

A MODULATION SENSITIVE SEARCH FOR NON-VIRIALIZED AXIONS
IN THE MILKY WAY HALO

By

JEFFREY D. HOSKINS

A DISSERTATION PRESENTED TO THE GRADUATE SCHOOL
OF THE UNIVERSITY OF FLORIDA IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY
UNIVERSITY OF FLORIDA

2014

© 2014 Jeffrey D. Hoskins

This thesis is dedicated to my parents, David and Carol Hoskins. Without your love and support I would not be where I am today.

ACKNOWLEDGMENTS

I would like to thank my advisor, David Tanner, for his advice and support which he provided with great alacrity. Even in the most stressful of times, he was always the first to point out where to go next, and always with a cheerful, reassuring smile. I also thank Pierre Sikivie for all of the knowledge he imparted to me on the theory side of my work. Additionally, I would like to thank all of my collaborators in the Axion Dark Matter eXperiment as well as the other members of my advisory committee, Neil Sullivan, Steven Detweiler, and Anthony Gonzalez.

Of my friends in Gainesville, none have aided in my graduate studies more than Chris "I still don't know why you were named 'The Cat'" Pankow. His assistance when I was learning about scripting was invaluable. The same can be said about all those times in which I encountered a "problem problem," for which I may still owe him a drink. Many thanks as well go out to my other dear friends, Hsin-Jung Lin, El Señor Arenas, Mr. Dr. John, and, of course, Richard "my brother from another mother" Ottens. I value the time I have spent with each of them, and the impact on my life each has made. I can't imagine what grad school would have been like had none of them been there at all.

Thank you once again, to everyone who has been there with me on this long and fulfilling journey, be they in Gainesville, Orlando, or even somewhere in Arizona. As with any endeavor, there have been difficult times and getting through them has been made possible only by the words and actions of my friends and colleagues. Yes, even if said action is just someone sitting there listening... to me... at length... make that great length. In fact, especially so in that case. You know who you are.

TABLE OF CONTENTS

	<u>page</u>
ACKNOWLEDGMENTS	4
LIST OF TABLES	7
LIST OF FIGURES	8
ABSTRACT	11
CHAPTER	
1 INTRODUCTION	13
2 BACKGROUND INFORMATION	16
2.1 Introduction	16
2.2 Motivation for Dark Matter	16
2.3 Dark Matter Properties	18
2.4 QCD and the Strong CP Problem	23
2.5 PQ Symmetry and the Axion	25
2.6 Astrophysical Considerations and Cosmology	28
2.7 Halo Structure	31
3 EXPERIMENTAL APPARATUS AND DATA ACQUISITION SYSTEM	36
3.1 Introduction	36
3.2 Axion Haloscopes	36
3.3 Cryogenic Hardware	38
3.4 Receiver Electronics	45
4 DATA ANALYSIS	54
4.1 Introduction	54
4.2 Axion Conversion and Signal Properties	54
4.3 Variable Resolution Considerations	64
4.4 Noise Properties	67
4.5 HR Data Set and Analysis Software	71
4.6 Data Quality Cuts	74
4.7 Removal of Systematic Effects	75
4.8 Candidate Search	82
4.9 Density Limits	85
4.10 Discussion	91
5 FUTURE WORK	94
5.1 Introduction	94
5.2 Threshold Selection	94

5.3	Hardware Upgrades	96
5.4	Higher Frequencies	98
6	SUMMARY AND CONCLUSIONS	107
APPENDIX		
A	A SEARCH FOR NON-VIRIALIZED AXIONIC DARK MATTER	110
B	ANALYSIS CODE	119
REFERENCES		121
BIOGRAPHICAL SKETCH		125

LIST OF TABLES

<u>Table</u>	<u>page</u>
4-1 A Summary of the resolutions examined by the HR channel. These include spectral resolutions, resolution reduction factor, effective thresholds, scale factor, and maximum associated velocity dispersions.	93

LIST OF FIGURES

<u>Figure</u>	<u>page</u>
2-1 Rotation curves for 7 of the 21 spiral galaxies measured in Rubin et al. [19]. Following the sharp rise caused by the core of the galaxy, these curves become asymptotically flat.	18
2-2 A NASA image of the Bullet Cluster merger. X-ray emission from gas is shown in pink, and the regions of highest density are shown in blue.	19
2-3 A NASA image of the CMB temperature map as measured by WMAP. On large scales the CMB shows a high degree of homogeneity and isotropy. The localized regions of higher and lower temperatures are indicative of small density perturbations in the early universe.	21
2-4 A) A 2D representation of phase-space winding in the rest frame of a strong gravitational potential. Here \dot{r} is the velocity of dark matter in the direction of r . As time progresses this clockwise winding continues, resulting in several degenerate populations of dark matter in the neighborhood of the potential. B) The projection of phase-space winding onto physical space. Note the significant increases in density at each turn-around point. In the limit of $\delta v = 0$ the density at these points would be infinite, though for small dispersions the increased densities remain finite.	33
3-1 A schematic of the ADMX cavity and receiver chain. Switch A allows for the selection of sending power to either the major or minor ports, or to not send power to the cavity at all. Switch B allows for power to be read out by a spectrum analyzer in lieu of being mixed down and saved to disk. When acquiring data, the switches are set as pictured.	37
3-2 A Feynman diagram of an axion decay within a haloscope. In this decay, the axion interacts with the applied B field and the decay photon at the loop-level.	39
3-3 A SolidWorks drawing of the ADMX cryostat. The central portion is referred to as the insert. From the top down the insert consists of radiation shields, the helium reservoir which contains the bucking magnet, and the ADMX cavity. RF cables and sensor wires run along the central axis and are read out at the top plate.	40
3-4 A picture of the ADMX cavity, opened, with tuning rods.	41
3-5 A picture of the copper plating flaking on top rim of the cavity wall. This created a poor electrical connection to the lid, which then degraded the Q	43
3-6 A diagram of the motion of tuning rods in the ADMX cavity from a top down prospective. Each rod (dark grey) is connected to an arm (light grey) which which can freely rotate 360°	44

3-7	A picture of the bucking magnet prior to installation.	46
3-8	The on-axis, magnetic field profile along the length of the bucking magnet.	47
3-9	A diagram of the MSA as used in ADMX. The SQUID, responsible for the actual amplification, consists of the Nb washer and Josephson junctions, while the microstrip, responsible for coupling power to the squid, is the Nb shown in blue.	48
3-10	The noise temperature profile of the MSA. Once above the quantum limit, the noise temperature is linearly dependant on the physical temperature.	50
3-11	A diagram of the balanced design for the HFET amplifiers.	51
4-1	A numerical simulation of the power measured in a signal with an average decoherence time of 2.38 seconds. For resolution selection in the frequency domain, the average loss of 38% is well below the theoretical maximum of 68.4%.	67
4-2	Power distribution for a sample of 42 mHz data (n=1).	69
4-3	Power distribution for a sample of 84 mHz data (n=2).	71
4-4	Power distribution for a sample of 168 mHz data (n=4).	72
4-5	Power distribution for a sample of 546 mHz data (n=13).	72
4-6	Power distribution for a sample of 1.09 Hz data (n=26).	73
4-7	A spectrum showing contamination from an external radio signal.	76
4-8	The 42 mHz power distribution for the data shown in Fig. 4-7. Large deviations from the exponential distribution caused by the radio signal can be seen at higher powers.	76
4-9	An example of an unnormalized, HR spectrum showing the shape imparted to the data by the crystal filter. The red line is the reference spectrum shown in Fig. 4-10.	77
4-10	A reference spectrum which is used to remove the crystal filter shape from HR data. This spectrum is the average of about 10,000 individual spectra.	78
4-11	The receiver response function fit to the residual structure of the truncated spectrum.	79
4-12	An example of a flat HR Power spectrum following truncation and the removal of the receiver response function.	80
4-13	A down-sampled HR power spectrum which has been fit with polynomials of degree 6, 7, and 9. The highest order of these three shows the best agreement with the residual shape of the spectrum.	81

4-14 Power contributed to f_j from a signal which is not at a basis frequency.	89
4-15 Exclusion limits on local density of non-virialized axions. The right axis is scaled to show density limits for DFSZ axions. All densities above the lines are excluded at 90% CL.	92
5-1 Close packing of 4 cavities within the old cavity volume. The smaller cavities have had their radii reduced by 50%. A) Top-down view showing empty space around the cavities. B) Side view showing full use of vertical space.	101
5-2 Close packing of 6 cavities within the old cavity volume. The smaller cavities have had their radii reduced by 54.2%. A) Top-down view showing less empty space around the cavities than for the 4-cavity configuration. B) Side view showing only partial use of vertical space.	102
5-3 Close packing of 21 cavities within the old cavity volume. The smaller cavities have had their radii reduced by 66.6%. A) Top-down view showing less empty space around the cavities than for either the 4-cavity or 6-cavity configurations. B) Side view once again showing full use of vertical space.	103
5-4 A picture of the split cavity prototype prior to being electro-polished.	105
5-5 A picture of the split cavity prototype after being electro-polished. At this stage the solder joint between the center plate and the wall was made for only one of the cells (top cell as pictured) in the cavity.	106
A-1 Diagram of the ADMX cavity and receiver chain. The power read out from the cavity is sent through both cold and room temperature amplification stages, is mixed down from radio frequencies to audio frequencies via two local oscillators (L.O.), and is ultimately saved to disk.	111
A-2 Monte Carlo simulations of non-virialized axion signals imposed on real HR data. The broader peak (velocity dispersion of 2×10^{-4}) is shown as an example of the narrowest signal that the MR channel can resolve. An excess power of 0 corresponds to power in that bin equal to the rms noise power at that frequency.	114
A-3 An example of a modulation-insensitive power spectrum after the crystal filter shape and amplifier chain effects have been divided out. The insert shows the raw power spectrum with a passband shape imparted by the crystal filter.	116
A-4 Density limits for the HR channel ($\Delta\beta \lesssim 3 \times 10^{-6}$) at 90% confidence from 812 MHz to 892.8 MHz. The scale for limits on KSVZ and DFSZ axions are shown on the left and right axes respectively. Density limits for the currently published MR channel data ($\Delta\beta \lesssim 2 \times 10^{-4}$)[14] are shown for comparison.	117

Abstract of Dissertation Presented to the Graduate School
of the University of Florida in Partial Fulfillment of the
Requirements for the Degree of Doctor of Philosophy

A MODULATION SENSITIVE SEARCH FOR NON-VIRIALIZED AXIONS
IN THE MILKY WAY HALO

By

Jeffrey D. Hoskins

May 2014

Chair: David B. Tanner

Major: Physics

Dark matter in the Milky Way could potentially be made of axions. If so, these axions would be detectable on Earth using an axion Haloscope. A search of this nature has been performed by the Axion Dark Matter eXperiment (ADMX). The high resolution (HR) data acquisition channel employed by ADMX is capable of searching for populations of non-virialized axions which are characterized by very low velocity dispersions. These populations are expected to lend structure to the Milky Way halo in the form of degenerate yet discrete flows.

The motivations for dark matter, axions as a dark matter candidate, and halo structure caused by non-virialized dark matter are considered, as are their expected properties. These properties guide the direction of the analysis of HR data. Details of the experimental apparatus are examined, and any analysis-limiting factors are noted. A discussion of axion signal properties follows, including the expected effects of signal modulation. Data from the HR channel is analyzed at multiple resolutions, with each resolution corresponding to a maximum velocity dispersion for dark matter axions. For each resolution, the data are corrected for systematic errors and searched for peaks above a candidate threshold in the thermal noise power spectrum. Frequency coincidence is looked for among all candidate signals at a given resolution, with an allowance made for signal modulation due to terrestrial motion. Failure of a signal to

be persistent, i.e. coincident in the majority of allowable power spectra, discounts it as being a real axion signal.

No candidate signals have shown statistically significant persistence over the frequency range of 812 to 892 MHZ. This corresponds to a mass range of $3.36 \mu\text{eV}$ to $3.69 \mu\text{eV}$. Exclusion limits are placed on the local density of non-virialized axions in this range. Lastly, hardware upgrades and changes to the search algorithm are considered as means of improving and/or expanding upon these limits.

CHAPTER 1

INTRODUCTION

Dark matter is an extra mass component found in large scale structures in the universe, such as galaxies and galaxy clusters, which interacts with the electromagnetic force feebly at best[1–3]. At present, there is a limited understanding of the fundamental nature of dark matter. Through experimental observations two facts have emerged. First, gravitational interactions heavily dominate both dark matter-dark matter interactions and dark matter-luminous matter interactions[4, 5]. Second, the mass of dark matter greatly exceeds that of luminous matter on the scale of both galaxies and galaxy clusters[6]. Beyond these two facts, the actual content of dark matter is still unknown, as is the detailed nature of its interactions. A wide range of particles have been suggested to fulfil in whole or part the roll of dark matter[7], each having unique production mechanisms, lifetimes, and physical interactions. Included in this list is the axion, which remains well motivated as a dark matter candidate and is the subject of the search presented in this work.

The Axion Dark Matter eXperiment (ADMX) is a terrestrial based experiment which searches for axions left over from production in the early universe. These relic axions would still be present today in the Milky Way dark matter halo in sufficient quantities to possibly be detected by current generation experiments[8–11]. ADMX uses the electromagnetic coupling of the axion, weak as it may be, for this detection, converting axions into photons within a cryogenic, microwave resonator. Power in the resonator is read out by a low-noise, high-gain radio receiver with the goal of seeing more power in the resonator than would be generated solely by black-body radiation[12].

The addition of a second data acquisition channel, the high resolution (HR) channel, in 2002 greatly expanded the capabilities of ADMX[13]. Halo axions which are a part of low-velocity-dispersion flows would generate a very narrow bandwidth signal in the ADMX receiver. The HR channel was designed with the intent of detecting these

narrow bandwidth signals. Reducing the bin width in the HR power spectrum has the effect of lowering the measured thermal noise power while preserving the signal power, provided that the bin width remains larger than the signal bandwidth. The HR channel complements the existing medium resolution (MR) data acquisition channel which searches for axions having a much larger velocity dispersion[14]. With the addition of the HR channel, ADMX gained the ability to probe not only the content of the dark matter halo but its structure as well.

This boon, however, is not completely without complications. The signal modulation in the HR channel due to terrestrial motion[15, 16] can easily cause an axion signal to drift out of the narrow power spectrum bins. This prevents the integration of signal power over multiple measurements; a procedure used by the MR channel to increase the signal to noise ratio of the data. As such, each power spectrum must be searched for axion signals independently. This is not to say that there is no correlation between spectra. The worst-case signal modulation is still smaller than the power spectrum bandwidth. Therefore, a real axion signal is expected to be found in the vast majority of spectra covering nearby frequencies. The range of this frequency neighborhood is determined by the elapsed time between the measurements of subsequent spectra.

The work presented here is that of a search for low-velocity-dispersion, or non-virialized, dark matter axions within the Milky Way halo. Following this introduction, Chapter 2 provides general background information on dark matter, axions, and potential halo structure. Motivations for and implications of each are considered. Then, Chapter 3 examines the experimental apparatus employed by ADMX, and the data acquisition system used for the HR channel. Chapter 4 discusses expected properties of axion signals, the HR data set, and the analysis procedure. Limits on the local density of non-virialized dark matter are derived from the HR data set. Next, Chapter 5 briefly describes improvements which can be made to the analysis procedure and to ADMX itself. Finally, Chapter 6 contains a summary and the conclusions of this work.

Included after the main body of this text are two appendices. Appendix A is a reprint of the body of reference [17]. In it, the same data set examined herein is analyzed at a spectral resolution of 10 Hz. Appendix B lists the code used in the present analysis.

CHAPTER 2

BACKGROUND INFORMATION

2.1 Introduction

ADMX is a direct detection dark matter experiment which seeks to identify axions as a part of Milky Way dark matter halo. This chapter aims to provide necessary background information on topics associated with ADMX. This will start with a discussion of the motivations for dark matter and observations which shed light on its known properties. Following this will be an overview of the axion including its origin and its roll as a dark matter candidate. Lastly, the possibility of low-velocity-dispersion flows being a part of the Milky Way halo structure will be considered, along with its implications.

2.2 Motivation for Dark Matter

Dark matter has been a topic of great interest to physicists and astronomers for many decades. This interest started when Fritz Zwicky published findings from galactic cluster observations in 1933. Using red shift measurements from the Coma Cluster, which is considered to be a typical example of his data set, he showed there to be a disparity between the measured spread of velocities of individual components of the cluster and that allowable given reasonable assumptions of the size and structure of the cluster. He assumed that the cluster had reached a steady state of motion and was subject to the Virial theorem such that

$$\overline{T} = -\frac{1}{2}\overline{V} \quad (2-1)$$

where \overline{T} is the average kinetic energy of the cluster and \overline{V} is its average potential energy. He further assumed a uniform mass distribution of 8×10^{11} solar masses, based on luminosity observations, with a cluster radius of 1 Mly. The value for \overline{V} calculated from these estimates led to an rms velocity spread of $\mathcal{O}(80 \text{ km/s})$ as opposed to the observed $\mathcal{O}(1000 \text{ km/s})$. From this, he concluded that were there only luminous matter

present, the Coma Cluster would lack the necessary mass to sustain its structure and would simply fly apart becoming a field of galaxies rather than a cluster [1]. Four years later, using more accurate measurements of the Coma Cluster, Zwicky calculated a lower limit on its average mass to light ratio, γ . He obtained a value of

$$\gamma = 500 \frac{M_{\odot}}{L_{\odot}}, \quad (2-2)$$

where M_{\odot} and L_{\odot} are the mass and luminosity of the sun respectively, which was found to be two orders of magnitude larger than expected[2]. This supported his initial conclusion that the Coma Cluster not only contained non-luminous matter, but that this "dark" matter comprised a nontrivial quantity of the cluster[1].

Despite using a value of 500 km/s/Mpc for the Hubble constant rather than the currently accepted value of 70 km/s/Mpc[3], Zwicky's conclusion still proved to be consistent with the findings of Vera Rubin and Kent Ford some 40 years later. Rubin and Ford measured the rate of rotation of spiral galaxies; first for just the Andromeda galaxy[18] in 1970 but soon after for dozens of other spiral galaxies[19, 20]. They found that the rate of rotation does not fall off as r^{-1} as one would expect would occur at radii extending past the galactic bulge. Instead, as can be seen in Fig. 2-1, these curves became asymptotically flat at large radii. This indicated that there must be the presence of some unseen additional mass component. From this it can be inferred that the dark matter in these galaxies not only extends well beyond the galactic bulge but also increases in abundance roughly as r^2 .

Recent observations of a galaxy cluster collision in the Bullet Cluster [4, 5] provide further indirect evidence for the existence of dark matter. In such a collision, the stellar components of individual galaxies, which constitute $\sim 1 - 2\%$ of the cluster mass, are essentially collisionless while intergalactic gasses, constituting $\sim 5 - 15\%$ of the cluster mass, instead behave like a fluid and bunch up in the geometric middle of the collision. Were there no additional mass component in these clusters and noting that the gas

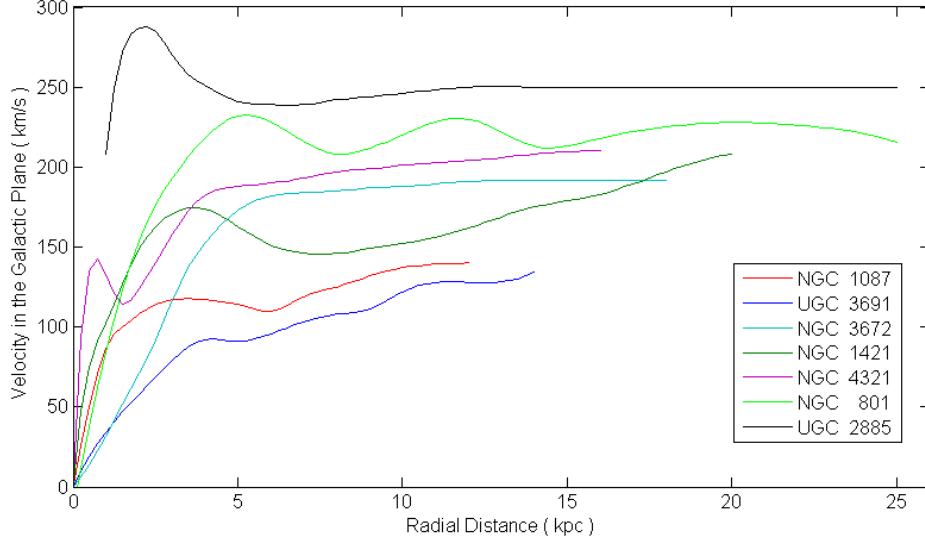


Figure 2-1. Rotation curves for 7 of the 21 spiral galaxies measured in Rubin et al. [19]. Following the sharp rise caused by the core of the galaxy, these curves become asymptotically flat.

component is about 5 times as massive as the stellar component [21–23], the center of the collision is where one would expect to see the center of mass of the system. Using gravitational lensing, this was shown not to be the case. The centers of mass of the two clusters remained distinct and were found to physically align with the stellar mass components of each. Figure 2-2 shows the Bullet Cluster with the locations of the centers of mass clearly marked. Outweighed by the intergalactic gas, the stellar component alone can not account for these center of mass locations. However, a collisionless dark matter component most certainly would. Additionally, these results can not be explained using Modified Newtonian Dynamics [24, 25], thus providing some of the most compelling evidence to date in favor of dark matter.

2.3 Dark Matter Properties

With indirect evidence of dark matter now established, consideration can be given to the properties expected of all dark matter candidates as they are pertinent to any sort of experimental search. These properties should be consistent with the current model for the make up of matter and energy in the universe, called the Λ Cold Dark Matter (



Dark Matter in Galaxy Cluster MACS J00254.4-1222
Hubble Space Telescope ACS/WFC • Chandra X-Ray Observatory

NASA, ESA, CXC, M. Bradač (University of California, Santa Barbara), and S. Allen (Stanford University)

STScI-PRC08-32

Figure 2-2. A NASA image of the Bullet Cluster merger. X-ray emission from gas is shown in pink, and the regions of highest density are shown in blue.

Λ CDM) model. In Λ CDM, luminous matter makes up about 5% of the universe; cold, i.e. non-relativistic, dark matter makes up about 22%; and dark energy, which is responsible for the accelerated expansion of the universe, is parameterized by the cosmological constant λ and makes up the remaining 73% [3]. Dark energy is beyond the purview of this work and will just be taken as a given for Λ CDM.

The first property for all dark matter candidates is that they must be considered "dark." That is, couplings between dark matter and force carriers, luminous matter,

or even other dark matter must be overwhelmingly dominated at large distances by gravitational interactions. Clearly, electromagnetic couplings must be extremely feeble as coupling of any notable strength would permit scattering of photons in quantities large enough to be directly detected. Weak scale couplings must also be quite small. This is not to say that all interactions via the weak force are forbidden; indeed, some dark matter searches rely heavily on such interactions [26, 27], just that they are grossly subdominant compared to gravity. Simply put, a dark matter candidate must be effectively collisionless.

Second, dark matter candidates must be non-relativistic. The criterion for cold dark matter is a result of evidence from observations made by the Wilkinson Microwave Anisotropy Probe (WMAP) which indicate a bottom-up formation for structure in the universe. WMAP found that galaxy clusters, the largest structures in the universe, formed more recently than their constituent galaxies[3]. In a bottom-up formation, tiny density perturbations in the more or less isotropic and homogeneous early universe, shown in Fig. 2-3, would slowly grow due to accretion. As the universe cools, gas would coalesce and larger structures such as stars would begin to form. These stars then form galaxies, which in turn form clusters. All large-scale structure is thus a result of small localized over-densities in the early universe. Were dark matter relativistic, or hot, this sort of structure formation would not occur. Hot dark matter, still collisionless from criterion one, would rapidly move from regions of high density to regions of low density. This redistribution of dark matter would occur faster than the accretion rate of matter near the, now short lived, over-density. Any density perturbations smaller than the free streaming length of hot dark matter would be erased, leaving no seeds from which the presently observed localized inhomogeneity of the universe could have formed [6].

The last property that constrains all dark matter candidates is that dark matter must be non-baryonic. This is inferred from the ratios of the relic abundances of light elements found in high redshift celestial objects [28]. Specifically, the abundances of Deuterium

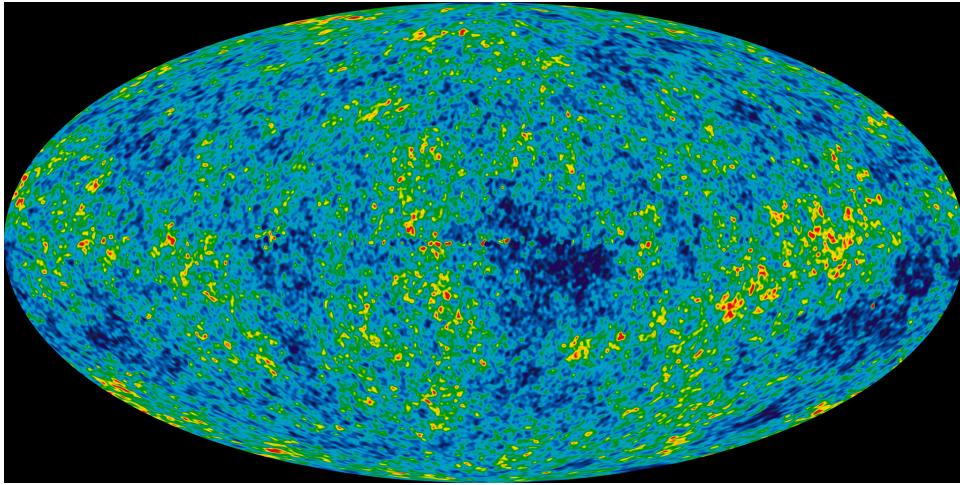
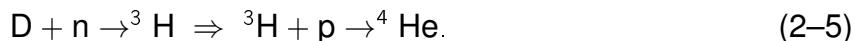


Figure 2-3. A NASA image of the CMB temperature map as measured by WMAP. On large scales the CMB shows a high degree of homogeneity and isotropy. The localized regions of higher and lower temperatures are indicative of small density perturbations in the early universe.

(D), ^3He , ^4He , and ^7Li are compared to that of ^1H . The abundances found in these distant objects are indicative of what was produced during primordial nucleosynthesis. Abundances of D and ^4He are highly sensitive to the proton and neutron densities in the primordial plasma during nucleosynthesis, with higher densities leading to a more complete conversion of D to ^4He via any of the following processes:



The expected values for the ratios of these light elements are in good agreement with measured values for a universe in which protons and neutrons are only a minority component at less than 5% of the total density of the universe. Assuming ΛCDM , the normalized total energy density of the universe given by

$$\Omega_T = \Omega_\nu + \Omega_B + \Omega_{DM} + \Omega_\Lambda, \quad (2-6)$$

where, listed in ascending order by contribution to Ω_T as calculated by WMAP, $\Omega_\nu < 0.0286$ at 95% CL is the neutrino density, $\Omega_B = 0.0449 \pm 0.0028$ is the baryon density, $\Omega_{DM} = 0.222 \pm 0.026$ is the dark matter density, and $\Omega_\Lambda = 0.734 \pm 0.029$ is the dark energy density. Further, WMAP reports $\Omega_T = 1.080^{+0.093}_{-0.071}$, indicating that the geometry of the universe is flat [3].

While numerous dark matter candidates both within and beyond the bounds of the standard model have been proposed, most of them can be excluded due to one or more of the criteria listed above [7]. Surviving this initial paring are the class of particles called Weakly Interacting Massive Particles (WIMPs) and the axion. The leading WIMP candidates are both stable partner particles predicted by Super-Symmetry (SUSY), the neutralino [6] and the photino [29], while axions are result of a spontaneous symmetry breaking in Quantum ChromoDynamics (QCD). The relic abundance of each is determined by the energy scale at which they thermally decouple, which is determined by mass, and their interaction cross-sections once decoupled. For sufficiently low cross-sections, relic densities consistent with $\Omega_{DM} \approx 0.22$ can be expected for WIMPs with a mass of order 100 GeV and for axions with mass $m_a \gtrsim 10^{-6}$ eV. Several experiments search for WIMPs by looking for nuclear recoils either in cryogenic solid-state detectors or in a liquid noble-gas. As of this writing, no data solidly supporting the existence of SUSY has been produced by the Large Hadron Collider [30] and no positive WIMP signal has been detected [31–34]. Events which can not be explained by expected backgrounds have been reported by CDMS and CRESST, though in both cases the statistical significance of the unexplained events is insufficient to claim detection [26, 27]. Not relying on SUSY, axions still remain a strong contender as a dark matter candidate. ADMX searches for axions of in the μ eV mass range, guided by different models for the structure of the Milky Way dark matter halo.

2.4 QCD and the Strong CP Problem

QCD was fueled by observations in 1970 which gave new insight into the nature of the Strong sector of the standard model and the rules which govern it. Inspired by the formalism of Quantum ElectroDynamics (QED), QCD was developed to be a similar formalism which describes the interactions between quarks and gluons. Whereas QED is a $U(1)$ gauge theory, invariant under unitary transformations in one dimension (ie. phase transformations), QCD is an $SU(3)$ gauge theory, invariant under transformations by three dimensional matrices which have a real determinant of value 1. The gauge bosons for this $SU(3)$ symmetry, gluons, are the 8 mixing states made from conserved color charges in QCD, labeled as red, blue, and green. Combining all of this, the QCD Lagrangian is

$$\mathcal{L} = -\frac{1}{4} G_{\mu\nu}^a G^{a\mu\nu} + \sum_j^N \left(i \bar{q}_j \gamma^\mu D_\mu q_j - \left(m_j q_{Lj}^\dagger q_{Rj} + H.C. \right) \right) + \frac{\theta g^2}{16\pi^2} G_{\mu\nu}^a \tilde{G}^{a\mu\nu}, \quad (2-7)$$

where $G_{\mu\nu}^a$ is the gluon strength tensor, q_j are the quark fields, γ_μ are the gamma matrices, D_μ is the covariant derivative, m_j are the quark masses, g is the color coupling, and θ is an angle which may be between 0 and 2π . The summation index j is representative of the 6 quark flavors, while the subscripts L and R respectively refer to the left-handed and right-handed chiralities of the quark fields. Lastly, the hermitian conjugate of the chiral term is denoted by H.C. While Eq. 2-7 dictates such things as the interaction strength and chirality of quarks, it is important to remember that QCD is still only a single part of the Standard Model. Consequently, the quark masses seen in Eq. 2-7 actually originate from symmetry breaking in the Weak sector.

The θ -term in Eq. 2-7 is a consequence of the non-abelian nature of QCD, which permits the existence of an infinite number of independent vacuum states, $|n\rangle$. Each of these states satisfies the least action principal for Eq. 2-7, and transitions between them can only occur via quantum tunneling events called instantons. However, a gauge invariant vacuum state can be constructed as a linear combination of these vacua. This

invariant state is given by,

$$|\theta\rangle = \sum_n e^{in\theta} |n\rangle, \quad (2-8)$$

and corresponds to the θ -term in Eq. 2-7. The angle θ in 2-8 can be of any value from 0 to 2π , and governs the degree to which tunneling occurs between vacua.

A classical treatment of the QCD lagrangian, which assumes all quarks to be massless, permits there to be a chiral $U(1)_A$ symmetry. However this symmetry is broken in a full quantum treatment due to The Adler-Bell-Jackiw anomaly [35, 36]. This anomaly breaks the $U(1)_A$ symmetry by introducing an axial current,

$$J_5 = \sum_j^N \bar{q}_j \gamma^\mu \gamma_5 q_j, \quad (2-9)$$

$$\partial_\mu J_5^\mu = 2 \frac{Ng^2}{16\pi^2} G_{\mu\nu}^a \tilde{G}^{a\mu\nu}. \quad (2-10)$$

This has the effect of changing θ by an amount $N\phi$ when a chiral rotation of $\frac{\phi}{2}$ is applied to the quark fields. Such a rotation is required to force the quark masses to be real.

Once the quark mass matrix, \mathbf{M} , is diagonalized it is of the form,

$$\mathbf{M}_{jk} = |m_{jk}| \delta_{jk} e^{i\phi}, \quad (2-11)$$

whereupon the phase ϕ must be removed. A chiral rotation of the quark fields by $\frac{\phi}{2}$ accomplishes this yielding,

$$q_j \rightarrow q_j e^{i\gamma_5 \frac{\phi}{2}} \quad (2-12)$$

$$m_j \rightarrow m_j e^{-i\phi} \quad (2-13)$$

$$\theta \rightarrow \theta + N\phi = \theta + \arg \det \mathbf{M} \quad (2-14)$$

The quark mass phases are absorbed by θ creating an invariant term,

$$\bar{\theta} = \theta + \arg \det \mathbf{M}, \quad (2-15)$$

which now replaces θ in the last term of 2-7.

For all non-zero values of $\bar{\theta}$ the θ -term introduces a violation of parity symmetry in strong interactions. Since charge conjugation symmetry is fully preserved in QCD, the combined charge-parity symmetry (CP) would be violated as well. The degree of CP violation dependant directly on the value of $\bar{\theta}$. Because $\bar{\theta}$ is invariant, it is observable and can thus be measured. Perhaps the best measurement of $\bar{\theta}$, and by extension the degree to which QCD violates CP, comes from the limits placed on the electric dipole moment, $|d_n| \lesssim 3 \times 10^{-26} e \text{ cm}$, where e is the electron charge. Since

$$|d_n| \simeq 10^{-16} \bar{\theta} e \text{ cm}, \quad (2-16)$$

a limit of $\bar{\theta} \lesssim 10^{-10}$ can be established [37]. Such a small value for $\bar{\theta}$, equivalently for CP violation, is rather peculiar. Both QCD and weak scale interactions permit CP violation through chiral transformations. So why should QCD preserve CP while the weak sector does not? This is the essence of the strong CP problem.

2.5 PQ Symmetry and the Axion

In 1977, an elegant solution to the strong CP problem was put forward by Peccei and Quinn. They proposed that by introducing a new spontaneously broken $U(1)_{PQ}$ chiral symmetry, θ would become a dynamical variable rather than just a parameter. The potential of the Higgs field becomes dependant on θ creating a unique minimum and forcing a particular set of phases on the quark masses. When the masses are made real via chiral rotation, $\arg \det \mathbf{M}$ naturally becomes $-\theta$. CP conservation is then allowed because $\bar{\theta}$ has automatically relaxed to 0 [38, 39].

Individually, both Weinberg[40] and Wilczek[41] pointed out that the Peccei-Quinn solution predicted the existence of a Nambu-Goldstone boson. Wilczek dubbed this new particle the axion. This electrically neutral particle is a result of the spontaneous breaking of the $U(1)_{PQ}$ symmetry. In addition to the spontaneous breaking of $U(1)_{PQ}$ which occurs at an energy scale f_a , an explicit breaking of the symmetry gives the axion a potential. Oscillations about the minimum of this potential result in the axion acquiring

a mass. The axion field can be written so as to absorb $\bar{\theta}$, such that $a(x) = f_a \bar{\theta}$, once again changing the θ -term in Eq. 2–7. Considering specifically the decay $a \rightarrow \gamma\gamma$, which is the interaction predominantly used for axion detection experiments, the θ -term takes the form

$$\mathcal{L}_{a\gamma\gamma} = g_{a\gamma\gamma} a(x) F_{\mu\nu} \tilde{F}^{\mu\nu}, \quad (2-17)$$

where $F_{\mu\nu}$ is the electromagnetic field strength tensor, and

$$g_{a\gamma\gamma} = \frac{\alpha g_\gamma}{\pi f_a}. \quad (2-18)$$

Here, α is the fine structure constant, and g_γ is a model dependant constant of order one. At present, the leading models are the Kim-Shifman-Vainshtein-Zhakharov (KSVZ) model with $g_\gamma = -0.97$ [42, 43], and the Dine-Fischler-Srednicki-Zhitnitsky (DFSZ) model with $g_\gamma = 0.37$ [44, 45]. Details of these models will be discussed later.

Independent of any models, simple relations to f_a can be written for some properties of the axion, namely its mass and lifetime. For the axion mass, m_a , the expression includes the masses and decay constants of other known particles and is given by

$$m_a = \frac{f_\pi m_\pi \sqrt{m_u m_d}}{f_a (m_u + m_d)}, \quad (2-19)$$

where $f_\pi = 93$ MeV is the pion decay constant, and $m_\pi = 135$ MeV, $m_u \simeq 2.3$ MeV, and $m_d \simeq 4.8$ MeV are the neutral pion, up quark, and down quark masses respectively[46, 47]. Evaluating Eq. 2–19 yields the more convenient form,

$$m_a \approx 6 \mu eV \left(\frac{10^{12} GeV}{f_a} \right). \quad (2-20)$$

While there are no a priori constraints on the mass of the axion, it remains dependent on f_a . The lifetime of the axion for the $a \rightarrow \gamma\gamma$ decay channel is given by

$$\tau_a = \left(\frac{10^5 eV}{m_a} \right)^5 \text{sec.} \quad (2-21)$$

From Eqs. A–3, 2–20, and 2–21 it is seen that a larger f_a would yield a very nice dark matter candidate as the axion would be an extremely light, extremely weakly interacting, and extremely long lived particle. On the other hand, a smaller f_a would yield an axion which, while more detection friendly due to stronger coupling, is clearly not a dark matter candidate.

A slight complication to the Peccei-Quinn solution is that one or more additions must be made to the standard model in order to accommodate the creation of a new particle. For the original Peccei-Quinn-Weinberg-Wilczek (PQWW) model axion, a second Higgs field was added [40, 41]. The two Higgs fields, Φ_u and Φ_d , would couple to up-type quarks and down-type quarks respectively. Electroweak symmetry breaking results in two charged and two neutral Nambu-Goldstone bosons. The charged bosons and one of the neutral bosons are consumed by the Higgs mechanism, contributing to the masses of W^+ , W^- , and Z . The remaining neutral boson is the axion. The Higgs fields would have vacuum expectation values,

$$\langle \Phi_u^0 \rangle = v_u \quad (2-22)$$

$$\langle \Phi_d^0 \rangle = v_d, \quad (2-23)$$

which satisfy

$$v = \sqrt{v_u^2 + v_d^2} \equiv f_a, \quad (2-24)$$

though the ratio, $\frac{v_u}{v_d}$, is left as a free parameter. For PQ symmetry breaking at the weak scale, ie. $f_a = (\sqrt{2}G_F)^{-\frac{1}{2}} \simeq 250$ GeV where G_F is the Fermi constant[46], the axion mass would be $m_a \simeq 100$ keV. With such a mass, the axion would have a lifetime of order seconds, and would be easily detectable in the laboratory[48]. Consequently, beam dump experiments quickly ruled PQWW axions [49], prompting the development of a different means of incorporating the axion into the standard model.

The current state of axion model evolution allows f_a to be orders of magnitude higher than the weak scale. This has the effect of drastically reducing the mass and

coupling constants of the axion as they are both inversely proportional to f_a . Being extremely difficult to detect, this version of the axion has been aptly named the invisible axion. As mentioned earlier, the two leading invisible axion models are the KSVZ model and the DFSZ model. In the KSVZ model, instead of using a second Higgs field to create the axion field, the accommodation in the standard model comes from the addition of a singlet scalar field and a heavy singlet quark [42, 43]. The KSVZ axion is the phase of this new scalar field. The energy scale of the singlet field is not constrained to the weak scale, thus permitting f_a to be arbitrarily large. Couplings for the KSVZ axion are generated by interactions with the heavy quark. This interaction scheme leads to KSVZ axions sometimes being referred to as hadronic. The DFSZ model also adds a singlet scalar field to the standard model, though in lieu of adding a heavy quark it repurposes the second Higgs field from the PQWW model [44, 45]. The DFSZ axion is, once again, the phase of the scalar singlet. The second Higgs field, as it no longer creates the axion field, is instead used to help generate couplings to the DFSZ axion. The scalar singlet interacts with other fields using both Higgs fields as intermediaries. The different values of g_γ stated earlier stem from the different ways by which couplings are generated. Despite being more difficult to detect, the DFSZ axion does have one advantage over the KSVZ axion. That is, the DFSZ model is compatible with current extensions to the standard model aimed at generating a Grand Unified Theory, whereas the KSVZ model is not.

2.6 Astrophysical Considerations and Cosmology

Current bounds on f_a for the KSVZ and DFSZ models come from a combination of astrophysical observations and cosmological constraints. Axion production in stars comes primarily from the Primakoff effect [50, 51], whereby a photon is converted to an axion by interacting with a strong magnetic field. These axions can then play a role in cooling the star or, as discussed below, the stellar remnant. This leads to a bound based on limiting the effectiveness of stellar cooling via axions. On the other end of the

spectrum, axions produced in the early universe, relic axions, would still be around today and would contribute to the total Ω as part or all of $\Omega_D M$ [8–11]. The requirement that $\Omega = 1$, leads to a bound based on the relic density of axions not exceeding current limits on $\Omega_D M$.

First, consider the lower bound on f_a , or the upper bound on m_a . This bound is set by the duration of the neutrino burst from SN1987a. Emission of neutrinos is the dominant means of cooling the collapsing iron core during a supernova. The time scale of this emission is governed by how long it takes neutrinos to migrate to the surface of the core. As the mean free path of neutrinos is much smaller than the size of the core, this diffusion time is not trivial. Neutrinos originating near the center of the core will end up leaving the star some noticeable time later than those originating closer to the surface. This process has been modeled for supernovae and is expected to take about 5-10 seconds. The neutrino burst measured from SN1987 matched the predicted result, consisting of 19 neutrino detections over about 10 seconds [48].

What would happen if axions played a role in the cooling of the core during its collapse? Clearly the time required for cooling would go down since there are now more particles carrying away energy. In this case, the neutrino burst too would be shorter. Because the measured data from SN1987a matches the predicted values, it is concluded that axions do not contribute significantly to cooling. This leaves two possibilities; axions free stream out of the core but are produced in very small quantities, or axions are produced copiously but are trapped since their mean free path is so small. The former case sets a bound of $f_a \gtrsim 10^9$ GeV or $m_a \lesssim 5$ meV. A lower value for f_a would increase electromagnetic coupling, and consequently axion production, to a point where they begin to take away an appreciable amount of energy when free streaming out of the core. The latter case, however, still allows for trapped axions where $f_a \lesssim 10^7$ GeV or $m_a \gtrsim 0.5$ eV. Fortunately, a similar argument can be applied to the cooling of horizontal branch red giants.

If axions were to play a role in the cooling of red giants then there would be a noticeable change in the helium-burning lifetimes of these stars. Specifically, these lifetimes would be shorter than in the case where axions do not contribute. Models of stellar evolution which include axion cooling predict lifetimes of order 10^7 years while those which ignore axion cooling predict lifetimes of order 10^8 years. The latter prediction is in good agreement with observed helium-burning lifetimes [52]. This leads to an identical structure for the excluded values of f_a . There is an excluded range bounded above by an allowed free-streaming range and below by an allowed trapped-axion range. This excluded range covers the trapped-axion range associated with SN1987a, while the trapped-axion range associated with red giants is excluded by the aforementioned laboratory experiments.

Now consider the upper bound on f_a . This bound is set by the requirement that the energy density of relic axions, Ω_a , should not force $\Omega_T > 1$. That is Ω_a , being a part of $\Omega_D M$, should not exceed a value of 0.222 set by WMAP. This value for Ω_a is, of course, the energy density of the axion field as it is today after having been effected by hubble expansion. When the universe has cooled to the QCD scale (~ 200 MeV), the axion field begins oscillating about the minimum of its potential, yielding a mass for the axion. For a zero momentum axion field, the equation of motion for these oscillations is given as

$$\ddot{a} + 3H(t)\dot{a} + m_a^2 \sin(a) = 0, \quad (2-25)$$

where $H(t)$ is the Hubble parameter[38, 39]. The energy density stored in the field oscillations is

$$\rho_a = \frac{1}{2}\langle \dot{a} \rangle^2 + \frac{1}{2}m_a^2\langle a \rangle^2 \simeq \frac{1}{2}m_a^2 f_a^2 \bar{\theta}_i^2, \quad (2-26)$$

where $\bar{\theta}_i$ is the initial misalignment angle of the axion field. Accounting for changes due to Hubble expansion, the present day energy density of the axion field is

$$\Omega_a \simeq 0.15 \left(\frac{f_a}{10^{12} \text{GeV}} \right)^{\frac{7}{6}} \bar{\theta}_i^2. \quad (2-27)$$

The exact value of $\bar{\theta}_i$ depends on whether the axion acquired mass before or after inflation, but in general can be taken as being of order unity. From this it is clear that for $\Omega_a < 0.222$, $f_a < 10^{12}$ GeV or $m_a \gtrsim 1 \mu\text{eV}$. Having now established an axion mass range of $1 \mu\text{eV} \lesssim m_a \lesssim 5 \text{ meV}$, the lifetime of the axion can be found from Eq. 2–21. At such a low mass, the axion lifetime is longer than the present age of the universe. Were $f_a > 10^{12}$, there would be increased axion production in the early universe; and, with the energy of their oscillations not dissipating significantly over their extremely long lifetime, the relic density of axions would be too large to be consistent with Λ CDM [48].

2.7 Halo Structure

The structure of dark matter halos, much like dark matter itself, leaves several questions unanswered. Is the structure trivial, ie. isotropic and homogeneous, or something more complex, containing regions of over- and under-densities? The total mass of a particular halo may be known, but what is known about the velocity distribution of dark matter in that halo? How much angular momentum does the halo posses? While there are many models of halo structure [53–56] which attempt to make predictions about how large scale structures have evolved to behave the way they do today, verification of these models proves to be rather difficult[57]. Often there are large unknowns in these models which must either be assumed or omitted entirely from halo formation simulations [58, 59].

The Λ CDM model serves as a base for halo structure models. So what does Λ CDM alone say about halo structure? It turns out that nontrivial halo structure is a natural consequence of Λ CDM, independent of what particle or particles make up dark matter. This can be inferred from two facets of Λ CDM. First, large scale structures formed due to accretion of mass about small density perturbations. Second, dark matter is collisionless. The first point clearly implies that large structures may have incorporated a number of smaller structures during accretion, while the second point requires a little more scrutiny before its significance becomes evident.

To expand on the first point, consider a number of small nearby density perturbations in the early universe. After some time, each has accreted matter and is, in its own right, a distinct structure. This, of course, does not preclude the mutual gravitational attraction of these structures. This group of structures begins to coalesce, accreting now about its center of mass. In the end, the resulting larger structure, say a galaxy sized halo, will contain regions of localized over-densities, or subhalos. Were one of these subhalos much more massive than its neighbor, perhaps due to being older or being seeded by a larger density perturbation, tidal forces would bleed off dark matter from the neighbor. The result is the creation of a flow of dark matter between the two subhalos. This logical argument can be scaled up, applying now to nearby galaxy sized halos having disparate masses.

Gaining insight into the second point starts by looking at the phase-space distribution of dark matter in the early universe [60]. Given the absence of a significant gravitational potential at that time, dark matter would have a continuous, nearly uniform spatial distribution and a velocity dispersion of $\delta v \lesssim 10^{-12} c$. This is best described as an extremely thin, 3-dimensional sheet in phase space which remains continuous as it evolves. Because dark matter is collisionless, the evolution of the phase space sheet is solely determined by gravitational interactions. Once structures start to form, their potentials cause the sheet to warp. In the case of larger structures like galaxies, the sheet may even fold over on itself and begin to wind up [61]. This sort of behavior is diagramed in Fig. 2-4.

A consequence of this winding is that at a given point in space there may be several distinct populations of dark matter, each having a different velocity vector. Each of these populations is a low-velocity-dispersion dark matter flow which is either falling into or out of the galaxy. A convenient labeling scheme is to denote whether the flow is entering or leaving the galaxy, and to count the number of times it has fallen in that direction. For example, dark matter which has just entered the galaxy is falling in for the first time. After

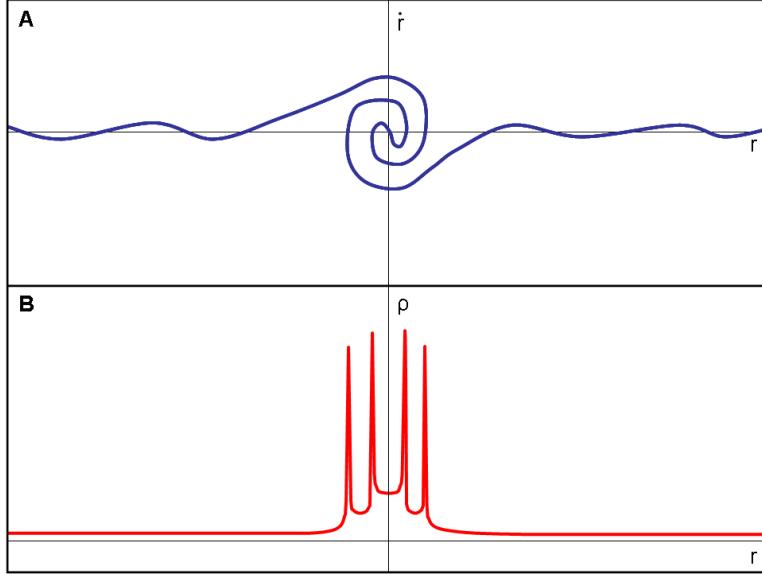


Figure 2-4. A) A 2D representation of phase-space winding in the rest frame of a strong gravitational potential. Here r is the velocity of dark matter in the direction of r . As time progresses this clockwise winding continues, resulting in several degenerate populations of dark matter in the neighborhood of the potential. B) The projection of phase-space winding onto physical space. Note the significant increases in density at each turn-around point. In the limit of $\delta v = 0$ the density at these points would be infinite, though for small dispersions the increased densities remain finite.

passing the galactic center, it is falling out for the first time. Upon reaching a turn-around point it begins falling in for the second time, and after passing the center again it is falling out for the second time. Consider also the case of dark matter falling in with a large amount of angular momentum. Such flows would fall in, reach a point of closest approach to the galactic center, and then fall out again. These flows would eventually turn around at some large distance and fall in once again.

There is no upper bound on the number of discrete flows this winding would cause, however in the limit of a large number of flows, such as near the center of a fairly old galaxy, the dark matter velocity distribution would approach that of a Maxwellian. These densely packed flows will have been in the galaxy for a long time and, through gravitational interactions, they will have reached a state of stable, dynamic equilibrium. Such dark matter is referred to as being virialized. In addition to virialized dark matter,

there would also be a dark matter population which has only recently entered the galaxy. This would consist of dark matter contained in flows which have fallen in to or out of the galaxy, at most, only a few times. This late-infall dark matter, so named due to entering the galaxy late in its evolution, will still appear in discrete flows having had insufficient time to reach the stable equilibrium of virialized dark matter. As such, late-infall dark matter is referred to as non-virialized.

The turn-around points of dark matter flows occur at the locations where the phase-space sheet actually folds over [55]. At both types of turn-around points, projecting phase-space onto physical space clearly shows the caustics generated by dark matter flows. This phase-space projection can be seen in Fig. 2-4. Caustics at the points of closest approach and those at distant turn-around points are called inner and outer caustics respectively. Outer caustics take the form of spherical shells, while inner caustics form tube-like rings with a tri-cusp, or elliptic umbilic catastrophe, shaped cross-section [62]. In the neighborhood of any of these caustics there would be a higher than average local dark matter density [63].

Both localized flows, caused by tidal disruptions, and galactic scale flows, caused by in-falling or out-falling dark matter, may be present in the Milky Way. The local density of late-infall dark matter flows has been estimated for the first 8 passes through the Solar system [64]. Each of these flows would have a density of roughly 2% of the average Milky Way dark matter density. With 8 passes, each going in 2 directions and having 2% of the average density, it is clear that these flows may constitute a significant fraction ($\sim 32\%$) of the local dark matter density. These estimates use an average local density of $0.52_{-0.17}^{+0.21} \text{ GeV/cm}^3$ as determined by Gates, Gyuk, and Turner [65] leading to flow densities of order 0.01 GeV/cm^3 . Note however, that ADMX chooses to use a more conservative estimate of average local density, 0.45 GeV/cm^3 . This does not effect the flow density estimates presented above, only the fractions of the average density they represent. Further, the earth may be in the neighborhood of an inner caustic caused

by flows which are falling into and out of the Milky Way for the fifth time [63]. These two flows would have enhanced densities in this region which are expected to be 0.95 GeV/cm³ and 0.08 GeV/cm³. The spatial degeneracy of the flows makes it impossible to assign specific flow directions (in or out) to these densities, though such an assignment is not required by ADMX. The larger of the two flows is referred to as the "Big Flow" and is expected to have velocity relative to the sun of ~ 480 km/s and velocity dispersion of $\lesssim 53$ m/s.

Designing a search which is tailored to looking for dark matter flows provides an interesting way to probe the structure of the Milky Way halo. This was the goal when developing and implementing the ADMX HR channel. The spectral resolution of the HR channel is roughly 3 orders of magnitude better than that of the medium resolution (MR) channel, used for detecting virialized axions. This allows for the identification of low-velocity-dispersion flows as distinct and independent peaks found within the HR power spectra. Because the widths of these peaks are proportional to δv , they are expected to form very sharp features. In MR spectra, the integrated thermal noise power contained in one frequency bin would be 1000 times higher, potentially masking the narrow peaks caused by dark matter flows. So long as the bin width of the HR power spectra is not significantly wider than the width of the sought after peak, there is a net gain in the ratio of signal power to thermal noise power. This makes the HR channel an effective tool in the search for non-virialized dark matter.

CHAPTER 3

EXPERIMENTAL APPARATUS AND DATA ACQUISITION SYSTEM

3.1 Introduction

ADMX searches for dark matter axions within the Milky Way halo using an axion haloscope kept at cryogenic temperatures and coupled to a low-noise amplifier chain. A diagram of the experiment showing the haloscope, cryogenics, and output electronics can be seen in Fig. 3-1. This chapter will cover the experimental setup and data acquisition system (DAQ) used for ADMX. First will be an overview of the general design and function of haloscopes. Then there will be a description of the cryogenic hardware which includes the ADMX cavity, the main magnet, and the bucking magnet. Finally, the components of the receiver chain will be discussed along with their effect on system noise and/or signal properties.

3.2 Axion Haloscopes

As mentioned earlier axions have a lifetime longer than the present age of the universe. Since we have no way to increase the number density of axions, relying solely on random chance for a decay into two photons is simply not practical. Any hope of detection requires a way to tip the scales in our favor. In the early 1980's Pierre Sikivie proposed the axion haloscope which would accomplish this using the inverse Primakoff effect [66, 67].

Physically, an axion haloscope consists of a resonant cavity of arbitrary shape within a static homogeneous magnetic field. The interaction of note for an axion haloscope is that of an axion decaying into two photons and is given by

$$\mathcal{L}_{a\gamma\gamma} = g_{a\gamma\gamma} a(x) \mathbf{E} \cdot \mathbf{B}, \quad (3-1)$$

where $g_{a\gamma\gamma}$ is given by Eq. A-3, $a(x)$ is the axion field, \mathbf{E} is the electric field of one of the decay photons, \mathbf{B} is the magnetic field of the other. Though, within the haloscope the static magnetic field acts as a virtual photon as shown in Fig. 3-2. This dependance on

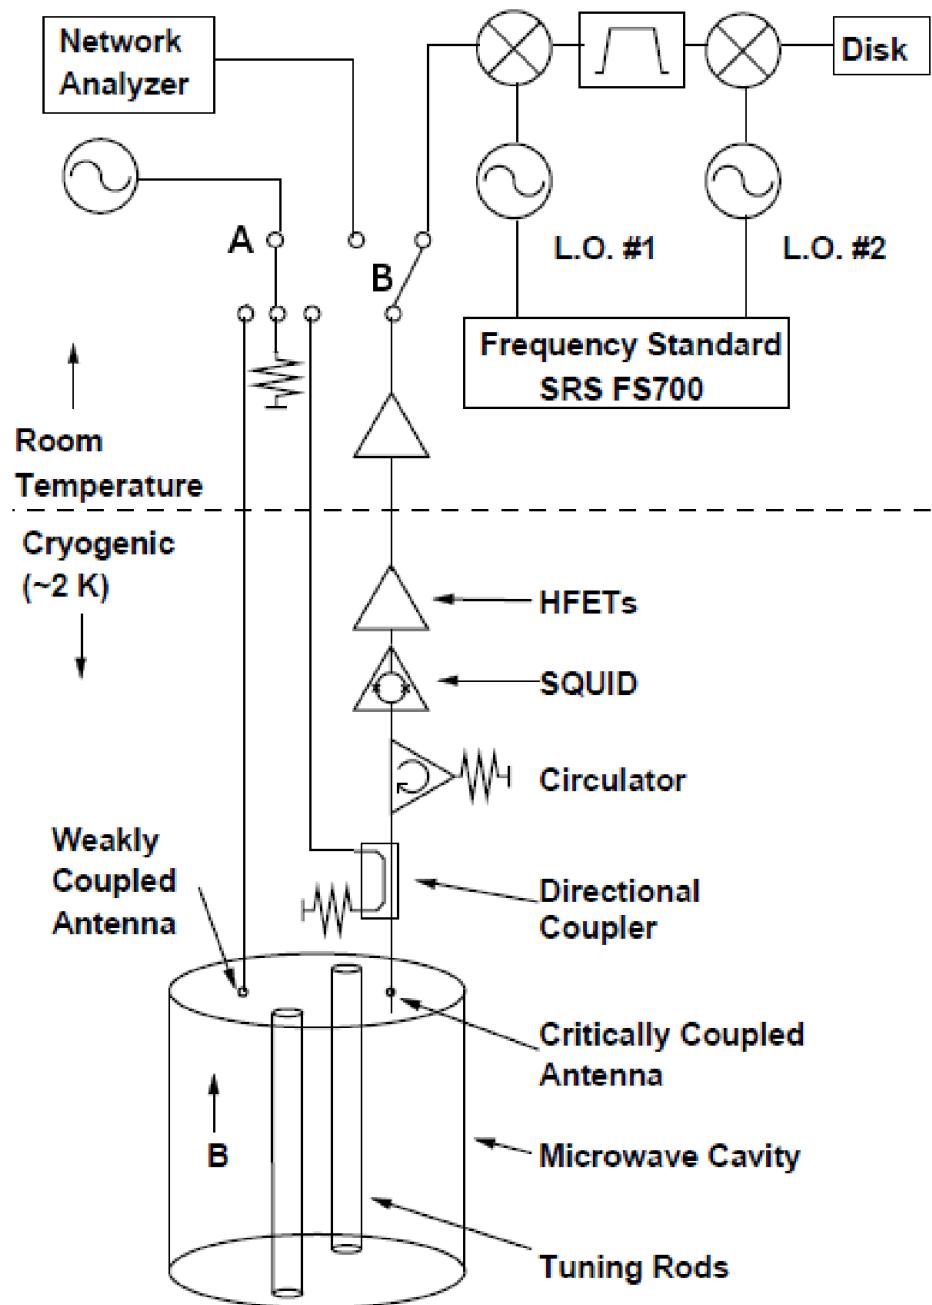


Figure 3-1. A schematic of the ADMX cavity and receiver chain. Switch A allows for the selection of sending power to either the major or minor ports, or to not send power to the cavity at all. Switch B allows for power to be read out by a spectrum analyzer in lieu of being mixed down and saved to disk. When acquiring data, the switches are set as pictured.

the magnetic field yields some control over the conversion process to the experimenter. The conversion rate of axions to photons can be increased merely by ramping up the magnetic field strength. The output of this interaction is now a single photon containing all of the energy previously held by the axion. That is to say that for an axion with mass m_a and velocity v relative to the Earth, the photon energy would be

$$E = hf = m_a c^2 \left(1 + \frac{1}{2} \left(\frac{v}{c} \right)^2 \right), \quad (3-2)$$

where c is the speed of light, h is Planck's constant, and f is the frequency of the decay photon. This process is enhanced when the frequency of the decay photon matches the resonant frequency of the cavity. When on resonance, axions can interact with the stored energy in the cavity, yielding a higher conversion rate. The scale of the enhancement is governed by the quality factor of the cavity, Q . The Q of any resonator is defined as

$$Q \equiv \frac{\text{Power stored in the resonator}}{\text{Power lost per cycle}}. \quad (3-3)$$

It is this resonance which provides the largest contribution to the increase in the conversion of axions into photons. Chapter 4 will discuss in greater detail the importance of Q .

3.3 Cryogenic Hardware

While the magnetic field strength and Q of the haloscope cavity weigh heavily on the conversion rate, less so is the impact of the cavity geometry. Any shape cavity is allowed, though some are easier to handle experimentally than others. When one considers the field produced by a commercially available superconducting solenoid, a cylindrical geometry stands out as being particularly convenient. As such, ADMX uses a right circular cylinder as the haloscope cavity. A diagram of the cryogenic volume used by ADMX, showing both the main magnet and the central insert, can be seen in Fig. 3-3.

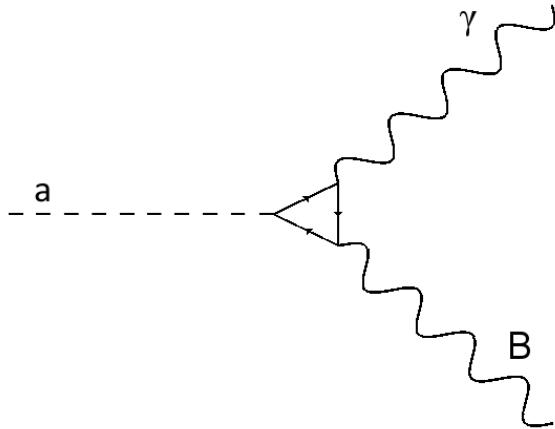


Figure 3-2. A Feynman diagram of an axion decay within a haloscope. In this decay, the axion interacts with the applied B field and the decay photon at the loop-level.

The ADMX cavity has a volume of 140 liters, is roughly 1 meter tall by 0.4 meters in diameter, and is kept under vacuum at pumped ${}^4\text{He}$ temperatures ($\sim 2\text{ K}$) during operation. Five main parts make up the cavity; the cylindrical wall, two removable end caps, and two tuning rods. The wall and end caps are copper plated stainless steel, while the tuning rods can be either a dielectric or copper plated stainless steel. Use of dielectric rods would lower the resonant frequency from that of the empty cavity, while metal rods would instead raise it. Metal rods were used when obtaining the data for this thesis, and the frequencies covered reached the top end of the tuning range (892.8 MHz) for the available rods. The end caps are affixed to the cylindrical wall by being firmly bolted in place. A knife edge connection ensures a good electrical connection at this joint which is necessary to support the surface currents generated by the fields in the cavity. Oxygen free copper was used for the plating of all parts. A picture of the cavity including tuning rods can be seen in Fig. 3-4.

For any right cylindrical cavity, such as the ADMX cavity, the theoretical value for Q is given by

$$Q = \frac{\mu_0}{\mu_c} \frac{V}{S\delta} \times G, \quad (3-4)$$

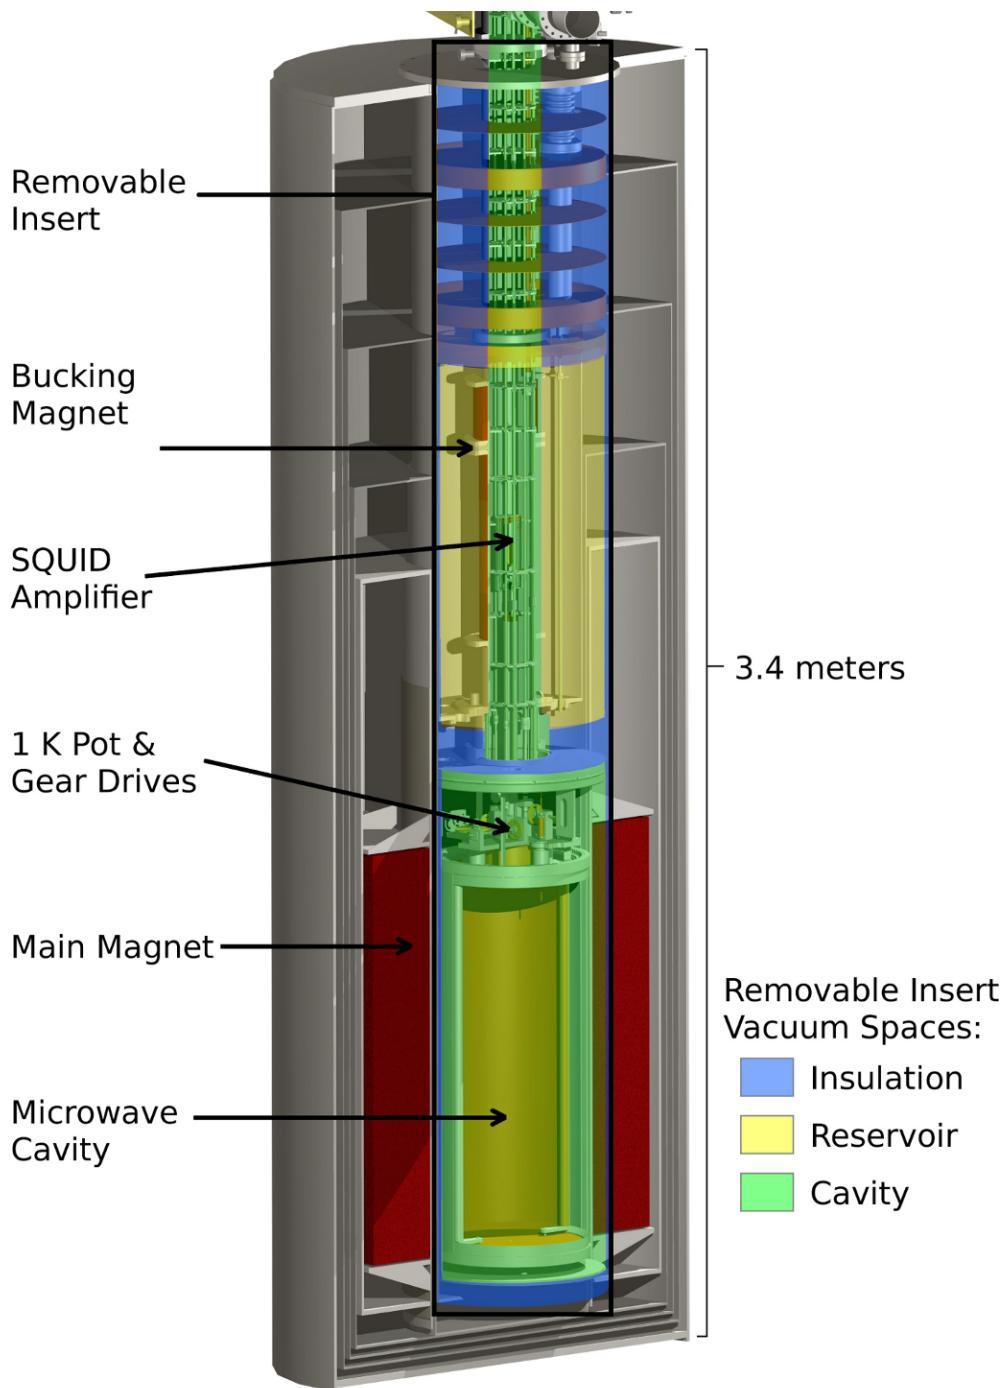


Figure 3-3. A SolidWorks drawing of the ADMX cryostat. The central portion is referred to as the insert. From the top down the insert consists of radiation shields, the helium reservoir which contains the bucking magnet, and the ADMX cavity. RF cables and sensor wires run along the central axis and are read out at the top plate.

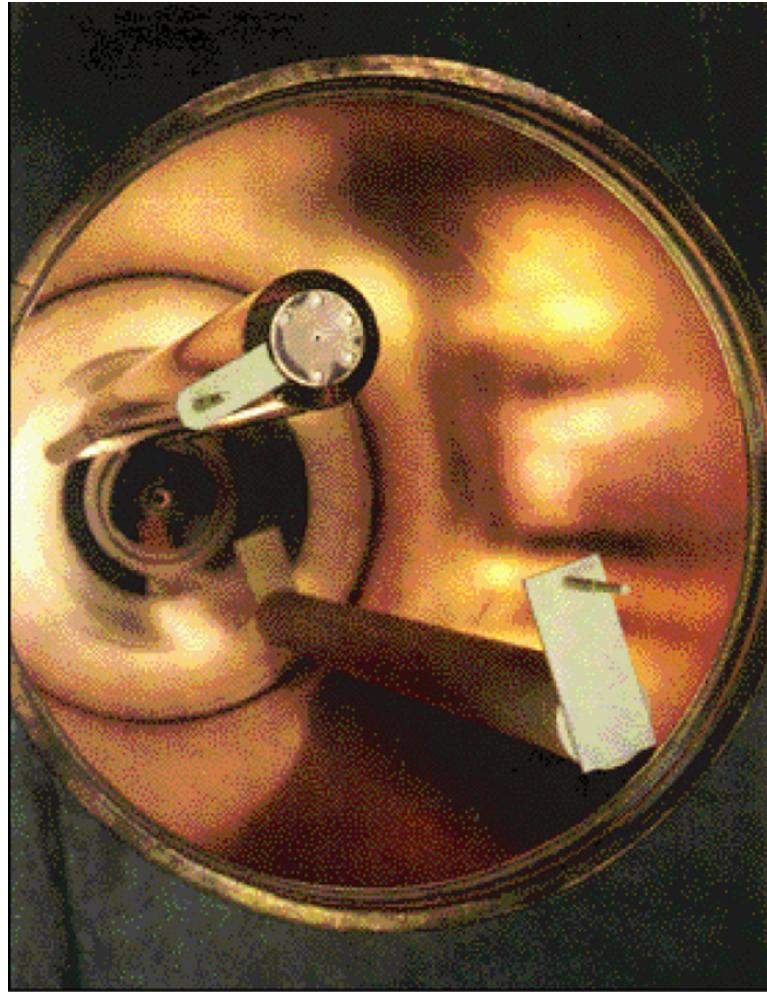


Figure 3-4. A picture of the ADMX cavity, opened, with tuning rods.

where μ_0 is the permeability of free space, μ_c is the permeability of the surface material, V is the volume of the cavity, S is the total area of the bounding surfaces, δ is the skin depth of the surface material, and G is a geometrical factor of order 1 which is dependant on the cross-sectional shape of the cylinder [68]. In practice, Q is measured just before the recording of each data stream. The skin depth of a material, that being the distance into which an oscillating electromagnetic field will penetrate before its amplitude is decreased by a factor of e^{-1} , strongly determines the power loss in Eq. 3-4 due to the Ohmic losses caused by this field penetration. Classically, δ can be

written in terms of the material and field properties as

$$\delta_{cl} = \sqrt{\frac{2}{\mu\sigma\omega}}, \quad (3-5)$$

where σ is the conductivity of the material, and ω is the angular frequency of the field oscillations. When the classical skin depth becomes significantly smaller than the mean free path of the electrons in the walls, Eq. 3-5 is no longer applicable. Instead the anomalous skin depth, given by

$$\delta_a = \left(\frac{\sqrt{3}c^2 m_e v_F}{8\pi\omega n e^2} \right)^{\frac{1}{3}}, \quad (3-6)$$

is used [69]. Here c is the speed of light, m_e is the electron mass, v_F is the Fermi velocity in the conductive surface (1.57×10^8 cm/sec), n is the electron number density (8.47×10^{22} /cm³), and e is the electron charge [70]. For room temperature copper ($\mu_c \approx \mu_0 = 1.256 \times 10^{-6}$ N/A², and $\sigma = 5.96 \times 10^7 \Omega^{-1} m^{-1}$) and frequencies consistent with the search range of ADMX (O(GHz)), the skin depth of the ADMX cavity would be $\delta_{cl} = 2.06$ microns. Though when cooled to liquid helium temperatures, the ADMX cavity falls into the anomalous anomalous skin depth regime resulting in $\delta_a = 285$ nm. With δ_a being less than a micron, small surface imperfections in the copper plating now cause significant degradation to the Q [71]. Figure 3-5 shows imperfections in the copper plating which led to a less than optimal Q while obtaining the data for this thesis.

To be effective, the cavity must be tunable over a range of frequencies. This tunability is achieved by means of two movable tuning rods. The tuning rods are 5 cm in diameter and run along the full length of the cavity. Each is mounted on a silica arm on opposite sides of the cavity and each can freely rotate through 360°. When changing the resonant frequency of the cavity between measurements, the typical step size is ~ 2 kHz. Figure 3-6 shows a diagram of the rod motion. The rods impose an additional boundary condition to Maxwell's equations by effectively excising part of the cavity. The fields no longer see this volume yielding a higher resonant frequency. When the rods

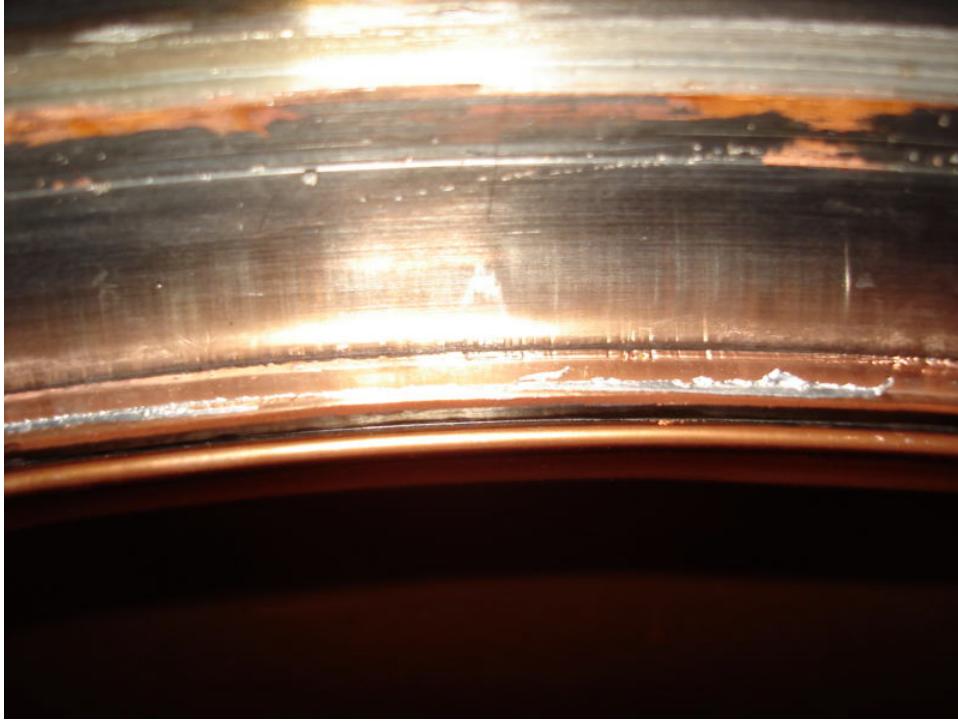


Figure 3-5. A picture of the copper plating flaking on top rim of the cavity wall. This created a poor electrical connection to the lid, which then degraded the Q .

are near the walls, where the field is already very small, the effect of their presence is minimal. Moving the rods closer to the cavity center, where the field strength is significantly higher, produces a much larger change to the resonant frequency.

As can be seen in Eq. 3-1, the resonant mode with the largest overlap between its electric field and the external magnetic field will yield the highest probability of axion to photon conversion. ADMX uses an 8 tesla superconducting solenoid manufactured by Wang NMR to generate its external magnetic field. The cryostat housing the cavity fits closely in the 60 cm diameter bore of the main magnet cryostat. Because the cavity rests at the same height as the magnet, the field is uniform and runs axially through it. A measure of the overlap between the electric field of the resonant mode, \mathbf{E} , and the main magnetic field, $\mathbf{B} = B_0 \hat{z}$, is given by a unitless form factor

$$C = \frac{\left| \int_V \mathbf{E} \cdot \mathbf{B}_0 d^3x \right|^2}{VB_0^2 \int_V |\mathbf{E}|^2 d^3x} \quad (3-7)$$

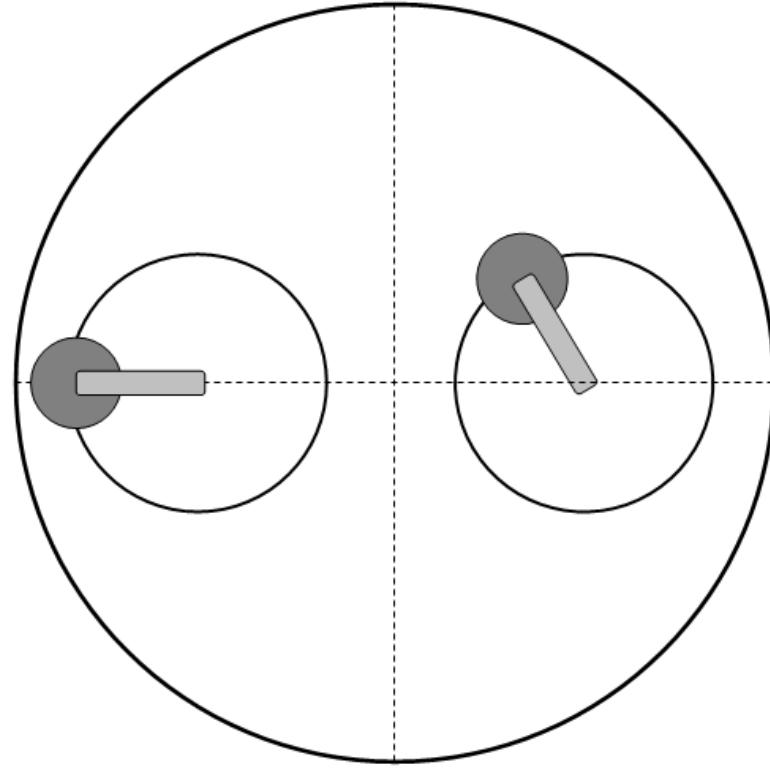


Figure 3-6. A diagram of the motion of tuning rods in the ADMX cavity from a top down prospective. Each rod (dark grey) is connected to an arm (light grey) which which can freely rotate 360°.

where V is the cavity volume and ϵ_0 is the permittivity of free space. The presence of tuning rods in the cavity makes \mathbf{E} non-analytical, forcing it to be solved numerically. The numerical solution is used in Eq. A-2 to determine C for a given rod configuration. This process is repeated for all rod configurations generated by changing the angle of the rod positions in 5° increments. A form factor map is then produced which can be interpolated for rod positions not at multiples of 5°. The TM_{010} mode has the largest net overlap of the axial component of the electric field and the external field, yielding a form factor ranging from 0.41 to 0.61. As such, ADMX tunes to the TM_{010} mode for all data taking.

The bucking magnet is a second magnet mounted in the helium reservoir in the cavity cryostat about 1 m above the cavity. This magnet, manufactured by American Magnetics Incorporated, was specially designed to nearly cancel a small region of the

main magnet's fringing field while minimizing interactions with the main magnet. This low-field region is mandatory for the effective operation of ADMX's first stage amplifier; a dc Superconducting QUantum Interference Device (SQUID). The SQUID is extremely sensitive to magnetic flux but is required to be near the cavity to avoid both signal degradation prior to amplification and large temperature differences between itself and the cavity.

It can be seen in Fig. 3-7 that the bucking magnet is actually arranged into two distinct coils. The bottom coil is the cancelation coil. It is in the center of this coil, about 1.6 m above the cavity, that the low field ($< 25 \mu\text{T}$) region sits. Figure 3-8 shows a field profile for this region. Above the cancelation coil is a secondary coil which is counter-wound to the cancelation coil. The two coils are affixed to each other but have opposing dipole moments. This results in very nearly no net mutual inductance and no net repulsive forces between the bucking magnet and the main magnet. Even with field cancelation from the bucking magnet, additional μ -metal shielding around the SQUID is necessary. Further still, The SQUID is placed inside a superconducting lead-lined brass box which diverts any flux which has managed to penetrate the μ -metal shields.

3.4 Receiver Electronics

RF power is picked up from the cavity through a critically coupled antenna. It is sent through a directional coupler and a signal isolator followed by a chain of amplifiers and image reject mixers before being digitized and saved to disk. The directional coupler itself is not a critical component during actual data taking. Instead it is used periodically to ensure that the antenna remains critically coupled to the cavity. Having the antenna and the cavity critically coupled results in half of the power resonating in the cavity being picked up by the antenna and half being dissipated in the walls of the cavity. The signal isolator which follows is a three port circulator connecting the directional coupler to the first stage amplifier with the third port being terminated by a 50Ω resistor. Any signal



Figure 3-7. A picture of the bucking magnet prior to installation.

reflection caused by an impedance mismatch between the transmission line and the first stage amplifier is instead sent to the terminated port and is attenuated by 30 dB prior returning to the cavity.

The switches shown in Fig. 3-1 are flipped to allow for adjustments to the antenna coupling and to measure both the Q of the cavity and the average power being transmitted along the receiver chain. Setting switch A to the right and switch B to the left, allows power to be sent down to the cavity along the middle transmission line and have its reflection be read out by a spectrum analyzer. When the antenna is critically coupled, nearly all of the power that is on resonance is dumped into the cavity. When the antenna is not critically coupled, much of this power is reflected at the cavity-antenna interface. While examining the reflected signal on the spectrum analyzer, the insertion depth of

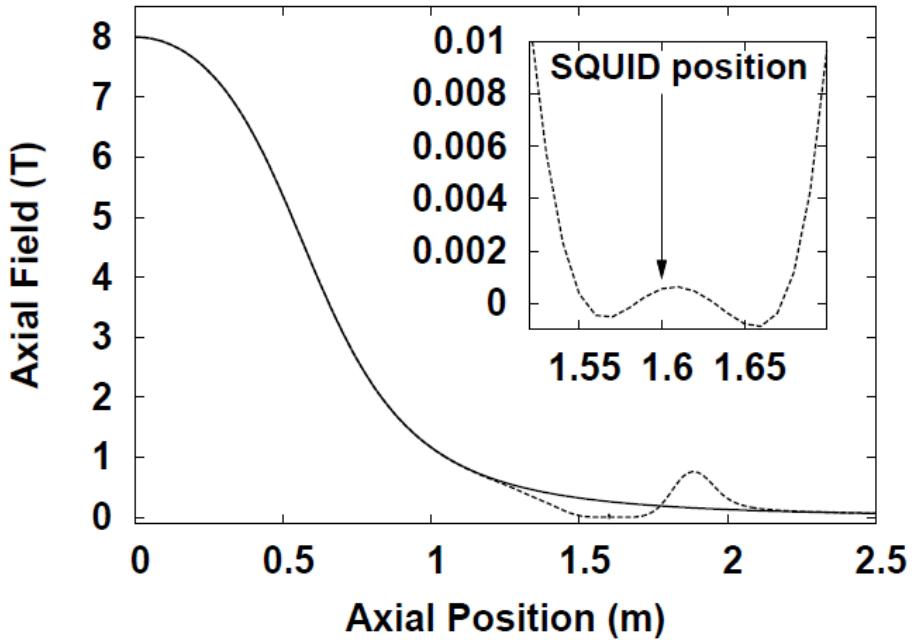


Figure 3-8. The on-axis, magnetic field profile along the length of the bucking magnet.

the antenna can be varied until the proper coupling is found. When the reflection is at its minimum, typically -30 dB, the antenna and cavity are considered critically coupled. To measure the Q of the cavity and the average transmitted power, both switches are set to the left. Power is then sent down to the cavity via the left-most transmission line. This power is swept over a range of frequencies. The power enters the cavity through a weakly coupled antenna and is read out as before. The shape of the signal as read out at the spectrum analyzer is given by the Lorentzian

$$P(f) = \frac{P_0}{1 + 4Q^2 \left(1 - \frac{f}{f_0}\right)^2}, \quad (3-8)$$

where P_0 is the maximum transmitted power, and f_0 is the resonant frequency of the cavity. The spectrum analyzer then fits this shape to determine Q . Lastly, the spectrum analyzer directly measures the average power transmitted along the receiver chain. This value is indicative of the total amplification provided by the receiver chain, and is saved as FFT_{avg} in the header of each data file. During normal data taking operations, this

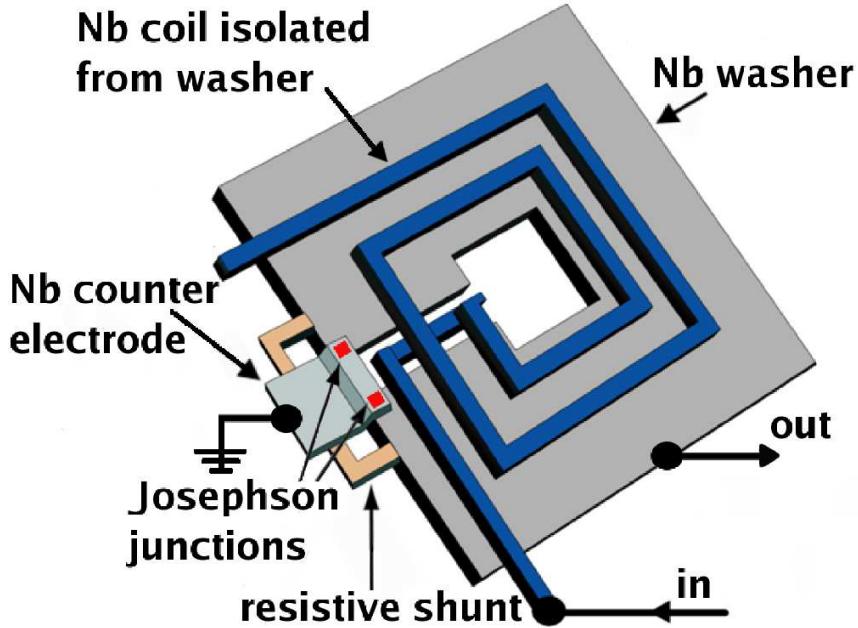


Figure 3-9. A diagram of the MSA as used in ADMX. The SQUID, responsible for the actual amplification, consists of the Nb washer and Josephson junctions, while the microstrip, responsible for coupling power to the squid, is the Nb shown in blue.

process of determining Q and $FFT_a vg$ occurs between data measurements and is the only time when power sent directly to the cavity.

The first of the amplifiers in the receiver chain is a low-noise microstrip SQUID amplifier (MSA), diagramed in Fig. 3-9 [72]. The base of the MSA is a dc SQUID consisting of a 1 mm by 1 mm square Nb washer, with a 0.2 mm by 0.2 mm hole in the middle and a gap on one side where the Josephson junctions are located. Power from the cavity is coupled to the SQUID through a Nb microstrip coil deposited over an insulation layer on top of the washer. One end of the coil is left free and the SQUID is used as its groundplane, turning the coil/SQUID system into a resonator. This yields a range of operational frequencies with the resonant frequency and the bandwidth of the MSA set by the length of the microstrip and the Q of the microstrip, respectively. Having a fairly large operational range permits data taking to continue for longer periods without needing to change out hardware. Amplification from the MSA is a result of flux biasing

the SQUID [73]. A current exceeding the critical current of the Josephson junctions is applied to the SQUID making the junctions resistive, resulting in a voltage, V , across the SQUID. Any magnetic flux, Φ , which threads the SQUID will cause V to vary sinusoidally with

$$V_{min} = V(0). \quad (3-9)$$

The period of this variation is given by

$$\Phi_0 = \frac{h}{2e} = 2.1 \times 10^{-15} \text{Wb}, \quad (3-10)$$

where h is Plank's constant and e is the electron charge. When $|\frac{dV}{d\Phi}|$ is large, any small change in Φ caused by the current in the microstrip will generate a large change in the voltage across the SQUID. The maximum value of $|\frac{dV}{d\Phi}|$ occurs $\frac{\Phi_0}{4}$ before or after an extremum of $V(\Phi)$. AMDX applies a constant current to the SQUID to maintain this biasing, maximizing the gain from the MSA ($\sim 10\text{-}12$ dB). The noise temperature of the MSA has been tested and found to be roughly half of the physical temperature, that is $T_{N,MSA} \approx 1$ K for typical operational temperatures. Figure 3-10 shows this trend continuing down to the quantum limit [74] of

$$T_{N,limit} = \frac{h\nu}{k_B} = 48 \left(\frac{\nu}{\text{GHz}} \right) \text{ mK}, \quad (3-11)$$

where k_B is Boltzmann's constant. Upgrades currently underway will lower the physical temperature to $\mathcal{O}(100 \text{ mK})$ in order to bring T_N closer to the quantum limit.

Most of the data used in this thesis utilized an MSA as the first stage amplifier, though near the end of data taking the SQUID was damaged beyond repair by a magnet quench. A replacement MSA would not have been available until after decommissioning was scheduled to occur in preparation for the relocation of the experiment from Lawrence Livermore National Laboratory to the Center for Experimental Nuclear Physics and Astrophysics at the University of Washington. In order to close out coverage of the tunable range of the cavity, data taking continued without this stage of amplification.

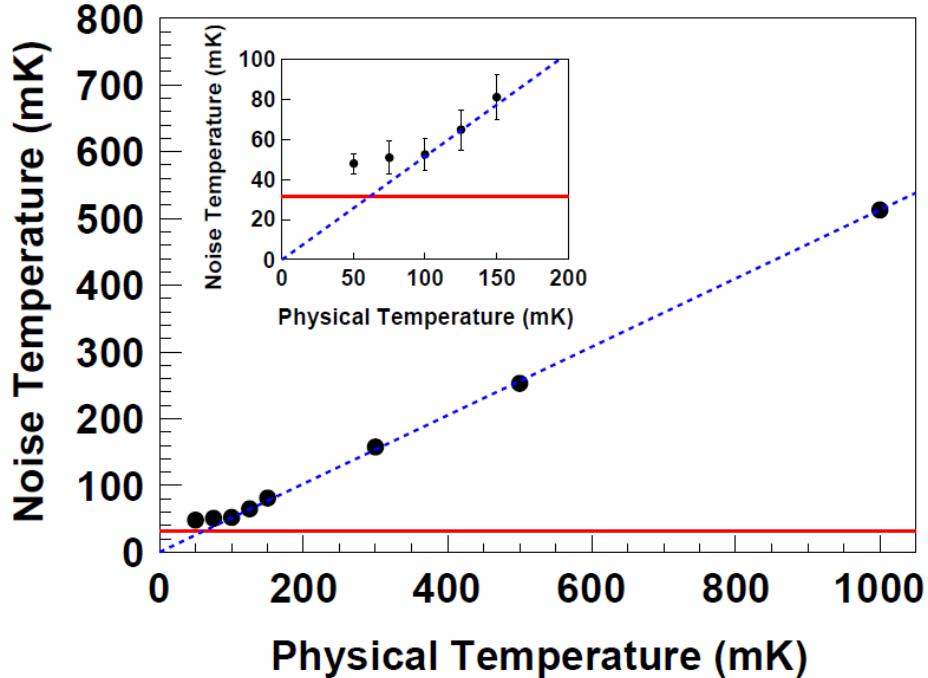


Figure 3-10. The noise temperature profile of the MSA. Once above the quantum limit, the noise temperature is linearly dependant on the physical temperature.

All data taken from February 2010 through the end of operations used a low-noise transistor amplifier as the first stage. The system noise temperature for this set of data is a little less than twice what it was previously, but this did not significantly affect the the limits produced by the analysis presented in Chapter 4.

Both the second and third stage amplifiers are balanced GaAs heterostructure field effect transistors (HFETs), diagramed in Fig. 3-11, designed and produced by the National Radio Astronomy Observatory (NRAO) [74, 75]. Each of these HFETs sends the signal through a hybrid coupler, followed by a pair of matched single-sided amplifiers, and finally through a second hybrid coupler. The first hybrid coupler splits the signal adding a 90° phase change to one of the legs of the split transmission line while the second hybrid coupler adds a 90° phase change to the other leg before recombining the signal. This design is referred to as being balanced and results in reflections off the single-sided amplifiers being recombined within the hybrid couplers having

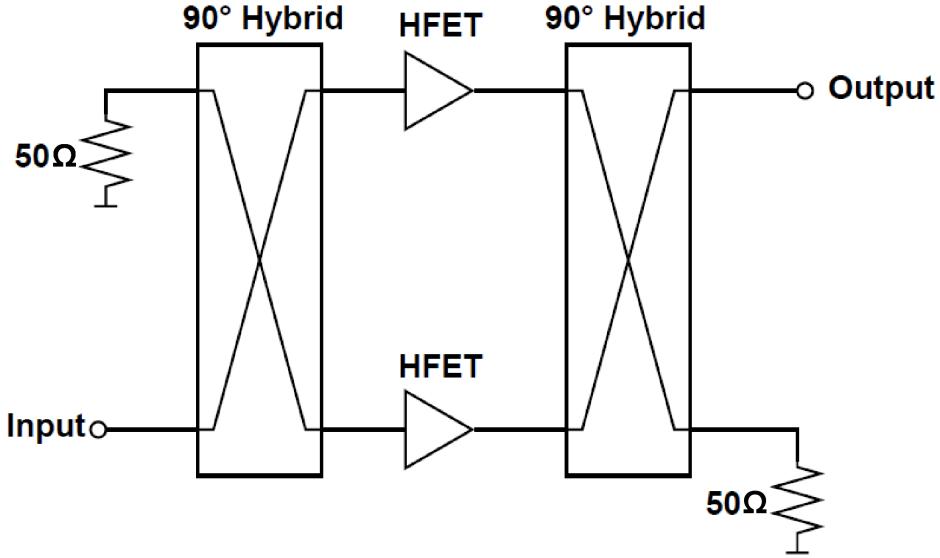


Figure 3-11. A diagram of the balanced design for the HFET amplifiers.

a 180° phase difference. Any remaining power from reflections is diverted to a port terminated with a 50Ω resistor, where this power is dissipated. The balanced design is required for the HFETs to match the 50Ω impedance of the RF transmission line. While an external form reflection isolation could be used at the HFET input, such as a second circulator like the one that precedes the SQUID, balancing the HFETs proves to be a more practical option. The main problem with external reflection isolation is that circulators used by ADMX do not work in strong magnetic fields as they contain permanent magnets and the limited amount of low-field volume provided by the bucking magnet is already claimed by the SQUID and its circulator.

Each HFET provides ~ 17 dB of gain and has a noise temperature of 4 K [73]. However, these noise temperatures are suppressed by the gain of all preceding amplifiers in the chain. The total system noise temperature is given by

$$T_{N,System} = T_P + T_{N,1} + \frac{T_{N,2}}{G_1} + \frac{T_{N,3}}{G_1 G_2} + \frac{T_{N,4}}{G_1 G_2 G_3} + \dots, \quad (3-12)$$

where T_P is the physical temperature of the cavity, $T_{N,i}$ is the noise temperature of the i^{th} amplifier, and G_i is the arithmetic gain of the i^{th} amplifier [76]. Conservatively, the

SQUID provided 10 dB of gain. This drops HFET-1's contribution to the system noise temperature from 4 K down to 0.4 K. HFET-2 follows a total of 27 dB of gain, or a factor of 500 arithmetically, reducing its contribution to 0.008 K. Any subsequent amplifiers will have negligible contributions to the system noise since they will be suppressed by at least 44 dB of gain, or a factor of $\sim 25,000$. This yields a system noise temperature of

$$T_{N, System} = T_P + T_{N, SQUID} + \frac{T_{N, HFET-1}}{G_{SQUID}} = 2K + 1K + \frac{4K}{10} = 3.4K. \quad (3-13)$$

When ADMX was forced to run without a squid, the HFETs became the first and second stage amplifiers. With no suppression of the noise temperature for HFET-1 and only 17 dB of gain to suppress HFET-2, the system noise temperature for this data is

$$T_{N, System} = T_P + T_{N, HFET-1} + \frac{T_{N, HFET-2}}{G_{HFET1}} = 2K + 4K + \frac{4K}{50} \approx 6K. \quad (3-14)$$

Amplifiers following HFET-2 continue to contribute only a negligible amount to the system noise temperature.

At room temperature the signal is amplified 60 dB by two commercial microwave-amplifiers, before being mixed down by an image-reject mixer (MITEQ IRM045-070-10.7) from GHz frequencies to an intermediate frequency (IF) of 10.7 MHz. The mixer is balanced in the same sense as the HFET amplifiers and contributes 7dB of loss to the signal. The IF signal is amplified 30 dB, then sent through an eight-pole crystal band-pass filter (3 dB of loss) with a bandwidth of 30 kHz, and further amplified another 30 dB. The IF signal is then mixed down, contributing another 7 dB of loss, to an audio frequency signal centered at 35 kHz whereupon it enters 15.5 bit Analog-to-Digital Converter (ADC) for the HR channel. The ADC reads the the signal at 80,000 samples/second for 23.8355 seconds, yielding a time series of voltage measurements that is 1,906,840 points long. The spectral resolution of this data stream is given as

$$b = \frac{1}{\text{Integration Time}} = 42\text{mHz}. \quad (3-15)$$

Each data stream is then time stamped and saved to disk along with data describing the state of the experiment at that time. The total amplification along the path from the cavity to the ADC, including a 6 dB loss in the transmission line cables, is 141 dB, or 131 dB when running without the MSA.

Despite having yoctowatt sensitivity, the receiver chain hardware is not perfect. The frequency standard (SRS FS700) used to set the mixing frequency has a short term stability of 10^{-10} . For an 850 MHz signal, a typical frequency for this set of HR data, a potential frequency drift of 85 mHz can be expected in the mixers over the course of one measurement. This is larger than the 42 mHz resolution of the HR power spectrum. Consequently, the stability of the SRS FS700 is what sets the finest resolution examined in this work. At a resolution of 84 mHz, a signal will have drifted by no more than one bin during one measurement. Note that this hardware decision was made long before the implementation of HR channel. At the time, the spectral resolution of ADMX data was $\mathcal{O}(100 \text{ Hz})$ so the $\mathcal{O}(10 \text{ mHz})$ drift for a typical signal was negligible.

CHAPTER 4

DATA ANALYSIS

4.1 Introduction

The HR channel produced roughly 2 TB of time series data during ADMX’s most recent science run. Potentially buried within this data set is an axion signal just waiting to be found. Discovering such a signal relies not only on knowing what to search for, but also on knowing what is obscuring it. Properly accounting for system noise, bad data, and systematic errors prior to searching is critical. Presented in this chapter is a study of the afore mentioned concepts, the means by which the search is performed, and the results generated from that search.

Considered first are axion signal properties, which serve to guide the direction of the analysis. These properties are the expected signal strength, spectral broadening, and signal modulation, with the latter being of particular importance as it directly affects the search algorithm. Second is the handling of the HR data at different resolutions. This includes the non-Gaussian noise statistics involved therein. Third would be a look at the HR data set itself and the quality control cuts applied to it. Finally, the processing required of the HR data prior to searching, the search algorithm, and the exclusion limits derived from the data set are examined.

4.2 Axion Conversion and Signal Properties

As mentioned in Chapter 3, a single axion conversion within the ADMX cavity will produce a photon with energy given by Eq. 3–2. However, ADMX doesn’t look for single conversion events. When the resonant frequency of the cavity matches the frequency of the decay photons, the rate of axion conversion increases drastically. Thus, the rate of energy deposition into the cavity also increases drastically. This produces a faint but noticeable change in power within the cavity at the resonant frequency. This power is given by

$$P_a = g_{a\gamma\gamma}^2 \frac{VB_0^2 \rho_a C}{m_a} \min(Q, Q_a), \quad (4-1)$$

where V is the cavity volume, B_0 is the magnetic field strength, ρ_a is the local axion density, and Q_a ($\sim 10^6$) is the quality factor of the axion energy distribution[67]. By extracting typical values for the experimental parameters used in ADMX, Eq. 4–1 can be rewritten in the more convenient form [77]

$$P_a \approx 5 \times 10^{-22} W \left(\frac{V}{500L} \right) \left(\frac{B_0}{7T} \right)^2 C_{010} \left(\frac{g_\gamma}{0.36} \right)^2 \times \left(\frac{\rho_a}{5 \times 10^{-25} g/cm^3} \right) \left(\frac{m_a}{5 \mu eV} \right) \left(\frac{\min(Q, Q_a)}{10^5} \right). \quad (4-2)$$

Because the radio receiver and the cavity are critically coupled, only half of P_a is picked up and sent through the receiver chain. This expected axion power is still roughly two orders of magnitude larger than the rms thermal noise power in the HR channel. For a spectral resolution of $b = 84$ mHz and a system noise temperature of $T_N = 3$ K, the rms noise power would be

$$P_N = T_N k_b b = 3.5 \times 10^{-24} W, \quad (4-3)$$

where k_b is Boltzmann's constant.

The term in P_a with the largest variability, and thus the largest potential to degrade an actual axion signal is Q . During the 2009 and 2010 operations, the Q of the cavity varied over a substantial range; from roughly 10,000 to roughly 80,000. Prior to 2010, the typical value for Q was between 10,000 and 20,000. Since the P_a is directly proportional to Q , the low values for Q bring the signal power down for this subset of data, which leads to slightly weaker limits for the frequencies covered prior to 2010. This problem was due to poor electrical contact between the cavity wall and the end caps caused by a degradation of the copper plating at the knife edge connection. The wear in the copper plating was discovered during a break in data taking at the end of 2009. Upon reassembling the cavity, indium wire was laid along the knife edge joint. Bolting the end caps onto the cavity wall compressed the indium which bridged the

electrical gap and restored the Q to a more reasonable value. 2010 data had a typical value for Q between 50,000 and 80,000.

In addition to the Q enhancing an axion signal which is on resonance, it also greatly effects those which are off resonance. When the signal is off resonance P_a is reduced by a Lorentzian factor, L , given by Eq. 3–8 and how far off resonance the signal is. For a signal at frequency f when the cavity is tuned to a resonant frequency f_0 , this reduction factor is given by,

$$L(f) = \frac{1}{1 + 4Q^2 \left(1 - \frac{f}{f_0}\right)^2}. \quad (4-4)$$

A signal near the edge of the usable bandwidth of the crystal filter, say 12 kHz off resonance, in a cavity with a Q of 70,000 and a resonant frequency of 850 MHz would be reduced by roughly a factor of 5 compared to a signal on resonance. It can be seen from Eq. 4–4 that spoiling the Q flattens $L(f)$ and effectively increases the usable bandwidth of the cavity. However, there is an overall loss of signal power as P_a is proportional to Q , and the extra cavity bandwidth is insufficient to compensate for this loss.

Other analyses have compensated for a low Q by averaging repeated measurements taken at the same frequency. Any real signal would consistently be present and would add coherently, while noise peaks would be stochastic and would add randomly. Given a steady signal and sufficient integration time, t , the Signal to Noise Ratio (SNR), given by the Dickie Radiometer equation [78]

$$SNR = \frac{P_a}{P_N} \sqrt{bt} = \frac{P_a}{k_b T_N} \sqrt{\frac{t}{b}}, \quad (4-5)$$

would eventually reach a statistically significant value. Unfortunately, at the resolutions attainable by the HR channel, significant assumptions regarding the velocity vectors of halo axions must be made in order to make use of this technique. As such, the HR power spectra are not capable of being averaged before they are searched. This puts a

premium on high Q measurements since low values of Q are potentially detrimental and could seriously weaken HR exclusion limits.

Before addressing any conditional statements regarding axion motion, it is important to know how a generic axion signal looks and behaves. As seen in Eq. 3–2, the frequency of the axion signal is determined by the axion's rest mass and its velocity relative to Earth. This is independent of the actual nature of the motion of the axion. Now consider a change to the axion velocity such that it differs from v by some small but otherwise arbitrary amount ϵ . Energy deposited in the cavity would then occur at a frequency

$$f' = \frac{m_a c^2}{h} \left(1 + \frac{v\epsilon}{c^2} + \mathcal{O}(\epsilon^2) \right) \approx f + \frac{m_a v \epsilon}{h}. \quad (4-6)$$

Rearranging 4–6 and dividing by f yields an expression for the fractional change in frequency,

$$\frac{f' - f}{f} = \frac{m_a v \epsilon}{m_a c^2 (1 + \frac{1}{2} v^2)} \approx \frac{v \epsilon}{c^2}. \quad (4-7)$$

Also useful is the absolute change in frequency,

$$\delta f = f' - f = \frac{f v \epsilon}{c^2}. \quad (4-8)$$

Eq. 4–8 provides a convenient platform for examining two very important properties of an axion signal; spectral line broadening and frequency modulation.

Descriptions of axion signals thus far have only referred to single frequencies. While true that a single axion will produce only a single photon, recall that many axions will be decaying simultaneously. Not every axion will have the same velocity, so power will necessarily be deposited at different frequencies. Let ϵ be the difference in velocity between any two groups of axions, where all axions within a given group have the same velocity. Clearly, one group would contribute to the power at f while the other would contribute at f' . In both cases the power deposited in the cavity is proportional to the number of axions in the appropriate group. A third group would also contribute to its own frequency based on its own population; and a fourth; and fifth; and so on. Ultimately, a

shape in the power spectrum of the cavity will be traced out that is identical to the shape of the axion velocity distribution.

This broadening of the spectral line shape is expected for any velocity distribution, where the mode of the distribution determines the center frequency of the axion signal. Since there is no a priori importance for one of the above groups over another, let the power at f come from a group having velocities equal to the mean of the velocity distribution and the power at f' come from a group having velocities which are one standard deviation away from the mean. In this case, ϵ represents the dispersion, δv , of the velocity distribution and Eq 4-8 yields the scale of the spectral broadening of the axion signal. The full width of the signal in the power spectrum is actually twice this value as it is determined by velocities varying both above and below the mean.

The second property an axion signal can exhibit is a displacement of the center frequency. Let ϵ now be a change in v which occurs over some length of time, say t seconds. The center frequency of a measurement made at one time would then be different from the center frequency for a measurement made t seconds later. Instead of giving the spectral broadening of the signal, Eq. 4-8 now gives the magnitude of the drift in the center frequency. Just as important as the magnitude is the rate at which this drift occurs, which is easily obtained by dividing Eq. 4-8 by t . But what if this variation in v were periodic? Then so too would be the drift in the center frequency. This frequency modulation would be governed by the oscillation in v perhaps written such that

$$\epsilon = v - \langle v \rangle = A \sin \left(\frac{2\pi t}{T} + \phi \right), \quad (4-9)$$

where $\langle v \rangle$ is the time averaged axion velocity, A is the amplitude of the oscillation, T is the period, and ϕ is an arbitrary phase. The maximum rate of change of ϵ is then given by

$$\dot{\epsilon}_{max} = \frac{2\pi A}{T}, \quad (4-10)$$

with a maximum short term modulation ($t \ll T$) of

$$\delta f_{short, max} = \frac{2\pi A t}{T}. \quad (4-11)$$

Prior knowledge of the phase of ϵ would permit a more accurate determination of instantaneous modulation rate and, by extension, the short term modulation, but for now no assumptions about ϕ will be made.

While nothing prevents the oscillations in v from being an inherent property of the motion of the axions, consider that the rest frame of the experiment is the Earth, which exhibits both rotational and orbital motion. Both modes have a period which is less than the lifetime of ADMX, making them noticeable and indicating the motion of the frame of reference a significantly more likely culprit. This, of course, would only affect axions whose velocities have some component which is in line with the motion in question. For axions whose velocity is perpendicular to one of these modes of terrestrial motion, that mode would not contribute to the signal modulation. However, there is a $\sim 23^\circ$ tilt between the orbital and rotational axes of the Earth. This prohibits there ever being an axion velocity which is always simultaneously perpendicular to both. Consequently, signal modulation of some sort is not only possible, it is expected. Once again, only by making assumptions about axion velocities can this expected modulation be corrected for in the HR spectra prior to searching for peaks. Note that the Earth's motion about the Earth-Moon barycenter is neglected as the amplitude of velocity oscillations is trivial. As well, the motion of the solar system through the galaxy is neglected as its period is too large to be noticed by ADMX in any realistic amount of time [16].

The degree to which these properties manifest themselves is highly dependant on the velocity distribution of the halo axions. As discussed in Chapter 2, the HR channel was created for the purpose of searching for a signal consistent with an organized low-velocity-dispersion flows of axions. These flows may be isolated local features of the halo substructure; or they may run the length of the galaxy, turning back on themselves

several times over, each time creating a region of higher than average density. In either cases, each flow would be characterized by a specific flow direction, its mean velocity relative to the Earth, its velocity dispersion, and its axion density. This work seeks to make as few assumptions as possible beyond those required by the halo models of interest, which are as follows:

- The Universe evolved following the Λ CDM cosmological model.
 - Dark matter is non-relativistic.
 - Dark matter is effectively collisionless.
 - Dark matter is non-baryonic.
- The KSVZ and DFSZ axion models are used.
 - $g_{KSVZ} = -0.97$
 - $g_{DFSZ} = 0.36$
- Late infall axions and axions caught in tidal disruptions will form flows.
- These flows will have lower velocity dispersions than virialized axions.
- The turn around points of late infall flows will cause caustics to form.

Beyond these required assumptions the following liberties are also taken:

- There is no preferred direction relative to the earth for axion flows.
- All axions considered move faster than the escape velocity of the solar system
 $v \gtrsim 42 \text{ km/s}$
- All axions considered move slower than the escape velocity of the Milky Way (i.e.
 $v \lesssim 600 \text{ km/s}$).
- Virialized axions have a Maxwellian velocity distribution with a velocity dispersion of 300 km/s.

With this in mind, the above descriptions of signal properties can now be examined further. Reasonable dispersion and mean velocity estimations can be used to give a more accurate picture of how axion flows might both look and behave when seen in the HR channel. Any calculations made to illustrate these effects will use 850 MHz and

300 km/s as a typical values of f and v respectively. Estimations will err on the side of caution assuming a worst-case scenario, though whenever possible accurate numbers will be used.

First consider again the spectral broadening of an axion signal. This phenomenon is central to the existence of the HR channel. Flows with low velocity dispersions would consequently have very little spectral broadening. This leads to significantly lower SNR when the signal width is much smaller than the spectral resolution of the power spectrum since P_N grows with increasing bin width but P_a does not. A finer resolution would cut out unwanted thermal noise power yielding a higher SNR. But how fine should one make the resolution? Were the resolution finer than the signal sought after, axion power would be spread across multiple bins in the power spectrum. In a search which averages successive power spectra, this feature would eventually become prominent enough to be clearly seen over the noise. With no averaging, as in the analysis presented here, such a signal would remain hidden in the noise. It is important to choose the resolution for the power spectrum that is neither too large nor too small. To ensure that an axion signal is not overlooked due to selecting the wrong resolution, a spread of resolutions will be used for this analysis.

As mentioned in Chapter 3, a hardware limitation constrains the finest resolution presented here to 84 mHz. Rearranging Eq. 4–8 and including the factor of 2 from the full width of the spectral line, the corresponding velocity dispersion is

$$\delta v = \frac{c^2(84\text{mHz})}{2(850\text{MHz})(300\text{km/s})} = 15\text{m/s.} \quad (4-12)$$

This dispersion is actually smaller than what is expected for the Big Flow, 50 m/s, though other flows could still be seen at this resolution. The coarsest resolution for this analysis was chosen to be 1.09 Hz, which corresponds to a dispersion of

$$\delta v = \frac{c^2(1.09\text{Hz})}{2(850\text{MHz})(300\text{km/s})} = 192\text{m/s.} \quad (4-13)$$

Coarser resolutions are not insignificant, but they are more efficiently handled elsewhere [14, 17] and are considered beyond the purview of this analysis. Resolutions of 168 mHz and 546 mHz are also examined filling the gap between 84 mHz and 1 Hz.

For sub-Hertz resolutions, the non-stationary nature of the detector's rest frame becomes quite significant. The Earth's rotational and orbital motion would respectively cause diurnal and annual modulations of the axion signal. It is the magnitude of these modulations which is of concern. The scale of diurnal modulation is primarily set by the rotational angular velocity of the Earth and to a lesser extent by the location of the detector on the Earth's surface. The latter component actually turns out to be rather favorable from an experimental standpoint. The Earth rotates 2π radians over one sidereal day (23.93 hours) yielding an angular velocity of

$$\omega_r = \frac{2\pi}{(23.93\text{hours})(3600\text{sec/hour})} = 7.29 \times 10^{-5}\text{rad/s}. \quad (4-14)$$

Multiplying this by the mean radius of the Earth gives

$$\delta v_r = \omega_r (6.37 \times 10^3\text{km}) = 465\text{m/s} \quad (4-15)$$

as the amplitude of velocity oscillations. However, the latitude of the location of the detector reduces this amplitude making

$$\delta v_r \Rightarrow = \delta v_r \cos(\ell). \quad (4-16)$$

ADMX operated in Livermore, CA ($\ell = 37.68^\circ$) resulting in a velocity amplitude due to rotational motion of $\delta v_r = 368$ m/s for an axion flow moving parallel to the Equator (worst-case motion). Substituting this into 4-8 gives a modulation amplitude of

$$\delta f = \frac{(850\text{MHz})(300\text{km/s})(368\text{m/s})}{c^2} = 1.04\text{Hz}, \quad (4-17)$$

immediately indicating the problem. With power spectra having order mHz resolutions, an axion signal will not be measured at the same frequency from scan to scan. Over

12 hours (one half period) an axion signal could drift by as much as twice this value. Further, not knowing the phase of the velocity oscillations means that this drift could be in either direction, giving a ± 2 Hz window for the signal to move in. This clearly demonstrates why averaging between spectra is prohibited as, more often than not, any axion signal would simply be averaged with incoherent noise.

Annual modulation follows the same logic using instead the Earth's orbital velocity and period, $\delta v_o = 29.8$ km/s and 1 year respectively. While the modulation amplitude due to orbital motion,

$$\delta f = \frac{(850\text{MHz})(300\text{km/s})(29.8\text{km/s})}{c^2} = 84.5\text{Hz}, \quad (4-18)$$

is significantly larger than the amplitude due to rotations, so too is the period. A maximum drift of ± 169 Hz would occur slowly over the course of six months for axions traveling parallel to the ecliptic (worst-case modulation). It is convenient to scale down the time interval to 12 hours to match the time interval examined for rotational modulation. Because this makes $t \ll T$, 4-11 can be used to obtain the the maximum short term drift of the signal,

$$\delta f_{12 \text{ hours}} = \pm 2\pi(84.5\text{Hz}) \frac{12\text{hours}}{1\text{year}} = \pm 0.73\text{Hz}. \quad (4-19)$$

For day to day operations this drift rate is considered to be linear. Even though 180 days of drift at this rate would yield $\delta_f = \pm 263$ Hz, which exceeds the previously established maximum drift, this overestimate is necessary due to the unknown phase of the velocity oscillations.

Extending the thought of scaling down modulations, consider the effects of signal modulation, both rotational and orbital, on the time scale of a single measurement. Much like a photograph of an object in motion, one measurement would take a snapshot of the integrated modulation. Consequently, it would then be treated as just another source of spectral broadening. Scaling the above modulation amplitudes down to a 23.8 second

time scale using 4–11 yields $\delta f_{r, 23\text{ sec}} = 1.8\text{ mHz}$ and $\delta f_{0, 23\text{s}} = 0.4\text{ mHz}$. Clearly, over the course of one measurement the combined modulations contribute a negligible amount of spectral broadening ($\lesssim 2\text{ mHz}$).

As a final note here, it is worth mentioning that by making assumptions regarding both the speed and direction of axion flows, the problems associated with signal modulation largely disappear. The phase of velocity oscillations as well as the angle between the flow direction and the detector motion would be known. Corrections could then be made to the power spectra permitting the use of an established averaging based analysis, though the limits generated from this analysis would necessarily be tied to the accuracy of the velocity assumption. For an analysis of this data set guided by such an assumption, see Hotz [79].

4.3 Variable Resolution Considerations

There are two ways to obtain a range of coarser resolutions from the time series data. One option is to perform a Fast Fourier Transform (FFT) on the entire time series data stream, then combine neighboring bins in the power spectrum (frequency domain) until the desired resolution is reached. The other option is to select the resolution by performing the FFT on a shorter length of the time series data (time domain). The following discussion, which will consider both coherent and decoherent signals, is intended to show the similarities and differences between the two methods.

Resolution adjustment in the frequency domain was the method of choice for all previous ADMX analyses. In the frequency domain, to broaden the resolution by a factor of n the power from each set of n neighboring bins is co-added¹. Any axion power that was previously spread over these n bins would add constructively increasing the

¹ Sometimes straight addition is used to combine the power instead of averaging. This a personal preference on the part of the analyst, since the division by n associated with averaging is applied equally to both signal power and uncertainty, resulting in no change in SNR.

signal power at a given frequency by a factor of n . Noise power on the other hand would add randomly, growing as \sqrt{n} instead. Assuming the resolution selected still does not exceed the signal bandwidth and that the signal is coherent over the entire duration of the measurement, the net result is an increase in the SNR by a factor of \sqrt{n} .

Resolution adjustment in the time domain is the method selected for this analysis. In the time domain, a resolution broadened by a factor of n is achieved by applying the FFT to a data stream which is shortened by a factor of n . In doing so the time series can be split into n independent sections, each generating its own subspectra. For ease of explanation, assume that the entire time series is accounted for and that the n subspectra are of equal length. Also assume that the signal is coherent over the measurement time of each individual section of the time series. This does not require that coherence be maintained from one section to the next. Because the binwidth and the integration time are scaled by the same factor, the measured power at a given frequency will be the same in each subspectrum as it was before for the full spectrum, again provided that the new resolution does not exceed the signal bandwidth. The subspectra can be recombined by summing the power at each frequency over all subspectra, increasing the total signal power by a factor of n . The noise power for each subspectrum, given by

$$P_N = k_b T_N b = k_b T_N \frac{n}{23.8 \text{sec}}, \quad (4-20)$$

also sees an increase by a factor of n . However, when recombining the subspectra, the noise power is added in quadrature. This reduces the factor of n in Eq. 4-20 to a factor of \sqrt{n} , preserving the \sqrt{n} increase in SNR.

For a signal which is coherent over the entire measurement time, both methods are valid and will produce the same SNR at the coarser resolution. If, however, coherence is lost part way through the measurement, this equality ceases to be true. In this case, the time scale on which decoherence occurs becomes critical. A signal which loses coherence over the full measurement duration can potentially remain coherent over the

individual sections of the time series. This would occur for decoherence times which are longer than the time series sections. Because power from a decoherent signal would not necessarily add constructively when the FFT is performed, resolution adjustment in the time domain may be preferable under these circumstances.

Consider a signal which losses coherence after roughly 10% of the full measurement time. For an entire 23.8 second measurement the signal would appear highly decoherent, while for measurements which are only 2.38 seconds long the signal would remain mostly coherent. Such a signal was modeled as a perfect sinusoid which acquires a random phase once coherence is lost. Thus, roughly every 2.38 seconds the phase of the signal changes. The exact points where the phase changes occur were also determined randomly having a gaussian probability distribution about each 2.38 second interval. This signal was added to randomly generated exponential noise (SNR=50) and both methods were tested. Figure [4-1](#) shows the average results of 10,000 iterations, with the power for a perfectly coherent signal being normalized to 1. For a perfectly decoherent signal, the power would add completely randomly just like the noise. For this test $n = 10$, leading to a maximum loss of 68.4% when compared to a perfectly coherent signal. The average power loss in the time domain and frequency domain are found to be 5% and 38% respectively.

As mentioned in Chapter 3, the finest resolution of the HR channel (84 mHz) is set by the stability of the frequency standard used in the receiver chain. While this has been looked at as the resolution at which a signal will have drifted by no more than 1 bin during a single measurement, it can also be looked at as a loss of coherence after about 12 seconds. It is for this reason that resolution selection in this analysis is performed in the time domain.

As a final thought on the differences between these two methods, consider once more a perfectly coherent signal. Being again equals in terms of SNR, the only real difference between the time domain and frequency domain would concern their ease of

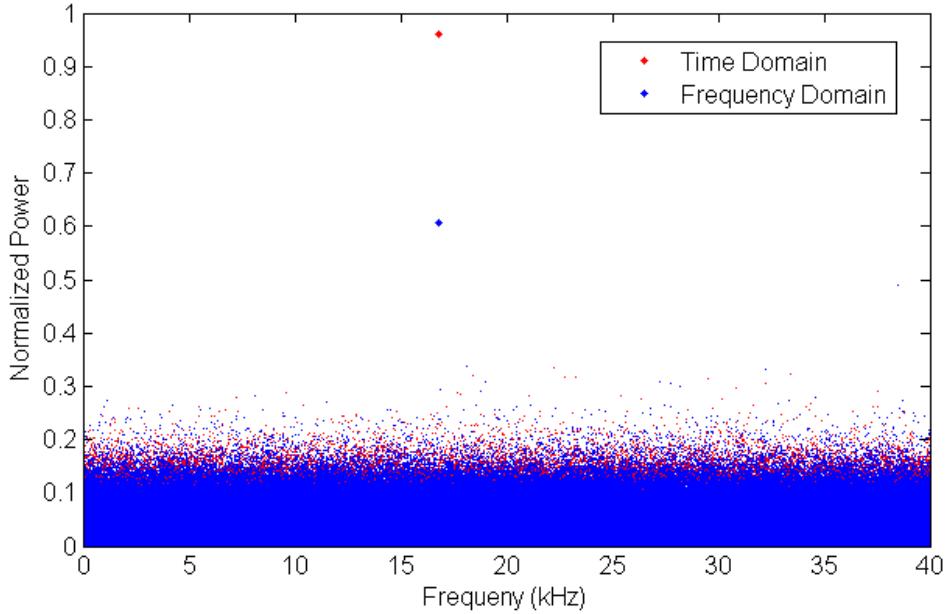


Figure 4-1. A numerical simulation of the power measured in a signal with an average decoherence time of 2.38 seconds. For resolution selection in the frequency domain, the average loss of 38% is well below the theoretical maximum of 68.4%.

use. Computationally, resolution selection in the frequency domain is far superior to that in the time domain. This is not due to any additional complexities in the latter, rather it is simply a function of the number of computations needed to achieve the same goal. Modern FFT algorithms require $\sim N \log_2 N$ floating point calculations to transform a data stream of length N . To change the resolution in the time domain, n FFTs of length $\frac{N}{n}$ are performed ($\sim N \log_2 \frac{N}{n}$ calculations). When compared to the frequency domain's single FFT of the full N point data set, there is a small reduction ($\sim N \log_2 n$) in the number of computations needed. Unfortunately, each resolution requires its own FFT in the time domain. Subtracting a few calculations simply does not make up for multiplicative increase in computational time.

4.4 Noise Properties

The conversion from the unnormalized units of the raw power spectrum into Watts is done via relation to the thermal rms noise power. An understanding of the nature

of the noise in the HR channel is important for a proper handling of the HR analysis. The unmodified 42 mHz resolution power spectrum will be the starting point of this discussion. Since no averaging is done between power spectra, a gaussian treatment of the 42 mHz resolution noise would be inaccurate. Instead, at this resolution, the HR channel is expected to exhibit an exponential distribution for thermal noise with a standard deviation equal to the rms noise power of the experiment, P_N . A derivation of why this exponential behavior is expected can be found in Duffy [13], though here it will be treated simply as a condition which must be verified. Upon doing so, the exponential nature of the 42 mHz noise will serve as the basis for ensuring that the noise statistics used for coarser resolutions are correct.

Confirming of the shape of the noise distribution is relatively straight forward, while relating this shape to the noise power requires a few calculations. For both, start with the probability density function (pdf) for any exponential noise distribution,

$$\frac{dP}{dp} = Ae^{(-Ap)} H(p). \quad (4-21)$$

Here A is a constant, P is the probability of the measuring a noise power of p or less, and $H(p)$ is the heaviside step function. A plot of this noise power shown on a semilog scale should be a straight line. Figure 4-2 shows a histogram for the 42 mHz resolution noise clearly demonstrating the appropriate behavior.

To relate Fig. 4-2 to the noise power, start with the mean value of p for Eq. 4-21, given by

$$\langle p \rangle = \int_0^{\infty} pAe^{(-Ap)} dp = \frac{1}{A}. \quad (4-22)$$

This leads to a variance in p of

$$\sigma^2 = \left(\int_0^{\infty} p^2 A e^{(-Ap)} dp \right) - \langle p \rangle^2 = \frac{1}{A^2}. \quad (4-23)$$

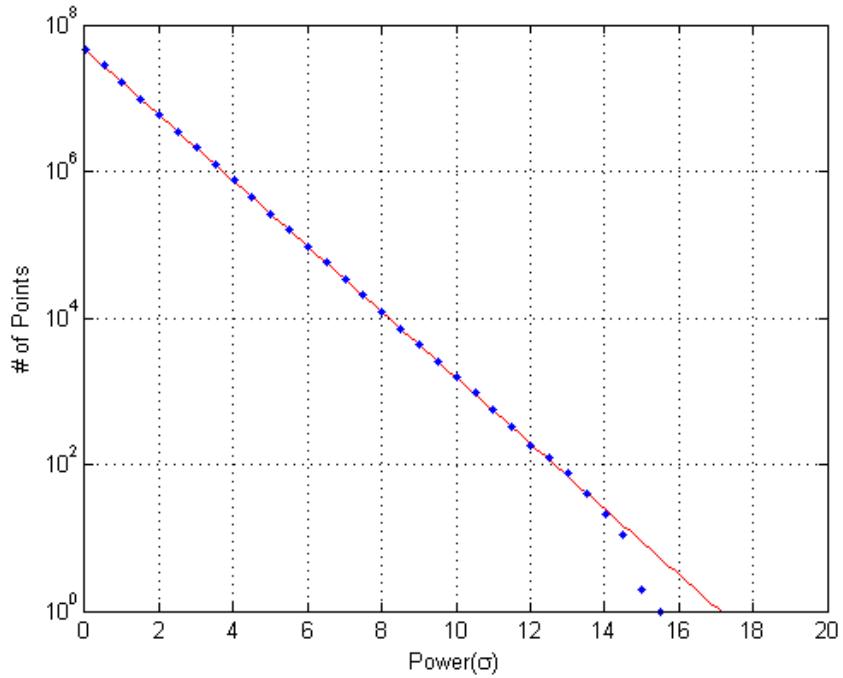


Figure 4-2. Power distribution for a sample of 42 mHz data (n=1).

The variance is, of course, the square of the standard deviation, σ . Setting σ equal to the rms noise power as stated above yields the relation

$$\sigma = \langle p \rangle = \frac{1}{A} = k_b T_N b. \quad (4-24)$$

Because 4-24 is a function of resolution, a slight adjustment to its notation should be made to reflect the different values it takes. From now on the rms noise power for a given resolution will be denoted as σ_n , where n is once again the number of subspectra generated from the time series for that resolution. Eq. 4-21 can then be rewritten for positive p as

$$dP = \frac{dp}{\sigma_1} e^{(-p/\sigma_1)} \quad (4-25)$$

for the 42 mhz pdf, though each subspectra for coarser resolutions would have similar distributions,

$$dP_{n,exp} = \frac{dp}{\sigma_n} e^{(-p/\sigma_n)}. \quad (4-26)$$

Discretizing Eq. 4–25 and multiplying by the total number of points, N , in the power spectrum gives an expression for the number of points with noise power between p and $p + \Delta p$ in the 42 mhz resolution spectrum,

$$N_p = \frac{N\Delta p}{\sigma_1} e^{(-p/\sigma_1)}. \quad (4-27)$$

Taking the natural log of Eq. 4–27 gives

$$\ln N_p = \ln \frac{N\Delta p}{\sigma_1} + \frac{-p}{\sigma_1}, \quad (4-28)$$

which perfectly describes the shape of Fig. 4-2, with σ_1 being the negative inverse of the slope. Note that while $\langle p \rangle$ and σ_1 should be equal, non-statistical noise sources can cause there to be a slight difference. Consequently, the formal assignment of noise power as the standard deviation is maintained.

For coarser resolutions, the noise no longer follows an exponential distribution, though each subspectra still does. Instead, the final pdf for a given resolution becomes the convolution of n exponential distributions, each having an rms noise power of σ_n . For example the 84 mHz noise, $n = 2$, has a pdf given by

$$\frac{dP_2}{dp} = \left(\frac{dP_{2,\text{exp}}}{dp} * \frac{dP_{2,\text{exp}}}{dp} \right) = \frac{1}{\sigma_2^2} \int_{-\infty}^{\infty} e^{(-p'/\sigma_2)} e^{(-(p-p')/\sigma_2)} H(p') H(p-p') dp'. \quad (4-29)$$

When evaluated, Eq. 4–29 reduces to

$$\frac{dP_2}{dp} = \frac{p}{\sigma_2^2} e^{(-p/\sigma_2)} H(p). \quad (4-30)$$

Figure 4-3 shows that the noise distribution for a subset of the 84 mHz data as well as a line given by Eq. 4–30 are in good agreement. Obtaining pdfs for larger n comes from chain convolving exponential distributions as in

$$\frac{dP_3}{dp} = \left(\left(\frac{dP_{3,\text{exp}}}{dp} * \frac{dP_{3,\text{exp}}}{dp} \right) * \frac{dP_{3,\text{exp}}}{dp} \right) = \frac{p^2}{2!\sigma_3^3} e^{(-p/\sigma_3)} H(p). \quad (4-31)$$

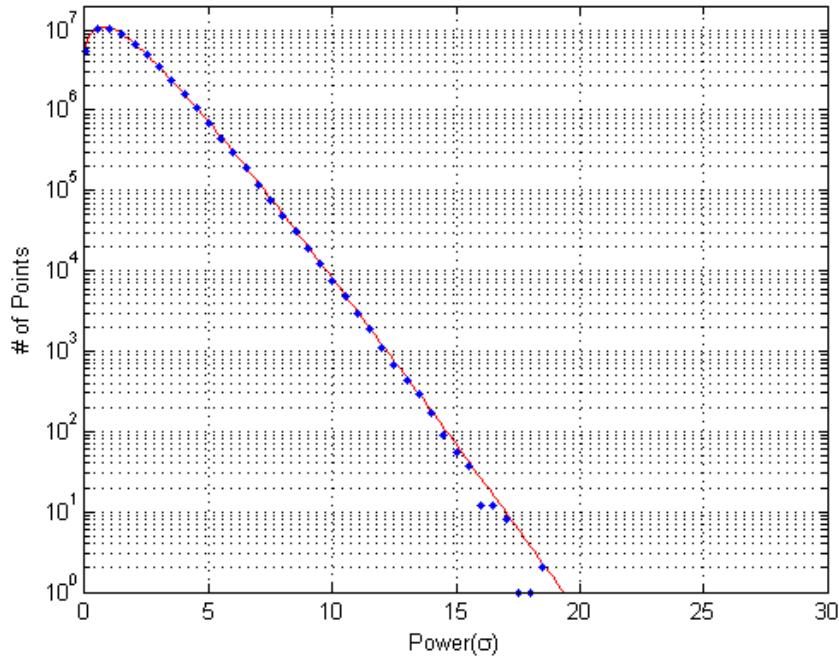


Figure 4-3. Power distribution for a sample of 84 mHz data (n=2).

After a few long but straight forward integrations, a generic result can be extrapolated.

The pdf for n combined subspectra is given by

$$\frac{dP_n}{dp} = \frac{p^{(n-1)}}{(n-1)!\sigma_n^n} e^{(-p/\sigma_n)} H(p). \quad (4-32)$$

Figures 4-4, 4-5, and 4-6 show the noise distributions for the 168 mHz, 546 mHz, and 1.09 Hz resolutions, as well as their respective pdfs. For each, the same subset of data was used as was used for Fig. 4-3. As dictated by the central limit theorem, the recombining of subspectra does causes the noise distribution to become ever so slightly more gaussian, though too few subspectra are ever used, even in the coarsest resolution, to make this an accurate approximation.

4.5 HR Data Set and Analysis Software

The data set used for this analysis is a subset of the data acquired during ADMX's most recent science run, which began in May 2008 and ended in April 2010. For roughly half of this duration the DAQ was not instructed to save the time series data for the

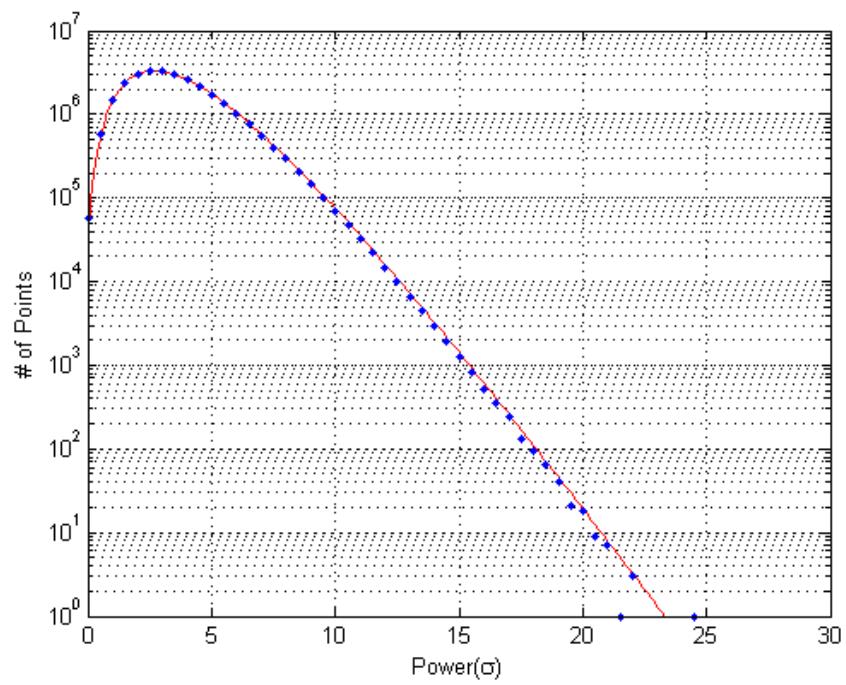


Figure 4-4. Power distribution for a sample of 168 mHz data (n=4).

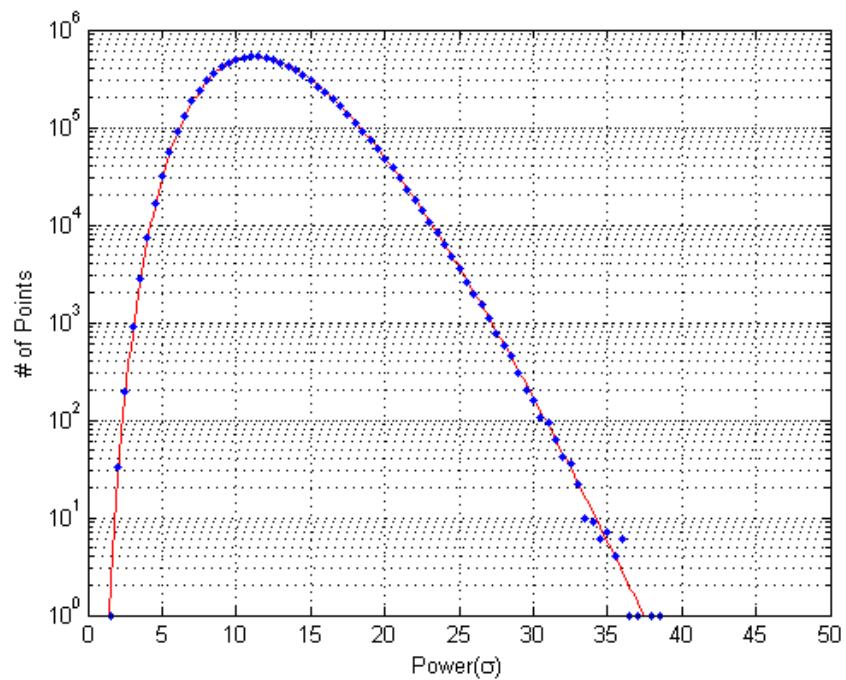


Figure 4-5. Power distribution for a sample of 546 mHz data (n=13).

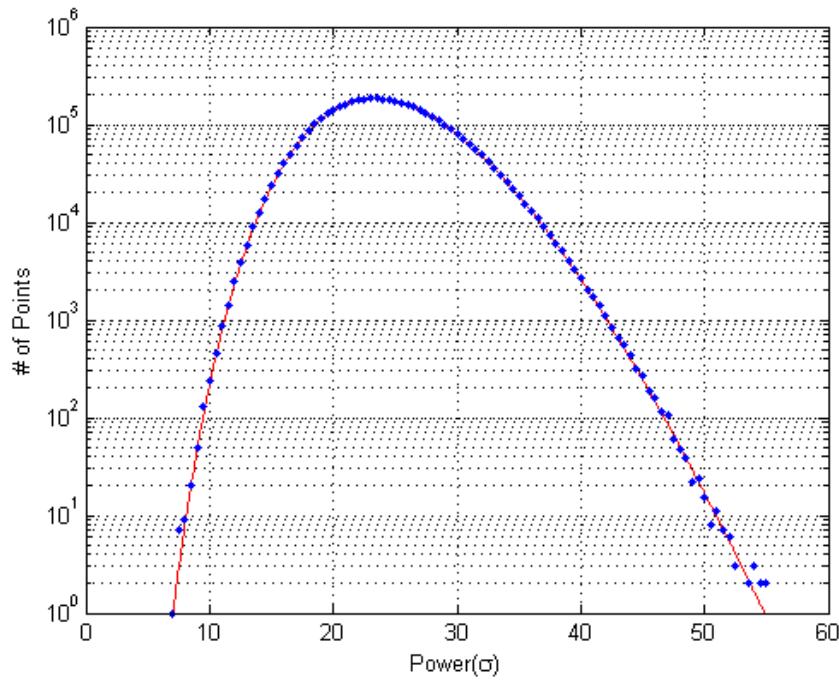


Figure 4-6. Power distribution for a sample of 1.09 Hz data (n=26).

HR channel. Instead, it only saved a composite, single-sided 42 mHz resolution power spectrum for every three spectra measured. This put the data acquisition time of the HR channel on the same scale as that of the MR channel, and reduced the disk space required for storing the raw data by a factor of six. Additional storage was acquired allowing both data formats to be saved to disk. The saving of the time series data began on March 31, 2009 and continued through the end of the science run on April 1, 2010. For this half of the science run, a total of 274,834 time series files were saved, covering a frequency range of 812 to 852 MHz and 858 to 892 MHz. The gap in frequency coverage from 853 to 857 MHz is due to having insufficient time series data in that range. This is a direct result of the late implementation of saving the time series data.

Each file was saved as a binary file consisting of a header (5,602 Bytes) containing the state data, the names and values of various experimental parameters, followed by a list of 1,906,840 single precision voltage measurements (7.6 MB). MatLab 7.4 is the software selected to handle the processing of these data files. The decision

to use this software package was based primarily on the ease with which MatLab handles matrix calculations, and the availability of built-in functions for input/output operations, performing FFTs, data fitting, plotting, plot editing, and even the generation of pseudo-data sets for a given pdf. This last function proved quite useful while testing the analysis code. Additionally, because MatLab is available for both Windows and Linux, the analysis code easily runs on either platform. After reading in and processing a file in the above binary format, the output is saved as a MatLab binary file (.mat extension) which preserves the structure and organization of variables in the MatLab workspace.

The HR data are recorded over 23.8355 seconds at a sampling rate of 80 kHz. Applying an FFT to the entire data stream yields a spectrum with a bandwidth of 40 kHz and a spectral resolution of 42 mHz. For each coarser resolution, the bandwidth remains at 40 kHz, while the spectral resolution is scaled by a factor of n . Due to the stability of the frequency standard used in the receiver chain causing an uncertainty in the frequency larger than 42 mHz, data at this resolution are not used for generating the final limits. However, these data are used to confirm the noise statistics of the HR data and to adjudicate some of the data quality cuts which are not related to frequency. As such, every step of the analysis up to limit generation is still carried out on the 42 mHz resolution power spectra.

4.6 Data Quality Cuts

A series of data quality cuts are applied to the data set to weed out files containing data unsuitable for analysis. In the most simple case, a file might have one or more voltage measurements recorded as a non-finite value, saved by the DAQ as "Not a Number" or NaN. The inclusion of a non-finite value in the data stream causes the FFT to return non-finite values as well, rendering the file useless. Other cuts applied prior to the FFT check that experimental parameters Q , B_0 , T_{Cavity} , and FFT_{avg} all fall within expected norms.

Files which pass these cuts are used to generate subspectra which are then processed to remove systematic errors. Described in more detail below, the removal of systematic errors includes the fitting of the overall shape of the subspectrum. How well this fit matches the subspectrum is determined by the reduced chi squared, χ_r^2 , of the fit. Large deviations of χ_r^2 from 1 indicate a poor fit to the subspectrum likely caused by some unknown systematic error. For this reason, data with $\chi_r^2 < 0.8$ or $\chi_r^2 > 1.2$ are cut.

Following the removal of systematic errors, all data should follow exponential statistics with few if any discrepancies. This is useful in determining if there is some outside source of power leaking into the data stream such as can be seen in Fig. 4-7. This sort of contamination is easily seen by looking at the distribution of points in the 42 mHz resolution. A semilog plot of the 42 mHz distribution for the data shown in Fig. 4-7 is shown in Fig. 4-8. The outside power shows up as a massive deviation from the expected linear shape of the plot. Data showing signs of outside contamination are cut. Numerous spectra taken between February 19, 2010 and March 1, 2010 show signs of contamination making this, by far, the largest cut applied to the data set. The source of this contamination was discovered to be a faulty RF cable in the receiver chain. No signs of contamination have been found following the replacement of the cable. Of the 274,834 unique time series files, 246,556 files pass all data quality cuts.

4.7 Removal of Systematic Effects

In addition to statistical noise, there are a number of systematic effects which must be removed from the data before it can be searched for an axion signal. These corrections are performed on each subspectra before recombination. All resolutions are treated the same way, following the steps now outlined. First the spectral shape of the crystal filter is removed. Then the edges of the power spectra are trimmed and the receiver response function is removed. Finally, the power is scaled to the system noise power for the appropriate resolution. At this point the subspectra are recombined to recover SNR.

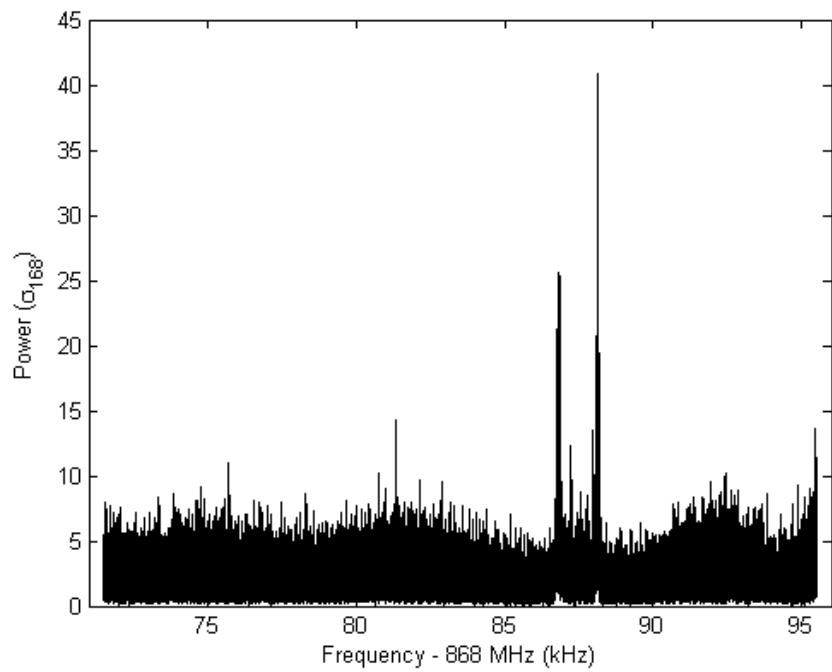


Figure 4-7. A spectrum showing contamination from an external radio signal.

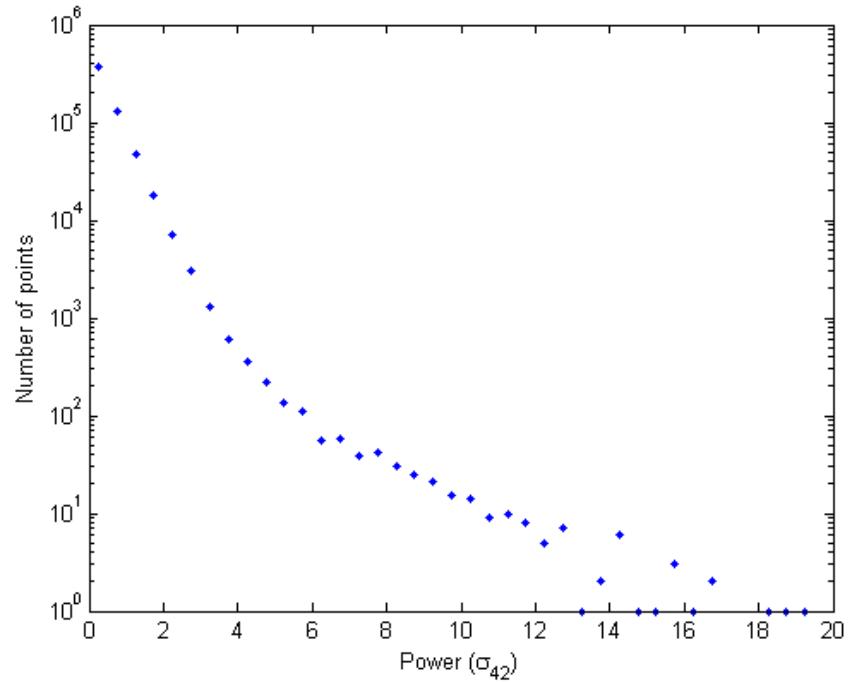


Figure 4-8. The 42 mHz power distribution for the data shown in Fig. 4-7. Large deviations from the exponential distribution caused by the radio signal can be seen at higher powers.

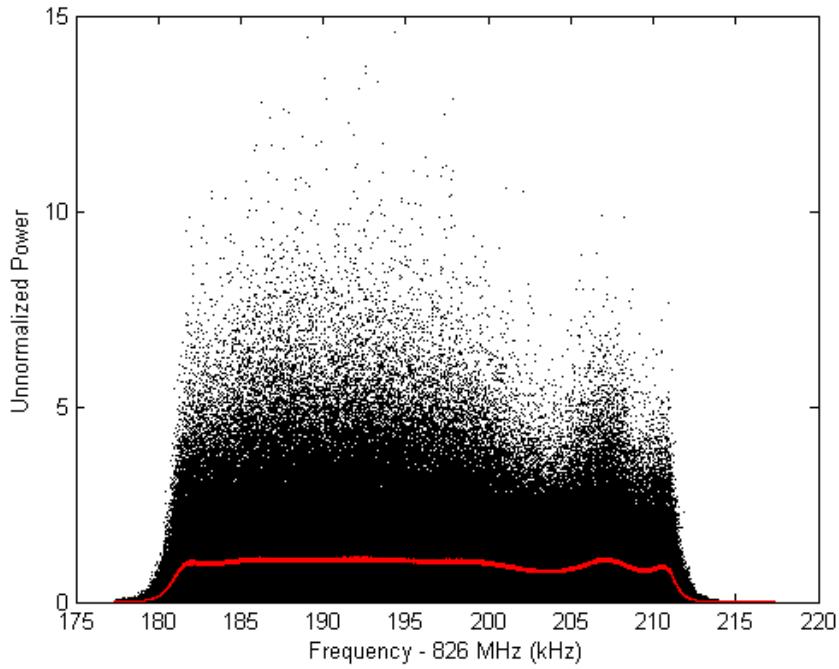


Figure 4-9. An example of an unnormalized, HR spectrum showing the shape imparted to the data by the crystal filter. The red line is the reference spectrum shown in Fig. 4-10.

The most noticeable of the systematic effects is the shape imparted to the power spectrum by the crystal filter, shown in Fig. 4-9. To facilitate the removal of this feature, a set of reference files is created for each resolution. These reference files are constructed by normalizing and averaging several days worth of power spectra, roughly 10,000 files, which would have had a range of center frequencies. This ensures that both statistical noise and systematic noise associated with the receiver response function are sufficiently averaged out. Three such files are made per resolution, each corresponding to a different operational time frame. The resulting shape, shown in Fig. 4-10, is representative of the spectral shape of the crystal filter and can now be divided out.

The bandwidth of the crystal filter is 30 kHz wide which is smaller than the 40 kHz wide power spectrum. As a result, there is dead space beyond the usable crystal filter bandwidth, clearly visible in Fig. 4-9, which must be removed. This is accomplished

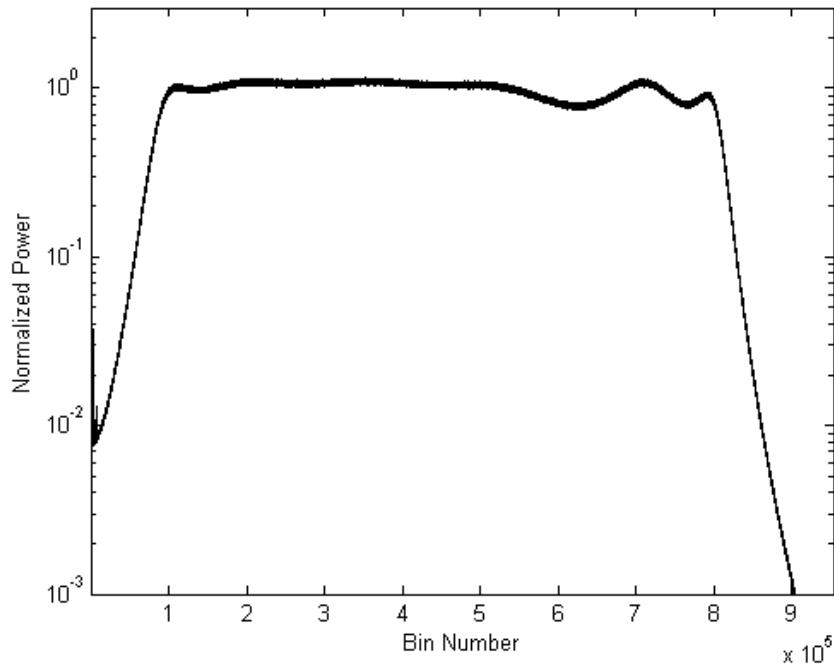


Figure 4-10. A reference spectrum which is used to remove the crystal filter shape from HR data. This spectrum is the average of about 10,000 individual spectra.

by simply truncating the beginning and end of the power spectrum. To ensure that the roll off at the edge of the 30 kHz bandwidth is not clipped, the first and last 20% of the power spectrum are removed. This leaves 24 kHz of usable bandwidth for each subspectrum.

Following truncation, there is usually some broad spectral structure, called the receiver response, which still remains. This structure is due to noise power which originates in the amplifiers and propagates towards the cavity where upon it is reflected back up the receiver chain. The receiver response varies with the frequency range covered by the power spectrum, thus affecting each data file slightly differently. The receiver response on each subspectra of a given data file should be identical, as they cover the same frequency range. To remove the receiver response, neighboring bins in the power spectrum are averaged to create a coarser resolution spectrum spectrum which is then fit to a polynomial of degree 9 called the receiver response function. Very

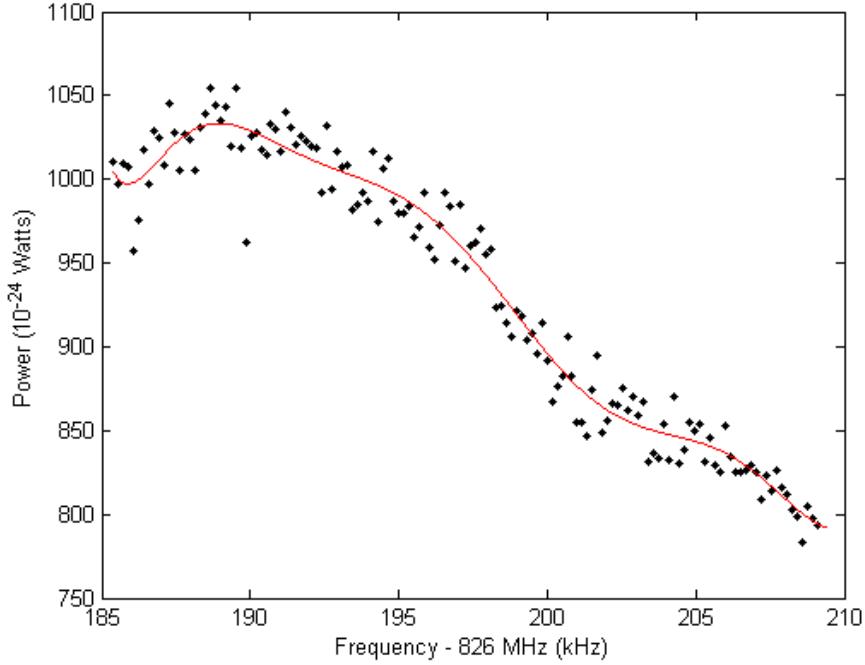


Figure 4-11. The receiver response function fit to the residual structure of the truncated spectrum.

large peaks ($> 18\sigma$) are ignored when creating the coarse spectrum so as to not bias the fit. The resolution of the coarse spectrum is 172 Hz for the 84 mHz and 168 mHz resolution subspectra, and 139.8 Hz for the 546 mHz and 1.09 Hz resolution subspectra. The coarse spectrum is sufficiently averaged such that the exponential nature of the individual subspectra does not affect the fitting of the receiver response function. For the fitting of the receiver response function, the 24 kHz bandwidth of the coarse spectrum is mapped to an interval of -0.5 to 0.5. It is for this fit that the χ^2_r used in the aforementioned cut is determined. An example of this fit is shown in Fig. 4-11. Once fit, the receiver response function is divided out of the original truncated subspectrum leaving it spectrally flat. This can be seen in Fig. 4-12.

Previous analyses [77] modeled an equivalent circuit for the receiver-chain/cavity coupling [80] which led to a receiver response function that was the ratio of two quadratic expressions having a total of five free parameters. While these free parameters

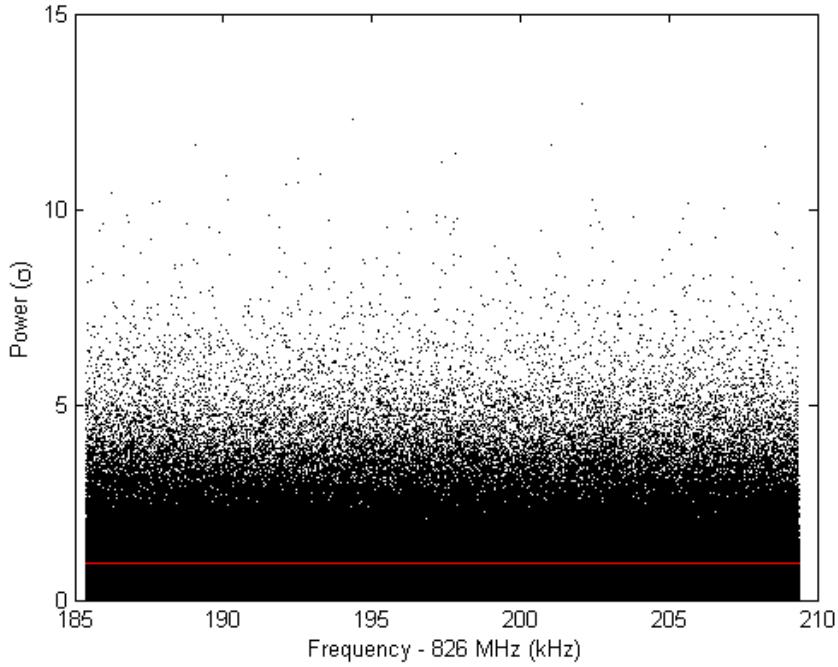


Figure 4-12. An example of a flat HR Power spectrum following truncation and the removal of the receiver response function.

had well defined relations to variables such as T_N , f_0 , Q , etc.; in practice the parameter values were still obtained via non-linear fitting. A simple polynomial fit was adopted for two reasons. First, it is easier to fit than the expression stemming from the equivalent circuit model. Second, the equivalent circuit model requires the interactions between the first stage amplifier and the cavity to be very well understood. When this model was used, the first stage amplifier was an NRAO HFET amplifier for which this condition was met. This is no longer the case, as the interactions between the MSA and the cavity are significantly more complicated. The degree of the polynomial now used was determined by trying different fits, and selecting the one that most often accounted for the entire structure of various power spectra. While the degree of the polynomial may seem high, the removal of an actual axion signal using this fit is not expected. This is due to an axion signal being at least two orders of magnitude narrower than the resolution of the coarse spectrum, and narrower still when compared to the scale of structure in the

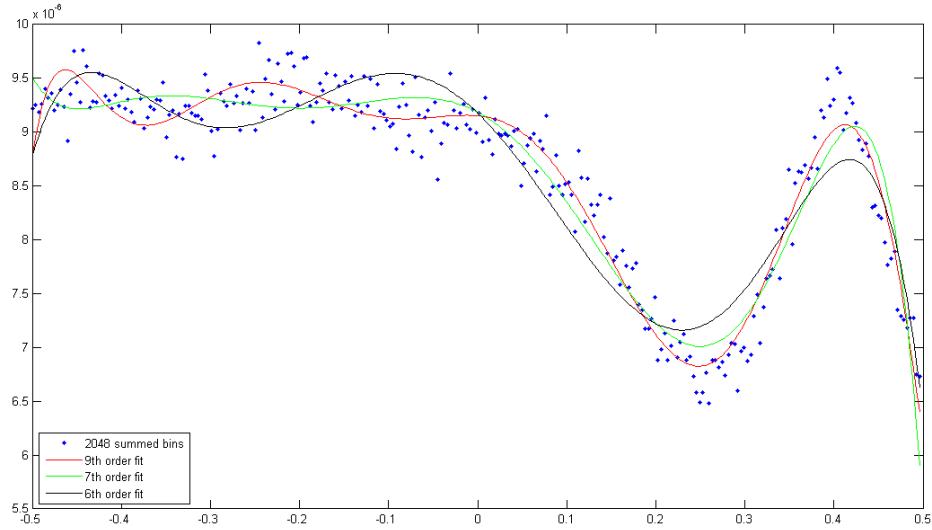


Figure 4-13. A down-sampled HR power spectrum which has been fit with polynomials of degree 6, 7, and 9. The highest order of these three shows the best agreement with the residual shape of the spectrum.

receiver response. Figure 4-13 shows an example of a coarse spectrum overlaid with three fits of varying degrees.

With the receiver response now removed, the subspectrum can be rescaled to units of σ_n . Having a noise distribution given by Eq. 4-26, a histogram of the subspectrum is plotted on a semilog scale and the slope, m , of the resulting line is determined via linear fit. Once again, very large peaks in the power spectrum are ignored in plotting this histogram so as to avoid biasing the fit. Consistent with the Eq. 4-28, this slope is the negative inverse of the standard deviation of the distribution, which was found earlier to be the rms noise power, σ_n . Additionally, the standard deviation is taken to be a measure of the uncertainty in power of each point in the subspectrum. All points in the subspectrum are then scaled by $\frac{-1}{m}$, putting both the power and the uncertainty in units of σ_n . Now appropriately scaled, all of the subspectra for a given file are recombined to produce a single spectrum for that file. This final spectrum is still in units of σ_n , and has a noise distribution given by Eq. 4-32. The uncertainties for each subspectra are added in quadrature, leading to an uncertainty in the final spectrum of \sqrt{n} in units of σ_n . The

value of σ_n is saved in the header of the processed file along with the time stamp, f_0 , Q , B_0 , and C . These values are used later to calculate exclusion limits.

4.8 Candidate Search

The recombined spectra can then be searched for points which exceed a certain power threshold, T . All points with power exceeding T are considered axion candidates, and must then be looked for in each other spectrum which overlapping frequency coverage. Because the comparison of candidates in different spectra must still account for the expected modulation from terrestrial motion, a simple matching up of frequencies is insufficient. A different means of comparing must be implemented.

Start by considering the candidates found in only one spectrum (S_1). Some time, δt , later, if they actually are axion signals, the candidates will necessarily have drifted by some amount. For these candidates, the maximum possible drift, δf_{max} , is set by δt and Eq. 4–11. A second spectrum (S_2), measured at this later time, may also contain candidates. Should any of the candidates in S_2 be at a frequency near any of the candidates from S_1 , specifically within $\pm\delta f_{max}$, there exists the possibility that they are in fact the same signal. Such candidates are considered to be coincident. Note that the sign of δt is irrelevant since the comparison is made within $\pm\delta f_{max}$. Consequently, it does not matter which spectrum was measured first. Were the indices of S_1 and S_2 swapped in this example, the magnitude of the maximum drift would be the same yielding identical results for coincidence.

For only two spectra, this comparison process is still relatively simple. However, each candidates in S_1 must be compared to each candidate in each spectra with overlapping coverage. Recall that the usable bandwidth of the power spectrum is 24 kHz and the the tuning step of f_0 is 2 kHz. This alone results in there being numerous spectra covering the same frequencies, not to mention spectra measured days or even weeks later which also overlap. Let the set of these spectra be denoted as

$$\mathcal{S} = \{S_1, S_2, S_3, \dots, S_N, \}. \quad (4-33)$$

Each spectra in \mathcal{S} has a different δt , thus a different δf_{max} , when compared to S_1 . Further, each member of \mathcal{S} must then be compared to S_2 , followed then by comparison to S_3 , then to S_4 , and so on. Each of these sets of comparisons has its own set of values for δt . There are two considerations that slightly reduce the number of comparisons required. First, candidates in the same spectrum need not be compared to each other. Because δt is 0 for this case, all candidates in the same spectrum are independent of each other. Second, only one comparison need be made between a given pair of spectra. That is, S_i need not be compared to $S_{j < i}$ because this comparison would have already been made when checking the candidates in S_j . Even with these conditions cutting the number of comparisons by roughly a factor of two, the number of calculations needed is still on the order of the square of the total number of candidates, causing this means of searching to become very cumbersome, very quickly.

Prior to searching for and comparing candidates, the spectra were organized chronologically and assigned an numerical index. An index file is then generated which contains the information saved in the headers of the processed spectra. Then a list of candidates from all spectra is constructed consisting of three numbers for each candidate. These numbers are the frequency of the candidate, the power of the candidate, and the index number of the spectrum containing the candidate. Because there can be multiple candidates per spectrum, the use of the index number permits header information to be written only once for all candidates in the same spectrum, thus greatly reducing the amount of memory required.

One at a time, each candidate is then considered. Let the candidate in question be denoted as C_0 . First the difference in frequency, δf_i , between C_0 and each other candidate, C_i , is calculated. This list of δf_i can be truncated in accordance with the two previously stated considerations. Also, a logical argument regarding widely separated candidates can be used to further shorten this list. The worst-case estimate for signal modulation is found from Eq. 4-21 as being ~ 200 Hz). Including a factor of 2 due

to the uncertainty associated with the mean velocity of the axion flow, a factor of 2 for a conservative estimate, and rounding up to further err on the side of caution, this maximum modulation becomes ~ 1 kHz. Any δf larger than this value could not be caused by terrestrial motion and can be dropped from the list. This drastically reduces the number of comparisons that need to be made.

For any surviving C_i , δt_i is calculated as the difference in time stamps between the spectra containing C_0 and C_i . The list of δt_i is used to calculate a list of maximum modulation values, $\delta f_{i, \max}$. All C_i with $\delta f_i < \delta f_{i, \max}$ are considered coincident with C_0 . A tally of the number of spectra containing at least one coincident candidate is recorded and the search repeats for the next candidate. Any real axion signal would be steady, thus it should be consistently present in spectra with overlapping frequency coverage. The frequency coverage of each spectrum is limited to the full width at half maximum ($FWHM$) of the Lorentzian cavity resonance about f_0 . This value is given by

$$FWHM = \frac{f_0}{Q}. \quad (4-34)$$

Because an axion signal beyond the $FWHM$ is suppressed and will likely not exceed T , the inclusion of such frequencies would bias the search by artificially inflating the number of coincident spectra required for statistically significance. Any C_0 which shows some coincidence but fails to appear in half of the overlapping spectra is determined to have statistically insignificant coincidence. The decision to use 50% of overlapping spectra for this cutoff is addressed below under Density Limits.

Some consideration must be given to the value selected for T . If T is too high, there may not be enough candidates to produce meaningful limits. If T is too low, the number of candidates increases exponentially, leading to unreasonable computational times. The thresholds for each resolution were set such that a pure noise peak would be counted only once for roughly every 10 spectra examined. The values of these T_n can be calculated by integrating Eq. 4-32. For ρ in units of σ_n and a spectrum length of N_n

points, T_n satisfies the expression

$$\int_0^{T_n} \frac{p^{(n-1)}}{(n-1)!} e^{(-p)} dp = 1 - \frac{n}{10N_n}. \quad (4-35)$$

When evaluated Eq. 4-35 becomes

$$e^{(-T_n)} \left(\sum_{j=0}^{n-1} \frac{T_n^j}{j!} \right) = \frac{n}{10N_n}. \quad (4-36)$$

Equation 4-36 was solved numerically for each resolution yielding target thresholds of 17.6, 21.8, 36.6, and 55.2 for $n = 2, 4, 13$, and 26 respectively. The actual thresholds used were set slightly lower than these values at 16, 20, 34, and 51 respectively, resulting in roughly 60,000 candidates per resolution. Setting slightly lower thresholds has the effect of improving the final limits while not severely bogging down the search algorithm.

No candidates at any of the examined resolutions in the frequency range from 812 to 852 MHz and 858 to 892 MHz have shown significant coincidence. As such, upper limits on the local non-virialized axion density are derived from this data set.

4.9 Density Limits

While candidates are defined as all points measured to have power over the threshold T_n , not all of this power is considered signal power. The signal power is actually the candidate power minus the noise power in that frequency bin. It is important to note that this is no longer the rms noise power. It is a single measurement of the noise power, and as such it must be treated in accordance with its pdf. To that end, the probability of the signal power and noise power summing to a value which exceeds T_n is called the confidence level (CL). This value is a free parameter in analysis and is chosen here to be 90%.

With the CL established, a determination about the noise power contribution can now be made. Consider noise powers given by the first 10% (1-CL) of the noise power distribution. There exists a value, $T_{\text{eff}, n}$, for the signal power such that the addition

of even the largest of these noise powers would be insufficient in reaching T . Noise powers in the remaining 90% of the distribution when added to $T_{\text{eff}, n}$ would then cause the candidate power to exceed T_n . This leads to the expression

$$T_{\text{eff}, n} = T_n - NP_n, \quad (4-37)$$

where NP_n is determined in a similar fashion to T_n . A slight modification to Eq. 4-36 gives

$$e^{(-NP_n)} \left(\sum_{j=0}^{n-1} \frac{NP_n^j}{j!} \right) = 90\%, \quad (4-38)$$

which when solved numerically gives the maximum allowed noise power, in units of σ_n , as 0.5, 1.7, 8.6, and 19.7 for $n = 2, 4, 13$, and 26 respectively. These values are subtracted from T_n yielding effective thresholds of 15.5, 18.3, 25.4, and 31.3 respectively. The effective thresholds are used below in calculating density limits for the HR data.

Consider once again the expression for P_a given by Eq. 4-1. In this expression, for a given mass, there are two free parameters. These are the local dark matter density, ρ , and the axion to two photon coupling, $g_{a\gamma\gamma}$. In order to set a limit on one of these parameters it is necessary to assume a value for the other. Lacking additional constraints on either parameter causes the decision of which to assume to be rather arbitrary. Traditionally, ADMX has assumed a value for the local dark matter density ($\rho = 0.45 \text{ Gev/cm}^3$) and set limits on $g_{a\gamma\gamma}$. While this choice is reasonable when placing limits on virialized axions, it becomes less so when applied to non-virialized axions. What fraction of the local density is comprised of non-virialized axions? Of that fraction, how much can be expected in each flow? The large uncertainties associated with the local density of non-virialized axions make an assumption on $g_{a\gamma\gamma}$ the optimal choice for HR data. Here, the assumptions of KSVZ and DFSZ axions are made, and limits are placed on the local density of non-virialized axions. Evaluating Eq. A-3 for $g_{\gamma, \text{KSVZ}} = -0.97$ and $g_{\gamma, \text{DFSZ}} = 0.37$ and converting f_a to m_a using Eq. 2-20 yields the

following values for $g_{a\gamma\gamma}$.

$$g_{KSVZ} = -2.25 \times \left(\frac{m_a}{6\mu\text{eV}} \right) \times 10^{-15} \text{GeV}^{-1} \quad (4-39)$$

$$g_{DFSZ} = 8.59 \times \left(\frac{m_a}{6\mu\text{eV}} \right) \times 10^{-16} \text{GeV}^{-1} \quad (4-40)$$

Because $g_{a\gamma\gamma}$ appears in all calculations as $g_{a\gamma\gamma}^2$, the difference in sign is insignificant.

The calculation of exclusion limits for each resolution starts by relating the measured power at that resolution's threshold to the power expected from an axion signal. This relation is not simply unity as there are three factors which reduce the amount of power actually seen in the cavity. Two of these factors have already been mentioned. First is the factor of 2 loss in signal power caused by the critical coupling of the cavity to the receiver chain. This loss is constant for all files. Second is the Lorentzian factor L caused by a signal being located away from the resonant frequency of the cavity. This loss is given by Eq. 4-4 for each file. Because the measured Q differs from file to file, there is a large variation in L over the data set. Further still L varies with frequency even within a single spectrum. What then should be the value of L used in determining the limit for a given spectrum. While a conservative treatment of L might use its minimum value which occurs when the frequency offset from resonance is maximized, this is perhaps too limiting. For any spectrum with a reasonably high Q , the value of L will fall off drastically long before reaching the edge of the spectral bandwidth. Recall as well that the effective frequency coverage of a given spectrum is limited to the $FWHM$ about f_0 . Even using the effective minimum of $L = 0.5$ is questionable. Because an axion signal has equal probability to be found at any frequency in the effective bandwidth, the mean value of L over this range should be used. For any spectrum where the $FWHM$ is smaller than the spectral bandwidth, this value is a constant and is given by

$$\bar{L} = \frac{Q}{f_0} \int_{f_0(1-\frac{1}{2Q})}^{f_0(1+\frac{1}{2Q})} L(f) df = \frac{\pi}{4} \quad (4-41)$$

For spectra with low values of Q , the $FWHM$ may be larger than the spectral bandwidth.

The effective bandwidth is thus unchanged and \bar{L} is instead given by

$$\bar{L} = \frac{Q}{f_0} \int_{f_0-12 \text{ kHz}}^{f_0+12 \text{ kHz}} L(f) df = \frac{f_0}{(24 \text{ kHz})Q} \arctan \left(\frac{(24 \text{ kHz})Q}{f_0} \right). \quad (4-42)$$

In both cases the probability of the signal being at a frequency with $L \geq \bar{L}$ is given by

$$Pr(L(f) \geq \bar{L}) = \frac{\text{Width at } \bar{L}}{\text{Min}(FWHM, 24 \text{ kHz})} \approx 50\%, \quad (4-43)$$

This is the justification for needing half of the overlapping spectra to show coincidence in order to be statistically significant.

The third reduction is an artifact of the discrete nature of the FFTs. An FFT will output each frequency's contribution to a signal, however this is only done for integer multiples of the spectral resolution, given by Eq. 3-15. These allowed frequencies form a basis set,

$$\mathcal{F} = \{0, \pm b, \pm 2b, \pm 3b, \dots\}, \quad (4-44)$$

which is then used to interpret frequencies which are not integer multiples of b . Each such frequency is written as a liner combination of these basis frequencies, causing the power to be spread across multiple nearby frequencies. This spread heavily favors the two closest frequencies in \mathcal{F} . For a signal at frequency $f_s = b(j + \Delta)$, the contribution at frequency $f_j = jb$ is given by

$$P_j = \left(\frac{\sin(\pi\Delta)}{\pi\Delta} \right)^2. \quad (4-45)$$

This power is shown in Fig. 4-14 for f_s being offset from a basis frequency by an amount which ranges from $-b$ to $+b$. At the extremes of Fig. 4-14 the signal is simply at the next higher or lower basis frequency, thus generating no contribution at f_j . By symmetry, it is clear that a signal exactly in the middle of to basis frequencies, ie. $\Delta = \frac{1}{2}$, contributes equally to both. This midpoint marks the boundary between one frequency bin and the next. Evaluating Eq. 4-46 at the bin edge yields a contribution at each of

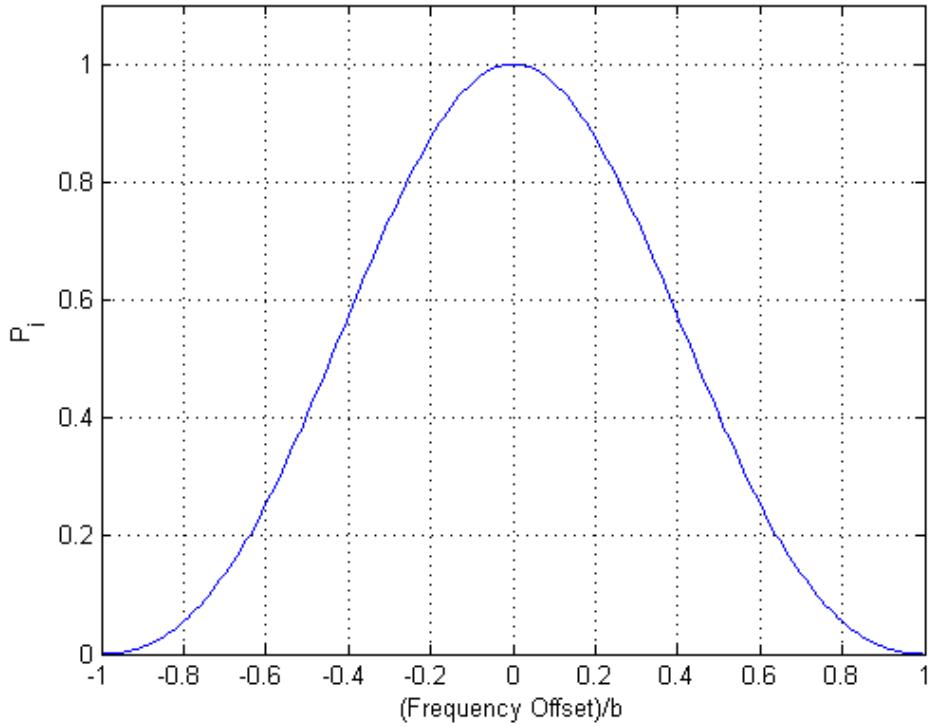


Figure 4-14. Power contributed to f_i from a signal which is not at a basis frequency.

the neighboring frequencies of 40.5% of the actual signal power. The remaining 19% is spread across all of the remaining basis frequencies.

Because an axion signal has equal probability to be at any frequency in the cavity bandwidth, it is expected that such a signal will not be exactly at a basis frequency, leading to some unavoidable power loss. This same argument applies equally well to the probability of having maximal power loss, which occurs when the signal sits at the edge of a frequency bin. As was done with \bar{L} , a more realistic measure of this loss is to use the mean value of Eq. 4-46 within the bounds of one frequency bin. This mean is given by

$$M = \frac{1}{b} \int_{-b/2}^{b/2} \left(\frac{\sin\left(\frac{\pi x}{b}\right)}{\frac{\pi x}{b}} \right)^2 dx = 0.774. \quad (4-46)$$

Conveniently, b drops out of the expression making this loss constant for all resolutions.

Incorporating these losses into Eq. 4–1 and accounting for power increases associated with the recombination of subspectra yields a power relation of

$$\frac{\bar{L}Mn}{2}g_{a\gamma\gamma}^2\frac{VB_0^2\rho_aCQ}{m_a} = T_{\text{eff}}\sigma_n. \quad (4-47)$$

Solving Eq. 4–47 for ρ gives

$$\rho_{\text{limit}} = T_{\text{eff}}\sigma_n \times \left(\frac{m_a}{g_{a\gamma\gamma}^2 VB_0^2 CQ} \right) \left(\frac{2}{0.774 \bar{L}n} \right), \quad (4-48)$$

as an expression for the upper limit on the local density of non-virialized axions as a function of m_a . Propagating the uncertainty in power for the recombined spectra through this calculation gives

$$\delta = \sigma_n \times \left(\frac{m_a}{g_{a\gamma\gamma}^2 VB_0^2 CQ} \right) \left(\frac{2}{0.774 \bar{L}\sqrt{n}} \right) \quad (4-49)$$

as the relative error in this limit. The loss factors are seen here to effectively raise the both the exclusion limit and its relative error.

In addition to m_a , Eq. 4–48 also shows dependance on $g_{a\gamma\gamma}$, spectral resolution via T_{eff} , and various experimental parameters which differ between spectra. Note that the n in the denominator of ρ_{limit} is canceled by the factor of n carried by σ_n , thus neither need be considered here. The first two variables are relatively easy to address. For axion model dependance, the limit for KSVZ axions is calculated first and then scaled by the ratio

$$\left(\frac{g_{\text{KSVZ}}}{g_{\text{DFSZ}}} \right)^2 = 6.87, \quad (4-50)$$

to produce the limit for DFSZ axions. A similar scaling is performed for the dependance on T_{eff} . Limits for the 84 mHz resolution are calculated which are then scaled by

$$\text{scale factor} = \left(\frac{T_{\text{eff}, n}}{T_2} \right) \quad (4-51)$$

to produce limits for the other resolutions. Finally, each spectrum has different values of σ_n , B_0 , C , Q , and \bar{L} . Substitution of these values into Eqs. 4–48 and 4–49 for a given

spectrum, yields a density limit and relative error specifically for that spectrum. This is then repeated for all spectra.

Limits over the frequency range from 812 to 852 MHz and 858 to 892 MHz are obtained by taking a weighted average of the limits for individual spectra. The total range is first subdivided into 250 kHz sections. Each spectrum is then associated with each section for which it has a frequency overlap. All spectra for a given section contribute to the final limit according to

$$\rho_{final} = \sum_i \frac{\delta_m^2 \rho_i}{\delta_i^2}, \quad (4-52)$$

where δ_m is defined by

$$\frac{1}{\delta_m^2} \equiv \sum_i \frac{1}{\delta_i^2}. \quad (4-53)$$

The effect of this weighting is that measurements with large errors, perhaps due to a poor Q or an abnormally high noise temperature, contribute less to the final limit than do measurements having a high Q and/or low noise temperature. These limits are then scaled by Eqs. 4-50 and 4-51 to obtain the limits for both axion models at each resolution. These limits are shown in Fig. 4-15. The limits displayed in the coverage gap are from a previous analysis which made use the power spectrum files saved prior to March 31, 2009[17].

4.10 Discussion

Over the range 812 to 892 MHz, or a mass range of 3.36 to 3.69 μeV , local axion flows are constrained at a CL of 90%. The limits for all frequencies above 860 MHz ($m_a \gtrsim 3.55 \mu\text{eV}$) as well as those near 835 MHz and 850 MHz are markedly better due to having higher Q values despite also having a higher system noise temperature from not using the MSA. In this range, densities above $\rho \simeq 0.1 \text{ GeV/cm}^3$ are excluded for KSVZ axions with $\delta v \lesssim 15 \text{ m/s}$. This limit becomes $\rho \simeq 0.7 \text{ GeV/cm}^3$ for DFSZ axions. Going to coarser resolutions weakens the limit relative to the 84 mHz resolution data, though they are applicable to a larger range of velocity dispersions.

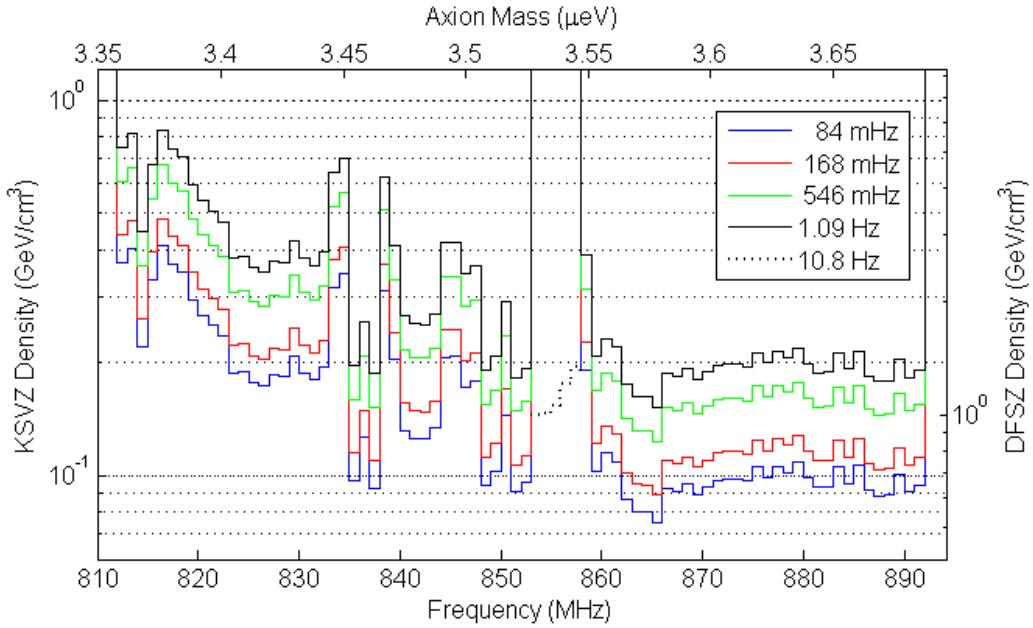


Figure 4-15. Exclusion limits on local density of non-virialized axions. The right axis is scaled to show density limits for DFSZ axions. All densities above the lines are excluded at 90% CL.

The maximum velocity dispersion for each resolution is a function of n , frequency, and flow velocity relative to earth. These maxima are calculated in the same manner as Eqs. 4-12 4-13 using $f = 850$ MHz and $v = 300$ km/s. Table 4-1 lists the maximum dispersion for each resolution as well as other information pertinent to interpreting the limits. Axion flows with this dispersion or less would deposit all of the power from an axion conversion in one frequency bin. Higher dispersion flows would be expected to spread their deposited power over multiple neighboring bins. These values should be kept in mind when applying the calculated limits to late-infall axions, the inner caustic of the big flow, and tidal flows.

Based on the predictions of Sikivie, Tkachev, and Wang, the densities of the first 8 late-infall axion flows are expected to be $\rho \sim 0.01$ GeV/cm³ [64]. Because this value is lower than the most stringent limits placed by this analysis, no signal from these flows is expected to be seen in the data set. Such flows for both KSVZ and DFSZ axions within the search range can not be excluded at this time.

Table 4-1. A Summary of the resolutions examined by the HR channel. These include spectral resolutions, resolution reduction factor, effective thresholds, scale factor, and maximum associated velocity dispersions.

Resolution (mHz)	δv (m/s)	n	T_{eff} (σ)	Scale Factor
84	15	2	15.5	1.00
168	30	4	18.3	1.18
546	86	13	25.4	1.64
1092	192	26	31.3	2.02

The Big Flow is currently estimated to have a density of $\sim 0.9 \text{ GeV/cm}^3$ in the neighborhood of its inner caustic and a velocity dispersion of $\delta v \simeq 50 \text{ m/s}$ [63]. The finest resolution that would be sensitive to the Big Flow is 546 mHz. For KSVZ axions, the 546 mHz limits are well under 0.9 GeV/cm^2 , excluding the big flow over the entire search range. However for DFSZ axions, the 546 mHz limits are not low enough to exclude the Big Flow at any frequency.

Estimates on the presence of tidal flows in the solar neighborhood are rather limited. Freese et al. suggest that the sun may currently be passing through a tidal flow originating from the Sagittarius dwarf galaxy. The velocity of the flow and its dispersion are estimated to be $v \simeq 300 \text{ km/s}$ and $\delta v \simeq 20 \text{ km/s}$ [81, 82]. Having a dispersion larger than the maximum value associated with the 1.09 Hz resolution limits, this specific flow is beyond the scope of this analysis. Note however that this does not preclude all tidal flows from being considered. Using the 1.09 Hz data, limits can be placed on tidal flows having $\delta v \lesssim 192 \text{ m/s}$. Flows of KSVZ axions with densities greater than 0.8, 0.4 and 0.2 GeV/cm^3 are excluded in the frequency ranges of 812-822 MHz, 822-852 MHz, and 859-892 MHz respectively.

CHAPTER 5 FUTURE WORK

5.1 Introduction

As with most research projects, one set of results tends to just be the beginning of something more. The exclusion limits set in Chapter 4 on the local density of non-virialized dark matter are just such a beginning. The inner caustic of the Big Flow is ruled out for KSVZ axions but its existence for DFSZ axions is left undetermined. Further, the existence of late-infall axions remains undetermined for both KSVZ and DFSZ axions. Resolution of these undetermined cases will only be achieved by lowering the density limits. Constrained by the current HR data set, this can only be done by reducing the candidate threshold. However, if at some time in the future further measurements are to be taken, there exists the option of fine tuning experimental parameters through hardware upgrades.

In addition to pushing for greater sensitivity in the HR channel, what can be done to expand these limits to higher frequencies? At the conclusion of the science run ending in April 2010, ADMX had reached the top end of the tuning range permitted by its current tuning rods. While larger tuning rods could still be used to push this top end a little higher, a more significant widening of the frequency range must come from elsewhere. This widening could be from coupling to a different mode of the current cavity, or even from redesigning the cavity itself. Clearly, these limits are just an invitation to investigate further. This chapter is intended to expand upon the above topics and to provide information about the present state of these new research avenues.

5.2 Threshold Selection

From Eq. 4–48 it can easily be seen that the limit is directly improved by lowering T_n , but doing so has already been noted to greatly increase the computational time required for comparing candidates. To mitigate this rather sizable complication, a more efficient means of searching must be written and employed. In designing a new search

algorithm is prudent to examine the shortcomings of the present one. In this case, the major problem is performing unnecessary computations.

The present algorithm was written to make use of the ease with which MatLab handles vector algebra. Typically, it is faster in MatLab to write a function which acts on vectors as a whole rather than iterating over each element of each vector. To this end, the frequency of the candidate C_0 is compared to a vector, \mathbf{F} , defined as

$$\mathbf{F}_i = f(C_i), \quad (5-1)$$

where $f(C_i)$ is the frequency of the i^{th} candidate. A vector, \mathbf{T} , containing the time stamps of each candidate is constructed in the same way. When comparing \mathbf{F} and \mathbf{T} to C_0 , a logical filter is applied which limits the indices used to those for nearby candidates only. The process of determining δt_i , δf_i and $\delta f_{i, \text{max}}$ is essentially done in parallel for all C_i . This is likely faster than iterating over all indices involved, but the truth is most of these computations need not be done in the first place. Once a candidate fails to appear in half of the overlapping spectra there is no need to compare it to candidates in the remaining spectra. This requires that the computations be run in series rather than in parallel.

Consider now the number of candidates present in each spectrum. Lowering T_n causes this number to grow exponentially. Thus comparing C_0 to each of the candidates found in a given spectrum also grows exponentially. However, computational time can be saved here as it is not always necessary to compare C_0 to all candidates in a given spectrum. Once a single candidate from an overlapping spectrum shows coincidence, comparison to the remaining candidates in that spectrum can be skipped and the next spectrum should be examined. This too must be a series calculation rather than parallel.

These two changes would likely be implemented as nested FOR loops containing BREAK statements for the above conditions. The outer loop would iterate over the overlapping spectra while the inner loop would iterate over candidates in those spectra.

Note that these conditions attempt to reduce the computational time for opposite circumstances. The inner loop is terminated when coincidence is present, while the outer loop is terminated when coincidence is absent. In the latter case, C_0 must still be compared to all candidates in a given spectrum in order to claim that coincidence is absent. Keeping this comparison as a vector calculation would definitely be faster when there is no coincidence, but there would be no option to break the inner loop if there is coincidence. This prompts a very important question. Which is greater, time saved by the inner loop when there is coincidence or the time lost when there is no coincidence?

One consideration which would favor keeping the inner loop would be to reduce the number of spectra which are allowed to contain no coincident candidates. Making this some low, constant value would minimize the number of times that the inner loop is forced to iterate over all candidates in a spectrum. On the other hand, this option has the side effect of somewhat increasing the density limits. The coincidence threshold is set by Eq. 4-43; thus, making it more strict would invalidate using \bar{L} . For a cutoff threshold which permits having only, say, two spectra which fail to show coincidence, the value used for L would end up being lowered to effectively $L = \frac{1}{2}$, down from $L \geq \frac{\pi}{4}$. This change would require a decrease in T_n of at least 33% just to maintain the present limits. As of this writing, there is no definitive answer on how to properly optimize the search algorithm for an exponentially growing set of candidates.

5.3 Hardware Upgrades

ADMX stopped data taking operations in April 2010 in order to relocate the experiment from Lawrence Livermore National Laboratory to the University of Washington in Seattle. In preparation for shipping, the experiment had to be completely disassembled, and upon arriving in Seattle work began on reconstruction and hardware upgrades. These upgrades will boost the sensitivity of future searches by improving upon two important experimental parameters; the cavity Q and the rms noise power.

As discussed in Chapter 4, maintaining high Q values is of great importance. Looking at Fig. 4-15, it is easy to see that the limits for data taken after the inclusion of indium wire at the knife edge joint (860-892 MHz) are much better than those for data taken prior. This clearly shows the destructive power of a low Q . The lack of a nickel strike beneath the copper plating on the old cavity resulted in poor adhesion between the copper plating and the stainless steel body of the cavity. This caused bubbling and flaking in the copper, which led to poor electrical connections at the end caps, and ultimately yielded low values for Q . To fix the low values of Q which plagued much of the last science run, a new cavity has been fabricated and plated. Including a nickel strike as an adhesion layer prior to the actual copper plating of this new cavity is critical to maintaining high values of Q . With this in mind, immense care has been taken to ensure that copper plating of the new cavity is handled properly.

Improvements to the rms noise power will come from directly reducing the system noise temperature. Recall from Chapter 3 that the noise temperature of the MSA is roughly half of its physical temperature and that the noise temperature of subsequent amplifiers is suppressed. When including the noise contribution of HFET-1, again assuming a gain of 10 dB for the MSA, the system noise temperature can be written as

$$T_{N,System} = T_{Physical} + T_{N,SQUID} + \frac{T_{N,HFET-1}}{G_{SQUID}} = \frac{3 T_{Physical}}{2} + 0.4 \text{ K.} \quad (5-2)$$

Equation 5-2 is valid down to the quantum limit for the MSA, which occurs at a physical temperature of about 100 mK. At 100 mK, the system temperature would be roughly 0.55 K, which is a factor of 6 improvement.

Attaining a physical temperature of 100 mK will require the installation of a dilution refrigerator. To accommodate this addition, a redesigned insert was built from the ground up which includes empty space that has been allocated for the future installation of the refrigeration system. Due to the one year lead time on the dilution refrigerator, the new insert design also includes a pumped ^3He system which will reduce the physical

temperature to about 300 mK. This will allow ADMX to resume data taking operations prior to the delivery of the dilution refrigerator while still improving upon the system noise. At a physical temperature of 300 mK, the system noise temperature is expected to be about 0.85 K, which is still a factor of 4 improvement.

5.4 Higher Frequencies

A number of options have been put forward on how to expand the frequency range of ADMX beyond its current bounds. The three most straightforward of these options are installing larger diameter tuning rods, coupling to a higher order resonant mode, and using a smaller diameter cavity. Implementation of each each option would be relatively simple, but that is not to say that they are without their drawbacks. All three methods increase the resonant frequency of the cavity, but at the same time, all three decrease P_a which weakens the final limits.

The option employed in the past has been that of using larger diameter tuning rods. The use of metal rods forces the tangential component of electric field to go to zero at the surface of the rod, weakening the average field strength as well as effectively excising the rod volume. Thus, using larger tuning rods further reduces the effective volume of the cavity which yields a slightly higher resonant frequency. The decrease in P_a due to the smaller volume is relatively small, being a loss of about 2.8% for the presently used 5 cm diameter rods. An increase to 6 cm or 7 cm diameter rods would only increase this loss to about 4% or 5.5% respectively. The major decrease in P_a comes from the form factor C .

Taking $\mathbf{B} = B_0 \hat{z}$, Eq. A-2 reduces to

$$C = \frac{\left| \int_V \mathbf{E}_z d^3x \right|^2}{V \int_V |\mathbf{E}|^2 d^3x}. \quad (5-3)$$

For an empty right circular cylinder of radius R , the TM_{010} mode has an electric field of

$$\mathbf{E}_{010} = E_0 J_0 \left(\frac{\chi_{01} \rho}{R} \right) \hat{z}, \quad (5-4)$$

where E_0 is the amplitude of the field strength, J_0 is the Bessel function of order 0, and $\chi_{01} = 2.405$ is the first root of J_0 [68]. While perturbations to this field caused by the presence of tuning rods may introduce components in the $\rho\phi$ -plane, the dominant portion of \mathbf{E} remains tangential to the surface of the rods. Increasing the radius of the rods, increases their vertical surface area, and with it the effective volume of the low field region of the cavity. For prospective, consider that an empty cavity would have a form factor of $C_{010} = 0.704$, and that the form factor for ADMX is already down to a value between 0.41 and 0.61. Regarding the use of larger, and larger tuning rods, ADMX may now be at a point where diminishing returns makes this less viable.

Coupling to higher order modes is an appealing option for two reasons. The first reason is that it greatly expands the searchable frequency range. Going back to an empty cavity, the base frequency of the TM_{mn0} mode is given by

$$f_{0n0} = \frac{c}{2\pi} \frac{\chi_{mn}}{R}. \quad (5-5)$$

For the ADMX cavity ($R = 0.2$ meters), the first two TM_{0n0} modes would have base frequencies of $f_{010} \approx 0.57$ GHz and $f_{020} = 1.31$ GHz. Expansion of the frequency range about the second order base frequency would then be accomplished the same way as before using metal and dielectric tuning rods. The second reason is that with only a few additions to the existing hardware, both modes can be measured simultaneously. This will allow ADMX to lower existing limits, thanks to the cryogenic hardware upgrades, while also setting limits for frequencies which have yet to be covered.

For this option, the loss in P_a once again stems from C . Evaluating Eq. 5-3 for

$$\mathbf{E}_{020} = E_0 J_0 \left(\frac{\chi_{02}\rho}{R} \right) \hat{z}, \quad (5-6)$$

where $\chi_{02} = 5.520$ is the second root of J_0 , yields a value of $C_{020} = 0.131$. This decrease in C by roughly a factor of 5 would seem to be as appalling as the increase to the frequency range is impressive. Fortunately, this loss is somewhat mitigated due to the

higher resonant frequency. Recall from Eq. 4–2 that

$$P_a \sim C_{0n0} m_a. \quad (5-7)$$

Doubling the frequency would also double m_a , resulting in a net decrease of P_a by a factor of only about 2.5. This option will be pursued by ADMX for future science runs. Coupled with a lower noise temperature, density limits at higher frequencies are expected to be at least comparable to those set for the present search range.

The last option is to simply use a smaller cavity. As seen in Eq. 5–5, the base frequency of a right circular cylinder is inversely proportional to its radius. So, halving the radius would double the frequency. In order to prevent more TE modes from bunching up near the TM_{010} mode, the aspect ratio of the current cavity should be maintained. Thus the height should be halved as well. At this point, the loss in P_a is painfully obvious. The volume of the cavity has just been reduced by a factor of 8! For a single cavity the loss in SNR is identical to the loss in P_a . In order for this option to remain viable, the SNR lost by shrinking the cavity must somehow be reclaimed.

Being half as wide two cavities could now fit in the horizontal space occupied by the old cavity. Being half as tall, the same is true for the vertical space. This four-cavity array could be used to regain some of the lost SNR. For simplicity, assume that the spectral resolution is coarse enough to permit a combined spectrum analysis, and that each cavity measures roughly the same signal power and has roughly the same noise power. Adding the power spectra from all four cavities would then recover a factor of 4 in signal power but only a factor of 2 in SNR. What if the output signals were added in phase prior to applying the FFT? In that case, the signal amplitude in the time domain would go up by a factor of 4 while the noise amplitude would only go up by a factor of 2. Converting to a power spectrum would yield a factor of 4 increase in SNR compared to that of a single small cavity. In general, N cavities operated in this way would have an SNR that is N times better than a single cavity. A multi-cavity array similar to the one

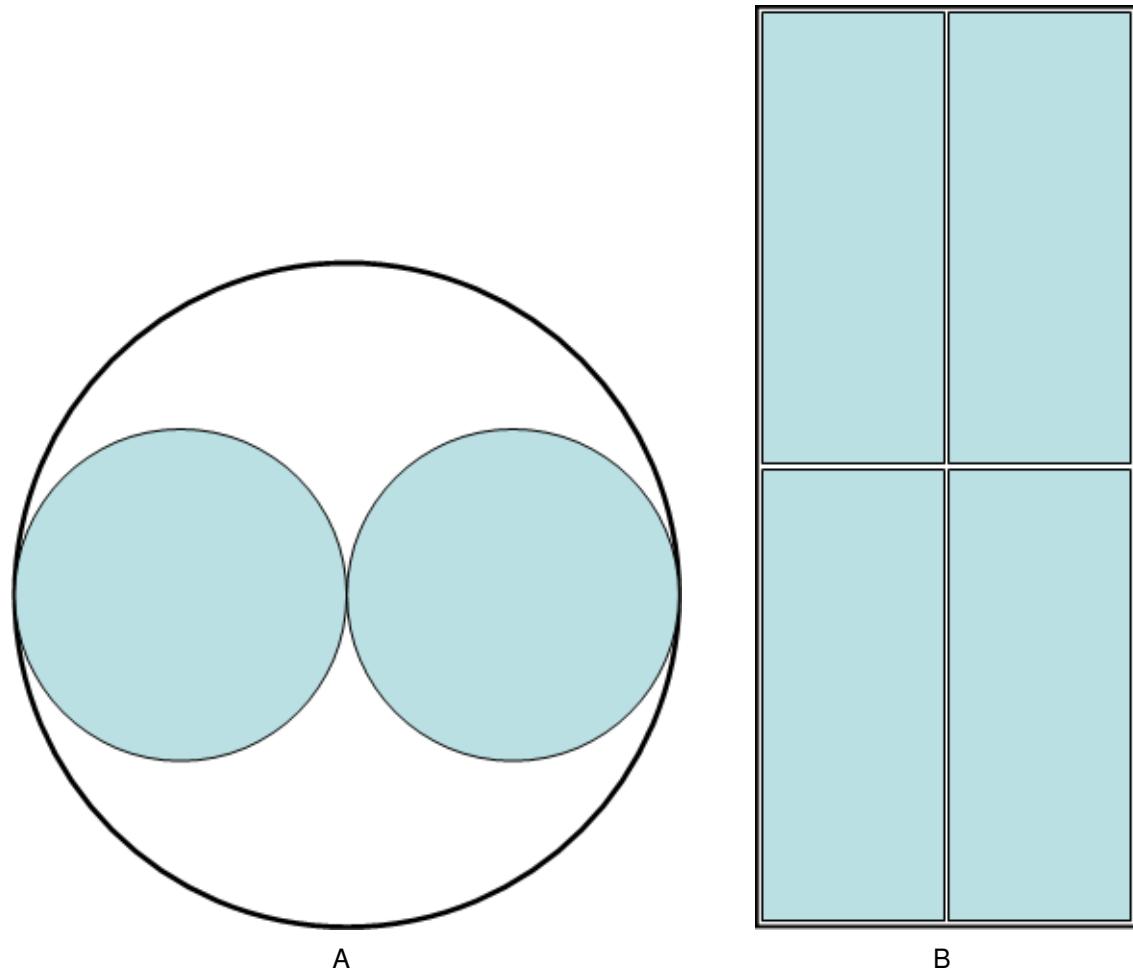


Figure 5-1. Close packing of 4 cavities within the old cavity volume. The smaller cavities have had their radii reduced by 50%. A) Top-down view showing empty space around the cavities. B) Side view showing full use of vertical space.

just described was constructed and successfully operated in the late 1990s as a proof of concept[83].

In the above example there is still a net loss of a factor of 2 in SNR. This is because half of the otherwise usable volume simply will not accommodate the addition of another cavity. This is easily seen in Fig. 5-1. Even fully utilizing the vertical space, there is still a 50% reduction in the space used. Reducing the size of the cavities a little bit more permits the use of a 6-cavity array. This array, diagrammed in Fig. 5-2, decreases the amount of empty distributed horizontally but increases the amount distributed vertically.

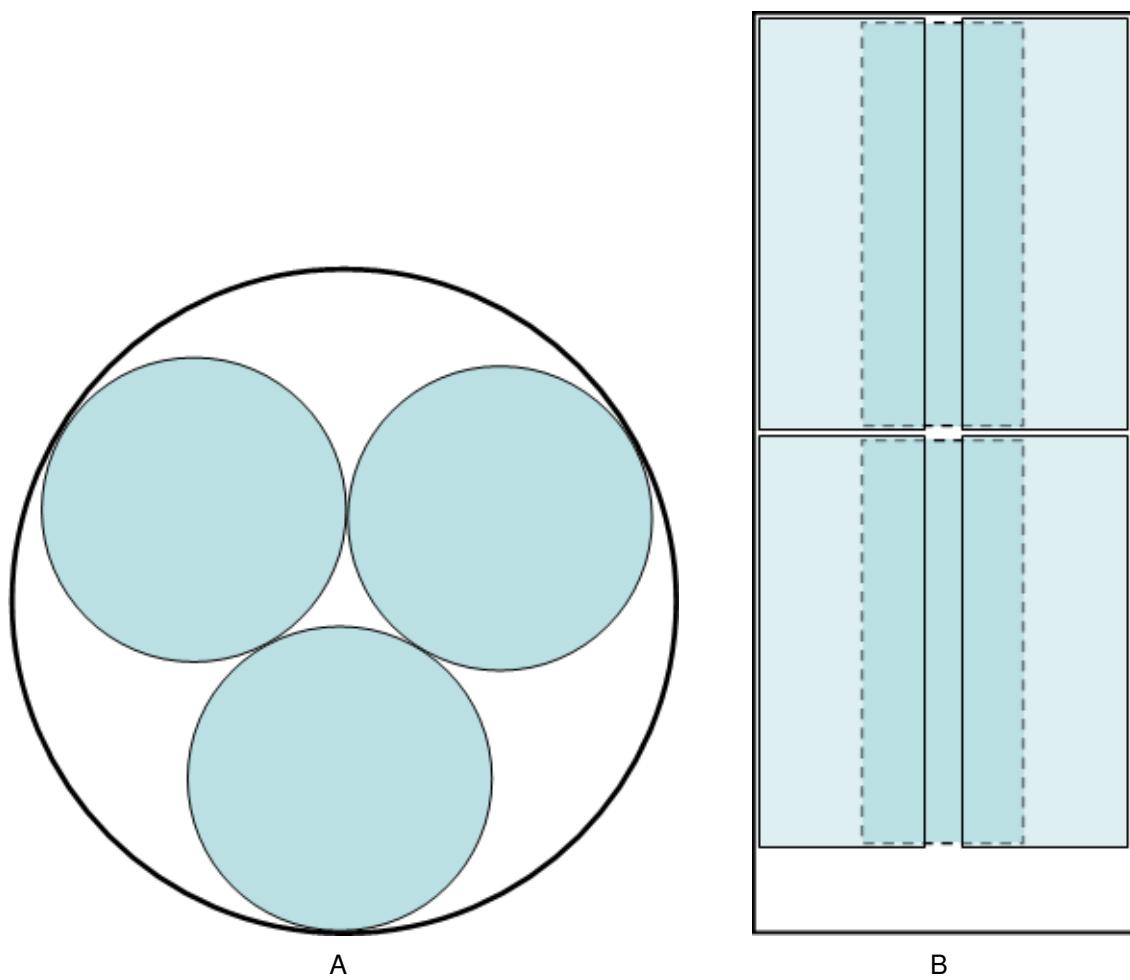


Figure 5-2. Close packing of 6 cavities within the old cavity volume. The smaller cavities have had their radii reduced by 54.2%. A) Top-down view showing less empty space around the cavities than for the 4-cavity configuration. B) Side view showing only partial use of vertical space.

The overall loss of volume for this configuration is 42.2%. As a final example of this trend, consider a 21-cavity array, seen in Fig. 5-3, where the cavities' radii have been reduced by 66%. Horizontal space is used even more efficiently than for the other arrays, and vertical space is once again fully utilized. For this configuration, there is still a loss of 22.2% of the useable volume. If the increase in m_a due to being at a higher frequency is included, this configuration would actually see a gain in SNR by factor of about 2.4. The deterrence to using such an array is in the logistics. While a 4-cavity

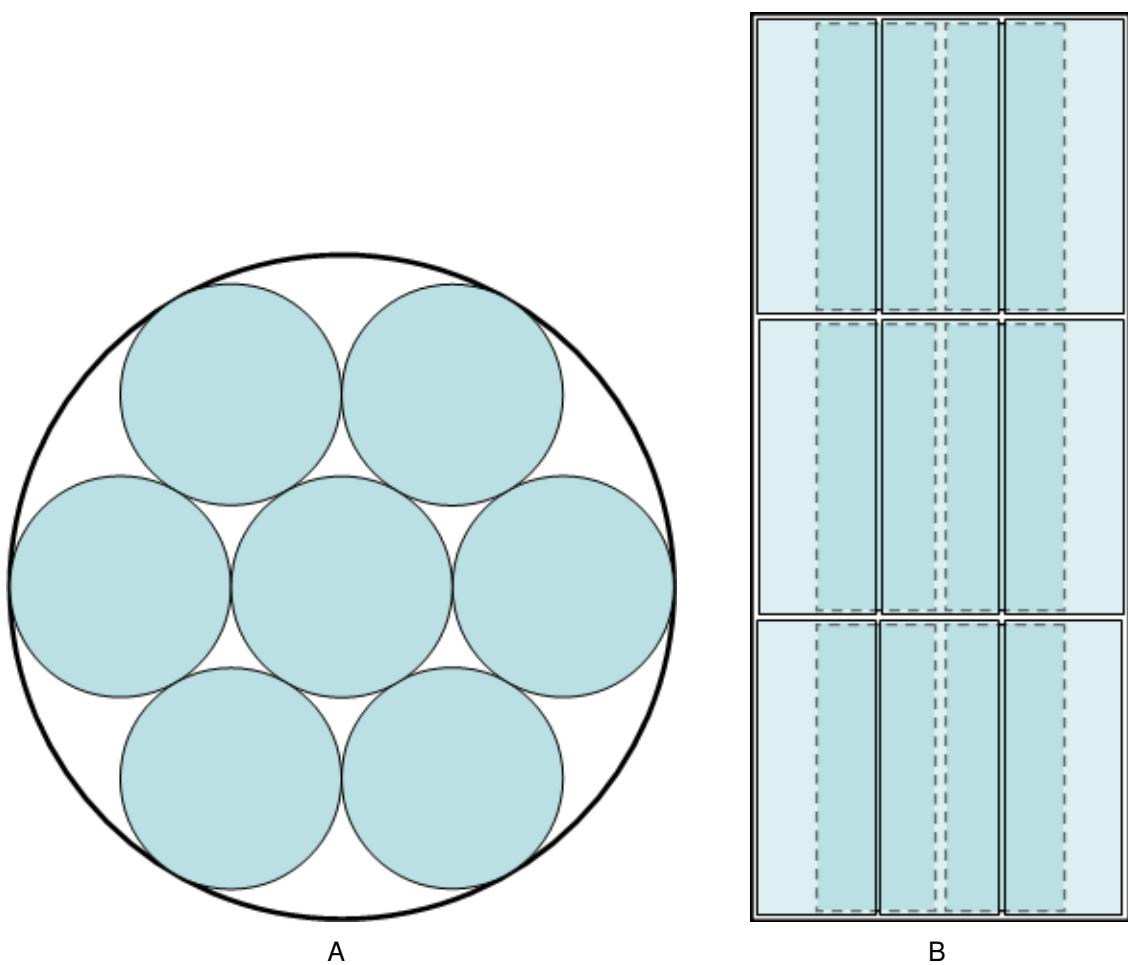


Figure 5-3. Close packing of 21 cavities within the old cavity volume. The smaller cavities have had their radii reduced by 66.6%. A) Top-down view showing less empty space around the cavities than for either the 4-cavity or 6-cavity configurations. B) Side view once again showing full use of vertical space.

array was managed relatively easily, a 21-cavity array would likely be overwhelming due to needing 21 independent tuning mechanisms.

Despite the complications of running a system with multiple cavities, the appropriate combination of cavity outputs causes the loss in SNR to depend less on the number of cavities and more on the efficiency of the volume used. The ideal case would be a cavity array which uses nearly all of the available space. Incorporating the idea of using a higher order mode may do just that. If the cavity were split into two D-shaped cells by affixing a vertical partition spanning the diameter of the cavity, then the two cells could

be treated as separate cavities and the loss in volume would just be governed by the thickness of the partition. The boundary condition required by the partition is that $\mathbf{E}_z = 0$ along the diameter. This kills the TM_{010} mode making the new lowest mode TM_{110} . For an empty cell of radius R , the electric field of the TM_{110} mode is

$$\mathbf{E}_{110} = E_0 J_1 \left(\frac{\chi_{11} \rho}{R} \right) \sin(\phi) \hat{z}, \quad (5-8)$$

where J_1 is the Bessel function of order 1 and $\chi_{11} = 3.832$ is its first root. In the full circular cavity, this mode would still exist but the form factor would be zero due to the integral over ϕ . For D-shaped cells, this integration would only span half of one period yielding non-zero form factors for both cells. Using Eq. 5-3, the form factor for an empty cell is found to be $C_{110D} = 0.644$.

A prototype cavity of this design has existed at the University of Florida for some time, though little work had been done on it in the past. The cavity body is made of solid copper, is 14 inches tall, and has an inner diameter of 5.375 inches. Brass flanges which are used to bolt on the end caps are soldered to the outside of the cavity, flush with each end of the cylinder. Both end caps and the 0.125 inch thick rectangular plate which acts as the partition are also made of solid copper. The partition is held in place by screws running along the full length of the cavity. A line of solder runs along each vertical edge of the partition to electrically connect the cylinder and the center plate. The solder at these joints provides negligible mechanical support to center plate. When work resumed on this R&D project, the cavity was found to have a dingy matte coloring caused by years of oxidation, as well as numerous defects in the surface finish of both the walls and the partition. A picture of the prototype cavity in this state can be seen in Fig. 5-4.

At room temperature, the unloaded Q of one cell in this cavity is expected to be $Q \sim 25,000$. Measurements at that time of the unloaded Q were an order of magnitude lower than expected. To rectify this problem, all parts of the cavity were first mechanically polished to remove any oxidation, and second electro-polished to remove

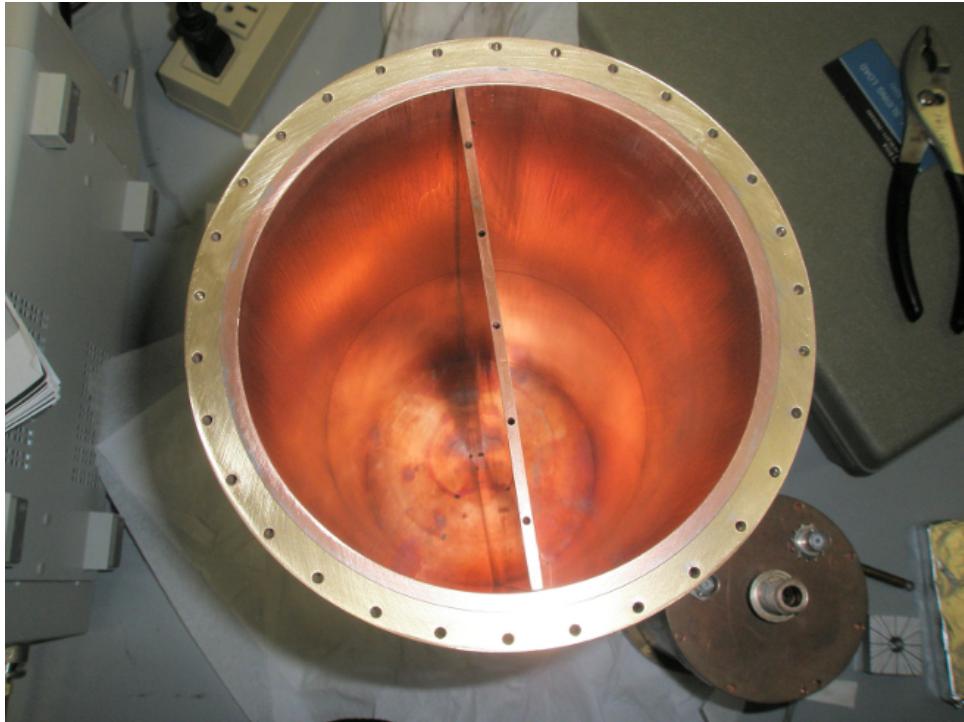


Figure 5-4. A picture of the split cavity prototype prior to being electro-polished.

surface defects. Once polished all parts were, the partition was screwed in place and once again soldered to the cavity wall. Figure 5-5 shows the cavity following this step. Measurements of the Q were taken at room temperature for each side of the now clean cavity. The Q for the cell which contained the solder joints increased to a value of about 22,000, while the Q for the other cell showed no signs of improvement. The vertical joints in the second cell were then soldered which brought that side's Q up to about 22,000 as well. At this time, power leakage between the two cells was tested by injecting power into one cell and measuring the transmitted power in the other. The cross-talk between the two cells was found to be negligible, having an attenuation of at least 60 dB.

The Woods metal solder used at these joints did not bond very well with the copper, leading to an eventual break in the electrical connection at these points. Both cells maintained high values of Q up to the point at which the solder joints failed. It was concluded that even though the surface currents which support the TM_{110} mode run only

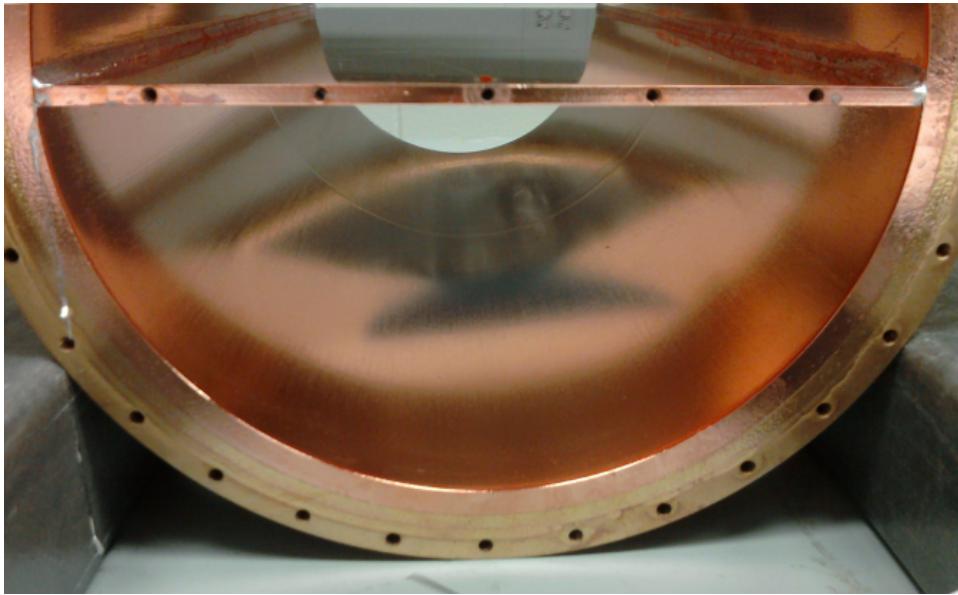


Figure 5-5. A picture of the split cavity prototype after being electro-polished. At this stage the solder joint between the center plate and the wall was made for only one of the cells (top cell as pictured) in the cavity.

vertically along the cavity walls and the partition, the electrical connection across these vertical joints is still quite important. All solder was removed from the inside of the cavity and the requisite electrical connections were made using silver paint instead. The first two applications of silver paint wicked in between the edge of the partition and cavity walls. A third coat was then applied, and another set of Q measurements were made. Both cells once again had Q 's around 22,000. Shortly after this point, work on high frequency cavity designs shifted to fellow ADMX collaborator Ian Stern, who has expanded this study to include alternate tuning mechanisms for circular cavities.

The prototype cavity has demonstrated that a split cavity design is relatively easy to construct, and that both cells can achieve room temperature values of Q of order 10^4 with negligible cross-talk. Further work on the prototype cavity would likely include testing the tunable frequency range of the cells, as well as the ease with which they can be frequency matched. Insight into optimizing the former may come from the work pursuant to alternate tuning mechanisms. Fabrication of a mechanism for tuning the cavity using thin metal vanes instead of rods is expected to begin in 2014.

CHAPTER 6 SUMMARY AND CONCLUSIONS

Dark matter continues to draw the interest of physicists and astronomers, more so now than ever given that experimental results are capable of placing meaningful exclusion limits on plausible dark matter candidates. The axion, which stems from the Peccei-Quinn solution to the strong CP problem, is one such plausible candidate. Within a mass range of $1 \mu\text{eV} \lesssim m_a \lesssim 5 \text{ meV}$, the axion would be produced copiously in the early universe, interact primarily through gravity, and have a lifetime sufficient to still be around today. Through gravitational interactions, many axions in the Milky Way would eventually reach a state of dynamic equilibrium and would be referred to as being virialized. Axions which have only recently entered the galaxy as well as those affected by tidal disruptions will have had insufficient time to reach dynamic equilibrium and would be referred to as non-virialized. These axions would move in discrete flows characterized by low velocity dispersions. Given the right conditions, these flows may even form caustics leading to localized regions of higher than average density.

Detection of dark matter axions is facilitated by the use of an axion haloscope. The interaction between axions and a strong magnetic field inside the cavity of the haloscope will lead to a resonant conversion of axions into photons when the resonant frequency of the cavity matches that of the axions. When read out by a low-noise radio receiver, these conversion photons would be seen as a spike in the power spectrum of the cavity. The amplitude of the spike is proportional to the number of converted axions, while the width is proportional to their velocity dispersion. Thermal noise in the power spectrum, which serves as the background to this signal, is optimized when the spectral resolution matches the signal width. Virialized axions would have a broad Maxwellian velocity distribution lending themselves to a search with a resolution of $\mathcal{O}(100\text{Hz})$. On the other hand, non-virialized axions would be more suited to a search using much finer resolutions.

A complication arises when the resolution is $\mathcal{O}(\text{Hz})$ or finer. At these resolutions, signal modulation due to terrestrial motion can cause an axion signal to drift in frequency from one measurement to the next. This prevents the measurements from being averaged prior to searching for an axion signal unless assumptions are made regarding the directions and magnitudes of the flow velocities. No such assumptions were made in this analysis, thus the actual search for axion signals was performed for each spectrum independently of the others. All signals in each spectrum with power exceeding a candidate threshold were then compared to those in each other spectrum. Candidate signals in different spectra which are separated in frequency by no more than the maximum modulation attributable to terrestrial motion are considered to be coincident. That is to say, they could potentially be the same signal. Repeatedly failing to show coincidence with other nearby signals eliminates a candidate as potential axion signal.

Uncertainties in the phase space structure of axion flows, favor the use of multiple searches at different resolutions. The selection of the spectral resolution was performed in the time domain, by adjusting the length of the FFT, rather than in the frequency domain, whereby neighboring bins in the power spectra are summed. While both methods yield the same signal to noise ratio for coherent signals, they differ when it comes to decoherent signals. For such signals, resolution selection in the time domain is preferred. Because resolution selection in the frequency domain is much less computationally intensive, it is recommended that future analyses use the frequency domain to set their resolutions if at all possible (ie. hardware permitting). Each resolution corresponds to a different maximum velocity dispersion for which the signal to noise ratio is maximized. Having searched the frequency range of 812 to 892 MHz, or $3.36 \mu\text{eV} \lesssim m_a \lesssim 3.69 \mu\text{eV}$, at resolutions of 84 mHz, 168 mHz, 546 mHz, and 1.08 Hz, no candidates have shown significant coincidence. Based on this, limits have been placed on the local density of non-virialized axions. The 84 mHz limits are at

best a factor of 2 lower than the limits generated from a modulation insensitive search performed at a resolution of 10 Hz and at worst a factor of 2 higher.

Further improvements to the sensitivity can be achieved by lowering either the candidate threshold or the system noise. The former option would require considerable effort in optimizing the search algorithm, but once this is done the algorithm would remain viable for future searches. The latter option is presently being pursued in the form of cryogenic hardware upgrades. Data taking operations for ADMX are scheduled to resume in 2014 with a base temperature of 300 mK following the installation of a pumped ^3He system.

APPENDIX A A SEARCH FOR NON-VIRIALIZED AXIONIC DARK MATTER

The Cold Dark Matter (CDM) component of the Milky Way halo may consist of low mass ($\mu\text{eV-meV}$) pseudoscalar particles called axions, which are a consequence of the Peccei-Quinn solution to the strong CP problem[38, 40, 41]. Axions of this mass range would have extremely weak couplings to standard-model particles as well as to both axionic and non-axionic dark matter, rendering them effectively collisionless[8–10]. Models of the structure of halos consisting of such particles predict discrete flows with low velocity dispersion and high density at special locations called caustics[55, 63, 84]. In a Sikivie type detector[66, 67], a low dispersion axion flow would appear as a narrow peak in the power spectrum. The Axion Dark Matter eXperiment (ADMX) high resolution (HR) data acquisition channel is sensitive to such signals, even more so than for signals expected from halo models lacking low dispersion flows[13]. Sensitivity to these signals is further improved by combining data from successive power spectra. This gain in the signal to noise ratio (SNR) relative to a single-scan analysis, such as in the previous ADMX HR analysis found in Ref. [13], allows for sensitivity to lower densities, though care must be taken to account for the change in frequency of the axion signal due to the Earth’s motion relative to the dark matter. Because any measurement of the axion’s kinetic energy will be subject to Doppler shifts, both spectral broadening and signal modulation are expected. In this paper we discuss the constraints on a combined-spectra analysis imposed by Doppler shifts, and present the results of the first ADMX combined-spectra HR analysis.

ADMX converts axions to detectable microwave photons via the inverse Primakoff effect within a tunable, high quality factor ($Q \simeq 50,000$) microwave cavity immersed in a strong magnetic field. An axion with mass m_a and velocity β would create a photon with energy $E \approx m_a c^2 + \frac{1}{2} m_a c^2 \beta^2$. A velocity dispersion $\Delta\beta$ yields spectral broadening

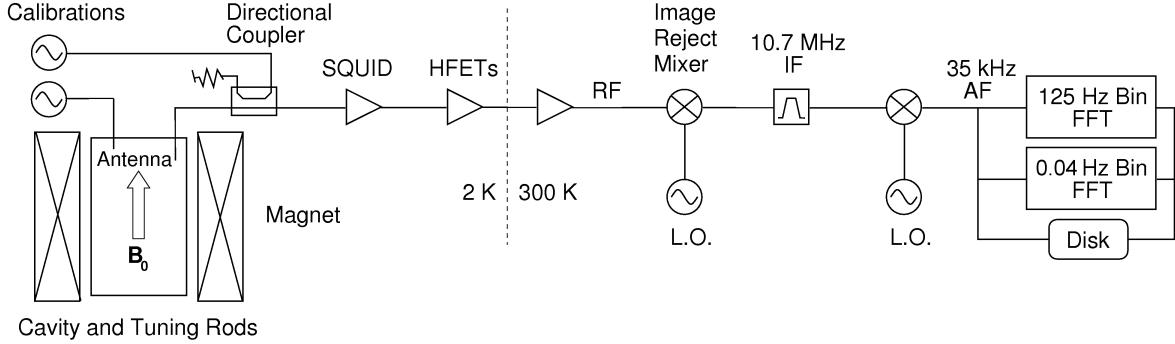


Figure A-1. Diagram of the ADMX cavity and receiver chain. The power read out from the cavity is sent through both cold and room temperature amplification stages, is mixed down from radio frequencies to audio frequencies via two local oscillators (L.O.), and is ultimately saved to disk.

of $\frac{\Delta f}{f} = 2 \frac{\Delta E}{E} \approx 2\beta\Delta\beta$. Axion conversion is expected to produce power in the cavity of [66, 67],

$$P = g_{a\gamma\gamma}^2 \frac{VB_0^2 \rho_a C}{m_a} \min(Q, Q_a), \quad (\text{A-1})$$

where V is the volume of the cavity, B_0 is the magnetic field strength, and ρ_a is the axion density in the detector. The mode dependent form factor of the cavity is given by

$$C = \frac{\left(\int_V \mathbf{E} \cdot \mathbf{B}_0 d^3x \right)^2}{VB_0^2 \int_V |\mathbf{E}|^2 d^3x}, \quad (\text{A-2})$$

where \mathbf{E} is the electric field of the tuned mode and \mathbf{B}_0 is the external magnetic field. The highest form factor obtainable ($C \approx 0.69$) corresponds to the TM_{010} mode[80] prompting its use over higher order modes. The quality factors of the cavity and the axions are $Q \lesssim 10^5$ and $Q_a \gtrsim 10^6$, respectively. Finally the axion to two photon coupling constant $g_{a\gamma\gamma}$ is given by

$$g_{a\gamma\gamma} = \frac{\alpha g_\gamma}{\pi f_a}, \quad (\text{A-3})$$

where α is the fine structure constant, f_a is the Peccei-Quinn symmetry breaking scale, and g_γ is a model dependent constant. In the Kim-Shifman-Vainshtein-Zakharov (KSVZ)

model $g_\gamma = -0.97$ [42, 43], and in the Dine-Fischler-Srednicki-Zhitnitskii (DFSZ) model $g_\gamma = 0.36$ [44, 45].

The ADMX detector consists of a 1-m tall, 0.5-m diameter, copper plated, stainless steel, right cylindrical cavity kept at 1.8 K and placed in a 7.6 T magnetic field. The TM_{010} mode of the cavity is swept through a range of frequencies, corresponding to a range of possible axion masses, by moving a metal rod, oriented parallel to the cylinder axis and extending along the full length of the cavity, from the wall to the center of the cavity space¹. For the experimental parameters of ADMX, Eq. A-1 predicts an expected signal power of order 10^{-22} W. The expected SNR is given by the Dicke radiometer equation [78]:

$$SNR = \frac{P}{P_N} \sqrt{bt} = \frac{P}{k_B T_N} \sqrt{\frac{t}{b}}, \quad (\text{A-4})$$

where P is the expected signal power, P_N is the noise power, b is the signal bandwidth, t is the integration time, k_B is Boltzmann's constant, and T_N is the total noise temperature which is equal to the physical temperature plus the noise temperature of the electronics.

ADMX has two data acquisition channels which use the receiver and amplifier chain shown in Fig. A-1. Power from the cavity is critically coupled to the receiver which sends the signal through a directional coupler followed by two cold ($\simeq 1.8$ K) amplification stages and then one room temperature amplification stage. The cold amplification stages are a DC Superconducting QUantum Interference Device (SQUID) amplifier[72] followed by two balanced GaAs heterostructure field-effect transistor (HFET) amplifiers. There is ~ 10 dB gain from the SQUID and ~ 34 dB combined gain from the HFETs[73–75]. At room temperature, the signal is amplified an additional 35 dB before being down-converted to a center frequency of 10.7 MHz and sent through

¹ There are two such rods for redundancy, and for tuning past mode degeneracies in the cavity.

an eight-pole crystal band pass filter. The signal is down-converted a second time to a center frequency of 35 kHz whereupon the medium resolution (MR) and HR channels diverge.

The MR channel takes 10,000 scans over 80 seconds, computes a fast Fourier transform (FFT) and power spectrum of each scan and averages the 10,000 spectra into one 400-point power spectrum with a resolution of 125 Hz. The HR channel takes data for 23.8 seconds with a sampling rate of 80 kHz resulting in a maximum resolution of 42 mHz. Three such spectra are taken during the integration time of the MR channel. The HR data are saved in two formats. First the raw time series data are written to disk with no signal processing or averaging having been applied. Analysis of these data will be the subject of a future publication. Second, the three scans taken during the MR integration are Fourier transformed, averaged, and saved as a single 42-mHz bin-width-power spectrum. After acquiring data at a given resonant frequency, f_0 , the tuning rods are moved, shifting $f_0 \sim 2$ kHz, and the data acquisition process is repeated. Because the tuning rod steps are much smaller than the 30-kHz bandwidth of the crystal filter and the resonant width of the cavity TM_{010} mode, the frequency coverage for successive scans overlap significantly. For both channels, the total exposure time for a given frequency bin is ~ 25 minutes.

The MR channel has recently produced limits for both virialized and non-virialized axions with $\Delta\beta \sim 10^{-3}$ and $\lesssim 2 \times 10^{-4}$, respectively[14]. While the MR and the HR channels both look for non-virialized axions, the difference comes in their sensitivity to models with differing velocity dispersions. The MR channel is equally sensitive to all signals with dispersions of $\lesssim 10^{-4}$, as all such signals would occupy a single 125 Hz bin. The HR channel potentially has a much smaller bin width, and therefore continues to gain in the SNR for signals with low velocity dispersion. Figure A-2 shows the expected signals in the HR channel for two simulated axion flows of equal density but differing velocity dispersion. While both signals would be seen in the HR channel, the signal

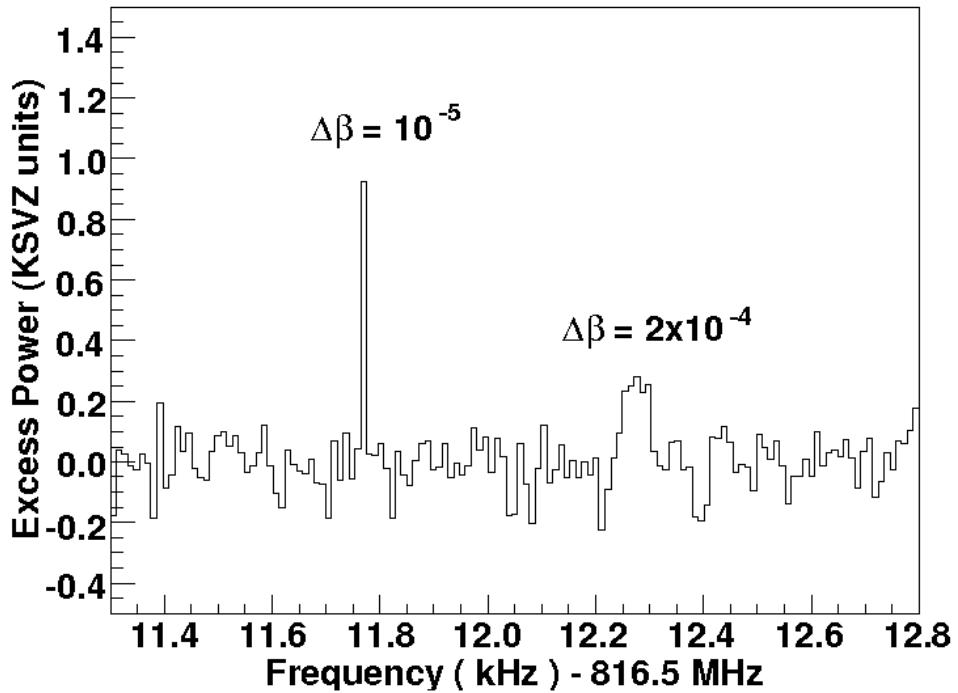


Figure A-2. Monte Carlo simulations of non-virialized axion signals imposed on real HR data. The broader peak (velocity dispersion of 2×10^{-4}) is shown as an example of the narrowest signal that the MR channel can resolve. An excess power of 0 corresponds to power in that bin equal to the rms noise power at that frequency.

having lower velocity dispersion stands dramatically above the noise. Note that at its highest resolution, the HR channel would have non-Gaussian noise. For a detailed discussion of the possible origins and the noise characteristics of low dispersion axion signals see Duffy *et al.*[13].

A combined-spectra analysis relies on the signal adding coherently while the noise adds randomly. Thus it is most effective if the signal remains in the same bin from spectrum to spectrum. Because the HR channel is affected by the Doppler shift of an axion signal due to the time varying velocity of the detector, this constraint imposes a lower limit on the bin width of this analysis. In a worst-case scenario, the combined daily signal modulation from both the orbital and rotational motions of the Earth is at most a few hertz[15, 16]. A bin width larger than twice the modulation amplitude is required

to minimize the probability of the signal drifting into a neighboring bin between scans. Therefore, each neighboring set of 256 bins in the 42 mHz spectrum was co-added, yielding a spectrum with a bin width of 10.8 Hz. The resulting spectrum is sufficiently averaged so as to possess Gaussian noise and to be mostly insensitive to the signal modulation caused by both orbital and rotational terrestrial motion on the time scale of a few days.

The modulation-insensitive spectra, shown in figure [A-3](#), were corrected for systematic effects, the most noticeable being the shape given to the power spectrum by the crystal filter. A reference spectrum representative of the crystal filter shape was created by averaging several days worth of data. This reference spectrum was divided out before each spectrum was cropped to the 30-kHz bandwidth of the crystal filter. Any lingering broad spectral structure was due to frequency dependent interactions within the amplifier chain. A sixth-order polynomial was fit to and divided out of each spectrum to remove this structure.

Frequencies with excessive power (i.e., bins containing power which is comparable to the expected power deposition from KSVZ axions) and/or excessive noise (i.e., SNR below 5) were rescanned within a day or two. These rescans were averaged with the previous data set to reduce noise caused by statistical fluctuations. For frequencies having too much power to exclude a KSVZ axion signal at 90% confidence, rescans totaling an additional 10 to 20 minutes of exposure time were taken to determine if the signal warranted further consideration as an axion candidate. All rescans extended above and below the frequency in question by several kHz. Because an actual axion signal must be persistent, the disappearance of a signal would indicate that its origin is either from statistical fluctuations or from an external transient source, rather than from an axion flow. In the frequency range of 812.0 MHz to 892.8 MHz (3.3 μ eV - 3.69 μ eV) we rescanned several frequencies with excess power, though no signals persisted through the rescan process.

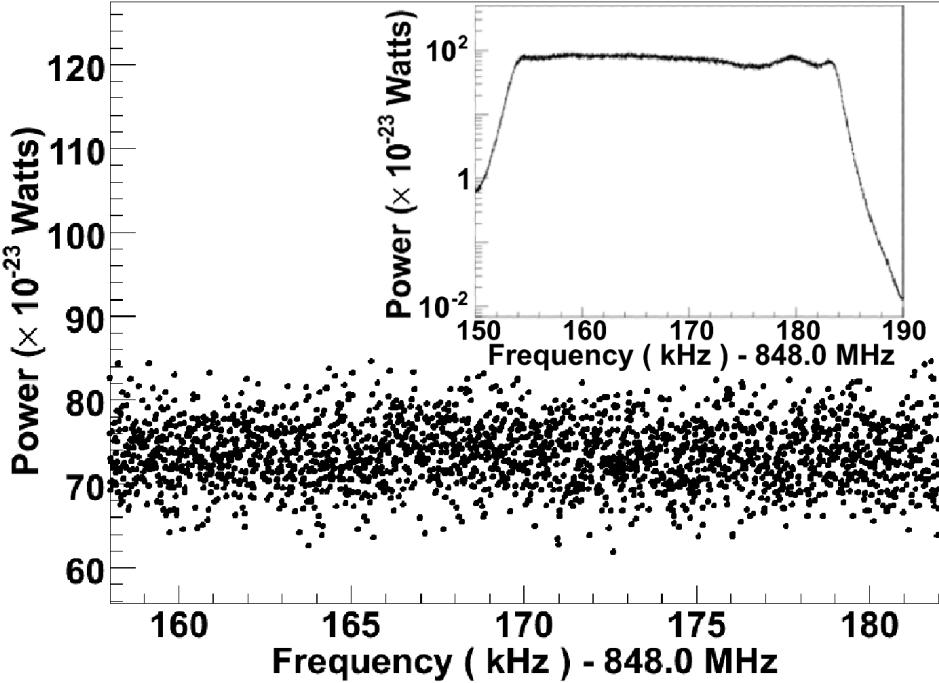


Figure A-3. An example of a modulation-insensitive power spectrum after the crystal filter shape and amplifier chain effects have been divided out. The insert shows the raw power spectrum with a passband shape imparted by the crystal filter.

Because no signals remained after performing rescans, the combined *SNR* and accumulated power were used to set limits on the local density of non-virialized axionic dark matter. Both KSVZ and DFSZ models are considered. In the 800 MHz range, a peak in the power spectrum from an axion flow with velocity relative to the detector of $\beta \sim 10^{-3}$ and dispersion of $\Delta\beta \lesssim 3 \times 10^{-6}$ would have a width of $\lesssim 5$ Hz. The limits presented here are for axion flows having this dispersion or lower. At 90% confidence, we limit the local density of non-virialized axions to exclude $\rho \gtrsim 0.2$ GeV/cm³ for KSVZ models and $\rho \gtrsim 1.4$ GeV/cm³ for the DFSZ model. Taking 0.52 GeV/cm³ as an estimate of the average local density[65, 85], non-virialized axions are excluded for KSVZ models over a mass range of $m_a = 3.3 \mu\text{eV}$ to $3.69 \mu\text{eV}$. Consider also the caustic ring halo model, presented in detail in Ref. [55], which predicts that most of the dark matter in the local neighborhood is in discrete flows. At the Earth's location in the Milky Way, one

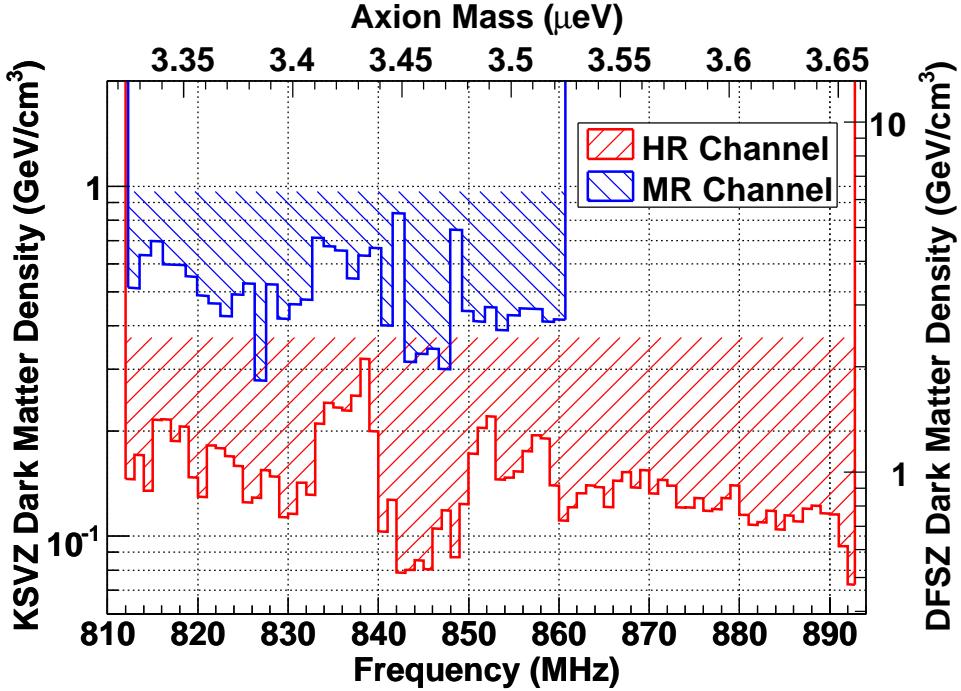


Figure A-4. Density limits for the HR channel ($\Delta\beta \lesssim 3 \times 10^{-6}$) at 90% confidence from 812 MHz to 892.8 MHz. The scale for limits on KSVZ and DFSZ axions are shown on the left and right axes respectively. Density limits for the currently published MR channel data ($\Delta\beta \lesssim 2 \times 10^{-4}$)^[14] are shown for comparison.

such flow is predicted to have a density as large as $\rho \approx 0.84 \text{ GeV/cm}^3$ which is well over our stated exclusion limit for KSVZ model axions. The HR density limits for non-virialized KSVZ and DFSZ axions are shown in Fig. A-4 with the MR limit (calculated from [14]) overlaid for comparison.

The existing time-series data will permit a future analysis of this same data set that will be capable of reaching finer resolutions while still accounting for terrestrial signal modulation. The analysis will return to a single initial scan analysis, where candidate signals found in one scan will be compared against any rescans having frequency coverage corresponding to where the signal could have moved to as a result of Doppler shifts. The time between scans will determine the maximum expected signal modulation for these candidates. To warrant further consideration, a candidate signal must be present in rescans within a window set by the maximum Doppler modulation about

the original candidate frequency. In addition, taking the FFT of smaller intervals of the time series data enables a range of dispersions to be tested. Further, it can be seen in Eq. A-4 that the SNR scales inversely as the square root of the signal bandwidth. Thus a finer resolution analysis would yield a higher sensitivity to $g_{\alpha\gamma\gamma}$. We expect approximately an order of magnitude improvement in sensitivity for the HR channel using such an analysis at a resolution of 100 mHz.

This work has been supported by the U.S. Department of Energy through grant numbers DE-FG02-97ER41029, DE-FG02-96ER40956, DE-AC52-07NA27344, and DE-AC03-76SF00098, as well as through the LDRD program at Lawrence Livermore National Laboratory.

APPENDIX B ANALYSIS CODE

The MatLab code used for this analysis is saved to a repository at the University of Florida. It is freely accessible from <http://ufdc.ufl.edu/IR00003620/00001>. What follows is a listing of the files contained therein. Included for each file is a brief description, and a list of its inputs and outputs. An input or output listed as "HR data structure" is a MatLab structure containing both the state data and the measured data for a single file.

- HRprocess.m - This is the "main" function. It removes systematic errors from a single file.
 - Input: File name
 - Output: HR data structure
- HRbatch.m - Calls HRprocess.m for a list of files, saving each output to disk.
 - Input: File containing a list of file names
 - Output: Spectra saved as .mat files
- HRread.m - Reads into memory a data file formatted as .mat, .txt, .psd, or .ttr.
 - Input: File name
 - Output: HR data structure
- HRwrite.m - Saves a spectrum as a .mat or .txt file.
 - Input: HR data structure, output file name
 - Output: A single spectrum saved as a .mat or .txt file
- HRcut.m - Cuts a faulty spectrum from the data set.
 - Input: HR data structure, cut code
 - Output: HR data structure
- HRfit.m - Fits a receiver response function to the truncated spectrum.
 - Input: HR data structure, degree of the polynomial fit
 - Output: Polynomial coefficients, reduced χ^2 of the fit

- HRsort.m - Tabulates the state data for all processed spectra of a given resolution.
 - Input: Resolution in mHz
 - Output: Tabulated state data saved to disk as a .mat file
- HRsearch.m - Finds candidates and looks for coincidence.
 - Input: Resolution in mHz, candidate threshold
 - Output: A list of candidates showing significant coincidence
- HRlimits.m - Calculates exclusion limits for local dark matter density.
 - Input: Resolution in mHz, start frequency, stop frequency, number of divisions
 - Output: A list of density limits and their errors

REFERENCES

- [1] F. Zwicky, *Helv. Phys. Acta* **6**, 110 (1933).
- [2] F. Zwicky, *Astro. Phys. J.* **86**, 217 (1937).
- [3] N. Jarosik *et al.*, ArXiv e-prints (2010), 1001.4744.
- [4] M. Markevitch *et al.*, *Astro. Phys. J.* **606**, 819 (2004), arXiv:astro-ph/0309303.
- [5] D. Clowe *et al.*, *Astro. Phys. J. Lett.* **648**, L109 (2006), arXiv:astro-ph/0608407.
- [6] G. Bertone, D. Hooper, and J. Silk, *Phys. Rep.* **405**, 279 (2005).
- [7] E. Kolb and M. Turner, *The Early Universe* Frontiers in Physics (Addison-Wesley Longman, Incorporated, 1990).
- [8] L. F. Abbott and P. Sikivie, *Phys. Lett. B* **120**, 133 (1983).
- [9] J. Preskill, M. B. Wise, and F. Wilczek, *Phys. Lett. B* **120**, 127 (1983).
- [10] M. Dine and W. Fischler, *Phys. Lett. B* **120**, 137 (1983).
- [11] J. Ipser and P. Sikivie, *Phys. Rev. Lett.* **50**, 925 (1983).
- [12] C. Hagmann *et al.*, *Phys. Rev. Lett.* **80**, 2043 (1998), arXiv:astro-ph/9801286.
- [13] L. D. Duffy *et al.*, *Phys. Rev. D* **74**, 012006 (2006), arXiv:astro-ph/0603108.
- [14] S. J. Asztalos *et al.*, *Phys. Rev. Lett.* **104**, 041301 (2010), 0910.5914.
- [15] F.-S. Ling, P. Sikivie, and S. Wick, *Phys. Rev. D* **70**, 123503 (2004).
- [16] M. S. Turner, *Phys. Rev. D* **42**, 3572 (1990).
- [17] J. Hoskins *et al.*, *Phys. Rev. D* **84**, 121302 (2011).
- [18] V. C. Rubin and W. K. Ford, Jr., *Astro. Phys. J.* **159**, 379 (1970).
- [19] V. C. Rubin, W. K. J. Ford, and N. . Thonnard, *Astro. Phys. J.* **238**, 471 (1980).
- [20] V. C. Rubin, D. Burstein, W. K. Ford, Jr., and N. Thonnard, *Astro. Phys. J.* **289**, 81 (1985).
- [21] C. S. Kochanek *et al.*, *Astrophys. J.* **585**, 161 (2003).
- [22] S. W. Allen, R. W. Schmidt, and A. C. Fabian, *Mon. Not. R. Astron. Soc.* **334**, L11 (2002).
- [23] A. Vikhlinin *et al.*, *Astrophys. J.* **640**, 691 (2006).
- [24] M. Milgrom, *Astro. Phys. J.* **270**, 365 (1983).

- [25] M. Milgrom, *Astro. Phys. J.* **270**, 371 (1983).
- [26] The CDMS II Collaboration, *Science* **327**, 1619 (2010), 0912.3592.
- [27] G. Angloher *et al.*, *The European Physical Journal C* **72**, 1 (2012).
- [28] L. Kawano, D. Schramm, and G. Steigman, *Astro. Phys. J.* **327**, 750 (1988).
- [29] G. R. Farrar and E. W. Kolb, *Phys. Rev. D* **53**, 2990 (1996).
- [30] CMS Collaboration, e. a. Chatrchyan, S., *Phys. Rev. Lett.* **107**, 221804 (2011).
- [31] LUX Collaboration *et al.*, ArXiv e-prints (2013), 1310.8214.
- [32] E. Armengaud *et al.*, *Phys. Rev. D* **86**, 051701 (2012).
- [33] P. Beltrame, ArXiv e-prints (2013), 1305.2719.
- [34] C. E. Aalseth *et al.*, *Phys. Rev. Lett.* **106**, 131301 (2011).
- [35] S. L. Adler, *Phys. Rev.* **177**, 2426 (1969).
- [36] J. Bell and R. Jackiw, *Nuovo Cimento A* **60**, 47 (1969).
- [37] C. A. Baker *et al.*, *Phys. Rev. Lett.* **97**, 131801 (2006).
- [38] R. D. Peccei and H. R. Quinn, *Phys. Rev. Lett.* **38**, 1440 (1977).
- [39] R. D. Peccei and H. R. Quinn, *Phys. Rev. D* **16**, 1791 (1977).
- [40] S. Weinberg, *Phys. Rev. Lett.* **40**, 223 (1978).
- [41] F. Wilczek, *Phys. Rev. Lett.* **40**, 279 (1978).
- [42] J. E. Kim, *Phys. Rev. Lett.* **43**, 103 (1979).
- [43] M. A. Shifman, A. I. Vainshtein, and V. I. Zakharov, *Nucl. Phys. B* **166**, 493 (1980).
- [44] M. Dine, W. Fischler, and M. Srednicki, *Phys. Lett. B* **104**, 199 (1981).
- [45] A. R. Zhitnitsky, *Sov. J. Nucl. Phys.* **31**, 260 (1980).
- [46] M. Peskin and D. Schroeder, *An Introduction to Quantum Field Theory* Advanced book classics (Addison-Wesley Publishing Company, 1995).
- [47] Particle Data Group, J. Beringer *et al.*, *Phys. Rev. D* **86**, 010001 (2012).
- [48] M. S. Turner, *Phys. Rep.* **197**, 67 (1990).
- [49] J. D. Bjorken *et al.*, *Phys. Rev. D* **38**, 3375 (1988).
- [50] J. E. Kim, *Phys. Rep.* **150**, 1 (1987).

- [51] G. G. Raffelt, *Phys. Rep.* **198**, 1 (1990).
- [52] G. G. Raffelt and D. S. P. Dearborn, *Phys. Rev. D* **36**, 2211 (1987).
- [53] J. I. Read, G. Lake, O. Agertz, and V. P. Debattista, *Mon. Not. R. Astron. Soc.* **389**, 1041 (2008).
- [54] A. Klypin, H. Zhao, and R. S. Somerville, *Astro. Phys. J.* **573**, 597 (2002), astro-ph/0110390.
- [55] L. D. Duffy and P. Sikivie, *Phys. Rev. D* **78**, 063508 (2008).
- [56] J. Einasto and U. Haud, *Astron. Astrophys.* **223**, 89 (1989).
- [57] C. W. Purcell, J. S. Bullock, and M. Kaplinghat, *Astrophys. J.* **703**, 2275 (2009).
- [58] M. Zemp *et al.*, *Mon. Not. R. Astron. Soc.* **394**, 641 (2009).
- [59] D. Stiff and L. M. Widrow, *Phys. Rev. Lett.* **90**, 211301 (2003).
- [60] P. Sikivie, I. I. Tkachev, and Y. Wang, *Phys. Rev. D* **56**, 1863 (1997).
- [61] T. Abel, O. Hahn, and R. Kaehler, (2011), arXiv/1111.3944.
- [62] P. Sikivie, *Phys. Rev. D* **60**, 063501 (1999).
- [63] P. Sikivie, *Phys. Lett. B* **567**, 1 (2003), arXiv:astro-ph/0109296.
- [64] P. Sikivie, I. I. Tkachev, and Y. Wang, *Phys. Rev. Lett.* **75**, 2911 (1995).
- [65] E. Gates, G. Gyuk, and M. Turner, *Astrophys. J.* **449**, L123 (1995).
- [66] P. Sikivie, *Phys. Rev. Lett.* **51**, 1415 (1983).
- [67] P. Sikivie, *Phys. Rev. D.* **32**, 2988 (1985).
- [68] J. D. Jackson, *Classical Electrodynamics Third Edition*, Third ed. (Wiley, 1998).
- [69] C. Kittel, *Quantum Theory of Solids* (Wiley, 1987).
- [70] N. Ashcroft and N. Mermin, *Solid State Physics* (Saunders College, 1976).
- [71] A. B. Pippard, *Proc. R. Soc. London A* **191**, 385 (1947).
- [72] M. Mück, M.-O. André, J. Clarke, J. Gail, and C. Heiden, *Appl. Phys. Lett.* **72**, 2885 (1998).
- [73] S. J. Asztalos *et al.*, *Nucl. Instrum. Methods A* (2011).
- [74] R. Bradley *et al.*, *Rev. Mod. Phys.* **75**, 777 (2003).
- [75] E. Daw and R. F. Bradley, *J. Appl. Phys.* **82**, 1925 (1997).

- [76] H. T. Friis, Proc. IRE **32**, 419 (1944).
- [77] L. D. Duffy, *High Resolution Search for Dark Matter Axions in Milky Way Halo Substructure*, PhD thesis, University of Florida, 2006.
- [78] R. H. Dicke, Rev. Sci. Instrum. **17**, 268 (1946).
- [79] M. Hotz, *A SQUID Based RF Cavity Search for Dark Matter Axions*, PhD thesis, University of Washington, 2013.
- [80] H. Peng *et al.*, Nucl. Instrum. Methods A **444**, 569 (2000).
- [81] K. Freese, P. Gondolo, H. J. Newberg, and M. Lewis, Phys. Rev. Lett. **92**, 111301 (2004).
- [82] K. Freese, P. Gondolo, and H. J. Newberg, Phys. Rev. D **71**, 043516 (2005).
- [83] D. Kinion, *First Results Form a Multiple-Microwave-Cavity Search for Dark-Matter Axions*, PhD thesis, University of California Davis, 2001.
- [84] P. Sikivie and J. R. Ipser, Phys. Lett. B **291**, 288 (1992).
- [85] P. Salucci, F. Nesti, G. Gentile, and C. Frigerio Martins, Astron. Astrophys. **523**, A23 (2010).

BIOGRAPHICAL SKETCH

Jeffrey David Vegna Hoskins was born on May 22, 1981 in Rochester, New York. At the age of one, his family moved to Orlando, Florida. After graduating from Dr. Phillips High School in 1999, he attended classes at Valencia Community College and eventually transferred to the University of Central Florida (UCF). He graduated with a B.S. in physics in 2005, and quickly enrolled in the graduate program at UCF. Following one year of graduate studies, he transferred to the University of Florida, despite being a life long Seminole fan. Once there, he was introduced to the fascinating world of dark matter research, eventually joining the Axion Dark Matter eXperiment where he would pursue his doctoral studies.