum uncertainty vs number of points per bin. As expected, the noise does indeed fall off as the inverse square root of the number of points per bin. The frequency dependence of the number of points per bin is the result of varying the scan speed during this data run.




Figure 5.11: Axion lineshape for a 1 kHz-wide axion.




Figure 5.12: Mini grand spectrum with synthetic axions. One axion with roughly 60x KSVZ axion power is injected with the SAG at approximately 1360.73 MHz, while another at 10x KSVZ axion power is digitally injected at exactly 1360.75 MHz. Both of these peaks are clearly visible over background. The added structure around the peaks is a result of the conditioning algorithm, though this structure is minor compared to the background.

Now I move to making the full grand spectrum for the 2014 data set.

5.3.3 2014 data set grand spectrum

The 2014 data set spans over 800 MHz, which visually makes it difficult to grasp at its native 1 kHz bin size. Furthermore, only about 1/4 of this span contains data, while the rest of the range is skipped over. Therefore I make two simplifications in presenting the grand spectrum. First, I average every 1000 1 kHz bins, weighted by the number of points in those bins, into 1 MHz bins. Second, I cut up the full spectrum into five sections, each of which is only tens of MHz wide and contains the regions in which data was collected. These plots are given in figures 5.13 - 5.17.




Figure 5.13: Grand spectrum for 590 - 630 MHz.




Figure 5.14: Grand spectrum for 720 - 760 MHz.

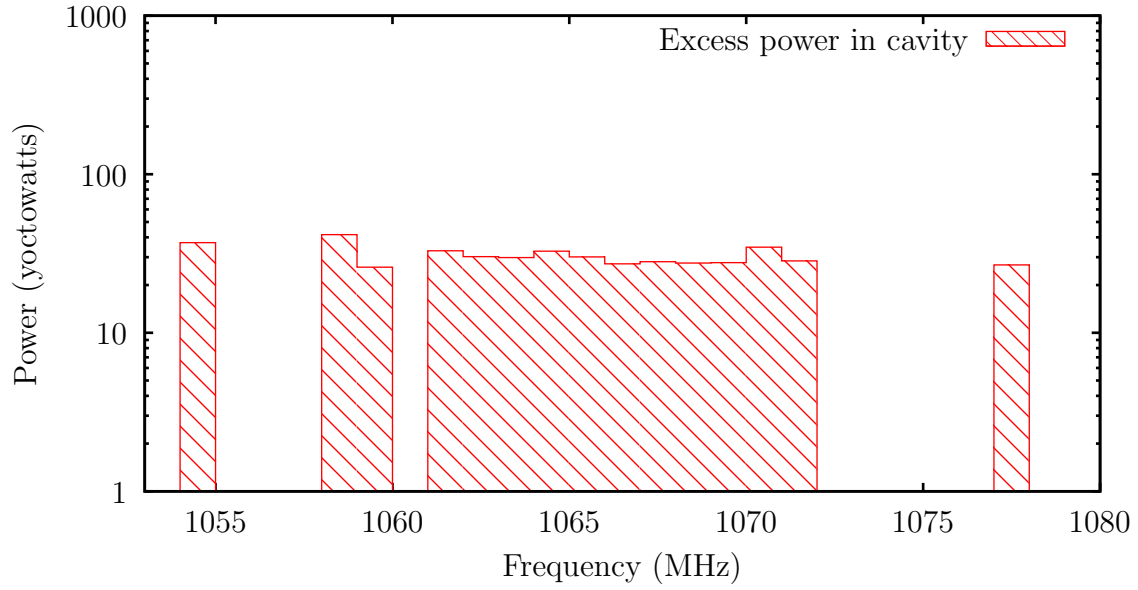


Figure 5.15: Grand spectrum for 1050 - 1080 MHz.




Figure 5.16: Grand spectrum for 1210 - 1280 MHz.




Figure 5.17: Grand spectrum for 1350 - 1390 MHz.

From the grand spectrum data I move on to actually searching for axions.

5.4 Axion candidate search

As mentioned earlier, axions would appear in the data as bins with excess power. This would mean that the entire grand spectrum is a potential sea of axion peaks.

I define an excess-power peak to be an axion candidate if

$$P_{bin} \geq 5\sigma_{bin}, \quad (5.3)$$

where P_{bin} is the excess power in a given bin, and σ_{bin} is the uncertainty in that bin. Under this condition, given my grand spectrum bin size (1 kHz) and the full bandwidth of the data set (GHz), I expect to find less than one candidate attributable entirely to statistical fluctuations. On the other hand, this condition does not preclude peaks of reasonable QCD-axion power. Figure 5.20 shows a plot of the mini grand spectrum with synthetic axions, along with a curve showing the 5σ threshold. Both synthetic axions would be tagged as candidates, whereas the noisy middle section of this spectrum would not.

In this data set no axion candidates were found, so I move to setting limits on $g_{a\gamma\gamma}$.

5.5 Limit generation

Since $P_a \propto g_{a\gamma\gamma}$, I can simply multiply the grand spectrum by the appropriate constant to generate a limit on $g_{a\gamma\gamma}$. To do this I need to calculate an upper limit for the power in each bin. The first step in this calculation is ascertaining how much my analysis attenuated the axion signal. This attenuation is caused by two factors, D and M, described below.

5.5.1 Dilution factor

To describe how much power is lost in an axion peak due to the two fits applied I define a dilution factor D as

$$D = \frac{P_{actual}}{P_{expected}}, \quad (5.4)$$




Figure 5.18: Mini grand spectrum with synthetic axions and 5σ curve. Both of the synthetic axions added would be tagged as candidates.

where $P_{expected}$ is the power I expect to see from an axion peak, while P_{actual} is the power observed. This definition of course requires using known synthetic axion peaks for calibration, and I do this digitally since using the SAG would be considerably slower.

I perform 50 different analyses on the mini grand spectrum, with each analysis featuring a digitally injected synthetic axion one bin size (1 kHz) advanced from the previous analysis. That is, analysis 1 has an axion injected at 1360.800 MHz, analysis 2 at 1360.801 MHz, and so on. Each synthetic axion is 1 kHz wide, for reasons discussed in the previous section. I make sure to line up the axion lineshape with the start of a bin for each analysis so that I get an accurate value for $P_{expected}$. If I expect axions to deposit 1000 yW in the cavity, I integrate the Maxwellian in figure 5.10 from 0 to 1000 Hz and multiply by 1000 yW to find that 668 yW should be deposited in bin 1. I pick out the power at the injected axion frequency for each analysis, and then average all 50 analyses to get 475 yW. The uncertainties in each bin are also sum-squared averaged to give 8.7 yW, with each individual bin in this range having

a standard deviation of roughly 60 yW. Finally, I obtain $D = .711$ and $\sigma_D = .013$.

Naively I expect my fits to have little effect on axion peaks, so the fact that roughly 30% of peak power disappears is quite surprising. However, I verify that axion power is indeed being fit out by running another series of 50 tests where I make the axion 1.1 kHz wide, but pump up the power to 1057 yW and offset the axion lineshape from the starting bin by -50 Hz. In this manner I still expect 668 yW to be deposited in the starting bin, though now the average power drops to 447 yW. Since a slightly-wider axion has such a large impact on P_{actual} , I am forced to conclude that the fitting algorithm negatively impacts this value.

5.5.2 *Binning a Maxwellian*

Another consideration I must make in looking at how my data analysis affects the observed power is how the axion lineshape actually lines up with my bins. The maximum amount of power I can get from a 1 kHz range of a 1 kHz-wide axion is about 68% when the axion lineshape starts at roughly -50 Hz relative to the bin. The maximum power I collect in a bin occurs when the lineshape is approximately 65 Hz delayed relative to the bin, whereas the minimum power occurs when the lineshape is approximately 370 Hz advanced of the bin, as shown in figures 5.18 and 5.19. In this case my analysis code picks up only 44% of the axion's power. I should note that it does not matter if the maximum power occurs in bin 1 or bin 2, as the analysis will pick out whatever bin has the maximum power.

I now define a scale factor M which arises from this binning of the Maxwellian, and which measures the ratio of maximum to minimum expected powers in a given bin. For a 1 kHz axion I get

$$M = \frac{0.68}{0.44} = 1.55. \quad (5.5)$$




Figure 5.19: Maximum 1 kHz axion power spread over two bins. When the axion lineshape is approximately 65 Hz delayed relative to the start of a bin, that initial bin contains the maximum possible power at 68% of the axion’s total power.

5.5.3 Upper limit on power per bin

Now that I have my attenuation factors, I can calculate the upper limit on P_{bin} . I define the excess amount of power that an axion signal would produce as P_{excess} , which is equal to the observed power scaled by the previously-mentioned factors D and M. Also, I do not want the uncertainty in D to play into the confidence level (CL) that I will set, so with 99.9% confidence I take the minimum value for D to be $D - 3.1\sigma_D$.

$$P_{excess} = P_{bin} \frac{M}{D - 3.1\sigma_D} = 2.3P_{bin} \quad (5.6)$$

I fix a 90% CL, which means my analysis should be able to detect 90% of axion signals. With a normal distribution like the thermal noise in the cavity signal, at 1.28σ from the




Figure 5.20: Minimum 1 kHz axion power spread over two bins. When the axion lineshape is approximately 370 Hz advanced of the start of a bin, both that starting bin and the following one equally contain 44% of the axion’s total power. Thus 44% of the total axion power is the minimum I expect to see in any one bin for a 1 kHz-wide axion.

mean the tail on each side the curve contains only 10% of the total data in the distribution. Here σ is the standard deviation of each bin. Almost all of the uncertainty in the analysis is taken to come from the thermal noise of the signal, while variables like the cavity Q , ω_0 , and f_{0n0} are assumed with minimal error. The greatest systematic uncertainty comes from determination of amplifier noise, which affects T_{system} and the consequent power scaling. I take this into account when I determine T_{system} , so this systematic does not factor into the 90% confidence limit. Defining $P_{90\%CL}$ as this upper limit, I have

$$P_{90\%CL} = P_{excess} - 1.28\sigma_{bin}. \quad (5.7)$$

5.5.4 Axion-to-photon limit

To calculate the upper limit on $g_{a\gamma\gamma}$ I start with

$$\frac{g_{a\gamma\gamma(90\%CL)}^2}{g_{a\gamma\gamma(KSVZ)}^2} = \frac{P_{excess}}{P_{KSVZ}}. \quad (5.8)$$

Rearranging the equation and substituting for $g_{a\gamma\gamma(KSVZ)}$ gives me

$$g_{a\gamma\gamma(90\%CL)} = 10^{-7}(\text{GeV}^{-1}) \frac{m_a}{0.62 \text{ eV}} \frac{\alpha g_\gamma}{\pi} \sqrt{\frac{P_{excess}}{P_{KSVZ}}}. \quad (5.9)$$

P_{KSVZ} varies from bin to bin as a result of changes in ω_0 , Q, B, and f_{0n0} between scans. The form factor used to calculate P_{KSVZ} is obtained from cavity simulations, and the local dark matter density is taken to be $0.3 \frac{\text{GeV}}{\text{cm}^3}$. As with derived cavity parameters, I do not consider cosmological uncertainties in my limit. The exclusion curve is given in figures 5.21 - 5.25, broken up into the same five sections as was the grand spectrum.




Figure 5.21: Exclusion plot for 590 - 630 MHz.




Figure 5.22: Exclusion plot for 720 - 760 MHz.




Figure 5.23: Exclusion plot for 1050 - 1080 MHz.




Figure 5.24: Exclusion plot for 1210 - 1280 MHz.




Figure 5.25: Exclusion plot for 1350 - 1390 MHz.

In the 1050 - 1400 MHz range these data represent new exclusion limits. Figure 5.26 shows limits generated from the TM_{020} modes as compared to limits from earlier cavity experiments done by the Rochester-Brookhaven-Fermilab (RBF) collaboration [29] and the University of Florida [30].




Figure 5.26: Exclusion limits on $g_{a\gamma\gamma}$ for 1050 - 1400 MHz compared to previous limits.




Figure 5.27: Exclusion limits on $g_{a\gamma\gamma}$ for 1050 - 1400 MHz compared to previous limits.

Chapter 6

CONCLUSIONS AND FUTURE WORK

I have described the motivation for, hardware and software aspects of, and first results from ADMX Phase II. This experiment has demonstrated the sensitivity and range of the TM_{020} cavity mode, as well as the ability of the ADMX experiment to take data consecutively in multiple cavity modes. New exclusion limits on $g_{a\gamma\gamma}$ were set in all frequency ranges studied.

The use of 3D modeling software to perform cavity simulations has also been very successful. Not only were the results of simulations able to guide ADMX data taking, but a previously-unappreciated capacitive effect was discovered. This suggests a strong need for continued 3D modeling, and the impact of this work on future cavity designs cannot be understated.

ADMX has a long history of running fruitful and groundbreaking experiments, and this run of Phase II fits that mold well. Currently we are hard at work ensuring that ADMX will continue its successes into 2015 and beyond with the Gen II experiment.

6.1 Future work

6.1.1 Dilution refrigerator

The most obvious and arguably most exciting step forward with the experiment is the addition of the dilution refrigerator. For the 2014 data runs the physical temperatures of the cavity and first stage amplifiers were around 1.5 K or greater. Dropping these temperatures to the 100 mK of the dilution refrigerator will increase ADMX's scan speed by a factor of

over 200. This means that with the dilution refrigerator ADMX would only need a single day of data taking to set limits as good as were obtained from the entire 2014 data set!

As of this writing (February 2015) the ADMX dilution refrigerator has been tested and is getting ready to ship. Figure 6.1 shows a picture of the refrigerator, taken during its acceptance test in January 2015.




Figure 6.1: Picture of the ADMX dilution refrigerator.

6.1.2 Use of JPAs

The balanced NRAO HEMT that was the 1st stage amplifier in the second channel was ideal for the 2014 data runs. Its large bandwidth, low noise, and relative insensitivity to physical temperature at 4 K made it a great amplifier over 1 - 1.5 GHz. However, that same temperature insensitivity ultimately invalidates HEMTs as 1st stage amplifiers for 2015 and beyond. With temperatures expected to drop precipitously, it is essential that we use an amplifier which shows a marked drop in noise temperature with decreasing physical temperature. Therefore, ADMX will use a Josephson Parametric Amplifier (JPA) as the channel 2 1st stage amp.

The basic idea of JPAs is that a harmonic oscillator can exhibit amplitude gain if the oscillator's parameters are themselves oscillated in a particular and nonlinear way. The classical analogy is to a child on a swing, who can increase the amplitude of her motion, provided she has some to begin with, by standing and squatting at various points in the pendulum's motion.

Physically a JPA is as simple as a SQUID hooked up in parallel with a capacitor. If the SQUID is considered a nonlinear inductor, then it is easy to imagine an equivalent nonlinear RLC circuit that serves as the oscillator. In this case the inductance would be the varying parameter. I will not go into the details of JPAs, though those interested may consult, for example, reference [31]. For now, it is important to note that JPAs provide high gain (20 - 30 dB) while adding only a slight amount of noise, very close to the quantum-limited value. Furthermore, they can be operated at several GHz and have large bandwidths (typically a few GHz), making them ideal for the Phase II second channel as well as the higher-frequency cavities of the future. Figure 6.2 shows the JPA inside its aluminum container.




Figure 6.2: At left, the JPA sits in its opened aluminum box. This device actually has four JPAs on one chip, all at different frequencies. ADMX uses only one JPA at a time, so the input to each of the other three amplifiers is terminated. The device has only two connections to external electronics: an RF cable to send in and extract power, as well as a twisted wire pair to bias the SQUID. A copper strip used for thermal sinking the device bolts to the chip and runs out of the aluminum box to a cold source. At right, the JPA in its aluminum shell is installed in the squidadel, right below the MSA, which sits in its own lead can. A few of the many switches and circulators that surround the amplifiers within the squidadel can be seen on the sides.

6.1.3 Improved thermal design

Of course, a dilution refrigerator and fancy amplifiers can only be useful if the rest of the system has a good thermal design. After all, a dilution refrigerator can easily be made

useless by an extraneous heat leak to the mK stage, and the fact that a JPA adds little noise temperature to the signal does not mean much if the wires carrying the signal are a major noise source themselves. To that end, we have considerably upgraded the insert Cold space. The Squidadel, 1K plate, and most of the other OFHC copper components around the 1K area have been gold plated to ensure that they will have the minimum possible thermal boundary (Kapitza) resistance. Further, all wires and RF cables going down to the cavity will have much better thermal sinking to the various temperature stages than they had before. Figure 6.3 shows off the new and improved 1K space.

6.1.4 *ADMX Sidecar*

In addition getting the experiment geared up to operate at mK temperatures, we will additionally be making a couple of other major changes. First, we will be running the ADMX Sidecar experiment in tandem with ADMX. Sidecar will use a 2.5" diameter, 6" long cavity to search for axions with a mass around 3-10 GHz. Figure 6.4 shows the prototype Sidecar cavity.

Perhaps the most ideal property of the Sidecar experiment is that it fits neatly on top of the main ADMX cavity, as shown modeled in figure 6.5. As Sidecar will be run alongside ADMX, we will be able to collect data from three frequency ranges (two channels from the main experiment plus Sidecar), for the price of running just one experiment.

6.1.5 *Ouroboros*

Another change that will possibly be made before the next round of data taking is the use of an active resonator system to increase SNR. The idea behind such a system is that a power splitter siphons off some of the ADMX signal, amplifies it, and then feeds it back to the cavity. When the conditions are set so that noise adds incoherently while the axion signal adds coherently, a considerably SNR improvement is observed. ADMX is studying active

resonator systems through the Ouroboros prototype. A plot showing recent results from Ouroboros is given in figure 6.6. More information about this experiment can be found in Appendix A.

6.1.6 New cavity designs

Finally, ADMX is actively engaged in R&D with eyes on future versions of the experiment. In general, the focus is on becoming sensitive to axions at higher and higher frequencies. One branch of this involves prototyping different-sized cavities with varied tuning mechanisms. A few examples include using mini cavities, vanes as tuning mechanisms, and even adding dielectric material inside the cavity to improve form factor. Figures 6.7 - 6.9 illustrate these examples.




Figure 6.3: Picture of improved Cold space. Many elements have been gold-plated to improve their thermal conductivity. A second 1K pot, which will be used to precool ^3He going to the dilution refrigerator, has also been integrated.




Figure 6.4: Prototype Sidecar cavity. The Sidecar experiment uses a one tuning rod, which is controlled by a piezo-electric drive. This drive is seen on the top plate of this cavity prototype.




Figure 6.5: Model of Sidecar cavity installed on main cavity. The Sidecar cavity is shown in a light-copper color bolted to the top plate of the main cavity.




Figure 6.6: Ouroboros recent results. Q is a read off of a network analyzer sweeping the cavity, while SNR is obtained from a power spectrum analyzer. The increase in Q is obtained by increasing the gain of power being fed back into the cavity. The linear increase in SNR with Q in this prototype suggests ADMX can use the same technique to improve axion search sensitivity.




Figure 6.7: Miniature cavity prototype. This cavity has a 10" long barrel and a 4" length inner diameter, with a tuning range of 3.5 - 5.8 GHz. It uses only one tuning rod, which rotates off center from the main cavity axis. The picture at left shows the cavity buttoned up hanging from its test stand, while at right the top plate is removed and the single tuning rod can be seen. This prototype is being tested at Lawrence Livermore National Lab by ADMX collaborator Gianpaolo Carosi, and images are courtesy of Gianpaolo.




Figure 6.8: UF vaned cavity. This prototype, which uses rotating vanes for turning, is being tested at the University of Florida. The top two pictures show COMSOL results of relative electric field strength within the cavity for two different positions of the vanes. The bottom picture is of the physical cavity with the top plate removed. Images courtesy of Ian Stern of ADMX.




Figure 6.9: Electric tiger prototype. At left is a cutaway CAD model of the experimental cavity. Dielectric spacers are positioned such that they minimize the form-factor-reducing component of the electric field. This is shown in a simulation result at right, which gives the electric field of the TE_{019} mode in the cavity. A constant magnetic field aligned with the minor axis of the cuboid cavity would produce a relatively good form factor (around 0.3), whereas if the spacers are removed the form factor drops down to zero. Furthermore, translating the spacers within the cavity gives it a tuning range of around 7 - 12 GHz. Images courtesy of Jim Sloan of ADMX.

Another area of active research involves working with entirely different resonator designs. ADMX's Gray Rybka leads the Orpheus experiment, which uses just two reflector plates to form an open-walled Fabry-Perot cavity. The hallmark of this experiment is the addition of a number of current-carrying wire planes interspersed between the reflectors. These wire planes are mobile and as such can be positioned in a way to line up the magnetic field generated by running current through the wires with the electric field modes of the cavity. Figure 6.10 illustrates this idea, while figures 6.11 and 6.12 show the experimental prototype and its recent results. For more information on the Orpheus experiment, see reference [32].





Figure 6.10: Orpheus experiment schematic. Two reflector plates form an open-walled Fabry-Perot cavity. Current-carrying wire planes are arranged between the reflectors such that the resulting magnetic field is always aligned with the electric field of the cavity mode; this maximizes the form factor for any given cavity mode. The reflectors and wire planes can be translated to tune the resonant frequency.

Wire Plane Power Supply

Amplification



Network Analyzer

Reflector

Wire Planes

Reflector

Spectrum Analyzer

Figure 6.11: Orpheus prototype setup. The Orpheus prototype is a bench top experiment, though in principle the experiment can be scaled much larger. Two polished aluminum plates serve as the reflectors, and 4 wire planes are used in this setup (though up to 8 were in use for other setups).

6.1.7 Final words

Overall this has been a great start for ADMX Phase II, with promises of much more to come. With ADMX continuing to make strides at this rate, the “invisible” axion is unlikely to stay invisible for long.




Figure 6.12: Orpheus results. The excluded region of axion-to-two-photon coupling $G_{a\gamma\gamma}$ versus putative axion mass is shown. The yellow band highlights the region of the KSVZ and DFSZ axion models. Despite being a prototype and a tabletop experiment, Orpheus has been able to set limits competitive with other very large-scale experiments, from which limits are also shown.

BIBLIOGRAPHY

- [1] W. Greiner, S. Schramm, and E. Stein. *Quantum Chromodynamics*. Springer, 2007.
- [2] C. A. Baker et al. Improved experimental limit on the electric dipole moment of the neutron. *Phys. Rev. Lett.*, 97(131801).
- [3] V. Baluni. *CP*-nonconserving effects in quantum chromodynamics. *Phys. Rev. D*, 19(7).
- [4] R.J. Crewther, P. Di Vecchia, G. Veneziano, and E. Witten. Chiral estimate of the electric dipole moment of the neutron in quantum chromodynamics. *Physics Letters*, 88B(1,2).
- [5] J. Kim. Weak-interaction singlet and Strong *CP* invariance. *Phys. Rev. Lett.*, 43(2).
- [6] R. D. Peccei and Helen R. Quinn. *CP* conservation in the presence of pseudoparticles. *Phys. Rev. Lett.*, 38(25).
- [7] S. Weinberg. A new light boson? *Phys. Rev. Lett.*, 40(4).
- [8] F. Wilczek. Problem of strong *P* and *T* invariance in the presence of instantons. *Phys. Rev. Lett.*, 40(5).
- [9] M.S. Turner E.W. Kolb. *The early universe*. Addison-Wesley, 1988.
- [10] K.A. Olive et al. (Particle Data Group). Quarks. *Chin. Phys.*, C38(090001).
- [11] Dicus et al. Astrophysical bounds on the masses of axions and Higgs particles. *Phys. Rev. D*, 18(6).
- [12] C. Edwards et al. Upper limit for $J/\Psi \rightarrow \gamma + \text{Axion}$. *Phys. Rev. Lett.*, 48(14).
- [13] M. Sievertz et al. Upper limit for axion production in radiative Υ decay. *Phys. Rev. D*, 26(3).
- [14] M. Shifman, A. Vainshtein, and V. Zakharov. Can confinement ensure natural *CP* invariance of strong interactions? *Nuclear Physics B*, 166.

- [15] M. Dine, W. Fischler, and M. Srednicki. A simple solution to the Strong CP problem with a harmless axion. *Phys. Lett.*, 104B(3).
- [16] A. Zhitnitsky. On possible suppression of the axion hadron interactions. *Sov. J. Nucl. Phys.*, 31(260).
- [17] P. A. R. Ade et al. Planck 2013 results. I. Overview of products and scientific results. *Astronomy and Astrophysics*, 571(A1).
- [18] P. Sikivie. Experimental tests of the invisible axion. *Phys. Rev. Lett.*, 51(16).
- [19] R. Dickie. *Rev. Sci. Instrum.*, 17(268), 1946.
- [20] Edward Daw. *A Search for Halo Axions*. PhD thesis, Massachusetts Institute of Technology.
- [21] Michael Hotz. *A SQUID-Based RF Cavity Search for Dark Matter Axions*. PhD thesis, University of Washington.
- [22] J. Clarke and A. I. Braginski. *The SQUID Handbook: Volume 2 Applications of SQUIDs and SQUID Systems*. Wiley, 2006.
- [23] R. Dickie. *Apl. Phys. Lett.*, 72(22), 1998.
- [24] S. J. Asztalos et al. Design and performance of the ADMX SQUID-based microwave receiver. *Nuclear Instruments and Methods in Physics Research A*, 656(1).
- [25] Mikulin et al. *Adv. Cryo. Eng.*, 31(629), 1984.
- [26] Fftw homepage, 2015.
- [27] Asa047 nelder-mead minimization algorithm, 2014.
- [28] Michael S. Turner. Periodic signatures for the detection of microwave axions. *Phys. Rev. D*, 42(10).
- [29] W. B. Fowler et al. Results of a laboratory search for cosmic axions and other weakly coupled light particles. *Phys. Rev. D*, 40(10).
- [30] C. Hagmann et al. Results from a search for cosmic axions. *Phys. Rev. D*, 42(4).

- [31] Manuel Beltran. *Development of a Josephson Parametric Amplifier for the Preparation and Detection of Nonclassical States of Microwave Fields*. PhD thesis, University of Colorado, 2002.
- [32] Rybka et al. Search for dark matter axions with the orpheus experiment. *Phys. Rev. D*, 91, 2015.
- [33] Gray Rybka. Improving dark matter axion searches with active resonators. *arXiv:1403.6720*, 2014.

Appendix A

OUROBOROS

The idea of the Ouroboros experiment is to artificially increase the Q of a resonant cavity. Integrated into ADMX, this would in turn increase the sensitivity of ADMX to the axion signal, as dictated by equation 3.9.

The basic idea of the Ouroboros setup is shown in figure B.1 below. Essentially, power

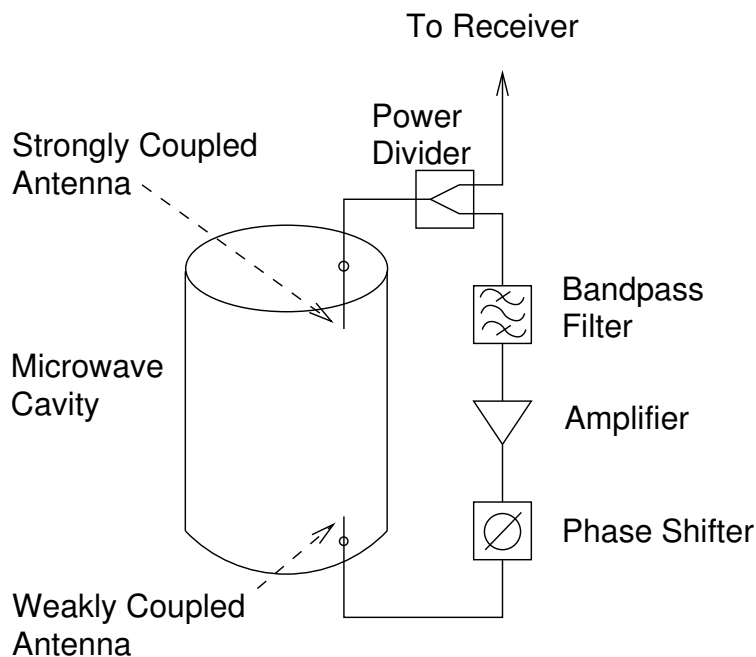


Figure A.1: Ouroboros schematic. The phase shifter is used to ensure that noise does not add coherently.

is siphoned from the cavity, amplified, and then fed back in. Noise adds incoherently while axion power adds coherently, resulting in an increase of SNR. Complete details of the theory

behind Ouroboros, including a derivation of increased SNR, can be found in reference [33].

This experiment is still ongoing, and results have not yet been published. However, there is preliminary data, as shown in figure 6.6, which is extremely promising. The setup from which this data was obtained is shown in figure B.2.




Figure A.2: Ouroboros setup. This prototype is more complicated than the simple schematic in figure B.1 would suggest only due to the addition of many switches, which are there for testing purposes. A network analyzer, at right in the picture, is used to measure the cavity Q . A signal generator, not seen, injects axion lineshapes into the cavity. A spectrum analyzer, also not shown, then extracts SNR by measuring the axion signal peak and the noise floor. The three power supplies are run the various electronic devices.