

©Copyright 2025

Michaela Guzzetti

Designing, Operating and Analyzing: The Quest for Axion Dark Matter with ADMX

Michaela Guzzetti

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

Doctor of Philosophy

University of Washington

2025

Reading Committee:

Gray Rybka, Chair

Alejandro Garcia

Miguel Morales

Program Authorized to Offer Degree:

Physics

University of Washington

Abstract

Designing, Operating and Analyzing: The Quest for Axion Dark Matter with ADMX

Michaela Guzzetti

Chair of the Supervisory Committee:
Gray Rybka
Department of Physics

The Axion Dark Matter eXperiment (ADMX), located at the University of Washington, is a world-renowned experiment and is at the forefront of the hunt for the elusive QCD axion. The QCD axion, originally proposed in the late 1970s to solve a problem in particle physics, was quickly identified as a promising dark matter candidate. While it was initially thought to be impossible to detect due to its extremely weak coupling to standard model particles, over the past couple of decades major advancements in experimental technology have allowed experiments, like ADMX, to become sensitive enough to detect such a particle. In this thesis I will begin by providing context on the history of and evidence for dark matter, the origin of the axion as a theoretical particle, and what makes the axion a good dark matter candidate. Next, I will explain how experiments like ADMX search for axions today as well as situate ADMX in the landscape of existing and proposed axion searches. The remainder of the thesis will cover the details of the most recent data taking run with ADMX, Run 1D. More specifically I will detail the hardware changes made since the last run, the most thorough noise calibration campaign done with ADMX to date, the operations and procedures involved in taking data, as well as the full analysis process. I will finish by reporting results including a 90% confidence level upper limit on axion-photon coupling between 1.088-1.315 GHz, as well as a discussion about the discovery ability of this data set.

TABLE OF CONTENTS

	Page
List of Figures	iii
Chapter 1: Introduction	1
1.1 Dark Matter	1
1.2 Axion Theory	7
1.3 Axion Searches	15
Chapter 2: Run 1D Hardware	23
2.1 Axion Dark Matter eXperiment Overview	23
2.2 Cavity	27
2.3 Cryogenic Receiver	33
2.4 DC Cables and Sensors	45
Chapter 3: Noise Calibration	61
3.1 Sources of Noise	61
3.2 Noise Models	66
3.3 Noise Measurements	74
Chapter 4: Data Taking Operations	97
4.1 Initial Scan	100
4.2 SNR Leveling and Candidate Rescan	101
Chapter 5: Axion Search Analysis	105
5.1 Parameter Extraction	106
5.2 Warm Receiver Removal	107
5.3 Raw Spectrum Preparation	110
5.4 Grand Spectrum Assembly	115

5.5	Candidate Identification and Elimination	117
5.6	Exclusion Limit	124
5.7	Results	125
Chapter 6:	Conclusions	132
Bibliography	136

LIST OF FIGURES

Figure Number	Page	
1.1	Rotation curves of 25 galaxies of various Hubble types. Figure and caption taken from Ref. [43].	3
1.2	Left: Color image of the Bullet Cluster taken using the 6.5 m Magellan telescope’s Inamori Magellan Areal Camera and Spectrograph. Right: X-ray image of the Bullet Cluster taken using the Chandra X-ray Observatory. Shown in green contours in both panels is the gravitational potential determined from gravitational lensing. Figure taken from Ref. [62].	4
1.3	The classic wine bottle potential that leads to spontaneous symmetry breaking. The lowest energy state is no longer at the point where ϕ (some scalar field) is equal to zero, which has become a local maximum, but instead is described by some randomly chosen point along the “rim” of equivalent global minima. Figure taken from Ref. [23].	9
1.4	Illustration of the evolution of the axion field potential over time. In the early Universe $T > f_a$, the PQ symmetry is unbroken, and $\bar{\theta}$ can take on any value. As the Universe cools to $T < f_a$, the PQ symmetry is spontaneously broken, and $\bar{\theta}$ takes on a random initial value, $\bar{\theta}_i$ between $\pm\pi$. At this point, the axion is a true, massless Goldstone boson. Finally, as the temperature of the Universe drops below Λ_{QCD} , the potential tilts, forcing the field into its CP conserving minimum at $\bar{\theta} = -\frac{\langle a \rangle}{f_a}$. Note that in this figure, $\theta = \bar{\theta} + \frac{a}{f_a}$ is the effective vacuum angle, which is why it shows θ rolling to 0 in the right most potential. Figure taken from Ref. [134].	13
1.5	Feynman diagram of axion decay into photons. Left: Axion to photon conversion in vacuum. Right: Axion to photon conversion in a static magnetic field (\mathbf{B}_0). Figure taken from Ref. [168].	16
1.6	Simplified diagram of a Sikivie haloscope. Axions passing through a resonant cavity that is immersed in a static magnetic field convert into observable photons. These photons are then picked up by an antenna and amplified by an ultra-low-noise amplifier. Figure taken from Ref. [37].	18

1.7	A simplified panorama of the axion-photon coupling parameter space. Constraints are roughly color coded, with direct searches for dark matter axions in the solar neighborhood shown in blue, searches for axions produced in astrophysical environments shown in green, and pure laboratory searches in red. The pure laboratory searches differ from direct detection experiments in that the pure laboratory searches do not make any assumptions about axions being dark matter, whereas the blue searches do. Figure taken from Ref. [133] . . .	22
2.1	Cutaway diagram of the ADMX insert inside the magnet bore. The insert, which contains the microwave cavity, as well as the cryogenic receiver chain and dilution refrigerator, is lowered into the magnet bore for data taking. Power is extracted from the cavity with an antenna and then amplified by a Josephson Parametric Amplifier (JPA) that lives in the field-free region marked with dashed red lines. The power is then amplified a second time by an HFET amplifier, before exiting to the room temperature receiver where it is processed further. Note that this diagram is slightly outdated as it shows the cavity as having two tuning rods. The Run 1D cavity only has one, larger tuning rod, but otherwise the diagram is representative of the current design.	24
2.2	Helium isotope phase transition diagram from Bluefors [6]. As the mixture of He-3 and He-4 is cooled below about 870 mK, it undergoes a spontaneous phase transition, splitting into a He-3 rich (concentrated) phase and a He-3 poor (diluted) phase.	27
2.3	A view inside the ADMX cavities for Runs 1A-1D, showing the differences in tuning rod size and quantity over time.	28
2.4	A series of photos showing the installation of the rod into the Run 1D cavity. A) Inside of the Run 1D cavity B) Lifting the rod to insert the bottom sapphire axle before installation C) Former ADMX undergraduate Peter Park and current LLNL staff scientist Nick Du are shown preparing to insert the rod into the cavity D) The rod and top sapphire axle after being installed into the cavity	29
2.5	A mode map showing the resonant frequencies of various cavity modes changing with rod position. The mode of interest for ADMX, TM_{010} , can be tuned over a range of roughly 0.9 – 1.4 GHz.	29

2.6	The axion power extracted by the antenna is proportional to $\beta/(1 + \beta)^2$, as shown in purple. The axion haloscope scan-rate on the other hand is proportional to $\beta^2/(1 + \beta)^3$, as shown in green. There is a dotted line at $\beta = 1$ to mark where the axion power extracted by the antenna is maximized with respect to β , which corresponds to critical coupling. There is a dashed line at $\beta = 2$ to mark where the axion haloscope scan rate is maximized with respect to β , which turns out to be slightly over coupled.	32
2.7	The cryogenic receiver chain for Run 1D. There are four different lines shown. The weak and bypass lines are primarily used for characterization of cavity parameters, while the pump line provides the input pump tone for parametric amplification with the Josephson Parametric Amplifier (JPA). Lastly, the output line is used to amplify and readout signals from the cavity. ADMX uses two amplifiers, a JPA and an HFET amplifier. The three small squares with circular arrows in them are unidirectional devices known as circulators which are used to direct the signal path and prevent downstream noise from reaching the cavity. There are 10 and 20 dB attenuators on weak, bypass and pump lines at the reservoir bottom and mixing chamber stages respectively. The black circles represent 0 dB “attenuators” used as thermal sinks. Connected to the weak line via a directional coupler is the hot load, which is a variable temperature stage used for noise calibration. It can be connected to the output line using an RF switch.	34
2.8	Top: standard transmission measurement of the TM_{010} cavity resonance. This type of transmission measurement is used to calibrate various experimental parameters such as the quality factor and receiver gain. Bottom: wide transmission measurement of many different cavity modes. This type of transmission measurement is used for making mode maps and tracking the TM_{010} mode.	35
2.9	Reflection measurement of the TM_{010} cavity resonance. This measurement is used to calibrate various experimental parameters such as the quality factor, receiver gain, and cavity coupling strength.	36

2.10	Figure taken from Ref. [111]. A JPA circuit diagram. The region outlined in purple contains an array of SQUID loops (right hand side) that are in series with a geometric inductance and shunted by a capacitor. The JPA can either be current-pumped or flux-pumped. In current-pumped operation (as was used for Run 1D), the pump tone at frequency ω_P enters the circuit via the main transmission line shown (red). In flux-pumped operation, the pump tone enters via the inductively coupled DC line (blue). In either case, the input signal a_{in} enters the main transmission line, mixes with the pump tone in the JPA, and then is reflected back down the transmission line and exits the JPA from the same port it entered in.	38
2.11	The cryogenic receiver chain for Run 1C. There are five different lines shown. The weak and bypass lines are primarily used for characterization of cavity parameters, while the pump line provides the input pump tone for parametric amplification with the Josephson Parametric Amplifier (JPA). The hot load line is used to monitor the gain of the receiver while it is connected to the hot load (left-most 20 dB attenuator in the diagram), which is a variable temperature stage used for noise calibration. The hot load can be connected to the output line using an RF switch. Lastly, the output line is used to amplify and readout signals from the cavity. ADMX uses two amplifiers, a JPA and an HFET amplifier. The three small squares with circular arrows in them are unidirectional devices called circulators which are used to direct the signal path and prevent downstream noise from reaching the cavity. There are 3 dB attenuators on weak, pump and hot load lines cooled baffle temperature stage. There is another set of 3 dB attenuators on the weak, bypass, pump and hot load lines at the 1K plate temperature stage as well. In addition, there are 20 dB (10 dB) attenuators on the weak and bypass (pump and hot load) lines at the reservoir bottom temperature stage. There are also 20 dB (30 dB) attenuators on the weak, pump and hot load (bypass) line at the mixing chamber temperature stage. Finally, there are 0 dB “attenuators” that serve as thermal sinks on the weak, bypass, pump and hot load lines at the ICP and still temperature stages.	44
2.12	Screenshot of the ADMX Run 1D wiring spreadsheet. The example shown here is the tab associated with Cable E.	50
2.13	ADMX Run 1D image gallery of sensors in Fig. 2.14. These images are on the ADMX wiki page that can be accessed by internal collaborators who are interested in locating a specific sensor. Not all sensors were able to be easily photographed, but those that were are included in the gallery.	51
2.14	ADMX Run 1D sensor map. Each number corresponds to a sensor in Fig. 2.15.	52

2.15	ADMX Run 1D sensor map key. Each number corresponds to a sensor in Fig. 2.14.	53
2.16	Design drawing for the upgrade dilution refrigerator cable, Cable F. This version of Cable F included more heat sinking bobbins in more strategic locations, different wire material, and more easily interchangeable connectors.	54
2.17	Pinout diagram for the upgraded dilution refrigerator cable, Cable F. Pins 13-16 were designated to be made with copper wire, while the rest of the pins are phosphor bronze. This is because the still heater, which uses more current than the temperature sensors, was designated to be on pins 14 and 15, leaving 13 and 16 as backups.	56
2.18	Left: the former “breakout box” for Cable F. Really, it was just a connector that could be disconnected from the stand. The left hand side where wires are soldered into connected to the old Cable F, and the right hand side with the white cables coming out were routed to the sensors and the still heater. This connector was not only dilapidated, but it was not as easily labeled and changes were difficult to keep track of because the wires on the cable itself were what needed to be re-soldered in the case of a break, rather than having some intermediate connection that changes could be made to. Right: new and improved breakout board. This board is what is housed inside of the box shown in Fig. 2.19. This breakout board has entirely fresh wiring and connectors, and is much more clearly labeled with every connector and every pin on each connector having a number associated with them. In addition, because Cable F plugs into this board, the changes made during rerouting are made to the board, rather than the cable itself.	57
2.19	The newly designed breakout box for Cable F mounted on top of the cavity next to the breakout box for Cable E (left edge of the image). The connectors for temperature sensors can be seen on the left hand side of the breakout box and there are 3 spares. Cable F is shown plugged into the top of the box. . .	58
2.20	Drawing showing the changes made to Cable D in the cryogenic space. Namely, the breakout box was relocated from the top of the reservoir to the cooled baffle and another heat sink was added to the top of the reservoir in its place. In addition, a filter box was added on the bottom of the reservoir to reduce electronic noise on the JPA bias line.	60

3.1	ADMX RF diagram in the cold space. Other than the HFET, all the components are connected to the milliKelvin stage directly. The blue arrows show the shared path for both noises and possible signal from the cavity. The brown arrows show the path of the attenuator A thermal noise to the cavity. Different stages (i.e. stage 1, stage 2, etc.) are labeled corresponding to the noise model in Sec. 3.2.	67
3.2	Left: Power data taken before beginning heating for a JPA-off-hot-load measurement to estimate the digitizer uncertainty with the JPA off (P_{off}). The example shown here corresponds to an uncertainty of $8\text{e}-10$ W. Right: Power data taken before beginning heating for a JPA-on-hot-load measurement to estimate the digitizer uncertainty with the JPA on (P_{on}). The example shown here corresponds to an uncertainty of $8\text{e}-9$ W.	75
3.3	Here I plot the hot load temperature and the output power changing together during the course of two JPA-off-hot-load measurements done at 1280 MHz. The temperature is smoothly increased using a DC power supply from $\sim 150 - 500$ mK and then smoothly decreased using the same voltage steps. The output power tracks the changing temperature closely, as expected from the linear relationship described in Eq. 3.21.	77
3.4	Here the hot load temperature is plotted against the output power for the two JPA-off-hot-load measurements done at 1280 MHz. The data are fit to Eq. 3.25 where $T = T_{\text{HL}}$ and $T_{\text{fit}} = T_{\text{HFET}}/\alpha_{\text{eff}} = \frac{T_{\text{circ}}(1-\alpha)}{\alpha} + \frac{T_{\text{HFET}}}{\alpha}$. The fits shown here used the entire data range, however the fits were performed multiple times using slightly different fitting windows to estimate systematic uncertainty introduced from a given choice of fitting window. After taking into account differences in the fit results caused by adjusting the fitting window, I found $T_{\text{HFET}}/\alpha_{\text{eff}} = 6.17 \pm 0.75$ (6.13 ± 0.20) K for the January 4th (January 30th) measurements. Both values are in agreement well within 1σ , indicating high stability in the HFET noise performance over time. Averaging the two effective HFET noise temperatures results in $T_{\text{HFET}}/\alpha_{\text{eff}} = 6.15 \pm 0.39$ K. . .	78
3.5	Here I plot the hot load temperature and the output power changing together during the course of two JPA-on-hot-load measurements done at 1280 MHz. The temperature is smoothly increased using a DC power supply from $\sim 150 - 200$ mK and then smoothly decreased using the same voltage steps. The output power tracks the changing temperature closely, as expected from the linear relationship described in Eq. 3.28.	81

- 3.6 Here the hot load temperature is plotted against the output power for the two JPA-on-hot-load measurements done at 1280 MHz. The data are fit to Eq. 3.29 where $T = 2T_{\text{HL}}$ and $T_{\text{fit}} = T_{\text{JPA,eff}} = 2\frac{T_{\text{circ}}(1-\alpha_1)}{\alpha_1} + \frac{T_{\text{JPA}}}{\alpha_1}$. The fit shown here used the entire data range, but the fit was performed multiple times using slightly different fitting windows. After taking into account differences in the fit result caused by adjusting the fitting window, I found $T_{\text{JPA,eff}} = 0.139 \pm 0.021$ (0.143 ± 0.019) K for the February (June) measurements. Both values are in agreement within error bars, indicating stability in the JPA noise both over time and under different bias settings. Averaging the two effective JPA noise temperatures results in $T_{\text{JPA,eff}} = 0.141 \pm 0.014$ K. 82
- 3.7 Here I plot the mixing chamber temperature and the output power changing together during the course of two JPA-off-cavity measurements done at 1280 MHz. The temperature is tracked during the course of a full system cool-down. The output power tracks the changing temperature closely, as expected from the linear relationship described in Eq. 3.23. 85
- 3.8 Here the mixing chamber temperature is plotted against the off-resonance output power for the two JPA-off-cavity measurements done at 1280 MHz. The data are fit to Eq. 3.25 where $T = T_{\text{mxc}}$ is the hot load temperature and $T_{\text{fit}} = T_{\text{HFET}}$. The fits shown here used the entire data range, but the fits were performed multiple times using slightly different fitting windows. In addition to adjusting the fitting window, the fits were performed using $T = T_{\text{circ}}$ instead because of the imperfect assumption that $T_{\text{mxc}} = T_{\text{circ}}$. After taking into account differences in the fit results caused by adjusting the fitting window and using a different temperature sensor, I found $T_{\text{HFET}} = 4.18 \pm 0.26$ (3.80 ± 1.02) K for the March (June) calibration. Both values are in agreement within 1σ , indicating stability in the HFET noise over time. The value used for calculating α is an average of the two: $T_{\text{HFET}} = 3.99 \pm 0.52$ K. 86
- 3.9 Here one can see the relevant temperatures and the output power changing together during the course of the JPA-on-cavity measurement done at 1280 MHz. The off-resonance fit uses $T = 2T_{\text{stage1}}$, so that is what is plotted here. Meanwhile, the on-resonance fit uses $T = T_{\text{stage1}}|\Gamma_{\text{cav}}|^2 + T_{\text{cav}}(1 - |\Gamma_{\text{cav}}|^2) + T_{\text{stage1}}$, so that is what is plotted for the on-resonance case. The temperatures are tracked during the course of a full system warm-up. The output power tracks the changing temperature closely in both cases, as expected from the linear relationship described in Eq. 3.27. The black dashed lines mark where data were removed due to rapidly changing gain resulting in poor gain correction. 89

- 3.10 Here the relevant temperatures are plotted separately against the off-resonance and on-resonance output power for the JPA-on-cavity measurement done at 1280 MHz. The data are fit to Eq. 3.29 where $T = 2T_{\text{stage1}}$ ($T = T_{\text{stage1}}|\Gamma_{\text{cav}}|^2 + T_{\text{cav}}(1 - |\Gamma_{\text{cav}}|^2) + T_{\text{stage1}}$) for the off-resonance (on-resonance) data and $T_{\text{fit}} = T_{\text{JPA,eff}} = 2\frac{T_{\text{circ}}(1-\alpha_1)}{\alpha_1} + \frac{T_{\text{JPA}}}{\alpha_1}$. The fits shown here used the entire data range, but the fits were performed multiple times using slightly different fitting windows. After taking into account differences in the fit results caused by adjusting the fitting window, I found $T_{\text{JPA,eff}} = 0.372 \pm 0.022$ (0.372 ± 0.018) K for the off-resonance (on-resonance) data. 90
- 3.11 Hot Load T_{sys} comparison at 1280 MHz. Here I plot the comparison between the two methods, using SNRI and $T_{\text{HFET}}/\alpha_{\text{eff}}$ (pink points) and using JPA fit (green points), over the course of both hot load measurements done at 1280 MHz. It is clear that for both the February and June data the system noise temperature we calculate is consistent between the two methods within error bars. 92
- 3.12 Cavity T_{sys} comparison at 1280 MHz. Here I plot the comparison between the two methods, using SNRI and $T_{\text{HFET}}/\alpha_{\text{eff}}$ (pink points) and using JPA fit (green points), over the course of the cavity calibration done at 1280 MHz. Data above 0.3(0.32) K have been randomly downsampled for plotting purposes due to the high density of data in that region. The discontinuity around this temperature was caused by a sharp increase in JPA gain during the course of the measurement, causing a sudden drop in T_{sys} . For both the off-resonance and on-resonance data the calculated system noise temperature is consistent between the two methods within error bars. 93
- 4.1 Example of a JPA bias plot. The x-axis is the JPA pump power, the y-axis the JPA bias current and the color map corresponds to the JPA gain in dB. This measurement was centered on a bias current of -2.01 mA with a span of 0.02 mA and a pump power of 11.3 dBm with a span of 1 dBm. There is a darker band of optimal bias parameters that have significantly higher gain than their surroundings ranging from the bottom left to the top right of the plot. 99
- 4.2 Top: grand spectrum SNR vs frequency for nibble 5 just before the initial pass was completed (scanning upward in frequency). Bottom: grand spectrum SNR vs frequency for nibble 5 after the SNR leveling pass. As you can see, the SNR was leveled out and nearly the entire spectrum is above our target SNR. Any narrow regions where the SNR hasn't reached the target value (like the small dip at 1238.4 MHz) are covered during the candidate rescan. 102

4.3	Here I plot the cavity resonant frequency (f_0) vs time to show the progression of data taking throughout Run 1D. As one can see, the run started up around 1300 MHz and the experiment more or less scanned downward in frequency ending around 1100 MHz. The inset plot shows the data taking cadence within a single nibble, nibble 7, as an example. One can see there is a slower initial scan, followed by a faster SNR leveling scan and two even quicker candidate rescans as described in the text.	104
5.1	Frequency dependence of the JPA SNRI. A few hours of data (blue) are averaged together for plotting purposes, but for the actual analysis no averaging was done. The data are fit using an SG filter (orange) to capture the general shape of the SNRI frequency dependence. The SG-filtered SNRI is then used to calculate a frequency dependent T_{sys} for each spectrum using Eq. 3.17. . .	107
5.2	Diagram of the ADMX warm receiver chain for Run 1D. From left to right: (A) DC amplifier (Low Noise Factory model LNF-LNR1_15B_SV), (B) DC amplifier (Mini-Circuits model ZX60-33LN-S+), (C) SPDT mechanical switch box (Mini-Circuits model RC-4SPDT-A18), (D) and (E) DC amplifiers (Mini-Circuits model ZX60-3018G-S+), (F) directional coupler (Mini-Circuits model ZX30-17-5-S+), (G) Polyphase Microwave image reject mixer (model IRM0622B), (H) low pass filter (Mini-Circuits model ZX75LP-50-S+), (I) directional coupler (Mini-Circuits model ZX30-17-5-S+), (J) 150 kHz wide custom band-pass filter centered on 10.7 MHz, (K) and (L) DC amplifiers (Mini-Circuits model ZFL-500+), (M) 2 MHz wide band-pass filter centered on 10.7 MHz (Mini-Circuits model SBP-10.7+). The directional couplers enable debugging before and after the mixing stage.	108
5.3	Run 1D warm receiver shape. The green points are the averaged spectrum, and the black line is the smoothed receiver shape that was fit using a Savitsky-Golay filter with order 3 and window length 17. Note that the frequency axis in units of bin number, not Hz.	110
5.4	Examples of three different individual power spectra after the warm receiver shape from Fig. 5.3 has been removed. Variations in the shape post-warm-receiver-removal are caused by frequency-dependent gain of the cryogenic receiver. There are further rounds of receiver removal to account for these differences in the cryogenic receiver shape, which is discussed further in this section and the results of which can be seen in Fig. 5.5.	112

5.5	The same three spectra shown in Fig. 5.4 after removing any residual background shape imposed by the cryogenic receiver. The process is described further in the text. Note that the y-axis is in units of SNR rather than power for visualization purposes. The spectra used for assembling the grand spectrum have the y-axis in units of power as expected.	113
5.6	The same three spectra shown in Fig. 5.4 and Fig. 5.5 after scaling the spectra to represent a signal inside of the cavity and converting the y-axis to units of Watts/Kelvin. This process is discussed further in the text.	114
5.7	Beginning in the top left, as the cavity tunes across the signal the signal experiences resonant enhancement at $f_{\text{sig}} = f_0$. This is how a “post-cavity” version of an axion-like signal would behave. On the contrary, looking at the top right we can see the behavior of a “post-cavity” version of an RFI-like signal. In this case, the signal amplitude is constant regardless of the difference between f_{sig} and f_0 . Moving to the bottom left, an “in-cavity” version of an axion-like signal would look much like the “post-cavity” version of an RFI-like signal, in the sense that it has a constant amplitude with respect to $f_{\text{sig}} - f_0$. In contrast, an “in-cavity” version of an RFI-like signal actually appears to minimize on-resonance with the cavity.	119
5.8	The top plot shows the persistent candidate from nibble 12 at 1089.9995 MHz. As one can see, it was undoubtedly a strong power excess that was present in both the original and rescan data. The bottom left photo shows Gray and I using a coaxial antenna connected to the signal analyzer (SA), centered on 1090 MHz, to check for ambient RFI on site. The bottom right photo shows the resulting RFI signal on the SA. This check was used to eliminate this candidate and close out nibble 12.	121
5.9	Here I plot signal amplitude vs signal frequency for the persistent candidate at 1130.6135 MHz. The black data show the signal while coupled to the usual TM_{010} mode, while the red data show the signal while coupled to the TM_{011} (which should not couple to axions). The signal was not only able to couple to the TM_{011} mode, but it actually had larger signal strength while coupled to this mode than it did to the TM_{010} mode. This is strong evidence that it is not an axion-like signal. Indeed, after confirming with the SAG control team at PNNL, this candidate was identified as a secondary/blinded SAG.	122

5.10	A global limit plot putting this work (purple) in context with other experiments, with an inset zooming in on this work's 90% C.L. upper limits on $g_{a\gamma\gamma}$ (as well as limits from Ref. [17] due to overlapping coverage). The dark matter density is assumed to be 0.45 GeV/cm^3 . Gaps in the limits are due to mode crossings, regions where axion search mode of the cavity intersected other static weakly tuning modes. KSVZ and DFSZ sensitivities are shown as dashed lines. Previous limits set by ADMX are shown in teal [26, 79, 46, 33, 91, 35, 44]. Limits from other experiments depicted include those set by University of Florida (UF) [98], Rochester-Brookhaven-Florida (RBF) [179], Center for Axion and Precision Physics (CAPP) [121, 105, 117, 122, 114, 182, 181, 115, 17], Haloscope At Yale Sensitive To Axion Cold dark matter (HAYSTAC) [186, 30, 106], Grenoble Axion Haloscope project (GrA-Hal) [94], Oscillating Resonant Group AxioN experiment (ORGAN) [126, 150, 149], QUaerere AXions experiment (QUAX) [20, 21, 19, 72, 152], Relic Axion Dark Matter Exploratory Setup (RADES) [189], Taiwan Axion Search Experiment with Haloscope (TASEH) [58], CAST-CAPP [13], and CERN Axion Solar Telescope (CAST) [22]. Caption taken from Ref. [63].	129
5.11	Top plot: Gaussian probability distribution function with a mean of $\mu = P_j + 1.28\sigma_j$ where $P_j = P_{\text{threshold}}$ and standard deviation $\sigma_j = 1/\text{SNR}_{\text{avg}}$ where SNR_{avg} is the average SNR for the whole nibble. Bottom plot: Gaussian complementary cumulative distribution function with the same mean and standard deviation as the PDF.	130
5.12	Top plot: Gaussian probability distribution function with a mean of $\mu = P_j/\sigma_j$ and standard deviation of 1 where $P_j = 3.6P_{\text{DFSZ}}$ and $\sigma_j = 1/\text{SNR}_{\text{avg}}$ where SNR_{avg} is the average SNR for the whole nibble. Bottom plot: Gaussian complementary cumulative distribution function with the same mean and standard deviation as the PDF.	131

ACKNOWLEDGMENTS

I would first like to sincerely thank my advisor Gray Rybka for his mentorship, for instilling confidence in me and my ability to succeed, and for his continued support and guidance throughout my graduate career. It was an honor to work with him on such an incredible experiment for so many years.

I would also like to thank the other members of my doctoral committee: Alejandro Garcia, Miguel Morales, Masha Baryakhtar, Marilena Loverde and Jeff Nivala.

Whether it was answering questions about axion theory, troubleshooting instruments with me on site, helping locate documentation, generally keeping the experiment afloat, or being someone I could turn to for support when things got too overwhelming or frustrating, I could not have succeeded in graduate school without the support of the various members, both past and present, of the local ADMX group at CENPA. As such, I must extend a huge thank you to: Leslie Rosenberg, Dan Zhang, Chelsea Bartram, Nicole Crisosto, Nick Du, Tom Braine, Jimmy Sinnis, Grant Leum, Charles Hanretty, Cyrus Goodman, Peter Park, and Christian Boutan.

I would also like to thank the CENPA electronic shop engineering technician David Peterson for always doing excellent work on the wiring projects I tasked him with. Similarly, I would like to thank some ADMX collaborators from Washington University at St. Louis, Jonah Hoffman, Jordan Russell and Jim Buckley, for their technical advice and for facilitating the manufacturing of Cable F.

Finally, I would like to thank my friends and family for their unwavering belief in me and for always putting life into perspective when things in the lab felt insurmountable. Thank you to my sweet angel of a cat, Alder, who I could always turn to for snuggles when I needed

a break from the tedium of writing. Last but not certainly least, I would especially like to thank my best friend Cole Man, and my partner Layne Perkins. Words cannot express the gratitude I have for them always being a shoulder to cry on, for constantly reminding me of my worth, for working by my side, for their love, for their loyalty and for their ability to always make me laugh even on an all-too-common dreary day in Seattle.

DEDICATION

To my partner and our cat. I truly could not have finished without the safety and happiness our lives together have brought me.

Chapter 1

INTRODUCTION

The axion is a theoretical particle originally proposed to solve a puzzle in particle physics known as the strong CP problem. However, over the last few decades physicists have turned their attention to the axion's potential to solve the dark matter mystery as well. Let us then explore these two seemingly unrelated phenomena, the history of the axion, and the various ways that physicists are searching for axions today.

1.1 Dark Matter

The modern history of dark matter can be traced back to Fritz Zwicky's application of the Virial Theorem to the Coma Cluster in response to its surprisingly large observed velocity dispersion. In doing so he found that the expected velocity dispersion derived using the Virial Theorem was significantly smaller than the observed velocity dispersion. In his words [188]:

“If this would be confirmed we would get the surprising result that dark matter is present in much greater amount than luminous matter.”

Observations of other galaxies and galaxy clusters around the same time further supported Zwicky's statement that there appears to be some substantial amount of dark matter dominating over regular, luminous matter [167, 158]. By the 1960s, more and more scientists were becoming convinced that this observed excess mass could indeed be real, but others remained skeptical.

At the same time that some astronomers like Zwicky were studying galaxy clusters, others were focusing their studies on the circular velocity profiles of the stars and gas in individual galaxies, also known as galactic rotation curves. Dating back to the 1930s, astrophysical

observations of the rotation curve of the Andromeda galaxy (M31) indicated the presence of large amounts of non-luminous (dark) matter [124, 29]. Despite these discoveries however, through the end of the 1950s there was no consensus that the discrepancies arising from observed rotation curves conflicted with the existing understanding of galaxies [41]. Over the next two decades however, this would change. In 1970, American astronomers Vera Rubin and Kent Ford published a higher quality rotation curve of M31 which showed that the mass of M31 continuously rose with distance from the galactic center up to the last measured region of 24 kpc [155]. In the same year, Australian astronomer Kenneth Charles Freeman compared the expected peak radius of the rotation curve based on photometric observations with the observed peak radius seen with 21 cm spectral observations [88]. In doing so, he found that the observed rotation curves peaked at larger radii than predicted concluding that [41]:

“... there must be in these galaxies additional matter which is undetected, either optically or at 21 cm. Its mass must be at least as large as the mass of the detected galaxy, and its distribution must be quite different from the exponential distribution which holds for the optical galaxy.”

In the years following this result, more studies were published that confirmed galactic rotation curves remained flat as far out from the galactic center as the observers could see [154, 176, 153]. Of particular importance was the addition of radio observations which extended rotation curves out beyond the optical radii of galaxies. For example, Albert Bosma’s 1978 Ph.D. thesis included the rotation curves of 25 different galaxies made using radio observations [43]. The rotation curves remained flat beyond the optical size of the galaxies (Fig. 1.1), indicating that the mass of the observed galaxies continued to grow past the realm of stars and gas. By the end of the decade, “the case for invisible mass in the Universe [was] very strong and getting stronger” [84].

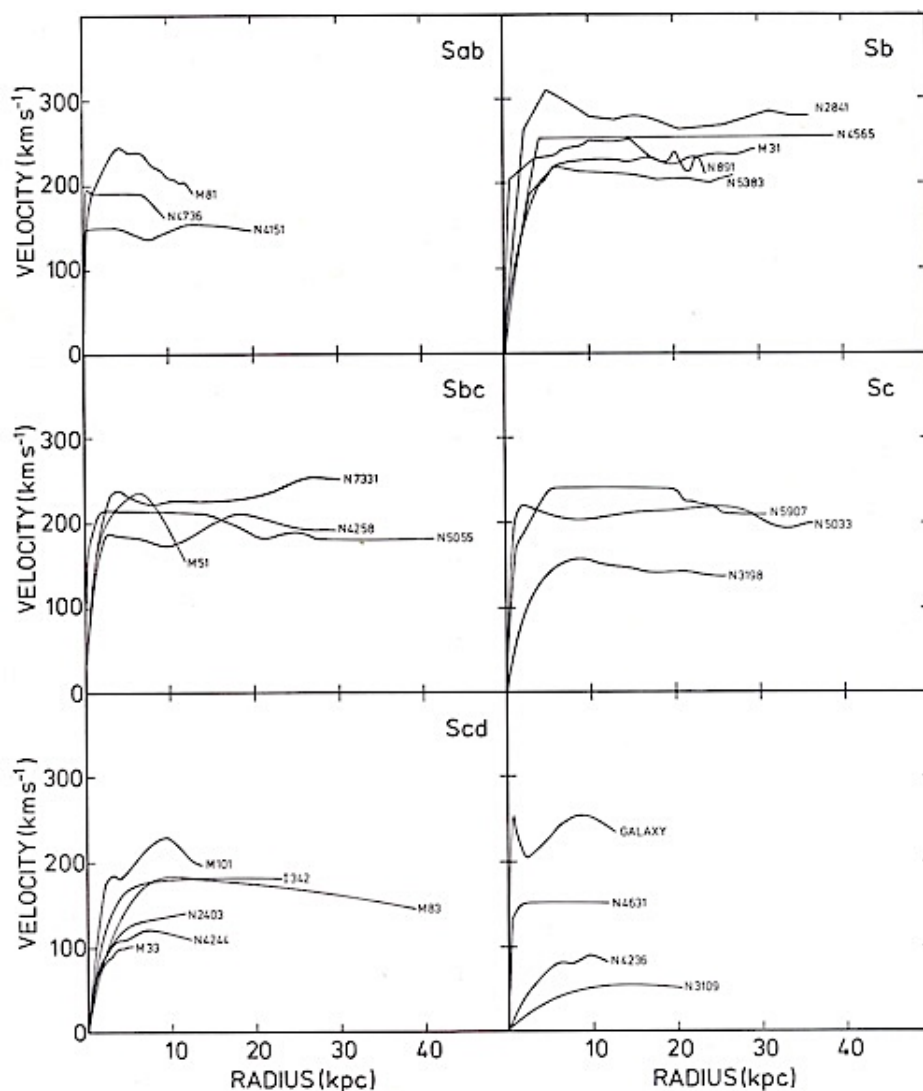


Figure 1.1: Rotation curves of 25 galaxies of various Hubble types. Figure and caption taken from Ref. [43].

However, in the early 1980s Mordehai Milgrom published a series of papers challenging the notion that dark matter was the only explanation for the observed galactic and stellar dynamics of the previous decades [130, 128, 129]. This theory was known as modified Newtonian dynamics (MOND), and, as the name suggests, modified Newton's second law to reproduce astrophysical observations in place of dark matter. However, there were multiple

issues with Milgrom’s first iteration of MOND including its inability to conserve momentum, angular momentum or energy as well as its incompatibility with general relativity. Various improvements have been made to MOND since then, resolving the issues above while introducing others. In particular, while the leading theory of MOND in the modern day is capable of explaining the observed dynamics of many spiral and elliptical galaxies, it breaks down on the scale of galaxy clusters. In 2006, observations of the Bullet Cluster further discredited MOND while bolstering the argument for dark matter [62]. As shown in Fig. 1.2, the gravitational potential of the Bullet Cluster determined from gravitational lensing does not spatially map to the majority of the baryon content of the bullet cluster (in the form of hot, x-ray emitting gas). Rather, the gravitational potential traces the distribution of galaxies, which had been separated from the hot gas during the recent collision of the two component clusters that make up the Bullet Cluster. The significant discrepancy between the location of the center of the *total* mass from the center of the *baryonic* mass could not be explained by modified theories of gravity. Only a substantial amount of non-baryonic dark matter could explain this observation.




Figure 1.2: Left: Color image of the Bullet Cluster taken using the 6.5 m Magellan telescope’s Inamori Magellan Areal Camera and Spectrograph. Right: X-ray image of the Bullet Cluster taken using the Chandra X-ray Observatory. Shown in green contours in both panels is the gravitational potential determined from gravitational lensing. Figure taken from Ref. [62].

So then the question is no longer whether or not dark matter exists, but rather what *is* dark matter? In many ways the simplest assumption at the time was that dark matter was composed of existing compact objects that are much less luminous than regular stars but otherwise relatively similar, known as massive astrophysical compact halo objects (MACHOs). These included objects such as planets, brown dwarfs, red dwarfs, white dwarfs, neutron stars and black holes. Direct searches for MACHOs using gravitational microlensing however, found that less than 8% of the Milky Way dark matter halo could be composed of MACHOs with masses between $0.6 \times 10^{-7} M_{\odot} - 15 M_{\odot}$ [119, 170]. The discovery of the cosmic microwave background (CMB) introduced another piece of evidence in the dark matter puzzle. More specifically, it enabled precise measurements of the cosmological baryon density, as well as the total cosmological matter density. In the 2018 Planck Collaboration results, the baryon density was found to be $\Omega_b h^2 = 0.0224 \pm 0.0001$ while the total matter density was found to be $\Omega_m h^2 = 0.14240 \pm 0.00087$, meaning that baryons only make up $15.73 \pm 0.12\%$ of the matter in the Universe [142]. This is further indication that MACHOs are not the full solution, and that instead we must look to non-baryonic dark matter candidates to account for the remaining $\sim 84\%$ of matter that is not composed of baryons.

Non-baryonic particle dark matter candidates can be split into three general categories, “hot”, “warm”, and “cold”, depending on their velocities in the early Universe. Hot dark matter (HDM) particles are predicted to have emerged from the Big Bang with ultrarelativistic velocities. This causes them to have a damping mass on the scale of superclusters, meaning primordial density fluctuations smaller than this scale would be suppressed in the early Universe [42, 148]. Consequently, HDM is expected to follow a top-down model for large scale structure formation. That is, objects on the scale of superclusters form first, and only later collapse into smaller scale structures like galaxies. A common hot dark matter particle candidate is the standard model neutrino. However, astrophysical observations indicate that supercluster collapse occurred more recently than galaxies began forming ($z_{\text{sc}} < z_{\text{galaxies}}$) which is inconsistent with the top-down model [175, 83, 109].

Warm dark matter (WDM) particles are predicted to have intermediate velocities and

corresponding damping mass on the scale of large galaxies [42, 148]. Therefore, primordial density fluctuations below this scale are suppressed and large galaxy sized structures are the first to form. Common WDM candidates include sterile neutrinos and gravitinos. Because of their intermediate velocities, WDM candidates have been predicted to follow a more complex model of large scale structure formation combining aspects of both top-down and bottom-up models [136]. While it is still an active area of study, astrophysical observations have placed tight constraints on what kind of WDM particle, if any, could be a dominant source of dark matter [110, 90, 132, 69].

Cold dark matter (CDM), particles are predicted to have very small velocities and, as a result, cosmologically negligible damping mass such that even small primordial density fluctuations are not suppressed [42, 148]. Therefore, CDM is expected to follow a bottom-up model for large scale structure formation where small over-densities are the first to collapse. These small objects then merge with other small objects becoming the sites for galaxy formation. These galaxies then merge into galaxy clusters, and so on. Common CDM particle candidates are axions (see Sec. 1.2.2) and weakly interacting massive particles (WIMPs). In their 1985 paper [67] Davis, Efstathiou, Frenk, and White performed the first numerical simulations of gravitational clustering in CDM-dominated universes to compare with observational data from the Center for Astrophysics Redshift Survey [68]. After comparing their simulations with observations, the authors concluded that [67]:

“...it is remarkable how many aspects of the observed galaxy distribution are reflected quite faithfully by the distribution of CDM [...]. This seems too good to be true, but perhaps it hints that we are at last approaching a correct resolution of the missing mass problem.”

This result bolstered the reputation of CDM as the model which most closely matched astrophysical data. Years later, precision measurements of the CMB with the Wilkinson Microwave Anisotropy Probe (WMAP) and Planck truly solidified the case for cold dark

matter. The authors of the 2009 paper detailing the results of the WMAP CMB data stated [81]:

“Considering a range of extended models, we continue to find that the standard Λ CDM model is consistently preferred by the data.”

Over the next decade the Planck collaboration made several measurements of the CMB, each one improving upon the last [141, 144]. In their most recent result [142], the authors found that:

“... the base- Λ CDM model provides a remarkably good fit to the Planck power spectra and lensing measurements, with no compelling evidence to favour any ... extended models [...]. The overall picture from Planck ... is one of remarkable consistency with the ... Λ CDM cosmology.”

Clearly, in the current day it is indisputable that dark matter exists in some form, and furthermore it seems exceptionally likely to be cold. Thus, in the remainder of this thesis I will consider non-relativistic velocity distributions an essential criteria for a viable dark matter candidate.

1.2 Axion Theory

1.2.1 Origin of the Axion

In order to understand the history of the axion, we must first discuss the strong CP problem. Fundamentally the strong CP problem is a question of why the neutron electric dipole moment (EDM) is so small. Quantum chromodynamics (QCD) is the field theory that describes the strong interaction between the fundamental particles, namely quarks and gluons, that form composite subatomic particles like protons and neutrons. The development of QCD in the 1970s illuminated a problem with the strong interaction dubbed the $U(1)$ problem. In short, QCD predicts a global chiral $U(1)$ symmetry, which when spontaneously broken,

should produce a fourth pseudo-Goldstone boson in addition to the three pions [173, 60]. However, no such particle has been observed. The resolution to the $U(1)$ problem, detailed by t'Hooft in a 1976 paper [169], lies in the true vacuum structure of QCD, which turned out to be more complicated than originally predicted. The updated theory of QCD to account for this complicated vacuum structure included an additional contribution to the Lagrangian:

$$\mathcal{L}_\theta = \frac{g^2}{32\pi^2} \theta_{\text{QCD}} F_a^{\mu\nu} \tilde{F}_{\mu\nu a}. \quad (1.1)$$

Here, g is the QCD coupling constant, θ_{QCD} is a fundamental parameter related to the QCD ground state, $F_a^{\mu\nu}$ is the QCD field strength tensor, and $\tilde{F}_{\mu\nu a}$ is its dual. If one takes electroweak interactions into account, θ_{QCD} is changed by an amount equal to $\arg(\det(m_q))$, where m_q is the quark mass matrix [139]. The result of this is that the term above in the QCD Lagrangian should actually be written as:

$$\mathcal{L}_\theta = \frac{g^2}{32\pi^2} \bar{\theta} F_a^{\mu\nu} \tilde{F}_{\mu\nu a}, \quad (1.2)$$

where $\bar{\theta} = \theta_{\text{QCD}} + \arg(\det(m_q))$. This term in the QCD Lagrangian that was added to resolve the $U(1)$ problem predicts that the strong force violates charge-parity (CP) symmetry in such a way that one would expect to observe a neutron EDM magnitude on the order of $d_n \approx 10^{-16} \bar{\theta} e \cdot \text{cm}$ [31, 66]. Naturally, one might expect the effective QCD vacuum angle, $\bar{\theta}$ to be of order 1. However, as the sensitivity of neutron EDM experiments has increased since the 1950s limits on the magnitude of d_n have only gotten smaller and smaller, with the current best measured limit of $d_n = (0.0 \pm 1.1)10^{-26} e \cdot \text{cm}$ [12]. This means that $\bar{\theta} < 10^{-10}$ — an unusually small number! This result requires that either θ_{QCD} and $\arg(\det(m_q))$ are both extremely close to zero, or that they just so happen to perfectly cancel each other out leading to the observation that $\bar{\theta} \approx 0$. This apparent fine tuning of a fundamental parameter necessary to achieve the observed CP conservation despite explicit CP violation being expected is what is known as the strong CP problem.

There have been several proposed solutions to the strong CP problem since its discovery,

but one is particularly elegant and, for that reason, is considered a leading explanation. This solution was first proposed by physicists Roberto Peccei and Helen Quinn in a pair of seminal papers describing what is now known as the Peccei-Quinn (PQ) mechanism [138, 137]. In these papers the authors proposed the introduction of a new global chiral $U(1)$ symmetry, known as the PQ symmetry ($U(1)_{\text{PQ}}$), which is necessarily spontaneously broken in the early universe at some energy scale f_{PQ} , above which CP is not conserved. The PQ mechanism effectively promotes $\bar{\theta}$ from a static parameter to a dynamical complex scalar field, $\phi(x)$.




Figure 1.3: The classic wine bottle potential that leads to spontaneous symmetry breaking. The lowest energy state is no longer at the point where ϕ (some scalar field) is equal to zero, which has become a local maximum, but instead is described by some randomly chosen point along the “rim” of equivalent global minima. Figure taken from Ref. [23].

At energies below f_{PQ} but above the QCD energy scale, Λ_{QCD} , the field follows a classic wine bottle potential like that of the Higgs field:

$$V_{\text{PQ}}(\phi) \propto \left(|\phi|^2 - \frac{f_{\text{PQ}}^2}{2} \right)^2, \quad (1.3)$$

where ϕ is the new scalar field and f_{PQ} is the spontaneous symmetry breaking energy scale. See Fig. 1.3 for a visual example of this type of potential. At energies above f_{PQ} , the

vacuum expectation value of the scalar field is $\langle\phi\rangle = 0$, meaning all values of $\bar{\theta}$ are equally likely and strong CP violation is present. At energies below f_{PQ} , $V_{\text{PQ}}(0)$ becomes a local maximum and the field is forced into the new minima corresponding to the “rim” of the wine bottle. That is, the vacuum expectation value of the axion field becomes $\langle\phi\rangle = \frac{f_{\text{PQ}}}{\sqrt{2}}$. In the classical potential, $\bar{\theta}$ could still take on a multitude of values corresponding to the rim of equivalent global minima, which does not provide a satisfactory reason for why we see $\bar{\theta} \approx 0$ today. Luckily, as the Universe cools further to energies below Λ_{QCD} , the $U(1)_{\text{PQ}}$ symmetry is explicitly broken by the QCD anomaly. This essentially tips the wine bottle potential in one direction such that it is minimized at a single value of ϕ , and introduces another term to the QCD Lagrangian:

$$\mathcal{L}_{\text{PQ}} = \xi \frac{\phi}{f_{\text{PQ}}} \frac{g^2}{32\pi^2} F_a^{\mu\nu} \tilde{F}_{\mu\nu a}, \quad (1.4)$$

where ξ is a model dependent parameter, g is the QCD coupling constant, and $F_a^{\mu\nu}$ and $\tilde{F}_{\mu\nu a}$ are the QCD field strength tensor and its dual. So now the full QCD Lagrangian becomes:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{SM,axions}} + (\bar{\theta} + \xi \frac{\phi}{f_{\text{PQ}}}) \frac{g^2}{32\pi^2} F_a^{\mu\nu} \tilde{F}_{\mu\nu a}, \quad (1.5)$$

where the second term is the sum of the two CP violating terms of the QCD Lagrangian ($\mathcal{L}_\theta + \mathcal{L}_{\text{PQ}}$). Peccei and Quinn show that if at least one fermion flavor acquires its mass by coupling to ϕ then the absolute minimum of $V_{\text{PQ}}(\phi)$ occurs when the effective QCD vacuum angle is equal to zero in the ground state ($\bar{\theta} + \xi \frac{\phi}{f_{\text{PQ}}} = 0$) [138]. Thus, after the QCD phase transition the axion potential will naturally relax to a vacuum expectation value of $\langle\phi\rangle = -\frac{f_{\text{PQ}}}{\xi} \bar{\theta}$, dynamically canceling whatever initial value $\bar{\theta}$ takes on, causing the CP violating terms in Eq. 1.5 to vanish. Thus the PQ mechanism has provided a dynamic solution to the strong CP problem.

Soon after the PQ mechanism was proposed, Steven Weinberg and Frank Wilczek noted that the introduction of a spontaneously, and explicitly, broken $U(1)$ symmetry would imply

the existence of a very light, but not massless, pseudo-Goldstone boson which they dubbed the “axion” [174, 177]. As such, from here on out I will replace $\phi(x)$ with $a(x)$ (the axion field), and will replace f_{PQ} with f_a (the axion decay constant). The mass of the axion is related to the axion decay constant [71] via the following equation:

$$m_a \simeq 5.7 \mu\text{eV} \left(\frac{10^{12}\text{GeV}}{f_a} \right). \quad (1.6)$$

The first realistic model for the QCD axion is known as the PQWW model, named after Peccei, Quinn, Weinberg and Wilczek. This model predicted an axion decay constant on the order of the electroweak scale $f_a = v_F \simeq 250 \text{ GeV}$, corresponding to an axion mass $\mathcal{O}(10 - 100 \text{ keV})$. This led to a surge in experimental activity in the years following because this type of axion would be easily detectable with existing accelerator and particle experiments of the time. However, the PQWW axion model was quickly ruled out by laboratory searches [78, 82, 38, 183, 165, 18]. This led to the development of so-called “invisible” axion models with $f_a \gg v_F$. Because the axion mass is inversely proportion to f_a , these models describe axions that are very light, very weakly coupled to standard model particles, and very long lived. Their exceptionally weak couplings make them difficult to detect, hence their reputation as “invisible”. The invisible axion (weakly) couples to a handful of particles, but the most relevant for this thesis is the axion-photon coupling in which axions decay into two photons. The axion-photon coupling strength is defined by:

$$g_{a\gamma\gamma} = \frac{\alpha}{2\pi f_a} \left[\frac{E}{N} - 1.92(4) \right] = \frac{\alpha g_\gamma}{\pi f_a}, \quad (1.7)$$

where the first term is the model dependent contribution proportional to the ratio of the electromagnetic (E) and QCD (N) anomaly coefficients of the PQ symmetry, the second term is the model independent minimal coupling to QCD, α is the fine structure constant and $g_\gamma = \frac{1}{2} \left[\frac{E}{N} - 1.92(4) \right]$. The two benchmark models for invisible axions are the Kim-Shifman-Vainshtein-Zakharov (KSVZ) [113, 159] and the Dine-Fischler-Srednicki-Zhitnitsky (DFSZ) [185, 75] models. For the KSVZ model $\frac{E}{N} = 0$, while for the DFSZ model $\frac{E}{N} = \frac{8}{3}$.

Plugging these into Eq. 1.7, we can see that $g_{\gamma,\text{KSVZ}} \simeq -0.97$ and $g_{\gamma,\text{DFSZ}} \simeq 0.36$. As we will see later in this thesis, the axion-photon conversion power is proportional to g_γ^2 so the expected axion-photon conversion power from a KSVZ axion is ~ 7.3 times stronger than it is for a DFSZ axion.

1.2.2 Axions as Dark Matter

Now that we have established the particle physics origin of the QCD axion, let us examine why it makes a particularly compelling dark matter candidate. Firstly, let us recall the crucial criteria that any viable dark matter candidate must meet:

1. At most feebly interacting with standard model particles and electromagnetism
2. Stable on cosmological timescales
3. Non-baryonic
4. Produced in sufficient quantities in the early Universe to account for the observed dark matter density today
5. Moving at non-relativistic speeds in the early Universe (“cold”)

As I have discussed in the previous section, invisible axions inherently satisfy conditions 1, 2, and 3. I will now discuss an athermal method of axion production that ensures axions are both produced in large numbers in the early Universe, and that they are born cold. That is, this production method, known as the misalignment mechanism, satisfies both condition 4 and condition 5. The misalignment mechanism is essentially the PQ mechanism in a cosmological context. We discussed in the previous section that the PQ symmetry is spontaneously broken at some time in the early Universe (determined by the axion decay constant f_a), which causes the axion field to randomly single out a particular initial value of $\bar{\theta}$ which we will call $\bar{\theta}_i$. After this, the potential takes on the form of the wine bottle potential

(Eq. 1.3) where $\phi = a$ and $f_{\text{PQ}} = f_a$. At this stage, the axion appears as a massless degree of freedom, a true Goldstone boson. As the Universe cools to a temperature around Λ_{QCD} , the PQ symmetry is explicitly broken singling out a true global minimum of the potential, $\langle a \rangle = -f_a \bar{\theta}$, such that the effective vacuum angle is equal to zero. Note that $\xi = 1$ for both the KSVZ and DFSZ models so I have dropped it from the definition of the global minimum for the axion field [139]. The energy density stored in the axion field due to the misalignment between $\bar{\theta}_i$, which can take on any value between $\pm\pi$, and the CP conserving value of $\bar{\theta} = -\frac{\langle a \rangle}{f_a}$ causes the axion to roll down the rim of the wine bottle and oscillate around the new global minimum resulting in the athermal production of very light *pseudo*-Goldstone bosons, namely axions. Fig. 1.4 shows an illustration of this process.




Figure 1.4: Illustration of the evolution of the axion field potential over time. In the early Universe $T > f_a$, the PQ symmetry is unbroken, and $\bar{\theta}$ can take on any value. As the Universe cools to $T < f_a$, the PQ symmetry is spontaneously broken, and $\bar{\theta}$ takes on a random initial value, $\bar{\theta}_i$ between $\pm\pi$. At this point, the axion is a true, massless Goldstone boson. Finally, as the temperature of the Universe drops below Λ_{QCD} , the potential tilts, forcing the field into its CP conserving minimum at $\bar{\theta} = -\frac{\langle a \rangle}{f_a}$. Note that in this figure, $\theta = \bar{\theta} + \frac{a}{f_a}$ is the effective vacuum angle, which is why it shows θ rolling to 0 in the right most potential. Figure taken from Ref. [134].

Now that we have a method of production that gives rise to non-relativistic axions in the early Universe, the question is how many axions can be produced in this way and is it enough to account for a significant portion of dark matter? It has been shown since the early 1980s that axions with decay constants $10^8 \text{ GeV} \lesssim f_a \lesssim 10^{12} \text{ GeV}$ can indeed account for the entire observed dark matter abundance [147, 11, 74]. However, the answer to the first half of this question really depends on two things: whether the PQ symmetry is broken before or after inflation, and what the axion mass is. Therefore let us instead ask the question: what axion mass would produce the correct abundance of dark matter today? Before we can do that however, we must address the first dependence. The relationship between the axion density today ($\Omega_a h^2$) and the initial misalignment angle is $\Omega_a h^2 \propto \bar{\theta}_i^2$. The choice of initial misalignment angle is where inflation comes into play.

If PQ symmetry is broken after inflation, a random ensemble of $\bar{\theta}_i$ values populates the observable universe, so it is safe to take an average of all possible values to find a single representative estimate for $\bar{\theta}_i$ when calculating $\Omega_a h^2$. In the post-inflation scenario where the average initial misalignment angle is $\bar{\theta}_i \approx 2.155$, invisible axions with masses of $\mathcal{O}(10 \mu\text{eV})$ could account for the entire observed dark matter density today [134]. However, this is not the full story for the post-inflation scenario. If PQ symmetry is broken post-inflation there are topological defects that arise during the phase transition that can have effects on the final dark matter abundance [37]. To elaborate, strings and domain walls that form during this time decay and annihilate producing axions as a result. Simulations and numerical approaches indicate a wide range of scenarios in which string and domain decay contribute either less, similar to, or more to the final axion abundance than the misalignment mechanism [52, 93, 156, 112, 39, 92]. Usually larger string contributions lead to higher masses, and as such these different scenarios point to axions having a wide range of possible masses, some as high as meV and some as low as 10s of μeV . Having said that, given the complexity of axion string network evolution, there are still considerable error bars on any predictions. Regardless, a discovery in the μeV range would not be surprising, and would encode information about inflation scale and other pre-CMB universe parameters.

On the other hand, if PQ symmetry is broken before inflation, the random ensemble of $\bar{\theta}_i$ values gets smoothed out during inflation, and taking an average is no longer appropriate as there would be a single value of $\bar{\theta}_i$ filling the horizon. Thus, we cannot make a solid prediction about what the initial misalignment angle should be — it could have been any value between $\pm\pi$! Using conventional $\mathcal{O}(1)$ numbers for $\bar{\theta}_i$ leads to an axion mass $\mathcal{O}(1 - 100 \mu\text{eV})$ [101], but there is nothing prohibiting us from living in a Universe where $\bar{\theta}_i \ll 1$, meaning that even axions with very small masses (down to below a neV) would not overproduce dark matter. However, this would require fine tuning of the initial misalignment angle, and for that reason it is often not considered as likely as a $\mathcal{O}(1)$ value, despite it being possible. On the other side of the spectrum, axions with masses larger than roughly $100 \mu\text{eV}$ would not be produced in large enough quantities to account for all of the dark matter [134]. Therefore, $1 - 100 \mu\text{eV}$ could be considered the “natural” range for experiments to focus their search efforts, and indeed many have done just that.

1.3 Axion Searches

As discussed in Sec. 1.2, the original PQWW axion was quickly ruled out by experiments due to its relatively strong predicted coupling to standard model particles and electromagnetism making it accessible in existing reactor and accelerator experiments. The negative results of these experiments indicated the PQ symmetry must have been broken well above the electroweak scale ($f_a \gg v_F$). Axions with large f_a would be very light, very weakly coupled, and very long lived. Therefore, these types of axions were deemed “invisible”, i.e. impossible to detect with existing technology, and experimental search efforts slowly down greatly. However, I will discuss in the following section how searching for the invisible axion is indeed possible, contrary to some physicists’ original pessimism.

1.3.1 Axion Haloscopes

In his 1983 paper Pierre Sikivie revolutionized the search for the QCD axion dark matter by proposing an experimental design known as the axion haloscope [161]. As we discussed in

Sec. 1.2.1, axions naturally decay into two photons (see Fig. 1.5).




Figure 1.5: Feynman diagram of axion decay into photons. Left: Axion to photon conversion in vacuum. Right: Axion to photon conversion in a static magnetic field (\mathbf{B}_0). Figure taken from Ref. [168].

Sikivie showed that the interaction of axions and photons can be described by [162]:

$$\mathcal{L}_{a\gamma\gamma} = -g_{a\gamma\gamma} a \mathbf{E} \cdot \mathbf{B}, \quad (1.8)$$

where $g_{a\gamma\gamma}$ is the model-dependent axion-photon coupling constant as described in Eq. 1.7, a is the axion field, and \mathbf{E} and \mathbf{B} are the electric and magnetic fields of the two photons. However, axions with masses $\mathcal{O}(1 - 100 \mu\text{eV})$ have lifetimes on the order of $10^{54} - 10^{44}$ s [40] which are vastly greater than the age of the Universe ($\sim 10^{17}$ s). Luckily, this interaction can be sped up by making use of a large static magnetic field. When axions are immersed in an external, static magnetic field, the interaction in Eq. 1.8 is realized with a virtual photon in the magnetic field taking the place of one of the two photons in the normal interaction (see Fig. 1.5). Effectively, this replaces \mathbf{B} in Eq. 1.8 with the external magnetic field \mathbf{B}_0 . Thus, the decay rate can be increased by increasing the strength of \mathbf{B}_0 . Because there is now only one photon produced in this interaction, and the axion's velocity is negligible, the energy of the resultant photon is equal to the rest mass of the axion (i.e. $hf = m_a$). Furthermore, Sikivie showed that galactic halo axions could couple to resonant modes of

a right-cylindrical electromagnetic cavity provided the cavity is much smaller than the de Broglie wavelength of the halo axions [162]. This effectively replaces \mathbf{E} in Eq. 1.8 with the electric field of a particular cavity mode, \mathbf{E}_{nlm} . Thus, the energy of the photons produced by the interaction between the axion field and the external magnetic field is deposited into a resonant mode of the cavity. The resultant conversion power is then enhanced when the photon frequency (f) is equal to the resonant frequency of the cavity mode (f_0), after which the signal can be extracted by an antenna and further amplified by low-noise electronics. A simple diagram of a standard axion haloscope can be seen in Fig. 1.6. In order to make use of this resonant enhancement over the broad range of possible axion masses, it is also important that haloscopes have some type of resonant frequency tuning mechanism. A more detailed discussion of which cavity mode is preferred for standard axion haloscopes, an example of a frequency tuning mechanism, and other experimental details of a typical axion haloscope are discussed further in Sec. 2.2.

This idea was quickly adopted by groups at Brookhaven National Lab (BNL) [70, 179] and the University of Florida (UF) [98] who developed small cavity experiments which, while not exceptionally sensitive, laid the groundwork for future generations of axion haloscopes.

The Axion Dark Matter eXperiment (ADMX) is the longest running haloscope experiment, with its initial operation beginning in the 1990s at Lawrence Livermore National Lab as a follow up to the UF experiment. During this initial run, it was able to achieve sensitivity to the KSVZ axion, but not the significantly more weakly coupled DFSZ axion [97]. Over the following two decades ADMX made improvements to reduce noise in the detector readout via the introduction of low-noise electronics [26] and eventually a dilution refrigerator [79], both of which were necessary for reaching the coveted DFSZ sensitivity, which ADMX finally achieved in 2018, making it the first axion haloscope to do so. Since then, ADMX has continued taking data at exceptionally high sensitivity over a wide range of possible axion masses. I will discuss the most recent data taking run with ADMX in great detail later in this thesis.




Figure 1.6: Simplified diagram of a Sikivie haloscope. Axions passing through a resonant cavity that is immersed in a static magnetic field convert into observable photons. These photons are then picked up by an antenna and amplified by an ultra-low-noise amplifier. Figure taken from Ref. [37].

While ADMX has been at the forefront of direct detection searches for the QCD axion for many years, it is by no means alone in its efforts. For example, the Haloscope at Yale Sensitive to Axion CDM (HAYSTAC) experiment is a smaller haloscope searching for higher frequency axions [49, 186]. As we will see in Chapter 2, the axion power as measured in a haloscope (Eq. 2.1) is proportional to the cavity volume. Therefore, reaching sensitivity to QCD axions with smaller cavities like HAYSTAC is quite a feat. Nonetheless, thanks to implementing vacuum squeezing in the receiver to reach below the standard quantum noise limit (see Chapter 3), HAYSTAC was able to achieve sensitivity to axion-photon couplings only a factor of 2 above KSVZ [30, 106]. These results made HAYSTAC the first group to successfully demonstrate that state-squeezing can be used to improve haloscope sensitivity.

If we extend our radius to include the international physics community, there are even more axion haloscope experiments to discuss. For example, there are a number of cavity experiments being done at the Center for Axion and Precision Physics Research (CAPP) in Korea [117]. CAPP joined ADMX in 2023 as one of two haloscopes in the world to achieve DFSZ sensitivity [182, 17]. The CAPP program has also reported results from a number of other R&D cavity experiments at couplings above DFSZ across a wide range of masses.

In addition to CAPP, there are a number of cavity experiments based out of Europe including the Relic Axion Dark matter Exploratory Set-up (RADES) experiment [189] and the QUaerere AXions (QUAX) experiment [32], both of which are searching for axions at higher masses ($\sim 30 - 50 \mu\text{eV}$). RADES is an axion haloscope housed in the the bore of the CERN Axion Solar Telescope (CAST) helioscope. The recent result from RADES [189] which reached 2 orders of magnitude below the CAST bound [22] is one of the most sensitive direct searches in this frequency range. The QUAX collaboration was originally designed to take advantage of the axion's coupling to electrons [65, 64], however they have also run searches as a standard haloscope and set limits on the axion-photon coupling [20, 21, 19, 72, 152]. These results demonstrated near-KSVZ sensitivity in a narrow band around $34 \mu\text{eV}$ placing them in a similar position as RADES in terms of achieving high levels of sensitivity for that mass range.

The last haloscope I will discuss is the Oscillating Resonant Group AxioN (ORGAN) experiment [126]. This haloscope is aimed at even higher frequencies than RADES or QUAX, specifically it is targeting axions with masses in the range $\sim 60 - 200 \mu\text{eV}$. While ORGAN's sensitivity is still many orders of magnitude away from the QCD axion band [150, 149], it serves as an example of how a haloscope can be used to probe high axion masses that were previously thought to be inaccessible.

Despite major advances in haloscope technology over the last few decades, there are still only two experiments that have been able to reach DFSZ coupling — a requirement for a complete axion search. Therefore, there is still much room for improvement in the realm of haloscopes. Some areas of active research and development include: increasing

magnetic field strength [95, 99], lowering noise temperatures [118, 184, 76, 180], as well as developing more complicated resonator designs such as multi-cavity systems [14], ultra high-Q cavities [145, 15, 16] and multi-wavelength resonators [56, 57, 36, 61, 120, 131, 51, 53, 24].

1.3.2 Other Direct Detection Searches for Axions

While traditional haloscopes remain uniquely sensitive to axion masses around $1 - 100 \mu\text{eV}$, there are significant challenges when moving to both lower and higher masses. Therefore it is worth discussing some more novel experimental techniques that are designed to overcome some of the challenges that traditional haloscopes face.

The challenge with moving to lower frequencies with traditional haloscopes is related to the cavity size. As I alluded to earlier, the volume of a cavity is inversely related to its resonant frequency ($V \propto f_0^{-3}$). Therefore, considering the ADMX cavity is ~ 1 m tall and ~ 0.5 m wide, one can imagine that going to lower frequencies will quickly lead to very large cavity sizes. Therefore, traditional cavity haloscopes are not best suited in this frequency range. Instead, some groups are opting to use LC circuits instead. This category of experiment makes use of the axion-induced oscillating effective current inside a static magnetic field. This oscillating current produces its own small magnetic field that can be amplified and detected using a cooled LC circuit and a very sensitive magnetometer [163]. The first experiment to successfully implement this technique for performing a low mass axion search is known as A Broadband/Resonant Approach to Cosmic Axion Detection with an Amplifying B-field Ring Apparatus (ABRACADABRA) [108]. As the name alludes to, ABRACADABRA consisted of a toroidal magnet as the external magnetic field source, resulting in an azimuthal induced current which produces a magnetic flux through a central pickup cylinder in the toroid bore. The benefit of using this magnet design is that the region where the pickup cylinder is placed is free from any external magnetic field. The experiment also has the benefit of being able to run in either broadband or resonant mode, however so far it has only been run in broadband mode [135, 157]. While this prototype experiment did not reach sensitivity to QCD axions, it served as an example of how this sort of apparatus

could be used to probe lower axion masses.

Building on the initial prototype of ABRACADABRA, the DM-Radio collaboration was formed in order to achieve sensitivity to QCD axions in the range of roughly 5 kHz - 200 MHz (~ 20 peV - $1 \mu\text{eV}$) using LC circuit based detectors [164, 77, 48]. The implementation of DM-Radio is planned to proceed in several stages. Firstly, DM-Radio 50L will be used to demonstrate this technology on a much larger scale than ABRACADABRA, accessing frequencies between 5 kHz - 5 MHz, but with sensitivity limited to orders of magnitude above the QCD axion band. The next phase of the experiment will be a larger, cubic meter volume apparatus named DM-Radio-m³ which will be used to probe axions between 10 MHz - 200 MHz. Finally, an even larger (roughly 10 m³) iteration of the experiment, DM-Radio-GUT, will be aimed at detecting QCD axions at GUT-motivated masses in the range 0.4-120 neV (100 kHz - 30 MHz).

In addition to LC circuit experiments, some groups have been pursuing a different broadband approach using dish antennas [103]. In particular the Broadband Reflector Experiment for Axion Detection (BREAD) collaboration has been at the forefront of this effort [123]. BREAD consists of a cylindrical metal barrel for converting axions to photons with a coaxial parabolic reflector used for focusing the signal photons onto a sensor located at the focal point. What sets BREAD apart from other dish antennas is that its geometry is optimized to fit inside standard cryostats and compact high-field solenoids. This broadband reflector design allows BREAD to search multiple decades of axion mass without tuning. The collaboration has recently reported limits for axion-photon coupling [104] that improve upon the CAST limits [22], but are still many orders of magnitude away from the QCD axion band, which may be reachable if necessary advancements in single-photon far infra-red and THz detector technology are made.

The experiments I have detailed here are by no means an exhaustive list of all existing axion searches. For those interested in learning more about the wide variety of searches for axions including other laboratory based experiments, as well as astrophysical and cosmological observations please see Refs. [14, 134].

To conclude this Chapter and set the stage for the remainder of this thesis let us take a look at Fig. 1.7. This plot is a simplified representation of the status of experimental searches, specifically it shows existing limits on the axion-photon coupling across many orders of magnitude in axion parameter space. As one can see, haloscopes are by far the most sensitive in the intermediate mass range ($1 - 100 \mu\text{eV}$), but they also result in the narrowest coverage due to their resonant nature. As such, there is still a massive amount of parameter space to be searched, and many more opportunities for the axion to be discovered.

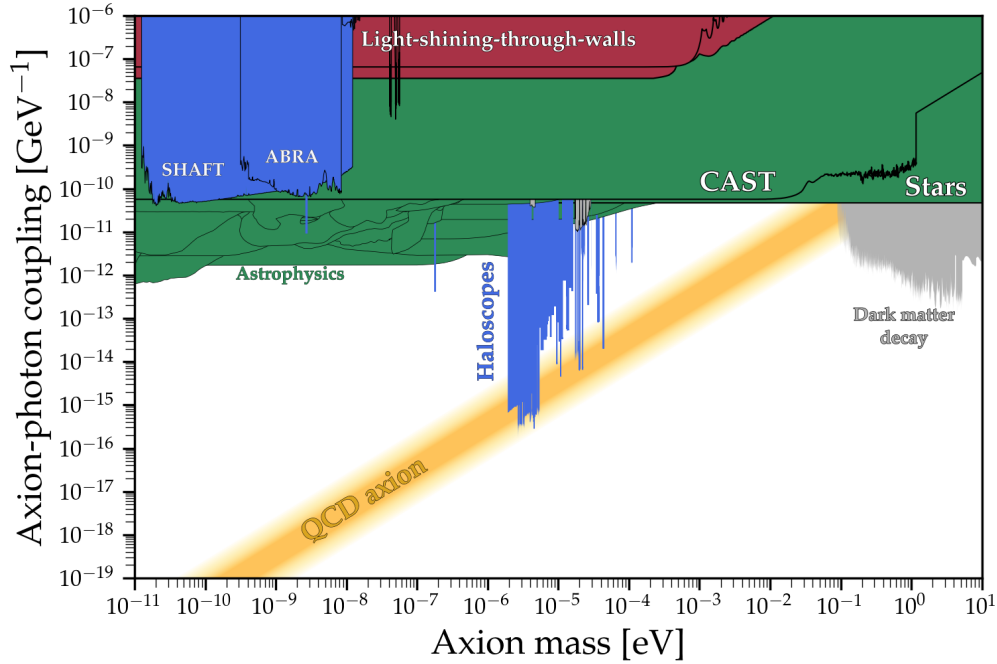


Figure 1.7: A simplified panorama of the axion-photon coupling parameter space. Constraints are roughly color coded, with direct searches for dark matter axions in the solar neighborhood shown in blue, searches for axions produced in astrophysical environments shown in green, and pure laboratory searches in red. The pure laboratory searches differ from direct detection experiments in that the pure laboratory searches do not make any assumptions about axions being dark matter, whereas the blue searches do. Figure taken from Ref. [133]

Chapter 2

RUN 1D HARDWARE

2.1 Axion Dark Matter eXperiment Overview

The Axion Dark Matter eXperiment is an axion haloscope currently based at the University of Washington in Seattle. A cutaway diagram of the ADMX insert inside the magnet bore can be seen in Fig. 2.1. Axions are converted to observable microwave photons as they pass through the experiment by the magnetic field produced by a ~ 7.6 T solenoidal magnet. The signal strength of the resultant photons are amplified when the photon frequency, f , matches the resonant frequency of the microwave cavity, f_0 . Given the large range of possible axion masses, the cavity must be tunable in order to make use of this effect. Thus, the cavity frequency is changed by altering the geometry inside the cavity with a single copper tuning rod. Power is extracted from the resonant cavity with an antenna, which is then amplified by a Josephson Parametric Amplifier (JPA). The JPA and other magnetic field sensitive components are housed in the field-free region of the insert. This region is kept field-free using a secondary magnet known as the bucking coil that produces a magnetic field opposite of that produced by the main magnet within this small region. The cavity and quantum amplifier package are kept at roughly 100 mK by the mixing chamber to greatly reduce thermal noise in the readout. The signal is then amplified by a Heterostructure Field Effect Transistor amplifier (HFET) before exiting to the room temperature receiver where it is further processed before the data is analyzed. To date, ADMX has excluded DFSZ (KSVZ) dark matter axions with masses between 2.66 - 3.34 and 3.9 - 4.1 μeV (1.9 - 4.2 and 4.54 - 5.41 μeV) [25, 26, 79, 46, 33, 91, 63].




Figure 2.1: Cutaway diagram of the ADMX insert inside the magnet bore. The insert, which contains the microwave cavity, as well as the cryogenic receiver chain and dilution refrigerator, is lowered into the magnet bore for data taking. Power is extracted from the cavity with an antenna and then amplified by a Josephson Parametric Amplifier (JPA) that lives in the field-free region marked with dashed red lines. The power is then amplified a second time by an HFET amplifier, before exiting to the room temperature receiver where it is processed further. Note that this diagram is slightly outdated as it shows the cavity as having two tuning rods. The Run 1D cavity only has one, larger tuning rod, but otherwise the diagram is representative of the current design.

The expected power coupled out of the ADMX haloscope is

$$P_a = 2.8 \times 10^{-23} \text{ W} \left(\frac{V}{106 \ell} \right) \left(\frac{B}{7.6 \text{ T}} \right)^2 \left(\frac{C_{010}}{0.4} \right) \left(\frac{g_\gamma}{0.36} \right)^2 \left(\frac{\rho}{0.45 \text{ GeV/cc}} \right) \left(\frac{f}{1.2 \text{ GHz}} \right) \left(\frac{\beta}{1 + \beta} \right) \left(\frac{Q_L}{30,000} \right) \left(\frac{1}{1 + (2\delta f / \Delta f)^2} \right), \quad (2.1)$$

where V is the effective volume of the cavity, B is the magnitude of the external magnetic field, C_{010} is the overlap between the cavity resonant mode and the magnetic field (also known as the cavity form factor), g_γ is the model-dependent constant that determines the axion-photon coupling ($g_{\gamma,\text{KSVZ}} = -0.97$ $g_{\gamma,\text{DFSZ}} = 0.36$), ρ is the local dark matter density, β is the coupling strength of the antenna to the cavity mode, Q_L is the loaded quality factor of the cavity ($Q_L = Q_0/(1 + \beta)$, where Q_0 is the unloaded quality factor), δf is the frequency offset between the cavity resonance (f_0) and axion signal frequency (f), and $\Delta f = \frac{f_0}{Q_L}$ is the cavity linewidth. Here, the equation has been normalized to represent typical experimental parameters for ADMX.

2.1.1 Dilution Refrigerator

Thermal noise is a primary contribution to the overall system noise temperature and can significantly limit the level of sensitivity ADMX is able to achieve. As such, keeping the system as cold as possible is paramount to performing a QCD-axion sensitive search. ADMX does this by employing a custom built dilution refrigerator (DR) developed at Janis Research Company (now part of Lake Shore Cryotronics) with the help of ADMX collaborators at the University of Florida. The DR works by manipulating the unique properties of two different isotopes of Helium: He-3 which is a fermion, and He-4 which is a boson. Because He-4 is able to undergo Bose-Einstein condensation, it can become a superfluid much more easily than He-3. In fact, He-4 becomes a superfluid at 2.17 K, whereas He-3 requires an extremely low temperature of only 2.5 mK, which is not achieved in the of operation of our DR. The superfluid He-4 and normal fluid He-3 are mixed inside of a pot called the mixing chamber.

When the mixture of these two isotopes is cooled below approximately 0.87 K, the mixture undergoes a spontaneous phase separation into a He-3 rich (concentrated) phase, and a He-3 poor (diluted) phase. As shown in Fig. 2.2, at very low temperatures the concentrated phase is almost pure He-3, and the diluted phase is about 6.6% He-3 and 93.4% He-4. Removing He-3 atoms from the diluted phase will cause He-3 atoms from concentrated phase to cross the phase-boundary. The enthalpy of He-3 in the diluted phase is larger than in the concentrated phase, thus the transition of He-3 from the concentrated to diluted phase in the mixing chamber is endothermic — meaning it removes energy from its environment. In this case, because the environment is well isolated, the temperature in the mixing chamber will go down. In order to continue this cooling process, the He-3 must be continuously circulated in and out of the mixing chamber. This is achieved by connecting the mixing chamber to a distiller (also known as the still) via a small stainless steel capillary. The diluted phase makes it way up to the still, where it is then heated to roughly 650-800 mK. This causes the He-3, which has a much higher vapor pressure than He-4 at these temperatures, to enter its gaseous phase while the majority of the He-4 remains a liquid. The mostly He-3 gas is then pumped out of the still and purified before being recirculated back into the mixing chamber, where it continues the cooling process.

Other components in the system are used to provide pre-cooling for the DR as well as establish other thermal stages in the insert. To begin, there is a large reservoir filled with liquid He-4 that is refilled periodically to replace He-4 that has boiled off. This reservoir supplies two 1K pots, small copper containers partially filled with liquid He-4. Pumping on the gaseous He-4 sitting above the liquid He-4 in the pots cools them from ~ 4 K to ~ 1 K (hence the name 1K pot). The small 1K pot is used to pre-cool the He-3 mixture before it enters the DR. The big 1K pot on the other hand is thermally sunk to a plate in the insert to establish a 1K thermal stage (the “1K plate”) that various components can be mounted to. Components in the insert that need to be cold, but not milliKelvin cold, are often mounted to the reservoir or the 1K plate since it can handle a much greater thermal load than the mixing chamber, which is reserved for only the coldest insert components.

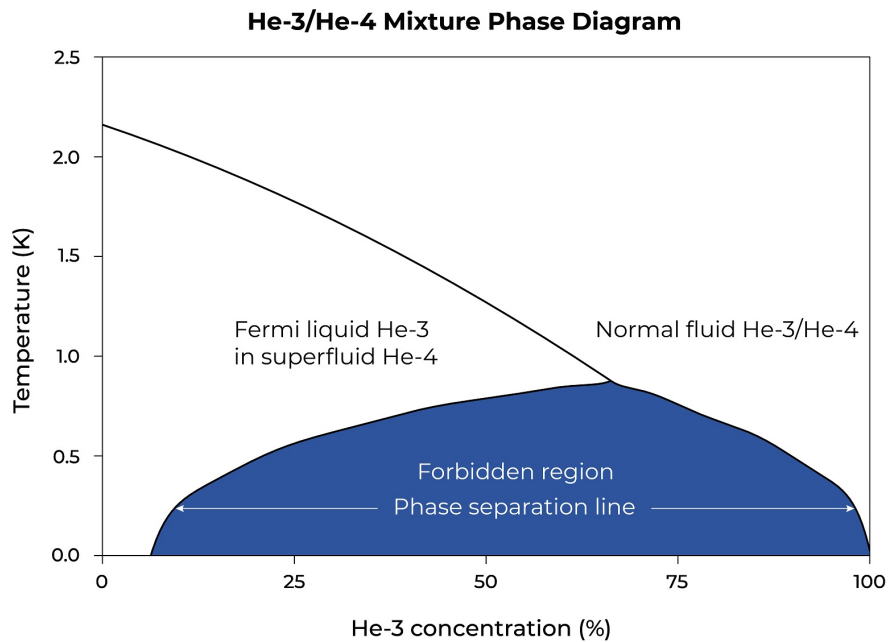


Figure 2.2: Helium isotope phase transition diagram from Bluefors [6]. As the mixture of He-3 and He-4 is cooled below about 870 mK, it undergoes a spontaneous phase transition, splitting into a He-3 rich (concentrated) phase and a He-3 poor (diluted) phase.

2.2 Cavity

Through Run 1A - Run 1C ADMX used the same copper-plated stainless steel right-circular cylindrical cavity with a volume of 136 l (before subtracting tuning rod volume) and a length of ~ 1 m. For Run 1A and 1B the cavity was equipped with two 0.05 m diameter copper tuning rods that run the length of the cavity which could be moved from near the walls of the cavity to near the center in order to adjust the resonant frequency over a range of ~ 580 – 890 MHz. Run 1C also used two copper tuning rods, but they were slightly larger with a diameter of 0.114 m allowing the experiment to cover a frequency range of ~ 0.8 – 1 GHz. For Run 1D, however, the desired frequency range of 1-1.4 GHz was outside of the available tuning range of the existing cavity. As a result, the microwave cavity was replaced prior to beginning Run 1D operations. The new design employs a 139 l (106.8 l effective volume after subtracting

tuning rod volume) copper-plated stainless steel right-circular cylindrical cavity that is ~ 1 m in length, but now contains only one, larger, tuning rod with a diameter of 0.203 m. The fabrication and testing of the Run 1D cavity was done by ADMX collaborators at Lawrence Livermore National Laboratory (LLNL). Photos of all the different tuning rods are shown in Fig. 2.3.




Figure 2.3: A view inside the ADMX cavities for Runs 1A-1D, showing the differences in tuning rod size and quantity over time.

Installation of the new cavity was done in June 2023 at the UW Center for Experimental Nuclear Physics and Astrophysics (CENPA) where ADMX is currently housed. The process involved removing the top plate of the cavity, lifting and inserting the new tuning rod, attaching the sapphire axles to both sides of the rod which allow it to move smoothly inside the cavity, securely sealing both end plates, and after some initial tests, mounting the entire cavity and rod system to the 1K plate using carbon fiber struts. A sample of photos documenting this procedure can be seen in Fig. 2.4.

A mode map of the Run 1D cavity showing its resonant frequency as the rod is tuned is shown in Fig. 2.5. A mode map is created by taking wide transmission measurements that simultaneously measure the resonant frequencies of multiple modes as the rods are moved. The process of taking transmission measurements is described further in Sec. 2.3, and an example of a wide transmission measurement can be seen in Fig. 2.8. Of particular interest is the lowest order, TM_{010} , mode because it is the mode that maximizes the observed

axion signal strength for reasons discussed later in this section. The new cavity and rod configuration can couple to the TM_{010} mode over a frequency range of $\sim 0.9 - 1.4$ GHz.




Figure 2.4: A series of photos showing the installation of the rod into the Run 1D cavity. A) Inside of the Run 1D cavity B) Lifting the rod to insert the bottom sapphire axle before installation C) Former ADMX undergraduate Peter Park and current LLNL staff scientist Nick Du are shown preparing to insert the rod into the cavity D) The rod and top sapphire axle after being installed into the cavity

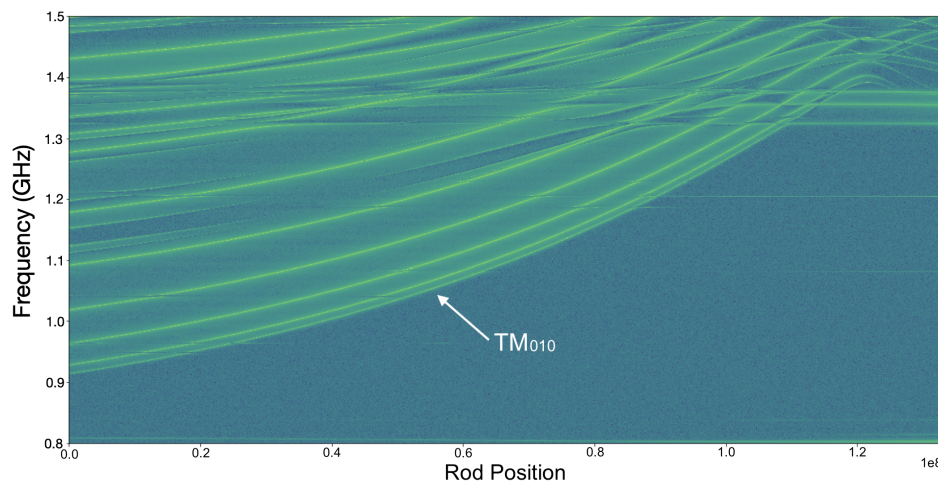


Figure 2.5: A mode map showing the resonant frequencies of various cavity modes changing with rod position. The mode of interest for ADMX, TM_{010} , can be tuned over a range of roughly $0.9 - 1.4$ GHz.

There is a strongly coupled antenna port at the top of the cavity and a weakly coupled antenna port at the bottom of the cavity. The strong antenna is used for extracting axion signal power from the cavity during data-taking. It is connected to a gearbox that allows operators to adjust the insertion depth thus varying the coupling strength to the cavity resonant mode. The weak antenna is set at a fixed, very weak, coupling strength. As I will discuss further in Sec. 2.3, the weak antenna is used for characterizing various cavity parameters.

The quality factor, Q , of a resonant cavity can be defined as $Q = f_0/\Delta f$, where f_0 is the resonant frequency of the cavity and Δf is the full width at half maximum (FWHM) of the cavity resonance. Q can be thought of as a measure of the sharpness of the resonance lineshape because a smaller FWHM corresponds to a narrower lineshape and higher Q . It is also commonly defined as the ratio of the energy stored in the resonator to the energy dissipated per cycle due to lossy elements in the resonator. That is to say, free oscillations will oscillate longer in a high- Q resonator than a low- Q one because a high- Q resonator will hold onto its stored energy longer.

The above discussion assumes a closed cavity system, but in reality an antenna must be inserted in the cavity in order to extract power, which acts as another source of power dissipation. That is to say, an antenna that is strongly coupled to the cavity will draw more power into the antenna than a weakly coupled antenna which will decrease the quality factor because of the increase in energy dissipated per oscillation. This decrease in Q referred to as the “loaded” quality factor (Q_L):

$$\frac{1}{Q_L} = \frac{1}{Q_0} + \frac{1}{Q_{\text{ant}}}, \quad (2.2)$$

where Q_0 is the unloaded quality factor and Q_{ant} is the associated Q of the antenna in the cavity. Using this equation, we can define the cavity coupling strength:

$$\beta = \frac{Q_0}{Q_{\text{ant}}}. \quad (2.3)$$

The relationship between the cavity coupling strength and the unloaded and loaded quality factors is therefore described by the following expression:

$$\beta = \left(\frac{Q_0}{Q_L} - 1 \right). \quad (2.4)$$

This relationship demonstrates that the larger β is, the more Q_L is decreased compared to Q_0 . There are three different regimes the coupling can fall under: under coupled, critically coupled, and over coupled. The first regime corresponds to $\beta < 1$. In this regime, the effect of the antenna on the cavity quality factor is small, as is the amount of power extracted by the antenna. When $\beta = 1$, the antenna is critically coupled. In this scenario, half of the power inside the cavity is dissipated in the walls of the resonator, and the other half is extracted by the antenna. Finally, when $\beta > 1$ the antenna is over coupled. In this regime, more power is extracted by the antenna than is lost to the walls of the cavity. This regime also sees the greatest reduction in Q_L compared with the intrinsic Q_0 .

Looking back to Eq. 2.1, the observed axion power is proportional to $\beta/(1 + \beta)$. This is because the total fraction of axion power exiting the cavity through the antenna is equal to $\kappa = \frac{1/Q_{\text{ant}}}{1/Q_L} = \beta/(1 + \beta)$. Therefore, one may assume that being highly over coupled is optimal since increasing β increases the overall power fraction exiting the cavity, κ . However, in addition to the factor of $\beta/(1 + \beta)$ from the antenna coupling, the power deposited into the cavity is dependent on Q_L , which contains another factor of $1/(1 + \beta)$ for a total factor of $\beta/(1 + \beta)^2$. Therefore, a value of $\beta = 1$ (critical coupling) will maximize the magnitude of the axion signal power extracted from the cavity (see Fig. 2.6). However, by increasing the cavity coupling which decreases the loaded quality factor, the effective search bandwidth of the cavity goes up which increases the rate at which the haloscope can scan over different frequencies. The combination of these two effects are summarized in the axion haloscope scan rate as defined in Eq. 4.1.1, where $df/dt \propto \beta^2/(1 + \beta)^3$. Given that there is a massive amount of axion parameter space that needs to be scanned over, the scan rate is the truly critical figure of merit for an axion haloscope. Differentiating this quantity to find the cavity coupling strength that maximizes the scan speed gives a value of $\beta = 2$ (see Fig. 2.6).

Therefore, ADMX aims to operate the experiment at a slightly over coupled value around $\beta = 2$ to maximize the axion search efficiency.




Figure 2.6: The axion power extracted by the antenna is proportional to $\beta/(1+\beta)^2$, as shown in purple. The axion haloscope scan-rate on the other hand is proportional to $\beta^2/(1+\beta)^3$, as shown in green. There is a dotted line at $\beta = 1$ to mark where the axion power extracted by the antenna is maximized with respect to β , which corresponds to critical coupling. There is a dashed line at $\beta = 2$ to mark where the axion haloscope scan rate is maximized with respect to β , which turns out to be slightly over coupled.

Another important cavity parameter is the cavity form factor, C_{lmn} , which is a measure of the overlap between the electric field of the cavity resonant mode and the external magnetic field. More specifically,

$$C_{lmn} = \frac{(\int dV E_{lmn} \cdot B_{\text{ext}})^2}{V B_{\text{ext}}^2 \int dV \epsilon_r E_{lmn}^2}, \quad (2.5)$$

where E_{lmn} is the electric field of the cavity mode, B_{ext} is the external magnetic field, and ϵ_r is the relative permittivity within the cavity. As one can see, the form factor (which is linearly related to the observed axion power in Eq. 2.1) is maximized when the dot product of the electric field of the cavity mode and the external magnetic field is maximized. Since the solenoidal magnet in ADMX produces a magnetic field that runs along the axis of the cavity, the cavity mode used to search for axions should have an electric field that also run along the axis of the cavity. This requires we use transverse-magnetic (TM) cavity modes which only have electric fields in the axial direction, meaning there is no magnetic field from the cavity mode in the axial direction. Specifically, we are interested in using TM_{0n0} modes, which have electric fields that are azimuthally symmetric ($l=0$) and constant in the longitudinal direction ($m = 0$). In particular, the TM_{010} mode, which has a peak electric field magnitude at the center of the cavity and decreases to zero at the walls of the cavity, has the greatest overlap with the external magnetic field and the largest form factor. Therefore, this is the mode used for performing an axion search in ADMX.

2.3 Cryogenic Receiver

Several upgrades were made to the cryogenic receiver chain for Run 1D, a diagram of which can be seen in Fig. 2.7. Before discussing the upgrades that were made, I will first introduce the various components and lines.

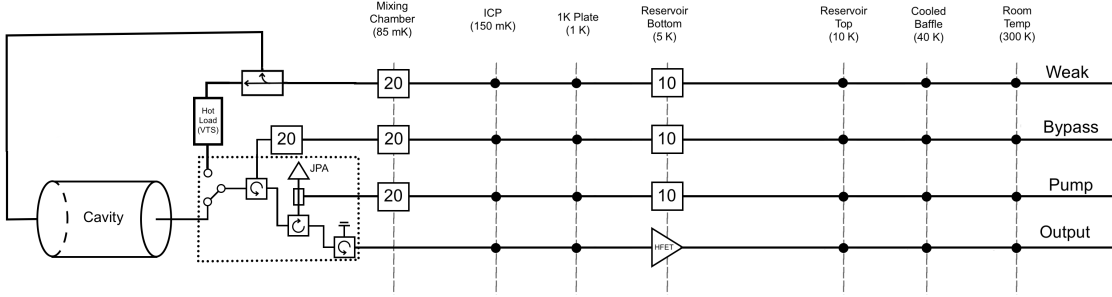


Figure 2.7: The cryogenic receiver chain for Run 1D. There are four different lines shown. The weak and bypass lines are primarily used for characterization of cavity parameters, while the pump line provides the input pump tone for parametric amplification with the Josephson Parametric Amplifier (JPA). Lastly, the output line is used to amplify and readout signals from the cavity. ADMX uses two amplifiers, a JPA and an HFET amplifier. The three small squares with circular arrows in them are unidirectional devices known as circulators which are used to direct the signal path and prevent downstream noise from reaching the cavity. There are 10 and 20 dB attenuators on weak, bypass and pump lines at the reservoir bottom and mixing chamber stages respectively. The black circles represent 0 dB “attenuators” used as thermal sinks. Connected to the weak line via a directional coupler is the hot load, which is a variable temperature stage used for noise calibration. It can be connected to the output line using an RF switch.

2.3.1 Weak Line

Starting from the top of Fig. 2.7, is the weak line. This line connects to the bottom of the resonant cavity via a weakly coupled antenna. It is primarily used for transmission (S_{21}) measurements in which a swept RF signal from the vector network analyzer (VNA) is routed through the weak line, up through the cavity and output line and then back to the VNA. Standard transmission measurements are used to calibrate various experimental parameters such as Q and receiver gain, and wide transmission measurements are used for tracking the TM_{010} mode as well as making mode maps (Fig. 2.5). Examples of both a standard and

wide transmission measurement are shown in Fig. 2.8.




Figure 2.8: Top: standard transmission measurement of the TM_{010} cavity resonance. This type of transmission measurement is used to calibrate various experimental parameters such as the quality factor and receiver gain. Bottom: wide transmission measurement of many different cavity modes. This type of transmission measurement is used for making mode maps and tracking the TM_{010} mode.

In addition to performing transmission measurements, the weak line is used to carry out blind injections with our synthetic axion generator (SAG) which are used as a way to check our detection capabilities. The SAG is controlled by ADMX collaborators at Pacific

Northwest National Lab and are only unblinded after we have completed the candidate identification and exclusion process. I will discuss the blind synthetic axion injections further in Chapter 5.

2.3.2 Bypass Line

Next is the bypass line, which is used to perform reflection measurements in which a swept RF signal from the VNA is routed through the bypass line, reflected off the cavity before going up the output line back to the VNA. In reality the path of the signal in this scenario is that of an S_{21} measurement rather than an S_{11} as one might expect, but we still refer to these as reflection measurements because they measure power that is reflected off the cavity. These measurements can be used to calibrate experimental parameters like the Q , cavity coupling and receiver gain. An example of a reflection measurement is shown in Fig. 2.9.




Figure 2.9: Reflection measurement of the TM_{010} cavity resonance. This measurement is used to calibrate various experimental parameters such as the quality factor, receiver gain, and cavity coupling strength.

2.3.3 Pump Line

Next is the pump line, which is connected to a signal generator that provides the input pump tone necessary for amplification with the Josephson Parametric Amplifier (JPA). To explain further, JPAs are ultra low-noise amplifiers based on the Josephson junction [107]. A JPA designed by Irfan Siddiqi's group at the University of California Berkeley [160, 100, 45] was used as ADMX's first stage amplifier for Run 1D because JPAs have excellent noise performance and high-gain. A JPA consists of a SQUID loop (or an array of SQUID loops) in parallel with a geometric capacitance (see Fig. 2.10). A DC SQUID is comprised of two Josephson junctions connected in parallel in a superconducting loop. The resonant frequency of a JPA is given by:

$$\omega_0 = \frac{1}{\sqrt{L_J + L_G C}}, \quad (2.6)$$

where L_J is the total Josephson inductance (a form of kinetic inductance) of the SQUID loop, L_G is the geometric inductance of the equivalent circuit, and C is the parallel capacitance.

The inductance of a single Josephson junction is modeled as

$$L_J = \frac{L_{J0}}{\sqrt{1 - (I/I_0)^2}} = L_{J0} \left(1 + \frac{1}{2}(I/I_0)^2 + \dots\right), \quad (2.7)$$

where I is the current through the junction, I_0 is the critical current of the junction, $L_{J0} = \Phi_0/2\pi I_0$ is the Josephson inductance in the absence of any supercurrent flow, and Φ_0 is the superconducting magnetic flux quantum. The critical current is the maximum dissipationless supercurrent that can flow through the Josephson junction, and it is dependent on the geometry of the junction and on the superconducting material used. Two Josephson junctions in a SQUID loop take on the properties of a single Josephson junction but with a magnetic flux-tunable critical current. Using the fact that I_0 is flux-tunable in a SQUID loop and that L_J depends on I_0 , the resonant frequency ω_0 (which depends on L_J) can be tuned by applying an externally applied magnetic flux through the SQUID loop.




Figure 2.10: Figure taken from Ref. [111]. A JPA circuit diagram. The region outlined in purple contains an array of SQUID loops (right hand side) that are in series with a geometric inductance and shunted by a capacitor. The JPA can either be current-pumped or flux-pumped. In current-pumped operation (as was used for Run 1D), the pump tone at frequency ω_P enters the circuit via the main transmission line shown (red). In flux-pumped operation, the pump tone enters via the inductively coupled DC line (blue). In either case, the input signal a_{in} enters the main transmission line, mixes with the pump tone in the JPA, and then is reflected back down the transmission line and exits the JPA from the same port it entered in.

A JPA amplifies the signal using the basic concept of parametric amplification: gain is achieved by varying a parameter of the system harmonically in time [166]. The variation of system parameters at the pump frequency (ω_P) causes some of the energy from the pump mode to be transferred to the signal mode (ω_S) and the remaining energy is transferred to another mode, called the idler (ω_I).

For Run 1D, ADMX used a four-wave mixing (or current-pumped) JPA, which takes two pump photons and converts them into one signal photon and one idler photon: $2\omega_P = \omega_S + \omega_I$. Under current-pumped operation, the pump tone (shown in red in Fig. 2.10) and signal tone are both sent through the main RF transmission line. Here, the bias current through the Josephson junctions in the SQUID loop is the parameter that is being varied at the resonant frequency of the device to provide parametric amplification. Additionally, Run 1D JPA was operated in phase-insensitive mode, meaning that the phase between the pump and

signal tones is random. Under phase-insensitive operation, both quadratures of the readout signal are amplified by a factor of \sqrt{G} where G is the gain of the JPA. Phase-insensitive amplification adds at least half a photon of noise to the readout signal from both the signal and idler modes, as required by quantum mechanics [55].

The JPA is optimized using three parameters, the DC bias current, the power of the pump tone, and the frequency of the pump tone. It is optimal to have the pump frequency close to the signal frequency for maximum amplification, however in order to protect against interference between the signal and the pump there needs to be a slight offset. For Run 1D, the JPA pump frequency was consistently set to $\omega_P = \omega_S + 320$ kHz to ensure both the pump and idler modes would be many Q-widths away from the cavity resonance, which had a bandwidth of ~ 56 kHz. So, in practice, the DC bias current and pump power were the two parameters used for optimizing the gain during data taking.

Finally, there is the output line which is used to extract signals from the cavity via a strongly coupled antenna, as well as amplify signals with a JPA from Siddiqi's group at Berkeley and HFET amplifier from Low Noise Factory, model (LNF-LNC0.6_2A) [85], before they are read out in the room temperature space. The HFET amplifier is a broadband device that exhibits gain of ~ 33 dB from roughly 0.6-2 GHz with an expected noise temperature in the range of 3-4 K. In addition to the amplifiers, there are three circulators from QuinStar [9] along the output line which are unidirectional devices that are used to direct the signal, as well as prevent downstream noise from entering the receiver. There is also a cryogenic RF switch from Radiall [7] between the cavity and the first circulator. The switch can be flipped to make it so the receiver chain, instead of being connected to the cavity, is connected to a variable temperature stage (VTS) [125] used for noise calibration, also referred to as the "hot load". The hot load and noise performance of the various components in the output line are discussed in great detail in Chapter 3.

There are 10 and 20 dB attenuators on the weak, bypass and pump lines at the reservoir bottom (5K) and mixing chamber (85 mK) temperature stages respectively. Additionally, there are 0 dB attenuators at the ICP (150 mK), 1K plate (1K), reservoir top (10 K), cooled

baffle (40 K) and room temperature (300 K) temperature stages on all lines. The 10 and 20 dB attenuators provide protection from downstream noise or spurious signals entering the receiver chain. Furthermore, the 0, 10 and 20 dB attenuators function as thermal sinks for the cables to aid in reducing the overall thermal load.

Now that we have covered the components that make up the Run 1D receiver chain, I will discuss the specific upgrades that were made for this run. To begin, one of the collaboration's goals for Run 1D was to decrease the mixing chamber base temperature to achieve a lower system noise temperature and improve the experiment's sensitivity. Therefore, upgrading the receiver chain to reduce the heat load on the dilution refrigerator was a priority. Additionally, we aimed to simplify the overall design, and switch out a few cables that had been identified as slightly magnetic after the end of the previous data taking run, Run 1C. This task was led by a visiting ADMX graduate student from the University of Sheffield, Claude Mostyn.

Fourier's law of conduction [87] tells us that given a cable with mean thermal conductivity λ , cross-sectional area A , length L , and a temperature difference in the direction of heat conduction ΔT , the steady state heat conduction flow is equal to

$$\dot{q}_{\text{cond}} = \frac{A}{L} \lambda \Delta T. \quad (2.8)$$

Let us discuss the upgrades one line at a time. To preface, I will denote the cable properties throughout in order of cable length, outer conductor material, inner conductor material. The conductor material abbreviations are as follows: stainless-steel (SS), copper (Cu), tin-plated copper (TC), silver-plated copper-clad stainless steel (CS) and niobium-titanium (NbTi).

Beginning with the weak line, working downward in temperature, the first change made was between the cooled baffle (40 K) and reservoir top (10 K). The old 20" SS-SS cable at this stage was replaced with a nearly identical but slightly longer, 24" SS-SS cable. As expected by the inverse relationship between \dot{q}_{cond} and L in Eq. 2.8, lengthening the cable slightly reduced the expected conductive heat load from 0.31 mW to 0.26 mW. The next change was made between the reservoir top (10 K) and reservoir bottom (5 K). In the old receiver,

there were 3 small cables that spanned across this temperature stage: a 36" TC-CS, a 32 cm SS-SS, and 24" TC-CS. In the new receiver, these three cables were replaced by a single, 60" SS-SS cable. This not only simplified the design but dramatically improved the heat load because tin, silver and copper all have significantly higher thermal conductivities than stainless steel. To be more specific, the expected conductive heat load for this temperature stage went from 0.70 mW to only 1.93 μ W. Between the reservoir bottom (5 K) and the 1K plate (1K), the previous 14 cm SS-SS cable was replaced with a 20 cm SS-SS cable. As before, lengthening the cable here lowered the expected heat load from 0.036 mW to 0.252 mW. Next, I will discuss two temperature stages at once, the 1K plate (1K) to the still plate (0.7 K) and the still plate (0.7 K) to the intermediate cold plate, or ICP, (0.15 K). In the old receiver, there was a 14 cm SS-Cu cable between the 1K plate and the still plate and a 32 cm SS-SS cable between the still plate and the ICP. In the new receiver, there is a single 20 cm SS-SS cable between the 1K plate and the ICP. As stated before, copper has greater thermal conductivity than stainless steel, so this change reduced the combined expected heat load of these two temperature stages from 3.30 μ W to 0.20 μ W. Finally, the 14 cm SS-SS cable between the ICP (0.15 K) and the mixing chamber, or MC, (0.085 K) was replaced with a slightly longer, 20 cm SS-SS cable. This reduced the expected heat load from 5.84 nW to 4.90 nW. Overall, the majority of the conductive heat load on the weak line comes from the room temperature (300 K) to cooled baffle (40 K) stage, 258 mW, which was not altered between Run 1C and Run 1D. However, by making these changes we were able to reduce the total expected conductive head load of the weak line from 259.05 mW to 258.28 mW. The JPA pump line had identical cable upgrades to the weak line, so it had the same expected reduction in heat load.

Next, let us discuss upgrades made to the bypass line. I will begin with two stages at once, the cooled baffle (40 K) to reservoir top (10 K) stage, and the reservoir top (10 K) to reservoir bottom (5 K) stage. The old receiver had a single, roughly 7 foot long Cu-CS cable between the cooled baffle and the reservoir bottom. In the new receiver, this cable was replaced by two cables: one 24" SS-SS between the cooled baffle and the reservoir top, and a

60" SS-SS cable between the reservoir top and reservoir bottom. This reduced the expected combined thermal load of these two stages from 1.24 mW to 0.26 mW, again due to the larger thermal conductivity of copper and silver compared to stainless steel. Next, from the reservoir bottom (5 K) to the 1K plate (1K), the previous 14 cm SS-Cu cable was replaced with a 20 cm SS-SS cable. Again, the thermal conductivity of copper is larger than stainless steel, and the new cable is slightly longer, both of which helped reduce the expected heat load from 2.15 mW to 25.20 μ W. Next, I will discuss both the 1K plate (1K) to still plate (0.7 K) and still plate to ICP (0.15 K) stages together. In the old receiver, there was one 14 cm SS-Cu cable from the 1K plate to the still plate and one 14 cm SS-SS cable from the still plate to the ICP. In the new receiver, these two cables were replaced by a single, 20 cm SS-SS cable from the 1K plate to the ICP. Again, by replacing the cable with copper with a fully stainless steel cable, the expected heat load was reduced from 3.40 μ W to 0.20 μ W. Finally, between the ICP (0.15 K) and the MC (0.085 K) the old receiver had two 14 cm SS-SS cables that were replaced by a single 20 cm SS-SS cable in the new receiver. While this actually slightly increased the expected heat load from 2.92 nW to 4.90 nW, the design is simpler and, ultimately, the heat loads at this stage are already so small this is not our greatest concern. As with the weak line, the room temperature (300 K) to cooled baffle (40 K) stage is the primary contribution to the overall heat load with a value of 168.4 mW, and this was not altered between Run 1C and Run 1D. However, by making the changes outlined above we were able to reduce the total expected heat load on the bypass line from 171.79 mW to 168.68 mW.

The output line remained the same between Run 1C and Run 1D above the cooled baffle (40 K) temperature stage. Between the cooled baffle (40 K) and the reservoir bottom (5 K) temperature stages, in the old receiver there were two cables, one roughly 7 foot long Cu-CS followed by a 24" long TC cable. In the updated receiver the order of the cables has been switched as well as the material of the shorter cable. To be more specific, there is now a 24" long SS-SS cable followed by the same roughly 7 foot long Cu-CS cable. By swapping the position of the longer Cu-CS cable and by replacing the TC cable with a SS-SS cable,

the expected heat load between the cooled baffle (40 K) and reservoir top (5 K) was greatly reduced from 18.97 mW to 0.98 mW. Next, between the reservoir bottom (5 K) and the MC (0.085 K) temperature stages the old receiver had a single, 30 cm long NbTi-NbTi cable. In the upgraded receiver, there is now an additional 30 cm long SS-BeCu cable between the reservoir bottom (5 K) and 1K plate (1K) temperature stages. The existing 30 cm long NbTi-NbTi cable now runs from the 1K plate (1K) to the MC (0.085 K) temperature stages. While the total combined expected heat load between the reservoir bottom (5K) and MC (0.085 K) temperature stages actually increased slightly from 12.50 μ W in Run 1C to 121.08 μ W in Run 1D, the expected heat load on the MC (0.085 K) temperature stage was decreased significantly by relocating the NbTi-NbTi cable from the reservoir bottom (5 K) to the 1K plate (1 K). Now, rather than the expected heat load being concentrated on the MC (0.085 K) stage which has limited cooling power, the heat load is concentrated on the 1K plate (1 K) temperature stage which has much greater cooling power from the big 1K pot. In fact, in the updated receiver the expected heat load between the 1K plate (1 K) and MC (0.085 K) temperature stages is only 83.4 nW, a small fraction (0.001%) of the total expected heat load between the reservoir bottom (5 K) and the MC (0.085 K).

For Run 1C a through line was added to connect the hot load to the top of the insert (see Fig. 2.11). This was done so the receiver gain could be monitored during noise calibrations as well as data taking. However, for Run 1D, the design was simplified by instead adding a directional coupler to the weak line that connects to the hot load. This enabled the removal of an entire through line while maintaining ability to monitor the gain when the receiver is connected to the hot load. This upgrade not only simplified the receiver chain, but also drastically lowered the heat load on the dilution refrigerator.

Changes were also made to the attenuators throughout to account for any increases in attenuation from the new set of cables. Specifically, the 3 dB attenuators at the cooled baffle (40 K) and 1K plate (1K) temperatures stages were replaced with 0 dB heat sinks and the 20 dB attenuators at the reservoir bottom (5 K) stage on the weak and bypass lines were replaced with 10 dB attenuators.

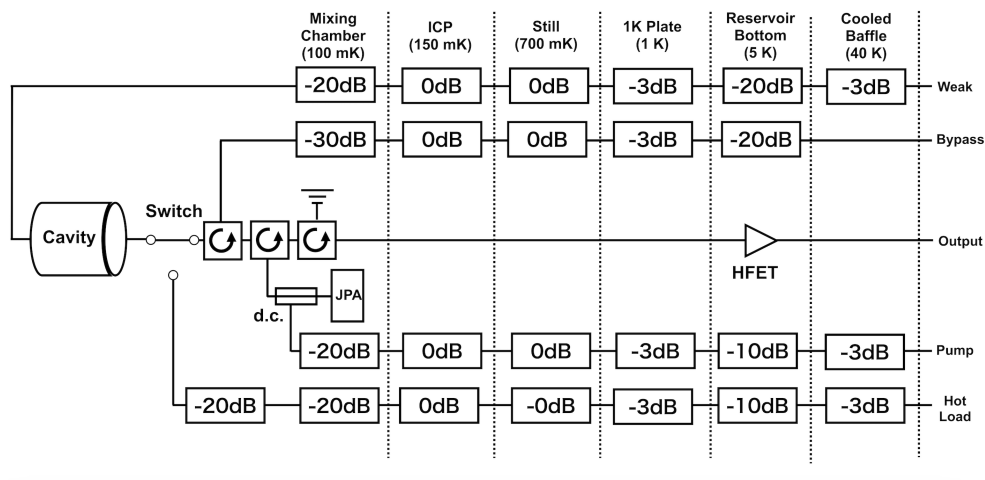


Figure 2.11: The cryogenic receiver chain for Run 1C. There are five different lines shown. The weak and bypass lines are primarily used for characterization of cavity parameters, while the pump line provides the input pump tone for parametric amplification with the Josephson Parametric Amplifier (JPA). The hot load line is used to monitor the gain of the receiver while it is connected to the hot load (left-most 20 dB attenuator in the diagram), which is a variable temperature stage used for noise calibration. The hot load can be connected to the output line using an RF switch. Lastly, the output line is used to amplify and readout signals from the cavity. ADMX uses two amplifiers, a JPA and an HFET amplifier. The three small squares with circular arrows in them are unidirectional devices called circulators which are used to direct the signal path and prevent downstream noise from reaching the cavity. There are 3 dB attenuators on weak, pump and hot load lines cooled baffle temperature stage. There is another set of 3 dB attenuators on the weak, bypass, pump and hot load lines at the 1K plate temperature stage as well. In addition, there are 20 dB (10 dB) attenuators on the weak and bypass (pump and hot load) lines at the reservoir bottom temperature stage. There are also 20 dB (30 dB) attenuators on the weak, pump and hot load (bypass) line at the mixing chamber temperature stage. Finally, there are 0 dB “attenuators” that serve as thermal sinks on the weak, bypass, pump and hot load lines at the ICP and still temperature stages.

2.4 DC Cables and Sensors

The nature of ADMX requires a significant number of cryogenic DC lines to power various components within the insert. These include temperature sensors, amplifier bias current sources, heaters, and more. There are a total of nine different DC cables that provide power to the insert during data taking operations. Five of the cables (A, B, C, D, E) have 51 pin micro-D connectors and the other four (F, L, R, SC) have 25 pin micro-D connectors. The steady, reliable operation of the various DC connections in the experiment is paramount to our ability to run. Testing, repairing and upgrading the DC cables and components made up a large part of my graduate career.

First, I will delve into more detail about what the various connections are. Cable A houses the temperature sensors used to monitor the temperature of the warmer insert components such as the warm and cooled baffles, and the coldhead. The warm component temperature sensors are platinum, PT-100, temperature sensors which provide excellent accuracy and stability above 30 K. Due to its use for monitoring the warmer temperature components, the breakout box for Cable A is mounted on and thermally sunk to the cooled baffle which sits at a temperature of about 40 K. Because ADMX does not have many components above 30 K that need temperature monitoring, Cable A also serves as overflow for other connections including the reservoir top (6), reservoir bottom (7), and 4K shield bottom (17) temperature sensors. These are Cernox temperature sensors from Lake Shore Cryotronics [1]. These sensors provide good sensitivity and stability over a broad range of temperatures ($\sim 0.1 - 325$ K), and have been proven very stable over repeated thermal cycling. However, below 1K the magnetoresistance is quite large ($\sim 70\%$), meaning they can't accurately measure temperatures below 1K when the field is ramped. Additionally, Cernox sensors do not follow a standard curve so we need to calibrate every uncalibrated Cernox sensor ourselves.

Cable B is intended for connections at the 1K stage. As such, the breakout box is mounted on and heat sunk to the 1K plate. The connections on Cable B include the HFET, Sorb, 1K

shield bottom, 1k shield top, big 1K pot, and rod gearbox temperature sensors. The sorb and 1k shield bottom sensors are Cernox, while the others are all Ruthenium Oxide (RuOx) sensors also from Lake Shore. RuOx sensors are designed specifically for operations within magnetic fields. Compared with Cernox sensors, the magnetoresistance of RuOx sensors at temperatures below 1K is small ($< 5\%$). In addition, RuOx sensors follow a standard calibration curve, so they can be purchased uncalibrated and still be used for monitoring temperature stages where the exact value is not hugely important.

Cable C is for higher current DC components. This includes the sorb, small 1K pot, big 1K pot, bore bottom, and mixing chamber heaters, as well as the hall probes. In addition, Cable C provides power to the HFET amplifier. The breakout box for Cable C is mounted on and thermally sunk to the bottom of the reservoir due to its large amount cooling power.

Cable D is for the various components in the quantum amplifier package. Due to the small volume of the quantum amplifier package, no breakout box is used for this cable. Instead, the output of the quantum amplifier package is two 25 pin micro-D connectors, one for temperature sensors and the other for the remaining DC connections. There are temperature sensors that monitor the JPA, the circulators, and the cold finger. These sensors are also Ruthenium Oxide sensors, but this time from Scientific Instruments [2]. The properties are much the same as before, with the sensors providing excellent magnetic field performance, interchangeability, stability and sensitivity to temperatures between 0.05 and 20 K. In addition, this cable supplies power for biasing the JPA, as well as for flipping the RF switch used for noise calibration measurements.

Cable E houses the milliKelvin space temperature sensors, and as such is mounted on and thermally sunk to the top of the cavity. There are many important sensors on this cable including temperature sensors on the ICP, hot load, bypass attenuator used for noise calibration, and the top and bottom of the cavity. One sensor on the top of the cavity and one sensor on the bottom of the cavity are Cernox, which due to their large calibrated temperature range, allows useful monitoring of the cavity temperature during cooldown. The other two cavity sensors, as well as all of the other sensors listed are RuOx sensors

from Lake Shore since we need them to perform well at very low temperatures inside a high magnetic field. Many of the sensors on this cable are critical to calibrating the system noise temperature, which allows us to determine our experimental sensitivity, so ensuring the stable operation of Cable E is vital.

Cable F supplies power to the DC power components of the dilution refrigerator. Specifically, it powers the small 1K pot, still, and mixing chamber temperature sensors as well as the still heater. Having timely, accurate readings of these temperatures is paramount to the operation of the dilution refrigerator. The still heater is another critical component of the dilution refrigerator because it is what allows us to increase the circulation of He3 through the mixing chamber, which allows us to decrease the mixing chamber temperature from ~ 250 mK to ~ 85 mK. Without still heating, the experiment would be significantly less sensitive. Unfortunately, throughout Run 1C we experienced regular disruptions to the temperature sensor readings from Cable F, and ultimately tracked the issue down to the deteriorating insert cable. As the DC cable lead for ADMX, I decided it was time we replace the cable, and build it its own breakout box like the other cables had to allow for quick and easy reroutes in the event that a pin breaks. I will discuss more about the Cable F upgrade in Sec. 2.4.2.

Cable L houses the temperature sensors for the leads box. This is a box that houses the leads for the bucking coil magnet and was intended to provide a thermal sink between the top of the high-TC leads and the cooled baffle (40 K). The temperature sensors were initially designed to ensure the leads were not getting too hot. Now that we now that is not a concern, these sensors are mostly used for monitoring the insert during cooldowns and magnet ramps. As such, the Cable L temperature sensors are all Cernox due to their wide temperature range and low magneto-resistance above 30 K. Since the sensors on this cable all go to the same place close to the top of the insert, there is no breakout box for Cable L.

Cable R supplies power to the various DC components in or around the helium reservoir. This includes 3 different Cernox temperature sensors on various parts of the bucking coil, the persistent switch for the bucking coil, a liquid helium level meter, and the reservoir heater.

The bucking coil creates a counter magnetic field to oppose the main magnetic field and create a field-free region of the insert where we house the magnetic-field sensitive quantum amplifier package. Being able to monitor the reservoir level is also critical to operations since we must have a constant supply of He4 available to ensure the insert remains cold during operations. Like Cable L, the sensors on this cable are all going inside the reservoir to the same place more or less, so there is no breakout box for Cable R either.

Lastly, there is the SC Piezo cable which powers the piezo motors that control the rod motion of the sidecar cavity, which is used as a test-stand for various R&D projects. More details about the status and results of sidecar from Run 1D can be found in former ADMX graduate student Tom Braine's dissertation [47].

2.4.1 Sensor Documentation and Testing

Given that experimental operations rely heavily on the reliability of the many DC connections described above, documentation and testing are paramount. I have thoroughly documented the existing DC connections inside the ADMX insert in a spreadsheet, as well as with a sensor map. The spreadsheet includes important information about each connection including what cable pins the component is occupying, what type of sensor it is and where it is from, the name of its calibration, a description of its location inside the insert, and its associated readout instrument. It also notes any known issues such as broken pins on the cables or flaky sensors. An example of part of the spreadsheet can be seen in Fig. 2.12.

Furthermore, I also created a sensor map of the insert as another, more visual, aid to be used for locating specific sensors inside the insert. Given the huge amount of components, the limited space in the insert, and the small size of the temperature sensors it can be quite difficult to locate them if one is not working with them regularly. In addition to the map itself, I also assembled a gallery of pictures of the sensors when possible to further aid in the documentation. This way, any collaborator or future graduate student should be able to step in and easily locate any sensor. The sensor map and its key can be seen in Fig. 2.14 and Fig. 2.15, and the image gallery is shown in Fig. 2.13.

In addition to documenting the insert DC connections, I also rigorously tested each connection in a procedural fashion. The first checks are done before the insert shields are put on prior to insertion into the magnet bore, the second checks are done after the shields are put on but before insertion, and the third checks are done after insertion. To begin the first set of checks, I used a digital multimeter to measure the resistances across the relevant pins for each connection at the top of the insert. The resistances were recorded in a spreadsheet where they were either confirmed to be the same as the expected value from the last run, or noted as problematic. Then, I checked that the resistance values at the appropriate readout instrument matched what was measured at the top of the insert to ensure the connections are sound all the way to the data acquisition (DAQ) rack. Any connections that are not reading out correctly at any stage were further evaluated to determine where the break was. For example, if a sensor read out at the top of the insert, but not at the DAQ rack, the issue is somewhere in the warm space. However, if the sensor doesn't even read out at the top of the insert, the issue is in the cold space, meaning the issue must be resolved prior to insertion. Cold space wiring issues therefore take higher priority than warm space issues. The exact reason for a cold space failure varies, but it is often due to either one of the thin, delicate sensor wires breaking at the connector, or because one of the pins on an insert cable was damaged. In the former case, that can often be resolved by soldering the broken wire back into the connector, as was done for a number of sensors prior to Run 1D. In the latter case, that pin is no longer usable and the sensor needs to be rerouted to a spare pin on the cable via the appropriate breakout box. This is the beauty of the breakout boxes, it allows for quick fixes when there are individual broken cable pins. Of course, there is a limited amount of space on each cable so this rerouting can only be done so many times before the entire cable must be replaced, but the ability to use spare pins maximizes the lifetime of the cables in the insert. Prior to the Run 1D insertion there were a number of connections on Cable E that required rerouting at the breakout box level.

The same set of checks described above are performed once the shields are put on to determine if we broke anything in the process of putting the shields on. Then, after insertion

into the magnet bore, this same set of checks is done once more to determine if anything was disturbed by the insertion process. The only difference here is that, with the insert in the bore, if any connections are not reading out at the top of the insert there is nothing that can be done to fix them at this stage. That is to say, only warm space issues are resolvable post-insertion. Therefore, this check determines whether or not operations can continue as planned or if we need to perform an extraction to repair a critical DC connection in the cold space. Luckily, for Run 1D, the pre-insertion checks and repairs were thorough enough that I did not see any broken connections post-insertion. In fact, throughout the course of this run there were no disruptions to operations due to sensor issues, as there had been during previous data taking runs.

Pin	Connection	Sensor Type	Calibration	Location	Readout Instrument (Channel)
1	Cavity Top Temp 1	Ruthenium Oxide (Lakeshore)	U09595 (done by Nick/LLNL + RX202a standard curve)	on cavity top plate near breakout box F	Lakeshore Resistance Bridge (Ch. 1)
2	Cavity Top Temp 1	Ruthenium Oxide (Lakeshore)	U09595 (done by Nick/LLNL + RX202a standard curve)	on cavity top plate near breakout box F	Lakeshore Resistance Bridge (Ch. 1)
3	Cavity Top Temp 1	Ruthenium Oxide (Lakeshore)	U09595 (done by Nick/LLNL + RX202a standard curve)	on cavity top plate near breakout box F	Lakeshore Resistance Bridge (Ch. 1)
4	Cavity Top Temp 1	Ruthenium Oxide (Lakeshore)	U09595 (done by Nick/LLNL + RX202a standard curve)	on cavity top plate near breakout box F	Lakeshore Resistance Bridge (Ch. 1)
5	Cavity Top Temp 2	Cernox	x76420 (from Lakeshore)	on cavity top plate near sidecar	Lakeshore Resistance Bridge (Ch. 2)
6	Cavity Top Temp 2	Cernox	x76420 (from Lakeshore)	on cavity top plate near sidecar	Lakeshore Resistance Bridge (Ch. 2)
7	Cavity Top Temp 2	Cernox	x76420 (from Lakeshore)	on cavity top plate near sidecar	Lakeshore Resistance Bridge (Ch. 2)
8	Cavity Top Temp 2	Cernox	x76420 (from Lakeshore)	on cavity top plate near sidecar	Lakeshore Resistance Bridge (Ch. 2)
9	Cavity Bottom Temp 1	Cernox	x76422 (from Lakeshore)	on cavity end plate (labeled)	Lakeshore Resistance Bridge (Ch. 3)
10	moved to pin 26 (broken)				
11	Cavity Bottom Temp 1	Cernox	x76422 (from Lakeshore)	on cavity end plate (labeled)	Lakeshore Resistance Bridge (Ch. 3)
12	Cavity Bottom Temp 1	Cernox	x76422 (from Lakeshore)	on cavity end plate (labeled)	Lakeshore Resistance Bridge (Ch. 3)

Figure 2.12: Screenshot of the ADMX Run 1D wiring spreadsheet. The example shown here is the tab associated with Cable E.

Images of sensors in map [\[edit\]](#)

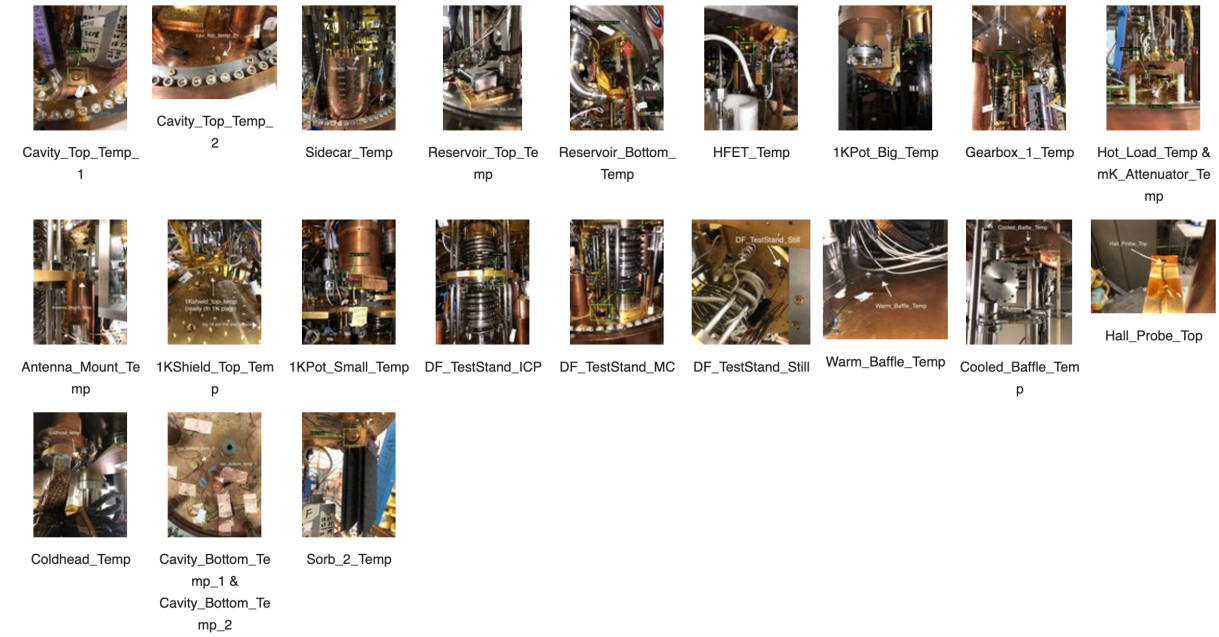


Figure 2.13: ADMX Run 1D image gallery of sensors in Fig. 2.14. These images are on the ADMX wiki page that can be accessed by internal collaborators who are interested in locating a specific sensor. Not all sensors were able to be easily photographed, but those that were are included in the gallery.




Figure 2.14: ADMX Run 1D sensor map. Each number corresponds to a sensor in Fig. 2.15.

SENSOR MAP KEY

- 1) Cavity-Top-Temp-1: on top of cavity, circular sensor in front of breakout box F
- 2) Cavity-Top-Temp-2: on top of cavity, flat rectangular sensor to the left of breakout box E, near sidecar
- 3) Cavity-Bottom-Temp-1: underneath cavity, pin-type sensor
- 4) Cavity-Bottom-Temp-2: underneath cavity, circular sensor
- 5) Sidecar-Temp: on top of sidecar cavity, circular sensor
- 6) Reservoir-Top-Temp: on top of reservoir on north side, next to breakout box C, circular sensor
- 7) Reservoir-Bottom-Temp: circular sensor on the bottom of the reservoir (4K plate), mounted on the "cheese board" above the antenna
- 8) HFET-Temp: small circular sensor on the HFET, located near Big 1K Pot
- 9) 1KPot-Big-Temp: near bottom of big 1K Pot on the side facing the center of the insert
- 10) Gearbox-1-Temp: on the underside of the rod 1 gearbox, near where the rod attaches (close to center of the insert)
- 11) 1KShield-Top-Temp: actually attached to the 1K plate, near the center of the insert on the same side as sidecar
- 12) 1KPot-Small-Temp: on the bottom of the small 1K pot, pin type sensor
- 13) DF-Teststand-ICP: on underside of the intermediate cold plate, small circular sensor
- 14) DF-Teststand-still: on underside of the still plate, small circular sensor
- 15) DF-Teststand-MC: on T-plate toward northwest side of insert, in front of mixing chamber
- 16) 1KShield-Bottom-Temp: on the bottom of the 1K shield (no photo)
- 17) 4KShield-Bottom-Temp: on the bottom of the 4K shield (no photo)
- 18) BC-Top-Temp: on top of bucking coil (no photo)
- 19) BC-Middle-Temp: in middle of bucking coil (no photo)
- 20) BC-Bottom-Temp: on bottom of bucking coil (no photo)
- 21) Warm-Baffle-Temp: on top of warm baffle, toward center of insert, flat piece of metal w/ black epoxy
- 22) Coldhead-Temp: attached to coldhead to the right of where the copper braids connect, flat piece of metal w/ black epoxy
- 23) Cooled-Baffle-Temp: on underside of cooled baffle, next to breakout box A, flat piece of metal w/ black epoxy
- 24) Hot-Load-Temp: mounted to bottom of hot load, on top of cavity on north side
- 25) mk-Attenuator-Temp: mounted directly under the hot load on the bypass attenuator block
- 26) Antenna-Mount-Temp: attached to the antenna mount, next to the ICP, behind antenna gear box
- 27) Sorb-Temp: inside burrito sorb above sidecar on south side of the insert (no photo)
- 28) Sorb-2-Temp: on inner face of Jimmy's plate sorb, near breakout box F on north east side of insert
- 29) SQUID-Circulator-Temp: mounted to the circulator closest to the cold finger on squidadel (no photo)
- 30) SQUID-Coldfinger-Temp: mounted to the squidadel coldfinger, near bottom of Hex frame (no photo)
- 31) SQUID-JPA-Temp: mounted to the quantum amplifier mount, below the JPA (no photo)
- 32) Hall-Probe-Top: inside reservoir, above squidadel, attached to long mounting piece

Figure 2.15: ADMX Run 1D sensor map key. Each number corresponds to a sensor in Fig. 2.14.

2.4.2 Cable F Redesign

As mentioned earlier in the text, Cable F was due for a replacement after experiencing regular disruptions to operations because its sensors kept failing. While the idea of purchasing a commercial cable was discussed, because of the specificity of the design needs we ultimately decided to pursue a custom cable designed by myself and assembled by collaborators at Washington University in St. Louis (WUSTL). The process began with removing the existing cable, and measuring how long the total length will need to be, which ended up being around 11 feet. However, in the process of removing the cable I noticed that it was not heat sunk very well to the different thermal stages in the insert. Considering this cable was used to supply power to some of the coldest components in the insert, this was alarming to me. Therefore, I went through and decided where in the insert the new cable should be heat sunk, and took measurements of each individual stage to give our collaborators at WUSTL a drawing which clearly marked how much room to leave between each copper bobbin. After adding extra slack to give us freedom in the location where the heat sinks would be mounted to the different temperature stages, the total length of the design ended up being just over 11.4 feet. The heat sink locations and section lengths are shown in Fig. 2.16.




Figure 2.16: Design drawing for the upgrade dilution refrigerator cable, Cable F. This version of Cable F included more heat sinking bobbins in more strategic locations, different wire material, and more easily interchangeable connectors.

The feedthroughs at the top of the insert for all of the cables are equipped to fit 50-pin micro-D connectors. However, for Cable F we only need 25 pins to actually reach the milliKelvin space. Therefore, at some point inside the insert we wanted to reduce the number of wires from 50 to only 25. Because the feedthrough itself required a 50-pin connector, we purchased a new 50-pin micro-D connector with a copper tail from MDC Precision [3] to connect to the inside of the vacuum flange on one end, and to the remainder of the cable on the other. At the end of the copper tail, a 25-pin micro-D connector was soldered to the first 25 wires, and the remaining 25 wires were terminated so that the rest of the cable would have only 25 wires while maintaining compatibility with the existing feedthrough. The rest of the cable past this point was primarily composed of quad twist phosphor bronze wire due to its high thermal resistivity, which aids in thermal isolation between the different temperature stages. However, because phosphor bronze also has high electrical resistance, it should not be used for high current components, like a heater. Due to this limitation, I chose to use copper wires for four of the pins (13,14,15,16), two of which (14 and 15) were designed to be used for the still heater. The other two pins serve as spares that the still heater could be rerouted to if pin 14 or 15 were to get damaged. The pinouts for the new Cable F are shown in Fig. 2.17.

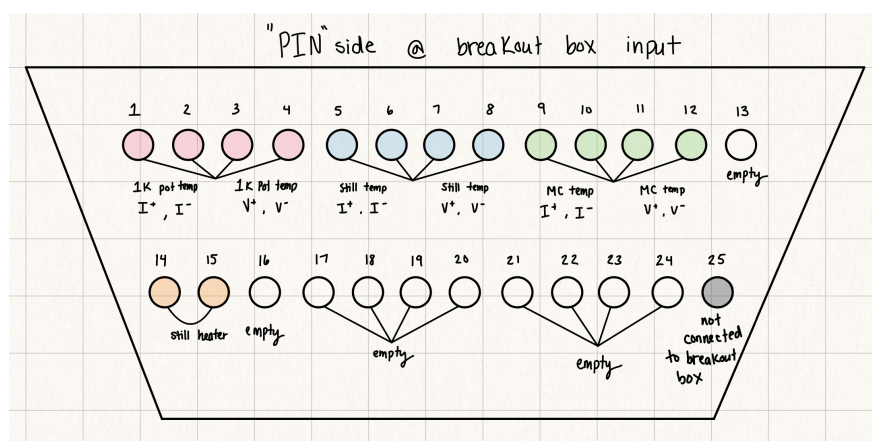


Figure 2.17: Pinout diagram for the upgraded dilution refrigerator cable, Cable F. Pins 13-16 were designated to be made with copper wire, while the rest of the pins are phosphor bronze. This is because the still heater, which uses more current than the temperature sensors, was designated to be on pins 14 and 15, leaving 13 and 16 as backups.

In addition to the cable itself, I designed a new breakout box to make rerouting pins easier than the existing connector. For temperature sensors we typically use four pin black TE Connectivity connectors, so the breakout board was designed to have the opposite type of connector allowing the sensors to plug directly in to the appropriate slot. The board itself has 24 pins that are arranged in sets of four, for a total of six black connector slots. After coming up with my design and purchasing the necessary components (PCB, black connectors and a 25-pin micro-D with a copper tail), the breakout board was assembled by the CENPA electronic shop engineering technician, David Peterson. Fig. 2.18 shows a comparison between the old connector and the new breakout board. The board itself was then housed inside of a small aluminum box that could be mounted directly to the mixing chamber plate. I machined the box to make openings for the connectors, enabling the breakout board to be contained and protected while still remaining accessible for electrical connections. A photo of the breakout box after being installed in the insert can be seen in Fig. 2.19.




Figure 2.18: Left: the former “breakout box” for Cable F. Really, it was just a connector that could be disconnected from the stand. The left hand side where wires are soldered into connected to the old Cable F, and the right hand side with the white cables coming out were routed to the sensors and the still heater. This connector was not only dilapidated, but it was not as easily labeled and changes were difficult to keep track of because the wires on the cable itself were what needed to be re-soldered in the case of a break, rather than having some intermediate connection that changes could be made to. Right: new and improved breakout board. This board is what is housed inside of the box shown in Fig. 2.19. This breakout board has entirely fresh wiring and connectors, and is much more clearly labeled with every connector and every pin on each connector having a number associated with them. In addition, because Cable F plugs into this board, the changes made during rerouting are made to the board, rather than the cable itself.



Figure 2.19: The newly designed breakout box for Cable F mounted on top of the cavity next to the breakout box for Cable E (left edge of the image). The connectors for temperature sensors can be seen on the left hand side of the breakout box and there are 3 spares. Cable F is shown plugged into the top of the box.

2.4.3 Cable D Filter Addition

Prior to Run 1D a cryogenic, combination Resistor-Capacitor(RC)/Inductor-Capacitor(LC) filter was installed on Cable D, which supplies power to the JPA as well as the RF switch used for noise calibration. The primary goal of adding this filter was to reduce electronic

noise on the JPA bias line in order to improve its noise performance and dynamic range. RC filters can be designed with very low cutoff frequencies, and the resistive elements assist in thermalizing wires at low temperatures. However, RC low pass filters become ineffective at frequencies > 100 MHz. Therefore, to eliminate frequencies above 100 MHz, multiple stages of LC filtering were used. Filtering at these higher frequencies is crucial because the signals used in wireless communications, which have frequencies on the order of hundreds of MHz to a few GHz, can easily couple into the antenna-like DC wiring entering the insert. All 25 lines on Cable D include the same LC filtering configuration: per line, there are two LC stages with a 105 MHz cutoff, one stage with a 1.45 GHz cutoff, and one stage with a 5 GHz cutoff. Each line also has two stages of RC filtering, with a cutoff frequency of 20 Hz for the JPA bias lines (plus several spares), and a low resistance high cutoff frequency of 48.8 kHz on the remaining lines to allow for high-current impulse signals to be sent to the RF switch. The filter box was largely designed and manufactured by our collaborators at WUSTL, especially Jordan Russell. I was responsible for integrating the new filter box into the existing cabling. In the process of adding the filter box, we also adding another stage of heat-sinking to Cable D, because we noticed that it had previously not been heat sunk anywhere until the top of the reservoir which is at 10 K. As seen in Fig. 2.20, this was achieved by relocating the breakout box for Cable D from the top of the reservoir to the cooled baffle which sits at around 40 K. Another heat sink was added at the top of the reservoir in place of the breakout box that had previous been there. The goal of this was to reduce the heat load on the reservoir and improve the helium boil off rate.




Figure 2.20: Drawing showing the changes made to Cable D in the cryogenic space. Namely, the breakout box was relocated from the top of the reservoir to the cooled baffle and another heat sink was added to the top of the reservoir in its place. In addition, a filter box was added on the bottom of the reservoir to reduce electronic noise on the JPA bias line.

To summarize, the hardware upgrades made prior to the start of Run 1D were a great success. Efforts to reduce the heat load on the system caused by components in the receiver chain allowed ADMX to reach lower temperatures than ever before. In addition, there were very few sensor failures throughout Run 1D, none of which were related to the newly designed DR cable, Cable F.

Chapter 3

NOISE CALIBRATION

Calibrating the receiver noise allows ADMX to determine at what level we can discriminate between a signal originating from the cavity and thermal noise after passing through a series of RF components. Some of these components attenuate, some amplify, and all have their own contributions to the overall system noise that must be considered. Additionally, in the event of discovery the noise calibration results will directly influence the uncertainty on the reported value of $g_{a\gamma\gamma} \times \rho_{DM}$. In this chapter, I will discuss in detail both the theoretical noise model based on the components in our receiver chain as well as the experimental and analytical procedures followed to measure the system noise. I will finish with a summary and discussion of the results from the Run 1D noise calibration campaign. I have already detailed most of the content from this chapter in a publication in Physical Review D [96], some of which is repeated here. The JPA-off-hot-load calibration on 01/04 and JPA-off-cavity calibration on 06/14, however, are new to this thesis.

3.1 Sources of Noise

The signal-to-noise ratio (SNR), which is the ratio of the signal power entering the cold receiver P_{sig} to the root-mean-square (RMS) uncertainty of the system noise power P_{sys} over some bandwidth b , is a fundamental measure of experimental sensitivity. In order to determine the RMS uncertainty on P_{sys} , we must understand the expected noise distribution. The spectra used to search for virialized halo axions with ADMX are designed to have a bin width of 100 Hz given the expected axion linewidth of ~ 1 kHz. Considering the native spectral resolution of ADMX is only 10 mHz, the combination of 10,000 bins into a single bin is necessary to create spectra with 100 Hz wide bins that will be optimally sensitive to

axions originating in the Milky Way halo. Due to this combination of many individual bins, the spectra analyzed in this thesis are expected to follow a Gaussian noise distribution [80]. Because the spectra follow a Gaussian noise distribution, the SNR can be defined via the radiometer equation (see Ref. [50] for a derivation), resulting in an RMS uncertainty of system noise power equal to $\sigma_{\text{sys}} = P_{\text{sys}}/\sqrt{bt}$ where b is the bandwidth at the output of the receiver and t is the integration time of the digitization. Thus, the noise temperature T_{sys} of a device is defined to be scaled relative to the input signal reference plane as

$$\text{SNR} = \frac{P_{\text{sig}}}{\sigma_{\text{sys}}} = \frac{P_{\text{sig}}}{k_B T_{\text{sys}} b} \cdot \sqrt{\frac{t}{b}}, \quad (3.1)$$

where k_B is the Boltzmann constant, $P_{\text{sys}} = k_B T_{\text{sys}} b$ and all other quantities are as defined above. For clarity, “reference plane” here defines the location in the receiver chain where the quantities of interest (T_{sys} and the resulting SNR, in this case) are being defined from. That is, the system noise temperature and corresponding SNR are defined as above (Eq. 3.1) from the specified reference plane to the output of the receiver. That being said, let us begin by examining each of the components that contribute T_{sys} .

3.1.1 Blackbody Noise

There are various sources of noise in the ADMX receiver that fall into two general categories, thermal and electronic. The thermal noise temperature of a blackbody is a function of its frequency f and physical temperature T_{phys} :

$$T_{\text{noise}}(f, T_{\text{phys}}) = \frac{hf}{2k_B} \coth\left(\frac{hf}{2k_B T_{\text{phys}}}\right), \quad (3.2)$$

where h is the Planck constant [54]. Noise power is generally related to noise temperature through the Boltzmann constant k_B and some bandwidth b at the RF receiver output:

$$P_{\text{noise}} = k_B b T_{\text{noise}}. \quad (3.3)$$

Depending on the relative magnitude of T_{phys} and $\frac{hf}{2k_B}$, the system is either thermally limited or quantum-limited. In the former case, $T_{\text{phys}} \gg \frac{hf}{2k_B}$ and $T_{\text{noise}} \approx T_{\text{phys}}$. In the

latter case, $T_{\text{phys}} \ll \frac{hf}{2k_B}$ and $T_{\text{noise}} = \frac{hf}{2k_B}$, which is equal to the Standard Quantum Limit (SQL) [55, 151].

3.1.2 Attenuator

Attenuators are used throughout the ADMX receiver chain to both aid in thermalization and to prevent noise from warmer temperature stages from reaching the coldest temperature stages. Kirchoff's law states that the emissivity and absorptivity of a passive object in thermal equilibrium must be equal [116]. In addition, by conservation of energy the sum of the absorptivity, reflectivity and transmissivity of a blackbody must be equal to 1. Here, we approximate our attenuators as non-reflective (that is, reflectivity = 0). Therefore, an attenuator that transmits a power fraction α and absorbs a power fraction $1 - \alpha$ would also radiate as a blackbody with emissivity equal to its absorptivity, $1 - \alpha$. Therefore, given a blackbody at a noise temperature T_A and an attenuator with the above transmissive and absorptive properties and noise temperature T_B , the noise power and temperature as measured downstream will be:

$$P_{\text{noise}} = k_B b (\alpha T_A + (1 - \alpha) T_B), \quad (3.4)$$

and

$$T_{\text{noise}} = T_A + \frac{1 - \alpha}{\alpha} T_B, \quad (3.5)$$

with the reference plane at blackbody A .

3.1.3 Circulator

Circulators are three-port, unidirectional devices that are used to control the direction of the signal path (see Fig.3.1). Namely, a signal entering through port 1 will exit port 2, a signal entering through port 2 will exit port 3 and a signal entering through port 3 will exit through port 1 [146]. For an ideal clockwise circulator, the transmissivity from port i to port

j , S_{ji} , is equal to zero except for $S_{21} = S_{32} = S_{13} = 1$. In reality circulators have some small amount of loss so these transmissivities are not exactly 1, but after measuring the losses, they can be treated as attenuators for the purpose of noise calibration.

3.1.4 Microwave Cavity

As discussed in Chapter 2, ADMX makes use of a resonant microwave cavity to enhance the strength of the photon signal produced by the axion field interacting with the static magnetic field. The resonant mode of interest, TM_{010} , has an associated unloaded quality factor Q_0 and is coupled to an antenna with coupling strength β . Power incident on the cavity at frequencies near some resonant frequency f_0 is reflected with reflectivity coefficient

$$|\Gamma_{\text{cav}}(f)|^2 = 1 - \frac{4\beta}{(1 + \beta)^2} \frac{1}{1 + 4Q_L^2 \left(\frac{f-f_0}{f_0}\right)^2}, \quad (3.6)$$

where $Q_L = Q_0/(1 + \beta)$ is the loaded quality factor [34, 146].

Generally, $|\Gamma_{\text{cav}}(f)|^2 \neq 0$, and therefore the noise temperature as seen from the antenna is a mixture of the cavity's physical temperature T_{cav} and noise temperature from power reflected off the antenna, T_{incident} :

$$T_{\text{noise}} = T_{\text{incident}}|\Gamma_{\text{cav}}|^2 + T_{\text{cav}}(1 - |\Gamma_{\text{cav}}|^2). \quad (3.7)$$

Only when the antenna is critically coupled ($\beta = 1$), and $f = f_0$ (on-resonance), does $|\Gamma_{\text{cav}}(f)|^2 = 0$. Under these conditions, the cavity appears as a blackbody radiating at T_{cav} .

3.1.5 Amplifier

In order to increase the signal to noise ratio, ADMX makes use of multiple amplifiers in its receiver chain. This is because not only do amplifiers increase the signal strength, but they also suppress the noise contribution from downstream receiver components [89]. Therefore, given a blackbody with noise temperature T_A amplified by an amplifier with gain G and

noise temperature T_B , followed by downstream components that introduce additional noise T_C , the noise power and temperature as measured at the reference plane of blackbody A is:

$$P_{\text{noise}} = k_B b(G(T_A + T_B) + T_C), \quad (3.8)$$

and

$$T_{\text{noise}} = T_A + T_B + \frac{T_C}{G}. \quad (3.9)$$

3.1.6 Parametric Amplifier

ADMX employs an ultra low-noise parametric amplifier as our first stage-amplifier so that it is the dominant contributor to the system noise, due to the suppression of noisier downstream components [89]. For Run 1D specifically, ADMX used a current-pumped (four-wave mixing) Josephson Parametric Amplifier (JPA) [111]. In the ideal case, a current-pumped JPA pumped at frequency f_P can achieve no additional noise beyond the zero-point fluctuations required by quantum mechanics at the signal and idler frequencies, f_S and $f_I = 2f_P - f_S$, corresponding to half a photon worth of noise at each frequency for a total of one photon worth of noise [125]. The noise power and temperature at the input of a parametric amplifier given a signal frequency with noise temperature T_S , an idler frequency with noise temperature T_I , and the gain of the parametric amplifier $G \gg 1$, followed by downstream components that introduce additional noise temperature T_D are:

$$P_{\text{noise}} = k_B b(G(T_S + T_I) + T_D), \quad (3.10)$$

and

$$T_{\text{noise}} = T_S + T_I + \frac{T_D}{G}, \quad (3.11)$$

with the reference plane at the input to the parametric amplifier. In reality, the parametric amplifier in ADMX adds noticeable extra noise above that required by quantum mechanics,

T_{JPA} , so the noise temperature becomes

$$T_{\text{noise}} = T_S + T_I + T_{\text{JPA}} + \frac{T_D}{G}. \quad (3.12)$$

3.2 Noise Models

3.2.1 Model Construction

The relevant part of the ADMX cold receiver chain for noise calibration is detailed in Fig. 3.1. The primary signal path (blue arrows) begins in the cavity, passes through two circulators (Circ1 and Circ2), is amplified by a four-wave mixing JPA, passes through two more circulators (Circ2 and Circ3), is amplified by an HFET amplifier, and then continues out to the warm receiver. In addition to the primary signal, the thermal noise path (brown arrows) begins at attenuator A, passes through Circ1 and reaches the antenna. Depending on the frequency f and coupling β , some of this incident noise is reflected, and some is replaced by thermal noise from the cavity (Eq. 3.7). The noise then passes through the same path as the signal, where additional electronic noise will be added by attenuation in the cables and circulators, by mixing with the idler frequency at the JPA, by the HFET amplifier, and by the post-amplifiers in the warm receiver.

In practice, the idler frequency for the JPA is always many Q widths away from the cavity resonance so its noise can be treated as independent from the cavity temperature. To be more specific, the measurements in Sec. 3.3 always detune f_p by 320 kHz higher than the cavity resonance (f_S) while the bandwidth of the cavity is 56 kHz. In addition, the HFET gain is quite high (~ 40 dB), so any additional noise from components downstream of the HFET is insignificant compared to the HFET noise and can be lumped in with it.

The system can be run with the JPA powered by a pump tone with stable gains up to 25 dB, or with the JPA inactive, in which case it behaves as an ideal reflector. The RF switch (S) can be flipped so that the noise originates in the “hot load” (VTS used in ADMX) for calibration. The cable length and composition between the hot load and the switch are designed to be the same as that between the cavity and the switch, so the attenuation in

both cases can be treated as nearly the same.

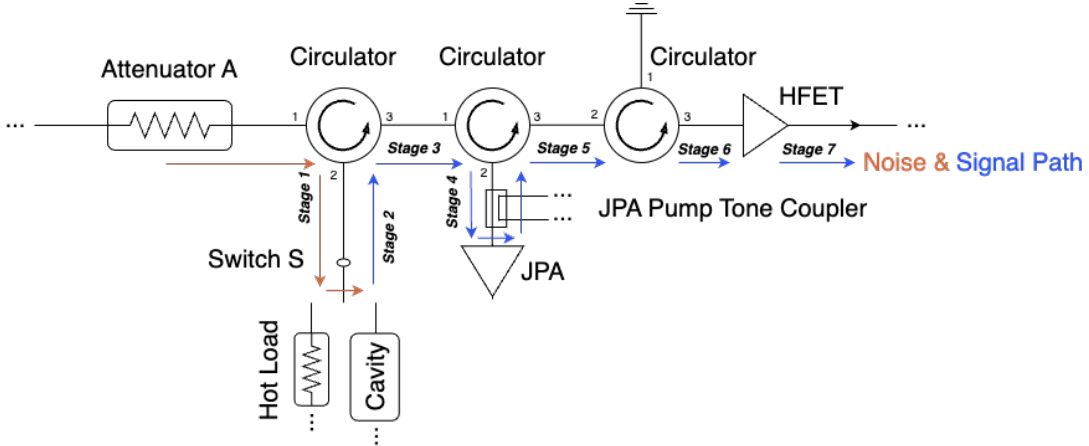


Figure 3.1: ADMX RF diagram in the cold space. Other than the HFET, all the components are connected to the milliKelvin stage directly. The blue arrows show the shared path for both noises and possible signal from the cavity. The brown arrows show the path of the attenuator A thermal noise to the cavity. Different stages (i.e. stage 1, stage 2, etc.) are labeled corresponding to the noise model in Sec. 3.2.

The thermal model is constructed by assuming temperature gradients among all the critical cryo-components including the cavity, attenuator A, the hot load, etc. (Fig. 3.1). Ideally, they should all be thermalized to the milliKelvin (mixing chamber) temperature stage, except for the HFET, but assuming temperature gradients is the most general starting point from which the model can be simplified later.

The cavity, attenuator A, and the hot load are each instrumented with temperature sensors that measure corresponding blackbody noise temperatures T_{cav} , T_A , and T_{HL} , respectively. The magnetic-field sensitive components including the circulators, switch and JPA are mounted to a cold finger in a field-free region of the insert, a system we call the quantum amplifier package, which is also instrumented with temperature sensors [27, 79]. For these components, I begin by assuming different physical temperatures for each of the circulators

in the interest of generality and later on I will simplify the model using the fact that all of the components in the quantum amplifier package are thermalized to the same temperature T_{circ} . The HFET noise temperature combined with any downstream receiver noise will be labeled T_{HFET} . The gains of the JPA, the HFET and the downstream post amplifiers are noted as G_{JPA} , G_{HFET} and G_{post} , respectively.

The noise power with the JPA off can be modeled by separating the cryo-space into different stages, where

$$\begin{aligned}
P_{\text{stage1}} &= k_B b T_A \alpha_{\text{circ1},21} + k_B b T_{\text{circ1}} (1 - \alpha_{\text{circ1},21}) \\
P_{\text{stage2}} &= P_{\text{stage1}} |\Gamma_{\text{cav}}|^2 + k_B b T_{\text{cav}} (1 - |\Gamma_{\text{cav}}|^2) \\
P_{\text{stage3}} &= P_{\text{stage2}} \alpha_{\text{circ1},32} + k_B b T_{\text{circ1}} (1 - \alpha_{\text{circ1},32}) \\
P_{\text{stage4}} &= P_{\text{stage3}} \alpha_{\text{circ2},21} + k_B b T_{\text{circ2}} (1 - \alpha_{\text{circ2},21}) \\
P_{\text{stage5}} &= P_{\text{stage4}} \alpha_{\text{circ2},32} + k_B b T_{\text{circ2}} (1 - \alpha_{\text{circ2},32}) \\
P_{\text{stage6}} &= P_{\text{stage5}} \alpha_{\text{circ3},32} + k_B b T_{\text{circ3}} (1 - \alpha_{\text{circ3},32}) \\
P_{\text{stage7}} &= G_{\text{HFET}} (P_{\text{stage6}} + k_B b T_{\text{HFET}}) \\
P_{\text{noise,out}} &= G_{\text{post}} P_{\text{stage7}}.
\end{aligned} \tag{3.13}$$

While the JPA is on, we rewrite the relation between P_{stage4} and P_{stage5} as

$$\begin{aligned}
P_{\text{stage5}} &= (G_{\text{JPA}} (P_{\text{stage4}} + P_{\text{JPA},S}) + (G_{\text{JPA}} - 1) (P_I + P_{\text{JPA},I})) \alpha_{\text{circ2},32} \\
&\quad + k_B b T_{\text{circ2}} (1 - \alpha_{\text{circ2},32}),
\end{aligned} \tag{3.14}$$

where $P_{\text{JPA},S}$ and $P_{\text{JPA},I}$ are extra noise at the signal and idler frequencies, respectively, due to imperfections in the JPA. P_I is the noise power at the idler frequency which can be traced up to P_{stage1} . As the idler frequency is always off-resonance with respect to the cavity, the reflection coefficient between stage 1 and stage 2 at the idler frequency is $|\Gamma_{\text{cav}}|^2 = 1$. More explicitly,

$$\begin{aligned}
P_I &= (P_{\text{stage1}}\alpha_{\text{circ1},32} + k_B b T_{\text{circ1}}(1 - \alpha_{\text{circ1},32}))\alpha_{\text{circ2},21} \\
&\quad + k_B b T_{\text{circ2}}(1 - \alpha_{\text{circ2},21}).
\end{aligned} \tag{3.15}$$

If there is a signal power P_{sig} coming out of the cavity, $T_{\text{cav}}(1 - |\Gamma_{\text{cav}}|^2)$ will be replaced with $(T_{\text{cav}}(1 - |\Gamma_{\text{cav}}|^2) + P_{\text{sig}})$ in Eq. 3.13. According to Eq. 3.1, P_{sig} should be compared to T_{sys} with the stage 2 as the reference plane because stage 2 is where the signal enters the receiver chain. Therefore, if all the attenuation and gain at different stages are known, T_{sys} can be calculated as

$$T_{\text{sys}} = \frac{P_{\text{noise,out}}}{k_B b \alpha_1 \alpha_2 G_{\text{total}}}. \tag{3.16}$$

Here, $\alpha_1 = \alpha_{\text{circ1},32}\alpha_{\text{circ2},21}$ is the transmissivity from the cavity (stage 2) to the JPA (stage 4), and $\alpha_2 = \alpha_{\text{circ2},32}\alpha_{\text{circ3},32}$ is the transmissivity from the JPA (stage 4) to the HFET (stage 6). While the JPA pump is on, $G_{\text{total}} = G_{\text{JPA}}G_{\text{HFET}}G_{\text{post}}$ and while it is off, $G_{\text{total}} = G_{\text{HFET}}G_{\text{post}}$. Although it is difficult to have a direct, accurate measurement of G_{post} during data taking, looking back to the definition of T_{sys} (Eq.3.16), it can be seen that both $T_{\text{sys,off}}$ and $T_{\text{sys,on}}$ contain G_{post} . Therefore, G_{post} can be canceled out by comparing the noise powers with the JPA turned off ($T_{\text{sys,off}}$) and turned on ($T_{\text{sys,on}}$) because

$$T_{\text{sys,on}} = \frac{T_{\text{sys,off}}}{\text{SNRI}}, \tag{3.17}$$

where SNRI refers to the signal-to-noise ratio improvement as

$$\text{SNRI} = \frac{G_{\text{JPA}}P_{\text{noise,out,off}}}{P_{\text{noise,out,on}}}. \tag{3.18}$$

The noise power coming out of the system with either the JPA on $P_{\text{noise,out,on}}$ or off $P_{\text{noise,out,off}}$ is measured regularly during data taking, and so is G_{JPA} , providing timely updates of the SNRI. To estimate $T_{\text{sys,on}}$ using the SNRI, however, $T_{\text{sys,off}}$ needs to be known first. Therefore, it is worthwhile to carefully trace both the JPA active and inactive models.

3.2.2 Model Simplification

As alluded to in the previous section, the noise model can be simplified because all of the circulators are thermalized to the same temperature T_{circ} , and we can combine the individual, identical, circulator attenuations into α_1 and α_2 as introduced in Eq. 3.16. In addition to Eq. 3.16, T_{sys} can also be decomposed as follows.

When the JPA is off,

$$T_{\text{sys,off}} = T_{\text{stage1}}|\Gamma_{\text{cav}}|^2 + T_{\text{cav}}(1 - |\Gamma_{\text{cav}}|^2) + \frac{T_{\text{circ}}(1 - \alpha_1\alpha_2)}{\alpha_1\alpha_2} + \frac{T_{\text{HFET}}}{\alpha_1\alpha_2}, \quad (3.19)$$

where the notations of the powers P_* from Eq. 3.13 are converted to the noise temperatures $T_* = P_*/k_B b$ for readability.

When the JPA is on and $G_{\text{JPA}} \gg 1$ is reached ($G_{\text{JPA}} \approx G_{\text{JPA}} - 1$),

$$\begin{aligned} T_{\text{sys,on}} = & T_{\text{stage1}}|\Gamma_{\text{cav}}|^2 + T_{\text{cav}}(1 - |\Gamma_{\text{cav}}|^2) + T_{\text{stage1}} + 2\frac{T_{\text{circ}}(1 - \alpha_1)}{\alpha_1} \\ & + \frac{T_{\text{JPA}}}{\alpha_1} + \frac{T_{\text{circ}}(1 - \alpha_2)}{\alpha_1\alpha_2 G_{\text{JPA}}} + \frac{T_{\text{HFET}}}{\alpha_1\alpha_2 G_{\text{JPA}}}. \end{aligned} \quad (3.20)$$

Here, we use T_{JPA} to denote the total extra noise introduced by the JPA which is equal to the sum at both the signal and idler frequencies ($T_{\text{JPA,S}} + T_{\text{JPA,I}}$) since the two sources of noise are not separable.

Hot Load Calibration Models

The primary noise calibration method used in ADMX is the hot load method. For this method, a VTS (from here on out simply called the hot load) with a 50Ω terminator (a typical reactance used for RF transmission lines) is connected into the system via an RF switch (S) as shown in Fig. 3.1. When the receiver is switched to the hot load configuration from the cavity, T_{stage1} is replaced with T_{HL} , and Γ_{cav} with $\Gamma_{\text{HL}} = 1$ in Eq. 3.19 and Eq. 3.20.

Therefore, when the JPA is off, T_{sys} becomes

$$T_{\text{sys,off,HL}} = T_{\text{HL}} + \frac{T_{\text{circ}}(1 - \alpha_1\alpha_2)}{\alpha_1\alpha_2} + \frac{T_{\text{HFET}}}{\alpha_1\alpha_2}. \quad (3.21)$$

Similarly, when the JPA is on, T_{sys} becomes

$$T_{\text{sys,on,HL}} = 2T_{\text{HL}} + 2\frac{T_{\text{circ}}(1 - \alpha_1)}{\alpha_1} + \frac{T_{\text{JPA}}}{\alpha_1} + \frac{T_{\text{circ}}(1 - \alpha_2)}{\alpha_1\alpha_2 G_{\text{JPA}}} + \frac{T_{\text{HFET}}}{\alpha_1\alpha_2 G_{\text{JPA}}}. \quad (3.22)$$

Cavity Calibration Model

The secondary noise calibration method used in ADMX is the cavity method. For this method, the cavity remains connected to the receiver chain, as it would be during normal data taking. When the system is thermalized to the same temperature from the attenuator (A) connected to Circ1 (stage 1) to the signal at the input of the HFET (stage 6), some terms in the T_{sys} models will cancel out and the equation will simplify. The receiver chain is thermalized in this way when the entire system is cooling down or warming up with respect to the same mixing chamber temperature (T_{mixc}). These simplifications require that one is fully off-resonance since the cavity typically takes more time to thermalize than the other components, so the assumptions that all the temperatures are the same are not as accurate on-resonance.

During a cavity calibration with the JPA off,

$$T_{\text{sys,off,mxc}} = \frac{T_{\text{mixc}} + T_{\text{HFET}}}{\alpha_1\alpha_2}. \quad (3.23)$$

During a cavity calibration with the JPA on,

$$T_{\text{sys,on,mxc}} = \frac{2T_{\text{mixc}} + T_{\text{JPA}}}{\alpha_1} + \frac{T_{\text{mixc}}(1 - \alpha_2) + T_{\text{HFET}}}{\alpha_1\alpha_2 G_{\text{JPA}}}. \quad (3.24)$$

Since the recorded physical temperature is T_{mixc} in Eq. 3.23 and Eq. 3.24, the quantities of interest that can be measured from cavity calibration data are the bare amplifier noise terms without the inverse scaling of transmissivity (α_1, α_2). That is, the cavity calibration models can be used to fit out just T_{HFET} and T_{JPA} , which can provide extra understanding of the

system when combined with the hot load fit results, which do include the inverse scaling of transmissivity. In other words, performing both types of calibration enables the calculation of the in-experiment transmissivity which can be compared with pre-experimental data.

However, it's common to have temperature gradients $\mathcal{O}(0.1T_{\text{mxc}})$ among the components that should be well-thermalized to T_{mxc} . For the JPA-off case, Eq. 3.23 is still practical when using the off-resonance data because T_{HFET} is more than an order of magnitude larger than the other noise contributions in Eq. 3.19. For the JPA-on case, the temperature gradients $\mathcal{O}(0.1T_{\text{mxc}})$ are so large that Eq. 3.24 fails the ideal assumption, and Eq. 3.20 is used for analysis instead.

Fit Functions and Gain Correction

Compared to previous ADMX runs [46, 33, 91], a stronger thermal link is used to connect the hot load to the milliKelvin space ($80 \sim 100$ mK) for Run 1D. The base temperature of the hot load reached $140 \sim 170$ mK, which was finally cold enough to perform JPA-on noise calibration measurements without saturating the JPA. Therefore, Run 1D is the first time a JPA-on noise calibration has been performed in the ADMX experiment [96, 63]. All the noise calibrations described in the following section are Y-factor measurements where the output powers are traced as a function of the physical temperatures, and a linear-fit is used to extract out the extra electronic noise introduced by the different components in the receiver chain.

For the JPA-off measurements the fit function is of the form

$$P_{\text{off}} = C(T + T_{\text{fit}}). \quad (3.25)$$

Here, P_{off} is the output power with the JPA off, T is the physical temperature that is being changed, T_{fit} is the noise temperature that is fitted out for each measurement and C is a constant.

For the JPA-on measurements, the power must be corrected due to inevitable

fluctuations in JPA gain during the course of the measurement, as the JPA is much more sensitive to changes in temperature than the HFET is. To achieve this, the left hand side of the fit function is changed from P_{on} to $(P_{\text{on}} - P_{\text{off}})/G_{\text{JPA}}$. Doing so gives

$$\begin{aligned} \frac{P_{\text{on}} - P_{\text{off}}}{G_{\text{JPA}}} \propto T_{\text{sys,on}} - \frac{T_{\text{sys,off}}}{G_{\text{JPA}}} &= (T_{\text{stage1}}|\Gamma_{\text{cav}}|^2 + T_{\text{cav}}(1 - |\Gamma_{\text{cav}}|^2))(1 - \frac{1}{G_{\text{JPA}}}) \\ &+ T_{\text{stage1}} + 2\frac{T_{\text{circ}}(1 - \alpha_1)}{\alpha_1} + \frac{T_{\text{JPA}}}{\alpha_1} \\ &- \frac{T_{\text{circ}}(1 - \alpha_1)}{\alpha_1 G_{\text{JPA}}}. \end{aligned} \quad (3.26)$$

For sufficiently high gain ($1 - \frac{1}{G_{\text{JPA}}} \simeq 1$ and $\frac{T_{\text{circ}}(1 - \alpha_1)}{\alpha_1 G_{\text{JPA}}} \simeq 0$), the model used for gain-corrected measurements is

$$\begin{aligned} T_{\text{sys,on}} - \frac{T_{\text{sys,off}}}{G_{\text{JPA}}} \simeq &T_{\text{stage1}}|\Gamma_{\text{cav}}|^2 + T_{\text{cav}}(1 - |\Gamma_{\text{cav}}|^2) + T_{\text{stage1}} \\ &+ 2\frac{T_{\text{circ}}(1 - \alpha_1)}{\alpha_1} + \frac{T_{\text{JPA}}}{\alpha_1}, \end{aligned} \quad (3.27)$$

which is the same as Eq. 3.20 without the final two terms.

For the JPA-on-hot-load measurement specifically, this equation simplifies further as was done earlier in this chapter, resulting in

$$T_{\text{sys,on,HL}} - \frac{T_{\text{sys,off,HL}}}{G_{\text{JPA}}} = 2T_{\text{HL}} + 2\frac{T_{\text{circ}}(1 - \alpha_1)}{\alpha_1} + \frac{T_{\text{JPA}}}{\alpha_1}, \quad (3.28)$$

which is the same as Eq. 3.22 without the final two terms.

That being said, the form of the fit function for the JPA-on measurements is nearly identical to the JPA-off function, aside from the gain correction to the output power on the left hand side:

$$\frac{P_{\text{on}} - P_{\text{off}}}{G_{\text{JPA}}} = C(T + T_{\text{fit}}), \quad (3.29)$$

where P_{off} (P_{on}) is the output power with the JPA off (on), and T , T_{fit} and C have the same definitions as in the JPA-off fit function.

3.3 Noise Measurements

In the following sections I go into more detail about the methods for taking and fitting data for each type of measurement, using examples taken at the same frequency, 1280 MHz, to easily compare between the different types. The calibrations done at 1280 MHz, as well as the temperatures of various components throughout the course of each measurement, are summarized Table 3.1. The error bars shown in Figures 3.4, 3.6, 3.8, and 3.10 are purely statistical to reflect data quality, and as such, do not include systematic error. The primary source of systematic error in each fitting is calibrated ruthenium oxide temperature sensors which have a known offset of about ± 5 mK, but this source of error gets canceled out in the calculation of T_{sys} so I do not include it here. I do however consider error introduced by varying the temperature windows used for fitting when reporting the final results in Table 3.3.6.

Type	Date	T_{cav} (K)	T_{circ} (K)	T_{A} (K)	T_{mxc} (K)	T_{HL} (K)
JPA-off-HL	01/04	0.147	0.106	0.090	0.091	0.154–0.518
JPA-off-HL	01/30	0.142	0.102	0.85	0.86	0.156–0.513
JPA-on-HL	02/14	0.130	0.095	0.079	0.079	0.152–0.210
JPA-on-HL	06/13	0.140	0.100	0.081	0.081	0.171–0.538
JPA-off-cav	03/30	1.5–0.72	1.09–0.63	1.06–0.62	1.07–0.60	N/A
JPA-off-cav	06/14	5–0.72	5–0.66	5–0.65	4.2–0.64	N/A
JPA-on-cav	04/30	0.151–0.195	0.119–0.180	0.102–0.162	0.102–0.162	N/A

Table 3.1: Summary of the dates and relevant insert temperatures associated with each of the noise calibration measurements for Run 1D.

3.3.1 JPA off Hot Load

A JPA-off-hot-load calibration is performed with the cryogenic switch flipped to connect the receiver chain to the hot load, with the JPA powered off. Under these circumstances the JPA behaves as a mirror, thus the first-stage amplifier is the HFET. Therefore, the HFET is the dominant contributor to the receiver noise. As such, the data are expected to follow the model described in Eq. 3.21, and the fit function will take the form of Eq. 3.25. Here, $T = T_{\text{HL}}$ and $T_{\text{fit}} = \frac{T_{\text{circ}}(1-\alpha)}{\alpha} + \frac{T_{\text{HFET}}}{\alpha}$ where $\alpha = \alpha_1\alpha_2$. Going forward, I will refer to this fit result as the *effective* HFET noise, $T_{\text{HFET}}/\alpha_{\text{eff}}$, due to the scaling of $1/\alpha$, as well as the inclusion of the circulator term, whereas the intrinsic HFET noise is simply T_{HFET} .

Now, I will detail the process of performing this type of calibration. First, the switch is flipped from the cavity state to the hot load state. Before beginning heating a few baseline digitizations and hot load temperature readings are taken that can be used for estimating the statistical uncertainty of the digitizer and temperature sensor, which will be used as inputs to the fitting software. This is typically done for a few minutes, until $\sim 5 - 10$ data points are collected. An example of this data is shown in Fig. 3.2.

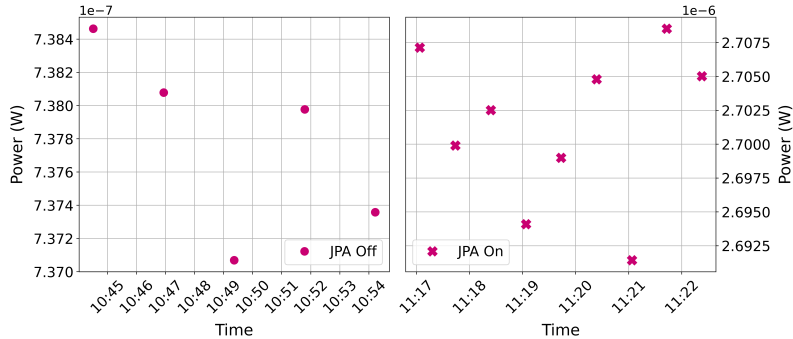


Figure 3.2: Left: Power data taken before beginning heating for a JPA-off-hot-load measurement to estimate the digitizer uncertainty with the JPA off (P_{off}). The example shown here corresponds to an uncertainty of $8e-10$ W. Right: Power data taken before beginning heating for a JPA-on-hot-load measurement to estimate the digitizer uncertainty with the JPA on (P_{on}). The example shown here corresponds to an uncertainty of $8e-9$ W.

After taking baseline measurements, heating can begin. For a given JPA-off-hot-load calibration, data are typically collected over a frequency range of 20-40 MHz with a spacing of 2 MHz between each data point. The number of points will influence how long the wait time between each heating step will be so that there is sufficient time to get at least a few (~ 3) digitizations for each frequency point during each heating step. For JPA-off-hot-load measurements, there are a total of 11 heater steps that bring the hot load smoothly from its base temperature of ~ 150 mK up to a maximum temperature of ~ 500 mK, and then back down. Examples of temperature and power vs. time data can be seen in Fig. 3.3. After collecting the data, the hot load temperature and the output power data are plotted against each other and a linear fit (Eq. 3.25) is performed. The JPA-off-hot-load data and corresponding fits are shown in Fig. 3.4. The January 4th (January 30th) JPA-off-hot-load fit resulted in $T_{\text{HFET}}/\alpha_{\text{eff}} = 6.17 \pm 0.75$ (6.13 ± 0.20) K, after accounting for systematic error introduced from adjusting the fitting window. Averaging the two effective HFET noise temperatures results in $T_{\text{HFET}}/\alpha_{\text{eff}} = 6.15 \pm 0.39$ K.

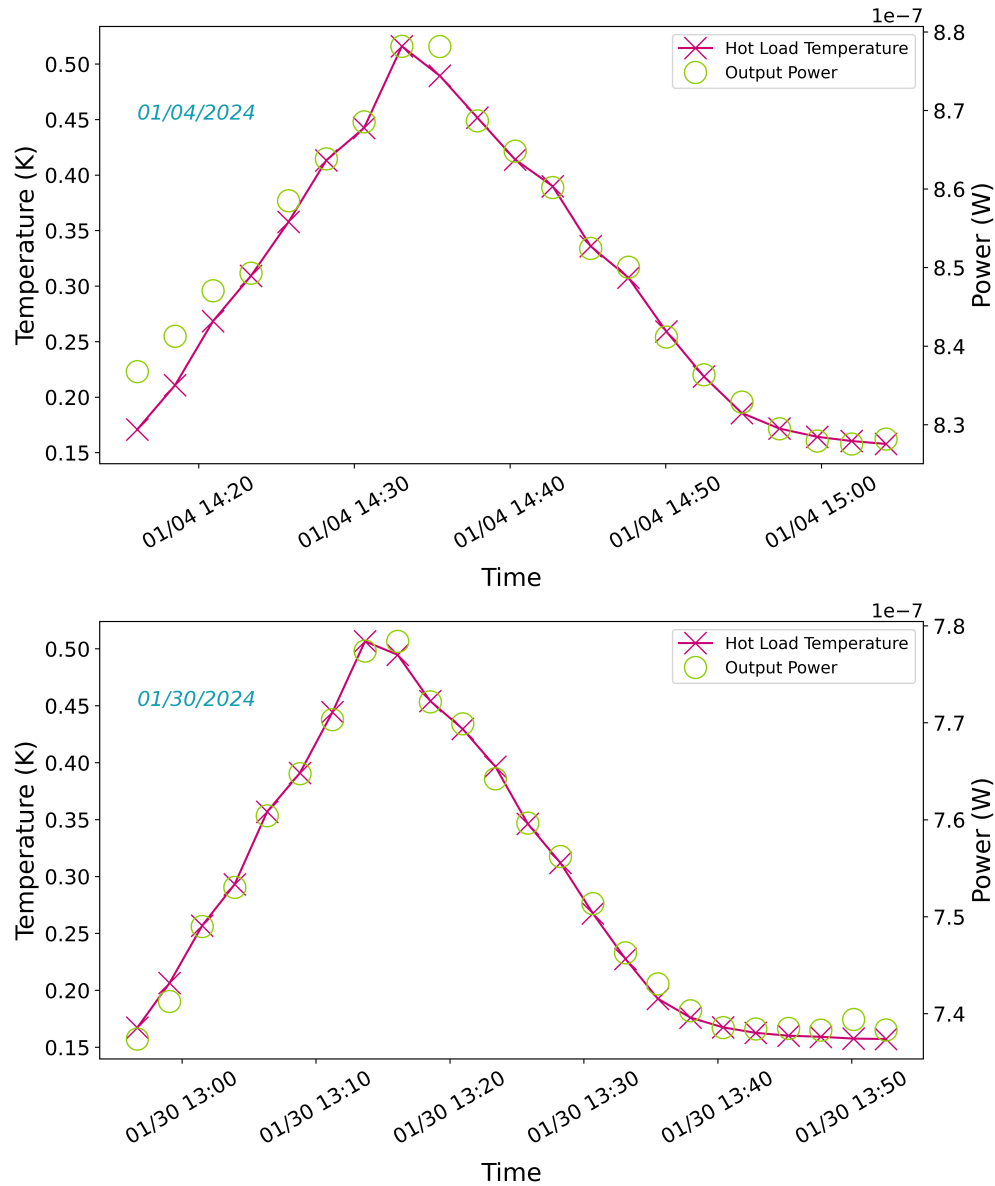


Figure 3.3: Here I plot the hot load temperature and the output power changing together during the course of two JPA-off-hot-load measurements done at 1280 MHz. The temperature is smoothly increased using a DC power supply from $\sim 150 - 500$ mK and then smoothly decreased using the same voltage steps. The output power tracks the changing temperature closely, as expected from the linear relationship described in Eq. 3.21.

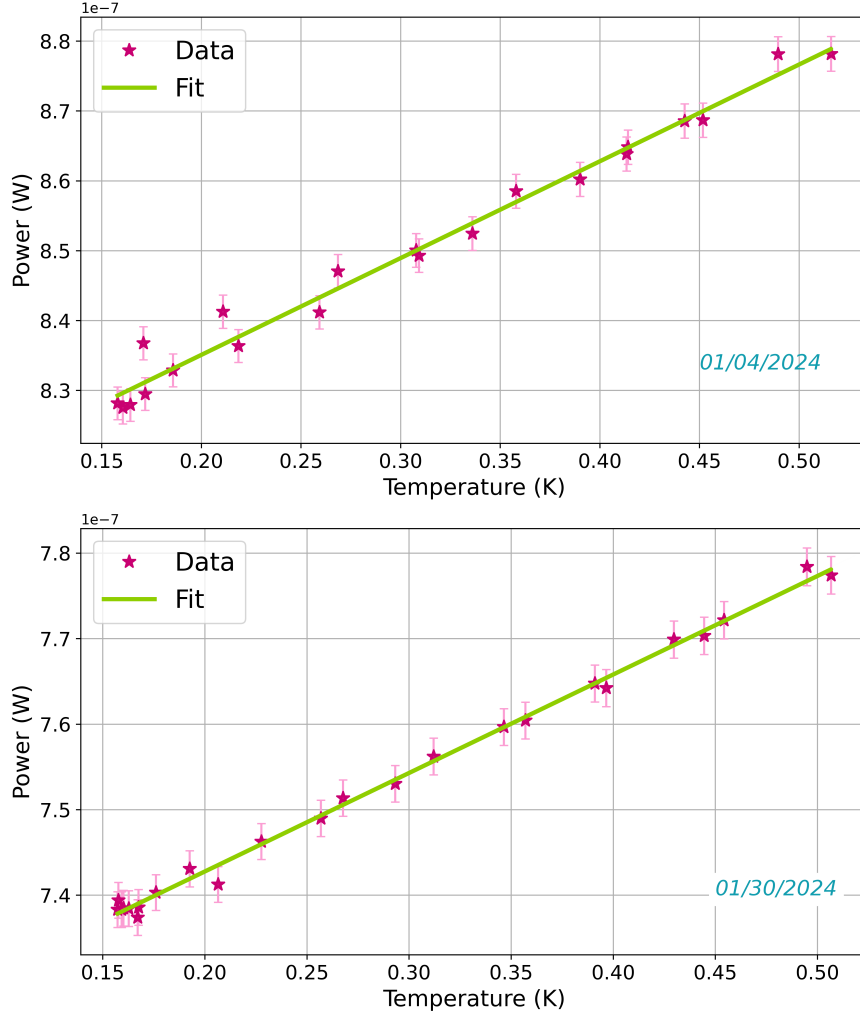


Figure 3.4: Here the hot load temperature is plotted against the output power for the two JPA-off-hot-load measurements done at 1280 MHz. The data are fit to Eq. 3.25 where $T = T_{\text{HL}}$ and $T_{\text{fit}} = T_{\text{HFET}}/\alpha_{\text{eff}} = \frac{T_{\text{circ}}(1-\alpha)}{\alpha} + \frac{T_{\text{HFET}}}{\alpha}$. The fits shown here used the entire data range, however the fits were performed multiple times using slightly different fitting windows to estimate systematic uncertainty introduced from a given choice of fitting window. After taking into account differences in the fit results caused by adjusting the fitting window, I found $T_{\text{HFET}}/\alpha_{\text{eff}} = 6.17 \pm 0.75$ (6.13 ± 0.20) K for the January 4th (January 30th) measurements. Both values are in agreement well within 1σ , indicating high stability in the HFET noise performance over time. Averaging the two effective HFET noise temperatures results in $T_{\text{HFET}}/\alpha_{\text{eff}} = 6.15 \pm 0.39$ K.

3.3.2 JPA on Hot Load

A JPA-on-hot-load calibration is performed with the cryogenic switch flipped to connect the receiver chain to the hot load, with the JPA powered on. Under these circumstances, the JPA is the first-stage amplifier. Due to the suppression of the HFET noise by $1/G_{\text{JPA}}$, the JPA is now the dominant contributor to the receiver noise. Since the JPA is a narrowband amplifier, unlike the HFET or the room temperature amplifiers, its optimal bias parameters are more sensitive to environmental disturbances such as changes in temperature or mechanical vibrations. As a result, the JPA-on measurements require some additional steps due to the instability of the JPA gain caused by heating. In addition to the procedure having more steps, the model itself is more complicated (Eq.3.28) because the output power must be corrected before fitting to account for inevitable gain fluctuations throughout the measurement. The fit function will now take the form of Eq. 3.29, where $T = 2T_{\text{HL}}$ and $T_{\text{fit}} = 2\frac{T_{\text{circ}}(1-\alpha_1)}{\alpha_1} + \frac{T_{\text{JPA}}}{\alpha_1}$. The factor of 2 present in the definition of T comes from the addition of the idler mode noise power when the JPA is on. For the same reason as described in Sec. 3.3.1, we refer to this fit result as the *effective* JPA noise, $T_{\text{JPA,eff}}$, where the intrinsic JPA noise is simply T_{JPA} as defined in Sections 3.1.6 and 3.2.2.

The procedure for a JPA-on-hot-load measurement is as follows. The switch is again flipped from the cavity to the hot load. The flipping of this switch introduces a non-negligible amount of heat to the JPA, which is sensitive to sudden changes in temperature. Therefore, the system is left to settle for a few minutes so that the JPA temperature can level out before adjusting the bias current and pump power to achieve a decent JPA gain. After finding parameters that give a reasonably high gain, the JPA is then tested for stability. This is done by varying the bias current and pump power over a small range and observing how much the gain changes. If the gain changes too much with slight variations in the optimization parameters it will likely be unstable during heating. In this case, the process of manually searching for a new set of parameters to test the stability of is repeated until a satisfactory level of stability is achieved. The most stable bias settings for each measurement

were: $G_{\text{JPA,February}} = 15.8 \pm 0.1$ dB, $I_{\text{bias,February}} = -0.183$ mA, $P_{\text{pump,February}} = -7.35$ dBm and $G_{\text{JPA,June}} = 18.1 \pm 0.8$ dB, $I_{\text{bias,June}} = -1.647$ mA, $P_{\text{pump,June}} = -8.47$ dBm. Performing the measurement at two different gains allows us to test whether the effective JPA noise varies under different bias settings, which is important to understand due to the ever-changing bias settings during real data taking. After settling on a set of bias parameters, the gain is monitored over time before introducing any heat to test the stability further. During this time, digitizations are also taken to get an estimate of the digitizer uncertainty with the JPA on as was done for the JPA-off-hot-load measurement (see Fig. 3.2). Finally, the load is heated in incremental steps as was done for the JPA-off-hot-load calibration. However, because the JPA is a narrowband amplifier, data are taken at only one frequency at a time. Therefore, the wait time can be shorter and there can be a higher density of points for the JPA-on-hot-load calibrations. A total of 9 heater steps are used to bring the hot load from its base temperature of ~ 150 mK to a maximum temperature of ~ 200 mK, and then back down. We do not take the hot load much higher than 200 mK due to the temperature sensitivity of the JPA. Above 200 mK the amplifier can easily become saturated and/or lose gain performance, making the noise calibration results useless. Fortunately, a temperature change of only ~ 50 mK is sufficient to measure the amplifier noise. Examples of temperature and power vs. time data can be seen in Fig. 3.5. After collecting the data, the hot load temperature and the output power data are plotted against each other and a linear fit (Eq. 3.29) is performed. The data and corresponding fits are shown in Fig. 3.6. The February (June) JPA-on-hot-load fit resulted in $T_{\text{JPA,eff}} = 0.139 \pm 0.021$ (0.143 ± 0.019) K, after accounting for systematic error introduced from adjusting the fitting window. The two values agree well within 1σ . This indicates that the JPA noise performance was extremely stable over the course of multiple months, as well as under different bias settings. Averaging the two effective JPA noise temperatures results in $T_{\text{JPA,eff}} = 0.141 \pm 0.014$ K.

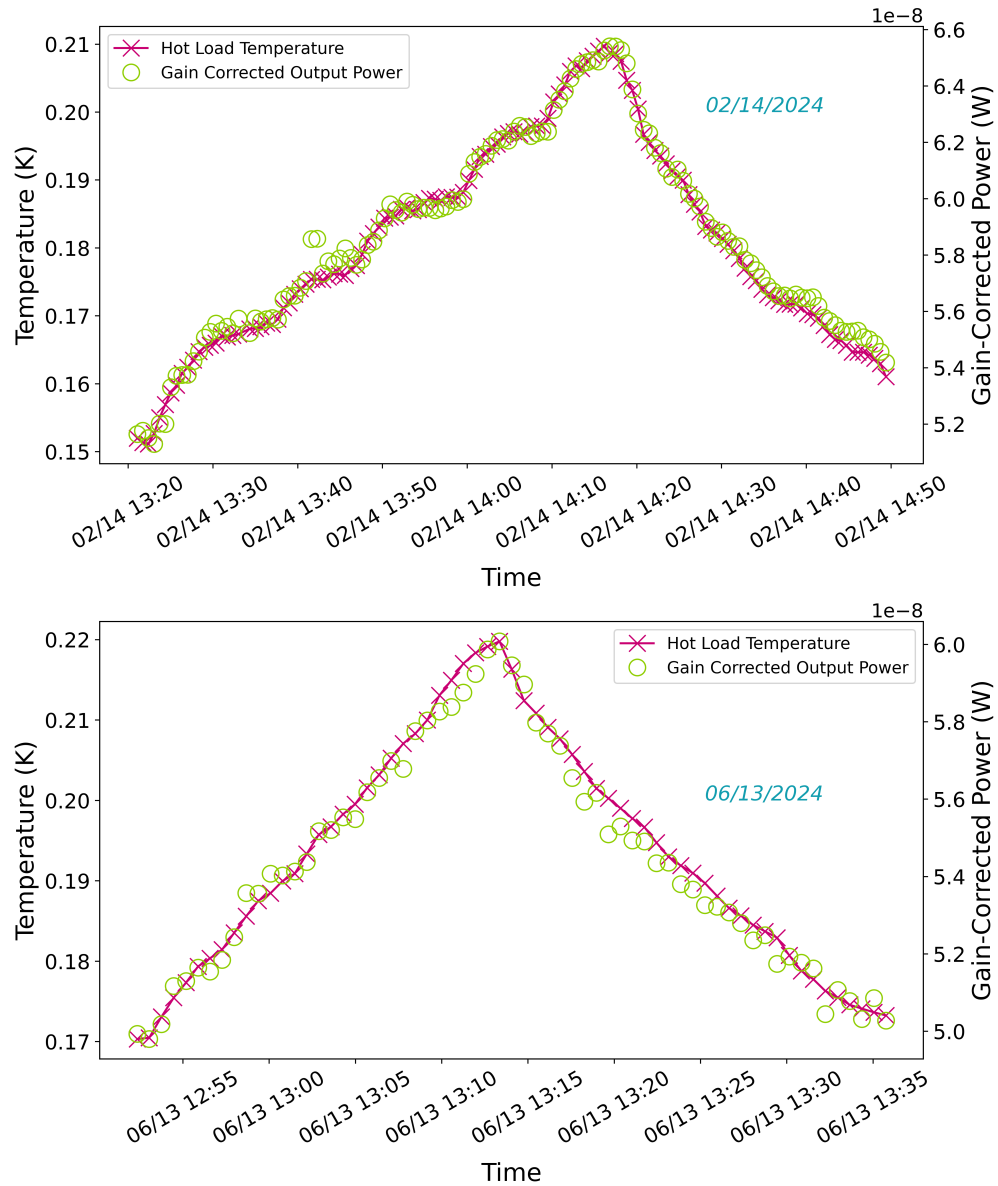


Figure 3.5: Here I plot the hot load temperature and the output power changing together during the course of two JPA-on-hot-load measurements done at 1280 MHz. The temperature is smoothly increased using a DC power supply from $\sim 150 - 200$ mK and then smoothly decreased using the same voltage steps. The output power tracks the changing temperature closely, as expected from the linear relationship described in Eq. 3.28.

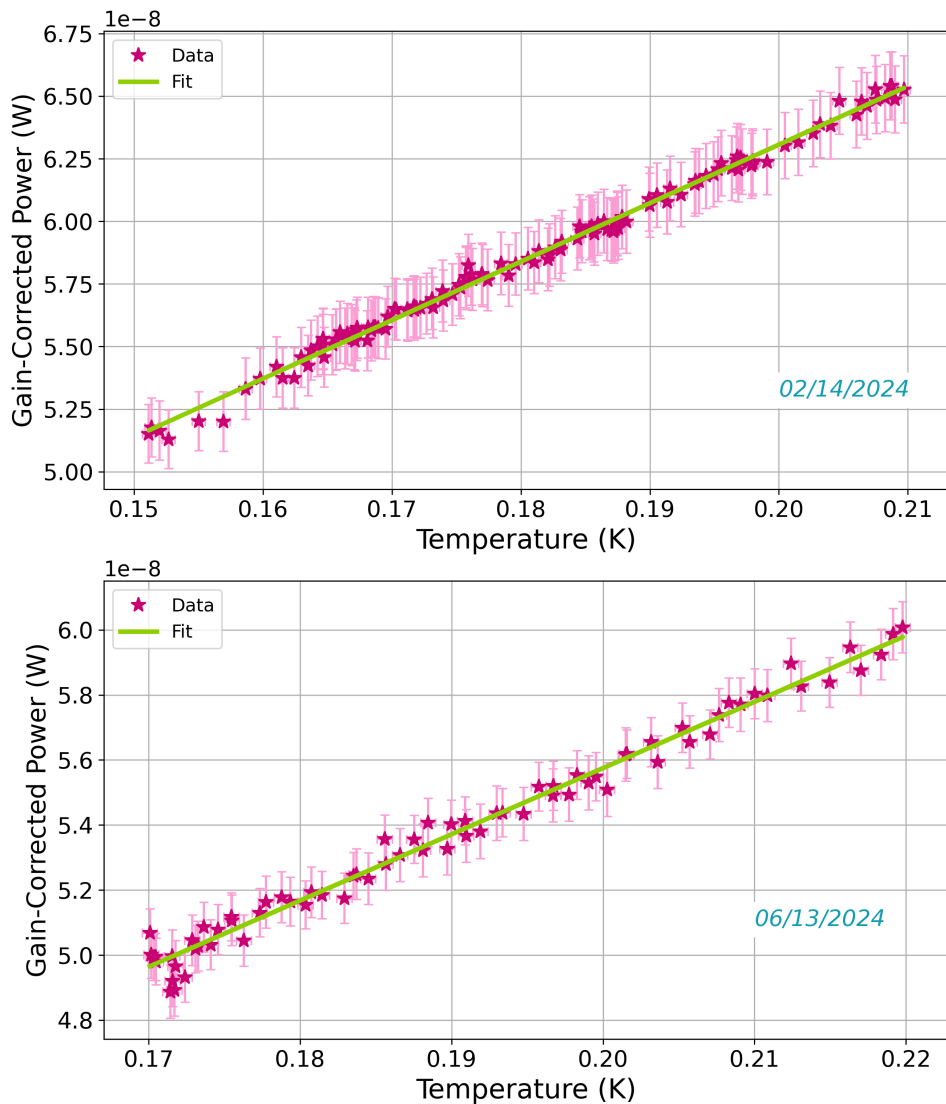


Figure 3.6: Here the hot load temperature is plotted against the output power for the two JPA-on-hot-load measurements done at 1280 MHz. The data are fit to Eq. 3.29 where $T = 2T_{\text{HL}}$ and $T_{\text{fit}} = T_{\text{JPA,eff}} = 2\frac{T_{\text{circ}}(1-\alpha_1)}{\alpha_1} + \frac{T_{\text{JPA}}}{\alpha_1}$. The fit shown here used the entire data range, but the fit was performed multiple times using slightly different fitting windows. After taking into account differences in the fit result caused by adjusting the fitting window, I found $T_{\text{JPA,eff}} = 0.139 \pm 0.021$ (0.143 ± 0.019) K for the February (June) measurements. Both values are in agreement within error bars, indicating stability in the JPA noise both over time and under different bias settings. Averaging the two effective JPA noise temperatures results in $T_{\text{JPA,eff}} = 0.141 \pm 0.014$ K.

3.3.3 JPA off Cavity

A JPA-off-cavity calibration is performed with the cryogenic switch (S) flipped to connect the receiver chain to the cavity, with the JPA powered off. Under these circumstances, the JPA behaves as a mirror, so the first-stage amplifier is the HFET. Therefore, as with the JPA-off-hot-load measurement, the HFET is the dominant contributor to the receiver noise. Thus, the data are expected to follow the model described in Eq. 3.23 and the fit function will take the form of Eq. 3.25. In this scenario, $T = T_{\text{mixc}}$ ($T = T_{\text{circ}}$) and $T_{\text{fit}} = T_{\text{HFET}}$. This type of calibration, unlike the hot load calibrations, cannot be done easily during regular data taking. This is because it requires the entire system to be cooling down or heating up at once, which typically only occurs at the beginning and end of each run, or if something goes awry during data taking warranting an emergency warm-up such as the magnet unexpectedly losing power. Therefore, the procedure is not as well defined, but the concept is similar to the hot load calibrations. That is, the output power and record relevant insert temperatures (T_A , T_{mixc} , T_{circ} , T_{cav}) are measured simultaneously as the entire system changes temperature. For a given JPA-off-cavity calibration, we typically take data over the entire frequency range (120 MHz for Run 1D) with a spacing of 4-5 MHz, resulting in ~ 30 individual frequencies. This is done because of the infrequent and slow nature of these measurements, enabling full coverage of the Run 1D frequency range while we have the opportunity. The temperature range covered was $\sim 0.6 - 1.1$ K and $\sim 0.6 - 5$ K for the March and June calibrations respectively. An example of temperature and power vs. time data can be seen in Fig. 3.7.

To reiterate and expand on the discussion in Sec. 3.2.2, the simplified fit function used here, $T_{\text{sys,off,mxc}} = \frac{1}{\alpha_1\alpha_2}(T_{\text{mixc}} + T_{\text{HFET}})$, requires the assumption that $T_{\text{mixc}} = T_{\text{cav}} = T_{\text{circ}}$ on-resonance and $T_{\text{mixc}} = T_A = T_{\text{circ}}$ off-resonance. Looking at Table 3.1 one can see that $T_{\text{mixc}} = T_A = T_{\text{circ}}$ is a reasonable assumption, however, $T_{\text{mixc}} = T_{\text{cav}} = T_{\text{circ}}$ is not as accurate. This is because the cavity typically takes longer to thermalize to the mixing chamber than other components in the system. Therefore, only the off-resonance JPA-off-cavity data are presented in this analysis to ensure the fit of T_{HFET} is as accurate as possible.

Likewise, the error for each power point is determined by calculating the standard error of the mean using only the off-resonance data for each spectrum. For calculating the errors in temperature sensors we choose a relatively stable period of time before the temperatures began changing to estimate the uncertainty on each temperature sensor. The remaining effects of the assumption that all the components are well thermalized are taken into account by monitoring both T_{mxc} and T_{circ} simultaneously, performing the fit twice using each sensor as the input, and averaging the fit results. This process is discussed further in the caption of Fig. 3.8.

As one can see from Eq. 3.23, doing this type of measurement makes it possible to fit out T_{HFET} on its own, without the $1/\alpha$ scaling present in the JPA-off-hot-load measurement (Eq. 3.21). This allows for two useful checks of the RF system. Firstly, the measured value of T_{HFET} can be compared to the datasheet from Low Noise Factory (LNF) [4] to confirm that the HFET is working as intended. Secondly, the results of this fit can be combined with the JPA-off-hot-load fit result (Sec. 3.3.3) to back out the total transmissivity, α , which can then be compared to measurements of α done prior to data taking. This is possible because the circulator term in $T_{\text{HFET}}/\alpha_{\text{eff}}$ is less than 1% of the magnitude of the HFET term in Eq.3.21, such that $T_{\text{HFET}}/\alpha_{\text{eff}} \approx T_{\text{HFET}}/\alpha$. The March (June) JPA-off-cavity fit resulted in $T_{\text{HFET}} = 4.18 \pm 0.26$ (3.80 ± 1.02) K, after accounting for systematic error introduced from adjusting the fitting window, as well as from using both $T = T_{\text{mxc}}$ and $T = T_{\text{circ}}$. These values agree within 1σ , indicating the HFET noise performance was stable over time, in agreement with the results for the JPA-off-hot-load calibration. The averaged result for the HFET noise temperature from both calibrations is $T_{\text{HFET}} = 3.99 \pm 0.52$ K, which is in agreement with the nominal noise level from the LNF datasheet of $T_{\text{HFET}} \approx 3.75$ K. Combining this value with the fit result from Sec. 3.3.1 ($T_{\text{HFET}} = 6.15 \pm 0.39$) gives a total transmissivity of $\alpha = 0.65 \pm 0.09$. The total transmissivity measured prior to data taking was $\alpha = 0.643 \pm 0.003$. These two independent measurements agree within 1σ , indicating that the models used in this analysis accurately represent the noise behavior of the receiver chain.

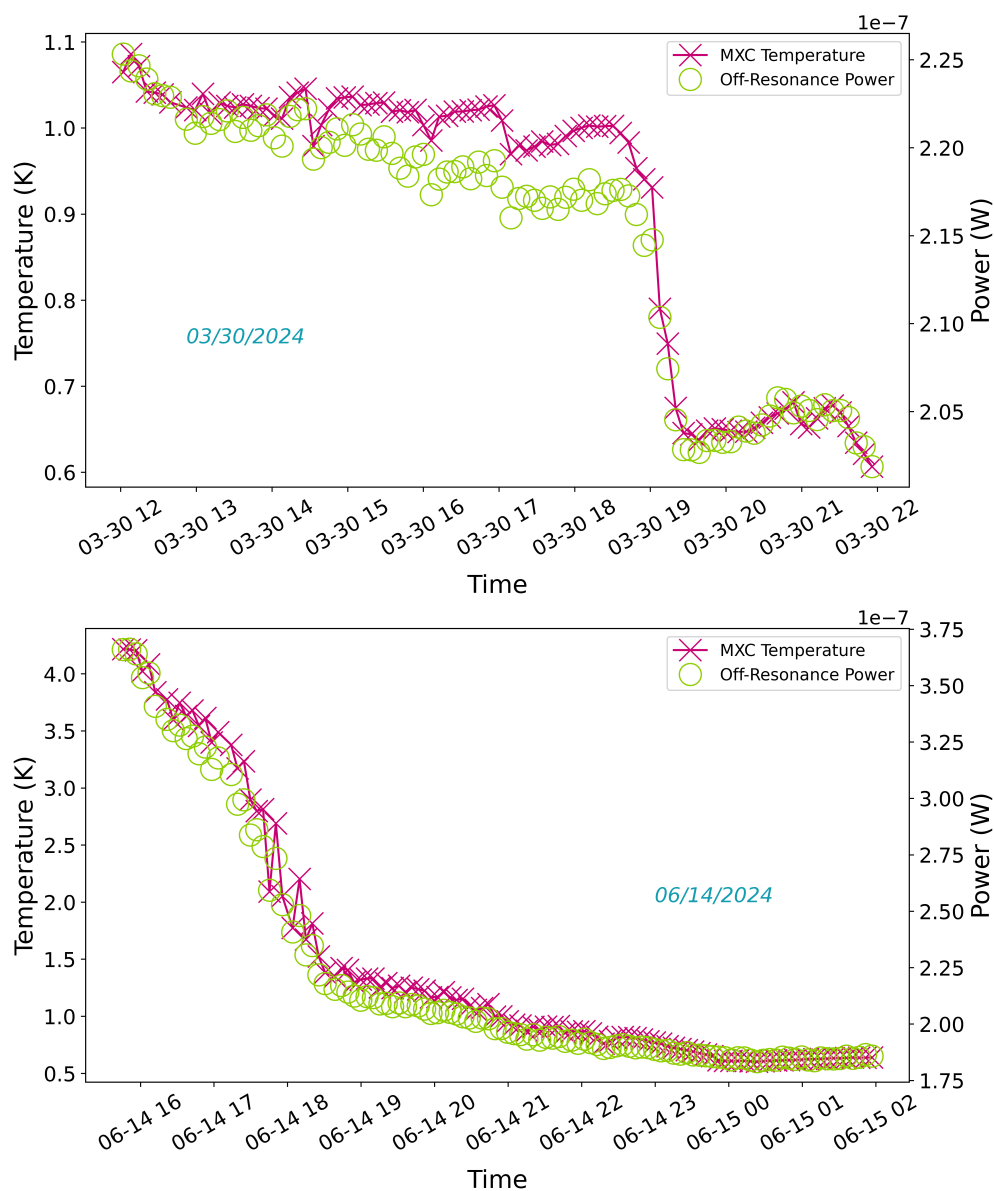


Figure 3.7: Here I plot the mixing chamber temperature and the output power changing together during the course of two JPA-off-cavity measurements done at 1280 MHz. The temperature is tracked during the course of a full system cool-down. The output power tracks the changing temperature closely, as expected from the linear relationship described in Eq. 3.23.

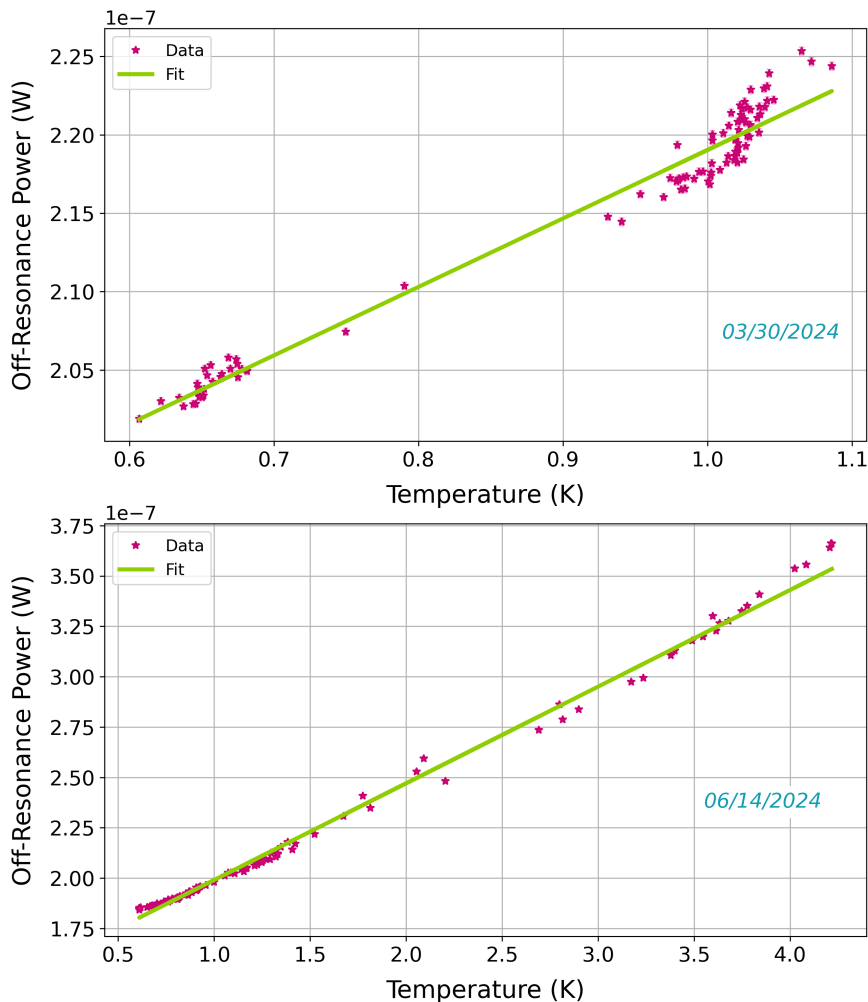


Figure 3.8: Here the mixing chamber temperature is plotted against the off-resonance output power for the two JPA-off-cavity measurements done at 1280 MHz. The data are fit to Eq. 3.25 where $T = T_{\text{mxc}}$ is the hot load temperature and $T_{\text{fit}} = T_{\text{HFET}}$. The fits shown here used the entire data range, but the fits were performed multiple times using slightly different fitting windows. In addition to adjusting the fitting window, the fits were performed using $T = T_{\text{circ}}$ instead because of the imperfect assumption that $T_{\text{mxc}} = T_{\text{circ}}$. After taking into account differences in the fit results caused by adjusting the fitting window and using a different temperature sensor, I found $T_{\text{HFET}} = 4.18 \pm 0.26$ (3.80 ± 1.02) K for the March (June) calibration. Both values are in agreement within 1σ , indicating stability in the HFET noise over time. The value used for calculating α is an average of the two: $T_{\text{HFET}} = 3.99 \pm 0.52$ K.

3.3.4 JPA on Cavity

A JPA-on-cavity calibration is performed with the cryogenic switch (S) flipped to connect the receiver chain to the cavity, with the JPA powered on. Under these circumstances, the JPA is the first-stage amplifier and, consequently, the dominant contribution to the receiver noise. Thus, the data are expected to follow the model described in Eq. 3.24. However, as discussed in Sec. 3.3.3, Eq. 3.24 requires the assumption that $T_{\text{mxc}} = T_{\text{cav}} = T_{\text{circ}}$ on-resonance and $T_{\text{mxc}} = T_{\text{A}} = T_{\text{circ}}$ off-resonance. For the JPA-off-cavity case, T_{HFET} is large enough compared to the $\mathcal{O}(0.1T_{\text{mxc}})$ temperature gradients between milliKelvin components for the simplified fit function (Eq. 3.23) to still be practical in the off-resonance case. However, for the JPA-on-cavity case, the $\mathcal{O}(0.1T_{\text{mxc}})$ temperature gradient is too large compared to T_{JPA} for Eq. 3.24 to be useful, so instead the full model (Eq. 3.20) must be used. In reality, Eq. 3.27 is used for fitting the JPA-on-cavity data after correcting for inevitable gain fluctuations throughout the measurement - see Sec.3.2.2 for further details. That being said, the fit function will take the form of Eq. 3.29. For the off-resonance data, $|\Gamma_{\text{cav}}|^2 = 1$ and, as a result, $T = 2T_{\text{stage1}} = 2(T_{\text{A}}\sqrt{\alpha_1} + T_{\text{circ}}(1 - \sqrt{\alpha_1}))$ and $T_{\text{fit}} = T_{\text{JPA,eff}}$ as defined in Sec. 3.2 and Sec. 3.3.2 respectively. Note that in the above definition of T_{stage1} it is assumed that the attenuation between attenuator A and the cavity is equal to half the attenuation from the cavity to JPA ($\alpha_{\text{circ1},21} = \sqrt{\alpha_1}$). For the on-resonance data, $0 < |\Gamma_{\text{cav}}|^2 < 1$ so the full definition for T is retained such that $T = T_{\text{stage1}}|\Gamma_{\text{cav}}|^2 + T_{\text{cav}}(1 - |\Gamma_{\text{cav}}|^2) + T_{\text{stage1}}$ and, again, $T_{\text{fit}} = T_{\text{JPA,eff}}$.

For reasons discussed in Sec.3.3.3, the procedure for the cavity calibrations is not as well-defined as the hot load calibrations. Nonetheless, the general concept remains the same: the output power, JPA gain and relevant insert temperatures (T_{A} , T_{mxc} , T_{circ} , T_{cav}) are simultaneously measured as the entire system changes temperature. As was done for the JPA-off-cavity calibrations, the on-resonance and off-resonance sections of the spectra are analyzed in two separate data pipelines. The procedure for estimating errors for the on (off) resonance data is the same as those outlined in Sec.3.3.3.

Due to the narrowband nature of the JPA, this calibration can only be done at a single frequency at a time. Therefore, the JPA-off-cavity calibration was performed at 1280 MHz to allow for direct comparison between the four different methods. The temperature range covered was $\sim 0.1-0.2$ K. An example of temperature and power vs. time data can be seen in Fig. 3.9. The gain varied from roughly 11-19.5 dB throughout the course of the measurement with the JPA bias settings held constant ($I_{\text{bias}} = -1.083$ mA and $P_{\text{pump}} = -6.46$ dBm) as the system warmed. Between 14:30 and 19:30 on 04/30 the gain was changing too rapidly to get a clean gain correction so that data has been removed from the analysis. After collecting the data, the temperature and output power are plotted against each other for both the on-resonance and off-resonance cases and linear fits (3.29) of both data sets are performed. The data and corresponding fits are shown in Fig. 3.10. The on-resonance (off-resonance) JPA-on-cavity fit resulted in $T_{\text{JPA,eff}} = 0.372 \pm 0.018$ (0.372 ± 0.022) K, after accounting for systematic error introduced from adjusting the fitting window. The two values agree well with each other, but, surprisingly, not with the JPA-on-hot-load $T_{\text{JPA,eff}}$ (see Table 3.3.6). Possible reasons for this discrepancy will be discussed further in Sec. 3.3.6.

3.3.5 System Noise Temperature Comparison

Before ADMX had the ability to perform a direct JPA-on noise measurement, the system noise temperature (T_{sys}) was calculated in two steps. The first step would be to directly measure the JPA-off noise $T_{\text{HFET}}/\alpha_{\text{eff}}$, as described in Sec. 3.3.1. Then, we measure the signal-to-noise ratio improvement (SNRI) as defined in Eq. 3.18. We can then combine the results of these two measurements to calculate our system noise temperature using Eq. 3.17. Now with the ability to measure the JPA effective noise, $T_{\text{JPA,eff}}$, directly, the full model defined in Eq. 3.20 can be used to calculate T_{sys} . In this section, I compare the two methods for calculating the system noise temperature.

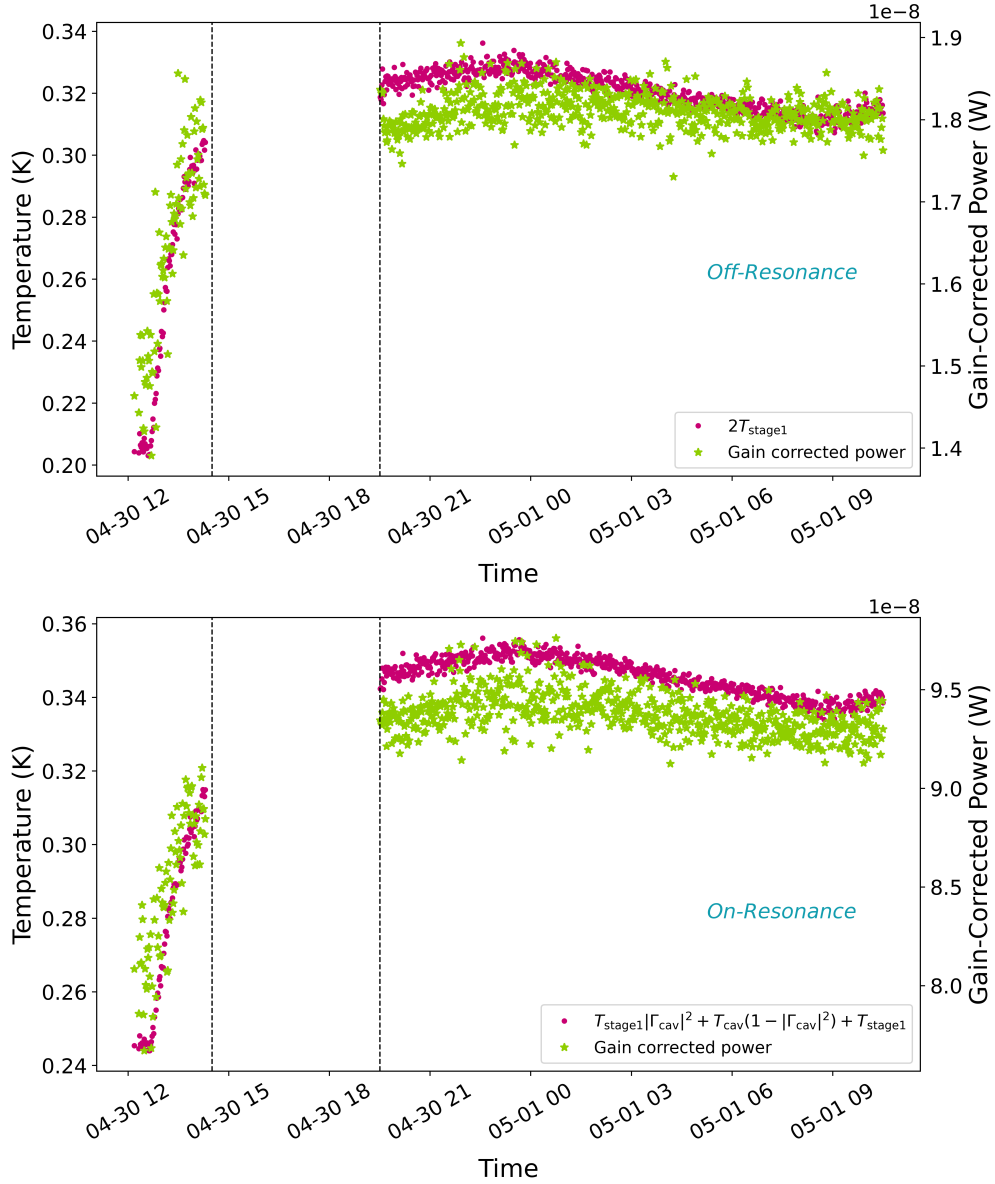


Figure 3.9: Here one can see the relevant temperatures and the output power changing together during the course of the JPA-on-cavity measurement done at 1280 MHz. The off-resonance fit uses $T = 2T_{\text{stage1}}$, so that is what is plotted here. Meanwhile, the on-resonance fit uses $T = T_{\text{stage1}}|\Gamma_{\text{cav}}|^2 + T_{\text{cav}}(1 - |\Gamma_{\text{cav}}|^2) + T_{\text{stage1}}$, so that is what is plotted for the on-resonance case. The temperatures are tracked during the course of a full system warm-up. The output power tracks the changing temperature closely in both cases, as expected from the linear relationship described in Eq. 3.27. The black dashed lines mark where data were removed due to rapidly changing gain resulting in poor gain correction.

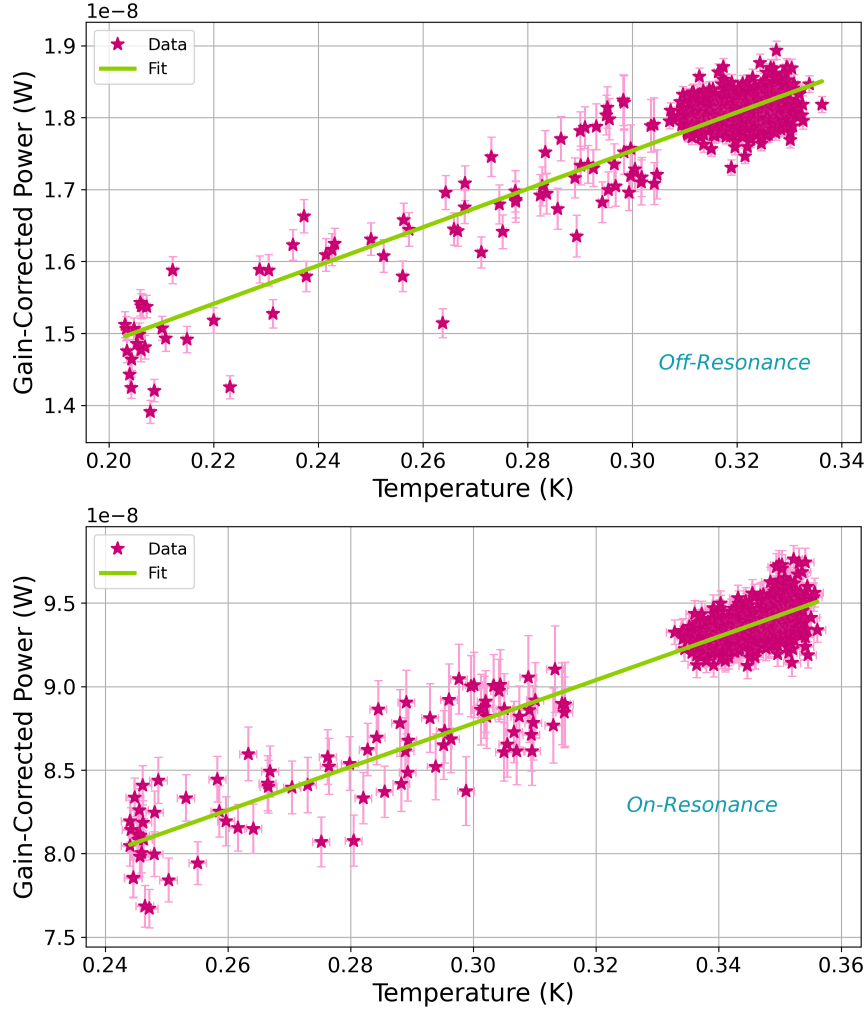


Figure 3.10: Here the relevant temperatures are plotted separately against the off-resonance and on-resonance output power for the JPA-on-cavity measurement done at 1280 MHz. The data are fit to Eq. 3.29 where $T = 2T_{\text{stage1}}$ ($T = T_{\text{stage1}}|\Gamma_{\text{cav}}|^2 + T_{\text{cav}}(1 - |\Gamma_{\text{cav}}|^2) + T_{\text{stage1}}$) for the off-resonance (on-resonance) data and $T_{\text{fit}} = T_{\text{JPA,eff}} = 2\frac{T_{\text{circ}}(1-\alpha_1)}{\alpha_1} + \frac{T_{\text{JPA}}}{\alpha_1}$. The fits shown here used the entire data range, but the fits were performed multiple times using slightly different fitting windows. After taking into account differences in the fit results caused by adjusting the fitting window, I found $T_{\text{JPA,eff}} = 0.372 \pm 0.022$ (0.372 ± 0.018) K for the off-resonance (on-resonance) data.

An example showing the SNRI and direct JPA-on fit methods for the hot load measurements at 1280 MHz can be seen in Fig. 3.11. Similarly, a comparison between the SNRI and direct JPA fit methods for the cavity measurement at 1280 MHz can be seen in Fig. 3.12. One can see that the results for the system noise temperature are consistent between the two methods for both the hot load and the off resonance cavity measurement, and nearly consistent for the on resonance cavity measurement as well. This serves as confirmation that the SNRI method (Eq. 3.17), which is how we monitor the system noise temperature during regular data taking, does an accurate job of calibrating $T_{\text{sys,on}}$ and properly represents the noise behavior of the receiver chain. Additionally, an estimate of the systematic uncertainty in T_{sys} to use for the axion search analysis of Run 1D can be determined from the slight differences between the two methods.

To trace T_{sys} during data taking, during which the cryo-switch is flipped to the cavity, ADMX uses the SNRI method for its promptness with $T_{\text{HFET}}/\alpha_{\text{eff}}$ as a calibrated input. The JPA-off-hot-load measurement fitting directly provides $T_{\text{HFET}}/\alpha_{\text{eff}}$ (Sec. 3.3.1). The JPA-on-hot-load and JPA-on-cavity measurements need further calculation where

$$T_{\text{HFET}}/\alpha_{\text{eff}} = \frac{1}{G_{\text{JPA}} - \text{SNRI}} \cdot \left(\text{SNRI} \cdot G_{\text{JPA}} \cdot (T_{\text{JPA,eff}} + T_{\text{JPA,on}}) - G_{\text{JPA}} \cdot T_{\text{JPA,off}} - \text{SNRI} \cdot T_{\text{circ}}(1 - \alpha_2)/\alpha_2 \right). \quad (3.30)$$

In the JPA-on-hot-load case, $T_{\text{JPA,on}} = 2T_{\text{HL}}$ and $T_{\text{JPA,off}} = T_{\text{HL}}$. In the JPA-on-cavity case, $T_{\text{JPA,on}} = T_{\text{stage1}}|\Gamma_{\text{cav}}|^2 + T_{\text{cav}}(1 - |\Gamma_{\text{cav}}|^2) + T_{\text{stage1}}$ and $T_{\text{JPA,off}} = T_{\text{stage1}}|\Gamma_{\text{cav}}|^2 + T_{\text{cav}}(1 - |\Gamma_{\text{cav}}|^2)$. For the off-resonance JPA-on-cavity case it is assumed that $|\Gamma_{\text{cav}}|^2 = 1$. Putting the $T_{\text{JPA,eff}}$ hot load calibration results at 1280 MHz reported in Sec. 3.3.2 into the equation above, $T_{\text{HFET}}/\alpha_{\text{eff}} = 6.21 \pm 0.34(6.16 \pm 0.26)$ K for the February (June) measurement. After averaging the two measurements separated by four months, I find $T_{\text{HFET}}/\alpha_{\text{eff}} = 6.18 \pm 0.21$ K. Similarly, putting the $T_{\text{JPA,eff}}$ cavity calibration results at 1280 MHz reported in Sec. 3.3.4 into the equation above, $T_{\text{HFET}}/\alpha_{\text{eff}} = 6.33 \pm 0.21 (6.72 \pm 0.17)$ K off (on) resonance. The inferred $T_{\text{HFET}}/\alpha_{\text{eff}}$ values from both the JPA-on-hot-load and JPA-on cavity calibrations are consistent with the JPA-off-hot-load result $T_{\text{HFET}}/\alpha_{\text{eff}} = 6.15 \pm 0.39$ K.




Figure 3.11: Hot Load T_{sys} comparison at 1280 MHz. Here I plot the comparison between the two methods, using SNRI and $T_{\text{HFET}}/\alpha_{\text{eff}}$ (pink points) and using JPA fit (green points), over the course of both hot load measurements done at 1280 MHz. It is clear that for both the February and June data the system noise temperature we calculate is consistent between the two methods within error bars.

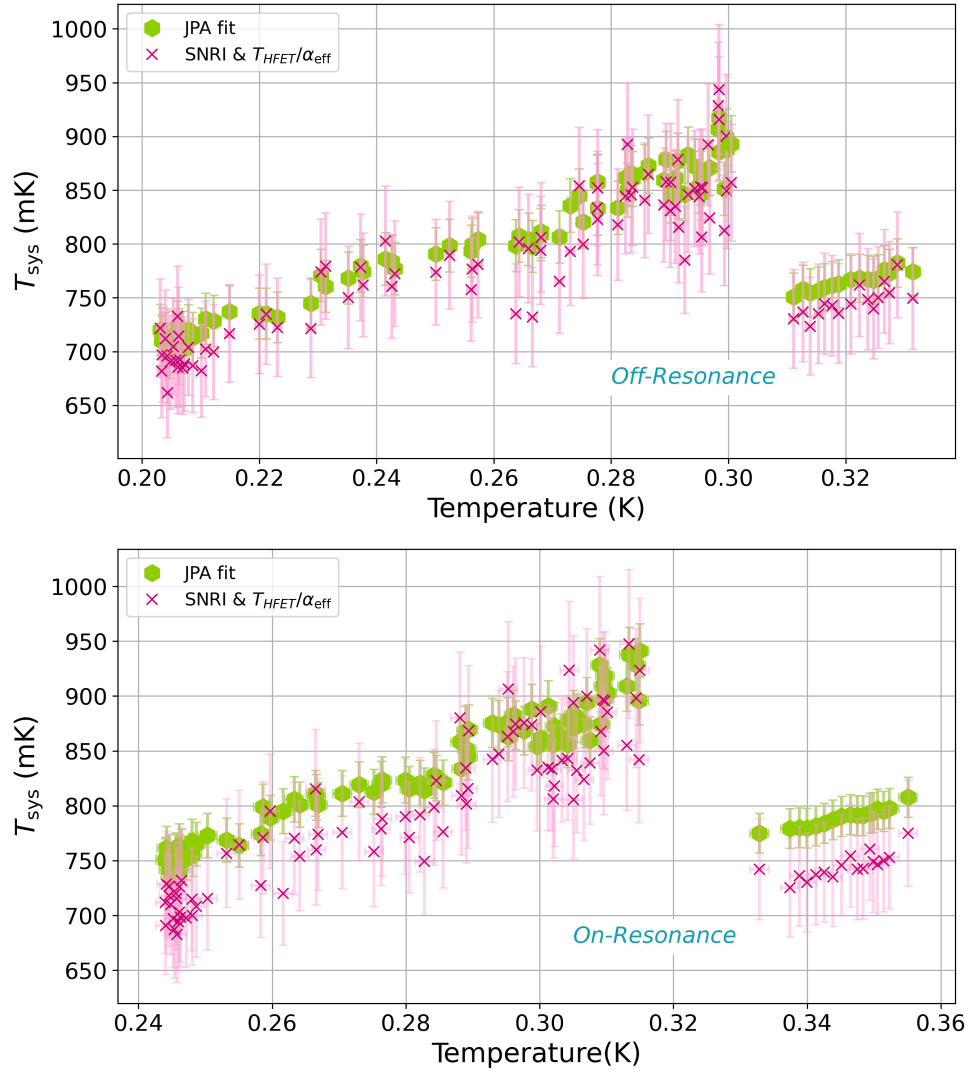


Figure 3.12: Cavity T_{sys} comparison at 1280 MHz. Here I plot the comparison between the two methods, using SNRI and $T_{\text{HFET}}/\alpha_{\text{eff}}$ (pink points) and using JPA fit (green points), over the course of the cavity calibration done at 1280 MHz. Data above 0.3(0.32) K have been randomly downsampled for plotting purposes due to the high density of data in that region. The discontinuity around this temperature was caused by a sharp increase in JPA gain during the course of the measurement, causing a sudden drop in T_{sys} . For both the off-resonance and on-resonance data the calculated system noise temperature is consistent between the two methods within error bars.

3.3.6 Summary and Discussion

ADMX Run 1D noise calibration results under different conditions at 1280 MHz are summarized in Table 3.3.6. By comparing the JPA-off-hot-load and JPA-off-cavity measurements, the transmissivity measured under real experimental conditions ($\alpha = 0.65 \pm 0.09$) between the cavity and the HFET was confirmed to be consistent with the pre-experiment measurement ($\alpha = 0.643 \pm 0.003$). Additionally, by comparing the JPA-off-hot-load $T_{\text{HFET}}/\alpha_{\text{eff}}$ (6.15 ± 0.39 K) and the inferred JPA-on-hot-load $T_{\text{HFET}}/\alpha_{\text{eff}}$ (6.18 ± 0.21 K), it was verified that the JPA-Y-factor and the SNRI & HFET-Y-factor give consistent results. Furthermore, as discussed in the previous section, the ability to perform JPA-on and JPA-off noise measurements in this calibration campaign provided us with an estimate of the systematic uncertainty on T_{sys} in the axion search analysis for this data-taking run. This will be discussed further in Chapter 5.

Quantity	Value (K)	Condition
T_{HFET}	3.99 ± 0.52	JPA-off-cavity
$T_{\text{HFET}}/\alpha_{\text{eff}}$	6.15 ± 0.39	JPA-off-hot-load
	6.18 ± 0.21	JPA-on-hot-load (inferred)
	6.33 ± 0.21	JPA-on-cavity off-resonance (inferred)
	6.72 ± 0.17	JPA-on-cavity on-resonance (inferred)
$T_{\text{JPA,eff}}$	0.141 ± 0.014	JPA-on-hot-load
	0.372 ± 0.022	JPA-on-cavity off-resonance
	0.372 ± 0.018	JPA-on-cavity on-resonance
α	0.65 ± 0.09	In-experiment
	0.643 ± 0.003	Pre-experiment

Table 3.2: Comparison of different noise calibration results at 1280 MHz for ADMX Run 1D.

Additionally, by performing two JPA-on-hot-load measurements at the same frequency, the JPA added noise at a given frequency was shown to be extremely stable over a long time span, and at two different gains. The gain of the JPA varies during regular data-taking, so it is useful to confirm that the amplifier noise performance is not affected by changes in gain on the order of a few dB. Furthermore, this data-taking run lasted for nearly a full year so it is important that the noise performance of the JPA did not degrade over time.

Unexpectedly, the JPA-on-cavity measurements present significantly higher $T_{\text{JPA,eff}}$ when compared to the JPA-on-hot-load $T_{\text{JPA,eff}}$. While one may think this could be due to some sort of mis-measurement, the accuracy of the measured $T_{\text{JPA,eff}}$ for the cavity calibrations is confirmed by the agreement of the T_{sys} calculations using both the direct JPA fit and the SNRI & JPA-off-hot-load fit in Fig. 3.12. That is to say, if the JPA-on-hot-load fit result ($T_{\text{JPA,eff}} = 0.141 \pm 0.014$ K) was used to calculate T_{sys} instead of using the JPA-on-cavity fit results ($T_{\text{JPA,eff}} = 0.372 \pm 0.022$ (0.018) K), the SNRI method and the JPA fit method would not agree in the slightest. One possible reason for this observed discrepancy between the hot load and cavity JPA effective noise is that the insertion loss between the antenna and the switch is larger (α_1 is smaller) than that between the hot load and the switch, which should only have a minor effect on the observed consistency of T_{sys} using either the direct $T_{\text{JPA,eff}}$ fit or SNRI & $T_{\text{HFET}}/\alpha_{\text{eff}}$ in Fig. 3.12. Another reason might be that the hot load cannot accurately represent the cavity when it comes to the interaction with the JPA, which is highly possible due to the impedance difference between the ADMX hot load (50Ω) and the ADMX cavity (highly-reflective at most frequencies). Lastly, some pre-run observations of the JPA used in these calibration measurements hint that $T_{\text{JPA,eff}}$ may have a slight temperature dependence, with higher physical temperatures of the RF components leading to higher effective JPA noise. As noted in Table 3.1, the component temperatures at the high end of the JPA-on-cavity measurement were $\sim 65 - 85$ mK higher than they were during the hot load measurements, so it is possible that this had an effect on the value of $T_{\text{JPA,eff}}$ that was measured. Importantly, without the addition of JPA-on noise calibration measurements, which are new to this data-taking run, this mysterious, non-negligible increase in the added

JPA noise under real experimental circumstances would have remained unidentified. This knowledge will help inform future upgrades for reducing this discrepancy.

That being said, the average system noise temperature for Run 1D was 0.59 ± 0.31 K. To further improve the noise behavior in ADMX, a lower physical temperature of the milliKelvin space is necessary before reaching the SQL. This could be achieved by using a dilution refrigerator with more cooling power or by implementing experimental design refinements that reduce the overall heat load of the system. A better JPA with lower added noise and higher stable gain can also bring down the system noise. Additionally, a set of circulators with lower insertion loss would be helpful because the transmissivity, α , would be larger. Circulators with better isolation would also be helpful for decreasing any standing waves between the cavity and the JPA and, as a result, potentially reduce the difference in $T_{\text{JPA,eff}}$ when switching between the cavity and the hot load. Finally, the consistency in $T_{\text{HFET}}/\alpha_{\text{eff}}$ between different noise measurements indicates that the receiver chain design can be simplified in future versions of ADMX by removing the cryogenic switch and the existing hot load [59]. This simplification could save precious cold and magnetic-field-free space for other electronic devices as well as further increase the transmissivity between the cavity and JPA and, ultimately, reduce the system noise temperature.

Chapter 4

DATA TAKING OPERATIONS

In this chapter I will detail the data taking procedure for ADMX Run 1D. The initial goal of ADMX Run 1D was to search for DFSZ axions between 1-1.4 GHz. However, after data taking began in late 2023, we soon realized there were issues with rod tuning and that the JPA had higher added noise than expected, both of which negatively impacted our scan rate. In the interest of covering this frequency range in a timely manner, the plan was adjusted to instead target sub-KSVZ axions, with the intention of returning for a second, DFSZ-sensitive pass at a later date after experimental upgrades could be made. Additionally, with the release of the Center for Axion and Precision Physics Research (CAPP) experiment’s exclusion limits in early 2024 [17], we adjusted the lower end of our search range to 1.088 GHz to avoid significant overlapping coverage. Furthermore, as can be seen in Fig. 2.5, there are many mode crossings above 1.3 GHz. These mode crossings occur when Transverse Electric (TE) and Transverse Electromagnetic (TEM) modes introduced by the addition of the tuning rod to the cavity mix with the primary TM_{010} cavity mode resulting in a complicated mode structure. When this occurs, the TM_{010} mode no longer aligns properly with the external magnetic field. Since we cannot take science data inside a mode crossing, we decided to adjust the upper end of the search range to 1.315 GHz.

Rather than covering the entire frequency range at once ADMX splits each of our data taking runs into multiple, smaller frequency segments called “nibbles”. These are typically 10-20 MHz wide, and the entire data taking process is performed for each nibble before moving onto the next. Each data taking cycle within a given nibble contains the following steps. First, if necessary, a wide transmission measurement is taken (see Fig. 2.8) to track the cavity resonant mode. This is typically done only every 10-20 data taking cycles since

the frequency is usually not changing quickly enough to warrant taking one every cycle. Next, a standard transmission measurement is taken (see Fig. 2.8) to determine the resonant frequency and loaded quality factor of the cavity. Next the JPA gain and SNRI (Eq. 3.18) are measured, and if necessary, the JPA is biased. While the JPA gain and SNRI are usually measured every cycle, biasing is typically done ever 5-10 cycles, or if the gain or SNRI drop below a specified level. Biasing the JPA involves repeatedly measuring the gain and SNRI while adjusting the pump power and bias current over some small range to determine the optimal combination of parameters. During regular JPA biasing while taking data, the bias current is typically varied over a span of 0.005 mA, while the pump power is typically varied over 0.05 dBm. If the gain or SNRI drop significantly, or if the cavity has moved to a new frequency range for the start of a new nibble, the pump power and bias current spans are typically widened at first to find the right region of parameter space to focus on. The process starts by taking reflection measurements with the JPA powered on and off to determine the gain at a pump power and bias current corresponding to the lower end of their full spans. Then, power measurements are taken with the JPA on and off to determine the SNRI (when combined with G_{JPA}). Next, while holding the pump power constant the bias current is increased incrementally, stopping to measure the gain and SNRI at each step, until the top of the bias current span is reached. Then, the pump power is increased by one step, the bias current is set back to its lowest value and the process of scanning over the bias current is repeated. This continues until every combination of pump power and bias current is tested. An example of the resulting gain measurements from a JPA biasing can be seen in Fig. 4.1.




Figure 4.1: Example of a JPA bias plot. The x-axis is the JPA pump power, the y-axis the JPA bias current and the color map corresponds to the JPA gain in dB. This measurement was centered on a bias current of -2.01 mA with a span of 0.02 mA and a pump power of 11.3 dBm with a span of 1 dBm. There is a darker band of optimal bias parameters that have significantly higher gain than their surroundings ranging from the bottom left to the top right of the plot.

Next, a reflection measurement is taken which again measures the resonant frequency and loaded quality factor as well as the cavity coupling strength. Next, the signal is digitized for 100 s. Finally, the rods are moved in order to change the resonant frequency. For Run 1D, because of the aforementioned rod instability, we typically only tuned the rod every 3 cycles in case the frequency jumped during one of the digitizations (individual spectra).

Sometimes the rods were moved more or less frequently depending on what type of data was being taken. This will be discussed further in the following sections.

4.1 Initial Scan

For a given nibble, the first stage of data taking is to perform an initial scan. This means that the entire frequency range of a given nibble is covered at a fixed scan rate. This initial pass is intended to reach (roughly) some desired SNR across the full nibble.

4.1.1 Calculating Scan Rate

The scan rate refers to how quickly the resonant frequency is changed during data taking. Minimizing the scan rate in an axion haloscope is paramount because the unexplored parameter space spans many orders of magnitude. Thus, scanning quickly increases the chances of discovering the axion. The scan rate can be derived by starting with the equation for the SNR of a single scan, Eq. 3.1, which I will repeat here for convenience:

$$\text{SNR} = \frac{P_{\text{sig}}}{k_B T_{\text{sys}}} \cdot \sqrt{\frac{t}{b}}. \quad (4.1)$$

Here, $P_{\text{sig}} = P_a$ as defined in Eq. 2.1, T_{sys} is determined as discussed in Chapter 3, k_B is the Boltzmann constant, t is the integration time, and b is some bandwidth over which the noise power is measured. P_a can be rewritten to separate out the effects on the signal strength from the cavity and antenna. Namely, let us define:

$$P_a = P_{a,\text{cav}} \left(\frac{\beta}{1 + \beta} \right) \left(\frac{1}{1 + (2\delta f / \Delta f)^2} \right), \quad (4.2)$$

where $P_{a,\text{cav}}$ is the power inside the cavity, δf is the offset between the cavity resonance (f_0) and the signal frequency (f), and $\Delta f = f_0/Q_L$ is the cavity linewidth. The equation for P_a assumes the full width of an axion signal, so I set the bandwidth b equal to the axion linewidth ($b = \Delta f_a = f/Q_a \approx 700$ Hz) so the noise power is calculated over the same bandwidth as P_a .

The total SNR from combining multiple scans can then be defined as:

$$\text{SNR}_{\text{total}}^2 = \sum \left(\frac{P_{\text{sig}}}{P_{\text{noise}}} \right)^2 = \frac{P_{a,\text{cav}}^2 \beta^2 t}{k_B^2 T_{\text{sys}}^2 (1 + \beta)^2 \Delta f_a} \sum \left(\frac{1}{1 + (2\delta f / \Delta f)^2} \right)^2. \quad (4.3)$$

Now, both sides of the equation can be multiplied by $(\frac{\delta f}{t})(\text{SNR}_{\text{total}}^{-2})$ to find an expression for the scan rate:

$$\frac{df}{dt} = \frac{\delta f}{t} = \frac{P_{a,\text{cav}}^2 \beta^2 \delta f}{k_B^2 T_{\text{sys}}^2 (1 + \beta)^2 \Delta f_a} \sum \left(\frac{1}{1 + (2\delta f / \Delta f)^2} \right)^2 \frac{1}{\text{SNR}_{\text{total}}^2}. \quad (4.4)$$

Next, the sum over the cavity lineshape must be calculated. Since the summation should cover the full cavity line width, $\Delta f = f_0 / Q_L$, the number of steps integrated will be $i = -K/2$ to $K/2$ where $K = \Delta f / \delta f$. Calculating the sum requires assuming a value for the step size, δf . Thus, the sum is calculated over a range of values of δf from 300-5000 Hz, resulting in $\delta f \sum \left(\frac{1}{1 + (2\delta f / \Delta f)^2} \right)^2 \approx 0.62 \Delta f$. Next, let us make the substitution: $\frac{\Delta f}{\Delta f_a} = \frac{Q_a}{Q_L}$ where $Q_a = 1.86 \times 10^6$ and $Q_L = Q_0 / (1 + \beta)$, the scan rate can be written as

$$\begin{aligned} \frac{df}{dt} = 36.2 \frac{\text{MHz}}{\text{day}} \cdot 0.62 \cdot \frac{\beta^2}{(1 + \beta)^3} \cdot \left(\frac{g_\gamma}{0.36} \right)^4 \left(\frac{\rho}{0.45 \text{ GeV/cc}} \right)^2 \left(\frac{f}{1.2 \text{ GHz}} \right)^2 \left(\frac{B}{7.6 \text{ T}} \right)^4 \\ \left(\frac{V}{106 \ell} \right)^2 \left(\frac{C_{010}}{0.4} \right)^2 \left(\frac{Q_0}{75000} \right)^2 \left(\frac{3.5}{\text{SNR}} \right)^2 \left(\frac{0.35}{T_{\text{sys}}} \right)^2. \end{aligned} \quad (4.5)$$

4.2 SNR Leveling and Candidate Rescan

Since the initial scan is performed at a fixed rate determined at the start of the nibble, when there are moments during which the JPA gain drops, or the antenna coupling changes, or the mixing chamber temperature increases unexpectedly, and so on, there is a resulting gap in SNR coverage. Therefore, after the initial scan of a given nibble the analysis is run and a grand spectrum is assembled. The details of grand spectrum assembly can be found in Sec. 5.4, but essentially it is a single power spectrum constructed by optimally combining all the individual spectra taken for a given nibble. After one has been assembled, the SNR of the grand spectrum is examined to determine frequency regions that need more data. The

target SNR for Run 1D was set to $\sim 1 - 1.2$ to achieve discovery ability to axions with $P_a \approx \frac{1}{2}P_{\text{KSVZ}}$. An example of the SNR vs frequency just before the initial scan of nibble 5 was complete can be seen in Fig. 4.2. In reality, due to how the SNR leveling script calculates the rod speed, the target SNR was typically set to slightly higher than necessary so that the rods tune more smoothly ensuring even coverage.




Figure 4.2: Top: grand spectrum SNR vs frequency for nibble 5 just before the initial pass was completed (scanning upward in frequency). Bottom: grand spectrum SNR vs frequency for nibble 5 after the SNR leveling pass. As you can see, the SNR was leveled out and nearly the entire spectrum is above our target SNR. Any narrow regions where the SNR hasn't reached the target value (like the small dip at 1238.4 MHz) are covered during the candidate rescan.

Therefore, once the initial scan's SNR is examined for low-SNR regions, an SNR-leveling pass is performed. During this pass, the rods are moved quickly over regions that have sufficient SNR, and no further digitizations are taken. On the other hand, the rods are moved more slowly over regions that require more data to reach the target SNR allowing for more data to be collected. The number of spectra needed for a given low-SNR region is determined using the difference between the SNR target and the SNR reached as well as the expected SNR of a single digitization. This can then be combined with the resonant frequency and loaded quality factor to determine what the step size in Hz should be, which is then translated to rod speed.

After this pass, the grand spectrum is checked for candidates. The process of how candidates are identified is discussed in more detail in Sec. 5.5. From a data-taking perspective, the analysis provides a list of candidate frequencies to be re-scanned. This pass is similar to the SNR-leveling pass in the sense that digitizations are not taken where there are no candidates. However, this pass typically has much narrower frequency regions to cover than the SNR-leveling pass, and the scan speed is even slower when near a candidate. This allows us to efficiently collect sufficient data at the candidate frequencies to determine if they are potential axion signals or not.

Occasionally, the first candidate rescan pass is not enough to eliminate all of the candidates for a given nibble. In this case, secondary candidate rescans may be performed until all candidates are either eliminated or, if we should be so lucky, the axion is discovered. Once all candidates are eliminated, this marks the end of a nibble. At this point, the rods are quickly moved to the next nibble frequency range and the process repeats beginning with the initial scan. The cavity resonant frequency progression/data taking cadence throughout Run 1D can be seen in Fig. 4.3.




Figure 4.3: Here I plot the cavity resonant frequency (f_0) vs time to show the progression of data taking throughout Run 1D. As one can see, the run started up around 1300 MHz and the experiment more or less scanned downward in frequency ending around 1100 MHz. The inset plot shows the data taking cadence within a single nibble, nibble 7, as an example. One can see there is a slower initial scan, followed by a faster SNR leveling scan and two even quicker candidate rescans as described in the text.

Chapter 5

AXION SEARCH ANALYSIS

In this chapter I will describe in detail the Run 1D analysis procedure. At a high level, the analysis involves the coaddition of many individual power spectra to create grand spectra. Each grand spectrum corresponds to a different frequency range which we call a nibble, as defined in Chapter 4. The SNR (Eq. 3.1) for each individual spectrum can be written as

$$\text{SNR} = \frac{P_{\text{sig}}}{\sigma_{\text{sys}}} = \frac{FP_a}{k_B T_{\text{sys}}} \sqrt{\frac{t}{b}}, \quad (5.1)$$

where F is the signal efficiency (ranges from 0 to 1), P_a is the axion signal power as defined in Eq. 2.1, k_B is the Boltzmann constant, T_{sys} is the system noise temperature (as discussed in depth in Chapter 3), t is the integration time and b is the bandwidth of the measured noise power. For Run 1D, $t = 100\text{s}$, $b = 100\text{ Hz}$ and $F = 0.9$. The mean value of T_{sys} was $0.59 \pm 0.31\text{ K}$. Each individual spectrum does not have a high enough SNR to achieve sensitivity to the benchmark QCD axion models. However, by coadding the raw spectra to create a grand spectrum, the overall SNR of the grand spectrum is increased to achieve the necessary level of sensitivity. This is possible because as N individual spectra are added together, the overall power increases to $P_{\text{sig,grand}} = NP_{\text{sig}}$, but the overall noise only increases to $\sigma_{\text{sys,grand}} = \sqrt{N}\sigma_{\text{sys}}$. Therefore, the SNR increases by a factor of \sqrt{N} : $\text{SNR}_{\text{grand}} = \frac{NP_{\text{sig}}}{\sqrt{N}\sigma_{\text{sys}}} = \text{SNR}\sqrt{N}$.

The following sections will detail each of the analysis steps necessary to create the final grand spectra that are ultimately used to search for axions and set limits on $g_{a\gamma\gamma}$.

5.1 *Parameter Extraction*

The first step of the analysis procedure is to extract the relevant experimental parameters from the ADMX database. The start and stop time for each nibble, as defined by experimental operators, are used to pull all data corresponding to the quantities of interest for analysis throughout the course of the nibble. Namely, for each digitization (individual spectrum) in a given nibble, the parameters extracted are: the timestamp, the loaded quality factor (Q_L), the resonant frequency (f_0), the cavity coupling strength (β), the form factor (C_{010}), the magnetic field strength (B), temperature sensor readings for the cavity (T_{cav}), circulators (T_{circ}), and attenuator A on the cavity bypass line (T_A), the JPA gain (G_{JPA}), and the JPA signal-to-noise ratio-improvement (SNRI). For Run 1D, it was determined that the SNRI had frequency dependence, even within a single digitization. That is to say, the system noise temperature was different on resonance vs. off resonance. To account for this, for every digitization the SNRI was calculated using power spectra with the JPA off and on as well as JPA gain spectra to calculate the SNRI for each frequency bin (see Eq. 3.18). The shape of the SNRI vs. frequency is then fit out using an SG filter with polynomial order 4 and window size of 701. Fig. 5.1 shows the structure of the SNRI frequency dependence with the SNRI minimized on-resonance and maximized off-resonance.

Using the SG-filtered binned SNRI values, the system noise temperature (T_{sys}) for each spectrum is calculated accounting for the observed frequency dependence. Because the SNRI is inversely proportional to the system noise temperature (Eq. 3.17), this means that T_{sys} is maximized on-resonance and minimized off-resonance.

If there are any digitizations that were designated to be cut ahead of time, for example digitizations that were taken during noise calibrations, or there was a mode crossing, the cut reasons are noted and saved. Digitizations that end up needing to be cut for unplanned reasons such as bad gain measurements or drastic/uneven rod movements causing relatively large frequency jumps are noted here as well. All of these parameters are then saved into a file to be accessed later on in the analysis.




Figure 5.1: Frequency dependence of the JPA SNRI. A few hours of data (blue) are averaged together for plotting purposes, but for the actual analysis no averaging was done. The data are fit using an SG filter (orange) to capture the general shape of the SNRI frequency dependence. The SG-filtered SNRI is then used to calculate a frequency dependent T_{sys} for each spectrum using Eq. 3.17.

5.2 Warm Receiver Removal

The next step is to remove the warm (room temperature) receiver shape from each of the individual spectra. The warm receiver shape should remain stable across the course of the data taking run, as long as no changes to the receiver hardware are made, so it is typically calibrated only once at the start of the data taking run. To accomplish this, the RF output line that connects the insert to the warm electronics was disconnected from the insert, and was reconnected to a flat noise source (Noisecom model number NC346B) with a frequency

range of 0.01-18 GHz [10]. A diagram detailing the components in the warm receiver chain is shown in Fig. 5.2.




Figure 5.2: Diagram of the ADMX warm receiver chain for Run 1D. From left to right: (A) DC amplifier (Low Noise Factory model LNF-LNR1_15B_SV), (B) DC amplifier (Mini-Circuits model ZX60-33LN-S+), (C) SPDT mechanical switch box (Mini-Circuits model RC-4SPDT-A18), (D) and (E) DC amplifiers (Mini-Circuits model ZX60-3018G-S+), (F) directional coupler (Mini-Circuits model ZX30-17-5-S+), (G) Polyphase Microwave image reject mixer (model IRM0622B), (H) low pass filter (Mini-Circuits model ZX75LP-50-S+), (I) directional coupler (Mini-Circuits model ZX30-17-5-S+), (J) 150 kHz wide custom band-pass filter centered on 10.7 MHz, (K) and (L) DC amplifiers (Mini-Circuits model ZFL-500+), (M) 2 MHz wide band-pass filter centered on 10.7 MHz (Mini-Circuits model SBP-10.7+). The directional couplers enable debugging before and after the mixing stage.

In the warm receiver the signal passes through multiple stages of amplification, is mixed with a local oscillator, passes through a low-pass filter and a narrow band-pass filter, after which it is further amplified before passing through another, wider, band-pass filter before being digitized by a 16-bit Alazar ATS 9462 digitizer [5] with a sampling rate of 50 MegaSamples per second.

Multiple digitizations were taken with the noise source powered on at a handful of fre-

quencies between 1180-1320 MHz. The warm receiver shape is not expected to change significantly with frequency, so averaging measurements at a few select frequencies in the Run 1D frequency range is sufficient to characterize the background for the full data set. After averaging the noise source digitizations together, a Savitsky-Golay (SG) filter is used to fit out the shape. The SG filter is popular for its ability to smooth the data while maintaining the integrity of the original signal. It works by fitting low-degree polynomials to sequential sub-sets of adjacent data points. The width of each sub-set is known as the window size. The window slides across the data, fitting a polynomial (of order of the analyst's choosing) to the points within this window along the way. The value of the fitted polynomial at the center of the window is then taken as the smoothed value. This process is repeated for every data point until the entire spectrum is smoothed. There are considerations to make when choosing the window size and the polynomial order. A smaller window size leads to less effective smoothing, but does capture small scale features. On the other hand, a larger window size does a better job at smoothing the data, but may miss narrow features in the power spectrum. Meanwhile, a lower order polynomial does a better job at capturing the general shape, but does not fit complex trends as well. Contrarily, a higher order polynomial may do a better job at fitting complex patterns but it also introduces the potential to overfit the noise. Considering that we want a fit that captures the general receiver shape, does not overfit the noise, but is still able to capture small scale features, a smaller window size and a lower order polynomial are ideal. In fact, because the warm receiver calibration data are taken using a separate noise source, we can be sure that no axion signal is present. Thus, a quite small window size can be used in order to capture any narrow features in the warm receiver shape. After testing out different configurations, we ultimately used a polynomial of order 3 and window length 17 to fit out the warm receiver shape. The fitted shape is then saved to a file to be divided out of each science data digitization later in the analysis. The warm receiver shape for Run 1D can be seen in Fig. 5.3.

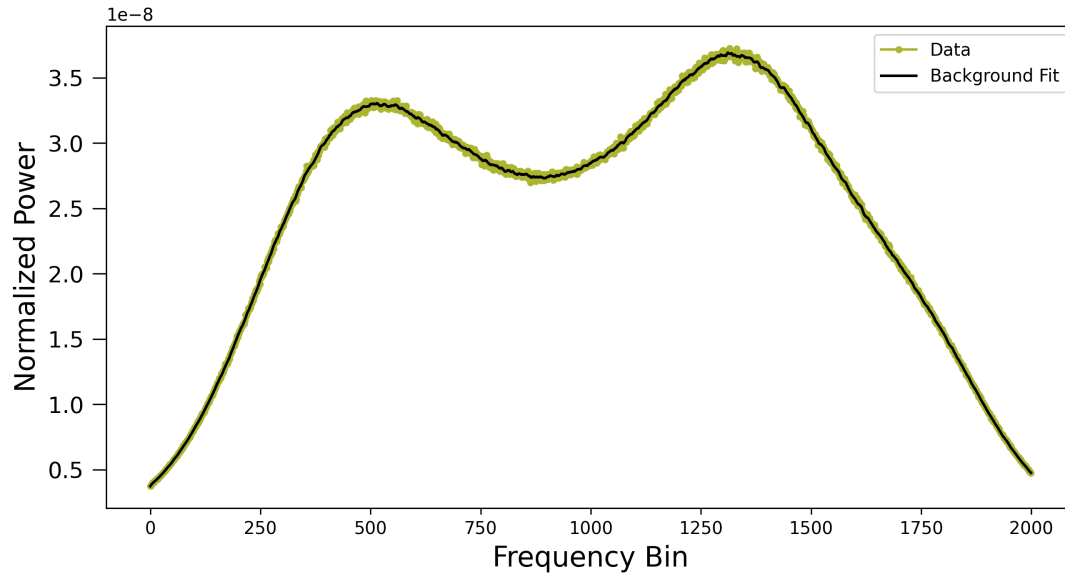


Figure 5.3: Run 1D warm receiver shape. The green points are the averaged spectrum, and the black line is the smoothed receiver shape that was fit using a Savitsky-Golay filter with order 3 and window length 17. Note that the frequency axis in units of bin number, not Hz.

5.3 Raw Spectrum Preparation

Before a grand spectrum can be assembled, we must prepare the individual raw spectra using the results from the previous two analysis steps.

The following procedures are done to each of the raw spectra. First, the warm receiver shape (Fig. 5.3) is removed from each raw spectrum. Examples of spectra after the warm receiver has been removed are shown in Fig. 5.4. Next, the remaining spectral shape caused by frequency-dependent gain of the cold receiver is removed by another pass of an SG filter, this time using a polynomial order of 4 and a window length of 401. Here, we use a larger window size to focus on fitting only large scale shapes to preserve narrow features since the axion signal could be present in this data. Even after two stages of SG filters, the observed

noise distribution of the grand spectrum appeared non-Gaussian. That is, there were still significant non-statistical deviations from the mean as a result of incomplete receiver shape removal. Therefore, the SG filter stages of background removal were followed by a third round of data-driven background removal. More specifically, the twice SG-filtered spectra are divided into groups consisting of at least 50 digitizations each. All of the scans in each segment are then averaged together into a single representative spectrum. By averaging 50+ spectra together, narrow features unique to each spectrum are smoothed over while any shared deviations from the mean caused by the receiver electronics are captured. This way, we do not risk losing any unique narrow features (like a potential axion signal), while still removing spectral features imposed by the receiver chain. The averaged spectrum is divided out of each corresponding science data spectra to remove any residuals remaining after the first two stages of background removal. The spectrum is then normalized by dividing by the mean before subtracting 1 to convert from power to power excess. This results in final processed spectra with powers centered around zero. Examples of filtered spectra can be seen in Fig. 5.5.

Finally, the spectra are divided by $\left(\frac{\beta}{(1+\beta)}\right) \times \left(\frac{1}{1+(2\delta f/\Delta f)^2}\right)$ where β is the cavity coupling strength, Q_L is the loaded quality factor of the cavity, δf is the frequency offset between the cavity resonance (f_0) and the axion signal frequency (f), and $\Delta f = \frac{f_0}{Q_L}$ is the cavity linewidth. These are the two terms in Eq. 2.1 that capture the effects the antenna and Lorentzian profile of the cavity resonance have on the resultant signal. Therefore, dividing by these terms allows us to back out what the signal power inside of the cavity looked like before these effects. In addition to this, the spectra are scaled by a factor of $k_B b$ where k_B is the Boltzmann constant and b is the bandwidth of the measured noise power ($b = 100$ Hz for Run 1D). This converts the spectra into units of Watts/Kelvin, which can later be multiplied by the system noise temperature, T_{sys} , to convert the spectra into units of Watts. Examples of scaled processed spectra are shown in Fig. 5.6.




Figure 5.4: Examples of three different individual power spectra after the warm receiver shape from Fig. 5.3 has been removed. Variations in the shape post-warm-receiver-removal are caused by frequency-dependent gain of the cryogenic receiver. There are further rounds of receiver removal to account for these differences in the cryogenic receiver shape, which is discussed further in this section and the results of which can be seen in Fig. 5.5.

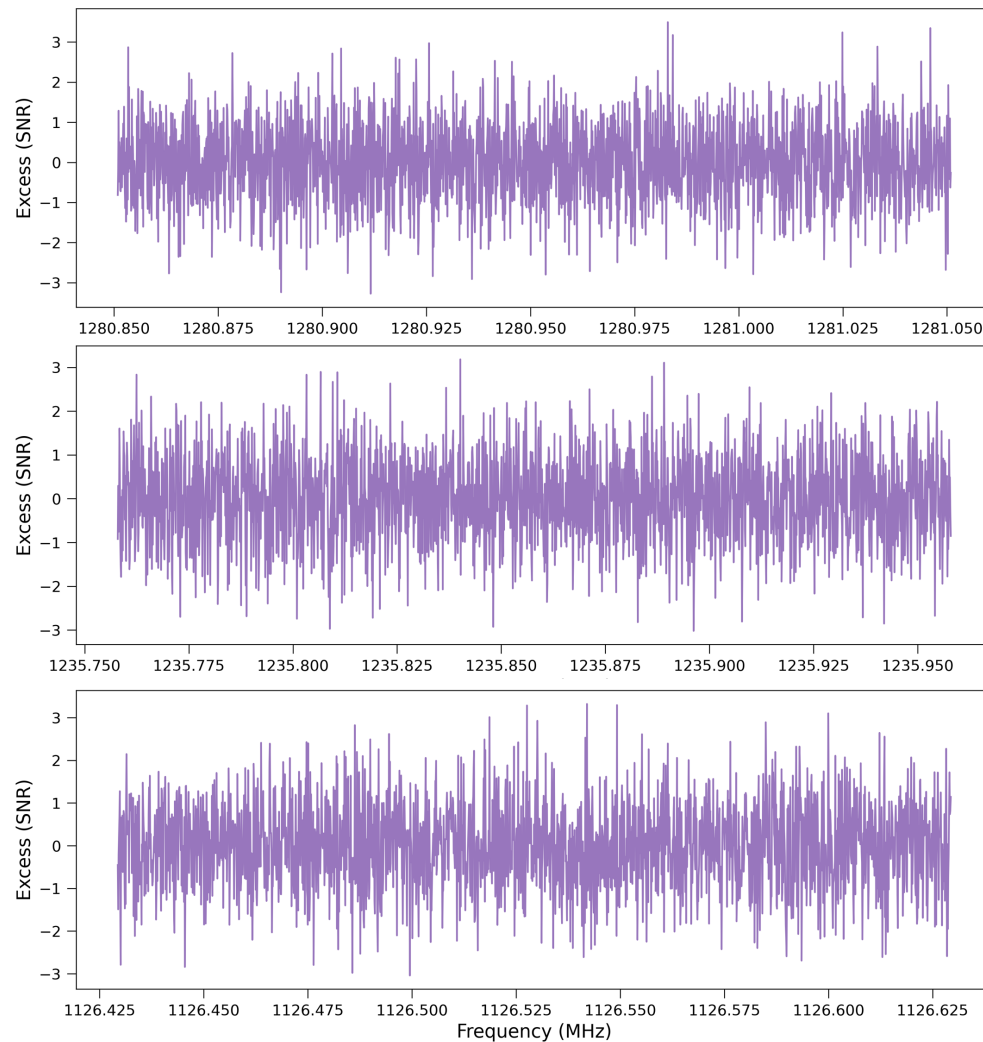


Figure 5.5: The same three spectra shown in Fig. 5.4 after removing any residual background shape imposed by the cryogenic receiver. The process is described further in the text. Note that the y-axis is in units of SNR rather than power for visualization purposes. The spectra used for assembling the grand spectrum have the y-axis in units of power as expected.




Figure 5.6: The same three spectra shown in Fig. 5.4 and Fig. 5.5 after scaling the spectra to represent a signal inside of the cavity and converting the y-axis to units of Watts/Kelvin. This process is discussed further in the text.

5.3.1 Software Synthetics

During the raw spectrum preparation stage of the analysis, there is the option to inject software synthetic axion signals. The software synthetic signals serve two purposes. The first purpose is to determine the signal efficiency, F , (see Eq. 5.1). This is essentially a

measure of how much of the signal power is lost during the process of analyzing the data. By injecting signals of known power into the raw spectra, we are able to determine how much power was lost by comparing the final observed power of the injections at the end of the analysis. For Run 1D, signals with powers of $3.5P_{\text{DFSZ}}$ and $4.4P_{\text{DFSZ}}$ were injected into each nibble with a density of 4 injections per MHz. The average signal efficiency across all nibbles for both the $3.5P_{\text{DFSZ}}$ and $4.4P_{\text{DFSZ}}$ injections was found to be 0.9, meaning 90% of the initial signal power is preserved after processing and analyzing the data. The software synthetic signals can also be used to determine discovery ability, i.e. at what power are we able to discover 90% of axion signals. However, I will not go into detail about this procedure at this point as the procedure will be discussed further in Sec. 5.7.

The frequencies of the software synthetic signals are determined by the desired number of injections, as well as the minimum spacing allowed. For Run 1D, the spectra could accommodate up to 4 software synthetic signals per MHz in a single raw spectrum, corresponding to a separation of 250 kHz. The injected signals are also assigned a power ratio (R) with respect to the expected KSVZ SNR (SNR_{KSVZ}). The signals are then generated by multiplying the Turner axion lineshape [171] by the chosen SNR ($R \times \text{SNR}_{\text{KSVZ}}$), as well as by the warm receiver shape and the Lorentzian profile of the cavity resonance. The software synthetic signals are then injected into the raw spectra after which the spectra undergo the same stages of background removal and scaling discussed earlier in this section.

5.4 Grand Spectrum Assembly

Grand spectrum assembly begins by initiating an empty power spectrum with start and stop frequency of the associated nibble and 100 Hz wide bins. Then, all of the processed spectra that fall within the given frequency range are looped over one by one. If a processed spectrum was designated as needing to be cut during the parameter extraction or raw spectrum preparation stages it is skipped over. If not, it is then checked to see whether it was part of the rescan data, in which case it is not included in the main grand spectrum. It is however, included in the rescan and combined (rescan + main scan) grand spectrum. Next, the

Turner axion lineshape kernel is established. This is a normalized version of the axion lineshape associated with the Turner axion energy distribution (see Equation 5a in Ref. [171]). This kernel is then convolved with each processed spectrum to enhance the SNR of signals whose form matches the expected axion lineshape. These filtered spectra are then saved and added to the grand spectrum bin by bin. Then, once all of the individual spectra have been filtered and added to the grand spectrum, a weighted average is calculated for each bin. For a given frequency bin (j) in the grand spectrum, the weighted average is calculated as

$$P_j = \frac{\sum_{i=1}^n w_i P_i}{\sum_{i=1}^n w_i}, \quad (5.2)$$

where n is the total number of digitizations that cover the frequency bin, $w_i = \frac{1}{\sigma_i^2}$ is the weight associated with each data point, P_i is the excess power of each data point.

The standard deviation of a given bin (j) in the grand spectrum is calculated similarly:

$$\sigma_j = \sum_{i=1}^n \sqrt{\frac{1}{w_i}}. \quad (5.3)$$

In addition to the standard grand spectrum which has been scaled to represent signals as they would appear inside of the cavity, a second “post-cavity” grand spectrum is created. This is achieved by multiplying the axion lineshape filtered spectra by the Lorentzian cavity lineshape: $\frac{1}{1+(2\delta f/\Delta f)^2}$. One may recall that during the raw spectrum preparation stage, the individual spectra were all divided by $\frac{1}{1+(2\delta f/\Delta f)^2}$ to scale the signal back to how it would have looked inside the cavity, so multiplying by the Lorentzian cavity lineshape results in “post-cavity” versions of the spectra. These spectra are then added to their own grand spectrum which is averaged in the same way as the standard (“in-cavity”) grand spectrum. The process of using the “post-cavity” spectrum to determine signal origin will be discussed further in Sec. 5.5.

Finally, the standard grand spectrum is converted into units of expected DFSZ axion power by dividing the entire spectrum by P_{DFSZ} . This is then used to calculate the SNR in DFSZ units ($\text{SNR}_{\text{DFSZ}} = 1/\sigma_{\text{DFSZ}}$) where σ_{DFSZ} are the uncertainties of the DFSZ scaled

grand spectrum.

If a candidate rescan has been performed, the rescan spectra are added to their own versions of the standard grand spectrum and “post-cavity” grand spectrum, in addition to those that contain all of the data. The reasons for this will be discussed further in the following section.

5.5 Candidate Identification and Elimination

Finally, the data have been processed and the search for axions can begin in earnest. There are three criteria that ADMX uses for identifying regions to be rescanned. The first criteria is that the SNR of a given bin (j) in the grand spectrum is below the target SNR. The second criteria is whether the 90% upper limit on the power excess of a given bin (j) in the grand spectrum, $P_j + 1.281\sigma_j$ exceeds $P_{\text{threshold}}$. That is, we would not be able to exclude an axion with power $P_{\text{threshold}}$ at that frequency. The third criteria is whether the power excess of a given bin (j) in the grand spectrum P_j exceeds $N\sigma_j$, where N is some scaling factor. The purpose of the first two criteria are to ensure that our data achieves the desired sensitivity across the full frequency range, and are mostly used when identifying rescan regions during data taking (see Sec. 4.2 for more details). The third criteria is intended to identify statistically significant power excesses in the grand spectrum. For Run 1D, each nibble had a slightly different value of N and $P_{\text{threshold}}$, however $N \approx 5$ on average and $P_{\text{threshold}} \approx 4.1P_{\text{DFSZ}}$ on average. The precise values of N for each nibble were chosen to produce a reasonable amount of candidates. The power thresholds were chosen to target axions with roughly half KSVZ power ($4P_{\text{DFSZ}} \approx 0.55P_{\text{KSVZ}}$).

The candidate identification scripts are run after the SNR leveling pass of a given nibble to identify candidate rescan regions. During this time, regions containing a candidate frequency are scanned over much more slowly than normal to collect sufficient data to eliminate candidates resulting from statistical fluctuations. If there are still candidates remaining after this step, they are checked for persistence. As alluded to in the previous section, during grand spectrum assembly a rescan-only grand spectrum is created in addition to an original

and a combined grand spectrum. This allows us to check if a candidate was present in both the rescan and original data. If it is not, this indicates the signal was transient. The signal from dark matter halo axions is expected to be present always, so this alone is enough to eliminate a candidate. However if it is present in both sets of data, further testing is required to eliminate the candidate. At this point, we can check whether a signal is enhanced on resonance with the cavity is by making use of the “post-cavity” versions of the grand spectra. Fig. 5.7 shows the effect that the Lorentzian cavity line shape scaling described in the previous section has on different types of signals as the experiment scans across them.

Axion-like signals should be maximized when $f_{\text{sig}} = f_0$, as can be seen in the axion-like, post-cavity plot (top left). Therefore, when the Lorentzian cavity lineshape is divided out during raw spectrum preparation, the “in-cavity” signal amplitude appears flat with respect to the the resonant frequency of the cavity, f_0 (bottom left). On the other hand, an RFI-like signal would not experience the same enhancement on-resonance, so its “post-cavity” amplitude should be flat with respect to f_0 (top right) . Therefore, when the Lorentzian cavity lineshape is divided out in raw spectrum preparation, the “in-cavity” signal amplitude is actually minimized on-resonance with the cavity (bottom right). Thus, during grand spectrum preparation when the spectra are multiplied by the Lorentz cavity lineshape this reverts the signals to their “post-cavity” forms. From here, we can calculate the weighted variance of both the “in-cavity” and “post-cavity” versions of the spectra:

$$\sigma^2 = \frac{\sum_{i=1}^n w_i (P_i - P_j)^2}{n} \approx \frac{\sum_{i=1}^n w_i (P_j^2 - 2P_i P_j + P_i^2)}{n}. \quad (5.4)$$

If the “in-cavity” version of a signal has smaller variance than the “post-cavity” version, it is more likely to be an axion signal than RFI. This is because the “in-cavity” version of an axion-like signal should have an amplitude that is consistent between scans, which would result in a smaller variance. In contrast, its “post-cavity” version which would have an amplitude that varied between scans, resulting in a larger variance. On the other hand, if the “post-cavity” version of a signal has smaller variance than the “in-cavity” version, it is more likely to be RFI, for the opposite reason as described above. That is, the “post-cavity”

version would have an amplitude that is consistent between scans, whereas the “in-cavity” version would have an amplitude that varied between scans, resulting in a larger variance for the “in-cavity” version of the signal.

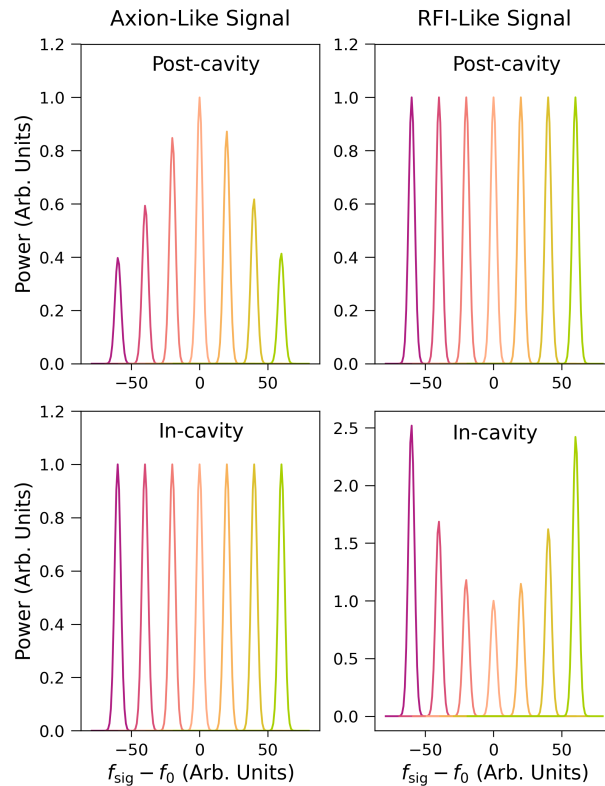


Figure 5.7: Beginning in the top left, as the cavity tunes across the signal the signal experiences resonant enhancement at $f_{\text{sig}} = f_0$. This is how a “post-cavity” version of an axion-like signal would behave. On the contrary, looking at the top right we can see the behavior of a “post-cavity” version of an RFI-like signal. In this case, the signal amplitude is constant regardless of the difference between f_{sig} and f_0 . Moving to the bottom left, an “in-cavity” version of an axion-like signal would look much like the “post-cavity” version of an RFI-like signal, in the sense that it has a constant amplitude with respect to $f_{\text{sig}} - f_0$. In contrast, an “in-cavity” version of an RFI-like signal actually appears to minimize on-resonance with the cavity.

If the signal appears to maximize on-resonance, or if the difference between the “in-cavity” and “post-cavity” variance is not clear enough, the candidate elimination process continues. At this stage in the elimination process, remaining candidates could very well be intentionally injected synthetic axions. As mentioned briefly in Sec. 2.3, throughout data taking hardware synthetic axion signals are injected into the cavity via the weak port using the SAG (synthetic axion generator). These signals are injected to test our detection capability and candidate identification and elimination procedures. After the first couple of checks for persistence and resonant enhancement, operators ask that our collaborators at PNNL who are in control of the SAG turn off the primary synthetic axion injections. After these are turned off, we take another set of candidate rescan data and perform a second persistence check.

If the signal still remains once the SAG has been turned off, the next natural step is to search for ambient RFI on site coming from something like a radio station. This is done by connecting a coaxial antenna to the signal analyzer (SA), setting it to look at the candidate frequency and seeing if a signal appears. Fig. 5.8 shows the process of testing a strong, persistent candidate at 1089.9995 MHz which remained after a candidate rescan. In this case, the candidate was very clearly RFI as we were able to see a strong signal using the coaxial antenna with the SA centered at 1090 MHz. However, if a candidate still remained at this point, as was the case for a candidate in nibble 8 at 1130.6135 MHz, there is still more to be done. The next check is to see whether the signal is present when the antenna is coupled to the TM_{011} mode. One may recall from Sec. 2.2 that the axion only meaningfully couples to TM_{0n0} modes. Therefore, if the candidate is indeed an axion, the signal should vanish while the antenna is coupled to the TM_{011} mode. A scan of the candidate at 1130.6135 MHz was done using the TM_{011} mode, and not only was the signal still present, but it was actually stronger in magnitude than its original observation with the TM_{010} mode (see Fig. 5.9).

This indicated that the signal was not from an axion, but something else. There are sometimes a second set of SAG injections that remain blinded longer than the first set to allow operators on site to exercise (most of) our candidate elimination procedures. Therefore,

at this stage operators ask the SAG control team at PNNL to turn off this second set of SAG injections for one, final pass. After this last pass, the signal disappeared and was confirmed to be a SAG injection. If, however, the candidate had passed the TM_{011} test and was confirmed to *not* be a SAG injection, the final candidate elimination procedure would have been performed. This involves ramping the magnet current down while measuring the signal power in order to test that the $P_{\text{sig}} \propto B^2$, as would be the case for a real axion signal. Ramping the magnet is a non-trivial procedure however, which is why we have many stages of candidate checks before resorting to this test.

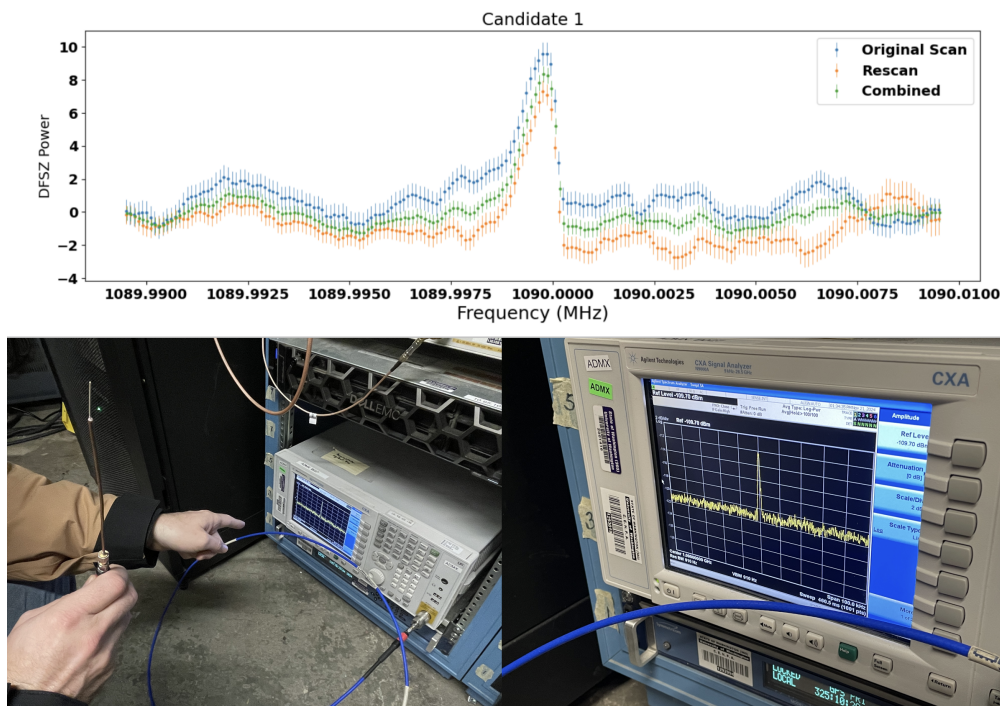


Figure 5.8: The top plot shows the persistent candidate from nibble 12 at 1089.9995 MHz. As one can see, it was undoubtedly a strong power excess that was present in both the original and rescan data. The bottom left photo shows Gray and I using a coaxial antenna connected to the signal analyzer (SA), centered on 1090 MHz, to check for ambient RFI on site. The bottom right photo shows the resulting RFI signal on the SA. This check was used to eliminate this candidate and close out nibble 12.




Figure 5.9: Here I plot signal amplitude vs signal frequency for the persistent candidate at 1130.6135 MHz. The black data show the signal while coupled to the usual TM_{010} mode, while the red data show the signal while coupled to the TM_{011} (which should not couple to axions). The signal was not only able to couple to the TM_{011} mode, but it actually had larger signal strength while coupled to this mode than it did to the TM_{010} mode. This is strong evidence that it is not an axion-like signal. Indeed, after confirming with the SAG control team at PNNL, this candidate was identified as a secondary/blinded SAG.

Table 5.1 summarizes the persistent candidates from Run 1D as well as the method by which they were ruled out. In the end, there were no persistent candidates that passed all of the candidate checks at the conclusion of Run 1D so we were able to set an upper limit on $g_{a\gamma\gamma}$. The process for setting an exclusion limit will be described in the next section.

Frequency (MHz)	Veto Reason
1247.5500	SAG
1247.3200	SAG
1245.6100	SAG
1193.7096	SAG
1192.1600	SAG
1138.8500	SAG
1131.8605	SAG
1128.1200	SAG
1130.6135	TM ₀₁₁
1089.9995	RFI

Table 5.1: Summary of the persistent candidates from Run 1D. Candidates with “SAG” as the veto reason were eliminated by turning off the primary set of hardware synthetic axion injections (a secondary/blinded SAG would not be turned off at this point). The candidate with “RFI” as the veto reason was eliminated by checking for ambient RFI on site. The candidate with “TM₀₁₁” as the veto reason was eliminated by taking data with the TM₀₁₁ mode, which revealed it to be a blinded SAG.

5.6 Exclusion Limit

The process for setting an exclusion limit is as follows. First, the different sources of systematic uncertainty must be determined. For Run 1D these were: the system noise temperature calibration, the cavity form factor simulation, the signal efficiency calculation, resonant frequency drift caused by rod instability, as well as variations in the loaded quality factor and cavity coupling strength due to antenna instability. The uncertainties related to each of these quantities were propagated to reflect their effects on the measured SNR for each frequency bin. Table 5.2 shows the percent uncertainties associated with each quantity.

Source	Percent Uncertainty on SNR
Q_L	0.6%
$\beta/(1 + \beta)$	0.3%
T_{sys}	9.9%
B^2VC	5%
f_0 Drift	1.5%
F	3%

Table 5.2: Summary of the average percent systematic uncertainties associated with the SNR from Run 1D.

After the systematic uncertainties are understood, the limit setting can begin. First, the individual grand spectra from each nibble are combined into a single Run 1D grand spectrum covering the full frequency range. The upper limit on power in units of DFSZ power for each bin (j) in the Run 1D grand spectrum is then calculated via the following equation:

$$\frac{P_{\text{exc}}}{P_{\text{DFSZ}}} = \frac{\sigma_{\text{totj}}}{F} \left[\frac{P_j}{P_{\text{DFSZ}}} - \left(\sqrt{2} \operatorname{erf}^{-1} \left[1 - \left(2C \left(1 - \frac{1}{2} \left[1 - \operatorname{erf} \left(\frac{P_j}{P_{\text{DFSZ}} \sigma_{\text{totj}} \sqrt{2}} \right) \right] \right) \right] \right) \right] \right], \quad (5.5)$$

where P_{DFSZ} is the expected power of a DFSZ axion in frequency bin j , P_j is the weighted average power excess in frequency bin j , $\sigma_{\text{tot}j} = \sqrt{\sigma_j^2 + \sigma_{\text{syst}j}^2}$ (where σ_j is the standard deviation and $\sigma_{\text{syst}j}$ is the systematic uncertainty of frequency bin j), F is the signal efficiency, $C = 0.9$ is the desired confidence level, and $\text{erf}(x)$ is the error function. This equation assumes a truncated normal distribution following the methodology outlined by Feldman and Cousins in Ref. [86] to exclude unphysical negative values of $g_{a\gamma\gamma}$. From here, the upper limit on power can be converted into an upper limit on axion-photon coupling using the following relationship:

$$\left(\frac{g_{a\gamma\gamma,\text{exc}}}{g_{a\gamma\gamma,\text{DFSZ}}} \right)^2 = \left(\frac{P_{\text{exc}}}{P_{\text{DFSZ}}} \right). \quad (5.6)$$

5.7 Results

The 90% confidence level upper limits on $g_{a\gamma\gamma}$ for ADMX Run 1D can be seen in context with other axion experiments in Fig. 5.10.

5.7.1 Discovery Ability

As alluded to earlier in the text, software synthetic injections can be used to determine our discovery ability. The exclusion limits answer the question: if we ran the experiment many times what power P_{exc} (and corresponding axion-photon coupling, $g_{a\gamma\gamma,\text{exc}}$) would 90% of my measurements fall below? This is essentially a way of quantifying how sensitive our data is in comparison with other experiments. However, a perhaps more scientifically interesting question to ask is: at what value of axion power (or, alternatively, axion-photon coupling) would my experiment be able to detect 90% of axions? In order to answer this question, we must consider what is necessary for a signal to be “detected” in ADMX. Most importantly, it must trigger one of the candidate definitions. As a reminder, the two primary candidate definitions are:

1. $P_j + 1.281\sigma_j > P_{\text{threshold}}$

$$2. P_j > N\sigma_j,$$

where P_j is the measured power excess of a given frequency bin (j) in the grand spectrum and σ_j is the associated standard deviation of bin (j). $P_{\text{threshold}}$ is some threshold power in units of P_{DFSZ} and N is some scaling factor that dictates the σ excess threshold, both of which are determined for each nibble. Therefore, the power and σ excess thresholds can provide some insight into what our discovery ability is.

The top plot in Fig. 5.11 shows a Gaussian probability distribution function (PDF) with a mean of $\mu = P_j + 1.28\sigma_j$ where $P_j = P_{\text{threshold}}$ and standard deviation $\sigma_j = 1/\text{SNR}_{\text{avg}}$ where SNR_{avg} is the average signal to noise ratio for the whole nibble. In reality, the SNR and therefore σ_j varies from bin to bin, but for estimation purposes using the average is sufficient. Using nibble 4 as an example, $\text{SNR}_{\text{avg}} \approx 1.7$ and $P_{\text{threshold}} = 3.9P_{\text{DFSZ}}$. The bottom plot is a Gaussian complementary cumulative distribution function (1 - CDF) with the same mean and standard deviation as the PDF. The complementary CDF shows what fraction of the area under the PDF lies to the right of a given value of $P_{\text{threshold}}$. In this case, since we have set $P_j = P_{\text{threshold}}$ 90% of the area under the PDF is greater than $P_{\text{threshold}}$. Given this information and the first candidate definition, $P_j + 1.281\sigma_j > P_{\text{threshold}}$, 90% of signals with power P_j will trigger the power threshold when $P_j = P_{\text{threshold}}$. Given that P_j is the measured signal power, which is smaller than the true signal power, $P_{j,\text{true}}$, by a factor $F = 0.9$, I expect nibble 4 to have 90% discovery ability for $4.3P_{\text{DFSZ}} = 0.59P_{\text{KSVZ}}$ strength axions, which should have measured signal strength $P_j = 3.9P_{\text{DFSZ}} = 0.54P_{\text{KSVZ}}$.

For nibble 4, the σ excess threshold was set to $4.8\sigma_j$. Using this threshold, a $3.9P_{\text{DFSZ}}$ strength axion would be discovered 97% of the time. Since this chosen threshold value for σ_j excess is less strict, it triggers more easily than the power threshold. Therefore, if we wanted to achieve 90% discovery ability to a $4.8P_{\text{DFSZ}}$ strength axion, we could set this threshold even higher to $5.3\sigma_j$. However, given that $4.8\sigma_j$ is ultimately what was used, what axion power should trigger this candidate definition 90% of the time? The top plot in Fig. 5.12 shows a Gaussian probability distribution function with a mean of $\mu = P_j/\sigma_j$ and standard

deviation of 1 where $P_j = 3.6P_{\text{DFSZ}}$ and σ_j is defined as it was for the power threshold. The bottom plot is a Gaussian complementary cumulative distribution function (1 - CDF) with the same mean and standard deviation as the PDF. The complementary CDF shows what fraction of the area under the PDF lies to the right of a given value of $\sigma_T = N\sigma_j$. In this case, given that $P_j = 3.6P_{\text{DFSZ}}$, 90% of the area under the PDF is greater than $4.8\sigma_j$. Given that P_j is the measured signal power, which is smaller than the true signal power, $P_{j,\text{true}}$, by a factor $F = 0.9$, I expect nibble 4 to have 90% discovery ability for $4P_{\text{DFSZ}} = 0.55P_{\text{KSVZ}}$ strength axions, which should have measured signal strength $P_j = 3.6P_{\text{DFSZ}} = 0.50P_{\text{KSVZ}}$.

Of course, there is cross over between these two candidate definitions, with many signals triggering both. Therefore, analytically calculating the combined discovery ability is difficult. Thus, the software synthetic axion injections (Sec. 5.3.1) can be used to calculate our discovery ability numerically. To do this, I split nibble 4 into 3 sub-nibbles, each with a width of 5 MHz across which the discovery ability could be calculated. That is, the first sub-nibble spans 1300-1305 MHz, the second from 1305-1310 MHz and the third from 1310-1315 MHz. Given the range of true signal powers $P_{j,\text{true}} = (4 - 4.3)P_{\text{DFSZ}}$ I found from the candidate definitions for nibble 4, I injected 1000 synthetic axion signals into each sub-nibble starting with the largest expected power, $P_{j,\text{true}} = 4.3P_{\text{DFSZ}}$, and working down in steps of $0.1P_{\text{DFSZ}}$ until 90% of injections were discovered for each sub-nibble. The percentage of the 1000 injected signals that were detected at every power level in each sub-nibble was recorded in Tab. 5.3. For the first sub-nibble spanning 1300-1305 MHz, 90% of the injected signals were discovered at a power between $(4.1 - 4.2)P_{\text{DFSZ}} = (0.56 - 0.58)P_{\text{KSVZ}}$. For the second sub-nibble spanning 1305-1310 MHz, 90% of the injected signals were discovered at a power of $3.5P_{\text{DFSZ}} = 0.48P_{\text{KSVZ}}$. For the third sub-nibble spanning 1310-1315 MHz, 90% of the injected signals were discovered at a power of $3.9P_{\text{DFSZ}} = 0.54P_{\text{KSVZ}}$. This simulation confirmed that the discovery ability predicted from the candidate definitions is in decent alignment with the true discovery ability. One caveat is that I used the average SNR across the full nibble to calculate the predicted discovery ability, but at the very edges of nibbles there tend to be drops in SNR caused by turning the rod around. I believe this could be why

the middle sub-nibble (1305-1310 MHz) reached 90% discovery ability at a lower power than either of the edge sub-nibbles. Something interesting to note is how this discovery ability compares to our upper limits in nibble 4. Between 1300-1315 MHz, the average discovery ability I found was $3.9P_{\text{DFSZ}}$, which is a little over 3 times worse than the average upper limit in the same frequency range of $1.2P_{\text{DFSZ}}$. However, the average discovery ability is marginally better than the worst upper limit in this nibble of $4.3P_{\text{DFSZ}}$. Therefore, it seems that the worst upper limit in a given nibble generally provides a better estimate of the discovery ability than the average upper limit.

P_{inj} (DFSZ)	P_{inj} (KSVZ)	% Discovered 1300-1305 MHz	% Discovered 1305-1310 MHz	% Discovered 1310-1315 MHz
3.4	0.47	65	88	73
3.5	0.48	70	90	77
3.6	0.5	75	93	80
3.7	0.51	78	94	84
3.8	0.52	82	95	87
3.9	0.54	84	96	90
4.0	0.55	87	97	92
4.1	0.56	89	97	94
4.2	0.58	92	98	96
4.3	0.59	94	98	96

Table 5.3: Summary of the discovery ability Monte-Carlo simulations. The first two columns are the injection power in units of DFSZ and KSVZ power respectively. The injections occur prior to spectrum processing, so the actual measured power is expected to be smaller than P_{inj} by a factor of $F = 0.9$. The following three columns are the percentage of the 1000 injected signals discovered for each of the three sub-nibbles: 1300-1305 MHz, 1305-1310 MHz, and 1310-1315 MHz.

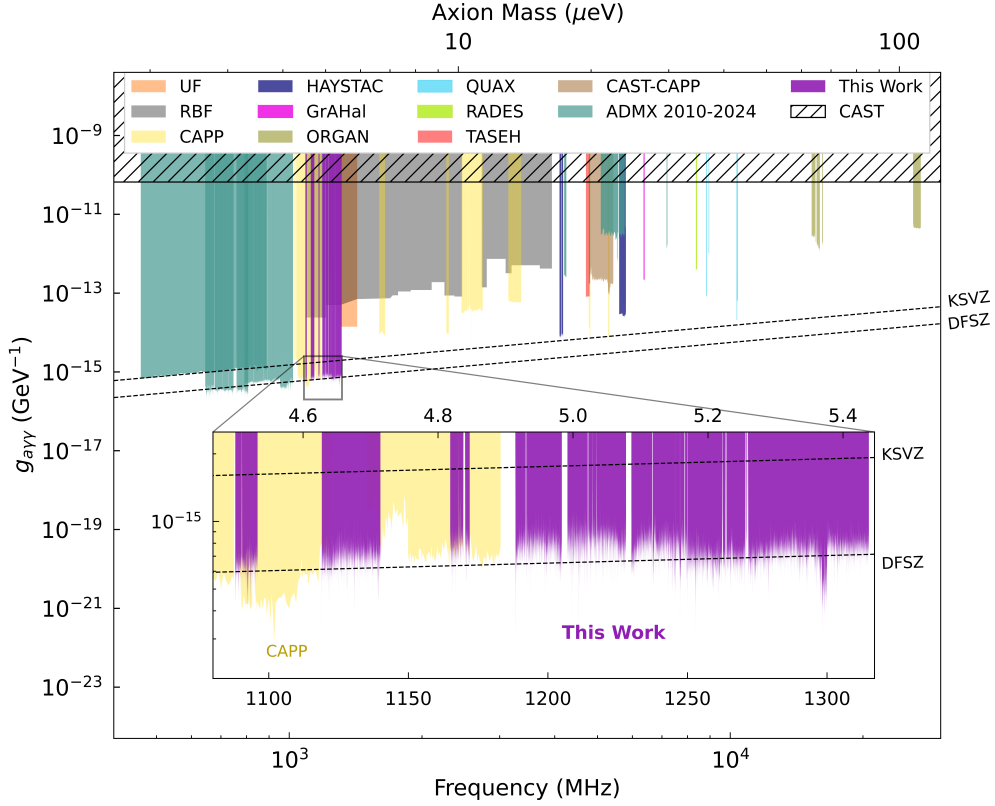


Figure 5.10: A global limit plot putting this work (purple) in context with other experiments, with an inset zooming in on this work’s 90% C.L. upper limits on $g_{a\gamma\gamma}$ (as well as limits from Ref. [17] due to overlapping coverage). The dark matter density is assumed to be $0.45 \text{ GeV}/\text{cm}^3$. Gaps in the limits are due to mode crossings, regions where axion search mode of the cavity intersected other static weakly tuning modes. KSVZ and DFSZ sensitivities are shown as dashed lines. Previous limits set by ADMX are shown in teal [26, 79, 46, 33, 91, 35, 44]. Limits from other experiments depicted include those set by University of Florida (UF) [98], Rochester-Brookhaven-Florida (RBF) [179], Center for Axion and Precision Physics (CAPP) [121, 105, 117, 122, 114, 182, 181, 115, 17], Haloscope At Yale Sensitive To Axion Cold dark matter (HAYSTAC) [186, 30, 106], Grenoble Axion Haloscope project (GrAHal) [94], Oscillating Resonant Group AxioN experiment (ORGAN) [126, 150, 149], QUaerere AXions experiment (QUAX) [20, 21, 19, 72, 152], Relic Axion Dark Matter Exploratory Setup (RADES) [189], Taiwan Axion Search Experiment with Haloscope (TASEH) [58], CAST-CAPP [13], and CERN Axion Solar Telescope (CAST) [22]. Caption taken from Ref. [63].




Figure 5.11: Top plot: Gaussian probability distribution function with a mean of $\mu = P_j + 1.28\sigma_j$ where $P_j = P_{\text{threshold}}$ and standard deviation $\sigma_j = 1/\text{SNR}_{\text{avg}}$ where SNR_{avg} is the average SNR for the whole nibble. Bottom plot: Gaussian complementary cumulative distribution function with the same mean and standard deviation as the PDF.




Figure 5.12: Top plot: Gaussian probability distribution function with a mean of $\mu = P_j/\sigma_j$ and standard deviation of 1 where $P_j = 3.6P_{\text{DFSZ}}$ and $\sigma_j = 1/\text{SNR}_{\text{avg}}$ where SNR_{avg} is the average SNR for the whole nibble. Bottom plot: Gaussian complementary cumulative distribution function with the same mean and standard deviation as the PDF.

Chapter 6

CONCLUSIONS

The success of ADMX Run 1D was enabled by a number of upgrades, both experimental and analytical. To be more specific, the cryogenic receiver was upgraded to reduce the heat load on the dilution refrigerator. This was done with the goal of lowering the mixing chamber base temperature to lessen the Johnson noise contribution to the overall system noise temperature. The receiver upgrades included changing out cables for those with better materials and/or different lengths, as well as by adding more heat sink locations when necessary. In addition, an entire RF through line was removed from the insert, which provided a significant amount of relief for the dilution refrigerator. These efforts allowed ADMX to achieve its lowest mixing chamber temperature during stable operations to date of ~ 80 mK.

Upgrades to the cryogenic DC cabling were also made alongside those to the receiver chain. In this case, the goal was more to improve the stability and continuous operation of critical DC connections in the insert. During the previous data taking run we had experienced critical dilution refrigerator sensor failures fairly often. To ameliorate this situation I designed a custom cryogenic DC cable to replace the existing DF cable (Cable F) that was manufactured by collaborators at WUSTL. The upgraded cable performed extremely well, with zero sensor disruptions during the entirety of Run 1D — which lasted for roughly one year. Another part of the effort to improve the reliability of DC connections was to improve internal documentation. Clear, extensive documentation makes it easier to fix issues when they arise and to do more thorough testing prior to the start of a data taking run. I created a significant body of documentation for the cryogenic DC cabling including a sensor map, gallery of sensor photos, as well as spreadsheets detailing the connections on each cable, as well as any broken pins. Lastly, a combination RC/LC filter was installed on Cable D, which

supplies power to the quantum amplifier package to reduce electronic noise on the JPA in order to improve its noise performance and dynamic range.

In addition to hardware upgrades, there were a few changes made to the analysis. Namely, an additional stage of data-driven cryogenic receiver removal was incorporated into the analysis to handle spectral fluctuations from the frequency dependent gain of the cold receiver. This ensured that the noise distribution was Gaussian as expected and properly prepared the spectra to be used to search for axions. In addition to this, the software synthetic injection capabilities were improved to enable us to numerically calculate the discovery ability.

I also calculated the discovery ability for one nibble using Monte Carlo methods, establishing a code base that can be built upon by future ADMX researchers in the remaining Run1D and future frequency ranges. Going forward, if the scripts are optimized to run more efficiently, the numerically determined discovery ability will be an interesting and scientifically informative addition to the ADMX science results, along with the standard 90% exclusion limits.

Finally, perhaps the most significant improvement made for Run 1D was to our noise calibration procedure. Firstly, the thermal sinking of the ADMX noise source, the hot load, was upgraded. This enabled us to reach a hot load base temperature of ~ 150 mK compared with the previous base temperature of ~ 475 mK. Since the base temperature of the noise source was so high during previous runs, we were not able to perform a Y-factor measurement with the JPA on. Reducing the base temperature of the hot load enabled us to perform this type of calibration without saturating the JPA (which would make our noise calibration results useless). Therefore, I was able to lead the most in-depth noise calibration campaign in ADMX to date, culminating in a journal publication [96]. This included four different types of y-factor noise calibration measurements: JPA-off-hot-load, JPA-on-hot-load, JPA-off-cavity, and JPA-on-cavity. All of these measurements were essential to fully characterize the noise performance of the cryogenic receiver chain. In addition, Dan Zhang and I thoroughly detailed the theoretical noise model based on the components in the receiver chain which we were then able to compare with the experimental results. In doing so we found that the

transmissivity between the cavity and the HFET (α) measured prior to the start of Run 1D was consistent with the transmissivity found by comparing the JPA-off-hot-load and JPA-off-cavity measurements. In addition, we verified that the JPA-on Y-factor and the SNRI & JPA-off Y-factor give consistent results for the effective HFET noise, $T_{\text{HFET}}/\alpha_{\text{eff}}$. This provided confidence in our established noise model and the effectiveness of using the indirect (SNRI) method for determining the JPA-on system noise temperature. Furthermore, the JPA added noise at a given frequency was shown to be exceptionally stable over the course of a few months and a two different gains. It was discovered that the JPA added noise while the receiver is connected to the cavity is appreciably higher than when the receiver is connected to the hot load. The precise reason for this remains a question, but we suspect it could be due to a difference in insertion loss between the hot load and switch and the antenna and the switch. The discrepancy could also be due to the impedance difference between the hot load and the cavity. Lastly, there was some evidence in pre-run tests of the JPA that its added noise exhibited a temperature dependence. Therefore, because the receiver chain is slightly warmer overall while connected to the cavity, this could have increase the measured JPA added noise in this case. Nonetheless, without the addition of JPA-on noise calibration measurements, this discrepancy would have remained undetected. Now, this knowledge can be used to inform future upgrades of the receiver and noise calibration procedure.

As ADMX moves forward further changes can be made to optimize the efficiency of the search for axions. From a mechanical perspective, improvements to the gearbox mechanisms would be useful to correct issues such as the resonant frequency drifts and sudden losses in coupling. In addition, given the largest systematic uncertainty for Run 1D was related to the system noise temperature calibration, changes can and should be made to help reduce this uncertainty. Given the consistency of the inferred value for $T_{\text{HFET}}/\alpha_{\text{eff}}$ from the JPA-on-cavity off-resonance measurement and the direct JPA-off-hot-load measurement, one idea worth pursuing is the implementation of a switchless VTS design [59]. That is, mounting the existing attenuator on the bypass line to a VTS that can achieve a base temperature within ~ 5 mK of the mixing chamber, while also being able to be heated by at least

~ 50 mK without disturbing the mixing chamber too much. This would enable the removal of the cryogenic RF switch between the cavity and the JPA, which would help increase the transmissivity on the output line. Lastly, some changes could be made to improve the experimental down time. For instance, updating the method for monitoring the JPA SNRI. Currently, this requires two types of measurements, one to measure power with the digitizer and one to measure gain with the VNA. However, this can be reduced to one type of measurement by injecting tones of known power into digitizations. This would not only cut down on the time required to monitor the JPA during data taking, but would also reduce error in the SNRI since the power and gain would be measured simultaneously. Lastly, switching from a current-pumped JPA to a flux-pumped JPA will reduce the time needed for biasing since the resonant frequency only depends on one parameter instead of two. Not to mention, utilizing a flux-pumped JPA also allows for the idler band to be included in the digitization bandwidth without interference from the pump tone. This could then be used to increase the SNR by optimally combining the uncorrelated portions of the idler and signal bands, as was done in Ref. [17].

To conclude, from December 2023 to November 2024 ADMX searched for QCD axion dark matter between 1.088-1.315 GHz, marking the longest continuous data taking run with ADMX to date. Over that frequency range ADMX excluded axions with an average power of $P_{\text{exc}} \approx 1.3P_{\text{DFSZ}} = 0.18P_{\text{KSVZ}}$ ($g_{a\gamma\gamma,\text{exc}} \approx 1.1g_{a\gamma\gamma,\text{DFSZ}}$) at 90% confidence. In addition, ADMX achieved 90% discovery ability to $(4.1 - 4.2)P_{\text{DFSZ}}$, $3.5P_{\text{DFSZ}}$, $3.9P_{\text{DFSZ}}$ strength axions (corresponding to coupling strengths of $(2.02 - 2.05)g_{a\gamma\gamma,\text{DFSZ}}$, $1.87g_{a\gamma\gamma,\text{DFSZ}}$, $1.97g_{a\gamma\gamma,\text{DFSZ}}$) between 1.3-1.305, 1.305-1.31, 1.31-1.315 GHz. Over the coming year, ADMX will take more data in the same frequency range after implementing hardware and analysis upgrades in order to achieve sensitivity to DFSZ axions in a timely manner.

BIBLIOGRAPHY

- [1] <https://www.lakeshore.com/products/categories/overview/temperature-products/cryogenic-temperature-sensors/cernox>.
- [2] <https://www.scientificinstruments.com/product/model-ro-600-1k/>.
- [3] <https://www.mdcprecision.com/9921034-vacuumsidecablekit>.
- [4] <https://lownoisefactory.com>.
- [5] <https://www.alazartech.com/en/product/ats9462/13/>.
- [6] How does a dilution refrigerator work? <https://bluefors.com/stories/how-does-a-dilution-refrigerator-work/>, 2023.
- [7] <https://www.radiall.com/>, Nov 2024.
- [8] <https://noisecom.com/>, Oct 2024.
- [9] <https://quinstar.com/>, May 2025.
- [10] https://noisecom.com/Portals/0/Datasheets/NC346_Datasheet.pdf, May 2025.
- [11] L. F. Abbott and P. Sikivie. A cosmological bound on the invisible axion. *Physics Letters B*, 120(1-3):133–136, January 1983.
- [12] C. Abel, S. Afach, N. J. Ayres, C. A. Baker, G. Ban, G. Bison, K. Bodek, V. Bondar, M. Burghoff, E. Chanel, Z. Chowdhuri, P.-J. Chiu, B. Clement, C. B. Crawford, M. Daum, S. Emmenegger, L. Ferraris-Bouchez, M. Fertl, P. Flaux, B. Franke, A. Fratangelo, P. Geltenbort, K. Green, W. C. Griffith, M. van der Grinten, Z. D. Grujić, P. G. Harris, L. Hayen, W. Heil, R. Henneck, V. H elaine, N. Hild, Z. Hodge, M. Horras, P. Iaydjiev, S. N. Ivanov, M. Kasprzak, Y. Kermaidic, K. Kirch, A. Knecht, P. Knowles, H.-C. Koch, P. A. Koss, S. Komposch, A. Kozela, A. Kraft, J. Kempel, M. Kuźniak, B. Lauss, T. Lefort, Y. Lemi ere, A. Leredde, P. Mohanmurthy, A. Mtchedlishvili, M. Musgrave, O. Naviliat-Cuncic, D. Pais, F. M. Piegsa, E. Pierre, G. Pignol, C. Plonka-Spehr, P. N. Prashanth, G. Qu em ener, M. Rawlik, D. Rebreyend,

- I. Rienäcker, D. Ries, S. Rocchia, G. Rogel, D. Rozpedzik, A. Schnabel, P. Schmidt-Wellenburg, N. Severijns, D. Shiers, R. Tavakoli Dinani, J. A. Thorne, R. Viot, J. Voigt, A. Weis, E. Wursten, G. Wyzynski, J. Zejma, J. Zenner, and G. Zsigmond. Measurement of the permanent electric dipole moment of the neutron. *Phys. Rev. Lett.*, 124:081803, Feb 2020.
- [13] C. M. Adair, K. Altenmüller, V. Anastassopoulos, S. Arguedas Cuendis, J. Baier, K. Barth, A. Belov, D. Bozicevic, H. Bräuninger, G. Cantatore, et al. Search for dark matter axions with cast-capp. *Nature Commun.*, 13(1):6180, 2022.
- [14] C. B. Adams, N. Aggarwal, A. Agrawal, R. Balafendiev, C. Bartram, M. Baryakhtar, H. Bekker, P. Belov, K. K. Berggren, A. Berlin, C. Boutan, D. Bowering, D. Budker, A. Caldwell, P. Carenza, G. Carosi, R. Cervantes, S. S. Chakrabarty, S. Chaudhuri, T. Y. Chen, S. Cheong, A. Chou, R. T. Co, J. Conrad, D. Croon, R. T. D’Agnolo, M. Demarteau, N. DePorzio, M. Descalle, K. Desch, L. Di Luzio, A. Diaz-Morcillo, K. Dona, I. S. Drachnev, A. Droster, N. Du, K. Dunne, B. Döbrich, S. A. R. Ellis, R. Essig, J. Fan, J. W. Foster, J. T. Fry, A. Gallo Rosso, J. M. García Barceló, I. G. Irastorza, S. Gardner, A. A. Geraci, S. Ghosh, B. Giaccone, M. Giannotti, B. Gimeno, D. Grin, H. Grote, M. Guzzetti, M. H. Awida, R. Henning, S. Hoof, G. Hoshino, V. Irsic, K. D. Irwin, H. Jackson, D. F. Jackson Kimball, J. Jaeckel, K. Jakovcic, M. J. Jewell, M. Kagan, Y. Kahn, R. Khatiwada, S. Knirck, T. Kovachy, P. Krueger, S. E. Kuenstner, N. A. Kurinsky, R. K. Leane, A. F. Leder, C. Lee, K. W. Lehnert, E. W. Lentz, S. M. Lewis, J. Liu, M. Lynn, B. Majorovits, D. J. E. Marsh, R. H. Maruyama, B. T. McAllister, A. J. Millar, D. W. Miller, J. Mitchell, S. Morampudi, G. Mueller, S. Nagaitsev, E. Nardi, O. Noroozian, C. A. J. O’Hare, N. S. Oblath, J. L. Ouellet, K. M. W. Pappas, H. V. Peiris, K. Perez, A. Phipps, M. J. Pivovarov, P. Quílez, N. M. Rapidis, V. H. Robles, K. K. Rogers, J. Rudolph, J. Ruz, G. Rybka, M. Safdari, B. R. Safdi, M. S. Safronova, C. P. Salemi, P. Schuster, A. Schwartzman, J. Shu, M. Simanovskaia, J. Singh, S. Singh, K. Sinha, J. T. Sinnis, M. Siodlaczek, M. S. Smith, W. M. Snow, A. V. Sokolov, A. Sonnenschein, D. H. Speller, Y. V. Stadnik, C. Sun, A. O. Sushkov, T. M. P. Tait, V. Takhistov, D. B. Tanner, F. Tavecchio, D. J. Temples, J. H. Thomas, M. E. Tobar, N. Toro, Y. D. Tsai, E. C. van Assendelft, K. van Bibber, M. Vandegar, L. Visinelli, E. Vitagliano, J. K. Vogel, Z. Wang, A. Wickenbrock, L. Winslow, S. Withington, M. Wooten, J. Yang, B. A. Young, F. Yu, K. Zhou, and T. Zhou. Axion Dark Matter. *arXiv e-prints*, page arXiv:2203.14923, March 2022.
- [15] Danho Ahn, Ohjoon Kwon, Woohyun Chung, Wonjun Jang, Doyu Lee, Jhinhwan Lee, Sung Woo Youn, HeeSu Byun, Dojun Youm, and Yannis K. Semertzidis. Biaxially textured $yba_2cu_3o_{7-x}$ microwave cavity in a high magnetic field for a dark-matter axion search. *Phys. Rev. Appl.*, 17:L061005, Jun 2022.
- [16] Danho Ahn, Ohjoon Kwon, Woohyun Chung, Wonjun Jang, Doyu Lee, Jhinhwan Lee,

- Sung Woo Youn, Dojun Youm, and Yannis K. Semertzidis. Maintaining high Q-factor of superconducting $\text{YBa}_2\text{Cu}_3\text{O}_{7-x}$ microwave cavity in a high magnetic field. *arXiv e-prints*, page arXiv:1904.05111, April 2019.
- [17] Saebyeok Ahn, JinMyeong Kim, Boris I. Ivanov, Ohjoon Kwon, HeeSu Byun, Arjan F. van Loo, SeongTae Park, Junu Jeong, Soohyung Lee, Jinsu Kim, Ça ğlar Kutlu, Andrew K. Yi, Yasunobu Nakamura, Seonjeong Oh, Danho Ahn, SungJae Bae, Hyoungsoon Choi, Jihoon Choi, Yonuk Chong, Woohyun Chung, Violeta Gkika, Jihn E. Kim, Younggeun Kim, Byeong Rok Ko, Lino Miceli, Doyu Lee, Jiwon Lee, Ki Woong Lee, MyeongJae Lee, Andrei Matlashov, Pallavi Parashar, Taehyeon Seong, Yun Chang Shin, Sergey V. Uchaikin, SungWoo Youn, and Yannis K. Semertzidis. Extensive search for axion dark matter over 1 ghz with capp’s main axion experiment. *Phys. Rev. X*, 14:031023, Aug 2024.
- [18] M. S. Alam, S. E. Csorna, A. Fridman, L. Garren, M. D. Mestayer, R. S. Panvini, D. Andrews, P. Avery, K. Berkelman, D. G. Cassel, J. W. DeWire, R. Ehrlich, T. Ferguson, R. Galik, M. G. D. Gilchriese, B. Gittelmann, A. M. Halling, D. L. Hartill, D. Herrup, S. Holzner, M. Ito, J. Kandaswamy, V. Kistiakowsky, D. L. Kreinick, Y. Kubota, N. B. Mistry, F. Morrow, E. Nordberg, M. Ogg, R. Perchonok, R. Plunkett, A. Silverman, P. C. Stein, S. Stone, D. Weber, R. Wilcke, A. J. Sadoff, C. Bebek, J. Hassard, M. Hempstead, J. M. Izen, K. Kinoshita, W. W. MacKay, F. M. Pipkin, J. Rohlf, Richard Wilson, H. Kagan, K. Chadwick, J. Chauveau, P. Ganci, T. Gentile, Joan A. Guida, R. Kass, A. C. Melissinos, S. L. Olsen, G. Parkhurst, R. Poling, C. Rosenfeld, G. Rucinski, E. H. Thorndike, J. Green, R. Hicks, F. Sannes, P. Skubic, A. Snyder, R. Stone, A. Chen, M. Goldberg, N. Horwitz, A. Jawahery, M. Jibaly, P. Lipari, G. C. Moneti, G. Trahern, and H. Van Hecke. Search for axion production in Υ decay. *Phys. Rev. D*, 27:1665–1667, Apr 1983.
- [19] D. Alesini, D. Babusci, C. Braggio, G. Carugno, N. Crescini, D. D’Agostino, A. D’Elia, D. Di Gioacchino, R. Di Vora, P. Falferi, et al. Search for Galactic axions with a high-Q dielectric cavity. *Phys. Rev. D*, 106(5):052007, 2022.
- [20] D. Alesini, C. Braggio, G. Carugno, N. Crescini, D. D’Agostino, D. Di Gioacchino, R. Di Vora, P. Falferi, S. Gallo, U. Gambardella, et al. Galactic axions search with a superconducting resonant cavity. *Phys. Rev. D*, 99(10):101101, 2019.
- [21] D. Alesini, C. Braggio, G. Carugno, N. Crescini, D. D’Agostino, D. Di Gioacchino, R. Di Vora, P. Falferi, U. Gambardella, C. Gatti, et al. Search for invisible axion dark matter of mass $m_a = 43 \mu\text{eV}$ with the QUAX- $a\gamma$ experiment. *Phys. Rev. D*, 103(10):102004, 2021.

- [22] K. Altenmüller, V. Anastassopoulos, S. Arguedas-Cuendis, S. Aune, J. Baier, K. Barth, H. Bräuninger, G. Cantatore, F. Caspers, J. F. Castel, et al. New upper limit on the axion-photon coupling with an extended cast run with a xe-based micromegas detector. *Phys. Rev. Lett.*, 133:221005, Nov 2024.
- [23] Luis Álvarez-Gaumé and John Ellis. Eyes on a prize particle. *Nature Physics*, 7(1):2–3, January 2011.
- [24] B. Ary dos Santos Garcia, D. Bergermann, A. Caldwell, V. Dabhi, C. Diaconu, J. Diehl, G. Dvali, J. Egge, E. Garutti, S. Heyminck, F. Hubaut, A. Ivanov, J. Jochum, S. Knirck, M. Kramer, D. Kreikemeyer-Lorenzo, C. Krieger, C. Lee, D. Leppla-Weber, X. Li, A. Lindner, B. Majorovits, J. P. A. Maldonado, A. Martini, A. Miyazaki, E. Öz, P. Pralavorio, G. Raffelt, J. Redondo, A. Ringwald, J. Schaffran, A. Schmidt, F. Steffen, C. Strandhagen, I. Usherov, H. Wang, and G. Wieching. First search for axion dark matter with a Madmax prototype. *arXiv e-prints*, page arXiv:2409.11777, September 2024.
- [25] S. J. Asztalos, R. F. Bradley, L. Duffy, C. Hagmann, D. Kinion, D. M. Moltz, L. J. Rosenberg, P. Sikivie, W. Stoeffl, N. S. Sullivan, D. B. Tanner, K. van Bibber, and D. B. Yu. Improved rf cavity search for halo axions. *Phys. Rev. D*, 69:011101, Jan 2004.
- [26] S. J. Asztalos, G. Carosi, C. Hagmann, D. Kinion, K. van Bibber, M. Hotz, L. J. Rosenberg, G. Rybka, J. Hoskins, J. Hwang, P. Sikivie, D. B. Tanner, R. Bradley, and J. Clarke. Squid-based microwave cavity search for dark-matter axions. *Phys. Rev. Lett.*, 104:041301, Jan 2010.
- [27] S.J. Asztalos, G. Carosi, C. Hagmann, D. Kinion, K. van Bibber, M. Hotz, L. J. Rosenberg, G. Rybka, A. Wagner, J. Hoskins, et al. Design and performance of the admx squid-based microwave receiver. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 656(1):39–44, 2011.
- [28] Jose Aumentado. Superconducting parametric amplifiers: The state of the art in josephson parametric amplifiers. *IEEE Microwave Magazine*, 21(8):45–59, 2020.
- [29] Horace W. Babcock. The rotation of the Andromeda Nebula. *Lick Observatory Bulletin*, 498:41–51, January 1939.
- [30] K. M. Backes, D. A. Palken, S. Al Kenany, B. M. Brubaker, S. B. Cahn, A. Droster, Gene C. Hilton, Sumita Ghosh, H. Jackson, S. K. Lamoreaux, et al. A quantum-enhanced search for dark matter axions. *Nature*, 590(7845):238–242, 2021.

- [31] Varouzhan Baluni. CP-nonconserving effects in quantum chromodynamics. *Phys. Rev. D*, 19:2227–2230, Apr 1979.
- [32] R. Barbieri, C. Braggio, G. Carugno, C. S. Gallo, A. Lombardi, A. Ortolan, R. Pengo, G. Ruoso, and C. C. Speake. Searching for galactic axions through magnetized media: The QUAX proposal. *Physics of the Dark Universe*, 15:135–141, March 2017.
- [33] C. Bartram, T. Braine, E. Burns, R. Cervantes, N. Crisosto, N. Du, H. Korandla, G. Leum, P. Mohapatra, T. Nitta, L. J. Rosenberg, G. Rybka, J. Yang, John Clarke, I. Siddiqi, A. Agrawal, A. V. Dixit, M. H. Awida, A. S. Chou, M. Hollister, S. Knirck, A. Sonnenschein, W. Wester, J. R. Gleason, A. T. Hipp, S. Jois, P. Sikivie, N. S. Sullivan, D. B. Tanner, E. Lentz, R. Khatiwada, G. Carosi, N. Robertson, N. Woollett, L. D. Duffy, C. Boutan, M. Jones, B. H. LaRoque, N. S. Oblath, M. S. Taubman, E. J. Daw, M. G. Perry, J. H. Buckley, C. Gaikwad, J. Hoffman, K. W. Murch, M. Goryachev, B. T. McAllister, A. Quiskamp, C. Thomson, and M. E. Tobar. Search for invisible axion dark matter in the $3.3 - -4.2 \mu\text{eV}$ mass range. *Phys. Rev. Lett.*, 127:261803, Dec 2021.
- [34] C. Bartram, T. Braine, R. Cervantes, N. Crisosto, N. Du, G. Leum, L. J. Rosenberg, G. Rybka, J. Yang, D. Bowering, A. S. Chou, R. Khatiwada, A. Sonnenschein, W. Wester, G. Carosi, N. Woollett, L. D. Duffy, M. Goryachev, B. McAllister, M. E. Tobar, C. Boutan, M. Jones, B. H. LaRoque, N. S. Oblath, M. S. Taubman, John Clarke, A. Dove, A. Eddins, S. R. O’Kelley, S. Nawaz, I. Siddiqi, N. Stevenson, A. Agrawal, A. V. Dixit, J. R. Gleason, S. Jois, P. Sikivie, J. A. Solomon, N. S. Sullivan, D. B. Tanner, E. Lentz, E. J. Daw, M. G. Perry, J. H. Buckley, P. M. Harrington, E. A. Henriksen, and K. W. Murch. Axion dark matter experiment: Run 1b analysis details. *Phys. Rev. D*, 103:032002, Feb 2021.
- [35] C. Bartram et al. Dark matter axion search using a josephson traveling wave parametric amplifier. *Review of Scientific Instruments*, 94(4):044703, 04 2023.
- [36] Masha Baryakhtar, Junwu Huang, and Robert Lasenby. Axion and hidden photon dark matter detection with multilayer optical haloscopes. *Phys. Rev. D*, 98:035006, Aug 2018.
- [37] Masha Baryakhtar, Leslie Rosenberg, and Gray Rybka. Searching for the QCD Dark Matter Axion. *arXiv e-prints*, page arXiv:2504.10607, April 2025.
- [38] D. J. Bechis, T. W. Dombeck, R. W. Ellsworth, E. V. Sager, P. H. Steinberg, L. J. Tieg, J. K. Yoh, and R. L. Weitz. Search for axion production in low-energy electron bremsstrahlung. *Phys. Rev. Lett.*, 42:1511–1514, Jun 1979.

- [39] Joshua N. Benabou, Malte Buschmann, Joshua W. Foster, and Benjamin R. Safdi. Axion Mass Prediction from Adaptive Mesh Refinement Cosmological Lattice Simulations. , 134(24):241003, June 2025.
- [40] J. Beringer, J. F. Arguin, R. M. Barnett, K. Copic, O. Dahl, D. E. Groom, C. J. Lin, J. Lys, H. Murayama, C. G. Wohl, W. M. Yao, P. A. Zyla, C. Amsler, M. Antonelli, D. M. Asner, H. Baer, H. R. Band, T. Basaglia, C. W. Bauer, J. J. Beatty, V. I. Belousov, E. Bergren, G. Bernardi, W. Bertl, S. Bethke, H. Bichsel, O. Biebel, E. Blucher, S. Blusk, G. Brooijmans, O. Buchmueller, R. N. Cahn, M. Carena, A. Cecucci, D. Chakraborty, M. C. Chen, R. S. Chivukula, G. Cowan, G. D’Ambrosio, T. Damour, D. de Florian, A. de Gouvêa, T. DeGrand, P. de Jong, G. Dissertori, B. Dobrescu, M. Doser, M. Drees, D. A. Edwards, S. Eidelman, J. Erler, V. V. Ezhela, W. Fetscher, B. D. Fields, B. Foster, T. K. Gaisser, L. Garren, H. J. Gerber, G. Gerbier, T. Gherghetta, S. Golwala, M. Goodman, C. Grab, A. V. Gribsan, J. F. Gri-vaz, M. Grünewald, A. Gurtu, T. Gutsche, H. E. Haber, K. Hagiwara, C. Hagmann, C. Hanhart, S. Hashimoto, K. G. Hayes, M. Heffner, B. Heltsley, J. J. Hernández-Rey, K. Hikasa, A. Höcker, J. Holder, A. Holtkamp, J. Huston, J. D. Jackson, K. F. Johnson, T. Junk, D. Karlen, D. Kirkby, S. R. Klein, E. Klempt, R. V. Kowalewski, F. Krauss, M. Kreps, B. Krusche, Yu. V. Kuyanov, Y. Kwon, O. Lahav, J. Laiho, P. Langacker, A. Liddle, Z. Ligeti, T. M. Liss, L. Littenberg, K. S. Lugovsky, S. B. Lugovsky, T. Man-nel, A. V. Manohar, W. J. Marciano, A. D. Martin, A. Masoni, J. Matthews, D. Mil-stead, R. Miquel, K. Mönig, F. Moortgat, K. Nakamura, M. Narain, P. Nason, S. Navas, M. Neubert, P. Nevski, Y. Nir, K. A. Olive, L. Pape, J. Parsons, C. Patrignani, J. A. Peacock, S. T. Petcov, A. Piepke, A. Pomarol, G. Punzi, A. Quadt, S. Raby, G. Raf-felt, B. N. Ratcliff, P. Richardson, S. Roesler, S. Rolli, A. Romaniouk, L. J. Rosenberg, J. L. Rosner, C. T. Sachrajda, Y. Sakai, G. P. Salam, S. Sarkar, F. Sauli, O. Schnei-der, K. Scholberg, D. Scott, W. G. Seligman, M. H. Shaevitz, S. R. Sharpe, M. Silari, T. Sjöstrand, P. Skands, J. G. Smith, G. F. Smoot, S. Spanier, H. Spieler, A. Stahl, T. Stanev, S. L. Stone, T. Sumiyoshi, M. J. Syphers, F. Takahashi, M. Tanabashi, J. Terning, M. Titov, N. P. Tkachenko, N. A. Törnqvist, D. Tovey, G. Valencia, K. van Bibber, G. Venanzoni, M. G. Vincter, P. Vogel, A. Vogt, W. Walkowiak, C. W. Walter, D. R. Ward, T. Watari, G. Weiglein, E. J. Weinberg, L. R. Wiencke, L. Wolfenstein, J. Womersley, C. L. Woody, R. L. Workman, A. Yamamoto, G. P. Zeller, O. V. Zenin, J. Zhang, R. Y. Zhu, G. Harper, V. S. Lugovsky, and P. Schaffner. Review of particle physics. *Phys. Rev. D*, 86:010001, Jul 2012.
- [41] Gianfranco Bertone and Dan Hooper. History of dark matter. *Rev. Mod. Phys.*, 90:045002, Oct 2018.
- [42] J. R. Bond, J. Centrella, A. S. Szalay, and J. R. Wilson. Dark matter and shocked pancakes. In J. Audouze and J. Tran Thanh Van, editors, *Formation and Evolution of*

Galaxies and Large Structures in the Universe, volume 117 of *NATO Advanced Study Institute (ASI) Series C*, page 87, January 1984.

- [43] A. Bosma. *The distribution and kinematics of neutral hydrogen in spiral galaxies of various morphological types*. PhD thesis, University of Groningen, Netherlands, March 1978.
- [44] C. Boutan et al. Piezoelectrically tuned multimode cavity search for axion dark matter. *Phys. Rev. Lett.*, 121:261302, Dec 2018.
- [45] Samuel Boutin, David M. Toyli, Aditya V. Venkatramani, Andrew W. Eddins, Irfan Siddiqi, and Alexandre Blais. Effect of higher-order nonlinearities on amplification and squeezing in josephson parametric amplifiers. *Phys. Rev. Appl.*, 8:054030, Nov 2017.
- [46] T. Braine et al. Extended search for the invisible axion with the axion dark matter experiment. *Phys. Rev. Lett.*, 124:101303, Mar 2020.
- [47] Thomas Braine. *Superconducting Resonator Development for the Axion Dark Matter eXperiment*. PhD thesis, University of Washington, Dept. of Physics, 2024.
- [48] L. Brouwer, S. Chaudhuri, H.-M. Cho, J. Corbin, C. S. Dawson, A. Droster, J. W. Foster, J. T. Fry, P. W. Graham, R. Henning, K. D. Irwin, F. Kadribasic, Y. Kahn, A. Keller, R. Kolevator, S. Kuenstner, A. F. Leder, D. Li, J. L. Ouellet, K. M. W. Pappas, A. Phipps, N. M. Rapidis, B. R. Safdi, C. P. Salemi, M. Simanovskaia, J. Singh, E. C. van Assendelft, K. van Bibber, K. Wells, L. Winslow, W. J. Wisniewski, and B. A. Young. Proposal for a definitive search for gut-scale qcd axions. *Phys. Rev. D*, 106:112003, Dec 2022.
- [49] B. M. Brubaker, L. Zhong, Y. V. Gurevich, S. B. Cahn, S. K. Lamoreaux, M. Simanovskaia, J. R. Root, S. M. Lewis, S. Al Kenany, K. M. Backes, I. Urdinaran, N. M. Rapidis, T. M. Shokair, K. A. van Bibber, D. A. Palken, M. Malnou, W. F. Kindel, M. A. Anil, K. W. Lehnert, and G. Carosi. First results from a microwave cavity axion search at $24 \mu\text{eV}$. *Phys. Rev. Lett.*, 118:061302, Feb 2017.
- [50] Benjamin M. Brubaker. First results from the HAYSTAC axion search. *arXiv e-prints*, page arXiv:1801.00835, January 2018.
- [51] P. Brun, A. Caldwell, L. Chevalier, G. Dvali, P. Freire, E. Garutti, S. Heyminck, J. Jochum, S. Knirck, M. Kramer, C. Krieger, T. Lasserre, C. Lee, X. Li, A. Lindner, B. Majorovits, S. Martens, M. Matysek, A. Millar, G. Raffelt, J. Redondo, O. Reimann, A. Ringwald, K. Saikawa, J. Schaffran, A. Schmidt, J. Schütte-Engel, F. Steffen, C. Strandhagen, and G. Wieching. A new experimental approach to probe QCD axion

- dark matter in the mass range above $\{ 40 \} \{ \mu \} \{ \text{eV} \}$. *European Physical Journal C*, 79(3):186, March 2019.
- [52] Malte Buschmann, Joshua W. Foster, Anson Hook, Adam Peterson, Don E. Willcox, Weiqun Zhang, and Benjamin R. Safdi. Dark matter from axion strings with adaptive mesh refinement. *Nature Communications*, 13:1049, February 2022.
- [53] Allen Caldwell, Gia Dvali, Béla Majorovits, Alexander Millar, Georg Raffelt, Javier Redondo, Olaf Reimann, Frank Simon, and Frank Steffen. Dielectric haloscopes: A new way to detect axion dark matter. *Phys. Rev. Lett.*, 118:091801, Mar 2017.
- [54] Herbert B. Callen and Theodore A. Welton. Irreversibility and generalized noise. *Phys. Rev.*, 83:34–40, Jul 1951.
- [55] Carlton M. Caves. Quantum limits on noise in linear amplifiers. *Phys. Rev. D*, 26:1817–1839, Oct 1982.
- [56] R. Cervantes, G. Carosi, C. Hanretty, S. Kimes, B. H. LaRoque, G. Leum, P. Mohapatra, N. S. Oblath, R. Ottens, Y. Park, G. Rybka, J. Sinnis, and J. Yang. Search for $70 \mu\text{eV}$ dark photon dark matter with a dielectrically loaded multiwavelength microwave cavity. *Phys. Rev. Lett.*, 129:201301, Nov 2022.
- [57] R. Cervantes, G. Carosi, S. Kimes, C. Hanretty, B. H. LaRoque, G. Leum, P. Mohapatra, N. S. Oblath, R. Ottens, Y. Park, G. Rybka, J. Sinnis, and J. Yang. Admx-orpheus first search for $70 \mu\text{eV}$ dark photon dark matter: Detailed design, operations, and analysis. *Phys. Rev. D*, 106:102002, Nov 2022.
- [58] Hsin Chang, Jing-Yang Chang, Yi-Chieh Chang, Yu-Han Chang, Yuan-Hann Chang, Chien-Han Chen, Ching-Fang Chen, Kuan-Yu Chen, Yung-Fu Chen, Wei-Yuan Chiang, et al. First Results from the Taiwan Axion Search Experiment with a Haloscope at $19.6 \mu\text{eV}$. *Phys. Rev. Lett.*, 129(11):111802, 2022.
- [59] Hsin Chang, Han-Wen Liu, Hien Thi Doan, and Yung-Fu Chen. Direct calibration of microwave amplification chain on an axion cavity haloscope. *Review of Scientific Instruments*, 96(1):014505, 01 2025.
- [60] Ta-Pei Cheng and Ling-Fong Li. *Gauge theory of elementary particle physics*. Oxford science publications. Clarendon Press, Oxford, reprinted with corrections edition, 2011.
- [61] Jeff Chiles, Ilya Charaev, Robert Lasenby, Masha Baryakhtar, Junwu Huang, Alexana Roshko, George Burton, Marco Colangelo, Ken Van Tilburg, Asimina Arvanitaki, Sae Woo Nam, and Karl K. Berggren. New constraints on dark photon dark matter

- with superconducting nanowire detectors in an optical haloscope. *Phys. Rev. Lett.*, 128:231802, Jun 2022.
- [62] Douglas Clowe, Maruša Bradač, Anthony H. Gonzalez, Maxim Markevitch, Scott W. Randall, Christine Jones, and Dennis Zaritsky. A direct empirical proof of the existence of dark matter*. *The Astrophysical Journal*, 648(2):L109, aug 2006.
- [63] ADMX Collaboration, G. Carosi, C. Cisneros, N. Du, S. Durham, N. Robertson, C. Goodman, M. Guzzetti, C. Hanretty, K. Enzian, L. J Rosenberg, G. Rybka, J. Sinis, D. Zhang, John Clarke, I. Siddiqi, A. S. Chou, M. Hollister, A. Sonnenschein, S. Knirck, T. J. Caligiure, J. R. Gleason, A. T. Hipp, P. Sikivie, M. E. Solano, N. S. Sullivan, D. B. Tanner, R. Khatiwada, L. D. Duffy, C. Boutan, T. Braine, E. Lentz, N. S. Oblath, M. S. Taubman, E. J. Daw, C. Mostyn, M. G. Perry, C. Bartram, J. Laurel, A. Yi, T. A. Dyson, S. Ruppert, M. O. Withers, C. L. Kuo, B. T. McAllister, J. H. Buckley, C. Gaikwad, J. Hoffman, K. Murch, M. Goryachev, E. Hartman, A. Quiskamp, and M. E. Tobar. Search for axion dark matter from 1.1 to 1.3 ghz with admx, 2025.
- [64] N. Crescini, D. Alesini, C. Braggio, G. Carugno, D. D’Agostino, D. Di Gioacchino, P. Falferi, U. Gambardella, C. Gatti, G. Iannone, C. Ligi, A. Lombardi, A. Ortolan, R. Pengo, G. Ruoso, and L. Taffarello. Axion search with a quantum-limited ferromagnetic haloscope. *Phys. Rev. Lett.*, 124:171801, May 2020.
- [65] N. Crescini, D. Alesini, C. Braggio, G. Carugno, D. Di Gioacchino, C. S. Gallo, U. Gambardella, C. Gatti, G. Iannone, G. Lamanna, C. Ligi, A. Lombardi, A. Ortolan, S. Pagano, R. Pengo, G. Ruoso, C. C. Speake, and L. Taffarello. Operation of a ferromagnetic axion haloscope at $\omega = 1.166667\text{eV}$ with a width of 0.166667em . *European Physical Journal C*, 78(9):703, September 2018.
- [66] R.J. Crewther, P. Di Vecchia, G. Veneziano, and E. Witten. Chiral estimate of the electric dipole moment of the neutron in quantum chromodynamics. *Physics Letters B*, 88(1):123–127, 1979.
- [67] M. Davis, G. Efstathiou, C. S. Frenk, and S. D. M. White. The evolution of large-scale structure in a universe dominated by cold dark matter. , 292:371–394, May 1985.
- [68] M. Davis, J. Huchra, D. W. Latham, and J. Tonry. A survey of galaxy redshifts. II. The large scale space distribution. , 253:423–445, February 1982.

- [69] Ariane Dekker, Shin'ichiro Ando, Camila A. Correa, and Kenny C. Y. Ng. Warm dark matter constraints using Milky Way satellite observations and subhalo evolution modeling. , 106(12):123026, December 2022.
- [70] S. DePanfilis, A. C. Melissinos, B. E. Moskowitz, J. T. Rogers, Y. K. Semertzidis, W. U. Wuensch, H. J. Halama, A. G. Prodell, W. B. Fowler, and F. A. Nezrick. Limits on the abundance and coupling of cosmic axions at $4.5 < m_a < 5.0 \mu\text{eV}$. *Phys. Rev. Lett.*, 59:839–842, Aug 1987.
- [71] Giovanni Grilli di Cortona, Edward Hardy, Javier Pardo Vega, and Giovanni Villadoro. The QCD axion, precisely. *Journal of High Energy Physics*, 2016:34, January 2016.
- [72] R. Di Vora, A. Lombardi, A. Ortolan, R. Pengo, G. Ruoso, C. Braggio, G. Carugno, L. Taffarello, G. Cappelli, N. Crescini, et al. Search for galactic axions with a traveling wave parametric amplifier. *Phys. Rev. D*, 108(6):062005, 2023.
- [73] R. H. Dicke. The Measurement of Thermal Radiation at Microwave Frequencies. *Review of Scientific Instruments*, 17(7):268–275, 07 1946.
- [74] Michael Dine and Willy Fischler. The not-so-harmless axion. *Physics Letters B*, 120(1-3):137–141, January 1983.
- [75] Michael Dine, Willy Fischler, and Mark Srednicki. A simple solution to the strong CP problem with a harmless axion. *Physics Letters B*, 104(3):199–202, August 1981.
- [76] Akash V. Dixit, Srivatsan Chakram, Kevin He, Ankur Agrawal, Ravi K. Naik, David I. Schuster, and Aaron Chou. Searching for dark matter with a superconducting qubit. *Phys. Rev. Lett.*, 126:141302, Apr 2021.
- [77] DMRadio Collaboration, A. AlShirawi, V. Ankel, C. Bartram, J. Begin, C. Bell, J. N. Benabou, L. Brouwer, S. Chaudhuri, H. M. Cho, J. Corbin, W. Craddock, S. Cuadra, A. Droster, J. Echevers, J. W. Foster, J. T. Fry, P. W. Graham, R. Henning, K. D. Irwin, F. Kadribasic, Y. Kahn, A. Keller, R. Kolevatov, S. Kuenstner, A. Kunder, N. Kurita, A. F. Leder, D. Li, N. Otto, J. L. Ouellet, K. M. W. Pappas, A. Phipps, N. M. Rapisdis, B. R. Safdi, C. P. Salemi, M. Simanovskaia, J. Singh, P. Stark, E. C. van Assendelft, K. van Bibber, K. Wells, J. Wiedemann, L. Winslow, W. J. Wisniewski, D. Wright, A. K. Yi, and B. A. Young. Electromagnetic modeling and science reach of DMRadio- m^3 . *arXiv e-prints*, page arXiv:2302.14084, February 2023.
- [78] T. W. Donnelly, S. J. Freedman, R. S. Lytel, R. D. Peccei, and M. Schwartz. Do axions exist? *Phys. Rev. D*, 18:1607–1620, Sep 1978.

- [79] N. Du et al. Search for invisible axion dark matter with the axion dark matter experiment. *Phys. Rev. Lett.*, 120:151301, Apr 2018.
- [80] L. D. Duffy, P. Sikivie, D. B. Tanner, S. J. Asztalos, C. Hagmann, D. Kinion, L. J. Rosenberg, K. van Bibber, D. B. Yu, and R. F. Bradley. High resolution search for dark-matter axions. *Phys. Rev. D*, 74:012006, Jul 2006.
- [81] J. Dunkley, E. Komatsu, M. R. Nolta, D. N. Spergel, D. Larson, G. Hinshaw, L. Page, C. L. Bennett, B. Gold, N. Jarosik, J. L. Weiland, M. Halpern, R. S. Hill, A. Kogut, M. Limon, S. S. Meyer, G. S. Tucker, E. Wollack, and E. L. Wright. Five-Year Wilkinson Microwave Anisotropy Probe Observations: Likelihoods and Parameters from the WMAP Data. , 180(2):306–329, February 2009.
- [82] C. Edwards, R. Partridge, C. Peck, F. C. Porter, D. Antreasyan, Y. F. Gu, W. Kollmann, M. Richardson, K. Strauch, A. Weinstein, D. Aschman, T. Burnett, M. Cavalliforza, D. Coyne, C. Newman, H. F. W. Sadrozinski, D. Gelpman, R. Hofstadter, R. Horisberger, I. Kirkbride, H. Kolanoski, K. Königsmann, R. Lee, A. Liberman, J. O’Reilly, A. Osterheld, B. Pollock, J. Tompkins, E. Bloom, F. Bulos, R. Chestnut, J. Gaiser, G. Godfrey, C. Kiesling, W. Lockman, M. Oreglia, D. L. Scharre, and K. Wacker. Upper limit for $\frac{J}{\Psi} \rightarrow \gamma + axion$. *Phys. Rev. Lett.*, 48:903–906, Apr 1982.
- [83] S. M. Faber. GALAXY FORMATION AND COSMOLOGY. In *1st ESO-CERN Symposium: Large Scale Structure of the Universe, Cosmology and Fundamental Physics*, 11 1983.
- [84] S. M. Faber and J. S. Gallagher. Masses and mass-to-light ratios of galaxies. *Annual Review of Astronomy and Astrophysics*, 17(Volume 17, 1979):135–187, 1979.
- [85] Low Noise Factory. 0.6-2 GHz cryogenic low noise amplifier. https://www.lownoisefactory.com/files/1215/2585/7504/LNF-LNC0.6_2A.pdf.
- [86] Gary J. Feldman and Robert D. Cousins. Unified approach to the classical statistical analysis of small signals. *Phys. Rev. D*, 57:3873–3889, Apr 1998.
- [87] Jean Baptiste Joseph Fourier. *The Analytical Theory of Heat*. Cambridge Library Collection - Mathematics. Cambridge University Press, 2009.
- [88] K. C. Freeman. On the Disks of Spiral and S0 Galaxies. , 160:811, June 1970.
- [89] H.T. Friis. Noise figures of radio receivers. *Proceedings of the IRE*, 32(7):419–422, 1944.

- [90] Daniel Gilman, Simon Birrer, Anna Nierenberg, Tommaso Treu, Xiaolong Du, and Andrew Benson. Warm dark matter chills out: constraints on the halo mass function and the free-streaming length of dark matter with eight quadruple-image strong gravitational lenses. , 491(4):6077–6101, February 2020.
- [91] C. Goodman, M. Guzzetti, C. Hanretty, L. J. Rosenberg, G. Rybka, J. Sinnis, D. Zhang, John Clarke, I. Siddiqi, A. S. Chou, M. Hollister, S. Knirck, A. Sonnenschein, T. J. Caligiure, J. R. Gleason, A. T. Hipp, P. Sikivie, M. E. Solano, N. S. Sullivan, D. B. Tanner, R. Khatiwada, G. Carosi, C. Cisneros, N. Du, N. Robertson, N. Woollett, L. D. Duffy, C. Boutan, T. Braine, E. Lentz, N. S. Oblath, M. S. Taubman, E. J. Daw, C. Mostyn, M. G. Perry, C. Bartram, T. A. Dyson, S. Ruppert, M. O. Withers, C. L. Kuo, B. T. McAllister, J. H. Buckley, C. Gaikwad, J. Hoffman, K. Murch, M. Goryachev, E. Hartman, A. Quiskamp, and M. E. Tobar. Admx axion dark matter bounds around $3.3 \mu\text{eV}$ with dine-fischler-srednicki-zhitnitsky discovery ability. *Phys. Rev. Lett.*, 134:111002, Mar 2025.
- [92] Marco Gorghetto and Edward Hardy. Post-inflationary axions: a minimal target for axion haloscopes. *Journal of High Energy Physics*, 2023(5):30, May 2023.
- [93] Marco Gorghetto, Edward Hardy, and Giovanni Villadoro. More Axions from Strings. *SciPost Physics*, 10:050, February 2021.
- [94] Thierry Grenet, Rafik Ballou, Quentin Basto, Killian Martineau, Pierre Perrier, Pierre Pognat, Jérémie Quevillon, Nicolas Roch, and Christopher Smith. The grenoble axion haloscope platform (grahal): development plan and first results. 10 2021.
- [95] Ramesh Gupta, Michael Anerella, John Cozzolino, Piyush Joshi, Shresht Joshi, Stephen Plate, William Sampson, Honghai Song, Peter Wanderer, Woohyun Chung, Jingeun Kim, Byeong Rok Ko, Sung Woo Youn, and Yannis K. Semertzidis. Status of the 25 t, 100 mm bore hts solenoid for an axion dark matter search experiment. *IEEE Transactions on Applied Superconductivity*, 29(5):1–5, 2019.
- [96] M. Guzzetti, D. Zhang, C. Goodman, C. Hanretty, J. Sinnis, L. J. Rosenberg, G. Rybka, John Clarke, I. Siddiqi, A. S. Chou, M. Hollister, S. Knirck, A. Sonnenschein, T. J. Caligiure, J. R. Gleason, A. T. Hipp, P. Sikivie, M. E. Solano, N. S. Sullivan, D. B. Tanner, R. Khatiwada, G. Carosi, N. Du, C. Cisneros, N. Robertson, N. Woollett, L. D. Duffy, C. Boutan, T. Braine, N. S. Oblath, M. S. Taubman, E. Lentz, E. J. Daw, C. Mostyn, M. G. Perry, C. Bartram, T. A. Dyson, C. L. Kuo, S. Ruppert, M. O. Withers, A. K. Yi, B. T. McAllister, J. H. Buckley, C. Gaikwad, J. Hoffman, K. Murch, J. Russell, M. Goryachev, E. Hartman, A. Quiskamp, and M. E. Tobar. Improved receiver noise calibration for admx axion search: 4.54 to $5.41 \mu\text{eV}$. *Phys. Rev. D*, 111:092012, May 2025.

- [97] C. Hagmann, D. Kinion, W. Stoeffl, K. van Bibber, E. Daw, H. Peng, Leslie J Rosenberg, J. LaVeigne, P. Sikivie, N. S. Sullivan, D. B. Tanner, F. Nezrick, Michael S. Turner, D. M. Moltz, J. Powell, and N. A. Golubev. Results from a high-sensitivity search for cosmic axions. *Phys. Rev. Lett.*, 80:2043–2046, Mar 1998.
- [98] C. Hagmann, P. Sikivie, N. S. Sullivan, and D. B. Tanner. Results from a search for cosmic axions. *Phys. Rev. D*, 42:1297–1300, Aug 1990.
- [99] Seungyong Hahn et al. 45.5-tesla direct-current magnetic field generated with a high-temperature superconducting magnet. *Nature*, 570(7762):496–499, 2019.
- [100] M. Hatridge, R. Vijay, D. H. Slichter, John Clarke, and I. Siddiqi. Dispersive magnetometry with a quantum limited squid parametric amplifier. *Phys. Rev. B*, 83:134501, Apr 2011.
- [101] Sebastian Hoof, Felix Kahlhoefer, Pat Scott, Christoph Weniger, and Martin White. Axion global fits with Peccei-Quinn symmetry breaking before inflation using GAMBIT. *Journal of High Energy Physics*, 2019(3):191, March 2019.
- [102] Anson Hook. TASI Lectures on the Strong CP Problem and Axions. In *Proceedings of Theoretical Advanced Study Institute Summer School 2018 'Theory in an Era of Data' — PoS(TASI2018)*, volume 333, page 004, 2019.
- [103] Dieter Horns, Joerg Jaeckel, Axel Lindner, Andrei Lobanov, Javier Redondo, and Andreas Ringwald. Searching for wispy cold dark matter with a dish antenna. *Journal of Cosmology and Astroparticle Physics*, 2013(04):016, apr 2013.
- [104] Gabe Hoshino, Stefan Knirck, Mohamed H. Awida, Gustavo I. Canelo, Simon Corrodi, Martin di Federico, Benjamin Knepper, Alex Lapuente, Mira Littmann, David W. Miller, Donald V. Mitchell, Derrick Rodriguez, Mark K. Ruschman, Chiara P. Salemi, Matthew A. Sawtell, Leandro Stefanazzi, Andrew Sonnenschein, Gary W. Teafoe, Peter Winter, and Gigabread Collaboration. First Axionlike Particle Results from a Broadband Search for Wavelike Dark Matter in the 44 to 10^5 μeV Range with a Coaxial Dish Antenna. *Physical Review Letters*, 134(17):171002, May 2025.
- [105] Junu Jeong, SungWoo Youn, Sungjae Bae, Jihngyeun Kim, Taehyeon Seong, Jihn E. Kim, and Yannis K. Semertzidis. Search for invisible axion dark matter with a multiple-cell haloscope. *Phys. Rev. Lett.*, 125:221302, Nov 2020.

- [106] M. J. Jewell et al. New results from haystac’s phase ii operation with a squeezed state receiver. *Phys. Rev. D*, 107:072007, Apr 2023.
- [107] B.D. Josephson. Possible new effects in superconductive tunnelling. *Physics Letters*, 1(7):251–253, 1962.
- [108] Yonatan Kahn, Benjamin R. Safdi, and Jesse Thaler. Broadband and resonant approaches to axion dark matter detection. *Phys. Rev. Lett.*, 117:141801, Sep 2016.
- [109] N. Kaiser. Constraints on neutrino-dominated cosmologies from large-scale streaming motion. , 273:L17–L20, October 1983.
- [110] Rachel Kennedy, Carlos Frenk, Shaun Cole, and Andrew Benson. Constraining the warm dark matter particle mass with milky way satellites. *Monthly Notices of the Royal Astronomical Society*, 442(3):2487–2495, 06 2014.
- [111] R. Khatiwada, D. Bowring, A. S. Chou, A. Sonnenschein, W. Wester, D. V. Mitchell, T. Braine, C. Bartram, R. Cervantes, N. Crisosto, N. Du, L. J. Rosenberg, G. Rybka, J. Yang, D. Will, S. Kimes, G. Carosi, N. Woollett, S. Durham, L. D. Duffy, R. Bradley, C. Boutan, M. Jones, B. H. LaRoque, N. S. Oblath, M. S. Taubman, J. Tedeschi, John Clarke, A. Dove, A. Hashim, I. Siddiqi, N. Stevenson, A. Eddins, S. R. O’Kelley, S. Nawaz, A. Agrawal, A. V. Dixit, J. R. Gleason, S. Jois, P. Sikivie, N. S. Sullivan, D. B. Tanner, J. A. Solomon, E. Lentz, E. J. Daw, M. G. Perry, J. H. Buckley, P. M. Harrington, E. A. Henriksen, K. W. Murch, and G. C. Hilton. Axion dark matter experiment: Detailed design and operations. *Review of Scientific Instruments*, 92(12):124502, 12 2021.
- [112] Heejoo Kim, Junghyeon Park, and Minh Son. Axion dark matter from cosmic string network. *Journal of High Energy Physics*, 2024(7):150, July 2024.
- [113] Jihn E. Kim. Weak-interaction singlet and strong CP invariance. *Phys. Rev. Lett.*, 43:103–107, Jul 1979.
- [114] Jinsu Kim, Ohjoon Kwon, Ca Kutlu, Woohyun Chung, Andrei Matlashov, Sergey Uchaikin, Arjan Ferdinand van Loo, Yasunobu Nakamura, Seonjeong Oh, HeeSu Byun, et al. Near-quantum-noise axion dark matter search at capp around 9.5 μeV . *Phys. Rev. Lett.*, 130:091602, Feb 2023.
- [115] Younggeun Kim, Junu Jeong, SungWoo Youn, Sungjae Bae, Kiwoong Lee, Arjan F. van Loo, Yasunobu Nakamura, Seonjeong Oh, Taehyeon Seong, Sergey Uchaikin, et al. Experimental search for invisible axions as a test of axion cosmology around 22 uev. *Phys. Rev. Lett.*, 133:051802, Jul 2024.

- [116] G. Kirchhoff. I. on the relation between the radiating and absorbing powers of different bodies for light and heat. *London Edinburgh Philos. Mag. & J. Sci.*, 20(130):1–21, 1860.
- [117] Ohjoon Kwon, Doyu Lee, Woohyun Chung, Danho Ahn, HeeSu Byun, Fritz Caspers, Hyoungsoon Choi, Jihoon Choi, Yonuk Chong, Hoyong Jeong, et al. First Results from an Axion Haloscope at CAPP around $10.7 \mu\text{eV}$. *Phys. Rev. Lett.*, 126(19):191802, 2021.
- [118] S. K. Lamoreaux, K. A. van Bibber, K. W. Lehnert, and G. Carosi. Analysis of single-photon and linear amplifier detectors for microwave cavity dark matter axion searches. *Phys. Rev. D*, 88:035020, Aug 2013.
- [119] T. Lasserre, C. Afonso, J. N. Albert, J. Andersen, R. Ansari, É. Aubourg, P. Bareyre, F. Bauer, J. P. Beaulieu, G. Blanc, A. Bouquet, S. Char, X. Charlot, F. Couchot, C. Coutures, F. Derue, R. Ferlet, J. F. Glicenstein, B. Goldman, A. Gould, D. Graff, M. Gros, J. Haissinski, J. C. Hamilton, D. Hardin, J. de Kat, A. Kim, É. Lesquoy, C. Loup, C. Magneville, B. Mansoux, J. B. Marquette, É. Maurice, A. Milsztajn, M. Moniez, N. Palanque-Delabrouille, O. Perdureau, L. Prévot, N. Regnault, J. Rich, M. Spiro, A. Vidal-Madjar, L. Vigroux, S. Zylberajch, and EROS Collaboration. Not enough stellar mass Machos in the Galactic halo. , 355:L39–L42, March 2000.
- [120] Matthew Lawson, Alexander J. Millar, Matteo Pancaldi, Edoardo Vitagliano, and Frank Wilczek. Tunable axion plasma haloscopes. *Phys. Rev. Lett.*, 123:141802, Oct 2019.
- [121] S. Lee, S. Ahn, J. Choi, B. R. Ko, and Y. K. Semertzidis. Axion Dark Matter Search around $6.7 \mu\text{eV}$. *Phys. Rev. Lett.*, 124(10):101802, 2020.
- [122] Youngjae Lee, Byeongsu Yang, Hojin Yoon, Moohyun Ahn, Heejun Park, Byeonghun Min, DongLak Kim, and Jonghee Yoo. Searching for Invisible Axion Dark Matter with an 18 T Magnet Haloscope. *Phys. Rev. Lett.*, 128(24):241805, 2022.
- [123] Jesse Liu, Kristin Dona, Gabe Hoshino, Stefan Knirck, Noah Kurinsky, Matthew Malaker, David W. Miller, Andrew Sonnenschein, Mohamed H. Awida, Peter S. Barry, Karl K. Berggren, Daniel Bowring, Gianpaolo Carosi, Clarence Chang, Aaron Chou, Rakshya Khatiwada, Samantha Lewis, Juliang Li, Sae Woo Nam, Omid Noroozian, and Tony X. Zhou. Broadband solenoidal haloscope for terahertz axion detection. *Phys. Rev. Lett.*, 128:131801, Mar 2022.
- [124] K. Lundmark. Über die Bestimmung der Entfernungen, Dimensionen, Massen und Dichtigkeit für die nächstgelegenen anagalaktischen Sternsysteme. *Meddelanden från Lunds Astronomiska Observatorium Serie I*, 125:1–13, January 1930.

- [125] M. Malnou, T. F. Q. Larson, J. D. Teufel, F. Lecocq, and J. Aumentado. Low-noise cryogenic microwave amplifier characterization with a calibrated noise source. *Review of Scientific Instruments*, 95(3):034703, 03 2024.
- [126] Ben T. McAllister, Graeme Flower, Eugene N. Ivanov, Maxim Goryachev, Jeremy Bourhill, and Michael E. Tobar. The ORGAN Experiment: An axion haloscope above 15 GHz. *Phys. Dark Univ.*, 18:67–72, 2017.
- [127] JOHN F. MEEKINS, GILBERT FRITZ, TALBOT A. CHUBB, H. FRIEDMAN, and RICHARD C. HENRY. Physical sciences: X-rays from the coma cluster of galaxies. *Nature*, 231(5298):107–108, May 1971.
- [128] M. Milgrom. A modification of the Newtonian dynamics - Implications for galaxies. , 270:371–383, July 1983.
- [129] M. Milgrom. A modification of the newtonian dynamics : implications for galaxy systems. , 270:384–389, July 1983.
- [130] M. Milgrom. A modification of the Newtonian dynamics as a possible alternative to the hidden mass hypothesis. , 270:365–370, July 1983.
- [131] Alexander J. Millar, Steven M. Anlage, Rustam Balafendiev, Pavel Belov, Karl van Bibber, Jan Conrad, Marcel Demarteau, Alexander Droster, Katherine Dunne, Andrea Gallo Rosso, Jon E. Gudmundsson, Heather Jackson, Gagandeep Kaur, Tove Klaesson, Nolan Kowitt, Matthew Lawson, Alexander Leder, Akira Miyazaki, Sid Morampudi, Hiranya V. Peiris, Henrik S. Røising, Gaganpreet Singh, Dajie Sun, Jacob H. Thomas, Frank Wilczek, Stafford Withington, Mackenzie Wooten, Jens Dilling, Michael Febraro, Stefan Knirck, and Claire Marvinnay. Searching for dark matter with plasma haloscopes. *Phys. Rev. D*, 107:055013, Mar 2023.
- [132] E. O. Nadler, A. Drlica-Wagner, K. Bechtol, S. Mau, R. H. Wechsler, V. Gluscevic, K. Boddy, A. B. Pace, T. S. Li, M. McNanna, A. H. Riley, J. García-Bellido, Y. Y. Mao, G. Green, D. L. Burke, A. Peter, B. Jain, T. M. C. Abbott, M. Agüena, S. Allam, J. Annis, S. Avila, D. Brooks, M. Carrasco Kind, J. Carretero, M. Costanzi, L. N. da Costa, J. De Vicente, S. Desai, H. T. Diehl, P. Doel, S. Everett, A. E. Evrard, B. Flaugher, J. Frieman, D. W. Gerdes, D. Gruen, R. A. Gruendl, J. Gschwend, G. Gutierrez, S. R. Hinton, K. Honscheid, D. Huterer, D. J. James, E. Krause, K. Kuehn, N. Kuropatkin, O. Lahav, M. A. G. Maia, J. L. Marshall, F. Menanteau, R. Miquel, A. Palmese, F. Paz-Chinchón, A. A. Plazas, A. K. Romer, E. Sanchez, V. Scarpine, S. Serrano, I. Sevilla-Noarbe, M. Smith, M. Soares-Santos, E. Suchyta, M. E. C. Swanson, G. Tarle, D. L. Tucker, A. R. Walker, W. Wester, and DES Collaboration. Constraints on Dark

- Matter Properties from Observations of Milky Way Satellite Galaxies. , 126(9):091101, March 2021.
- [133] Ciaran O’Hare. cajohare/axionlimits: Axionlimits. <https://cajohare.github.io/AxionLimits/>, July 2020.
- [134] Ciaran O’Hare. Cosmology of axion dark matter. *PoS, COSMICWISPers:040*, 2024.
- [135] Jonathan L. Ouellet, Chiara P. Salemi, Joshua W. Foster, Reyco Henning, Zachary Bogorad, Janet M. Conrad, Joseph A. Formaggio, Yonatan Kahn, Joe Minervini, Alexey Radovinsky, Nicholas L. Rodd, Benjamin R. Safdi, Jesse Thaler, Daniel Winklehner, and Lindley Winslow. Design and implementation of the abracadabra-10 cm axion dark matter search. *Phys. Rev. D*, 99:052012, Mar 2019.
- [136] Sinziana Paduroiu. *Structure formation in warm dark matter cosmologies*. PhD thesis, University of Geneva, Dept. of Astronomy, 2015.
- [137] R. D. Peccei and Helen R. Quinn. Constraints imposed by CP conservation in the presence of pseudoparticles. *Phys. Rev. D*, 16:1791–1797, Sep 1977.
- [138] R. D. Peccei and Helen R. Quinn. CP conservation in the presence of pseudoparticles. *Phys. Rev. Lett.*, 38:1440–1443, Jun 1977.
- [139] Roberto D. Peccei. *The Strong CP Problem and Axions*, pages 3–17. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.
- [140] A. A. Penzias. Free Hydrogen in the Pegasus I Cluster of Galaxies. , 66:293, March 1961.
- [141] Planck Collaboration, P. A. R. Ade, N. Aghanim, C. Armitage-Caplan, M. Arnaud, M. Ashdown, F. Atrio-Barandela, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, J. G. Bartlett, E. Battaner, K. Benabed, A. Benoît, A. Benoit-Lévy, J. P. Bernard, M. Bersanelli, P. Bielewicz, J. Bobin, J. J. Bock, A. Bonaldi, L. Bonavera, J. R. Bond, J. Borrill, F. R. Bouchet, F. Boulanger, M. Bridges, M. Bucher, C. Burigana, R. C. Butler, E. Calabrese, J. F. Cardoso, A. Catalano, A. Challinor, A. Chamballu, H. C. Chiang, L. Y. Chiang, P. R. Christensen, S. Church, D. L. Clements, S. Colombi, L. P. L. Colombo, C. Combet, F. Couchot, A. Coulais, B. P. Crill, A. Curto, F. Cuttaia, L. Danese, R. D. Davies, R. J. Davis, P. de Bernardis, A. de Rosa, G. de Zotti, J. Delabrouille, J. M. Delouis, F. X. Désert, C. Dickinson, J. M. Diego, H. Dole, S. Donzelli, O. Doré, M. Douspis, J. Dunkley, X. Dupac, G. Efstathiou, F. Elsner, T. A. Enßlin, H. K. Eriksen, F. Finelli, O. Forni, M. Frailis, A. A. Fraisse, E. Franceschi, T. C. Gaier, S. Galeotta, S. Galli, K. Ganga, M. Giard, G. Giardino, Y. Giraud-Héraud,

E. Gjerløw, J. González-Nuevo, K. M. Górski, S. Gratton, A. Gregorio, A. Gruppuso, J. E. Gudmundsson, F. K. Hansen, D. Hanson, D. Harrison, G. Helou, S. Henrot-Versillé, C. Hernández-Monteaagudo, D. Herranz, S. R. Hildebrandt, E. Hivon, M. Hobson, W. A. Holmes, A. Hornstrup, W. Hovest, K. M. Huffenberger, G. Hurier, A. H. Jaffe, T. R. Jaffe, J. Jewell, W. C. Jones, M. Juvela, E. Keihänen, R. Keskitalo, K. Kiiveri, T. S. Kisner, R. Kneissl, J. Knoche, L. Knox, M. Kunz, H. Kurki-Suonio, G. Lagache, A. Lähteenmäki, J. M. Lamarre, A. Lasenby, M. Lattanzi, R. J. Laureijs, C. R. Lawrence, M. Le Jeune, S. Leach, J. P. Leahy, R. Leonardi, J. León-Tavares, J. Lesgourgues, M. Liguori, P. B. Lilje, M. Linden-Vørnle, V. Lindholm, M. López-Cañiego, P. M. Lubin, J. F. Macías-Pérez, B. Maffei, D. Maino, N. Mandolesi, D. Marinucci, M. Maris, D. J. Marshall, P. G. Martin, E. Martínez-González, S. Masi, M. Massardi, S. Matarrese, F. Matthai, P. Mazzotta, P. R. Meinhold, A. Melchiorri, L. Mendes, E. Menegoni, A. Mennella, M. Migliaccio, M. Millea, S. Mitra, M. A. Miville-Deschênes, D. Molinari, A. Moneti, L. Montier, G. Morgante, D. Mortlock, A. Moss, D. Munshi, J. A. Murphy, P. Naselsky, F. Nati, P. Natoli, C. B. Netterfield, H. U. Nørgaard-Nielsen, F. Noviello, D. Novikov, I. Novikov, I. J. O'Dwyer, F. Orioux, S. Osborne, C. A. Oxborrow, F. Paci, L. Pagano, F. Pajot, R. Paladini, D. Paoletti, B. Partridge, F. Pasian, G. Patanchon, P. Paykari, O. Perdereau, L. Perotto, F. Perrotta, F. Piacentini, M. Piat, E. Pierpaoli, D. Pietrobon, S. Plaszczynski, E. Pointecouteau, G. Polenta, and N. Ponthieu. Planck 2013 results. XV. CMB power spectra and likelihood. , 571:A15, November 2014.

- [142] Planck Collaboration, N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo, S. Basak, R. Battye, K. Benabed, J. P. Bernard, M. Bersanelli, P. Bielewicz, J. J. Bock, J. R. Bond, J. Borrill, F. R. Bouchet, F. Boulanger, M. Bucher, C. Burigana, R. C. Butler, E. Calabrese, J. F. Cardoso, J. Carron, A. Challinor, H. C. Chiang, J. Chluba, L. P. L. Colombo, C. Combet, D. Contreras, B. P. Crill, F. Cuttaia, P. de Bernardis, G. de Zotti, J. Delabrouille, J. M. Delouis, E. Di Valentino, J. M. Diego, O. Doré, M. Douspis, A. Ducout, X. Dupac, S. Dusini, G. Efstathiou, F. Elsner, T. A. Enßlin, H. K. Eriksen, Y. Fantaye, M. Farhang, J. Fergusson, R. Fernandez-Cobos, F. Finelli, F. Forastieri, M. Frailis, A. A. Fraisse, E. Franceschi, A. Frolov, S. Galeotta, S. Galli, K. Ganga, R. T. Génova-Santos, M. Gerbino, T. Ghosh, J. González-Nuevo, K. M. Górski, S. Gratton, A. Gruppuso, J. E. Gudmundsson, J. Hamann, W. Handley, F. K. Hansen, D. Herranz, S. R. Hildebrandt, E. Hivon, Z. Huang, A. H. Jaffe, W. C. Jones, A. Karacki, E. Keihänen, R. Keskitalo, K. Kiiveri, J. Kim, T. S. Kisner, L. Knox, N. Krachmalnicoff, M. Kunz, H. Kurki-Suonio, G. Lagache, J. M. Lamarre, A. Lasenby, M. Lattanzi, C. R. Lawrence, M. Le Jeune, P. Lemos, J. Lesgourgues, F. Levrier, A. Lewis, M. Liguori, P. B. Lilje, M. Lilley, V. Lindholm, M. López-Cañiego, P. M. Lubin, Y. Z. Ma, J. F. Macías-Pérez, G. Maggio, D. Maino, N. Mandolesi, A. Mangilli, A. Marcos-Caballero, M. Maris, P. G. Martin, M. Martinelli, E. Martínez-González, S. Matarrese, N. Mauri, J. D. McEwen, P. R. Meinhold, A. Melchiorri, A. Mennella, M. Migliaccio,

M. Millea, S. Mitra, M. A. Miville-Deschênes, D. Molinari, L. Montier, G. Morgante, A. Moss, P. Natoli, H. U. Nørgaard-Nielsen, L. Pagano, D. Paoletti, B. Partridge, G. Patanchon, H. V. Peiris, F. Perrotta, V. Pettorino, F. Piacentini, L. Polastri, G. Polenta, J. L. Puget, J. P. Rachen, M. Reinecke, M. Remazeilles, A. Renzi, G. Rocha, C. Rosset, G. Roudier, J. A. Rubiño-Martín, B. Ruiz-Granados, L. Salvati, M. Sandri, M. Savelainen, D. Scott, E. P. S. Shellard, C. Sirignano, G. Sirri, L. D. Spencer, R. Sunyaev, A. S. Suur-Uski, J. A. Tauber, D. Tavagnacco, M. Tenti, L. Toffolatti, M. Tomasi, T. Trombetti, L. Valenziano, J. Valiviita, B. Van Tent, L. Vibert, P. Vielva, F. Villa, N. Vittorio, B. D. Wandelt, I. K. Wehus, M. White, S. D. M. White, A. Zacchei, and A. Zonca. Planck 2018 results. VI. Cosmological parameters. , 641:A6, September 2020.

- [143] Planck Collaboration, N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo, S. Basak, K. Benabed, J. P. Bernard, M. Bersanelli, P. Bielewicz, J. J. Bock, J. R. Bond, J. Borrill, F. R. Bouchet, F. Boulanger, M. Bucher, C. Burigana, R. C. Butler, E. Calabrese, J. F. Cardoso, J. Carron, B. Casaponsa, A. Challinor, H. C. Chiang, L. P. L. Colombo, C. Combet, B. P. Crill, F. Cuttaia, P. de Bernardis, A. de Rosa, G. de Zotti, J. Delabrouille, J. M. Delouis, E. Di Valentino, J. M. Diego, O. Doré, M. Douspis, A. Ducout, X. Dupac, S. Dusini, G. Efstathiou, F. Elsner, T. A. Enßlin, H. K. Eriksen, Y. Fantaye, R. Fernandez-Cobos, F. Finelli, M. Frailis, A. A. Fraisse, E. Franceschi, A. Frolov, S. Galeotta, S. Galli, K. Ganga, R. T. Génova-Santos, M. Gerbino, T. Ghosh, Y. Giraud-Héraud, J. González-Nuevo, K. M. Górski, S. Gratton, A. Gruppuso, J. E. Gudmundsson, J. Hamann, W. Handley, F. K. Hansen, D. Herranz, E. Hivon, Z. Huang, A. H. Jaffe, W. C. Jones, E. Keihänen, R. Keskitalo, K. Kiiveri, J. Kim, T. S. Kisner, N. Krachmalnicoff, M. Kunz, H. Kurki-Suonio, G. Lagache, J. M. Lamarre, A. Lasenby, M. Lattanzi, C. R. Lawrence, M. Le Jeune, F. Levrier, A. Lewis, M. Liguori, P. B. Lilje, M. Lilley, V. Lindholm, M. López-Cañiego, P. M. Lubin, Y. Z. Ma, J. F. Macías-Pérez, G. Maggio, D. Maino, N. Mandolesi, A. Mangilli, A. Marcos-Caballero, M. Maris, P. G. Martin, E. Martínez-González, S. Matarrese, N. Mauri, J. D. McEwen, P. R. Meinhold, A. Melchiorri, A. Mennella, M. Migliaccio, M. Millea, M. A. Miville-Deschênes, D. Molinari, A. Moneti, L. Montier, G. Morgante, A. Moss, P. Natoli, H. U. Nørgaard-Nielsen, L. Pagano, D. Paoletti, B. Partridge, G. Patanchon, H. V. Peiris, F. Perrotta, V. Pettorino, F. Piacentini, G. Polenta, J. L. Puget, J. P. Rachen, M. Reinecke, M. Remazeilles, A. Renzi, G. Rocha, C. Rosset, G. Roudier, J. A. Rubiño-Martín, B. Ruiz-Granados, L. Salvati, M. Sandri, M. Savelainen, D. Scott, E. P. S. Shellard, C. Sirignano, G. Sirri, L. D. Spencer, R. Sunyaev, A. S. Suur-Uski, J. A. Tauber, D. Tavagnacco, M. Tenti, L. Toffolatti, M. Tomasi, T. Trombetti, J. Valiviita, B. Van Tent, P. Vielva, F. Villa, N. Vittorio, B. D. Wandelt, I. K. Wehus, A. Zacchei, and A. Zonca. Planck 2018 results. V. CMB power spectra and likelihoods. , 641:A5, September 2020.

- [144] Planck Collaboration, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, J. G. Bartlett, N. Bartolo, E. Battaner, K. Benabed, A. Benoît, A. Benoit-Lévy, J. P. Bernard, M. Bersanelli, P. Bielewicz, J. J. Bock, A. Bonaldi, L. Bonavera, J. R. Bond, J. Borrill, F. R. Bouchet, F. Boulanger, M. Bucher, C. Burigana, R. C. Butler, E. Calabrese, J. F. Cardoso, A. Catalano, A. Challinor, H. C. Chiang, P. R. Christensen, D. L. Clements, L. P. L. Colombo, C. Combet, A. Coulais, B. P. Crill, A. Curto, F. Cuttaia, L. Danese, R. D. Davies, R. J. Davis, P. de Bernardis, A. de Rosa, G. de Zotti, J. Delabrouille, F. X. Désert, E. Di Valentino, C. Dickinson, J. M. Diego, K. Dolag, H. Dole, S. Donzelli, O. Doré, M. Douspis, A. Ducout, J. Dunkley, X. Dupac, G. Efstathiou, F. Elsner, T. A. Enßlin, H. K. Eriksen, J. Fergusson, F. Finelli, O. Forni, M. Frailis, A. A. Fraisse, E. Franceschi, A. Frejsel, S. Galeotta, S. Galli, K. Ganga, C. Gauthier, M. Gerbino, M. Giard, E. Gjerløw, J. González-Nuevo, K. M. Górski, S. Gratton, A. Gregorio, A. Gruppuso, J. E. Gudmundsson, J. Hamann, F. K. Hansen, D. L. Harrison, G. Helou, S. Henrot-Versillé, C. Hernández-Monteagudo, D. Herranz, S. R. Hildebrandt, E. Hivon, W. A. Holmes, A. Hornstrup, K. M. Huffenberger, G. Hurier, A. H. Jaffe, W. C. Jones, M. Juvela, E. Keihänen, R. Keskitalo, K. Kiiveri, J. Knoche, L. Knox, M. Kunz, H. Kurki-Suonio, G. Lagache, A. Lähteenmäki, J. M. Lamarre, A. Lasenby, M. Lattanzi, C. R. Lawrence, M. Le Jeune, R. Leonardi, J. Lesgourgues, F. Levrier, A. Lewis, M. Liguori, P. B. Lilje, M. Lilley, M. Linden-Vørnle, V. Lindholm, M. López-Cañiego, J. F. Macías-Pérez, B. Maffei, G. Maggio, D. Maino, N. Mandolesi, A. Mangilli, M. Maris, P. G. Martin, E. Martínez-González, S. Masi, S. Matarrese, P. R. Meinhold, A. Melchiorri, M. Migliaccio, M. Millea, S. Mitra, M. A. Miville-Deschênes, A. Moneti, L. Montier, G. Morgante, D. Mortlock, S. Mottet, D. Munshi, J. A. Murphy, A. Narimani, P. Naselsky, F. Nati, P. Natoli, F. Noviello, D. Novikov, I. Novikov, C. A. Oxborrow, F. Paci, L. Pagano, F. Pajot, D. Paoletti, B. Partridge, F. Pasian, G. Patanchon, T. J. Pearson, O. Perdereau, L. Perotto, V. Pettorino, F. Piacentini, M. Piat, E. Pierpaoli, D. Pietrobon, S. Plaszczynski, E. Pointecouteau, G. Polenta, N. Ponthieu, G. W. Pratt, S. Prunet, J. L. Puget, J. P. Rachen, M. Reinecke, M. Remazeilles, C. Renault, A. Renzi, I. Ristorcelli, G. Rocha, M. Rossetti, G. Roudier, B. Rouillé d'Orfeuil, J. A. Rubiño-Martín, B. Rusholme, L. Salvati, M. Sandri, D. Santos, M. Savelainen, G. Savini, D. Scott, P. Serra, L. D. Spencer, M. Spinelli, V. Stolyarov, and R. Stompor. Planck 2015 results. XI. CMB power spectra, likelihoods, and robustness of parameters. , 594:A11, September 2016.
- [145] S. Posen, M. Checchin, O.S. Melnychuk, T. Ring, I. Gonin, and T. Khabiboulline. High-quality-factor superconducting cavities in tesla-scale magnetic fields for dark-matter searches. *Phys. Rev. Appl.*, 20:034004, Sep 2023.
- [146] David M Pozar. *Microwave engineering*. Wiley, Chichester, 4th edition edition, 2012.

- [147] John Preskill, Mark B. Wise, and Frank Wilczek. Cosmology of the invisible axion. *Physics Letters B*, 120(1-3):127–132, January 1983.
- [148] J. R. Primack and G. R. Blumenthal. What is the dark matter ? Implications for galaxy formation and particle physics. In J. Audouze and J. Tran Thanh Van, editors, *Formation and Evolution of Galaxies and Large Structures in the Universe*, volume 117 of *NATO Advanced Study Institute (ASI) Series C*, page 163, January 1984.
- [149] Aaron Quiskamp, Ben T. McAllister, Paul Altin, Eugene N. Ivanov, Maxim Goryachev, and Michael E. Tobar. Exclusion of axionlike-particleogenesis dark matter in a mass window above 100 μeV . *Phys. Rev. Lett.*, 132:031601, Jan 2024.
- [150] Aaron P. Quiskamp, Ben T. McAllister, Paul Altin, Eugene N. Ivanov, Maxim Goryachev, and Michael E. Tobar. Direct search for dark matter axions excluding alpogenesis in the 63- to 67- μeV range with the organ experiment. *Sci. Adv.*, 8(27):abq3765, 2022.
- [151] M. Renger, S. Pogorzalek, Q. Chen, Y. Nojiri, K. Inomata, Y. Nakamura, M. Partanen, A. Marx, R. Gross, F. Deppe, and et al. Beyond the standard quantum limit for parametric amplification of broadband signals. *npj Quantum Information*, 7(1), Nov 2021.
- [152] A. Rettaroli, D. Alesini, D. Babusci, C. Braggio, G. Carugno, D. D’Agostino, A. D’Elia, D. Di Gioacchino, R. Di Vora, P. Falferi, et al. Search for axion dark matter with the QUAX–LNF tunable haloscope. *Phys. Rev. D*, 110(2):022008, 2024.
- [153] M. S. Roberts and A. H. Rots. Comparison of Rotation Curves of Different Galaxy Types. , 26:483–485, August 1973.
- [154] D. H. Rogstad and G. S. Shostak. Gross Properties of Five Scd Galaxies as Determined from 21-CENTIMETER Observations. , 176:315, September 1972.
- [155] Vera C. Rubin and W. Kent Ford, Jr. Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions. , 159:379, February 1970.
- [156] Ken’ichi Saikawa, Javier Redondo, Alejandro Vaquero, and Mathieu Kaltschmidt. Spectrum of global string networks and the axion dark matter mass. , 2024(10):043, October 2024.
- [157] Chiara P. Salemi, Joshua W. Foster, Jonathan L. Ouellet, Andrew Gavin, Kaliroë M. W. Pappas, Sabrina Cheng, Kate A. Richardson, Reyco Henning, Yonatan Kahn, Rachel Nguyen, Nicholas L. Rodd, Benjamin R. Safdi, and Lindley Winslow. Search

- for low-mass axion dark matter with abracadabra-10 cm. *Phys. Rev. Lett.*, 127:081801, Aug 2021.
- [158] M. Schwarzschild. Mass distribution and mass-luminosity ratio in galaxies. , 59:273, September 1954.
- [159] M. A. Shifman, A. I. Vainshtein, and V. I. Zakharov. Can confinement ensure natural CP invariance of strong interactions? *Nuclear Physics B*, 166(3):493–506, April 1980.
- [160] I. Siddiqi, R. Vijay, F. Pierre, C. M. Wilson, M. Metcalfe, C. Rigetti, L. Frunzio, and M. H. Devoret. RF-driven josephson bifurcation amplifier for quantum measurement. *Physical Review Letters*, 93(20), Nov 2004.
- [161] P. Sikivie. Experimental tests of the "invisible" axion. *Phys. Rev. Lett.*, 51:1415–1417, Oct 1983.
- [162] P. Sikivie. Detection rates for "invisible"-axion searches. *Phys. Rev. D*, 32:2988–2991, Dec 1985.
- [163] P. Sikivie, N. Sullivan, and D. B. Tanner. Proposal for axion dark matter detection using an *lc* circuit. *Phys. Rev. Lett.*, 112:131301, Mar 2014.
- [164] Maximiliano Silva-Feaver, Saptarshi Chaudhuri, Hsiao-Mei Cho, Carl Dawson, Peter Graham, Kent Irwin, Stephen Kuenstner, Dale Li, Jeremy Mardon, Harvey Moseley, Richard Mule, Arran Phipps, Surjeet Rajendran, Zach Steffen, and Betty Young. Design Overview of DM Radio Pathfinder Experiment. *IEEE Transactions on Applied Superconductivity*, 27(4):2631425, June 2017.
- [165] M. Sivertz, J. Lee-Franzini, J. E. Horstkotte, C. Klopfenstein, R. D. Schamberger, L. J. Spencer, P. M. Tuts, T. Böhringer, P. Franzini, K. Han, S. W. Herb, G. Mageras, D. Peterson, E. Rice, J. K. Yoh, R. Imlay, G. Levman, W. Metcalf, V. Sreedhar, G. Blunar, H. Dietl, G. Eigen, E. Lorenz, F. Pauss, and H. Vogel. Upper limit for axion production in radiative Υ decay. *Phys. Rev. D*, 26:717–719, Aug 1982.
- [166] Daniel Huber Slichter. *Quantum Jumps and Measurement Backaction in a Superconducting Qubit*. PhD thesis, University of California, Berkeley, Dept. of Physics, 2011.
- [167] Sinclair Smith. The Mass of the Virgo Cluster. , 83:23, January 1936.
- [168] Ian P. Stern, ADMX, and ADMX-HF collaborations. Axion dark matter searches. *AIP Conference Proceedings*, 1604(1):456–461, 06 2014.

- [169] G. 't Hooft. Symmetry breaking through bell-jackiw anomalies. *Phys. Rev. Lett.*, 37:8–11, Jul 1976.
- [170] Tisserand, P., Le Guillou, L., Afonso, C., Albert, J. N., Andersen, J., Ansari, R., Aubourg, É., Bareyre, P., Beaulieu, J. P., Charlot, X., Coutures, C., Ferlet, R., Fouqué, P., Glicenstein, J. F., Goldman, B., Gould, A., Graff, D., Gros, M., Haissinski, J., Hamadache, C., de Kat, J., Lasserre, T., Lesquoy, É., Loup, C., Magneville, C., Marquette, J. B., Maurice, É., Maury, A., Milsztajn, A., Moniez, M., Palanque-Delabrouille, N., Perdureau, O., Rahal, Y. R., Rich, J., Spiro, M., Vidal-Madjar, A., Vigroux, L., and S. Zylberajch (The EROS-2 collaboration). Limits on the macho content of the galactic halo from the eros-2 survey of the magellanic clouds ***. *AA*, 469(2):387–404, 2007.
- [171] Michael S. Turner. Periodic signatures for the detection of cosmic axions. *Phys. Rev. D*, 42:3572–3575, Nov 1990.
- [172] Barry E. Turnrose and Herbert J. Rood. On the Hypothesis that the Coma Cluster is Stabilized by a Massive, Ionized Intergalactic Gas. , 159:773, March 1970.
- [173] Steven Weinberg. The $u(1)$ problem. *Phys. Rev. D*, 11:3583–3593, Jun 1975.
- [174] Steven Weinberg. A new light boson? *Phys. Rev. Lett.*, 40:223–226, Jan 1978.
- [175] S. D. M. White, C. S. Frenk, and M. Davis. Clustering in a neutrino-dominated universe. , 274:L1–L5, November 1983.
- [176] Robert N. Whitehurst and Morton S. Roberts. High-Velocity Neutral Hydrogen in the Central Region of the Andromeda Galaxy. , 175:347, July 1972.
- [177] F. Wilczek. Problem of strong p and t invariance in the presence of instantons. *Phys. Rev. Lett.*, 40:279–282, Jan 1978.
- [178] N. J. Woolf. On the Stabilization of Clusters of Galaxies by Ionized Gas. , 148:287, April 1967.
- [179] Walter Wuensch, S. De Panfilis-Wuensch, Y. K. Semertzidis, J. T. Rogers, A. C. Melissinos, H. J. Halama, B. E. Moskowitz, A. G. Prodell, W. B. Fowler, and F. A. Nezrick. Results of a Laboratory Search for Cosmic Axions and Other Weakly Coupled Light Particles. *Phys. Rev. D*, 40:3153, 1989.
- [180] K. Wurtz, B.M. Brubaker, Y. Jiang, E.P. Ruddy, D.A. Palken, and K.W. Lehnert. Cavity entanglement and state swapping to accelerate the search for axion dark matter. *PRX Quantum*, 2:040350, Dec 2021.

- [181] Byeongsu Yang, Hojin Yoon, Moo Hyun Ahn, Youngjae Lee, and Jonghee Yoo. Extended axion dark matter search using the capp18t haloscope. *Phys. Rev. Lett.*, 131:081801, Aug 2023.
- [182] Andrew K. Yi, Saebyeok Ahn, C Kutlu, JinMyeong Kim, Byeong Rok Ko, Boris I. Ivanov, HeeSu Byun, Arjan F. van Loo, SeongTae Park, Junu Jeong, et al. Axion dark matter search around $4.55 \mu\text{eV}$ with dine-fischler-srednicki-zhitnitskii sensitivity. *Phys. Rev. Lett.*, 130:071002, Feb 2023.
- [183] A. Zehnder, K. Gabathuler, and J.-L. Vuilleumier. Search for axions in specific nuclear gamma-transitions at a power reactor. *Physics Letters B*, 110(5):419–422, 1982.
- [184] Huaixiu Zheng, Matti Silveri, R. T. Brierley, S. M. Girvin, and K. W. Lehnert. Accelerating dark-matter axion searches with quantum measurement technology. *arXiv e-prints*, page arXiv:1607.02529, July 2016.
- [185] A. R. Zhitnitsky. On Possible Suppression of the Axion Hadron Interactions. (In Russian). *Sov. J. Nucl. Phys.*, 31:260, 1980.
- [186] L. Zhong, S. Al Kenany, K. M. Backes, B. M. Brubaker, S. B. Cahn, G. Carosi, Y. V. Gurevich, W. F. Kindel, S. K. Lamoreaux, K. W. Lehnert, et al. Results from phase 1 of the HAYSTAC microwave cavity axion experiment. *Phys. Rev. D*, 97(9):092001, 2018.
- [187] L. Zhong, S. Al Kenany, K. M. Backes, B. M. Brubaker, S. B. Cahn, G. Carosi, Y. V. Gurevich, W. F. Kindel, S. K. Lamoreaux, K. W. Lehnert, S. M. Lewis, M. Malnou, R. H. Maruyama, D. A. Palken, N. M. Rapidis, J. R. Root, M. Simanovskaia, T. M. Shokair, D. H. Speller, I. Urdinaran, and K. A. van Bibber. Results from phase 1 of the haystac microwave cavity axion experiment. *Phys. Rev. D*, 97:092001, May 2018.
- [188] F. Zwicky. Republication of: The redshift of extragalactic nebulae. *General Relativity and Gravitation*, 41(1):207–224, January 2009.
- [189] A. Álvarez Melcón, S. Argüedas Cuendis, J. Baier, K. Barth, H. Bräuninger, S. Calatroni, G. Cantatore, F. Caspers, J. F. Castel, S. A. Cetin, et al. First results of the cast-rades haloscope search for axions at $34.67 \mu\text{eV}$. *JHEP*, 21:075, 2020.