

©Copyright 2024

Thomas Braine

Superconducting Resonator Development for the Axion Dark Matter eXperiment

Thomas Braine

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

Doctor of Philosophy

University of Washington

2024

Reading Committee:

Gray Rybka, Chair

Gianpaolo Carosi

Alejandro Garcia

Program Authorized to Offer Degree:

Physics

University of Washington

Abstract

Superconducting Resonator Development for the Axion Dark Matter eXperiment

Thomas Braine

Chair of the Supervisory Committee:

Gray Rybka

Physics

The Axion Dark Matter eXperiment (ADMX) is the first axion haloscope search to reach DFSZ sensitivity in any mass range for the QCD axion. The QCD axion is a well-motivated dark matter candidate that additionally solves the strong CP problem in nuclear physics. A haloscope has three necessary components; a very strong external magnet, a high Q cavity resonator embedded in this field, and an ultra-sensitive RF read-out system. This dissertation mainly reports on the future development of the resonators for ADMX, specifically superconducting RF (SRF) cavities. It starts with an overview of axions, haloscopes, and the current ADMX experiment, followed by the work done at Lawrence Livermore National Laboratory (LLNL) on SRF cavities, and finally the preliminary results of the ADMX "hybrid" superconducting-copper Sidecar cavity in run 1D.

TABLE OF CONTENTS

	Page
List of Figures	iv
List of Tables	xx
Chapter 1: Introduction	1
Chapter 2: Axion theory and cosmology	5
2.1 The strong CP problem	5
2.2 Peccei-Quinn solution	8
2.3 The original and invisible axions	10
2.4 Axions as dark matter	12
2.5 Axion energy spectrum lineshapes	15
2.6 The current axion parameter space	16
Chapter 3: Overview of haloscopes	24
3.1 Components	24
3.2 Axion power in a haloscope	33
3.3 Scan rate	37
Chapter 4: Current ADMX experimental site	42
4.1 Brief history of ADMX	42
4.2 Site infrastructure	45
4.3 Magnet	46
4.4 Cavity	50
4.5 Helium dilution refrigerator	54
4.6 Cryogenic receiver	56
4.7 Helium liquefaction plant	67

Chapter 5:	Data-taking and analysis techniques for haloscopes	73
5.1	Data-taking cadence	73
5.2	Hot load measurements/Y-factor method 1	75
5.3	Y-factor method 2	78
5.4	SNRI	82
5.5	Axion search analysis procedure	83
Chapter 6:	Cavity Resonators	90
6.1	Formal introduction to cavities and modes	90
6.2	Quality factor	99
6.3	Geometric factors for multiple surfaces	103
6.4	A simple hybrid cavity model	106
6.5	Multi-mode decomposition	107
6.6	Antenna coupling	112
6.7	Measuring Q in the lab	118
6.8	High frequency axion Searches	131
Chapter 7:	SRF Cavity Theory	142
7.1	The Meissner effect	143
7.2	London theory	144
7.3	Ginzberg-Landau theory	146
7.4	BCS theory	149
7.5	Type-I and type-II superconductors	151
7.6	Fluxons in SRF cavities	154
Chapter 8:	Experimental SRF cavity work at LLNL	161
8.1	PPMS system	161
8.2	DC measurements	163
8.3	PPMS cavity design	173
8.4	Results of a bulk NbTi clamshell surface decomposition	178
8.5	Design and status of a Hakki-Coleman PPMS Cavity	184
Chapter 9:	ADMX Sidecar run 1D preliminary results	190
9.1	Previous Sidecar runs overview	190

9.2	Preparations and design changes for run 1D	196
9.3	Hybrid cavity characterization	209
9.4	Noise temperature calibrations	222
9.5	Axion search analysis	231
9.6	Future prospects for Sidecar	248
	Bibliography	251

LIST OF FIGURES

Figure Number	Page
2.1 A diagram of the neutron with its charge distribution and spin resulting in aligned electric and magnetic dipole moments. Applying a time-reversal symmetry (T), preserves the charge distribution but reverses the spin, causing the magnetic moment to reverse directions as well; By CPT theorem, this is the equivalent of a CP operation. Under just a parity operation, the charge distribution can be thought of as flipping with the respect to the spin axis, resulting in the electric dipole changing direction; a subsequent (C) operation will lead to the same result as the time reversal operation. One can see in the end, the neutron now has electric and magnetic dipole moments that are anti-aligned, meaning CP symmetry is not preserved in this system. Credit: public domain work from Andreas Knecht.	7
2.2 The "wine-bottle" potential of the field ϕ introduced by Peccei-Quinn. Credit: CERN	9
2.3 Feynman diagrams of axion coupling to photons and fermions.	11
2.4 N-Body Simulations from Romulus25 showing the axion frequency spectra for various Milky-Way-like galaxies (color lines), containing an axion of $m_a = 1\mu eV$. The solid black line shows a fit to the new line width model shown in Equation 2.17, whereas the dotted line shows the Maxwell-Boltzmann distribution (Equation 2.16) [75].	17
2.5 The axion-photon coupling exclusion limit plot for a variety of experimental techniques. The KSVZ and DFSZ theory lines are plotted in orange, with a yellow band around them marking the "discovery" range of where the invisible axion is expected to be. Laboratory/dedicated axion searches are highlighted in red colors, with ADMX and other haloscopes being the most sensitive and able to exclude axion within the yellow "discovery" range. Astrophysical constraints using measurements of various stars, galaxies etc. are highlighted in green; These measurements do not assume the axion is dark matter. More astrophysical measurements such as constraints set by the CMB are highlighted in blue; these assume the axion is the entirety of dark matter.	19

2.6	The axion parameter space centered around the haloscope exclusion region with various haloscope experiments, including ADMX, highlighted in red. The yellow region highlights the invisible axion discovery region, centered around the predicted DFSZ and KSVZ mass-coupling lines.	20
3.1	A toy model of an axion haloscope’s main components and detection process [28].	25
3.2	An equivalent circuit for a resonant cavity being driven by axion photonic power.	33
4.1	The ADMX superconducting solenoid magnet arriving at UW CENPA, Seattle, WA, from LLNL, Livermore, CA.	44
4.2	The ADMX experimental site.	46
4.3	A cut-away diagram of the ADMX insert within the solenoid magnet.	47
4.4	Extraction of the ADMX insert from magnet after run 1B.	48
4.5	Mounting ADMX insert to clean room ceiling after run 1B.	49
4.6	A cut-away diagram of the ADMX main magnet windings.	50
4.7	A sketch diagram of the bucking coil.	51
4.8	A diagram of the Knife-Edge clamp mechanism for ADMX run 1 cavities. A sharp knife edge is constructed off of an inset cut a couple of inches within the end of the tube. This inset will have the knife edge on the interior edge with a perpendicular face above it. The copper-plated end cap is placed on top of the knife edge, and the stainless steel compression ring inserts into the space above it within the inset of the cavity tube end. Compression screws create separation between the ring and end cap, causing the end cap to compress against the copper knife edge, and the compression ring against the opposing surface. This creates a strong seal between the copper end cap and tube, creating a solid current path for a good TM_{010} resonance.	52
4.9	The various ADMX tuning rod arrangements within the run 1 cavity.	53
4.10	A simple schematic of a generic ‘wet’ helium dilution refrigerator [3].	56
4.11	The ADMX dilution refrigerator mixing chamber mounted on top of cavity with heat exchangers (the coils) above providing pre-cooling to the $^4He/^3He$ mixture input.	57
4.12	The exterior ADMX dilution refrigerator components.	58

4.13	The schematic of the ADMX cryogenic receiver. " C_n " represents the circulators. "S" is a solenoid switch. "A" represents the attenuators. "JPA" is the Josephson parametric amplifier. "HFET" is a heterostructure field-effect transistor amplifier. The pump line is used for supplying a pump tone to the JPA. The weak port and cavity bypass port are used for transmission and reflection measurements respectively, which allow for in situ cavity characterization and calibration.	59
4.14	Left: The SQUIDadel chassis with circulators and coaxial switches and amplifier core. Top right: The amplifier core with amplifiers mounted and shielding removed. Bottom right: the JPA itself used in Run 1B.	63
4.15	The circuit schematic of the Radiall coaxial switch.	64
4.16	The pulsed power supply circuit design for a single coaxial switch.	65
4.17	The author holding the coaxial switch mounted to paddle for cryogenic testing. The cryogenic vessel located behind.	67
4.18	An outline of the ADMX liquid helium plant.	68
4.19	The old 500 LL ADMX helium dewar.	70
4.20	An electronic liquid helium fill pump. 4 of these pumps allow for automated fills of the magnet cryostat and its reservoir from the mother dewar.	71
4.21	The ADMX cryogenic control system display. This allows for operators to adjust automatic helium fill parameters and monitor helium inventory.	71
4.22	Installation of 2500 LL helium dewar.	72
5.1	Data-taking decision tree. Typically a first nibble pass-through always results in several re-scan triggers in the analysis due to the statistics associated with the chosen tuning rate. A persistence check means verifying that a potential signal appears in all the spectra acquired for it so far, and it is not intermittent. Primary synthetic axion injections are un-blinded by the operations team, whereas secondary injections are only known by an off-site SAG operator to check the operations teams' detection efficiency. RFI checks are made with a signal analyzer outside the cavity, ruling out any external signals such as local radio stations. Final checks include trying to modulate the signal power according to Equation 3.24 by looking for suppression of the signal off-resonance and modulating the magnet field strength. Figure credit from Ref. [21].	76
5.2	The run 1C hot load.	80
5.3	The run 1D hot load.	81

5.4	The filtered background shape ("Filtered BG") and average base-line fitted shape from the warm electronics during run 1B [21].	84
5.5	The raw spectrum preparation process.	86
5.6	The ADMX Gen 2 exclusion limit plot from October 2021. The light blue is from the run 1A 2018 paper, Orange is run 1B concluded in 2019, and in red is run 1C part 1 that concluded in 2021. Run 1C continued into 2022, to exclude the full range to DFSZ level, but these results have not been reported at the time of writing yet.	89
6.1	The electric field for different TM and TE modes in an empty Sidecar cavity [28]. The cross-sectional color plots show the magnitude of \vec{E} , while the vector plots represent \vec{E} , both for a single phase in the oscillation. In this way, red indicate extrema points in the electric field that can be used to label the mode.	97
6.2	A schematic of a coaxial antenna inserted into a cavity antenna port. The antenna itself only has two major parameters, L_{tip} , the length of exposed inner conductor and L_{outer} , the length of outer conductor before the coaxial connector. Based on the width of the cavity port, W_{port} , and the length of exposed outer conductor outside the cavity, d_{out} , the insertion depth, d_{in} , can be calculated.	113
6.3	A schematic of a typical 1-port cavity measurement with a vector network analyzer (VNA).	115
6.4	A series of network analyzer reflection measurements on a single ADMX run 2A cavity TM_{010} for varying antenna depths. They are indexed in the legend by how many thousandths of an inch the antenna is inserted into its port from the antenna tip. Additionally, a peak trend-line is interpolated across the series of Lorentzian dips, and a vertical line marks the critically coupled point.	117
6.5	A series of network analyzer transmission measurements on a single ADMX run 2A cavity TM_{010} for varying antenna depths. In this case, port 1 is being inserted into the cavity, while port 2 is fixed in a weakly coupled position. They are indexed in the legend by how many thousandths of an inch the antenna is inserted from its tip. Additionally, a peak trend-line is interpolated across the series of Lorentzian peaks, and a vertical line marks the critically coupled point.	121
6.6	An ideal S_{21} Q-circle. In principle, an ideal reflection Q-circle is the same, however in practice it is a much noisier fit [77].	124
6.7	A practical S_{21} Q-circle fit that reflects the effects of the transmission line and antenna setup [77].	125

6.8	An equivalent circuit model of a two-port cavity measurement with antennae and coaxial cables connected to a VNA: an ideal resonator in parallel with two transmission lines with complex impedance [77].	125
6.9	A Q-circle with an impedance mismatch [76].	130
6.10	The run 2 cavity system assembled at LLNL.	133
6.11	The run 2A cavity array original coarse tuner design.	134
6.12	Comparison of the TM_{010} mode weakly coupled reflection measurement for each cavity of the run 2 system. Drifts in the dip are due slightly different couplings. Only cavity C had an appreciably different frequency > 1 kHz. . .	135
6.13	Diagram of the new EFR site at Fermilab. The resonator system and amplifier systems are stored and cooled separately. This allows the cold electronics to be cooled further and minimize the magnetic shielding requirements from the 9.4 T magnet. The system is horizontal to accommodate the MRI magnet design.	136
6.14	The proposed ADMX-EFR 18 cavity array. Each cavity is a hexagonal clamshell to maximize packing efficiency and allow for superconducting coatings. Every cavity is topped with a rotor piezoelectric motor and geared tuning arm, as well as a linear piezoelectric motor for individual adjustment of the antenna coupling.	137
6.15	An open aluminum reduced length EFR prototype cavity being tested at LLNL.	138
6.16	The instantaneous scan rate estimated for full ADMX-EFR run. The orange line represents our conservative estimate (<i>threshold</i>). The green line represents our most optimistic estimate (<i>objective</i>) assuming optimal performance of all components, superconducting cavities and squeezing.	139
6.17	Diagram of a dielectrically-loaded cavity operating at a multi-wavelength mode. The dielectrics (grey boxes) are placed at every anti-node such that they suppress parts of the mode wavelength (\vec{E}) so that the axion form factor is non-zero ($\propto \vec{E} \cdot \vec{A}$) [35].	139
6.18	CAD model diagram for Orpheus Cavity [35].	140
6.19	The project Orpheus axion limit assuming a 1.5 T dipole magnet[35]. Figure adapted from [87].	141
7.1	The periodic table of elements with the critical temperature of known superconductors listed. Some elements only exhibit superconductivity above atmospheric pressure, known as high-pressure superconductors [53].	143

7.2	The phase diagram of a type-I superconductor. The transition line represents the parabolic relation of B_c , the critical field, with temperature. If $B < B_c$ and $T < T_c$, the critical temperature, the object is superconducting. Note that B_c and H_c are often used interchangeably [25].	152
7.3	Type-II superconducting film in the intermediate vortex phase, gray arrows represent the magnetic field penetrating through the quantized flux vortices, which generate their own circular surface currents (green arrows). The red arrow represents the Meissner surface current [109].	153
7.4	A scanning SQUID microscopy image of a 200 nm thick YBCO film (Type-II superconductor) after field cooling at $6.93 \mu\text{T}$ to 4 K. The black dots are flux vortices or fluxons. [117].	154
7.5	The phase diagram of a type-II superconductor. Both critical fields follow a parabolic relation with temperature. In zone a, $B < B_{c1}$, the object is completely superconducting. In zone b, $B_{c1} < B < B_{c2}$, the object is the mixed superconducting state, where the external field penetrates through quantized flux vortices, called fluxons. In zone c, $B > B_{c2}$, the field penetrates fully and the object is in the normal state [26].	155
8.1	Quantum Designs Physical Properties Measurement System (PPMS) from the outside.	162
8.2	Diagram of the PPMS bore.	163
8.3	Absolute field deviation of the 16 T PPMS magnet as function of distance from field center. Data was taken at three different magnet field strengths: 0.5, 1, and 2 T.	163
8.4	Percent field deviation of the PPMS magnet as function of distance from field center.	164
8.5	A typical DC resistance versus temperature graph for determining T_c . Resistance measurements are periodically taken as the sample is cooled through its critical temperature, At the critical temperature, the superconductor's resistance abruptly drops, making it easy to read off T_c . The residual resistance is from the leads connecting to the sample, in this case, being only a two-wire measurement. This particular sample is a coating of NbTiN made at LLNL that had an abnormally low T_c of ≈ 2.8 K; typically it is between 10-12 K. This transition temperature was also not particularly stable; upon cycling the sample back through critical temperature, several resistance measurements imply it was superconducting up to 3.1 K. All of this indicated that the coating was compromised by impurities in the deposition process.	165

8.6	A typical DC resistance versus magnetic field strength graph for determining B_{c2} . Resistance measurements are periodically taken as the magnet is ramped through the samples' critical fields. For type II superconductors, such as NbTiN in this case, the resistance will start to increase and plateau into a fully normal state above B_{c2} . The residual resistance is from this being only a two-wire measurement. This particular sample is a coating of NbTiN, the same used in Figure 8.5, made at LLNL. It had an abnormally low B_{c2} of ≈ 5.3 T; typically it is in excess of 15 T. The sample was cycled back through its critical field when the magnet was ramped down, generating the hysteresis loop structure; this is from trapped flux from the initial ramp shifting and increasing the measured resistance on the ramp back down. Both the low B_c and large hysteresis loop, indicate that the coating was compromised by impurities in the deposition process.	166
8.7	The circuit diagram for a 2-wire resistance measurement. A fixed current source and voltmeter are connected in parallel internally within the DMM. Two wire leads are then connected in series with the sample load of interest.	167
8.8	The circuit diagram for a 4-wire resistance measurement. One set of leads connects the sample to a fixed current source, while another connects to a voltmeter reading the voltage drop solely across the sample.	167
8.9	A PPMS puck for 4-wire resistance measurements of samples. This particular puck has 3 sets of 4 solder contacts ($V+$, $V-$, $I+$, and $I-$) to handle 3 simultaneous sample characterizations. The central square where samples would be placed is about 1 cm^2	169
8.10	A plot of the PPMS operating temperature and magnetic field versus time during a 33 hour measurement of a bulk NbTi sample. Samples are typically cooled overnight, and magnet operations begin the following day. Sample may lose connectivity due to thermal or magnetic strains, at which point the system is warmed and sample is removed for re-mounting (This was the case after the 2nd magnet ramp in the plot). Warming above the critical temperature in between magnet ramps can also be useful for expelling any trapped flux in the sample.	170
8.11	Measuring the B_{c2} of a bulk NbTi sample. $R1$ and $R2$ denote two different resistance bridges in the PPMS, that each were connected to a NbTi sample. Because the magnet was limited to 14 T, it can only be said that $H_{c2} > 14$ T. This is in agreement with similar alloys that exhibit an $H_{c2} \approx 14 - 15$ T. . .	171
8.12	Nb ₃ Sn coating deposited on bronze substrate. This deposition work was done at Florida State University by Lance Cooley's group, see Ref. [121] for details.	172

8.13	The Quantum Designs PPMS rotating puck stage. This allows for samples to be rotated with respect to the magnetic field axis from its typical perpendicular arrangement.	173
8.14	The resistance versus magnetic field of Nb ₃ Sn on bronze substrate for multiple angles with respect to the field axis; 0 degrees corresponds to the sample being perpendicular to the field axis, while at 90 degrees the sample surface is parallel. Successive Magnet ramps at 5 different angles were made. Note that the parallel arrangement does not begin its superconducting transition, while the other angles with perpendicular components begin to transition above ≈ 12.5 T.	174
8.15	PPMS clamshell cavity machined from a bulk NbTi square stock. A) The NbTi cavity assembled next to aluminum prototype. B) A profile of both NbTi halves side by side. C) One half of the copper clamshell close-up. D) A close-up of the cavity top without and with fixed pin antennae attached; additional holes are for mounting to probe.	175
8.16	A field overlay of the cavity modes' surface currents in Ansys™ HFSS software. A) The TM_{010} mode. B) The TM_{011} mode. C) The TM_{012} mode. Scale bar for the magnitude of the surface currents picture left in A/m.	176
8.17	The estimated fractional error in geometric factor by mode and subsurface. Note that the end caps factor is the sum of the top and bottom end cap contributions, and the total geometric factor is the sum of the end caps and walls. A) The error in the TM_{010} mode geometric factors. B) The error in the TM_{011} mode geometric factors. C) The error in the TM_{012} mode geometric factors.	177
8.18	The PPMS readout probe with an aluminum clamshell cavity mounted after a test. A) The entire probe with NbTi cavity mounted on bottom. B) The probe top with DC and SMA vacuum feed-through connections. C) The probe bottom with cavity, SMA to SMP connections, and temperature sensor line. D) Cavity bottom puck that ensures thermal connection to PPMS internal temperature sensor.	178
8.19	NbTi versus copper clamshell in PPMS system: The cavity quality factor for the first three TM modes as function of applied magnetic field. Data was taken at a temperature of 2 K.	180
8.20	NbTi versus copper clamshell in the PPMS system: The total mode surface resistance for the first three TM modes as function of applied magnetic field.	181

8.21	The mean surface resistance contributions of the walls and end caps of the NbTi cavity based on two different two-mode decompositions: TM_{010}/TM_{011} and TM_{010}/TM_{012} . This is compared to the mean copper cavity resistance, averaged over surface type for the same mode combinations.	182
8.22	The two two-mode surface resistance decompositions for the pure copper clamshell cavity vs field. A) Surface resistance of the cavity sub-surfaces for the TM_{010}/TM_{011} combination. B) The same plot for the TM_{010}/TM_{012} combination. C) The corresponding resistance per unit surface area for the TM_{010}/TM_{011} combination. D) The same plot for the TM_{010}/TM_{012} combination. These bottom plots account for the differing surface areas of the walls and end caps. Note the differing error bars are due to the higher uncertainty in the end cap geometric factors.	185
8.23	An aluminum clamshell dielectric loaded cavity for precise surface impedance measurements of samples. A) The two halves have circular insets for the sapphire cylinder and rectangular inset below for inserting sample blocks. The bottom circular inset acts as a mask to remove edge defects from the exposed surface. B) The cavity loaded with sapphire rod and aluminum sample block. Superconducting coatings would be sputtered to a substrate block of similar size and can be exchanged easily. C) The cavity contains the same exterior features for compatibility with the PPMS. D) A rising cylinder lifts the shorter cavity to center the sample surface in the PPMS magnet.	186
9.1	The ADMX Sidecar Magnetic Field Region. The magnitude of the magnetic field is plotted for each cross-section passing through the center of the Sidecar cavity. The radial coordinate using the ADMX insert center as zero. The axial 'Z' coordinate uses the center of the magnet coils as its zero. The 'Y' coordinate is the only coordinate that uses the Sidecar center as its zero. Note that this plot doesn't display the direction of the field vector, which is predominantly in the axial direction; off-axis components are quite small with the average $B_z/ B $ over the region being ≈ 0.981	192
9.2	The Sidecar JTWPA. A) The circuit diagram of the amplifier. B) A diagram of the amplifier micro structure. C) A picture of the entire amplifier chip. . .	193
9.3	An RF schematic of the run 1C and 1D Sidecar receiver. The JTWPA is placed between two circulators to prevent power being reflected back into the cavity and to allow for periodic RF calibrations. Although a completely separate chain, these magnetic-sensitive cold electronics resided in the SQUIDadel along side the main cavities' sensitive electronics protected by the bucking coil magnet.	194

9.4	The run 1C sidecar clam-shell cavity disassembled with 2" copper tuning rod and alumina axles.	195
9.5	The Nb ₃ Sn coated tuning rod from a side and top viewpoint. The rod is 1.75" in diameter, and features an inset for mounting the copper sapphire axle tabs. The center and right holes are tapped with 8-32 threads for mounting these tabs. The center hole goes through the length of the rod so that it could be hung by a wire in the coating furnace.	198
9.6	A close up of the run 1D Sidecar cavity tuning rod port profile as seen from one clam-shell half. The extension increases the length of port, minimizing RF power leakage out.	199
9.7	The simulated real and ideal axion form factor for the TM_{010} mode as a function of the mode resonant frequency from COMSOL simulations. The ideal form factor assumes a perfectly straight magnetic field, whereas the real form factor data interpolates the magnetic field map data shown in Figure 9.1. Because the real field isn't straight and constant, it is less than the ideal form factor. Both form factors drop to zero in mode crossing regions, of which there are 5 crossings visible in this plot. The real form factor data is used later in the axion search analysis for calculating expected axion signal power for each digitization scan.	200
9.8	The simulated TM_{010} mode geometric factors for the rod and exterior cavity as a function of mode resonant frequency over the entire tuning range of the cavity. The geometric factor tends to drop to zero near mode crossings, which makes it indistinguishable with TE modes, and therefore these regions are cut from the data set, producing the gaps in the plot data. When making a run plan, it was best to choose a region between mode crossings with relatively constant, well-separated geometric factors, such as the region between 5.2-5.6 GHz.	201
9.9	The simulated Sidecar exterior cavity geometric factors with error bars for each mode of interest in the target run range. The error was calculated by simulating $3^5 = 243$ possible geometries based on the machine tolerances of the design at each tun rod angle; because of the large computational time necessary, this was only done by moving the tuning rod in steps of 2.5 degrees from 10 to 50 degrees. This median value set of G factors is then plotted against the initial base geometry that used half degree steps. It agrees well in between mode crossings, but the uncertainty increases greatly near mode crossings.	203

9.10	A close up of the run 1D Sidecar cavity tuning port. The cavity is opened with tuning rod and axle resting inside and ball bearing inserted in place above. Special care had to be taken to ensure there was no rubbing of the rod with the cavity interior, and no rubbing of the axle assembly within the port. . . .	204
9.11	The Sidecar tuning rod with copper axle tab installed. The sapphire axle was epoxied to the copper tab with stycast. Vented, silver-plated 8-32 screws attached the tab to the Nb-substrate tuning rod.	205
9.12	The Sidecar cavity with rod, axles, ball-bearings and stainless steel axle coupler all installed. The copper clamp on the cavity bottom was used for rod thermalization, whereas the axle coupler on top connected to the rotor motor. Both of these locked the gaps on either side of the rod in place, ensuring no rubbing, before the other cavity half was added.	206
9.13	A room temperature mode map of the run 1D Sidecar taken at LLNL. Because of the high resistance of room temperature Nb ₃ Sn, Q_0 was only around 500-2500 throughout the tuning range.	207
9.14	The run 1D Sidecar cavity installed in the ADMX insert with piezoelectric motors mounted.	208
9.15	The copper tuning rod cooling strap attached to the bottom axle. This strap is then directly bolted to the mixing chamber plate. The sapphire rod keeps the rod electrically isolated, while still providing thermal conductivity for cooling.	208
9.16	The measured mode frequencies during cool down plotted onto the simulated mode map of which simulated geometric factors are derived. One hopes that all mode frequencies would align closely with a single tuning rod angle, but that is not the case here; the TM_{012} is at an anomalously lower frequency than the other 3, giving it a corresponding tuning angle of about 33.7 degrees, whereas the other lay much closer to the TM_{010} average angle of 26.7 degrees. This θ_{010} angle is plotted as a vertical line, whereas the θ_{avg} for all four modes is plotted with the upper and lower 1-sigma bounds plotted as well.	212
9.17	The unloaded quality factors versus the Sidecar temperature during the insert cool down for each of the relevant modes: A) TM_{010} , B) TM_{011} , C) TM_{012} , and D) TM_{020} . In addition, horizontal lines of the expected cryogenic quality factor for each mode with varying tuning rod surface resistance is shown. The $R_{rod} = R_{Cu}$ line corresponds to the expected cryogenic Q for an all copper cavity, whereas the $R_{rod} = 0$ would be the upper bound where the rod contributes zero surface resistance to the Q.	214

9.18	The unloaded quality factors of the TM_{010} mode versus Sidecar cavity temperature. The horizontal lines are the same as Figure 9.17, however they introduce an associated Q_{leak} value. A) shows the values with no associated leakage, $Q_{leak} \rightarrow \infty$. B) $Q_{leak} = 720$, which was the expected leakage Q for Sidecar during its previous run. One can see that a leakage Q this low is not possible because the experimental data greatly exceeds the expectation for an ideal hybrid cavity. C) $Q_{leak} = 10000$ represents a prescribed lower bound on the leakage Q where the experimental Q settles into a good hybrid cavity value between an all copper cavity and ideal superconducting hybrid. However, nothing prevents the actual Q_{leak} from being much higher than this, and the cavity surface resistances actually behaving less than ideally.	216
9.19	The ratio of the tuning rod resistance to copper during the Sidecar cool down. This is calculated using Equation 6.55, with the experimentally measured quality factors, and the expected quality factors for an all copper cavity. It doesn't account for any RF leakage, or sub-quality copper resistance within the cavity. This method will also produce different values than the implied resistance values in Figure 9.17.	217
9.20	The total mode resistance versus Sidecar cavity temperature for each mode of interest. The horizontal lines represent the expected total mode resistance for possible resistance values of the tuning rod; the total mode resistance is a weighted average based on the geometric factors of the superconducting surfaces and copper surfaces.	219
9.21	The Sidecar cavity temperature versus the average Sidecar magnetic field during the ramp. The initial ramp was done before the system had fully cooled to base temperature, whereas the later two points the insert had fully cooled to its base value at around 115 mK. This temperature shift was well correlated with a shift in the Sidecar mode frequencies with the TM_{010} frequency displayed on a separate y-axis here. This is also seen as an increase in Q within most of the subsequent ramp plots.	220
9.22	The unloaded quality factors for the modes during the magnet ramp. All of them seem to show signs of Q-degradation, indicative of superconducting behavior. The small rises in Q for the 3.6 T and 3.8 T values could be because of the much lower temperature of the insert when they were taken.	221
9.23	The unloaded quality factors for the modes during the magnet ramp compared to expected hybrid cavity Q values.	223

9.24	The TM_{010} mode unloaded quality factor versus magnetic field strength compared to the expected quality factor for varying rod resistances and different leakage Qs. A) shows the expected values for a cavity with no leakage ($Q_{leak} \rightarrow \infty$). B) shows a very low $Q_{leak} = 720$ that was estimated in the previous Sidecar run 1C. As one can see, this value is too low, as it implies the experimental data exceeds the ideal hybrid cavity case, which is nonphysical. C) shows a $Q_{leak} = 10000$, this was a prescribed lower bound estimate for Q_{leak} in the cool down data (see Figure 9.18) because it set the lowest temperature experimental value near the ideal hybrid cavity expected Q. In this case, the zero field data point does not fall near the $R_{rod} = 0$ value but below the $R_{rod} = R_{Cu}$ this might be because the cavity temperature was slightly warmer (700 mK versus 666 mK), changing the resistance of the rod, but one doesn't expect Q_{leak} to change with temperature at this regime because copper has already fully thermally contracted.	224
9.25	The expected ratio of tuning rod resistance to copper resistance for each mode during the magnet ramp. This value is derived using Equation 6.55 and the other analysis outlined in section 6.4.	225
9.26	The Sidecar total mode resistances versus the magnetic field strength during the run 1D magnet ramp. The expected total mode resistance for a hybrid cavity with varying tuning rod resistance is also shown as colored horizontal lines.	226
9.27	The relevant temperature sensor values versus time during the December magnet ramp cool down noise characterization period. This period was divided into to separate measurements because the frequency steps taken during the wide-power measurement changed part way through. The points that occurred after the gap in data at about 30,000 seconds constitute the temperatures used for the second measurement.	227
9.28	The relevant temperature sensor values versus time during the March anomalous fridge crash cool down noise characterization period. This crash heated the system above 1 kelvin, whereas the previous magnet ramp cool down only reached about 0.75 kelvin. This gave this measurement a much larger swing in temperature, which is better for the fit integrity.	228

9.29	An example fit from the 2nd half of the December magnet ramp cool down. The blue error bars represent the systematic error due to the temperature difference between the attenuator and cavity, which is carried over as an uncertainty in the measured power. Because this temperature uncertainty is relatively constant, the fractional uncertainty increases at lower temperatures, causing the increase in the power uncertainty near 0.1 K. The orange 1σ lines represent the error in the fit itself. These two errors are then combined when setting the error in the HFET noise temperature.	229
9.30	An example fit from the 1st half of the December magnet ramp cool down, which has significantly reduced errors compared to the 2nd half. The blue error bars represent the systematic error due to the temperature difference between the attenuator and cavity, which is carried over as an uncertainty in the measured power. The orange 1σ lines represent the error in the fit itself. These two errors are then combined when setting the uncertainty in the HFET noise temperature.	230
9.31	An example fit from the March fridge cool down. The blue error bars represent the systematic error due to the temperature difference between the attenuator and cavity, which is carried over as an uncertainty in the measured power. This measured temperature difference was roughly 20-30 mK higher than in the December cool down, increasing the overall error. The orange 1σ lines represent the error in the fit itself. These two errors are then combined when setting the error in the HFET noise temperature.	231
9.32	The registered temperature difference between the attenuator and cavity during the December magnet ramp cool down measurement.	232
9.33	The registered temperature difference between the attenuator and cavity during the March cool down measurement.	232
9.34	The HFET noise temperature versus frequency for each of the three noise calibration measurements, as well as the combined result for all three. Note that the December cool down part 2 measurement was taken at a different set of frequencies than the other two. Since all three measurements seem to disagree with one another, a conservative combined result was made to be used in the subsequent axion search analysis.	233
9.35	The Sidecar warm receiver shape and its fit for background extraction. . . .	235
9.36	The Sidecar warm receiver residue after the background fit is removed. . . .	237

9.37	The Sidecar loaded quality factor (transmission) over the course of the 14 hour data set used for analysis. Evident in this plot is a small shift in Q (≈ 150) due to a helium reservoir fill that took place roughly at hour 4. The value recovers after about 2 hours, but it is shifted by ≈ 75 ; This is probably within the uncertainty in Q for any individual fit, so it is not considered a problem at the moment.	238
9.38	The Sidecar TWPA gain performance over the course of a 14 hour data set. .	239
9.39	An example raw spectrum shown in blue with its cold receiver Padé fit shown in orange.	240
9.40	An example raw spectrum with the warm receiver shape removed, but no cold receiver fit applied; there is still residual structure unaccounted for.	241
9.41	The raw spectrum after the cold receiver fit being removed. It is now normalized, and should be flat and centered about zero power.	241
9.42	The final Lorentzian weighted spectrum sent for co-adding into a grand spectrum. The Lorentzian weighting applied matches the associated resonant frequency and quality factor at the time of the digitization.	242
9.43	The standard deviation increase versus the spectrum index of the 408 spectra included in the data set. The data is centered about the expected mean with a ± 0.01 variance as expected.	243
9.44	The individual bin deviations of raw spectra about the mean fitted to a Gaussian distribution.	243
9.45	The number of raw spectra that contributed to each frequency bin in the grand spectrum.	244
9.46	The DFSZ SNR of the grand spectrum versus frequency. Its Lorentzian shape is a product of the Lorentzian weighting of the individual raw spectra. The off-center peak is due to a mismatch between the spectra centers and the measured resonant frequency of the cavity.	245
9.47	The grand power excess spectrum versus frequency for the Sidecar run 1D sample dataset. It is apparent that there is significant noise from the receiver leaking through due to the compounding error from individual spectra fits. Ideally, this spectrum should be Gaussian for a static frequency dataset. . .	246
9.48	A histogram of the SNR excess in the grand spectrum.	247
9.49	Preliminary axion exclusion limits for Sidecar run 1D, based on a 14 hour dataset. This work is shown in green against the previous Sidecar run 1A limits in this region (red), and the CAST broadband limit in this region (blue).	248

9.50 A zoomed in exclusion limit plot of the Sidecar run 1D preliminary dataset shown in green. The previous Sidecar run 1A limit is shown in red, and the CAST experimental limit shown in blue. 249

LIST OF TABLES

Table Number	Page	
4.1	A summary of the minimum energy and initial pulse power to fire the Radial coaxial switch at various temperatures. The RC time constant was fixed at 4.7 ms.	68
5.1	A Summary of the frequency and fraction of time each process takes in a data-taking cycle of ADMX based on data from run 1B. Axion search data is only acquired during the data digitization cycle. Since the JPA bias procedure only occurs every 5-7 cycles the average fraction of time spent digitizing is kept at 98%. SAG stands for Synthetic Axion Generator.	75
6.1	A summary of the empty cavity characteristics and base frequencies for the various ADMX resonator systems. In the case of the ADMX run 2 and EFR arrays, these are the measurements of a single empty cavity within the multi-cavity system. All lengths and frequencies are listed in units of centimeters and gigahertz.	96
8.1	Critical temperature and fields of various materials measured at LLNL. The NbTiN coating was produced at LLNL. The NbTi was acquired from American Elements. The Nb ₃ Sn was supplied by Lance Cooley's group at Florida State University (FSU).	171
8.2	Geometric factors for clamshell cavity design for first three TM modes by interior surfaces. These were calculated from simulation in the HFSS software package and are expressed in Ohms.	176
8.3	PPMS clamshell cavity multi-mode decomposition: determinant of weight matrix and associated error for the three possible mode combinations.	182
8.4	A summary of the maximum estimated uncertainties for the surface decomposition measurement setup. They are expressed as percentages of their nominal values with the exception of surface resistance, since many values measured were effectively zero while the cavity was superconducting. Note that ΔC_{ms}^{-1} and ΔR_s does not include the TM_{011}/TM_{012} decomposition because of near-zero determinant.	183

8.5	A comparison of geometric factors expressed in Ohms between the LLNL PPMS-compatible clamshell Hakki-Coleman Resonator versus a previous, more traditionally designed Hakki-Coleman resonator used in Ref. [8]. This highlights a design flaw in our resonator; the sample surface is much less significant than the other surfaces in the resonator, therefore making accurate surface impedance measurements of the sample very difficult.	189
9.1	The geometric factors in Ohms of the Sidecar cavity for a tuning rod angle of 26.74 degrees. This angle corresponds to the TM_{010} mode frequency that the Sidecar cavity was left at for the start of run 1D, and therefore they are the relevant values used for characterizing the cavity in cool down and magnet ramp. Unfortunately, this was near a mode-crossing for the TM_{020} mode so some errors are quite large for that mode. There is anomalous behavior in the values for the TM_{012} mode as well, which has much larger values and smaller errors than the other modes. TM_{012} mode also had a measured mode frequency that did not correspond to this rod angle.	202
9.2	The average value of run parameters and their uncertainties for a random 14 hour data-set used for initial Sidecar run 1D results.	236

ACKNOWLEDGMENTS

First and foremost, I would like to thank my University of Washington advisor, Gray Rybka, and my LLNL advisor, Gianpaolo Carosi, who are also both co-spokespersons for ADMX. It was a great opportunity to research at Lawrence Livermore National Laboratory with GP on future resonator development, while also getting to travel back to UW to help with the current experimental operations at CENPA.

Thank you to the other members of my committee: Alejandro Garcia, David Kaplan, and Ben Marwick.

Thank you to Leslie Rosenberg for being my first contact with ADMX, and giving me the initial opportunity join the collaboration as first-year graduate student.

I'd also like to thank the senior graduate students and postdocs that mentored me along the way: Nick Du, Raphael Cervantes, Chelsea Bartram, Nicole Crisosto, Nathan Woollett, Stephan Knirck, and Dan Zhang. Additionally, the ADMX graduate students that came after me, Michaela Guzzetti and Jimmy Sinnis, who helped with on-site Sidecar operations and analysis questions towards the end of my graduate tenure.

Thank you to Sam Posen and the others at the SQMS center at Fermilab for not only producing the superconducting tuning rod for Sidecar, but providing SRF accelerator cavity expertise.

Thank you to Lance Cooley, Wenura Winthrage, and Andre Juliao for the superconducting samples and material science expertise.

Thank you to Alex Baker for not only getting us time on the PPMS system but all the help with mounting samples.

Thank you to the LLNL technicians that machined and fixed the many parts that made

this research possible: Sean Durham and Nathan 'Eric' Robertson.

Thank you to the many intrepid helium liquefier and dilution refrigerator operators that kept ADMX running at UW: Nick Force, Seth Kimes, Grant Leum, Charles Hanretty, and Cyrus Goodman.

Thank you to the members of the Rare Event Detection group at LLNL for welcoming me in at the lab as well.

Thank you to my good friends, Kian Alden, Trevor Thomas, and Chance Pryor for the well-needed phone calls and trips along the way.

Thank you to my wonderful cats, Higgs and Boson, who stayed right by my side throughout writing this dissertation.

Finally, I'd like to thank my parents, Tim Braine and Judith Ivey, and my grandfather, Dr. Nathan Ivey, who've supported me to the fullest in pursuit of this degree.

DEDICATION

To my grandmother, Dorothy Lee Lewis Ivey, 1922-2023.

"I finally got it done Grammy!"

Chapter 1

INTRODUCTION

One of the bigger ironies when reflecting on the immense progress in science and physics of the last century is that physicists actually know less about what the majority of the universe is comprised of than what it was thought to comprise of a century earlier. What I allude to is the perplexing discovery of dark matter from astronomical and cosmological measurements. The dark matter problem arises when one tries to account for all the matter that can be seen by interactions with the electromagnetic force, and compare it to the amount of matter that can be "seen" by interactions with the gravitational force; it becomes apparent that the vast majority of the matter in our galaxy, 85% at current measurements, only interacts via gravitation appreciably, and is otherwise 'dark' and non-luminous when trying to detect it electromagnetically. Furthermore, scientists are confident this matter is non-baryonic and unlike any particles we know [10, 22]. The crown jewel of physics theories, the standard model, seems to only explain the interactions of a small fraction of the universes' constituents; an experimental observation pointing not to just a small hole in the model, but a sea of darkness around it. However, it isn't to say that this dark matter isn't interacting with the electromagnetic force entirely, but perhaps it is such a weak interaction as to avoid detection by current technologies. It is the job of physicists to ask what new particle could have the right exotic properties if a few adaptations to the standard model could be made?

Another lesser known mystery surrounding the standard model of particle physics is known as the "Strong CP Problem." Although its easy to explain dark matter in plain terms as "missing" matter to a laymen, the strong CP problem remains lesser known because it requires understanding physical law in the language of symmetries. In essence, the standard model describes the interactions of the electromagnetic, weak nuclear, and strong nuclear

forces in terms of the symmetries they preserve and violate. A force that is symmetrical under a given transformation, such as charge conjugation (C), will produce the same phenomenon after such a transformation (changing the signs of all the charges in the case of charge conjugation). Quantum chromodynamics (QCD), the theory that describes the nature of the strong nuclear force, contains a term in its Lagrangian that is explicitly "CP"-violating, meaning a transformation of charge conjugation and parity (changing the charge and spin signs of all the particles in the system) should not produce the same system. The problem, however, is that experiments show that the strong nuclear force doesn't seem to be CP-violating at all; a classic piece of evidence for this is the measurement of the neutron's electric dipole moment (nEDM) that is consistent with zero. Original estimates based on QCD would have the nEDM be many orders of magnitude larger than what its upper value is measured to be. This has forced theorists to lower the strength parameter of the CP-violation, θ , to a value of below 10^{-10} . This is what constitutes a 'fine-tuning' problem; one is prescribing an incredibly specific number to make the theory work without any mechanism as to why that number is that specific value.

In 1977, Helen Quinn and Roberto Peccei proposed a mechanism that would explain the near zero value of θ , known now as Peccei-Quinn theory [90]. If true, it implied the existence of a new particle dubbed the axion, named after a popular detergent because it cleaned up the strong CP problem. At the time, it seemed like new particles were being discovered everyday; it was only a matter of time until this axion was found in an accelerator. And yet, this never happened, which pushed the energy scales higher, and lowered the expected mass of the axion. It was then realized that this lighter axion had a lot of properties that a dark matter particle would have; it would have mass, feebly interact with photons, and would be created in abundance during the early universe. They are produced non-thermally so that they are naturally cold (unlike neutrinos), which allows them to easily cluster around galaxies early in their formation. What makes them exotic, however, is that these particles would be so light that they would have de Broglie wavelengths on the order of kilometers, and would behave more like waves than particles. Because of these properties, physicists

initially thought that these particles could never be detected, and coined them "invisible axions."

An experimental apparatus was then proposed by Pierre Sikvie in 1983, to detect this axion field within our galaxies' dark matter halo [106]: a resonant cavity embedded in a strong static magnetic field with an ultra-low noise receiver. The static magnetic field provides a large density of virtual photons to stimulate the axion to photon conversion. If the resonant cavity is tuned to the frequency of these photons, the build up of power could be enough to detect the signal over the noise via the ultra-low noise receiver. This was dubbed an axion *haloscope*, and plans were made to one day construct such a device.

Just over 40 years later, the world now has several axion haloscopes, and many more planned to be built. As large WIMP dark matter experiments continue to show no events, the experimental axion field is gaining momentum. However, there is no doubt who were the pioneers in this field: the Axion Dark Matter eXperiment (ADMX). Initially formed in the 1990s, for over thirty years, ADMX has paved the way to creating the first discovery-sensitive haloscope. This has involved the development of several novel technologies to be possible.

For the past 6 years, it has been my privilege to work alongside the small team of scientists that make up the ADMX collaboration, developing yet another novel axion haloscope technology. My graduate work has been focused on the development and testing of the current and future cavity resonators used in ADMX. More specifically, this dissertation focuses on efforts to replace the current copper cavities used with cavities that are made of superconductors; this could potentially increase the haloscope's sensitivity to axions. As part of this project, I spent about 4 years of my graduate career stationed at Lawrence Livermore National Laboratory (LLNL), working under Dr. Gianpaolo Carosi, testing superconducting materials, developing cavity measurement techniques, and designing prototype cavity experiments. This work has culminated in the first axion search that uses a superconducting tuning rod within the resonant cavity; This is an upgrade to an in-situ R&D testbed cavity system known as ADMX Sidecar, which runs in parallel with the main ADMX cavity.

The structure of this dissertation is as follows: Chapter 2 will formally introduce the axion

theory through the solution to the strong CP problem, and ultimately how the axion would act as the dark matter particle within our universe. Chapter 3 will introduce the experimental detector concept, the haloscope. Chapter 4 gives the specifics of the main ADMX haloscope and experimental site. Chapter 5 describes the analysis techniques that we use to report axion detection events or lack thereof. Chapter 6 is a detailed description of resonant cavities and the techniques used to characterize them. Chapter 7 is a formal introduction to the theory of superconductors. Chapter 8 reviews my work on superconducting cavities at LLNL. Finally, Chapter 9 reports performance of and the preliminary axion search analysis of the hybrid superconducting Sidecar cavity during the currently on-going ADMX data run.

Chapter 2

AXION THEORY AND COSMOLOGY

In this chapter, I will cover the mathematical foundation needed to understand why the axion may be the invisible dark matter particle. This will start with introducing the strong CP problem, the Peccei-Quinn solution, and introduction of the axion. I will then move toward the initial experimental attempts to find the axion and the subsequent implications of its characteristics. Finally, I will motivate the axion as a dark matter candidate, and describe theories of how axions would have become the dark matter, and which axions would be preferred for dark matter production.

2.1 The strong CP problem

As mentioned in the introduction, the strong CP problem originates in the quantum chromodynamics' (QCD) prediction of CP-violation. The minimal form of the QCD Lagrangian can be written:

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

where the sum of f is over the 6 flavors of quarks, up (u), down (d), charm (c), strange (s), top (t) and bottom (b), q_f is the quark field of the given flavor, γ^μ are the gamma matrices, D^μ is the covariant derivative, m_f is the given quark mass, and $G_{\mu\nu}^a$ is the gluon field strength tensor.

This Lagrangian poses a problem because it is symmetric under $U(1)_A$ global axial rotations, which would imply the existence of hadron parity doublets, meaning there would be another form of the proton with opposite parity for instance; this particle has not been observed in nature. One can propose that this symmetry is broken in the early universe, but that implies the existence of a pseudo-goldstone boson similar to η , but with pion mass or-

der; no particle like this has been discovered either, when it should have already appeared in accelerator experiments. This means the $U(1)_A$ symmetry must be broken by other means; it can be avoided if it is anomalously broken by an axial current. This is what the Adler-Bell-Jackiw (ABJ) anomaly does, originally introduced to explain the observed decay of pions. It, however, introduces the CP-violating term to the QCD Lagrangian:

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

The $\bar{\theta}$ term can be thought of as the strength parameter of the CP-violation, but it is also made up of two sub-terms, one that describes the QCD vacuum, and the other based on the quark masses; $\bar{\theta} = \theta + \text{argdet}(M)$, where θ describes the vacuum violation, and M is the quark mass matrix. This could also be written in terms of the electric and magnetic components of the gluon field:

$$\mathcal{L}_{\bar{\theta}} = \frac{\bar{\theta}}{16\pi^2} E^a \cdot B^a \quad (2.3)$$

Taking a step back, this term points to the non-abelian nature of the vacuum gluon field, and indicates that it contains degenerate, yet topologically distinct vacuum states [92]. Each topologically distinct state can be labeled by an integer n , a topological winding number. The actual QCD vacuum then can be described as a superposition of these states:

$$|\Theta\rangle = \sum_n e^{-in\theta} |n\rangle \quad (2.4)$$

In this way, there are an infinite number of vacuum states in QCD, and the real vacuum is continuously 'quantum-tunneling' through the energy barriers between the various vacuum states. θ is an arbitrary phase-offset that describes the strength of the quantum-tunneling between vacuum states; hence it is an angle between 0 and 2π .

There is no reason why θ or $\text{argdet}(M)$ should be zero or should precisely cancel each other out; one expects the value to be order 1 as they both vary from 0 to 2π . Yet if it is order 1, this would imply several observable effects, such as the existence of electric dipole moments in neutral particles with non-degenerate ground states, such as the neutron. This is outlined in Figure 2.1, where a time-reversal operation results in the dipole moments being

anti-aligned, a clear sign of CP-violation. However, the neutron EDM has been measured extensively over decades only to further constrain it closer to a zero value, with a currently accepted value of $d_n = (0.0 \pm 1.1_{stat} \pm 0.2_{sys} \times 10^{-26} e \text{ cm})$, which further constrains $\bar{\theta} < 10^{-10}$ [1].




Figure 2.1: A diagram of the neutron with its charge distribution and spin resulting in aligned electric and magnetic dipole moments. Applying a time-reversal symmetry (T), preserves the charge distribution but reverses the spin, causing the magnetic moment to reverse directions as well; By CPT theorem, this is the equivalent of a CP operation. Under just a parity operation, the charge distribution can be thought of as flipping with the respect to the spin axis, resulting in the electric dipole changing direction; a subsequent (C) operation will lead to the same result as the time reversal operation. One can see in the end, the neutron now has electric and magnetic dipole moments that are anti-aligned, meaning CP symmetry is not preserved in this system. Credit: public domain work from Andreas Knecht.

The heart of the strong CP problem is why this number $\bar{\theta}$ would be so close to zero. There is no reason θ , originating from strong force vacuum, and $argdet(M)$, originating from the weak interaction, should cancel each other out as they are unrelated quantities. They seem to be 'fine tuned' to the precise value we see in nature. One way to solve this

quickly is to introduce a massless quark, which would cause the $\text{argdet}(M)$ term to become undefined, and therefore make $\bar{\theta}$ not observable; such a quark has not been discovered, and has other conflicting implications for lattice QCD. There are also other proposed solutions that can make θ and $\text{argdet}(M)$ both go to zero [85, 18], but this is beyond the scope of this dissertation. We will focus on the Peccei-Quinn solution to the strong CP problem, because it is a minimal extension to the standard model and produces the axions we hope to find.

2.2 Peccei-Quinn solution

In 1977, Roberto Peccei and Helen Quinn published the first paper on what came to be known as the Peccei-Quinn mechanism of solution to the strong CP problem [90]. It is popular because it is quite simple and sensible, and most of all, produces the axion as a by-product for physicists to look for. Their solution was to add a new global $U(1)$ chiral symmetry known as the Peccei-Quinn symmetry $U(1)_{PQ}$ that is broken in the early universe. This would introduce a new complex, pseudoscalar field, ϕ . $\bar{\theta}$ is promoted from a constant to a variable in this theory equal to $\text{arg}(\phi)$. For temperatures greater than Λ_{QCD} , this symmetry takes on a classic "wine bottle" potential:

$$V(\phi) = \mu^2(|\phi|^2 - f_{PQ}^2/2)^2 \quad (2.5)$$

where μ is a dimensionless parameter and f_{PQ} is the symmetry-breaking energy scale above which CP is not conserved. This potential is shown in Figure 2.2. Above the symmetry breaking scale, f_{PQ} , the vacuum expectation of ϕ is zero, and all values of $\bar{\theta}$ are equally probable, and strong CP violation exists. This means the neutron would have had an EDM in the early universe. However, as the universe cools below this energy, the field is forced to the minima of the potential, which is $\phi = f_{PQ}/\sqrt{2}$. In a classical potential, $\bar{\theta}$ could continue to take on a variety of values around the minima 'rim' of the wine bottle, but when $T \propto \Lambda_{QCD}$, quantum instanton effects tip the wine bottle potential off axis (an anomaly not described here), forcing $\bar{\theta}$ to a global minima. A new PQ Lagrangian term arises:

$$\mathcal{L}_{PQ} = \xi \frac{\phi}{f_{PQ}} \frac{g_s^2}{32\pi^2} \tilde{G}_b^{\mu\nu} G_{b\mu\nu} + \mathcal{L}_{interactions} \quad (2.6)$$

where ξ is a model dependent parameter, and g_s is the strong force coupling constant.




Figure 2.2: The "wine-bottle" potential of the field ϕ introduced by Peccei-Quinn. Credit: CERN

This additional CP-violating term can be grouped with the others, and $\bar{\theta}$ can be redefined as:

$$\bar{\theta} = \theta + \text{argdet}(M) - \xi \frac{\phi}{f_{PQ}} \quad (2.7)$$

One can then see that $\bar{\theta}$ naturally relaxes to zero when the potential is minimized at:

$$\langle \phi \rangle = \frac{f_{PQ}}{\xi} (\theta + \text{argdet}(M)) \quad (2.8)$$

They had found a mechanism that naturally gives $\bar{\theta}$ a zero value! The strong CP problem is solved, but what are the side-effects? Within a year after the Peccei-Quinn paper being published, Steven Weinberg [116] and Frank Wilczek [119] had published papers showing that the PQ mechanism postulates the existence of a new pseudo-Nambu-Goldstone boson that, because of the anomalous axial symmetry breaking in the potential, would have non-zero mass. Weinberg called this particle the higglet, because of the similarities between the

PQ mechanism and the Higgs mechanism. Frank Wilczek dubbed it the 'axion', after the laundry detergent, because it had cleaned up the strong CP problem; the name also works because it arises from the axial symmetry breaking in the wine-bottle potential. Wilczek's name caught on, and the hunt for the axion began!

2.3 The original and invisible axions

The original axion, sometimes called the Peccei-Quinn-Weinberg-Wilczek (PQWW) axion, had an associated symmetry breaking scale with the electroweak scale, that is $f_{PQ} \approx 10^2 - 10^3 GeV$. This would place the mass of the axion around 10-100 keV. These axions should have been detected in a variety of accelerator experiments via several particle channels: $a \rightarrow \gamma\gamma, a \rightarrow e^-e^+, J/\Psi \rightarrow \gamma a, K^+ \rightarrow \pi^+ a, \Upsilon \rightarrow \gamma a$, and $N^* \rightarrow Na$. When they didn't appear in experiments, this ruled out the PQWW axion, and it pushed the f_{PQ} to higher energy values [49, 108]. The result of this was a much lighter axion given by the equation:

$$m_a = \frac{(m_u m_d)^{1/2}}{m_u + m_d} \frac{f_\pi}{f_{PQ}/N} m_\pi \approx 0.62 eV \frac{10^7 GeV}{f_{PQ}/N} \quad (2.9)$$

Where m_u and m_d are the up and down quark masses respectively, f_π is the pion decay constant, N is the PQ color anomaly (related to the ξ term earlier), and m_π is the pion mass. This can be even more simply written for the experimenter:

$$m_a \approx 5.691 \mu eV \frac{10^{12} GeV}{f_{PQ}} \quad (2.10)$$

Because the axion would now have an extremely light mass, trying to detect and measure the mass of axions via meson decay channels ($K^+ \rightarrow \pi^+ a, \Upsilon \rightarrow \gamma a$) or nuclear de-excitations ($N^* \rightarrow Na$) would be a similar to, but even harder than, measuring the neutrino mass; the axion would be a tiny drop of energy and momentum in comparison with its partner in the interaction, being probably a far lighter particle than a neutrino. The $a \rightarrow e^-e^+$ channel would only occur in some axion models, as discussed later in this section, so it is usually not of interest. The decay channel, $a \rightarrow \gamma\gamma$, is of interest, pictured in Figure 2.3, because it would be directly correlated with the axion energy, and would most likely produce measurable microwave/RF photons based on its predicted mass.




Figure 2.3: Feynman diagrams of axion coupling to photons and fermions.

The interaction Lagrangian describing this decay of the axion into two photons is given by,

$$\mathcal{L}_{a\gamma\gamma} = -g_{a\gamma\gamma}\phi_a\vec{E}\cdot\vec{B} \quad (2.11)$$

where ϕ_a is the axion field, \vec{E} and \vec{B} are the electric and magnetic fields, and $g_{a\gamma\gamma}$ is the coupling constant of the decay. This constant is given by

$$g_{a\gamma\gamma} = \frac{\alpha}{2\pi f_{PQ}}\left(\frac{E}{N} - 1.92\right) = \alpha\frac{g_\gamma}{\pi f_{PQ}} \quad (2.12)$$

where α is the fine structure constant, g_γ is a model-dependent coupling, and E and N are the electromagnetic and color anomalies of the axial symmetry. The model dependent parameter, g_γ , arises from the fact that there is a fermionic loop in this interaction; which fermions have interactions with the axion in this loop are dependent on the model.

They are two benchmark "invisible" axion models referred to as the Kim-Shifman-Vainshtein-Zakharov (KSVZ) model [70, 104], and the Dine-Fischler-Srednicki-Zhitnitsky (DFSZ) model [44, 124]. The KSVZ model only permits the axion to couple to heavy

quarks, whereas the DFSZ model allows coupling to all fermions, quarks and leptons. The KSVZ model has $g_\gamma = -0.97$, whereas the DFSZ has coupling, $g_\gamma = +0.36$. DFSZ axions can also be easily implemented into grand unified theories that contain the $SU(5)$ group.

2.4 Axions as dark matter

The invisible axion quickly became recognized as a potential dark matter candidate. For the purposes of keeping this dissertation to a reasonable length, I will assume the reader is familiar with the dark matter problem and its origins, and has accepted the numerous pieces of evidence for its existence; galactic rotation curves, gravitational lensing measurements, and signatures in the cosmic microwave background to name a few. I will point to two previous ADMX dissertations that go over the supporting evidence well [28, 47]. What is important for the purposes of ADMX is what we know about dark matter's characteristics based on the supporting evidence. It must be:

1. Very stable
2. Very weakly coupled to photons
3. Non-baryonic
4. Mostly cold or non-relativistic
5. Numerous enough to make up 85% of the mass in the universe

From what we know about the invisible axion already, it matches the first three criteria pretty well; a non-baryonic particle with an extremely light mass that would very weakly couple to photons and have few other decay channels once created. The real question is how do enough axions get produced in the early universe and cool off in time to make up the dark matter.

There are three primary ways in which axions could have been produced in the early universe: thermally in quark and nucleon interactions when the universe was a hot, dense

plasma, string or domain wall decays, and a process called the "vacuum realignment" mechanism. In the case of thermal axions, the mass would need to be greater than 1 meV for its cross-sections to be large enough to make a significant amount of matter, and the axions produced in these interactions would be very hot, which is not in agreement with the structure of dark matter we observe. Furthermore, measurements of red giant stars and supernova rule out most axions above 1 meV, as is discussed later in the chapter. The axions produced in string or domain wall decay are outside the scope of this dissertation, but also have several constraints that indicate they are not the dominant production mechanism for dark matter. In the case of domain wall decay, over production of axions is very likely unless symmetry-breaking parameters are finely tuned to get the desired axion energy density [64]. String decays have similar dependencies to become a dominant dark matter production mechanism [59]. The vacuum realignment (VR) mechanism is, however, a very compelling process for the creation of dark matter.

This mechanism originates from the fact that before PQ symmetry is broken at its energy scale, θ can take on any value from 0 to 2π at any point in space. At the point where it cools below the f_{PQ} scale, there will be some initial mis-alignment value, θ_i , and the value will begin to drop down the wine-bottle potential towards its minimum, following the equation of a damped harmonic oscillator:

$$\ddot{\theta} + 3H(t)\dot{\theta} + m_a^2 \sin(\bar{\theta}) = 0 \quad (2.13)$$

where $H(t)$ is the Hubble parameter and m_a is the axion mass. For $H < \theta$, the oscillations will be under-damped and lose energy via "Hubble Friction." These oscillations correspond to a zero-momentum, pressureless condensate of axions that would be cold immediately. This would allow them to then fall into the gravitational wells produced by the anisotropies in the hot baryonic matter, seeding the large scale structure we see today as the dark matter. The fractional cosmic mass density from this VR mechanism is:

$$\Omega_A^{VR} h^2 = 0.12 \left(\frac{6 \mu eV}{m_A} \right)^{1.165} F \bar{\theta}_i^2 \quad (2.14)$$

where h is the Hubble expansion parameter, $\bar{\theta}_i$ is the initial misalignment angle relative to the CP-conserving value between $-\pi$ and π , and F is a factor accounting for the anharmonicities in the axion potential [24, 14, 15]. What is important to note here is that the leading value, if the rest is kept to $\mathcal{O}(1)$, is at the energy density value expected for cold dark matter, $\Omega_{cdm}h^2 = 0.12$, from cosmological measurements. This means if $F\bar{\theta}_i$ was kept to $\mathcal{O}(1)$, then an axion of order a μeV would account for 100% of the dark matter observed in the universe.

This is the point where one must consider the cosmological event known as inflation, and its effect on this production process. Inflation is another well accounted for cosmological event in the early universe, whereby the universe expanded rapidly, 'faster' than the speed of light, expanding its volume by at least a factor of 10^{78} [58]. This event resulted in separate casual volumes within the universe, and our casual volume (a light-sphere) is often called our visible universe. This means if PQ symmetry is broken after inflation, θ_i would take on different values in different casual volumes, and would average out to the root-mean-square average of all the volumes, giving a value of $\bar{\theta} = \pi/\sqrt{3}$. This simplifies the fractional cosmic mass density of the axions:

$$\Omega_A^{VR}h^2 = 0.12\left(\frac{30\mu eV}{m_A}\right)^{1.165} \quad (2.15)$$

One can then see a plausible range of axion masses is around $25\mu eV < m_a < 1000\mu eV$ for a post-inflationary scenario.

In the pre-inflationary scenario, the symmetry would break first at a random value of θ_i , and from that tiny volume expand rapidly with inflation, making θ_i the same random value for all casual volumes. Nonetheless, it is still constrained to be $\mathcal{O}(1)$, such that axions much lighter than μeV are produced too abundantly in the early universe, making too much dark matter, and causing the universe to over-close. There is one argument, using the pre-inflationary scenario, where the universe randomly gets a near zero θ_i , which allows for much lighter axions; the anthropic argument is that only these small θ_i scenarios allow for life to form, and therefore we must be in a small θ_i universe. Whether or not one believes this argument, there are other reasons to want ultra light axions, as they can help explain some discrepancies astronomers see in the large-scale structure of dark matter, by use of

'fuzzy' dark matter, where the axion has de Broglie wavelengths on the order of light years. There are also several grand unified theories that are preferential to it. There are several axion experiments targeting this ultralight mass range as will be discussed at the end of this chapter.

2.5 Axion energy spectrum lineshapes

In this section I will cover the velocity distribution axions take on after they are generated by the VR mechanism. Since these axions are generated initially as a pressureless zero-momentum gas, it is fair to assume that the axions around us fell into the gravitational well that became our galaxy, and became thermalized in the process. This simple model is called the "isothermal sphere" model and would take on the classical Maxwell-Boltzmann thermal velocity distribution. The axion line shape based on a Maxwell-Boltzmann distribution, following the standard halo model would be [114]:

$$g(f) = \frac{2}{\sqrt{\pi}} \sqrt{f - f_a} \left(\frac{3}{f_a \frac{c^2}{\langle v^2 \rangle}} \right)^{3/2} e^{-\frac{3(f-f_a) \frac{c^2}{\langle v^2 \rangle}}{f_a}} \quad (2.16)$$

where f is the measured frequency, f_a is the axion rest mass frequency, and the rms velocity of the dark matter halo is $\langle v^2 \rangle \approx (270 \text{ km/s})^2$. Note that this equation averages out the orbital motion of the Earth around the Sun and rotational motion of the detector about the earth's axis; these would produce a small spectral modulation in the line shape on a scale smaller than what's needed to just detect the signal. High resolution analysis of this modulation, once an axion is found, would be highly beneficial for characterizing the local distribution of dark matter in our solar system. As it stands, this would result in an axion virial velocity of $\beta = v/c = 10^{-3}$, and a fractional line width of 10^{-6} .

The isothermal sphere model is a bit naive however, and unable to explain features of our galaxy. A primary mechanism that can produce errors in the model is galactic mergers; two galaxies that pass between each other can have a tendency to pull some dark matter on to the more massive partner. This results in a 'dark disc' on top of the isothermal sphere for some galaxies [97]. It is also possible that axions have entered our halo and not fully thermalized,

resulting in "cold flows" as they fall towards the galactic core. These flows support the formation of "caustic rings" at points where the number of flows present change value [48]. The rings can be thought as stationary points in the velocity phase space, producing a very high density region of axions, with diminishing velocity dispersion, making an extremely small energy line width of the axion photon signal locally.

Developments in galaxy formation simulations have enabled calculation of alternative velocity distributions for the axion that can account for these unique phenomenon. Ref. [75] outlines simulation work done at the University of Washington to better approximate the axion linewidth for ADMX. ADMX now reports its exclusion limits (more on this later), using both the Maxwell-Boltzmann shape, and the predicted "N-body" shape from these simulations. These simulations describe galaxies using the N-body+smooth-particle-hydrodynamics (N-body+SPH) method, in lieu of the standard halo model (SHM). The axion line shape based on these N-body simulations would be:

$$g(f) \approx \left(\frac{(f - f_a)}{m_a \kappa} \right)^\alpha e^{-\left(\frac{(f - f_a)}{m_a \kappa} \right)^\beta} \quad (2.17)$$

Again, f is the measured frequency, f_a is the axion rest frequency, and m_a is the axion rest mass. The best fit parameters were computed from simulation to be $\alpha \approx 0.36 \pm 0.13$, $\beta \approx 1.39 \pm 0.28$, and $\kappa \approx (4.7 \pm 1.9) \times 10^{-7}$ [75]. Figure 2.4 shows that the axion linewidth could be narrower than the Maxwell-Boltzmann prediction. This however doesn't effect detection very much, because one is simply looking for any peak in the beginning and looking for an energy excess in a wider shape is a more conservative strategy. Once the axion signal is found in a wider bin, it is relatively easy to narrow in on the specifics of its line shape. Throughout the rest of this dissertation, I will assume the Maxwell-Boltzmann shape unless stated otherwise.

2.6 The current axion parameter space

As of now, no axions have ever been detected, but through a variety of experimental techniques, areas of the axion parameter space have been able to be excluded, meaning an axion




Figure 2.4: N-Body Simulations from Romulus25 showing the axion frequency spectra for various Milky-Way-like galaxies (color lines), containing an axion of $m_a = 1\mu eV$. The solid black line shows a fit to the new line width model shown in Equation 2.17, whereas the dotted line shows the Maxwell-Boltzmann distribution (Equation 2.16) [75].

with those properties could not exist. As the axion theory has been outlined, one can see that there are only two open theory parameters that define the axion: f_{PQ} and g_γ . f_{PQ} determines the axion mass, m_a , by Equation 2.10, and both f_{PQ} and g_γ determine $g_{a\gamma\gamma}$ by Equation 2.12. For most experimenters dealing with measuring the power of a given axion interaction coupling, it is better to report experimental sensitivity in terms of g (y-axis) vs. m_a (x-axis), thus we will be reporting the haloscope sensitivities in axion-photon coupling, $g_{a\gamma\gamma}$, although this could also be for coupling with any fermion, g_{af} . One can see that, in the case of the KSVZ and DFSZ invisible axion theories, $g_{a\gamma\gamma} \propto m_a$, and each theory produces a predicted coupling-mass line through a plot of $g_{a\gamma\gamma}$ vs m_a based on its value of g_γ . Figure 2.5 shows such a plot, with the KSVZ and DFSZ predicted lines overlay in a yellow/orange

band; this band marks the 'discovery' range where one expects to find invisible axions. Any axions found outside of this yellow band would be considered 'Axion-Like Particles' or ALPs, and would not solve the strong CP problem. The rest of the plot contains exclusion limits from three major categories: astrophysical measurements where axions aren't the dark matter, astrophysical measurements where they are the dark matter, and dedicated laboratory searches. One can see that the upper end of the mass range has been well excluded by astronomical and cosmological measurements, making KSVZ and DFSZ axions impossible above 100 meV; as we already discussed, however, these higher mass axions from 1-100 meV are more likely made thermally, and therefore 'hot', so they would most likely not be the dark matter. Another important contribution to this plot, with a lot of real estate in the axion parameter space, is the CAST experiment in dark red; It is an axion helioscope that looks for solar axions, and is able to set limits over a wide range of masses, but only to a limited value of $g_{a\gamma\gamma}$. Finally, one should notice the series of red rectangles reaching down to the KSVZ/DFSZ yellow region starting at about $m_a = 1\mu eV$ that start to get smaller and sparser above $10\mu eV$; These are the haloscope experiments including CAPP, HAYSTAC, ORGAN, RBF+UF, and of course ADMX! These are dedicated instruments designed to reach this discovery region, and the goal for haloscopes is to fully exclude the discovery region up to a 1 meV axion mass. Figure 2.6 highlights this haloscope discovery region, from $1\mu eV$ to $0.1meV$, more clearly and shows all the various experiments. I will briefly outline how each of these techniques work in the following sections, with the exception of the haloscope as that will be covered in the next chapter in detail.

2.6.1 Astrophysical constraints

If axions exist, depending on their mass and coupling, they would be a source of energy dissipation in astrophysical objects. Similar to a neutrino, if axions couple strongly enough to regular matter, they would be produced enough in stars and supernovae to affect their energy loss rates and lifetimes. However, at the same time, if they are too strongly coupled, they would be trapped within the star endlessly colliding off the dense matter within. By

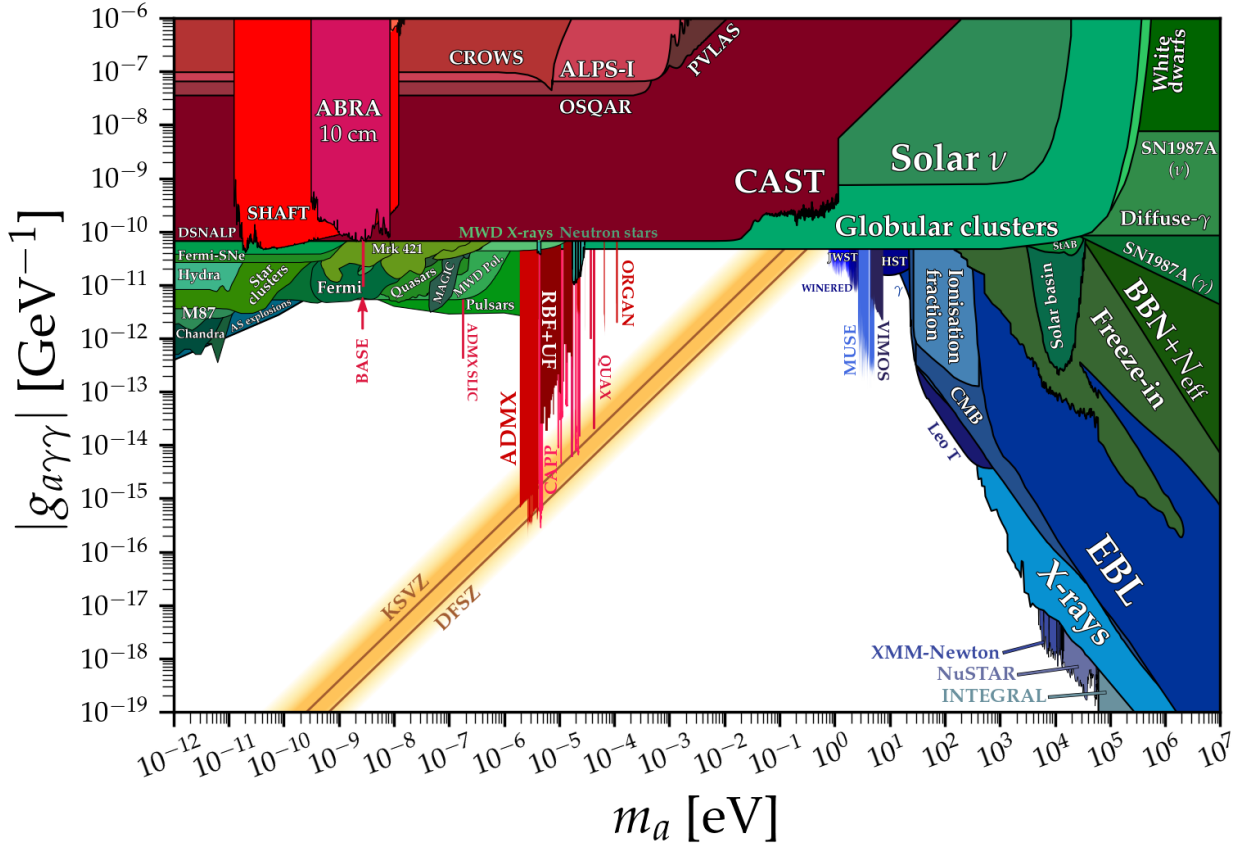


Figure 2.5: The axion-photon coupling exclusion limit plot for a variety of experimental techniques. The KSVZ and DFSZ theory lines are plotted in orange, with a yellow band around them marking the "discovery" range of where the invisible axion is expected to be. Laboratory/dedicated axion searches are highlighted in red colors, with ADMX and other haloscopes being the most sensitive and able to exclude axion within the yellow "discovery" range. Astrophysical constraints using measurements of various stars, galaxies etc. are highlighted in green; These measurements do not assume the axion is dark matter. More astrophysical measurements such as constraints set by the CMB are highlighted in blue; these assume the axion is the entirety of dark matter.

looking at the lifetime of stars and their energy budgets, constraints on the axions mass and coupling can be made.

In the case of our sun, which is only burning hydrogen currently, axions would be primarily generated through a Primakoff process, $\gamma + Ze \rightarrow a + Ze$, The additional energy loss from this extra channel would cause the sun to contract faster, in turn raising its temperature,




Figure 2.6: The axion parameter space centered around the haloscope exclusion region with various haloscope experiments, including ADMX, highlighted in red. The yellow region highlights the invisible axion discovery region, centered around the predicted DFSZ and KSVZ mass-coupling lines.

and thereby increasing the fusion rate in its core. This increase in fusion would result in increased neutrino production, which can be measured by experiments on Earth. The Sudbury Neutrino Observatory (SNO) was able to set limits in this way, citing $g_{a\gamma\gamma} < 7 \times 10^{-10} \text{GeV}^{-1}$ [55].

A stronger limit than the Solar ν is the constraints set by globular clusters in Figure 2.5. Globular clusters are not full-form galaxies, but spheroidal conglomerations of stars caused

by gravity. Most importantly, they do not have active, appreciable star formation taking place within them. This means most stars within them are all nearly the same age, and vary by their initial masses. This means the population sizes of the various phases of stars within them is proportional to the time the star is spending in such a phase. Because axions would have the effect of shortening the helium burning phase of stellar evolution, the horizontal branch stars (HB), one can compare the populations of these stars to the hydrogen-burning phase stars, the Red Giant Branch (RGB), within globular clusters, and set limits on axions accordingly [43]. This has ruled axions out over a very wide mass range to a coupling of $g_{a\gamma\gamma} < 6.6 \times 10^{-11} GeV^{-1}$ [115].

Limits based on supernovae observations can also be set. If axions exist, the primary means of axionic cooling would be from nucleon-nucleon axionic brehmsstrahlung ($NN \rightarrow aNN$). Its important to note however, that axion masses $< 1 meV$ would have not taken away enough energy to be observable, and axions with mass $> 2 eV$ would not have escaped the supernova, giving their energy back to the event. But if the axion has mass in the intermediate range, this would have resulted in a significantly shortened neutrino burst during the SN1987A event, which was not observed [33, 113].

There are many other astrophysical constraints that I do not have time to discuss here; neutron star lifetimes, black-hole super radiance, and optical telescope measurements for instance [105, 11, 57]. If one assumes axions are the majority of dark matter, constraints on the CMB, Active Galactic Nuclei (AGN), and other phenomena can also be made. Exclusion limits from the James Webb Telescope have even been made recently [68]. In general, these other astrophysical methods do not exclude axions in the KSVZ and DFSZ coupling ranges, either by not being sensitive enough to the low couplings, or by looking at axion masses much less than $1 \mu eV$.

2.6.2 Laboratory Searches

Axion helioscopes are the high-energy sibling to the axion haloscope [106]. They aim to convert solar axions originally produced in the sun back into X-ray photons via a large

magnet. A large-bore, long, high-field magnet is pointed at the sun where solar axions would enter and along its length decay into photons. An X-ray optics package focuses these potential photons onto a detector. Because of this setup, axions can be a variety of masses below a limiting upper mass which is constrained by other astrophysical backgrounds. The most famous of these axion telescopes is the CERN Axion Solar Telescope (CAST), which has set a limit of $g_{a\gamma\gamma} < 6.6 \times 10^{-11} GeV^{-1}$ for $m_a < 0.02 eV$ [9]. The physics world is greatly anticipating the follow-up experiment, the International Axion Observatory (IAXO), which has a target sensitivity of $g_{a\gamma\gamma} < 5 \times 10^{-12} GeV^{-1}$.

Another major category of laboratory searches are "light-shining-through-walls" experiments. These experiments contain two Fabry-Perot cavities; an axion generator cavity and receiver. The generator cavity is embedded in a dipole magnetic field and contains a laser that is pumping photons within it. The laser photons are stimulated to convert to axions, with an enhancement factor based on the cavity optical path length. Once axions are generated, they could travel through an optical barrier (the "wall") and end up in the receiver cavity, which is also embedded in a magnet. The axion in the receiver cavity then decays back into a photon that would be detected by a CCD detector. These experiments include ALPS, ALPS-II, CROWS, PVLAS, and OSQAR. OSQAR has set the most stringent limit at $g_{a\gamma\gamma} < 3.5 \times 10^{-8} GeV^{-1}$ for $m_a < 0.3 meV$ [16].

Another technique gaining popularity looks for axions in the low mass range via LC-resonant circuits, such as the ADMX-SLIC, ABRACADABRA and SHAFT experiments [56, 102]. They consist of a toroidal magnet that causes a very tiny oscillating current within the loop of the toroid due to axions; a superconducting sheath acts as the wire for this current to travel on. This oscillating current would produce a magnetic flux in the center of the toroid, which could be detected by an ultra-sensitive magnetometer. This inductor at the center of the toroid is usually a tunable SQUID sensor. These can be thought of as haloscopes that simply replace the resonant cavity with lumped circuit elements. ADMX-SLIC, sited at the University of Florida, uses this technique and was able to exclude axions with coupling $g_{a\gamma\gamma} < 10^{-12} GeV^{-1}$ in the $174.98 - 175.19 neV$ and $177.34 - 177.38 neV$ mass ranges [39].

ABRACADABRA was able to recently set a new limit of $g_{a\gamma\gamma} < 3.2 \times 10^{-11} \text{GeV}^{-1}$ for axions in the $0.41 - 8.27 \text{ neV}$ mass range [102]. The DM-radio collaboration will soon push the limits of this technique, with construction of DM Radio 50 L underway at Stanford, and DM-radio m^3 in the design phase.

As you can see from Figure 2.5, the through-line for all these other techniques, is none of them have been able to exclude significant parts of the DFSZ and KSVZ region except the haloscopes; I will cover how these work in the next chapter.

Chapter 3

OVERVIEW OF HALOSCOPES

This chapter will cover what an axion haloscope is, its main components, and how these components affect the sensitivity of the detector system. The concept of an axion haloscope was first proposed by Pierre Sikivie in 1983 [106, 107]. Invisible axions are made visible by exploiting the coupling of axions to the electromagnetic field with a large static magnetic field. In this arrangement, the magnetic field will act as a virtual photon, and axions from the galactic halo will convert to photons within the high static magnetic field via the two photon coupling illustrated in Figure 2.3 (hence the name 'halo'-scope). The experimental apparatus must then have a way of detecting the resultant photons and distinguish them from any background sources; this is done with a resonant cavity, low-noise readout, and cooling system.

3.1 Components

The 4 main components of any axion haloscope as alluded to above are:

1. High field strength magnet to stimulate axion to photon conversion
2. An embedded, frequency tunable, resonant cavity or cavities within that magnet's field to collect axion photon power
3. A readout system for the resonant cavity system to detect resultant photon power
4. A cooling system to minimize external photon noise (primarily thermal photons) within the cavity and readout systems




Figure 3.1: A toy model of an axion haloscope’s main components and detection process [28].

These components are pictured in Figure 3.1 and will be used as a visual guide for this section. The specifics of these components build on one another and dictate what kind of axion masses and couplings can be searched for. From a fiscal standpoint, the magnet, because it is usually the most expensive equipment of these components, is the starting point for what kind of axion search is possible within a given budget.

3.1.1 Magnet

The magnet and the static field it produces creates the potential detection volume for axions as pictured by the field lines in Figure 3.1. A higher peak field strength will stimulate more axion to photon production, specifically proportional to the interaction Lagrangian as shown in Equation 2.11. Therefore the axion signal power is proportional to the magnet field strength, specifically by B^2 as will become clear in the axion power equation section.

Additionally, the volume of the external field will cast a larger detection footprint. It is important to note, however, because axions have a de Broglie wavelength much larger than the length scale of most magnet volumes, a classical particle picture isn't quite accurate; more volume isn't casting a larger target to 'catch' more axions like in many WIMP experiments, rather it is covering a larger portion of the axion field and that will drive more power into the embedded resonant cavity system. Simply said, a magnet's figure of merit for a search would be $\oint \vec{B} dV$ where B is the magnet field strength and V is the field volume.

Superconducting solenoid magnets are the standard for generating the highest static magnetic fields in any laboratory, and thus have dictated the geometry and design of axion haloscope experiments so far. The detection volume in this case is the internal cylindrical bore dimensions, which subsequently dictates the size of a resonant cavity system that can be inserted within it. These magnets are commonly wound with NbTi or NbSn superconducting wires, with the latter being used for higher strength magnets that exceed the critical field of NbTi (≈ 15 T). NbSn magnets, because of their winding materials' critical fields and flux creep, are limited to $\approx 10 - 20$ T depending on construction details (these superconducting physics terms are discussed in Chapter 7).

These superconducting magnets commonly make use of a closed loop design with a persistent switch, where the power supply is only connected when ramping the magnet current. This saves on electricity costs when operating, leaving cooling of the magnet the principle expense. These magnets are cooled with a pumped liquid helium system, and if operating for a long duration like axion searches do, often have an associated helium liquefaction plant and recovery system.

This limit can be exceeded by the use of a Bitter electromagnet design, with the record being 37.5 T at Radboud University High Magnetic Field laboratory in the Netherlands [118]. The strongest continuous, static magnet field ever generated in a laboratory was of a hybrid design: a Bitter magnet inside a superconducting solenoid. This was constructed at the National High Magnetic field laboratory (NHMFL) in Tallahassee, FL, USA. The Bitter magnet generated 33.5 T and the outer solenoid generated the remaining 11.5 T for a total

peak field of 45 T [84]. Although incredible feats of engineering, these magnets do not make sense for current axion searches because of their prohibitive energy and cooling expenses. In the case of NHMFL’s record magnet, it consumed around 33 megawatts of power for its liquid helium system and surrounding infrastructure, which at the time cost approximately \$1500 per hour; An axion search must operate for months at a time, making use of these ultra-high field magnets prohibitively expensive at current budgets. Because of the requirement of a large bore, most current and planned axion searches stick to conventional NbTi or NbSn magnets with very few experiments having magnets over 10 T.

3.1.2 Cavity

Similar to the magnet, there is no axion signal to be detected without a cavity resonator to be driven by the axion field. As previously mentioned the resonant conversion of halo axions to photons is proportional to $\vec{E} \cdot \vec{B}$, see Equation 2.11, where \vec{E} , is the electric field of the cavity mode, and \vec{B} is the external magnetic field. This means for a solenoid magnet with a cylindrical cavity, having an axial electric field for as much of the internal cavity volume as possible, maximizing the $\vec{E} \cdot \vec{B}$. The TM_{010} resonant cavity mode does this, with its electric field entirely along the cylindrical cavities’ axial direction (see Figure 3.1). The specifics of resonant cavity modes will be discussed in detail during Chapter 6. In order for this cavity mode resonance to be excited, it would have to be tuned precisely to the resultant photons frequency:

$$f_a \approx \frac{m_a c^2}{h} \quad (3.1)$$

For axion rest masses of 10^{-5} eV, this corresponds to a frequency of 2.4 GHz or wavelength of 12 cm. As mentioned in Chapter 2, the local cold dark matter density is approximately 0.45 GeV/cc, therefore, if axions make up the entirety of the dark matter, the number density will be over 10^{13} cm^{-3} . Also discussed earlier, these axions will be cold, with $\beta \approx 10^{-3} c$, corresponding to a de Broglie wavelength of 1000 m (its also of note that this is why the approximation in Equation 3.1 works). This means that on the length scale of resonant

cavity, ≈ 12 cm, an axion halo will behave like a classical field, coherent over a much larger scale than the cavity, $100 - 1000$ m. Thus, one can view the axion field as a coherent classical source of radiation that is driving the resonant cavity mode. Because of the very small velocity dispersion of the axion field, the fractional width of the excitation peak will be around 10^{-6} within the cavity power spectrum. Therefore a resonator with a high quality factor, as close to 10^{+6} as possible, is ideal for initial detection (more on this in Chapter 6).

This is all to say that the aim of any haloscope experiment is to detect this resonant peak by monitoring the power spectrum of the cavity near the TM_{010} mode or other search mode of interest (see the exaggerated resonant peak pictured in Figure 3.1). To do this, antennae are carefully inserted into the cavity space to sample this power spectrum precisely. Because one does not know the axion rest mass precisely, this resonant mode frequency must also be adjustable, or "tunable" in order to line it up with a range of potential axion photon frequencies. This is why tuning rods are used as pictured in Figure 3.1.

There are several factors that go into the design of an optimum cavity resonator. One wants to maximize the volume of axion field driving the resonant system, therefore maximizing V_{cav} within the magnet bore volume is ideal. However, this effects the resonant frequency of the TM_{010} mode, f_{010}^{TM} , which must be set according to the desired axion search frequency range. For an empty right cylindrical cavity the analytical solution to this mode frequency is:

$$f_{010}^{TM} = \frac{x_{01}}{2\pi\sqrt{\mu\epsilon}R} \quad (3.2)$$

where μ and ϵ are the interior permeability and permittivity, x_{01} is the first root of the Bessel function $J_0(x)$, and R is the radius of the cavity (the derivation of this expression is covered in Chapter 6). Therefore, at least for an empty cavity, the radius sets the resonant frequency, and the length can be maximized to the length of the magnet bore. When tuning rods are introduced, the axial symmetry is broken, and there are no analytical solutions; one is forced to use numerical modelling methods to calculate frequency. The cavity length in this case will have a perturbative effect on the frequency. Although it will be discussed more in Chapter 6, this connection between volume and frequency can be circumvented by the

use of multi-cavity systems that maximize detection volume while still targeting a specific resonant frequency; this is especially useful for higher axion mass searches.

The factor that includes the axion-photon interaction Lagrangian, $\propto \vec{E} \cdot \vec{B}$, we refer to as the cavity form factor. It is dimensionless with a range of 0 to 1, and is maximized for the TM_{010} mode, being 0.7 for an empty cavity, and typically between 0.3-0.5 for a cavity with tuning rods. Similar to frequency, it can be calculated explicitly for an empty cavity, but only through numerical modelling for tunable cavities. Although it will be derived explicitly in the next section, I leave it here for reference:

$$C_{lmn} = \frac{(\int \vec{E}(\vec{x}, t) \cdot \vec{B}_{ext}(\vec{x}) dV)^2}{\int \vec{B}_{ext}^2 dV \int \epsilon_r \vec{E}^2 dV} \quad (3.3)$$

where $\vec{E}(\vec{x}, t)$ is the electric field of the lmn indexed cavity mode (such as 010 in the case of TM_{010}), V is over the interior cavity volume, ϵ_r is the relative permittivity of the interior cavity, and $\vec{B}_{ext}(\vec{x})$ is the externally applied magnetic field. One can see that only modes of the form TM_{0n0} will have axial electric fields that overlap with the axial magnetic field in order for form factor to be appreciably non-zero.

Finally, the quality factor of the cavity, the primary parameter of this dissertation, should be maximized in order to achieve unity with the very narrow axion line shape and maximize resonant power exchange into the mode. Although it will be discussed extensively in Chapter 6, the quality factor is a measure of how many cycles the resonator mode can hold before energy is dissipated through the cavity's interior surfaces, antennae, and other radiated losses. This definition can be written as:

$$Q = \omega_0 \frac{U}{P_d} \quad (3.4)$$

where ω_0 is the angular frequency of the mode ($2\pi f_0$), U is the energy stored in the mode, and P_d is the power dissipated via both material and radiated losses.

In summation, a cavity system must be designed to target a certain axion mass frequency range, and its figure of merit for sensitivity to axions would be $\approx C^2 V^2 Q$, where C is the

form factor, Q is the quality factor, and V is cavity system volume (the powers of two for C and V will become apparent later in this chapter).

3.1.3 Readout System and Cooling System

The readout and cooling system is designed to detect an ultra-low power resonant signal generated by the cavity over the pre-dominant thermal noise background. I combine these two systems together, because in terms of operations and analysis they are very closely connected. This is because the dominant background within the cavity as well as the readout system itself is Johnson-Nyquist thermal noise, due to the non-zero temperature of these components. Therefore, the cavity is cooled to the lowest cryogenic temperatures feasible and the lowest noise electronics available are employed to readout the very cold power spectrum. This includes low-loss superconducting lines, and quantum amplifiers to boost the power signal as it is sent to room temperature electronic components, also known as the 'warm receiver'. According to Nyquist theorem, the noise power, P_{Noise} deposited over a given bandwidth of frequency space, Δf , from a system with a total system noise temperature, T_{sys} is:

$$P_{Noise} = k_B T_{sys} \Delta f \quad (3.5)$$

where k_B is the Boltzmann constant. The system noise temperature is not only the physical temperature of the cavity, but takes into account the temperature of the electronic components along the chain. Since the cavity readout signal is passed through a series of amplifiers, each with an associated gain that multiplies the input signal, the subsequent up-stream added noise is suppressed by the gains of the down-stream components. In this arrangement, it is easier to associate noise introduced by cables and other electronic components with the amplifier stage it is contained within by a single noise temperature. For a chain of N amplifiers this translates to the total system noise temperature expression being:

$$T_{sys} = T_{phys} + T_1 + \frac{T_2}{G_1} + \frac{T_3}{G_1 G_2} + \dots + \frac{T_N}{G_1 G_2 G_3 \dots G_{N-1}} \quad (3.6)$$

Where T_{phys} is the physical temperature of the cavity, T_1, \dots, T_N are the noise temperatures of each amplifier stage including the down-stream cable component of each stage, and G_1, \dots, G_N

are the linear gains of each amplifier in the chain. This means even if a cavity is cooled to a 100 mK, a typical solid state amplifier with a noise temperature of 2 K if used as a primary amplifier ($n = 1$), will add its full noise temperature to your system. However, if it is used just a stage above an amplifier with 10 dB gain, it only adds 200 mK to the system. Attenuation of the coaxial line between the cavity and primary amplifier can be made near-zero by the use of superconducting cables. For this reason, the primary amplifier is the most critical component to any haloscope RF chain.

Most current experiments, including ADMX, make use of Superconducting Quantum Interference Device (SQUID) type amplifiers where the noise temperature scales with its physical temperature, and is only limited by the uncertainty principle where,

$$T_{SQUID} \geq \frac{hf}{k_B} \quad (3.7)$$

This is typically referred to as the standard quantum limit. For an amplifier operating at 2 GHz, this is about 100 mK.

However, this is much easier said than done, especially since these SQUIDs are sensitive to even a single magnetic flux quanta ($h/2e$), yet must sit not far from an 8 T solenoid magnet so it is still cooled by the same fridge system as the cavity. It's for this reason that another application of SQUIDS is extremely sensitive magnetometers for condensed matter physics and material science. For our purposes, we have to protect this amplifier as best as possible from magnetic overload. To do this, ADMX employs a 'bucking' magnet that is counter-wound to have zero mutual inductance with the main magnet, but cancels the main magnet's field to < 100 Gauss within its bore. Within this "bucking" magnet bore, the quantum-sensitive electronics are protected (a citadel for the SQUIDs, or 'SQUIDadel' so to speak). However, to maintain the generality of this section, the quantum amplifier package could be stored in a far-field region and cooled separately, or in some other way, to ensure optimal performance.

These SQUID amplifiers are a field of research and development in themselves, and the axion haloscope community likes to think we are driving the need for new emerging SQUID

amplifier technologies. In the case of ADMX, this includes implementing Microstrip SQUID Amplifiers (MSA) initially, Josephson Parametric Amplifiers (JPA) on the current main experiment, and most recently a Travelling Wave Parametric Amplifier (TWPA) on the ADMX pathfinder experiment, Sidecar. Increasing the operational bandwidth of these amplifiers is the main incremental improvement, which allows the haloscope to scan a larger frequency range before needing to be tuned or even replaced.

This leaves the physical temperature of the cavity, T_{phys} , in Equation 3.6 as the main driver of the system noise. In its simplest form, the cavity is already being inserted into a superconducting magnet cooled by liquid helium, which will bring the bore temperature to 4 K, but this can be improved upon. The first generation of ADMX implemented a pumped liquid helium system which could cool the cavity to 1.3 K. In its current Gen 2 configuration today, it uses a He_3/He_4 dilution refrigerator that can achieve a base temperature of 80 mK, but it typically operates in the 100 mK range. This seems to be the most popular, cost effective technology of the day driven by the demand of quantum computing research. These refrigerators can go as low as 2 mK, just not for a large thermal mass with moving components like the ADMX insert. Beyond the 2 mK limit, magnetic cooling technologies would be necessary. However, as stated above, 2 mK is already below the quantum limit for the typical axion haloscope frequency ranges, thus would only be relevant to low-mass axion searches below ≈ 42 MHz; it would lower T_{sys} , but the primary amplifier would continue to sit at the quantum limit. There are ways of circumventing the quantum thermal limit, such as quantum squeezing, but this is out of the scope of this dissertation.

In summary, a haloscopes' dominant background is thermal noise, therefore the cavity followed by a high gain, quantum-limited, primary amplifier are cooled as close to the quantum limit as possible for a low system noise temperature. Up-stream amplifiers can be significantly warmer because they will be suppressed by the gains of down-stream amplifiers.

3.2 Axion power in a haloscope

In this section, I will review an equivalent circuit derivation of the axion signal power in a haloscope, however, there are several ways to arrive at this equation, and this derivation is not at all original, but follows ref. [40, 99, 71, 123]. It can also be derived by introducing an axion source term to Maxwell's equations, which is done in ref. [72, 61, 81, 32]. What I would call a combination of these methods can be found in ref. [66, 80]. I have chosen to include this derivation here because it introduces the equivalent circuit picture and methods that will be useful in Chapter 6 for thinking about resonant cavity characterization.




Figure 3.2: An equivalent circuit for a resonant cavity being driven by axion photonic power.

The resonant cavity being driven by the axion field decaying into photons can be written as an RLC circuit with a time dependent voltage source term, as pictured in Figure 3.2. A resonant cavity stores both electrical and magnetic field energies, giving it an inherent capacitance and inductance, while also dissipating this energy through its associated surface resistances and radiation losses, acting as the resistor. The time-dependent voltage source is

the axion photon power being deposited in the cavity. This circuit has a classical Lagrangian of:

$$\mathcal{L} = \frac{1}{2}L\dot{q}^2 - \frac{1}{2C}q^2 + qV(t) \quad (3.8)$$

As well as the dissipative power for the resistor:

$$P_d = \frac{1}{2}R\dot{q}^2 \quad (3.9)$$

We can relate the inductive and capacitive energy terms in this equivalent circuit Lagrangian back to Maxwell's equations by the magnetic and electric field Poynting energy integrals for a given cavity mode:

$$\frac{1}{2}L\dot{q}^2 = \frac{\mu}{2} \int |\vec{H}|^2 dV \quad (3.10)$$

$$\frac{1}{2C}q^2 = \frac{\epsilon}{2} \int |\vec{E}|^2 dV \quad (3.11)$$

In the non-dissipative and non-driven case ($P_d = 0$, $V(t) = 0$), the Euler-Lagrange equation will simply result in simple harmonic oscillations as expected:

$$\frac{d}{dt}\left(\frac{\partial\mathcal{L}}{\partial\dot{q}}\right) - \frac{\partial\mathcal{L}}{\partial q} = 0 \quad (3.12)$$

$$L\ddot{q} + \frac{q}{C} = 0 \rightarrow \ddot{q} = -\omega_0 q, \quad \omega_0 = \frac{1}{\sqrt{LC}} \quad (3.13)$$

The Euler-Lagrange equation in the driven, dissipative case is:

$$\frac{d}{dt}\left(\frac{\partial\mathcal{L}}{\partial\dot{q}}\right) - \frac{\partial\mathcal{L}}{\partial q} + \frac{\partial P_d}{\partial\dot{q}} = 0 \quad (3.14)$$

with equation of motion:

$$L\ddot{q} + R\dot{q} + \frac{q}{C} = V(t) \quad (3.15)$$

We expect for the dissipative, driven cavity case to behave like an under-damped, continuously driven oscillation; the mode will be excited for more than a single cycle from the deposited axion energy. In this way, we can assume that the mode energy will be transferred back and forth between the electric and magnetic fields over the course of a single oscillation

cycle, making the time-averaged energy stored in the mode for the initial oscillation equal to either Equations 3.11 or 3.10:

$$U_0 = \frac{1}{2}L\dot{q}^2 = \frac{\mu}{2} \int |\vec{H}|^2 dV = \frac{1}{2C}q^2 = \frac{\epsilon}{2} \int |\vec{E}|^2 dV \quad (3.16)$$

Next, we can relate the power dissipated in the circuit picture to the power dissipated in the field equation picture:

$$P_d = \frac{1}{2}R\dot{q}^2 = \frac{R}{2} \int |\vec{H}|^2 dS \quad (3.17)$$

Combining Equations 3.4, 3.16, and 3.17 together, we can arrive at a circuit expression for quality factor in the circuit picture (note this assumes no antenna or radiation losses):

$$Q = \frac{1}{R}\sqrt{\frac{L}{C}} \quad (3.18)$$

Next we must relate the axion-photon coupling to the voltage source, $V(t)$. To do this, we take the Lagrangian term describing the axion fields coupling to the electromagnetic fields (Equation 2.11) and integrate this interaction of the cavity volume (V_{cav}). This result should be equal to the energy of the voltage source $q(t)V(t)$:

$$q(t)V(t) = -g_{a\gamma\gamma}c\epsilon_0 a(t) \int \vec{E}(\vec{x}, t) \cdot \vec{B}(\vec{x}, t) dV_{cav} \quad (3.19)$$

We can then substitute the form factor expression from the last section, Equation 3.3, as well as Equation 3.11 into the expression for form factor. Finally, I will assume a perfectly axial magnetic field, $\vec{B}(\vec{x}, t) = B_0\hat{z}$ which simplifies the expression for voltage source to:

$$V(t) = -g_{a\gamma\gamma}cB_0\sqrt{\frac{\epsilon_0 C_{lmn} V_{cav}}{C}} a(t) \quad (3.20)$$

Note that C_{lmn} is cavity form factor, C is capacitance, $V(t)$ is voltage, and V_{cav} is cavity volume. Next, sticking to the circuit picture, the power, P , deposited in the cavity due to this axion-photon conversion voltage source would be given by:

$$P = \frac{\langle V^2(t) \rangle}{R} = g_{a\gamma\gamma}^2 \epsilon_0 c^2 B_0^2 V_{cav} C_{lmn} \frac{1}{RC} \langle a^2(t) \rangle \quad (3.21)$$

We can then use the circuit picture expression for quality factor, Equation 3.18, and the circuit resonant frequency, Equation 3.13, to get rid of the circuit parameters:

$$P_{a \rightarrow \gamma} = g_{a\gamma\gamma}^2 \epsilon_0 c^2 B_0^2 V_{cav} C_{lmn} \omega_0 Q \langle a^2(t) \rangle \quad (3.22)$$

The square amplitude of the axion field can be related to the mass density of axions in the region:

$$\langle a^2(t) \rangle = \frac{\rho_a \hbar^2}{m_a^2 c} \quad (3.23)$$

This gives the expression:

$$P_{a \rightarrow \gamma} = \frac{2\pi \epsilon_0 \hbar^2 c}{m_a^2} g_{a\gamma\gamma}^2 \rho_a f_a B_0^2 V_{cav} C_{lmn} Q \quad (3.24)$$

Using the Chapter 2 theory expressions for m_a and $g_{a\gamma\gamma}$, Equations 2.10 and 2.12, m_a can be eliminated and g_γ can be used in place of $g_{a\gamma\gamma}$. Additionally, we have ignored the antenna that is sampling this cavity power; this introduces a factor of $\frac{\beta}{1+\beta}$, where β is the coupling coefficient. This coefficient will be covered in Chapter 6, and technically varies from 0 to ∞ , but is typically kept at a value of 1-2 for axion searches. Substituting in the appropriate constants and normalizing the equation to typical values seen in the laboratory, one gets an expression more useful to the experimenter:

$$P_{a \rightarrow \gamma} = 6.94 \times 10^{-23} W \left[\left(\frac{g_\gamma}{0.36} \right)^2 \left(\frac{\rho_a}{0.45 \text{ GeV/cm}^3} \right) \left(\frac{f}{1 \text{ GHz}} \right) \right] \times \left[\left(\frac{B}{7.6 \text{ T}} \right)^2 \left(\frac{V_{cav}}{136 \text{ L}} \right) \left(\frac{C_{lmn}}{0.4} \right) \left(\frac{\beta}{1+\beta} \right) \left(\frac{Q_0}{70000} \right) \right] \quad (3.25)$$

where values of the terms grouped in the left brackets are fixed by nature (the experiment is designed around the frequency range) and terms in the right brackets are parameters that can be controlled by the scientist. The value g_γ here is the DFSZ model value, and the axion mass density is assumed to be the local halo dark matter density; that is to say axions make up the entirety of dark matter. One can adjust these values based on their assumed axion model and dark matter halo composition. By far the most striking feature of this expression is the incredibly feeble size of this signal power: $10^{-23} W$ or a 10 'yotta' watts

order. Because of this, one needs to 'listen' at each frequency long enough to ensure they could detect such a weak power excess over the thermal background before moving onto the next. That is to say, there is a proper frequency scan rate to exclude axions to a certain sensitivity. Next I will combine this expression with the expected thermal background noise, obtain a signal-to-noise ratio, and derive the instantaneous frequency scan rate.

3.3 Scan rate

As illustrated in Figure 3.1, a power spectrum of the cavity near its mode resonance is taken and searches for a power excess over the noise floor that could be from an axion. Once no power excess is seen, the tuning rods are moved to change the mode resonant frequency, and a new power spectrum is taken near that new resonance; this is the 'scanning' process of the experiment.

More precisely, a time-domain digitization is taken over a period of time of the cavity antennas' field power response. This time-series data is then Fourier transformed into a frequency domain power spectrum. Power excesses above the mean noise floor can then be characterized by a signal-to-noise ratio (SNR) according to the Dicke Radiometer Equation [42]:

$$SNR = \frac{P_{a \rightarrow \gamma}}{k_B T_{sys}} \sqrt{\frac{t}{B}} \quad (3.26)$$

where $P_{a \rightarrow \gamma}$ is the axion power deposited according to Equation 3.24, T_{sys} is the system noise temperature described by Equation 3.6, t is the integration time over which the Fourier transform is taken, and B is the integrated bandwidth of the Fourier transform. That is to say, for a given noise temperature, by taking a longer time sitting on a given cavity resonance, or using a smaller frequency bandwidth for the transform, the SNR can be increased. However, spending more time on one frequency limits how much frequency space can be scanned in a given amount of time. ADMX typically tries to keep an SNR of about 3.5 across its scanning range, so I will assume that to be the default value going forward.

It is important to make a note about bandwidth B here. Effectively it is the width of our

frequency 'bins' as we take multiple transforms centered at different frequencies. Yet, the cavity resonant mode peak has an effective frequency width itself Δf_c , and so does the axion signal lineshape, Δf_a . This a good time to introduce the next definition of quality factor:

$$Q = \frac{f_0}{\Delta f} \quad (3.27)$$

where f_0 is the resonant frequency and Δf is the full-width at half maximum (FWHM) of the resonant peak (more on this in Chapter 6). As discussed in Chapter 2, the axion lineshape is dictated by its expected velocity dispersion, and a safe assumption is $\Delta f_a \leq 1 \text{ kHz}$ for an axion at 1 GHz, producing an associated axion quality factor, $Q_a \approx 10^6$. In the case of an axion search, where one's goal is to just find the axion signal, and not try to measure the line-shape structure itself, one would not want to take an integration bandwidth smaller than Δf_a , because it would not capture the full axion signal power; our purposes here are just to detect the axion and not capture the signal structure. ADMX does perform high resolution analysis of its data that uses smaller integration bandwidths to search for small signal modulations, but this is outside the scope of this proof. Additionally, any axion signal that is outside the cavity bandwidth Δf_c will not have the resonant amplification and therefore won't be detectable. In fact, having a larger bandwidth will contribute extra thermal noise from those other included frequencies to the transformed power, according to Equation 3.5. Therefore, it is clear that one wants $\Delta f_c \gtrsim B \gtrsim \Delta f_a$, and in the best case scenario, $B = \Delta f_a$, where extraneous frequency noise is excluded. Therefore I will assume for the rest of this section that $B = \Delta f_a$, is the ideal case; it is important to note that Δf_a is an estimate based on the velocity dispersion of the axion field, but it turns out this value will cancel out later in the proof. The rest of this scan rate proof is heavily inspired from Ref. [32].

Since we are searching over multiple cavity resonant frequencies, another vary important parameter is the optimal tuning step size, δf_c . One can imagine the following search scheme: the cavity starts at a center mode frequency, $f_c = f_0$, time-series data is taken for a fixed t , a single scan, then the center cavity frequency is tuned up by δf_c and the cycle is repeated,

such that the entire frequency range is covered ($B = \Delta f_a$, but we aren't performing the Fourier transforms yet, just taking time-series data). One can see in this scheme, assuming no down time between scans, that the frequency scan rate would be $\delta f_c/t$. For simplicity, assume the other major signal and noise parameters such as T_{sys} , C_{lmn} , and Q_L are frequency independent for now. The cavity is sensitive to axions between $f_c + \Delta f_c > f_a > f_c - \Delta f_c$ for a given tuning arrangement, therefore if $\delta f_c > \Delta f_c$, there will be gaps, or at least significant drops, in axion sensitivity between scans; in truth, our analysis procedure accounts for the drop off in sensitivity outside Δf_c based on the Lorentzian shape of the cavity resonance, and it is actually the bandwidth of power spectrum that determines this, but if this bandwidth is significantly larger than Δf_c , one will see proportional drop-offs in sensitivity between scans. If $\delta f_c \leq \Delta f_c$ then there will be no gaps, but some overlap even. I will therefore define $\delta f_c = \Delta f_c/F$, where $F \geq 1$. In this regime, measurements at multiple consecutive tuning steps will contribute to the SNR at each potential axion frequency f_a in the tuning range, because cavity bandwidths at each step overlap with each other to some degree by a factor of F . That is to say, there is some $K > F$ value of tuning step scans in either direction of a given f_a , which itself is within a central f_c scan, that contribute to the SNR at f_a . To complicate things just a little bit more, we know that if f_a is not on the central cavity frequency, f_c , for a given scan, the signal power will degrade slightly according to the cavity resonance peak shape, a Lorentzian. How far off the f_a is from the central frequency in a given scan is $\delta f_a = |f_c - f_a|$.

One can see that adding the SNR of multiple integrations of the same frequency would add in quadrature. This can be shown by splitting the integration time of a single scan into two scans. The original scan has $SNR = R_0$ with t_0 , and it will become two scans with times t_1 and $t_2 = t_0 - t_1$. Although the SNR will obviously go down for each individual scan, The SNR should remain R_0 when we combine these two scans again. Using the Dicke Radiometer equation, Equation 3.26, one can derive the expression $R_i = R_0 \sqrt{t_i/t_0}$ for $i = 1, 2$. Therefore, if we want to combine a set of K scans in each direction of a central scan centered at f_0 , we

can write assuming $f_a = f_0$:

$$SNR(f_0)^2 = \frac{1}{(k_B T_{sys})^2} \frac{t}{\Delta f_a} \sum_{k=-K}^K P_{sig}^2(\delta f_a = k\delta f_c) \quad (3.28)$$

Where $P_{sig}(\delta f_a = k\delta f_c)$ would represent the signal power that is off from the central peak cavity frequency by $k\delta f_c$. This signal power would follow the Lorentzian shape $P_{sig} = P_0/(1+x^2)$ where $x = \delta f_a/\Delta f_c$ and P_0 is the peak power (more on Lorentzians in Chapter 6). This expression can be converted to the scan rate by dividing out the SNR, and multiplying by $\delta f_c/t$:

$$\frac{df}{dt} = \frac{\delta f_c}{t} = \frac{1}{SNR(f_0)^2} \frac{1}{(k_B T_{sys})^2} \frac{\Delta f_c}{\Delta f_a} P_{sig}^2(\delta f_a = 0) \frac{1}{F} \sum_{k=-K}^K \left(\frac{1}{1 + (2k/F)^2} \right)^2 \quad (3.29)$$

This sum at the end of the scan rate expression can be solved numerically for a variety of values, but I will use the convenient approximation:

$$Z(F, K) = \frac{1}{F} \sum_{k=-K}^K \left(\frac{1}{1 + (2k/F)^2} \right)^2 \approx \frac{4}{5} \quad (3.30)$$

for any $F \geq 2$ and any $K \gtrsim F$, which we already required for K , and can easily be satisfied for F . Considering antenna losses again, δf_c will degrade by a factor $1 + \beta$ when the Q is loaded. Finally, if we substitute Equation 3.25 in for signal power, pull constants to the front (and remember $Q_a = f_a/\Delta f_a \approx 10^6$), and then normalize to some base experimental values, we arrive at a much more useful expression for experimenter:

$$\frac{df}{dt} = \left(\frac{3.43 \text{ GHz}}{\text{year}} \right) \left(\frac{0.36}{g_\gamma} \right)^4 \left(\frac{\rho_0}{0.45 \text{ GeV/cm}^3} \right)^2 \left(\frac{f}{1 \text{ GHz}} \right)^2 \left(\frac{3.5}{\text{SNR}} \right)^2 \left(\frac{0.2K}{T_{sys}} \right)^2 \left(\frac{B}{8 \text{ T}} \right)^4 \left(\frac{V}{100 \text{ L}} \right)^2 \left(\frac{C_{lmn}}{0.5} \right)^2 \left(\frac{\beta^2}{(1 + \beta)^3} \right) \left(\frac{Q_0}{10^5} \right) \quad (3.31)$$

As you can see from this expression, g_γ and B have the very strongest dependence on scan rate, nonetheless any improvement in Q would directly translate to speeding up the experiment by the same factor.

In the above derivation, it was assumed many parameters were frequency independent, while in reality, the experimental parameters, mainly T_{sys} , f , C_{lmn} , and Q , are changing as

one tunes the cavity. This is okay as long as they are being measured periodically in between scans (or simulated beforehand in the case of C_{lmn}), as they all are relatively frequency independent on the scale of a single scan width. Therefore, this scan rate expression can be used to determine the frequency "speed limit" to tune the cavity at any given time for a given sensitivity. These periodic measurements, calibrations, and repairs all introduce dead time to this data taking cycle, and this expression is sometimes adapted to include a dead-time fraction to account for that additional time. In actual data-taking, SNR based on the expected DFSZ power is assigned to every frequency bin in the "grand spectrum" where individual overlapping spectra have been added together. Areas of power excess with high SNR can be flagged as axion candidates, and are subsequently rescanned to search for persistent axion signals. In the absence of persistent signals, this SNR is ultimately translated into a minimum value of $g_{a\gamma\gamma}$ that can be set for the frequency bin, creating the exclusion limits shown in Figure 2.6.

In conclusion, scan rate is often the key parameter looked at when designing a haloscope experiment; based on the expected experimental values, one can estimate how fast the targeted axion frequency space can be searched at a given $g_{a\gamma\gamma}$, and therefore how long the experimental run would be. It is also a key operations parameter because it signals to the operator if tuning needs to be slowed or if it can be sped up through certain regions of the cavity tuning range. Scan rate is also one of the best figures of merit for comparing different haloscopes, because although one haloscope may have a much higher signal power, another may have a much lower noise; both are accounted for in the scan rate. The haloscope scan rate expression is also indicative of a big problem with higher frequency haloscopes, which will be discussed in Chapter 6.

Chapter 4

CURRENT ADMX EXPERIMENTAL SITE

The Axion Dark Matter eXperiment (ADMX) is the worlds' premier axion haloscope experiment, and, at the time of writing, has performed the highest sensitivity axion searches published so far. This chapter will outline the specifics of the ADMX haloscope and experimental site at the Center for Nuclear, Particle, and Astrophysics (CENPA) at the University of Washington as of the time of writing, spring 2024. The haloscope is currently taking data for our run 1D search. This "run" nomenclature will be explained in this section. It will start with a brief history of the experiment and outlines the key steps and equipment that lead to the current experimental configuration and site. The following subsections will briefly outline the technical specifications of each sub-system.

4.1 *Brief history of ADMX*

ADMX was originally conceived to be the first Sikivie haloscope with enough sensitivity to detect axions of masses and couplings predicted by theoretical invisible axion models, initially the KSVZ model, and subsequently the DFSZ model after several upgrades (see Chapter 2). It was known that this would involve multiple developmental stages, as technology improved, and be a multi-decade project. The collaboration was formed and work began on the phase "0" detector in 1993 (the same year that the writer of this dissertation was born!).

The experiment was sited initially at Lawrence Livermore National Laboratory (LLNL), in Livermore, CA, where the superconducting magnet was sourced and located. It operated with a pumped helium refrigeration system with a 1.3 K base temperature, and used 2 cryogenic HEMT amplifiers supplied by the National Radio Astronomy Observatory (NRAO). The copper resonant cavity was 50 cm in diameter and 1 meter length, using two 5 cm

diameter tuning rods. The first results of the phase "0" experiment were reported in 1998, excluding axions to the KSVZ limit between $2.9 - 3.3 \mu\text{eV}$; it was the first haloscope experiment to reach the theoretical benchmark model [60, 40]. Subsequent searches ending in 2005 expanded this range to $1.9 - 3.3 \mu\text{eV}$ with this system.

The next major upgrade phase of the experiment was to replace the HEMT amplifiers with Superconducting Quantum Interference Devices (SQUIDs) that would greatly lower the noise temperature of the primary amplifier stage. This required development of SQUID amplifiers that could operate at the higher microwave frequency range of the cavity, which was significantly above the resonant range of SQUIDs at the time. These amplifiers would be developed at UC Berkeley by John Clarke and become known as DC micro-strip SQUID amplifiers (MSAs). Because SQUIDs operate via transfer of single magnetic flux quanta, they would be overwhelmed in the magnetic field of ADMX; therefore a field-free region had to be developed to store these sensitive electronics. This was solved with the installation of a bucking magnet into the insert. This culminated with the first SQUID search from 2008-2010 setting new limits KSVZ axions between $3.36 - 3.69 \mu\text{eV}$ [13, 66].

In 2010, the 9 Ton, 8.5 T magnet was transported to CENPA at the University of Washington in Seattle, WA from LLNL in California (see Figure 4.1). Around this time, DOE and NSF were conceptualizing "Generation 2" dark matter detectors with a significant increase in sensitivities, and, although they primarily had WIMP searches in mind, ADMX was slotted in as the single experiment to search for axions. The primary ADMX "Gen-2" upgrade was the installation of the helium dilution refrigerator, which lowered the physical temperature to $\approx 150 \text{ mK}$, significantly speeding up the scan rate, and finally opening the possibility of searching for the much more elusive DFSZ model axions.

This goal was accomplished in 2016, when ADMX became the first axion haloscope to exclude DFSZ axions, with limits set in the $2.66 - 2.81 \mu\text{eV}$ mass range; this run within the collaboration is referred to as the Gen-2 run 1A [46]. This run was limited to a small mass range because of the gain bandwidth of the MSA used. The subsequent run in 2017, Run 1B, implemented a Josephson Parametric Amplifier (JPA), with significantly larger bandwidth,



Figure 4.1: The ADMX superconducting solenoid magnet arriving at UW CENPA, Seattle, WA, from LLNL, Livermore, CA.

allowing for DFSZ exclusion of axions between $2.81 - 3.31 \mu\text{eV}$ [29].

Running parallel to these main experimental searches, the ADMX "Sidecar" experiment was installed in 2016 in what was otherwise wasted free space in the insert. The Sidecar is a high frequency ($\approx 4 - 7 \text{ GHz}$), piezoelectrically tuned cavity, meant to test budding technologies for future high mass axion detectors. It isn't meant to achieve KSVZ or DFSZ sensitivity, but does occupy a mass range that hasn't entirely been excluded to either sensitivity, meaning it can produce new science limits as a by-product. The run 1A-C Sidecar papers can be found in Refs. [27, 20], and the preparations and results of the Sidecar run 1D will be the primary topic of Chapter 9.

Run 1C started in January of 2020, and required installing larger tuning rods to increase

the frequency range of the copper resonator, among other RF component exchanges to accommodate the new mass range. In March 2020, the COVID-19 pandemic caused a halt to all non-essential on-site operations. It was then observed that the system noise temperature was too high to exclude DFSZ axions in a reasonable amount of time, and extraction, inspection, and reparation of the insert was needed. Unable to do so during the pandemic, the initial report of run 1C includes both a DFSZ range and KSVZ exclusion limits between $3.3 - 4.2 \mu\text{eV}$. This entire range was subsequently excluded to DFSZ level concluding in December 2022. Run 1D began in mid-December 2023 and is currently on-going, with target mass exclusion range of $4.2 - 5.75 \mu\text{eV}$ or $1.02 - 1.39 \text{ GHz}$. To achieve this frequency range with the same diameter cavity, a single large diameter tuning rod was implemented. Increasing the rod's diameter further than this will cause significant losses in sensitivity via the volume loss, and therefore it will be the last configuration with the current single copper cavity system. The future haloscopes, ADMX Run 2 and ADMX-EFR, will use multi-cavity arrays that will be discussed at the end of Chapter 6.

4.2 *Site infrastructure*

The ADMX Gen 2 experiment resides at the end of the CENPA accelerator hall shown in Figure 4.2. At the center of the site, within a 11' concrete pit, putting its top at ground level, is the solenoid magnet. The heart of the experiment is referred to as the insert, pictured in Figure 4.3, which contains the main resonant cavity, dilution refrigerator, 'SQUIDadel' cold electronics package within its bucking coil, Sidecar cavity, and a myriad of cabling and sensors. Next to the magnet is a scaffold clean room capable of mounting the insert to its ceiling for hands-on repairs and upgrades (see Figure 4.5). At the beginning of a run, the insert is moved by crane and inserted into the magnet to cool down before ramping up the magnet (see Figure 4.4). At the conclusion of each run, the magnet is ramped down, refrigerator helium mixture collected, and the insert is extracted from the magnet to be mounted in the clean room. The rest of the site is almost entirely made up of the various helium gas handling systems, including a helium liquefier. There are two separate closed-loop

cryogenic systems: the ^4He magnet cooling system that the liquefier is providing its output to, and the $^4\text{He}/^3\text{He}$ dilution system.



Figure 4.2: The ADMX experimental site.

4.3 Magnet

The main magnet and its cryostat was manufactured by WANG NMR in Livermore, CA in 1993. It has niobium titanium (NbTi) windings forming a 1.12 m tall solenoid coil with a 60 cm inner diameter. It is actually made up of four concentric solenoids connected in series with progressively higher critical current ratings for the wire as pictured in Figure 4.6. In total, there is 99 km of copper coated NbTi wire that is wound around the stainless steel spool frame. The magnet itself weighs 6000 kg and the cryostat is about 3000 kg. Its inductance is 534 H and is rated to have a peak field strength of 8.5 T, corresponding to 249 A. In terms of stored energy, this is ≈ 16.5 MJ, about the explosive energy of sixteen and half sticks of dynamite. However it is normally operated at 7.6 T during data-taking in order to mitigate any risk of quenching (only 13 sticks of dynamite stored!). This field strength drops off quickly along the length of the coil, with only about 70% of the center field strength at either end of the coil. After 30 years, it still typically consumes 65 liquid




Figure 4.3: A cut-away diagram of the ADMX insert within the solenoid magnet.

liters of helium every day under normal operations, necessitating the helium liquefier.

Part of the reason for running below its rated strength is because its persistent switch is no longer operational increasing Ohmic heating and a risk of quench. A persistent switch usually works by having the power supply connected in parallel to a small portion of the superconducting winding, and the switch operates by a heating source to that winding sec-




Figure 4.4: Extraction of the ADMX insert from magnet after run 1B.

tion. When in "non-persistent" operation or a ramping phase of the magnet, the winding section is forced into a normal state by the heat source raising it above its critical temperature, and current flows to the power supply instead, allowing the power supply to adjust the current within the windings. When ramping operations are done, it can be set to be persistent by turning off the heat source, allowing that winding portion between the leads to



Figure 4.5: Mounting ADMX insert to clean room ceiling after run 1B.

go superconducting, creating a closed current loop. This cuts down on helium consumption by lowering Ohmic heating from the power supply leads. The ADMX magnet's switch failed at Wang NMR during commissioning and was delivered as is. LLNL engineers decided the cost-effective repair was to permanently install vapor-cooled leads instead of the removable configuration.

The second 'bucking-coil' magnet resides within the insert (see Figure 4.3) and generates




Figure 4.6: A cut-away diagram of the ADMX main magnet windings.

a field opposite the main magnet, creating a net-zero field region for the field-sensitive electronics. Figure 4.7 outlines the magnet's construction, consisting of one coil to cancel the main magnet's field and another counter-wound coil that prevents a non-zero mutual inductance between the main magnet and the first coil. This prevents any significant net force between the main magnet and bucking coil. The bore of the bucking coil has net field of less than 100 Gauss, and the sensitive electronics within can be passively shielded by mu-metal for further field reduction from there.

4.4 Cavity

The ADMX Gen-2 run 1 cavity has been the same base design through all the runs 1A-C, with only the tuning rod system changing between. It consists of a stainless steel tube that is




Figure 4.7: A sketch diagram of the bucking coil.

101 cm and has an inner diameter of 42 cm, giving it an empty volume of 136 L. Within the tube on top and bottom is a sharp, knife-edge lip, as well as a blunt lip above for applying counter-pressure against the end plates (or end-caps) seating them firmly against the knife-edge (see Fig 4.8). The top and bottom end plates are also stainless steel and contain port fixtures for the antennae and tuning rods. These components are all copper plated with oxygen free copper. Upon assembly, the end-caps are placed to rest onto the knife edge, and a stainless steel mounting ring is placed on top. Several stainless section pieces are inserted under the blunt lip and affixed to the mounting ring, then 96 bolts pass through clearance holes in the section pieces and thread into the mounting ring. As the bolts are torqued it creates a separation between the blunt lip and mounting ring, pushing the end-cap firmly against the knife edge (see Fig 4.8). This is because currents for the TM_{010} mode run across this joint, so it is very important that the knife-edge to end cap electrical connection is solid.




Figure 4.8: A diagram of the Knife-Edge clamp mechanism for ADMX run 1 cavities. A sharp knife edge is constructed off of an inset cut a couple of inches within the end of the tube. This inset will have the knife edge on the interior edge with a perpendicular face above it. The copper-plated end cap is placed on top of the knife edge, and the stainless steel compression ring inserts into the space above it within the inset of the cavity tube end. Compression screws create separation between the ring and end cap, causing the end cap to compress against the copper knife edge, and the compression ring against the opposing surface. This creates a strong seal between the copper end cap and tube, creating a solid current path for a good TM_{010} resonance.

The run 1A-C systems used two different double tuning rod systems. Having two rods allows for different tuning arrangements that can be advantageous for shifting the frequencies of mode crossings, which will be discussed further in Chapter 6. They are made of hollow OHFC copper tubes with the end-caps welded on. They run the length of the interior cavity, with a $\approx 0.020''$ clearance gap on top and bottom. This is because they must electrically disconnect from the cavity itself, in order for the mode structure to be maintained and tuned (see Chapter 6). The axles are connected off-center of the rod such that they will turn the rod center to various radii within the cavity, shifting the resonant frequency of the TM

modes. The axles extend through axle ports in the end-caps where they attach to a gearbox system above the cavity top. This gearbox system runs the length of the insert where it is powered by room temperature servo motors with a gear ratio of 19,600 to 1. Again, to maintain electric disconnect from the cavity, and minimize power losses out of the cavity, the axles were made of alumina in earlier runs, and, more recently, of sapphire. Sapphire and alumina are low-loss dielectric material that prevent the axles from acting like antennae themselves, dissipating power out of the cavity and spoiling the quality factor. Sapphire is preferred over alumina because it has a higher thermal conductivity, making cooling of the tuning rods faster and more homogeneous with the cavity. The run 1A-B rods were 5 cm in diameter, whereas the run 1C tuning rods were 11.1 cm. The run 1D system implements a single 20.5 cm diameter rod and thus necessitating a different end cap design with only 1 tuning rod axle port. This rod weighs 51.2 lbs, and therefore is installed via crane unlike the much lighter predecessors that could be installed by hand. These are all pictured in Figure 4.9.




Figure 4.9: The various ADMX tuning rod arrangements within the run 1 cavity.

The cavity has 3 antenna ports: 2 "weak" or fixed coupling ports, and 1 "strong" or variable coupling port. The strong port antenna is the data-taking antenna for power time-series data, and has its own gearbox and room temperature stepper motor that allows the operator to adjust its insertion depth into the cavity, which allows for variable coupling to the

cavity (coupling will be discussed more in depth in Chapter 6). The other two antenna ports have fixed antennae that are set to be very weakly coupled to the cavity. These weak ports are only used for periodic cavity characterization periods: monitoring the quality factor, resonant frequency, and coupling between tuning cycles. For this reason, the coupling is kept weak to prevent it from dissipating potential axion power. In practice, only one of these weak ports was actively used, and the other port was kept as a back-up. The strong port, being the axion data antenna, is kept critically coupled (much stronger coupling) to maximize the potential axion power going into the receiver system. This data-taking cadence will be discussed more in depth in the next section of this chapter.

4.5 Helium dilution refrigerator

The helium dilution refrigerator is a custom high-flow design developed and manufactured by Janis Research (part of which has more recently been bought by Lakeshore Cryogenics) with collaborative help from the ADMX group at University of Florida. It is a 'wet' dilution system and has a cooling power of $800\ \mu\text{W}$ at an operating temperature of 100 mK; this is the physical temperature, not to be confused with system noise temperature that determines the noise floor of the experiment.

A helium dilution refrigerator makes use of a phase separation between ^3He and ^4He for mixtures below 870 mK: The concentration of ^3He can either be highly concentrated (essentially pure ^3He) or a dilute mixture of about 6.6% (this ratio is fixed by Van Der Waals forces in a stable equilibrium). This mixture exists in the part of the refrigerator called the mixing chamber, with the pure phase on the top of this chamber, while the fixed dilution is below. By pumping from the dilution side of the mixing chamber, ^3He is drawn into the dilution to maintain the fixed concentration ratio. Because this transition is going from an orderly to more disorderly phase, entropy increases, and heat must be absorbed from the system to maintain the dilution ratio. This endothermic reaction is what removes heat from the cavity and surrounding components.

As outlined in Figure 4.10, there are many other important components to the standard

fridge system that pre-cool the helium to make this interaction possible. First of all, typically a liquid nitrogen bath surrounds the system to provide base cooling to 77 K, connected to some heat shielding system, as well as a liquid ^4He bath to then cool to 4 K. In the case of ADMX, the magnet cryostat serves this purpose. The "4-K stage" consists of a helium reservoir that was filled about once a day. This reservoir supplies helium to the 1-K "pots" in the experiment. As shown in Figure 4.10, the 1-K pot is connected to the helium reservoir, as well as a 1-K bath pump, and by pumping in the vapor above the ^4He in the pot, the boiling point is lowered to 1 K. This pot is then connected via a heat exchanger to the pure ^3He input line for pre-cooling before it enters the mixing chamber. In the ADMX system, this ^3He pre-cooling pot is referred to as the "small 1-K pot", while another, "big 1-K pot", is an auxiliary pot to provide cooling for shielding and other components to 1 K. After pre-cooling the pure helium via the small 1-K pot, the pure ^3He continues through the main impedance line that has a high flow resistance and is connected via heat exchangers to the "still", described below, which cools it below the phase transition point to 500 – 700 mK. This is where the nearly pure ^3He then enters the mixing chamber, where it mixes and energy is pulled from the system as it dilutes, cooling the system to the 130 mK cavity operating temperature. As the pumped dilution moves out of the mixing chamber, it provides cooling to the incoming ^3He , and then enters the "still". The still is the area where ^3He is then separated again from the ^4He . Pressure is still kept low (about 10 Pa) via a ^3He pump, and because ^3He has a lower boiling point than ^4He , the vapor in the still is basically pure ^3He , whereas the ^4He remains liquid and stationary. A still heater controls this distillation process, keeping a steady flow of ^3He out, where it is then subsequently compressed again and fed back into the system on the input end, thus completing the fridge cycle.

As pictured in Figure 4.11, The ADMX dilution fridge mixing chamber is attached to the "T-plate" or "mK stage" which is directly on top of the top cavity end cap. This is also where the bottom of the Squidadel arm (more on this later) is attached to, as well as the Sidecar cavity that will be discussed in Chapter 9. The pots and still sit above this stage, and are not pictured here. The output plumbing of all this comes out of the top of the insert,

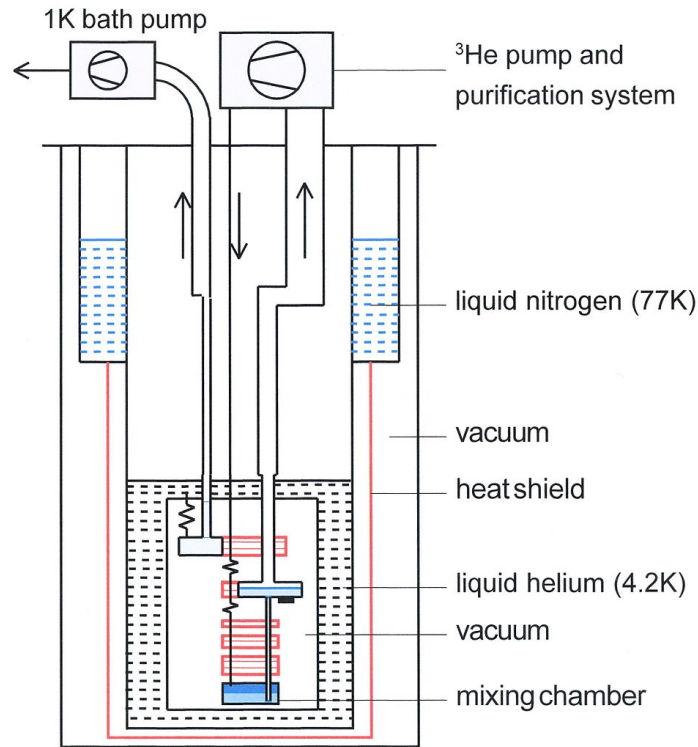


Figure 4.10: A simple schematic of a generic 'wet' helium dilution refrigerator [3].

and is then fed back down into an area known as "the pit" where the magnet itself sits. This is where the pumps and control systems are kept for the dilution fridge, as well as traps for storing the $^4\text{He}/^3\text{He}$ mixture when performing maintenance (see Figure 4.12).

4.6 Cryogenic receiver

The RF circuit layout of the signal receiver within the insert is picture in Figure 4.13 with the various temperature stages outlined.

As mentioned in the cavity section, the cavity has two antennae, one strongly coupled antenna for extracting axion data, and one weak antenna for inputting power during calibration measurements. For instance, this weak port is what allows for the in-situ measurement of the transmission cavity quality factor. Within the schematic, the weak antenna is the




Figure 4.11: The ADMX dilution refrigerator mixing chamber mounted on top of cavity with heat exchangers (the coils) above providing pre-cooling to the $^4\text{He}/^3\text{He}$ mixture input.

connection on the bottom, and the strong antenna is on top of the cavity. In addition to calibration, some "synthetic axion" signals are injected into the weak port periodically for testing the haloscope and its operators ability to detect signals, although this sub-system won't be discussed in depth in this dissertation.

In addition to the transmission measurements through the weak port, the cavity bypass line allows for reflection measurements off the strong port: Power is injected through the bypass line. It then travels through the circulator to bounce off the strong port antenna, back



(a) Fridge control.

(b) Fridge pump.



(c) Helium traps.

Figure 4.12: The exterior ADMX dilution refrigerator components.

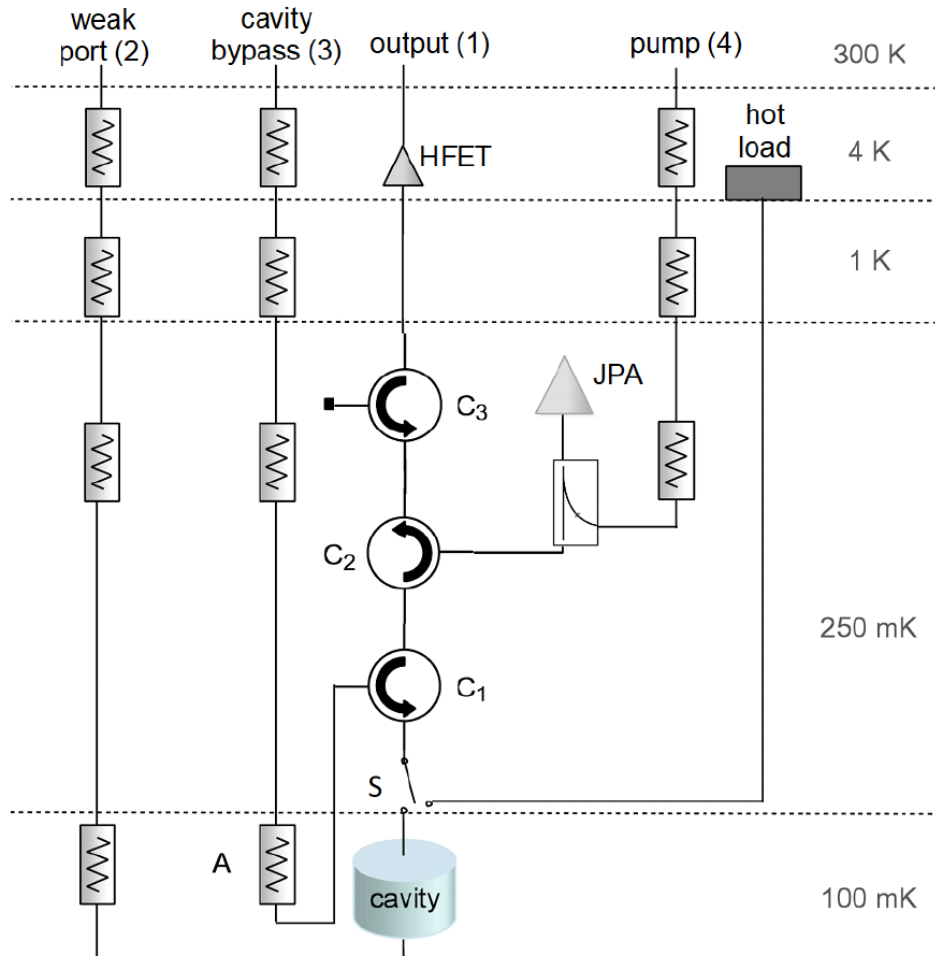


Figure 4.13: The schematic of the ADMX cryogenic receiver. " C_n " represents the circulators. "S" is a solenoid switch. "A" represents the attenuators. "JPA" is the Josephson parametric amplifier. "HFET" is a heterostructure field-effect transistor amplifier. The pump line is used for supplying a pump tone to the JPA. The weak port and cavity bypass port are used for transmission and reflection measurements respectively, which allow for in situ cavity characterization and calibration.

into the circulator and out the output line (ignoring the upstream amplification components). These reflection measurements determine the coupling coefficient (β) of the strong port, which is a measure of the fraction of cavity mode power being absorbed by the antenna (more on this in Chapter 6). The cavity bypass line can also be used to assess the overall

health of the receiver and amplification components separate from the cavity as discussed in the next chapter.

As shown in Figure 4.13, both of these lines have several attenuators along the various thermal stages of the insert for two reasons: These act as heat sinks for the connecting coaxial cables themselves, as well as act as noise suppression for the warm thermal noise coming into the cold space. These coaxial cables for the two lines are made of stainless steel flashed with copper, in order to minimize the thermal conduction between the temperature stages.

The circulators along the output line are critical for directing the RF signals through the right components and stopping them from entering the wrong components. Circulators are non-reciprocal devices that act as a one-way roundabout or rotary for RF signals. In our case, they have 3 ports. The key difference between these circulators and a roundabout you'd find driving, is that the RF signals have to exit at the next port from when they entered. For example, a signal entering on port 1 will exit via port 2 and never from port 3 ($1 \rightarrow 2$, $2 \rightarrow 3$, $3 \rightarrow 1$). This one-way directionality is represented by the counter-clockwise arrow in the schematic. Its important to note, however, these are real-world devices, so there is some signal loss between ports, and isolation isn't perfect, but it is characterized. Being that this is the output line, coaxial cables used here were also superconducting NbTi, in order to minimize all RF loss, up until the HFET.

Each circulator has a specific purpose. C_1 is primarily for directing the cavity bypass reflection signal as described above; the input signal on the bypass line can only travel to the strong port antenna first, and, vice-versa, the reflected power can only travel out the output line, not the bypass line. C_2 directs the raw cavity power signal to the primary amplifier, the JPA, where it is amplified by being reflected off the JPA. This amplified reflected power is then directed only up the output line by C_2 and not back down to the cavity. Additionally, any thermal noise from upstream output line is then dumped into C_1 where it is directed into the by-pass line, preventing any noise getting into the cavity. C_3 serves a similar isolation purpose: the amplified signal coming out of C_2 is directed to the HFET, secondary amplifier,

and any noise from the HFET side is directed to a terminator rather than back down the output line.

The JPA itself requires what's called pump tone for operation, hence the 4th RF line labeled "pump" in Figure 4.13. Josephson parametric amplifiers work similarly to how you swing on a swing set; you pump your legs at a given frequency back and forth, which drives the amplitude of the swinging to increase. This only happens because the frequency of your leg pumping motion is in the vicinity of the resonant frequency of that swing (thinking of it as a pendulum), and the resonance makes each leg pump more efficiently transfer energy into increasing the amplitude to your swing. However, if you pump too fast (or too slow) the swing amplitude stays relatively low. The JPA pump tone acts as the legs in this case, and instead of the pendulum frequency of the swing, you have your cavity output power spectrum. In electronics terms, the JPA acts as a resonant LC circuit here, where the inductance is adjusted by varying the critical current of a Josephson junction, which is done by applying a magnetic flux via a separate squid loop threaded through the junction. The pump tone is a modulation of bias current supplied to this threaded squid loop which then adjusts the resonant frequency of the JPA. As in the swing set example, the gain of the JPA is optimal when the pump tone matches the frequency of the incoming cavity signal (in the case of current-biased amplifiers. It is twice this frequency for flux-pumped amplifiers). However, this isn't optimal because it puts a separate signal right over where one wants to look for axions, therefore to decouple the pump tone from the cavity signal, it is offset from the signal frequency by 375 kHz. This spacing was chosen to allow use of band pass filters in the warm space to remove the pump tone before digitizing the power from receiver.

The JPA provides at least 20 dB of gain in most runs over a 10 – 20 MHz bandwidth. Since this bandwidth is considerably smaller than the frequency space covered in a run, it is adjusted by changing the DC bias current and power of the pump tone signal roughly every 10 minutes of operation.

In Figure 4.13, the switch, S, just above the cavity strong port is a solenoid operated coaxial switch that allows for switching between the cavity and a "hot-load" on the 4-K stage.

The hot-load acts as a variable temperature noise source, which allows for measurement of the JPA noise temperature, a process that will be discussed later in Chapter 5 . The switch was actuated by mA order DC current applied for about 1 second originally, which supplied a considerable amount of heat to the system. In the next subsection, I will talk about a small project I did to fix this issue, and minimize the heat to actuate.

What Figure 4.13 doesn't show you is what we call the "SQUIDadel", which houses all the magnetic-field sensitive components within the bucking coil: The ferromagnetic circulators, the SQUID amplifiers, and the solenoid switches. It consists of a long gold-plated copper shaft, called the "cold-finger", with a bolt head on one end for bolting it to the T-plate, thermally sinking it to the dilution fridge mixing chamber for maximum cooling. The other end is a chassis for the package itself containing switches and circulators on the outside and in the center a μ -metal shielded core for the ultra-sensitive amplifiers as shown in Figure 4.14. This shaft raises the SQUIDadel up into the field-free region, while also allowing it to be cooled by the mixing chamber. The coaxial cables are wrapped around the cold finger down to the mK-stage where they then break out and move up the insert in thermal stages. In order to prevent a thermal touch and minimize radiation imposed on the package from the relatively warm Bucking coil space (4 K), nylon spacers have been added to prevent any leaning of the cold finger into the bucking coil. I will also note that although this section only covered the cold receiver for the main experiment, a parallel receiver exists for the Sidecar system and is also housed in the same SQUIDadel cold finger. The other components, mainly HFET, attenuators, and hot-load, are mounted in the field region, on their appropriate temperature stages.

4.6.1 *A pulsed RF switch box*

As mentioned in the last section, the coaxial switch, used for hot-load calibrations, was hand-actuated by an operator with a DC power supply by applying around 160 mA for one second, which could overheat the refrigerator. One of my first projects on ADMX was designing a pulsed power supply system that would remove the operator from turning the




Figure 4.14: Left: The SQUIDadel chassis with circulators and coaxial switches and amplifier core. Top right: The amplifier core with amplifiers mounted and shielding removed. Bottom right: the JPA itself used in Run 1B.

direct connection to the power supply on and off. The idea was to implement a simple RC circuit that would be charged up by the power supply, disconnected, and then reconnected to the coaxial switch where it would discharge, sending a pulse to actuate the switch.

The solenoid coaxial switch used is from Radiall, model R585433210, and operates from 0 – 18 GHz, with a nominal switching voltage of 28 V at 25 C. The schematic diagram of this switch is shown Figure 4.15. It is a Single-Pole Double-Throw (SPDT) switch, where 'C' in the diagram refers to common pole and ports 1 and 2 are each of the "throw" ports.

The left side of the diagram is the coaxial connections and the right side is the DC power input terminals. In the diagram, the switch is currently configured with the common pole connected to port 1, and port 2 is connected to a $50\ \Omega$ terminator. By applying a positive voltage to the common input and the port 2 power input to its ground, the temporary solenoid magnetic field exerts a force on a ferromagnet that tilts the angled sections into the reciprocal configuration where the common coax is connected to the port 2 coax, and port 1 is terminated by the other load. What was discovered about this switch configuration is one cannot simply reverse the polarity of the voltage applied to switch back, but one needs to connect the port 1 input to the negative connection and disconnect the port 2 input. If not, simple reversing the polarity can result in a incomplete throw where all 3 coaxial ports are weakly connected with high loss, until a voltage with the correct polarity is applied (or worse an absurdly high negative voltage forces it into place).

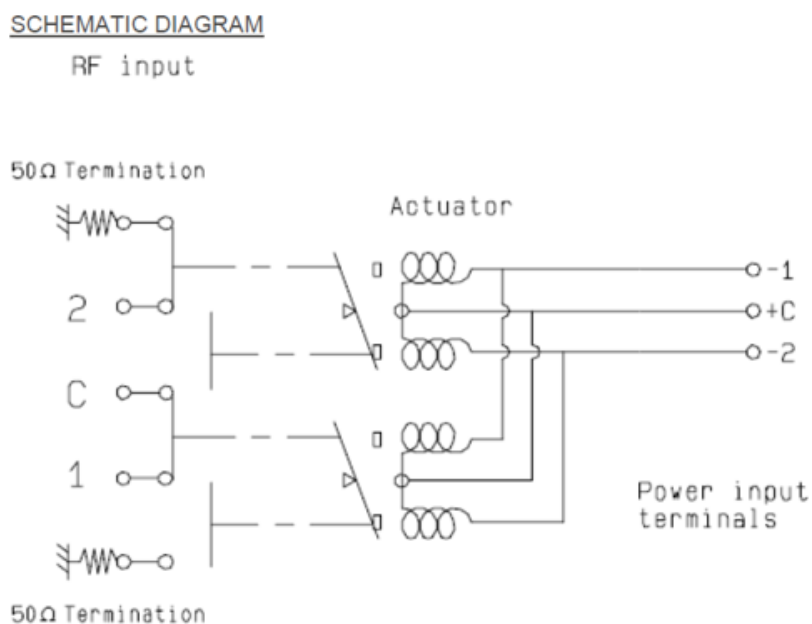


Figure 4.15: The circuit schematic of the Radial coaxial switch.

The pulsed switch system designed is shown in Figure 4.16, and was made using an Agilent switch box module (model 349980A) that consisted of an array of DC SPDT switches. The

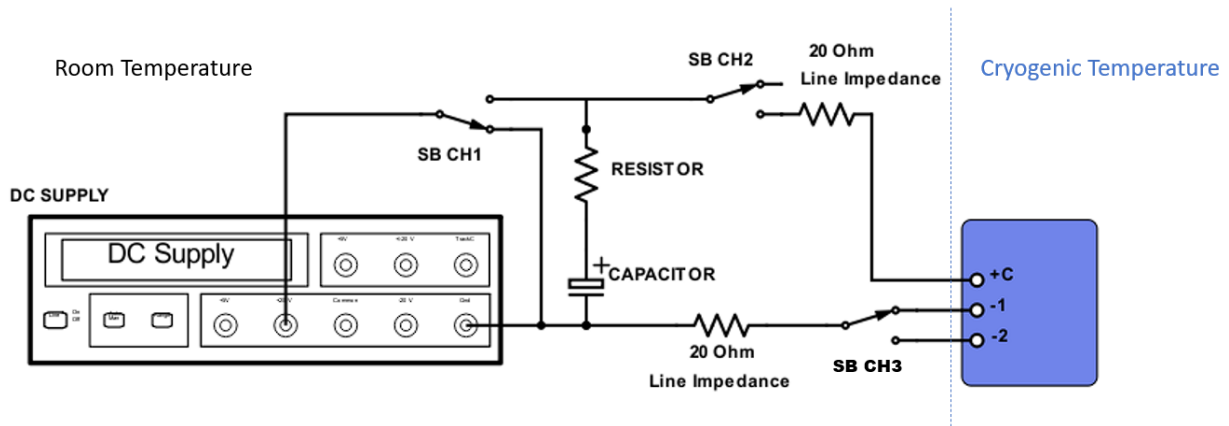


Figure 4.16: The pulsed power supply circuit design for a single coaxial switch.

CH1 switch connected the power supply to the RC circuit during the charging phase, while the CH2 switch disconnected the common input of the Radial switch. Once the capacitor had charged, CH1 would disconnect from the power supply and CH2 would connect to the common pole, discharging the capacitor into the switch and actuating it. The CH3 switch allowed for going back and forth between the port 1 and 2 ground connections, depending on which port one needed to switch to. The main parameters in this circuit then were the capacitance, C , resistance, R , and the applied voltage V_0 . These determined the key parameters of stored energy and time constant for the discharge:

$$E = \frac{1}{2}CV_0^2, \tau = RC \quad (4.1)$$

The voltage supplied then to the common pole as a function of time during a discharge was thus:

$$V(t) = V_0e^{-t/RC} \quad (4.2)$$

The key was to supply the nominal 28 V switching voltage with the correct RC time constant to actuate the switch, while still minimizing the energy deposited as heat in the system.

Test were first done at room temperature to figure out the optimal resistance and capacitance for a variety of operating voltage conditions. Initially, the thought was to match the

pulse time length to the cited actuation time of the switch from Radiall, $\tau_{pulse} \approx 3RC \approx 10ms$, and use the cited nominal voltage of 28 V. A 47Ω resistor and $100\mu F$ electrolytic capacitor gave us an $RC = 4.7ms$. This did work, but one couldn't lower the voltage anymore to minimize the energy. Although one could lower the energy stored by lowering the capacitance while increasing the resistance accordingly, capacitors are sold with capacitance varying in factors of 10, which lowered the stored energy beyond the minimum to fire. Therefore, it was settled to lower the time constant by a factor of 10, by lowering the resistance, to increase the power supplied to the switch for the same energy stored. With this strategy, the nominal switching voltage was lowered to 20 V, for an $RC = 0.47ms$ and $E = 20mJ$. It was this resistance, $RC = 4.7\Omega$, and capacitance, $C = 100\mu F$, that was settled on for cryogenic testing.

Cryogenic testing was done in both liquid nitrogen and liquid helium. The switch was mounted to a paddle that could be inserted into a cryogenic vessel holding either liquid nitrogen or helium, pictured in Figure 4.17. Coaxial cables were connected to the switch and went up and out the top of the paddle so that transmission through the switch could be verified on a network analyzer. The DC wiring for actuating the switch also came out of the top of the paddle and connected to the RC circuit kept at room temperature. Because the conductivity of the wiring and the switch's internal components goes up at cryogenic temperatures, unsurprisingly, the minimum energy to actuate the switch goes down. The results are reported in Table 4.1.

The switch box was then finalized and installed on the ADMX site for use during Run 1C. A python program handled the switching process by automating the various channel switches in the pulse firing process. Unfortunately, the operations team blew out the polar electrolytic capacitor by running the circuit with reversed polarity very early in the run, rather than using the CH3 switch. This was due to poor communication on how to operate the system and the complexity of the sequence of switches that need to be made to fire a pulse. If CH1 is not disconnected before CH2 or vice versa, the system fails. If the voltage leads are switched, the circuit destroys itself. Although frustrating as the designer of this



Figure 4.17: The author holding the coaxial switch mounted to paddle for cryogenic testing. The cryogenic vessel located behind.

system it is to see, I understand now that these are very key weak points for its usability. The system has been fixed since this incident, and now employs software that prevents such a situation from occurring as easily again.

4.7 Helium liquefaction plant

The helium liquefaction plant at the experimental site provides cost-effective cooling for the superconducting magnet, as well as cooling of the insert components to 4 K. As pictured in Figure 4.18, the magnet cryostat reservoir is a good starting point for understanding the cooling cycle: this is filled with liquid helium and supplies liquid helium to the magnet superconducting coils with periodic magnet fills, keeping the coils fully submerged at all

Temperature (K)	Minimum Pulse Energy (mJ)	Initial Pulse Power (W)
300	20	85
77	12	31
4	3	2

Table 4.1: A summary of the minimum energy and initial pulse power to fire the Radial coaxial switch at various temperatures. The RC time constant was fixed at 4.7 ms.

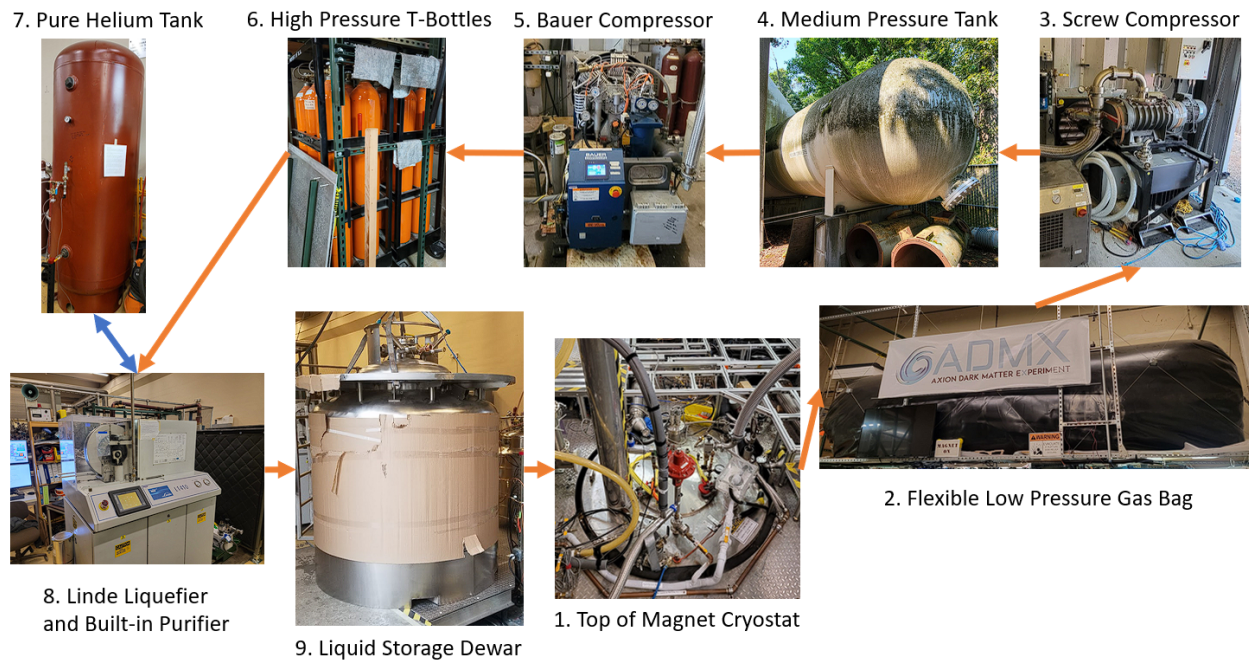


Figure 4.18: An outline of the ADMX liquid helium plant.

times. Helium vapor boils off over time in the magnet cryostat (about 20% of its capacity per day) and is then transferred to a flexible gas bag (Figure 4.18-2) that is at roughly atmospheric pressure; this bag slowly inflates over the period of about 12 hours, and upon 92% capacity is pumped down by the Kaeser Quantum Designs screw compressor (Figure 4.18-3) into the medium pressure helium gas tank that is kept outside the accelerator hall

(Figure 4.18-4). This tank fluctuates between 2 – 12 Bar with a period of about 12 days, averaging out to about 5 Bar or 72 PSI. Two Bauer compressors (Figure 4.18-5) transfer helium gas from the medium pressure tank to a set of high pressure T-bottles (Figure 4.18-6) kept at 120 Bar. The Linde helium liquefier (Figure 4.18-8) requires certain purity level of helium to operate, therefore it contains a purifier on its gas input. However, it needs a certain amount of pure helium to start the liquefier and the purification process; therefore a separate high purity helium tank is connected to the inlet as well (Figure 4.18-7). Once the liquefier and purifier have started, it then starts to pull helium from the high pressure T-bottles to liquefy (it will also fill the pure helium tank back up off the outlet of the purifier). The produced liquid helium is then transferred via a Remote Delivery Tube (RDT) to 2500 LL (Liquid Liter) storage dewar, often called the 'Mother Dewar'(Figure 4.18-9). This dewar then contains automated liquid helium pumps and transfer lines that periodically fill the magnet cryostat reservoir to maintain liquid level, thus completing the cycle.

4.7.1 2500 LL mother dewar installation at CENPA

During my first year as a research assistant for ADMX, the liquid helium plant had several key upgrades to smooth out operations that I assisted with. The primary upgrade was replacing the older and much smaller 500 LL mother dewar (Figure 4.19) with the current 2500 LL model. The other two upgrades included installation of the automated liquid fill system from this dewar to the magnet cryostat (Figure 4.20) , as well as a new computerized cryogenic control system for the entire plant (Figure 4.21).

The primary motivation for this upgrade was to increase the time between liquefaction runs by increasing storage capacity, and lightening the weekly duties of the liquefier operators. Previously, the liquefier had to be run 1-2 times per week to maintain liquid inventory for the magnet, as well as manual magnet fills nearly every day and reservoir fills every 3 days. This meant a liquefier operator had to be monitoring and controlling the system around the clock with very little room for error or contingency time for maintenance. With the increase in liquid storage by a factor of 5, it reduces liquefier runs to once every 10 days approximately,



Figure 4.19: The old 500 LL ADMX helium dewar.

and the automatic fills simply need to be monitored via a control website, rather than an in-person operation.

The installation process of the dewar is outlined in Figure 4.22. The installation of this new dewar required use of a forklift to remove it from the shipping trailer onto the deck of the accelerator hall. From there, it was transferred by crane to its current location next to the helium liquefier. A ladder and railing was installed on top of it for safety when working on top of it. The dip tube was installed by crane, followed by the RDT that connected to the liquefier. Finally, the automatic pumps were installed for transferring liquid to the magnet cryostat.



Figure 4.20: An electronic liquid helium fill pump. 4 of these pumps allow for automated fills of the magnet cryostat and its reservoir from the mother dewar.




Figure 4.21: The ADMX cryogenic control system display. This allows for operators to adjust automatic helium fill parameters and monitor helium inventory.



Figure 4.22: Installation of 2500 LL helium dewar.

Chapter 5

DATA-TAKING AND ANALYSIS TECHNIQUES FOR HALOSCOPES

This chapter will cover the necessary data-taking operations of a haloscope including the data-taking cycle as well as less periodic calibrations and operations. It will then lead into an analysis overview, focused on how the data could translate into an axion signal or the more common exclusion limit plot that is the key figure reported in every ADMX run paper so far (in the absence of a positive axion detection). Chapter 9 will cover the analysis results of the run 1D Sidecar, which will provide more details of what a real analysis looks like. The bulk of this chapter follows Refs. [21, 20].

5.1 *Data-taking cadence*

As discussed in Chapter 3, an axion haloscope scans through cavity resonant frequencies looking for power fluctuations above an average thermal noise background that could be due to an axion-photon decay. These power fluctuations can often be false axion signals such as RF interference, statistical fluctuations, and even intentionally injected "synthetic" axion signals. Therefore, certain frequency regions must be re-scanned to eliminate non-persistent signals. A given data-taking run covers the larger tuning range of the cavity, usually on the order of 100 MHz for the main experiment, while individual raw spectral centers of data were usually 2 kHz apart. To deal with this large separation in frequency span and organize the re-scanning process, runs are broken up into "nibbles", sections of the run on the order of 10 MHz wide, that take roughly a week to initially scan through.

A given nibble is first scanned through at the fastest rate possible while still achieving the desired sensitivity (this is DFSZ sensitivity for the main experiment today), but if run

conditions are not ideal, this might not be the case over the entire nibble. The goal of this first pass is to identify candidates for the re-scanning process. A data-taking cycle in this first pass starts with tracking and characterizing the resonant frequency of the TM_{010} mode after the cavity has been tuned by moving the rods at the end of the last cycle. This involves a transmission measurement using the cavity weak port, which provides the resonant frequency of the cavity and its quality factor, and then a reflection measurement via the cavity bypass line, which provides the antenna coupling coefficient for the given tuning arrangement (these network analyzer measurements will be discussed more in depth in Chapter 6). After this, a primary amplifier re-biasing procedure might take place, which would optimize the gain for the target frequency, but this is only done every 5-7 cycle iterations as it is quite time-intensive. In the case of 'un-blinded' Synthetic Axion Generated injected signals (SAG injections), there is a check if any such injections are in the target frequency band. The actual axion data taking then commences with a 100 second power digitization. Finally, the cavity is tuned to the next frequency by stepping the rods to a new position, and the next cycle begins. An additional process of manually re-coupling the strong port antenna could occur if the last coupling coefficient was not ideal. However, this requires stopping the automated data-taking software, so it is not usually included in the typical cycle. The frequency and fraction of time these processes take in a given cycle are outlined in Table 5.1.

Ideally, the first nibble pass-through can run with relative autonomy, after which the nibble re-scan procedure begins. Re-scan regions are flagged by the first pass through in the analysis software. The cavity is tuned back to these regions to repeat the same data-taking cycle, and skipping the non-flagged regions entirely. After this re-scan, any candidate signals that are persistent through both digitization cycles are then subjected to other tests that would evaluate if it had an "axionic" nature. These tests include RFI checks outside the cavity, checking for dependence on the TM_{010} mode, and ultimately, seeing if the signal power scaled with B^2 via ramping the magnet. Once all the candidates in a nibble are eliminated, the experiment then moves on to the next nibble, repeating the process until all nibbles in

Process (K)	Frequency (mJ)	Fraction of Time per Cycle
Transmission Measurement	Every Cycle	< 1%
Reflection Measurement	Every Cycle	< 1%
JPA Re-bias	5-7 Cycles	25%
Check for SAG injecion	Every Cycle	< 1%
Data Digitization	Every Cycle	98%
Move Rods	Every Cycle	< 1%

Table 5.1: A Summary of the frequency and fraction of time each process takes in a data-taking cycle of ADMX based on data from run 1B. Axion search data is only acquired during the data digitization cycle. Since the JPA bias procedure only occurs every 5-7 cycles the average fraction of time spent digitizing is kept at 98%. SAG stands for Synthetic Axion Generator.

the run are complete. This decision tree is outlined in Figure 5.1.

5.2 Hot load measurements/ Y -factor method 1

The "hot load" shown in Figure 4.13 is critical to measuring the system noise temperature of the haloscope. As alluded to before, the system noise temperature is critical to understanding the noise level in the receiver and quantifying the SNR of potential signals. Critical to the noise temperature calculation, is the receiver temperature as a whole. These measurements are done only a few times throughout the run, as they tend to produce stable results; in run 1B, this procedure was only performed 4 times. These are also done sparingly because it requires halting the data-taking process for several hours at least. They are referred to as "Hot-Cold" measurements (hence the "hot load" name) or "y-factor" measurements [120].

In abstract terms, this is done by connecting the receiver to a known variable temperature resistive source at its input and as its incremented in temperature by some ΔT , the receiver

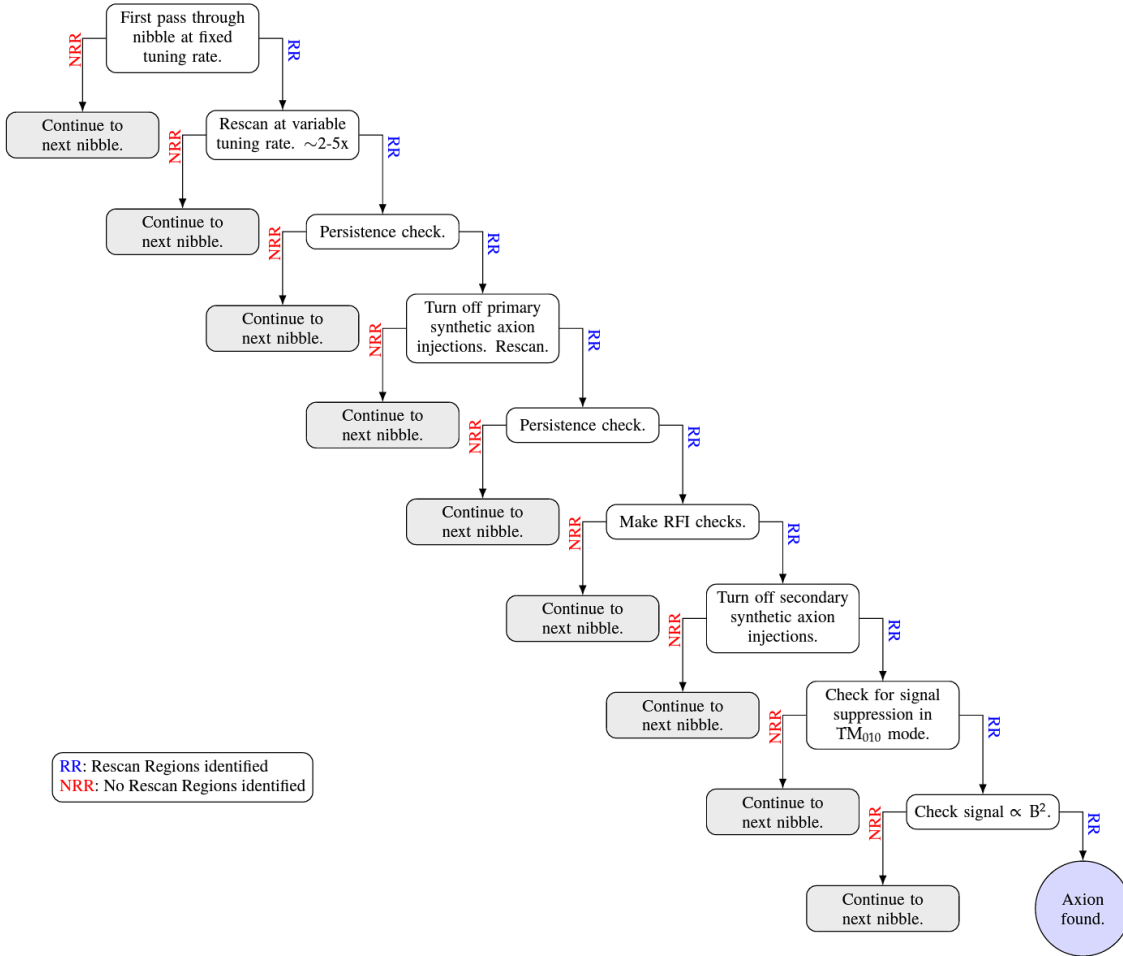


Figure 5.1: Data-taking decision tree. Typically a first nibble pass-through always results in several re-scan triggers in the analysis due to the statistics associated with the chosen tuning rate. A persistence check means verifying that a potential signal appears in all the spectra acquired for it so far, and it is not intermittent. Primary synthetic axion injections are un-blinded by the operations team, whereas secondary injections are only known by an off-site SAG operator to check the operations teams' detection efficiency. RFI checks are made with a signal analyzer outside the cavity, ruling out any external signals such as local radio stations. Final checks include trying to modulate the signal power according to Equation 3.24 by looking for suppression of the signal off-resonance and modulating the magnet field strength. Figure credit from Ref. [21].

output power will change by some ΔP (in the case of coherent receivers, with a "DC", linear response). If we label the temperature of the hot and cold states of the load, T_H and T_L

respectively, then we can write down the power output of the receiver in each case, P_H and P_L :

$$\begin{aligned} P_L &= (T_L + T_{rx})G \\ P_H &= (T_H + T_{rx})G \end{aligned} \tag{5.1}$$

defining:

$$y = \frac{P_H}{P_L} \tag{5.2}$$

the resultant receiver temperature would then be:

$$T_{rx} = \frac{T_H - T_L y}{y - 1} \tag{5.3}$$

This measurement assumes, however, that this variable temperature source doesn't change the operation of the receiver components, which is easier said than done for ADMX. In the case of ADMX, an ideal noise measurement would be done with the JPA pump enabled allowing for full characterization along the entire receiver chain. However, since the JPA can easily be saturated by thermal noise and the hot load was kept on the 4 – K stage during run 1B, the JPA was kept disabled (the hot load has since been moved to a 175 mK arrangement for run 1C-D). When its not enabled, the JPA conveniently acts as a passive RF mirror, reflecting signals with minimal attenuation back down the output line. During a measurement, the RF switch is set to point to the hot load in this arrangement (see Figure 4.13) and a heater on the hot-load raises the temperature as a temperature sensor monitors this warming process. A wide-bandwidth power measurement is then taken, which has the expected power output per unit bandwidth:

$$P = G_{HFET} k_B [T_{JPA}(1 - \epsilon) + \epsilon T_{load} + T_{HFET}] \tag{5.4}$$

where T_{JPA} and T_{load} are the physical JPA and hot load temperatures measured by sensors, G_{HFET} is the gain of the HFET secondary amplifier, and T_{HFET} is the noise temperature of the HFET and all down stream electronics. The transmission efficiency of the quantum amplifier package is given by ϵ , and is simply the linear version of the attenuation, α , in

dB ($\epsilon = 10^{-\alpha/10}$). It was quantified ex-situ by measuring the dominant loss source, the two up-stream circulators from the JPA, but also in-situ by a more complicated Y-factor measurement, or using a JPA signal-to-noise ratio improvement (SNRI) measurement (more on SNRI in the next section). In short, Equation 5.4 was used to fit for G_{HFET} and T_{HFET} by holding the hot load (and as a result, the JPA too) at a variety of temperatures.

5.3 Y-factor method 2

The second method for determining the receiver noise temperature essentially replaces the hot-load with the final cavity bypass line attenuator (labeled 'A' in Figure 4.13). In this method, the RF switch is still connected to cavity. Off-resonant thermal photons originating from the bypass line will travel through the first circulator (labeled 'C1' in Figure 4.13), reflect off the cavity, then reflect off the disabled primary amplifier package, and out through the output line through the post-amplifiers. Since these thermal photons pass through an extra length of the first circulator, they will be slightly more attenuated than a photon originating from the hot load; this was approximated 0.5 dB based on ex-situ measurements of the circulator, hence $\epsilon_c = 10^{-(\alpha+0.5\text{ dB})/10}$. The form of Equation 5.4 still applies here, however one replaces ϵ with ϵ_c , and T_{load} with T_{cav} since the bypass attenuator is mounted directly to the cavity top (there is an assumption that $T_{attenuator} = T_{cav}$). Additionally, it was approximated that T_{cav} and T_{JPA} were the same, therefore Equation 5.4 can be re-written as:

$$P = G_{HFET}k_B[T_{JPA} + T_{HFET}] \quad (5.5)$$

In this case, the fit parameters were G_{HFET} and T_{HFET} whereas T_{JPA}/T_{cav} was the varied and measured quantity. Because this temperature must still be varied, there must be a way of heating the JPA. This was done in run 1B by applying a small current to the RF switch such that it heated, but did not actuate. As will be discussed in Chapter 9, the run 1D Sidecar HFET was characterized using the magnet ramp as a heating source; The ramping of the magnet heats the insert through eddy currents, and therefore the mK attenuator. Operators in anticipation of this heating and cooling process, turn on wide-bandwidth power

measurements. These wide-bandwidth power measurements are done by taking a series of digitizations across the tuning range of the cavity; This is because the bandwidth of the digitizer is limited to 2 MHz. The idea here is to capture off-resonant photons across the search frequency range (on-resonance thermal photons will be absorbed by the cavity), so that one can give an associated HFET noise temperature for each target search frequency during normal data-taking. Gaps in between these power spectra are okay, as the noise temperature can be interpolated as a function of frequency (it should follow a black-body distribution).

There are several advantages to this method. Primarily, this gives a separate confirmation of T_{HFET} independent of the attenuation through the SQUIDadel, which improves confidence in the analysis. In terms of operations, this method can be less intensive and time-consuming, or even the only option; the RF switch can fail to actuate, rendering a hot-load measurement impossible. Conveniently, an RF switch is usually said to fail because it is heating the system towards the limits of the dilution fridge before actually actuating; one can then use this heat to perform this back-up method. It is also adaptable; any anomalous heating or cooling of the cavity and amplifiers can be a moment to perform this operation. Finally, in some cases, this can be done remotely, whereas a hot-load measurement requires on-site personnel to flip switches and monitor. See Chapter 9 for details of the Sidecar HFET characterization.

5.3.1 Designing the run 1C part 2 "hot-cold" load

One of my projects during my tenure as a graduate student was redesigning the hot-load for run 1C, where it would be moved from the 4-K stage to the intermediate cold plate that normally operated at 150 mK. This move would bring the y-factor fit measurement closer to the operating temperature regime of the JPA. The hot load would have to be heated from this base temperature up to about 500 mK without overwhelming the dilution fridge, but also cool down in a reasonable amount of time. This is done by making the hot load thermally disconnected from the mounting plate except for a well-known cooling wire connection that controls the cooling power of the load.

The original run 1C hot load is pictured in Figure 5.2. Two nylon standoffs thermally disconnect the main OHFC copper body, on which 3 attenuators are mounted to each by a single screw. One attenuator is for the main experiment receiver, one for Sidecar, and one extra was budgeted in. Within the copper body, 3 through-holes are made where $10\text{ k}\Omega$ resistors are embedded and epoxied in place; these resistors act as the heaters for the load. The resistors are connected to superconducting wire that leads to power supplies, thermally disconnecting them from the warm space above by the superconductor, but electrically connected to act as a heater when in use. A temperature sensor is mounted on the underside to monitor the heating and cooling process. Finally, two copper cooling wires were connected underneath each nylon standoff, and attached to the intermediate cold plate below.




Figure 5.2: The run 1C hot load.

This original design has been iterated since, and the current configuration is pictured in Figure 5.3. First, the extra attenuator was removed because it was an unnecessary, unknown thermal load on the piece. The nylon standoffs, being plastic threading, did not provide

enough torque to properly connect the cooling wire to the body, and thus it was given its own threaded connection with a metal screw in place of the 3rd attenuator. Additionally, the attenuators were mounted with two screws from the top, to ensure a more solid thermal connection to the body. This was motivated by the hot load being hotter than predicted, and cooling slower than predicted, which is indicative of bad thermal contacts. Key to this




Figure 5.3: The run 1D hot load.

design is understanding how to calculate the thermal conductive power of each component touching the hot load, as well as the heat capacity of the load itself. The conductive power of a wire/cylinder connecting two bodies can be written:

$$G = kT \frac{r^2}{L} \quad (5.6)$$

where k is the thermal conductivity of the wire material, T is the temperature of the wire, r is the radius of the wire, and L is the length of wire between the two bodies. In the case of the hot load, we calculate the conductive power of the two nylon stand offs, cooling wire,

heater wire, 4 coaxial cable connections, and temperature sensor connection. Next one wants to know the heat capacity of the hot load body, which is simply $C = cM$ where c is the specific heat and M is the mass of the body of copper. The screws and attenuator cubes also add to the heat capacity, but are not included in this calculation. From this we calculate the time constant for cooling as:

$$\tau_{cool} = \frac{C}{G_{cool}} \quad (5.7)$$

where G_{cool} is just the combined power of all connections to the intermediate cold plate, not the warm space connections like the heater wire and temperature sensor wire; these, if anything, should be subtracted from the cooling power, but being superconducting, they are minimal sources of heat. The heating time constant can simply be written based on the current supplied to the resistor:

$$\tau_{heat} = \frac{C}{I^2 R} \quad (5.8)$$

where R is the total resistance of the heaters and I is the current from the power supplies. The other warm and cold space connections are negligible compared to this supplied heater power, so they are not included in this calculation.

5.4 SNRI

As mentioned earlier, signal-to-noise improvement (SNRI), commonly used to quantify performance of amplifiers, can be used to circumvent directly measuring the attenuation of the quantum amplifier package. SNRI can be defined as:

$$SNRI = \frac{T_{sys,off}}{T_{sys,on}} = \frac{G_{on}P_{off}}{G_{off}P_{on}} \quad (5.9)$$

where $T_{sys,on/off}$ refers to the noise temperature with the JPA powered and off, $G_{on/off}$ is the gain of the JPA powered and off, and $P_{on/off}$ is the measured noise power when powered and off. Essentially, it gives the ratio of improvement expected from having the primary amplifier operating in the receiver. Hot load measurements have given us T_{HFET} with the amplifier powered off. We can write down the total system noise temperature with the amplifier off

as:

$$T_{sys,off} = T_{cav} + T_{JPA,off} + \frac{T_{HFET}}{G_{off}} \quad (5.10)$$

Since $T_{HFET} \approx 11K$ while $T_{cav/JPA} \approx 0.1K$, we can make the assumption $T_{HFET} \gg T_{cav} + T_{JPA,off}$. Additionally, an amplifier that isn't powered shouldn't have significant gain, and therefore can be approximated as just 1. This means we can write the total system noise temperature with the JPA powered ON just in terms of T_{HFET} and SNRI :

$$T_{sys,on} = \frac{T_{HFET}}{\epsilon SNRI} \quad (5.11)$$

where ϵ is the transmission efficiency from the cavity to the primary amplifier. Since it has been determined that T_{HFET} is pretty stable throughout the run, this means periodic SNRI measurements are all that is needed to update the system noise temperature measurements. This can be done when re-biasing the amplifier, and uses the latter half of Equation 5.9, measuring the noise power and gain both on and off. This can be done using the cavity bypass line; a known power off-resonant signal inputted through the bypass will reflect off the cavity, then be amplified by JPA, giving the gain both on and off. The noise power level can just be sampled from the output with the amplifier on and off. This makes periodic checks of the noise temperature much less time-intensive using SNRI.

5.5 Axion search analysis procedure

This section is an overview of the ADMX "medium" resolution analysis procedure, and how it ultimately produces limits on the axion parameter space when there are no persistent candidates left. As described in the data-taking cadence section, after time-series digitization data is recorded, "medium"-resolution analysis is performed on the data, separate from the data-taking cycle, to look for the presence of axion candidates, determining re-scan regions for a given nibble. Once all these candidates have been reviewed and eliminated, a limit on $g_{a\gamma\gamma}$ over the frequency range can be set. This process can be divided into spectra background processing, analysis cuts, and grand spectrum construction.

5.5.1 Spectra Background Processing

Each spectrum in the main experiment during run 1B consisted of 512 bins each with a width of 95 Hz, for a total spectrum width of 48.8 kHz; an example spectrum is shown in Figure 5.5(a). A single spectrum is the result of 100 seconds of time-series data being chopped up into 10^4 Fourier transforms of each 10 millisecond segment. There is frequency overlap in the spectra, as was mentioned in the scan rate section, meaning we will have to account for multiple spectra having the same frequency bins. Therefore it's easiest to denote the power in the "j"th spectra within the "i"th bin as P_i^j .

The first step in processing the raw spectra is removing the fixed background shape due to the warm electronics from all the raw spectra. This is primarily due to frequency dependent gain variations after mixing. The shape of these variations is somewhat predictable by the filters in the receiver chain. Subdominant contributions are gain variations before mixing that originate in the cold space, as well as frequency dependent noise variations. The background shape is pictured in Figure 5.4, alongside its fit which was averaged and smoothed using a Savitzsky-Golay filter [103, 81]. This baseline filtered shape is then divided from all the raw spectra.




Figure 5.4: The filtered background shape ("Filtered BG") and average base-line fitted shape from the warm electronics during run 1B [21].


After the baseline shape is removed, there is still more filtering and normalization to

be done on each raw spectra. A six-order Padé filter was convoluted on each spectra to remove the residual shape from the cryogenic receiver transfer function. This is motivated by a derivation of the power spectrum shape at the output of the last-stage cold amplifier [40]. The power in each bin is then normalized by dividing by the mean power of the whole spectrum. This means in the absence of an axion signal, one would expect a Gaussian distribution around a mean of one, which is the case if one creates a histogram of the data. Since we are looking for power excesses above one, one is then subtracted from all the power bins, so it is more intuitively looking for excesses above zero. This filtered result is pictured in Figure 5.5(b).


Another feature that must be accounted for in this raw data is the Lorentzian shape of the cavity resonance. If an axion signal is towards the edges of the spectrum it will not be enhanced fully by the Q of the cavity, whereas a signal that is directly centered will be. This Lorentzian shape is determined by the cavity resonant frequency and loaded quality factor at the time of the scan. Each bin is then further normalized based on its distance from the cavity resonant frequency, and where it would lay on this shape. Evident after this processing is the increase in the excess on the edges of a Lorentzian weighted spectrum, as pictured in Figure 5.5(c).

5.5.2 Analysis Cuts

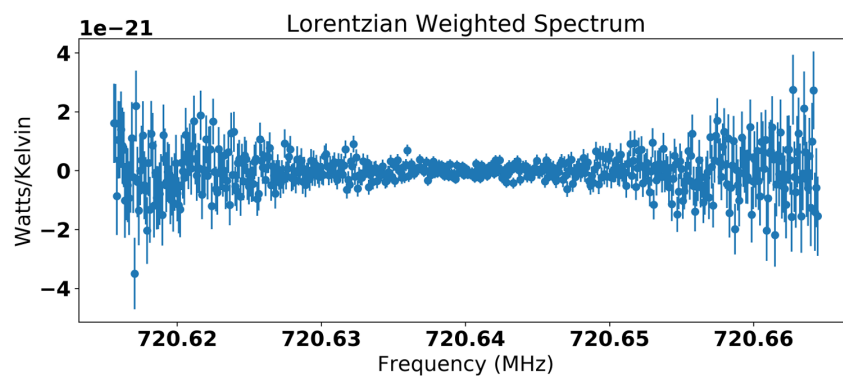
After the spectra are processed, there are a series of cuts to the data to be made for quality control. The vast majority of the cuts made are referred to as "timestamp" cuts. These are simply scans taken during known, aberrant run conditions. This could result from manually adjusting the dilution fridge, manual biasing of the JPA, digitizer failures, software malfunctions, and various other engineering studies that could compromise the data quality. During run 1B, these made up about 57% of all the analysis cuts made. The 2nd most common analysis cut, about 36% of cuts in run 1B, are scans that have an associated system noise temperature that is either less than $0.1K$, which is below the quantum limit and nonphysical, or greater than $2K$; these both likely result from incorrectly measuring the SNRI and



(a) Example of a raw spectrum from run 1B.



(b) Example of a filtered spectrum from run 1B.



(c) Example of a Lorentzian weighted spectrum from run 1B.

Figure 5.5: The raw spectrum preparation process.

therefore not trusted to be included in the data set. The remaining 6% of cuts are similarly on associated data values being badly measured; quality factors less than 10,000 or greater than 120,000 are omitted because they are likely from a bad transmission fit measurement. The six-order Padé fit had to have a χ^2 per degree of freedom less than 2 to be included, as well as a few additional cuts on the warm receiver filter shape. All in all, in run 1B, these cuts reduced the total run scans from 197,680 spectra to 185,188, or about 7% of the data set.

5.5.3 Grand Spectrum Construction

The final step in the analysis process to finding a candidate axion is to construct a 'grand spectrum' that combines all of the processed, normalized raw spectra together, usually for just a single nibble. To do this, we must account for the varying potential axion signal height as well as noise floor across spectra, based on their associated run conditions at the time of digitization; scaling the spectra based on their measured Q, f, C_{010} , and T_{sys} . Each filtered P_i^j is converted to a $P_{i_{scaled}}^j$ based on this. To further increase sensitivity, a filtering process based on the axion line shape is also performed, by convolving each bin power with a filter associated with either the Maxwell-Boltzmann line shape (Equation 2.16) or N-body line shape (Equation 2.17) discussed in Chapter 2.

The combining process uses a well-established 'optimal weighting procedure' outlined in Ref. [12, 31]. This process finds weights for the individual power excesses that results in the optimal SNR for the grand spectrum. The weights are chosen such that the maximum likelihood estimation of the true mean value, μ , is the same for all contributing bins; this can be outline mathematically on a bin-by-bin basis using the equation:

$$P_w = \frac{\sum_{j=0}^N \frac{P_{scaled}^j}{\sigma_j^2}}{\sum_{j=0}^N \frac{1}{\sigma_j^2}} \quad (5.12)$$

where N is the total number of spectra for a given frequency bin, and P_w is the weighted power for an individual RF bin in the grand spectrum. The standard deviation for each bin

in the grand spectrum is then calculated with

$$\sigma_w = \sqrt{\frac{1}{\sum_{j=0}^N \frac{1}{\sigma_j^2}}} \quad (5.13)$$

The grand spectrum is then completely defined by power excess P_w and its standard deviation, σ_w . It is expected that this grand spectrum will be relatively level and flat with a Gaussian distribution if the background has been properly subtracted. From this grand spectrum, sharp power excesses above the noise can be flagged as axion candidates; more precisely, when the weighted power at that frequency is in excess of 3σ . A re-scan would also be triggered if a candidate's power exceeded that of a DFSZ axion by 0.5σ . Regions in the grand spectrum that didn't meet the target SNR or target DFSZ sensitivity are also flagged to be re-scanned. Once all candidates are deemed not to be axion-like, and the desired sensitivity is met over the grand spectrum, an exclusion limit can be set.

5.5.4 *Setting Limits*

Once all data-taking is completed, a final limit is set on $g_{a\gamma\gamma}$ over the frequency range that was scanned. A frequency bin that contained an axion signal, if scanned multiple times, would have a Gaussian distribution about a mean, $\mu = g_\gamma^2 \eta$, where η is the SNR for the given measurement. Whereas, a frequency bin not containing a signal would be Gaussian about a mean $\mu = 0$. A limit is set by computing μ for each bin that gave a 90% confidence limit that the measurement did not contain an axion. To deal with negative power values, a cumulative distribution function for a truncated normal distribution is used when determining the 90% confidence μ value [50]. For each bin, a normal distribution is generated using the weighted power as the mean and the uncertainty as the standard deviation. The generated distribution is then randomly sampled 100 times, clipping to zero if a negative value is generated; this list of generated μ values is then sorted, and the 90% confidence limit is value closest to 90% to the top of the sorted list. This results in a very jagged limit plot, with each bin value being around 100 Hz wide, thus some combining and smoothing is also done; A small number of bins, around 200 representing a single plot pixel, were combined into a single limit value.

The most recent limit plot for the main ADMX experiment is shown in Figure 5.6. These limits are also shown, albeit zoomed out, in Figure 2.6.




Figure 5.6: The ADMX Gen 2 exclusion limit plot from October 2021. The light blue is from the run 1A 2018 paper, Orange is run 1B concluded in 2019, and in red is run 1C part 1 that concluded in 2021. Run 1C continued into 2022, to exclude the full range to DFSZ level, but these results have not been reported at the time of writing yet.

Chapter 6

CAVITY RESONATORS

This chapter will cover the physics and techniques of resonant cavities. This will not only be in the context of axion haloscopes, but a general characterization of cavities with a variety of geometries. This will start with a formal field formalization for cavity resonant modes, and how to label their geometries. This will then transition to defining quality factor and the techniques surrounding Q measurements. From there, a brief overview of antenna coupling. The final sections will cover considerations for multi-cavity systems, and give an overview of the ADMX multi-cavity systems in development: Run 2 and ADMX Extended Frequency Range (EFR).

6.1 Formal introduction to cavities and modes

A cavity is simply an empty space within a solid object; make this solid object a conductive material, and electromagnetic energy within the empty space will be trapped within as the conductor shields any external radiation and constrains the wave structure of the electromagnetic fluctuations within the cavity. The goal of this section is to solve for the structure of the preferred oscillations within a given cavity structure, which will be called the resonant modes of the cavity.

6.1.1 Solving Maxwell's equations

For this section I will assume a cylindrical geometry and coordinate system for most of the proofs, but this analysis can be done for a rectangular or spherical, or general cavity structure. This analysis stems mostly from Ref. [67]. An easy starting point is to consider a cylindrical cavity as a cylindrical wave guide with plane end faces imposed on either side,

separated by some finite distance. We will further impose that these "end-caps" will be flat and perpendicular to the cylindrical cross-section of the wave guide. Furthermore, this cavity is made entirely of a perfect electric conductor (PEC) material.

Cylindrical Wave-guide

Before assessing the boundary conditions of the end caps, let us write down Maxwell's equations for an arbitrary cylindrical wave-guide. I can first make an assumption that the propagating waves have a sinusoidal time-dependence $e^{-i\omega t}$ for fields inside the cylinder. Furthermore, I can see for a infinite cylindrical wave-guide, there is no restriction along its z -direction yet, thus it will have a spatially sinusoidal variation, $e^{\pm ikz}$, with k as a free, unknown, possibly complex parameter. I can write the fields down as such:

$$\left. \begin{array}{l} \vec{E}(x, y, z, t) \\ \vec{B}(x, y, z, t) \end{array} \right\} = \left\{ \begin{array}{l} \vec{E}(x, y)e^{\pm ikz - i\omega t} \\ \vec{B}(x, y)e^{\pm ikz - i\omega t} \end{array} \right. \quad (6.1)$$

This wave equation under these conditions reduces to the two-dimensional form:

$$[\nabla_t^2 + (\mu\epsilon\omega^2 - k^2)] \left\{ \begin{array}{l} \vec{E} \\ \vec{B} \end{array} \right\} = 0 \quad (6.2)$$

where μ and ϵ are the permeability and permittivity of the interior, non-conducting material (presumably vacuum or some dielectric material). ∇_t^2 is the transverse part of the Laplacian operator $\nabla^2 - \partial_z^2$. It is thus convenient to separate the fields into transverse and parallel components to the z -axis, such that $\vec{E} = E_z \hat{z} + \vec{E}_t$, where $\vec{E}_t = (\hat{z} \times \vec{E}) \times \hat{z}$. We can then write out Maxwell's equations as such:

$$\frac{\partial \vec{E}_t}{\partial z} + i\omega \hat{z} \times \vec{B}_t = \nabla_t E_z, \quad \hat{z} \cdot (\nabla_t \times \vec{E}_t) = i\omega B_z \quad (6.3)$$

$$\frac{\partial \vec{B}_t}{\partial z} + i\mu\epsilon\omega \hat{z} \times \vec{E}_t = \nabla_t B_z, \quad \hat{z} \cdot (\nabla_t \times \vec{B}_t) = i\mu\epsilon\omega E_z \quad (6.4)$$

$$\nabla_t \cdot \vec{E}_t = -\frac{\partial E_z}{\partial z}, \quad \nabla_t \cdot \vec{B}_t = -\frac{\partial B_z}{\partial z} \quad (6.5)$$

Evident from the first two sets of Equations, 6.3 and 6.4, is that if the z components are known, the transverse components are completely determined, with only the sign of k , the direction of the propagation, changing the fields' direction. These transverse fields are:

$$\vec{E}_t = \frac{i}{(\mu\epsilon\omega^2 - k^2)} [\pm k \nabla_t E_z - \omega \hat{z} \times \nabla_t B_z] \quad (6.6)$$

$$\vec{B}_t = \frac{i}{(\mu\epsilon\omega^2 - k^2)} [\pm k \nabla_t B_z - \mu\epsilon\omega \hat{z} \times \nabla_t E_z] \quad (6.7)$$

This leaves us with several options for the types of waves that can propagate in this arbitrary cross-section wave-guide. The simplest case would be where both E_z and B_z are zero, therefore completely transverse to the propagation direction; these are called transverse electromagnetic (TEM) waves. In this case, $\vec{E}_t = \vec{E}_{TEM}$, simplifies using equations 6.3 and 6.5 to:

$$\nabla_t \times \vec{E}_{TEM} = 0, \quad \nabla_t \cdot \vec{E}_{TEM} = 0 \quad (6.8)$$

This means that \vec{E}_{TEM} is a solution to an electrostatic problem in two dimensions. This imposes several restrictions on this mode. First, for Equation 6.2 to be satisfied, $k = \sqrt{\mu\epsilon}\omega$, thus the wave number or wavelength is determined by the properties of the propagation material for an infinite variety of frequencies. Additionally, Equation 6.4 sets $\vec{B}_{TEM} = \pm \sqrt{\mu\epsilon} \hat{z} \times \vec{E}_{TEM}$, which is just the same connection for plane waves in open infinite space. Most importantly, this means that the mode cannot exist on a single cylindrical conducting surface because that conducting surface must be an equipotential. If there are two or more separated conducting surfaces, then this TEM oscillation can propagate electrostatically by one surface supporting a positive equipotential and another operating as the negative or ground potential. Therefore, this mode is only dominant for familiar coaxial cables, or parallel-wire transmission line geometries. Cavities that contain tuning rods will also have TEM modes. Important to note, is it has no cutoff frequency in a wave-guide, and k is real for all ω , which won't be true for hollow cylinders of 1 surface.

The next more generalized classification of modes would be if only one of the field components was transverse, meaning either $E_z = 0$, Transverse Electric (TE) waves, or $B_z = 0$,

Transverse Magnetic (TM) waves. In this case, we still have to satisfy some surface boundary conditions at the surface of the conductor to ensure the fields vanish inside. If $B_z = 0$, E_z will still have to vanish at the surface of the conductor, so $E_z|_S = 0$. If $E_z = 0$, we must ensure the component of \vec{B} parallel to the surface must vanish, requiring $\frac{\partial B_z}{\partial n}|_S = 0$ where \hat{n} is the normal unit vector to the cylindrical surface. One notices that when pairing these two boundary conditions with 6.2, it forms two distinct eigenvalue problems. For a given value of k , only certain ω frequencies can be supported, and vice-versa. Since TE and TM have distinct boundary conditions, their eigenvalues are not related in general, and form a natural way of dividing the type of fields propagating in the wave-guide. With these 3 categories, TEM, TE, TM waves, one can construct any arbitrary electromagnetic field disturbance within the wave-guide; they form the basis of waves within the wave-guide or cavity.

In the case of TE and TM modes, we can see that the eigenvalue problem formed implies cutoff frequencies for the various waves. We can write this eigenvalue problem in terms of a scalar potential function, ψ , such that for TM waves:

$$\vec{E}_t = \pm \frac{ik}{\gamma^2} \nabla_t \psi \quad (6.9)$$

or for TE waves:

$$\vec{H}_t = \pm \frac{ik}{\gamma^2} \nabla_t \psi \quad (6.10)$$

where $\gamma^2 = \mu\epsilon\omega^2 - k^2$. The wave equation, Equation 6.2, then simplifies to the eigenequation:

$$\nabla_t^2 \psi = -\gamma^2 \psi \quad (6.11)$$

with the boundary conditions for TM or TE modes:

$$\psi|_S = 0 \text{ or } \frac{\partial \psi}{\partial n}|_S = 0 \quad (6.12)$$

Thus we expect a spectrum of eigenfunction solutions, ψ_λ , with corresponding eigenvalues γ_λ^2 , where $\lambda = 1, 2, 3, \dots$, which form an orthogonal set of modes (note that γ^2 must be non-negative to satisfy Equation 6.12). This means that there will be k_λ that is only a real value above a cutoff frequency:

$$\omega_\lambda = \frac{\gamma_\lambda}{\sqrt{\mu\epsilon}} \quad (6.13)$$

The wave number can then be written in terms of frequency:

$$k_\lambda = \sqrt{\mu\epsilon}\sqrt{\omega^2 - \omega_\lambda^2} \quad (6.14)$$

This is all to show that below the cutoff frequency, certain frequencies cannot propagate within the wave-guide. The common practice is to set the dimensions of the wave-guide accordingly such that only the lowest mode can occur at the desired operating frequency.

Cylindrical Resonant Cavity

When we add conductive end-caps to either end of the infinite cylindrical wave-guide from the last section, a boundary condition is imposed in the z-direction such that only a standing wave structure, $A \sin kz + B \cos kz$, with value zero at either end cap, can propagate. For a cavity of length L , this means $k = n\pi/L$ for $n = 0, 1, 2, \dots$ and the resultant fields can be written in terms of ψ . For TM Modes these fields are:

$$\begin{aligned} \vec{E}_z &= \psi(x, y) \cos\left(\frac{n\pi z}{L}\right) \\ \vec{E}_t &= -\frac{n\pi}{L\gamma^2} \sin\left(\frac{n\pi z}{L}\right) \nabla_t \psi(x, y) \\ \vec{H}_t &= -\frac{i\epsilon\omega}{\gamma^2} \cos\left(\frac{n\pi z}{L}\right) \hat{z} \times \nabla_t \psi(x, y) \end{aligned} \quad (6.15)$$

and for TE Modes the fields are:

$$\begin{aligned} \vec{B}_z &= \psi(x, y) \sin\left(\frac{n\pi z}{L}\right) \\ \vec{E}_t &= -\frac{i\omega\mu}{\gamma^2} \sin\left(\frac{n\pi z}{L}\right) \hat{z} \times \nabla_t \psi(x, y) \\ \vec{H}_t &= -\frac{n\pi}{L\gamma^2} \cos\left(\frac{n\pi z}{L}\right) \nabla_t \psi(x, y) \end{aligned} \quad (6.16)$$

where γ^2 is adapted to:

$$\gamma^2 = \mu\epsilon\omega^2 - \left(\frac{n\pi}{L}\right)^2 \quad (6.17)$$

and the cutoff frequency is now further adapted by n :

$$\omega_{\lambda n}^2 = \frac{1}{\mu\epsilon} \left[\gamma_\lambda^2 + \left(\frac{n\pi}{L}\right)^2 \right] \quad (6.18)$$

Again, a cavity designer would try to adjust L and γ_λ such that the intended operating mode frequency was well-separated from other modes, to improve stability.

The important special case for ADMX, and all the test cavities used at LLNL, is the right cylindrical cavity. So far the analysis has been agnostic to this geometry to maintain a level of generality, but the only step left is the solve for $\psi(x, y)$ for this circular cross-section. It is best to work in cylindrical coordinates ρ and ϕ in this case. If the cavity has a radius, R , then $E_z(\rho = R) = 0$ for a TM mode, and ψ has solution:

$$\psi(\rho, \phi) = E_0 J_l(\gamma_{lm}\rho) e^{\pm il\phi} \quad (6.19)$$

where:

$$\gamma_{lm} = \frac{x_{lm}}{R} \quad (6.20)$$

x_{lm} is the m th root of the equation, $J_l(x) = 0$, where $J_l(x)$ is the l th Bessel function. E_0 is a normalization constant. The expression for ψ works just as well for TE modes, however $\gamma_{lm} = x'_{lm}/R$ where x'_{lm} is the m th zero root of the derivative Bessel function $J'_l(x)$. In this way, we now can denote cylindrical modes by 3 indices, lmn , with l representing the number of half wavelengths in the ϕ direction, m for the radial direction, and n for the axial direction. The frequencies of this empty cylindrical cavities' modes will then be:

$$f_{lmn}^{TM(TE)} = \frac{1}{2\pi\sqrt{\mu\epsilon}} \sqrt{\frac{x_{lm}^{(')2}}{R^2} + \frac{n^2\pi^2}{L^2}} \quad (6.21)$$

The lowest frequency mode then for TM modes is the TM_{010} because $x_{01} \approx 2.405$, is the lowest value root, and less than π for the $n = 1$ case. This is referred to as the fundamental TM mode for the cavity, and, remembering Chapter 3, is the main search mode for axion searches because of its axion form factor. In the case of TE modes, notice that $n \neq 0$ for this solution to work, and the lowest root, x'_{11} , is $x'_{11} \approx 1.841$, thus the fundamental TE mode is TE_{111} with a frequency:

$$f_{111}^{TE} = \frac{1.841}{2\pi\sqrt{\mu\epsilon}R} \left(1 + 2.912 \frac{R^2}{L^2} \right)^{1/2} \quad (6.22)$$

Cavity Name	Radius	Height (cm)	F_{010}^{TM}	F_{111}^{TE}	F_{011}^{TM}	F_{020}^{TM} (GHz)
ADMX Main	21.0	101	0.547	0.444	0.567	1.25
ADMX Sidecar	4.60	12.1	2.50	2.28	2.79	5.73
ADMX run 2	8.5	100	1.35	1.04	1.36	3.1
ADMX EFR	6.4	100	1.79	1.38	1.8	4.12

Table 6.1: A summary of the empty cavity characteristics and base frequencies for the various ADMX resonator systems. In the case of the ADMX run 2 and EFR arrays, these are the measurements of a single empty cavity within the multi-cavity system. All lengths and frequencies are listed in units of centimeters and gigahertz.

This means that the ratio of cavity length to cavity radius determines whether the TM_{010} or TE_{111} is the lowest mode frequency in the cavity, and taking the ratio of the two frequency expressions, one finds that if $L \gtrsim 2.03R$, then $f_{111}^{TE} < f_{010}^{TM}$. This is an important note because we will want long length cavities that maximize our detection volume for a given R that sets the base f_{010}^{TM} search frequency, thus there may be many TE modes at lower frequencies before we come across the first TM modes. One doesn't want to increase this ratio too much, however, because it will increase the number of TE mode crossings for the TM_{010} as one shifts the TE spectrum lower. If one didn't have to maximize volume, then you could decrease the length of the cavity such that they are pushed far above the TM spectrum. I'll end this section with Table 6.1 that lists the base height and radius of the various ADMX cavities as well as their several important mode frequencies when they are empty. These frequencies are significantly lower than their loaded tuning rod configurations.

6.1.2 Identifying Modes

Now that we've formally derived the fields and frequencies of an empty cavity analytically, this can inform us on how one can identify resonant cavity modes by field structure and form factors. This is typically done in RF simulation software, such as Ansys HFSS, CST

microwave studio, or COMSOL, using an eigenmode solver package. A cavity geometry can be drawn or imported from a CAD file, and the resonant mode fields can be numerically solved for rather than analytically. One then typically plots several cross-sections of the field magnitude at every point in the interior cavity for a given mode, with \vec{E}^2 usually being the norm. This is done in Figure 6.1 for several important modes.

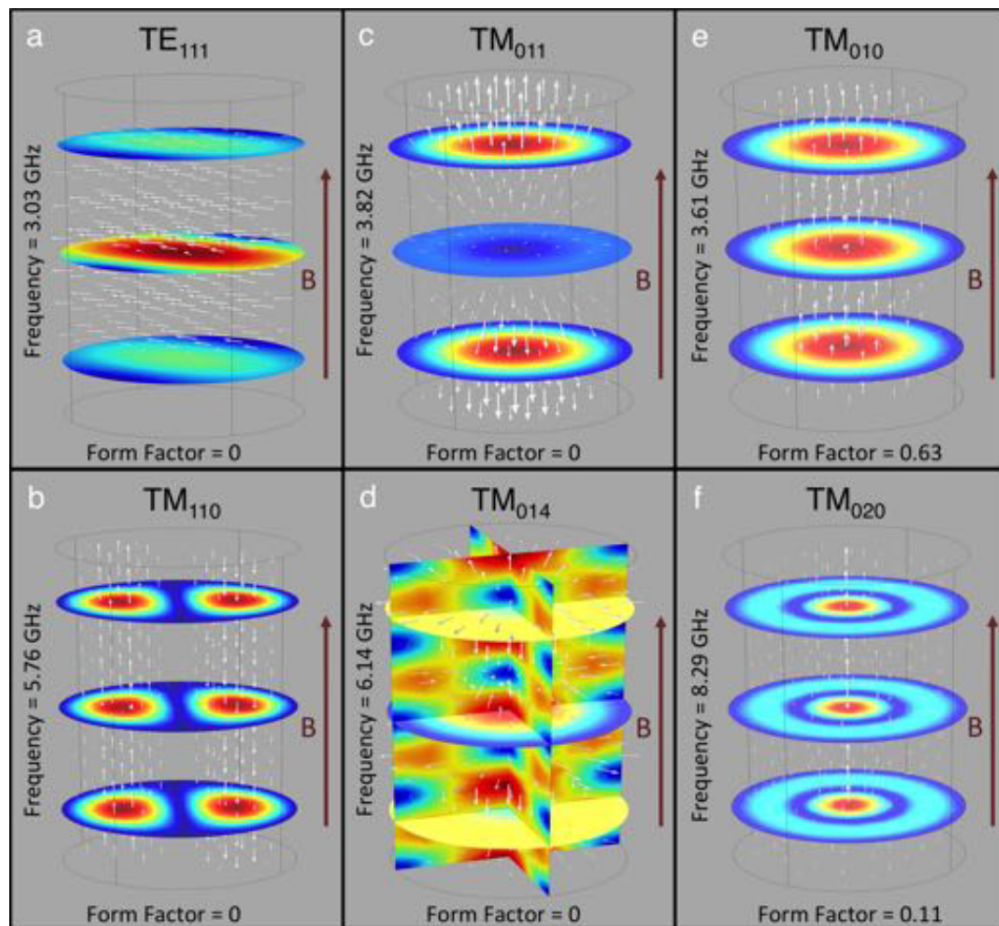


Figure 6.1: The electric field for different TM and TE modes in an empty Sidecar cavity [28]. The cross-sectional color plots show the magnitude of \vec{E} , while the vector plots represent \vec{E} , both for a single phase in the oscillation. In this way, red indicate extrema points in the electric field that can be used to label the mode.

The best way visually to assess a mode from its field structure is to count the number of extrema, representing the nodes of half wave-lengths, in each of the cylindrical coordinate

directions, ϕ , ρ , and z , which would directing correspond to l , m , and n respectively. Although this is easy for a human that can read plots and identify extrema by eye, its not the best for a computer. In this case, using a form factor that picks out the particular mode structure is a better route. In the case of the TM_{010} mode, the axion form factor (Equation 3.3) works perfectly because it is maximized for this mode and near-zero for most others. In fact, because TE modes have $E_z = 0$, it will be exactly zero for a perfect solenoid magnet ($\vec{B} = B_0\hat{z}$). TM modes such as the TM_{011} or TM_{110} may have $E_z \neq 0$ components, but they're opposing extrema will cancel each other once the integral over the volume is done. Thus, one can quickly see that the axion form factor is only non-zero for modes TM_{0n0} with increasing n corresponding to decreasing C_{0n0} . One can easily adapt this form factor to pick out all of the TM modes by squaring the numerator within the integral, such that:

$$C_{TM} = \frac{\int_V E_z^2 dV}{\int_V \vec{E}^2 dV} \quad (6.23)$$

This factor will be non-zero for all TM modes, and typically will decrease with increasing number of indices, lmn , due to more zero magnitude points in the field structure, but that is a trend rather than a hard rule. More importantly, because TE modes by definition have $E_z = 0$, this form factor will also be zero. In this way, the above form factor is great for separating the TM from TE and TEM mode spectra. Similarly, you could define a TE/TEM form factor based off the tangential field:

$$C_{TE/TEM} = \frac{\int_V \vec{E}_t^2 dV}{\int_V \vec{E}^2 dV} \quad (6.24)$$

However, this will simply switch for which modes the form factor is zero vs. non-zero. To distinguish TEM modes, one needs to ensure both E_z and B_z are zero for the mode, hence one could look at the parallel mode magnetic field form factor:

$$C_{TE} = \frac{\int_V B_z^2 dV}{\int_V \vec{B}^2 dV} \quad (6.25)$$

This form factor would be zero for TM and TEM modes, and would only be non-zero for TE modes, hence it can be used as a follow up after separating out TM modes via Equation

6.23. Be careful not to confuse this mode magnetic field with the external magnetic field used in the original axion form factor Equation 3.3.

Although the above derivation was for an empty, right cylindrical cavity, much of the analysis holds true for a cavity with tuning rods, other internal structures, or different geometries. The three mode indices will still be reflective of the number of half wavelengths in the base axes of the geometry. Introducing a tuning rod, there will still be TM_{lmn} -like and TE_{lmn} -like modes that are perturbations of the empty cavity fields amongst other modes that are possibly unique to the tuning rod geometry. In these cases, we are using simulation software to numerically solve and approximate the field solutions, so $E_z = 0$ might not be exactly true for what is obviously a TE mode by eye, but using the TM form factor Equation 6.23 as a follow-up check should produce a near-zero result most of the time.

6.2 *Quality factor*

In this section, I will cover the many definitions and ways of measuring cavity quality factor. Several definitions have already been shown in Chapter 3, such as Equations 3.4, 3.18, and 3.27, but hopefully this section will illuminate the connection between these equations and add a few more. This will start with some cavity quality factor basics for a classical copper cavity, then a derivation of the Lorentzian shape a resonant mode makes when taking a power spectrum of the cavity, which will highlight the connection between 3.4 as the physical definition of Q and 3.27 as how we measure it. Since actually measuring Q in lab will require an understanding of antennae and coupling, I reserve a later section for this process.

6.2.1 *Quality factor Basics*

The physical definition of quality factor is Equation 3.4 in terms of U , P_d , and ω_0 . By conservation of energy, the time rate of change in U should be equal to the negative of P_d , which allows us to solve for $U(t)$:

$$\frac{dU}{dt} = -\frac{\omega_0}{Q}U \quad (6.26)$$

which easily solves for:

$$U(t) = U_0 e^{-\omega_0 t/Q} \quad (6.27)$$

In this way, one can see that an initial U_0 just decays away with a constant inversely proportional to Q . We can write out Q using the field integral expressions for U and P_d as shown:

$$Q = \frac{\omega_0 \mu \int |\vec{H}|^2 dV}{R \int |\vec{H}|^2 dS} \quad (6.28)$$

we could have written U in terms of the time-average electric field as well. This equation is often grouped in terms of a geometric factor that absorbs all the cavity mode geometric information, and separates it from the material properties of the cavity, R . Therefore, cavity geometric factor is written:

$$G = \frac{\mu \omega_0 \int |\vec{H}|^2 dV}{\int |\vec{H}|^2 dS} \quad (6.29)$$

Where V is over the total cavity interior volume, and S is over the total interior surface. Quality factor in this case is simply:

$$Q = \frac{G}{R} \quad (6.30)$$

One can then write the resistance in terms of a classical surface conductivity:

$$R = \sqrt{\frac{\mu_c \pi}{\sigma}} \quad (6.31)$$

Where μ_c is the permeability of the cavity material, which is essentially μ_0 for copper. Often the skin depth is the preferred parameter for thinking about the finite conductivity of normal metals like copper. Effectively its the length parameter measuring the penetration of the cavity field into the conductor before its fully shielded by internal currents; a perfect conductor has no field penetration and therefore zero skin depth. Skin depth of a normal conductor like copper is written:

$$\delta_{norm} = \sqrt{\frac{2}{\mu_c \omega_0 \sigma}} \quad (6.32)$$

If we pullout a factor of V/S and μ from the geometric factor expression 6.29, another common expression for Q in terms of skin depth can be derived:

$$Q = \frac{\mu}{\mu_c} \left(\frac{V}{S\delta} \right) \times (\text{Geometry Factor} \approx 1) \quad (6.33)$$

It is important to note that this geometry factor is not the same as the geometric factor in Equation 6.29; this expression is intended to illuminate the importance of volume to surface ratio and skin depth in the base cavity design, but its not great for serious calculations. Using the field results from the right cylinder, the geometric factor can be solved analytically, and the result expressed in terms of skin depth, R , L , and lmn mode indices. This assumes no other external sources of power dissipation from antennae or other radiative losses.

$$Q_{lmn} = \frac{\lambda \sqrt{x_{lm}^2 + (n\pi R/L)^2}}{\delta \cdot 2\pi(1 + 2R/L)} \quad (6.34)$$

where λ is the free space wavelength of the cavity mode.

It is important to note that although the Equation 6.32 is valid for room temperature, at cryogenic temperatures the skin depth will start to shrink and conductivity will increase. When the mean free path of electrons in the metal exceeds the skin depth, surface electron currents will be more effected by the surface finish of the cavity surface itself, enough that Equation 6.32 is no longer valid [28, 66]. The 'anomalous' skin depth of this new regime can be written [61]:

$$\delta_{anom} = \left(\frac{\sqrt{3}c^2 m_e v_F}{8\pi^2 \omega_0 n e^2} \right)^{1/3} \quad (6.35)$$

where m_e is the electron mass, e is the electron charge, v_F is the fermi velocity of the material, n is the conduction electron density. One can estimate most of these parameters for copper to write the anomalous skin depth as a function of frequency; using $v_F = 1.57 \times 10^8$ cm/s, $n = 8.50 \times 10^{22}$ cm⁻³, and the typical electron mass and charge. This expression simplifies to:

$$\delta_{anom,Cu} = 2.84 \times 10^{-4} \text{ m} \left(\frac{1}{f} \right)^{1/3} \quad (6.36)$$

where frequency is in Hz. Often simulation software only allows for the adjustment of conductivity, assuming the classical regime for skin depth, therefore its often useful to represent this anomalous skin depth as an effective classical conductivity. This also allows us to use Equation 6.31 to assign a conductivity to a given R , and subsequently a given Q . The

effective anomalous conductivity would be:

$$\sigma_{eff,anom,Cu} = \left(\frac{2}{\mu_c 2\pi f} \right) \left(\frac{1}{2.84 \times 10^{-4} \text{ m} \left(\frac{1}{f} \right)^{1/3}} \right)^2 = 3.14 \times 10^{12} \text{ S/m} \cdot f^{-1/3} \quad (6.37)$$

This means the conductivity is from $\sigma \approx 3.95 - 2.41 \times 10^9 \text{ S/m}$ between $0.5 - 2.2 \text{ GHz}$ at cryogenic temperatures, as opposed to a room temperature value of $\sigma = 5.8 \times 10^7 \text{ S/m}$. The important takeaway is copper and normal metals will start to degrade in quality factor if used for higher and higher frequency searches. Between $1 - 10 \text{ GHz}$, the quality factor drops by $\approx 42\%$ for copper, slowing down the potential scan rate. By switching to a superconducting material with even marginal improvement in conductivity could avoid this degradation issue.

6.2.2 Derivation of the Lorentzian in a RLC circuit

Let us return to the equivalent circuit picture. Figure 3.2, that was used to calculate axion signal power in a cavity, and where we derived Equation 3.18. We can define the impedance, Z , of an *RLC* network at an arbitrary frequency, ω , to be:

$$Z = R + i \left(\omega L - \frac{1}{\omega C} \right) \quad (6.38)$$

This means at the resonant frequency $\omega_0 = 1/\sqrt{LC}$, the imaginary part vanishes. This impedance can be rewritten:

$$Z = R + iL \left(\frac{\omega^2 - \omega_0^2}{\omega} \right) \quad (6.39)$$

We can define the de-tuning of the arbitrary frequency from the resonant frequency as $\Delta = \omega - \omega_0$. If we make the safe assumption that $\omega_0 \gg \Delta$, the impedance then reduces further to:

$$Z = R + 2iL\Delta \quad (6.40)$$

This assumption means we are looking within the vicinity of the resonant mode. The magnitude of this complex impedance will then be:

$$|Z|^2 = R^2 + 4\Delta^2 L^2 \quad (6.41)$$

Even though it will be complex off resonance, only the real part of this impedance will contribute to the power dissipated in this circuit, V^2/R . Playing with the complex numbers a bit, one can determine the $\cos \theta$ of Z would be equal to $1/\sqrt{1 + 4\Delta^2(L/R)^2}$. The power dissipated can then be determined to be:

$$P(\Delta) = \text{Re}\left(\frac{V^2}{Z}\right) = V_0^2 \text{Re}\left(\frac{1}{Z}\right) = V_0^2 \frac{1}{|Z|} \cos \theta = \frac{V_0^2}{R} \frac{1}{1 + 4\Delta^2 \left(\frac{L}{R}\right)^2} \quad (6.42)$$

This is the Lorentzian power spectrum we are looking for. Nonetheless, this expression can be cleaned up a bit by recognizing two things; V_0^2/R is simply $P(f_0)$ the power dissipated on resonance. Secondly, the Full-Width-at-Half-Maximum (FWHM) can be recognized, $\Delta f_c = R/L$, within the Lorentzian. Using Equation 3.27, we can then re-write this expression in terms of Q and frequency:

$$P(f) = \frac{P(f_0)}{1 + 4Q^2 \left(\frac{f-f_0}{f_0}\right)^2} \quad (6.43)$$

Note that an assumption underlying this spectrum function is that the RLC network is being driven by a voltage source, $V_0 e^{i\omega t}$; One needs an external signal source to excite the mode. By inputting a series of signals at varying f about f_0 , this peak-shape can be constructed. This will be how one ultimately measures Q in the lab; by analyzing the shape, specifically the width, of this peak.

6.3 Geometric factors for multiple surfaces

Geometric Factors are critical for understanding how the different sources of power loss in the cavity breakdown amongst its design components. Equation 6.29 describes the total surface geometric factor for a given mode. One of the novel techniques of this dissertation is extending this definition to sub-surfaces of the cavity. The motivation being for estimating and understanding hybrid, multi-material cavities' quality factor. A hybrid cavity will have different resistances for each material's surface as well as each surface having different overlap with the mode fields. This can be done by breaking up the surface integral of Equation 6.29 into a sum of surfaces, $S = S_1 + S_2 + \dots + S_N$, we then define a sub-surface geometric factor

as:

$$G_n = \frac{\mu\omega_0 \int |\vec{H}|^2 dV}{\int |\vec{H}|^2 dS_n} \quad (6.44)$$

One can then write the total geometric factor in terms of its sub-surface factors as:

$$\frac{1}{G_{total}} = \frac{1}{G_1} + \frac{1}{G_2} + \dots + \frac{1}{G_N} \quad (6.45)$$

We can then write an associated sub-surface quality factor that takes into account the individual resistance of that surface:

$$Q_n = \frac{G_n}{R_n} \quad (6.46)$$

Similarly, one combines these sub-surface quality factors in the same way as geometric factors:

$$\frac{1}{Q_{total}} = \frac{1}{Q_1} + \frac{1}{Q_2} + \dots + \frac{1}{Q_N} \quad (6.47)$$

This represents a sum of resistance weighted by the geometric factors. It is not $R_{total} \neq R_1 + R_2 + \dots + R_N$, but some linear combination of them with weights C_1, C_2, \dots, C_N . If we take Equations 6.45 and 6.47 and input them into 6.30, one can solve for the weights to be:

$$C_n = \frac{\frac{1}{G_n}}{\sum_{i=1}^N \frac{1}{G_i}} \quad (6.48)$$

The total resistance for a given mode could then be written down in terms of these weights and individual resistances of the cavity surfaces:

$$R_{total} = C_1 R_1 + C_2 R_2 + \dots + C_N R_N \quad (6.49)$$

In the simplest case, $S = S_1 + S_2$, the weights are:

$$C_1 = \frac{G_2}{G_1 + G_2}, \quad C_2 = \frac{G_1}{G_1 + G_2} \quad (6.50)$$

This two dimensional case is important because it can be used to compare any specific feature on the cavity with the remainder of the total surface area in terms of resistance contribution. By using Equation 6.45 for different combinations of sub-surfaces, one can always reduce the problem to just two dimensions by combining surfaces back together. In fact, one can write

pretty easily a python function that can combine geometric factors in iterative steps of two surfaces for a set of N surfaces to result in any composite two surfaces. For instance I could combine the first half of the set together ($l = 1 + 2 + 3 + \dots + N/2$) and the second half ($k = N/2 + (N/2 + 1) + \dots + N$):

$$\begin{aligned}
 \frac{1}{G_{total}} &= \left(\frac{1}{G_1} + \frac{1}{G_2} \right) + \left(\frac{1}{G_3} + \frac{1}{G_4} \right) + \dots + \left(\frac{1}{G_{N-1}} + \frac{1}{G_N} \right) \\
 &= \left(\frac{1}{G_{1+2}} + \frac{1}{G_{3+4}} \right) + \left(\frac{1}{G_{5+6}} + \frac{1}{G_{7+8}} \right) + \dots + \left(\frac{1}{G_{(N-3)+(N-2)}} + \frac{1}{G_{(N-1)+N}} \right) \\
 &\dots \tag{6.51} \\
 &= \left(\frac{1}{G_{1+\dots+(N/4)}} + \frac{1}{G_{(N/4+1)+\dots+(N/2)}} \right) + \left(\frac{1}{G_{(N/2+1)+\dots+(3N/4)}} + \frac{1}{G_{(3N/4+1)+\dots+N}} \right) \\
 &\qquad\qquad\qquad = \frac{1}{G_l} + \frac{1}{G_k}
 \end{aligned}$$

If starting with an odd N you will instead combine the $N - 2$ and $N - 1$ factors and carry the N th factor into the next iteration, and it will eventually be combined. Since the combined expression for two G factors is a simple geometric expression, $\frac{G_1 G_2}{G_1 + G_2}$ this avoids using more complex solving methods. Note that I could have picked any l or k list of surfaces to combine, and just re-arranged terms. In fact, often l will be a single surface index while $k = \sum_{j \neq l} n$, when we are comparing a single surface to compare to the rest of the cavity.

These geometric factors are usually calculated using RF simulation software like Ansys HFSS or COMSOL, as mentioned earlier for identifying mode structure. The eigen-mode solver field data can be numerically integrated to calculate the field volume and surface integrals in the geometric factor as well as the frequency for that mode. One can define many sub-surfaces to calculate a G value for within the model. The most typical features are each end cap and the walls/tube of both the cavity and tuning rod, totally to 6 surfaces. Often one will also define antenna port regions as well. These geometric factors are then usually exported as a data file to be manipulated with python scripts along with each mode frequency and mode identifying form factors.

6.4 A simple hybrid cavity model

For several reasons discussed in Chapter 7, it is not necessarily ideal for some surfaces of a superconducting cavity to be made of superconductor when operating in high magnetic fields, and a hybrid cavity with those surfaces made of copper instead can be preferable. We can consider this 2-material hybrid cavity case without any superconducting physics to start; Just consider that each material has an associated surface resistance, R_{s_1} and R_{s_2} , applied to the surfaces that are made of that material. To simplify this further, since we know we will be using a well-known normal metal, most likely copper, we will use this as the reference resistance value, and judge the resistance of the superconducting metal such that $R_{SC} = \kappa R_{Cu}$, where κ is simply the ratio of the two resistances. We can write down the Q of this simply hybrid cavity, by considering the geometric factors of the surfaces made of each material, G_{Cu} and G_{SC} ; we will assume there is minimal dissipation from other surfaces such that $S_{tot} = S_{Cu} + S_{SC}$.

$$\frac{1}{Q_{hybrid}} = \frac{R_{Cu}}{G_{Cu}} + \frac{R_{SC}}{G_{SC}} \quad (6.52)$$

If we consider the Q of this same cavity but made entirely of copper, one can find a convenient expression for the Q ratio of hybrid to a normal copper cavity in terms of the resistance ratio $\kappa = R_{SC}/R_{Cu}$ and geometric factor ratio, $g = G_{Cu}/G_{SC}$:

$$\frac{1}{Q_{Cu}} = \frac{R_{Cu}}{G_{Cu}} + \frac{R_{Cu}}{G_{SC}} \quad (6.53)$$

$$\frac{Q_{hybrid}}{Q_{Cu}} = \frac{1 + g}{1 + \kappa g} \quad (6.54)$$

One can then also solve for the κ from a given Q ratio, $q = Q_{Cu}/Q_{Hybrid}$, if one has a measurements from each cavity type:

$$\kappa = q + \frac{q - 1}{g} \quad (6.55)$$

The ideal case for the hybrid cavity is when $\kappa = 0$ and therefore the superconducting surfaces have no resistance at all, giving Q its maximum boost:

$$\frac{Q_{ideal\ hybrid}}{Q_{Cu}} = 1 + g \quad (6.56)$$

Therefore, the hybrid cavity improvement factor is limited by the geometric factor even in the best case scenario. Any other value of κ from 0 to 1 will increase the Q, albeit by not as much, and predictably, any value where $\kappa > 1$ will decrease the Q, because the enhancing metal is more resistive than the copper reference metal. Superconductors are often terrible normal metal conductors at room temperature, and much more resistive than copper at room temperature, so hybrid cavity will have a worse Q than an all copper version until it transitions, so this formulation is useful for tracking that as is done in Chapter 9 for the ADMX Sidecar cool-down. In this way, one expects the κ to drop as it cools from being order 10^2 or more at room temperature, to $\kappa = 1$ as the lower bound for the cavity starting to work as intended, up to $\kappa \rightarrow 0$. Often, a $\kappa \approx 0.1$ is already well-within 90% of the ideal hybrid Q value.

6.5 Multi-mode decomposition

The previous section on sub-surface geometric factors isn't very useful in lab for a single cavity mode, because all one can really measure is the total resistance for any given mode. Typically, we can only estimate the sub-surface resistances based off models of each cavity materials' properties. But if we have multiple modes, backed up with geometric factors for each, we can construct multiple equations of the form in 6.49. Each equation forms a row $R_m = C_{m1}R_1 + C_{m2}R_2 + \dots + C_{mN}R_N$, for a single mode indexed m instead of lmn for simplicity. One can imagine constructing a matrix equation from this data:

$$\vec{R}_m = C_{m,s}\vec{R}_s \quad (6.57)$$

where \vec{R}_m is the list of multiple mode resistances corresponding to $\frac{G_m}{Q_m}$, the total surface geometric factor and total surface quality factor for the mode. The \vec{R}_s is the list of corresponding surface resistances for each sub-surface defined; note that this relation assumes R of a surface is independent of mode, while in reality this might not be true. Different modes can have very different frequencies, making the resistance of the material also very different. It was already shown for copper and other normal metals that there is some frequency

dependence to account for in Equation 6.37. Nonetheless, if they are similar frequencies or the frequency dependence can be accounted for, this is still useful. If one can get N mode measurements of Q , forming a list \vec{R}_m , with a corresponding rectangular matrix, C_{ms} with N sub-surfaces, this equation can be inverted, and can solve for a corresponding N -length list of \vec{R}_s :

$$\vec{R}_s = C_{ms}^{-1} \vec{R}_m \quad (6.58)$$

where C_{ms}^{-1} is the inverse of the weight matrix. This will also only hold true if an inverse exists for C_{ms} , therefore $\det(C_{ms})$ must not be zero. In the 2x2 case, this occurs if $G_{1i} = G_{2i}$, when the modes do not have distinct geometric factors. For higher dimensions, this pattern will hold true if all modes are identical, $G_{1i} = G_{2i} = \dots = G_{Ni}$, but there also may be some non-trivial solutions. Nonetheless, this suggests the best modes to use for performing this decomposition should have distinct geometric factors for the set of surfaces you are using.

6.5.1 Solving for the 2 mode decomposition case

The entirety of this dissertation's actual analysis will only make use of the 2-mode, 2-surface case of decomposition. This is because finding modes that have both similar frequencies and distinct geometric factors can be tough. Often modes of similar field structure will have similar frequencies but therefore not have distinct geometric factors. Also as we will see in the error analysis section, even if they are distinct, if the uncertainties in G overlap, error in the inverse matrix will greatly increase.

In this case, there will be four geometric factors, G_{11}, G_{12}, G_{21} , and G_{22} , and two Q measurements Q_1 and Q_2 . First the Q measurements are converted to an \vec{R}_m by using the total mode geometric factors:

$$G_n = \frac{G_{n1}G_{n2}}{G_{n1} + G_{n2}} \quad (6.59)$$

where $n = 1, 2$ for each of the two modes. The \vec{R}_m is then:

$$\vec{R}_m = \begin{bmatrix} G_1/Q_1 \\ G_2/Q_2 \end{bmatrix} \quad (6.60)$$

The weight matrix in this case comes out to be:

$$C_{ms} = \begin{bmatrix} \frac{G_{12}}{G_{11}+G_{12}} & \frac{G_{11}}{G_{11}+G_{12}} \\ \frac{G_{22}}{G_{21}+G_{22}} & \frac{G_{21}}{G_{21}+G_{22}} \end{bmatrix} \quad (6.61)$$

The determinant of this weight matrix is:

$$\det(C_{ms}) = \frac{G_{12}G_{21} - G_{11}G_{22}}{(G_{11} + G_{12})(G_{21} + G_{22})} \quad (6.62)$$

Notice this is where we get the condition that geometric factors must be unique; If I use the same mode for two different surfaces or the same surface for two different modes, this numerator will be zero. The final inverted weight matrix is then:

$$C_{ms}^{-1} = \frac{1}{G_{12}G_{21} - G_{11}G_{22}} \begin{bmatrix} G_{21}(G_{11} + G_{12}) & -G_{11}(G_{21} + G_{22}) \\ -G_{22}(G_{11} + G_{12}) & G_{12}(G_{21} + G_{22}) \end{bmatrix} \quad (6.63)$$

We then multiply \vec{R}_m by this inverted weight matrix to get the resultant \vec{R}_s of length 2 corresponding to the resistances of the two surfaces we defined. This method will be used several times Chapters 8 and 9.

6.5.2 Error analysis for decomposition technique

In addition to all of the multi-mode decomposition technique described above is the associated analysis of propagating the error into a final surface resistance. The error comes from two sources; the quality factor measurements, and geometric factor simulations. The quality factor uncertainty is going to come from laboratory measurement usually, which is fitted to a Lorentzian, of which the process is described later in this chapter. Typically a $\Delta Q/Q \approx 1 - 5\%$ is very achievable. The geometric factor error is a little harder to estimate; theoretically it comes from the expected error within the dimensions of your cavity. This means any vibrations of the cavity that cause the radius, length, or other dimensions to significantly change would effect the geometric factor. However, what is usually more dominant than vibrations is machining tolerances of the parts themselves, which are typically anywhere from 0.001"-0.010" depending on the quality. The G factors are then calculated for

variations of the geometry within this chosen length uncertainty Δx ; The simulation software calculates the G factors for all combinations of each cavity dimension varied by 3 points, $x_0 - \Delta x, x_0, x_0 + \Delta x$. Computationally, this can be quite intensive, because the number of cavity design arrangements will be 3^D where D is the number of dimensions in the cavity design; for the Sidecar cavity in Chapter 9, 5 dimensions were incorporated, resulting in 243 arrangements. From there, one has a set of G factors for each cavity design arrangement. What's important is that each design corresponds to a set of corresponding geometric factors; one cannot simply average each geometric factor together, but has to find the most average set of geometric factors corresponding to a real design. To do this, we must calculate the average set of geometric factors, and then find the real set that is the least different. Each term of \vec{G}_{avg} can be defined by $G_{i,avg} = (1/N) \sum_n^N G_{i,n}$ where N is the number of cavity design arrangements. The most average design arrangement minimizes the magnitude of the vector difference with \vec{G}_{avg} :

$$\Delta \vec{G} = \vec{G}_j - \vec{G}_{avg} \quad (6.64)$$

The set we then choose to represent the average real set, even though not strictly a median, is:

$$\vec{G}_{median} = \vec{G}_j(\min(|\Delta \vec{G}|^2)) \quad (6.65)$$

On the other end, the set with the most error would be the design arrangement that maximized $|\Delta \vec{G}|^2$, $\Delta \vec{G}_{max}$. The final estimated uncertainty in the geometric factors is determined by taking the difference in these two vectors that represent real cavity designs and are all correlated with each other.

$$\Delta \vec{G}_{final} = |\vec{G}_{max} - \vec{G}_{median}| \quad (6.66)$$

This produces typically a fractional error, $\Delta G/G \approx 5\%$. This list of ΔG 's can then be combined with errors in quality factors, propagating through the analysis steps of a multi-mode decomposition. The total mode resistance error is simply the fractional errors of G

and Q combined in quadrature:

$$\left| \frac{\Delta R_m}{R_m} \right| = \sqrt{\left(\frac{\Delta G_m}{G_m} \right)^2 + \left(\frac{\Delta Q_m}{Q_m} \right)^2} \quad (6.67)$$

where $G_m = \sum_s G_{ms}$, the total mode geometric factor, and its corresponding error is $\Delta G_m = \sqrt{\sum_s \Delta G_{ms}^2}$. Up until now, this error analysis hasn't assumed a dimension. But going forward one has to solve for the error in each weight index, C_{ms} , which depends on the matrix dimension. For simplicity, I will assume the 2x2 case to calculate ΔC_{ms} . To make things easier to read, I will denote each index in the matrix, $C_{ix} = \frac{G_{iy}}{G_{ix} + G_{iy}}$, where if $x = 1$, then $y = 2$ and vice-versa:

$$\left| \frac{\Delta C_{ix}}{C_{ix}} \right| = \sqrt{\left(\frac{\Delta G_{iy}}{G_{iy}} \right)^2 + \frac{\Delta G_{ix}^2 + \Delta G_{iy}^2}{(\Delta G_{ix} + \Delta G_{iy})^2}} \quad (6.68)$$

Taking the inverse of the 2x2 matrix gives us an expression for ΔC_{ms}^{-1} in terms of error in C_{ms} components as well as error in the determinant of the matrix, $\Delta(\det(C))$, which can be written:

$$\left| \frac{\Delta(\det(C))}{\det(C)} \right| = \sqrt{C_{11}^2 \Delta C_{22}^2 + C_{22}^2 \Delta C_{11}^2 + C_{12}^2 \Delta C_{21}^2 + C_{21}^2 \Delta C_{12}^2} \quad (6.69)$$

$$\left| \frac{\Delta C_{ix}^{-1}}{C_{ix}^{-1}} \right| = \sqrt{\left(\frac{\Delta C_{xy}}{C_{xy}} \right)^2 + \left(\frac{\Delta(\det(C))}{\det(C)} \right)^2} \quad (6.70)$$

I use the indexing xy here, because these are no longer specifically mode and surface indexes after the inversion. Note that if the determinant is near zero, it will inflate this inversion error significantly. This can make or break your ability to perform a decomposition because error bars can quickly grow to multiples of the resistance values you are trying to solve for. Later measurements will show that even for good mode combinations, the fractional error can commonly be about 10% for the inverse matrix. Since this expression is entirely dependent on the simulated G factors, one can check the fractional error in the determinant for a variety of mode combinations to find pairs that work well together. The final step is to perform the matrix multiplication with \vec{R}_m to make \vec{R}_s , so the error will incorporate ΔC_{xy}^{-1}

as well as ΔR_m

$$\left| \frac{\Delta R_s}{R_s} \right| = \sqrt{\sum_{x=1}^2 (R_x^2 (\Delta C_{sx}^{-1})^2 + (C_{sx}^{-1})^2 \Delta R_x^2)} \quad (6.71)$$

This covers all the equations needed to propagate the uncertainties for a 2-mode decomposition. See chapters 8 and 9 for examples of this method in action.

6.6 Antenna coupling

Up until now, this analysis has assumed a closed cavity with no antennae or losses to the outside. In reality, one needs at least a single antenna to pull some power out of the cavity to detect a resonant mode in the lab. An antenna simply acts as another power loss source for the cavity. The complication is we are actually measuring the power deposited in this loss source from the cavity, not the intrinsic power of the cavity. The more coupled an antenna is to the cavity, the more power will be drawn from the cavity into the antenna, which in turn will spoil the cavity quality factor with a high P_d . This is referred to as the "loaded" cavity quality factor Q_L vs. the "unloaded" or intrinsic Q of the cavity, Q_0 , that is independent of antenna coupling.

6.6.1 Calculating coupling coefficient

Similarly to how we defined sub-surface quality factors for dealing with multiple surfaces, we can define a Q_e that is the associated Q of the antenna within the cavity, which implies the loaded quality factor would be:

$$\frac{1}{Q_L} = \frac{1}{Q_0} + \frac{1}{Q_c} \quad (6.72)$$

This expression implies that if $Q_c \gg Q_0$, it will have a small effect on the Q_L , whereas if $Q_c \leq Q_0$, it will be a greater contribution. To characterize this more precisely, one defines the coupling coefficient:

$$\beta = \frac{Q_0}{Q_c} \quad (6.73)$$

This creates a convenient expression for Q_0 , the intrinsic cavity value, in terms of Q_L and β , the measurable values:

$$Q_0 = Q_L(1 + \beta) \quad (6.74)$$

This means if $\beta \approx 0$, that antenna coupling doesn't have a significant effect on Q and the measured $Q_L \approx Q_0$.

6.6.2 Antenna Design




Figure 6.2: A schematic of a coaxial antenna inserted into a cavity antenna port. The antenna itself only has two major parameters, L_{tip} , the length of exposed inner conductor and L_{outer} , the length of outer conductor before the coaxial connector. Based on the width of the cavity port, W_{port} , and the length of exposed outer conductor outside the cavity, d_{out} , the insertion depth, d_{in} , can be calculated.

ADMX typically uses a simple coaxial antenna design for its searches because it couples well to the electric field of the TM_{010} mode. Since the TM_{010} has an axial electric field,

an axially oriented segment of the conductor will have surface currents moving along its length, which are then transferred into a coaxial cable and subsequent transmission line. This obviously means putting the antenna port on an end cap of the cavity for the right orientation. Since this field is maximal towards the cavities' radial center, there is somewhat of a preference to put the antenna port towards the center, but it tends not to be that important. Figure 6.2 outlines what this antenna setup looks like: A semi-rigid coaxial cable is cut such that a small length of the inner conductor is exposed, L_{tip} , leaving a length, L_{outer} , of the outer conductor before connectorized, usually with an SMA-type connection. This antenna is inserted into the port with a threaded, silver-plated feed-through to minimize any radiation losses through gaps in the port. This also grounds the outer conductor to the outer cavity itself. Based on the width of the end plate, W_{port} , the port is drilled through and the antenna is inserted with an exposed length of outer conductor above, d_{out} , and the insertion depth of the antenna can be estimated, d_{in} , which is useful for quantifying geometrically how far the antenna is inside of the cavity. This parameter can also be used for simulating antenna coupling.

6.6.3 1-port Measurements

Measurements with a single antenna into the resonant cavity can be classified as a 1-port measurement. Up until now, I have not formally introduced the primary instrument for making cavity measurements in the lab: The network analyzer. A network analyzer is essentially two instruments working in conjunction with one another: A signal generator makes known RF signals to input into the "RF network", our antenna and cavity setup, via one of its ports, while a spectrum analyzer measures the power spectrum coming from that network into a network analyzer port. In the case of 1-port measurements, the network analyzer will only have 1 port in use, but typically they come with 2 ports. Fig 6.3 is a simple schematic of this setup. SMA connections are typical, and connect the transmission line, coaxial cable, to the network analyzer and antenna.

The only measurement one can make in this set up is known as reflection measurement.




Figure 6.3: A schematic of a typical 1-port cavity measurement with a vector network analyzer (VNA).

This is a swept measurement, where a fixed power signal is generated and sent down the transmission line at various frequencies (this is often referred to as a "sweep"), and the resultant amount of power that is reflected back down the input line is measured with the spectrum analyzer. The input signal will not only be attenuated in magnitude, but there could be an associated phase shift in the signal as well, therefore the response will be a complex value Γ at each frequency. The ability to measure this phase information as well as magnitude makes this network analyzer a vector network analyzer (VNA). One could do this measurement without the input signal, but it won't necessarily have the same power being inputted at each frequency hence $\Gamma(f)$ will not be accurate. Nonetheless, an axion search digitization simply removes the input signal and looks for ambient axion power exciting the modes (after one has characterized the cavity behavior). The magnitude represents the ratio of reflected to input power, $|\Gamma|$, a reflection coefficient. In this way, $\Gamma(f) = |\Gamma(f)|e^{i\angle\Gamma(f)}$ where $\angle\Gamma(f)$ is the associated phase shift for that frequency. Unless the input power is near a resonant frequency, most of the power will be reflected back up the transmission line resulting in a reflection ratio $|\Gamma| \approx 1$ or $S \approx 0$ dB. However, if the antenna is coupled to the cavity, some portion of the power will be absorbed by the cavity near the resonant mode

frequency, producing a Lorentzian dip! This will make $0 < |\Gamma| < 1$ and therefore S will be a negative dB value. This Lorentzian dip can be related to β , the coupling coefficient, by the relation:

$$\beta = \frac{1 + \text{sign}(\angle\Gamma(f_0) - \pi)|\Gamma(f_0)|}{1 - \text{sign}(\angle\Gamma(f_0) - \pi)|\Gamma(f_0)|} \quad (6.75)$$

Where f_0 is the mode resonant frequency and therefore $\Gamma(f_0)$ refers to the values of Γ at the bottom of the Lorentzian dip. The $\text{sign}(x)$ function simply outputs the sign of the argument: $\text{sign}(x) = +1$ for $x > 0$, $\text{sign}(x) = -1$ for $x < 0$, and $\text{sign}(0) = 0$. This is how one measures the β and couples antennae in lab; look at the mode reflection dip using a network analyzer, write down the $|\Gamma(f_0)|$ and $\angle\Gamma(f_0)$, and calculate β . One can see that from Equation 6.75, β is to some degree proportional to the dip depth, but if $|\Gamma(f_0)| \approx 1$, no power is coupled into the cavity, could respond to either $\beta = 0$ or $\beta \rightarrow \infty$, depending on the phase shift.

6.6.4 Coupling regimes

The coupling coefficient β can be used to characterize 3 regimes of antenna coupling. In Figure 6.4, a successive number of reflection measurements from a network analyzer are made on a cavity as an antenna is inserted at fixed length intervals. One starts on the right side of the plot where insertion depths are lowest and the antenna isn't very coupled to the cavity. If the antenna isn't inserted at all, one would expect a flat spectrum, and as it is inserted slowly, small dips would appear around the resonant frequencies of the modes. The coupling coefficient would be $\beta \approx 0$ for these first couple sweeps, when the dip is $> -1\text{dB}$, but will start to increase towards $\beta = 1$ as the dip gets deeper. This regime where $\beta < 1$, is known as being weakly coupled; the antenna is only a perturbation on the cavity. Notice that there is some downward shift in the resonant frequency as $\beta \rightarrow 1$, but it is less than 1 MHz over the range for a 1.35 GHz cavity.

When $\beta = 1$, the antenna will have roughly the same Q as the cavity itself, and theoretically will be taking half the power out of the cavity compared to the surface resistance losses. In terms of the reflection dip, this should be the maximum dip depth one sees; for




Figure 6.4: A series of network analyzer reflection measurements on a single ADMX run 2A cavity TM_{010} for varying antenna depths. They are indexed in the legend by how many thousandths of an inch the antenna is inserted into its port from the antenna tip. Additionally, a peak trend-line is interpolated across the series of Lorentzian dips, and a vertical line marks the critically coupled point.

$\beta > 1$, the antenna Q will dominate over the cavity Q , meaning a smaller fraction of inputted power will transmit into the cavity, lowering the dip depth. In Figure 6.4, you can see this maximum depth right around 1.356 GHz, marking the $\beta = 1$ insertion depth. This regime, where $\beta \approx 1$, is called critically coupled; typically β values near 1 are still considered critically coupled even out to $\beta = 2$ (more on this in the next section).

The third regime is known as being over-coupled, and is when $\beta > 1$. This is represented by the entire left side of Figure 6.4 past the critically coupled minimum. Not only does the dip depth degrade, but one can see that the width of the Lorentzian dip starts to widen significantly, indicative of the Q_L degrading from the antenna coupling. The antenna starts to be a considerable object within the internal cavity geometry in this regime, causing the resonant frequency to shift on the order of several MHz for this cavity.

6.6.5 Axion search antenna coupling

Equations 3.25 and 3.31, both include the effect of antenna coupling on axion power and scan rate, but we did not cover where these factors come from or its optimization. Clearly the cavity quality factor at a given frequency will degrade inversely, according to Equation 6.72, introducing the first factor of $\frac{1}{1+\beta}$, as β is increased. Increasing the coupling will also increase the fraction of power being pulled out of the cavity and into the digitizer; this introduces a factor β . The overall signal power is therefore enhanced by a factor of $\frac{\beta}{\beta+1}$, implying that an over-coupled antenna will draw the most power out of the cavity. What this doesn't include though is, by increasing β towards the critically coupled regime, one also effectively spreads out their cavity search bandwidth, δf_c by a factor of $1 + \beta$, which will allow one to take bigger frequency steps during an axion search. These two effects are combined in Equation 3.31, enhancing the scan rate by a factor of $\frac{\beta^2}{(\beta+1)^3}$. This factor can be differentiated to find critical points, one of which is a non-negative maximum at $\beta = 2$. For this reason, the optimal coupling for an axion search is $\beta = 2$, slightly over-coupled. When calibrating an antenna, its easiest to first critically couple it, and then incrementally increase it a tiny bit more to a $\beta = 2$; therefore its often said that an axion search is best performed with a critically coupled antenna, rather than the more nuanced $\beta = 2$ answer.

6.7 Measuring Q in the lab

This section will cover how one typically measures cavity Q in a lab using a network analyzer. In the antenna coupling section we covered 1-port measurements for determining a coupling coefficient of a single antenna, but typically one uses a 2-port setup for cavity characterization.

6.7.1 2-port Measurements

Adding a second port increases the number of possible measurements from 1 to 4; The signal generator can input power on either port 1 or port 2, and the spectrum analyzer can

measure the resultant signal at either port. This forms a 2×2 matrix known as the scattering parameters, or S-matrix, for the system, and they are indexed by the port measured and power inputted respectively; S_{ij} is produced from measuring at port i and inputting a signal at port j .

The two possible reflection measurements are S_{11} and S_{22} , indexed by each antenna port; They will each have an associated β_1 and β_2 . This doesn't change much to the analysis of the 1-port measurement except that Q will degrade more based on both antennae. A simple proof, following the analysis of how we defined β in terms of Q_0/Q_c , will show that $\beta_{total} = \sum_{i=1}^N \beta_i$ when combining N antenna ports, therefore in the 2 port case, $\beta_{1+2} = \beta_1 + \beta_2$. As outlined in Chapter 3, haloscopes typically run with a strong, variably coupled port antenna ($\beta_{strong} \approx 1 - 2$) that takes the axion search data, and a fixed, weakly coupled port antenna where $\beta_{weak} \approx 0$ for periodic cavity characterization. This means the strong antenna is the majority contributor to the total coupling most of the time.

In the 2-port scenario, transmission measurements are now possible by looking at the S_{21} and S_{12} parameters. In this case, the cavity will act like a filter for many frequencies; unless an input frequency is near a resonant mode, no power will be transmitted through the cavity from the input port to output port. This makes the transmission coefficient, sometimes called gain $|G|$, with linear units, near-zero off resonance. When the input signal frequency is in the vicinity of a cavity mode, the cavity will resonate and transmit the signal through, creating a non-zero $|G|$; this generates a Lorentzian peak on the transmission S-parameter plots, unlike the dip in reflection measurements. Since both of the antennae aren't necessarily coupled in the same way, S_{21} and S_{12} might not be reciprocal to each other. In the case of a strong vs weak port, a signal inputted from the weak port will not transmit as much of the signal into the cavity because β is much lower than the strong port. A strong port that is measuring the transmitted signal from the weak port will be more sensitive because of its higher β than if measuring with the weak port. In the case of a low power axion search, it is then best to input power from the weak port, minimizing the external RF power put into the cavity system, and measure with the strong port that is more coupled to the cavity. If one is

just characterizing a cavity, and doesn't need a critically coupled antenna, for simplicity, one often just sets both β_1 and β_2 to be as weakly coupled as possible, to minimize the antennae effecting the cavity system; when $\beta_1 = \beta_2$, assuming the antenna ports are also placed in symmetric positions to one another, the network should then be reciprocal and $|S_{21}| = |S_{12}|$.

The reason one wants to take transmission measurements is that they generally produce much more consistent Lorentzian shape to measure and calculate quality factor from. In principle, one could use the following methods on reflection data alone to calculate Q, but typically reflection measurements will contain more noise, standing waves, etc. from the transmission line and upstream RF components of the measuring port. Additionally, the Lorentzian dip changes drastically depending on the associated coupling coefficient, whereas a transmission measurement, by depending on two β values, can be less sensitive. This can be seen in Figure 6.5, which is the partner plot to Figure 6.4. The peak height in this case is much stabler than the dip depth in reflection measurements, and the widening of Δf_c is much easier to see in the weakly to critically coupled regime. However, one can see that picking out the critically coupled scan is much less intuitive in this plot compared to reflection measurements. For this reason, transmission measurements are used for getting a Q_L , and reflection measurements are used to get β as well as a secondary, less reliable $Q_{L,reflection}$. We are now in a position to talk about various methods of measuring Q.

6.7.2 3-dB method

This method is by far the simplest method, which often can be more stable than the more complex methods, but less robust. First, one wants to window their VNA scan such that the resonant peak or dip is the only feature. In this way, the maximum S_{21} or S_{12} for the scan corresponds to $G(f_0)$, the peak value at the resonant frequency. In terms of reflection measurement, the minimum S_{11} or S_{22} corresponds to $\Gamma(f_0)$. One can then simply find the frequency that has the maximum S value in the scan, and f_0 is determined. This can be tricky for reflection measurements that can more often have large standing waves that put




Figure 6.5: A series of network analyzer transmission measurements on a single ADMX run 2A cavity TM_{010} for varying antenna depths. In this case, port 1 is being inserted into the cavity, while port 2 is fixed in a weakly coupled position. They are indexed in the legend by how many thousandths of an inch the antenna is inserted from its tip. Additionally, a peak trend-line is interpolated across the series of Lorentzian peaks, and a vertical line marks the critically coupled point.

the minimum S-value within the noise region, and not on the resonant dip.

The Lorentzian shape defines Δf_c to be the Full-Width-at-Half Maximum (FWHM) of the resonant peak or dip of interest. Therefore the goal is to find the two points on either side of the peak (or dip) that correspond to half the peak value, $G(f_0)/2$, (or twice the dip value, $2 \times \Gamma(f_0)$). Because the S-parameters are usually expressed in dB units, rather than linear units, this means looking for the two points $\approx -3dB$ from the peak (or $+3dB$ from the dip). This can be done by splitting the original scan at the measured resonant frequency into left and right arrays, and looking for the frequency value that corresponds closest to $S_{21}(f_0) - 3dB$ (or $S_{12}(f_0) - 3dB$, $S_{11}(f_0) + 3dB$, $S_{22}(f_0) + 3dB$) in each array. From there, $\Delta f_{3dB} = |f_R - f_L|$ where f_R is the right side 3 dB frequency and f_L is the left side 3dB

frequency. From there the Q_{3dB} is simply:

$$Q_{3dB} = \frac{f_0}{\Delta f_{3dB}} \quad (6.76)$$

The main issue with this method is it doesn't take into account other factors within the Lorentzian equation, like the noise level, phase, etc. Additionally, there is some tiny error introduced by 3 dB not exactly corresponding to $\frac{1}{2}$ in linear units, but this can be corrected for by simply converting the set into linear units through out this whole process. If all 4 S-parameters are taken, theoretically you have 4 equivalent values of Q assuming a reciprocal network. The subsequent fitting methods often take the Q_{3dB} as the initial guess for the quality factor, and re-fit the value accordingly.

6.7.3 Lorentzian Fit

The next step up in complexity would be to perform a least-squares fit according to the power spectrum function in Equation 6.42. However, this function is adapted to have an arbitrary constant noise level, P_{noise} , usually. This fit then has 4 parameters to fit to; P_{noise} , P_{peak} , f_0 , and Q . If one has done the 3-dB method before this, they will already have on hand values for $P(f_0)$, f_0 , and Q which are convenient for giving a fitting routine software an initial guess. In this setup, one must be careful to recognize $P(f_0) = P_{peak} + P_{noise}$; this can even be reduced to a single parameter, ΔP , the prominence of the peak, but it still involves making some separate guess or fit to the noise level beforehand. The best initial guess for the noise level is usually done by taking the average of either end of the scan, which should be the most off-resonant values in the set. This can be finicky at times, because it is correlated with the width of the VNA scan. If the VNA scan is a bit narrow, it will over-estimate the noise level by incorporating part of the Lorentzian into its estimate. If the scan is too wide, one might not have enough resolution for fitting the Lorentzian itself. This can be corrected for by taking VNA scans at multiple window span sizes; one wider scan that captures the noise level estimate, and a narrow scan that captures the Lorentzian parameters. Because Q depends on f_0 , Δf is sometimes used instead of Q directly for the fit, and Q is

calculated subsequently from the fit results; this is just personal preference, and shouldn't make a difference. Typically, the python package we use is `Scipy.optimize.curve_fit`; this will not only output the fit parameters but a covariance matrix from the fit, that will give uncertainties for all the fit parameters, including Q . Although the 3 dB method may get you a value for Q very quickly, and even if this Lorentzian fit doesn't change the Q value by that much, this fit method produces a smooth and refined function for the cavity Lorentzian that can be used in subsequent analysis.

6.7.4 Complex Q -circle Fit

The final Q -fitting method that is used by the ADMX collaboration is a complex fit, and the most rigorous assessment of the cavity's performance. This method makes use of the shape of the complex plot of the S-parameters for the resonant cavity: a circle. This complex data is often plotted on a smith chart, which is a convenient tool for visualizing and measuring the related parameters to this measurement: impedance, admittance, reflection coefficients, etc. How this circle looks plotted in the complex plane also captures many of the exterior phenomenon to the cavity: phase delay and losses from transmission lines, connectors and antennae, impedance mismatch, and cross-talk [76, 77]. Each of these has an effect on shifting the Q -circle from the ideal case. Figure 6.6 represents the ideal case where none of these effects take place: The real and imaginary parts of the S-parameter near a resonance form a circle that passes through the origin at the edges of the scan (the detuned point), whereas the S-parameter at the resonant frequency passes through the real-axis at a point ($Re(S_{21}(f_{res})) \neq 0, Im(S_{21}(f_{res})) = 0$); notice the center of the Q -circle will also be along the real axis as well.

In practice, a Q -circle more often looks like Figure 6.7. The cross-talk between the two antennae (coupling loops) offsets the circle away from the origin. A phase shift due to the transmission lines between the VNA and resonator causes the Q -circle to be rotated about the origin. Losses in the connectors, cables, and antennae have the effect of reducing the diameter of the circle. Any noise in the data translates into the roughness of the circle; In




Figure 6.6: An ideal S_{21} Q-circle. In principle, an ideal reflection Q-circle is the same, however in practice it is a much noisier fit [77].

general, reflection Q-circles are much rougher than transmission measurements, hence why we use transmission measurements to measure quality factor.

These effects can be better quantified by expanding the equivalent circuit model for the resonator (Figure 3.2) to include the impedance of the two transmission lines, $Z_{si} = R_{si} + iX_{si}$, and a load representing the VNA itself, R_c , as pictured in Figure 6.8. The admittance, Y_{exp} , the reciprocal of impedance, as well as conductance, G_{exp} , the reciprocal of resistance, are useful to define here for ports $p = 1, 2$:

$$Y_{exp} = \frac{(R_c + R_{sp}) - iX_{sp}}{(R_c + R_{sp})^2 + X_{sp}^2} = G_{exp} + iB_{exp} \quad (6.77)$$

$$G_{exp} = \text{Re}(Y_{exp}) \quad (6.78)$$

$$B_{exp} = \text{Im}(Y_{exp}) \quad (6.79)$$

Additionally, the conductance of the cavity, $G_0 = 1/R_0$, can be used to define the coupling coefficients of each port:

$$\beta_p = \frac{G_{exp}}{G_0} \quad (6.80)$$




Figure 6.7: A practical S_{21} Q-circle fit that reflects the effects of the transmission line and antenna setup [77].




Figure 6.8: An equivalent circuit model of a two-port cavity measurement with antennae and coaxial cables connected to a VNA: an ideal resonator in parallel with two transmission lines with complex impedance [77].

It is worth noting that although $G_0 \neq Q_0$ exactly, $Q_0 \propto G_0$ and will only differ by a constant that is completely independent of the transmission lines. The cavity itself also will have an

admittance:

$$Y_0 = G_0 \left[1 + i2Q_0 \frac{\omega - \omega_0}{\omega_0} \right] \quad (6.81)$$

The $ABCD$ matrix formalism described in most microwave engineering textbooks, such as Ref. [96], is useful for describing this circuit as 3 cascaded networks:

$$\begin{bmatrix} A & B \\ C & D \end{bmatrix} = \begin{bmatrix} 1 & R_{s1} + iX_{s1} \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ Y_0 & 1 \end{bmatrix} \begin{bmatrix} 1 & R_{s2} + iX_{s2} \\ 0 & 1 \end{bmatrix} \quad (6.82)$$

The S-parameters can be derived from the ABCD parameters using conversion tables from Ref. [96]. In the case of S_{21} :

$$S_{21} = \frac{2R_c}{AR_c + B + CR_c^2 + DR_c} \quad (6.83)$$

Which follows that:

$$S_{21} = \frac{2R_c Y_{ex1} Y_{ex2}}{Y_{ex1} + Y_{ex2} + Y_0} \quad (6.84)$$

We can define the total admittance of the loaded resonator, Y_L :

$$Y_L = Y_0 + Y_{ex1} + Y_{ex2} = G_0 + G_{ex1} + G_{ex2} + i(2Q_0 G_0 \frac{\omega - \omega_0}{\omega_0} + B_{ex1} + B_{ex2}) \quad (6.85)$$

Comparing equation form of Equation 6.81 to Equation 6.85, one can see that the loaded resonator will have an effective shift in the resonant frequency based on the susceptance of the two transmission lines. This can be solved for by remembering that the resonance point will be when the impedance is entirely real, and therefore the susceptance must be zero:

$$\omega_L = \omega_0 \left[1 - \frac{B_{ex1} B_{ex2}}{2Q_0 G_0} \right] \quad (6.86)$$

What is important to note here is the denominator of the shift term containing $Q_0 G_0$; this can effectively be thought of as $\propto Q_0^2$, therefore, this shift term is negligible for $Q_0 \gg 1$. This means for high-Q cavities, usually above 1000, $\omega_L \approx \omega_0$. Further examination of Equation 6.85 shows that we can also define the total conductance of the loaded circuit, $G_L + G_0 + G_{ex1} + G_{ex2}$, which will, similarly to G_0 , be proportional to the loaded quality factor, Q_L such that:

$$Q_L = Q_0 \frac{G_0}{G_L} \quad (6.87)$$

The complex transmission coefficient for the mode of the cavity, S_{21r} , can then be derived from 6.84,6.80,6.85, 6.86, and 6.87 as:

$$S_{21r} = \frac{2R_c Y_{ex1} Y_{ex2}}{G_0(1 + \beta_1 + \beta_2) \left[1 + i2Q_L \frac{(\omega - \omega_L)}{\omega_0} \right]} \quad (6.88)$$

If we use the high-Q approximation, $\omega_L \approx \omega_0$, Equation 6.88 can be parameterized into the form,

$$S_{21} = \frac{a_1 t + a_2}{a_3 t + 1} \quad (6.89)$$

Where:

$$t = 2 \left(\frac{\omega - \omega_0}{\omega_0} \right) \quad (6.90)$$

and in this case:

$$a_1 = 0, \quad a_2 = \frac{2R_c Y_{ex1} Y_{ex2}}{G_0(1 + \beta_1 + \beta_2)}, \quad a_3 = iQ_L \quad (6.91)$$

Equation 6.89 represents a fractional linear transformation in the complex plane. This is a class of transformations that maps equations of lines to circles and vice-versa. In this case, the parameter, t , is the horizontal line in the complex plane, $y = 0$, the x-axis, being that ω and ω_0 are fully real. This special case of coefficients will always map this line to a circle centered on a point on the x-axis, with $S_{21r} \rightarrow 0$ for $f \rightarrow \pm\infty$, and $Im(S_{21r}(\omega_0)) = 0$, just as Figure 6.6 shows. The radius of this circle is the measure of the system loss, because $Re(S_{21r}(\omega_0)) = 2R_{Q-circle}$ will be entirely resistive. The point that S_{21} approaches as $f \rightarrow \pm\infty$, in this case the origin, is generally called the detuned point, S_{ppd} , where p is the port numbers and d just represents "detuned."

This analysis can be repeated for the reflection S-parameters as well, and will result in the same fractional linear form as Equation 6.89, but with different coefficients:

$$a_1 = iQ_L S_{ppd}, \quad a_2 = S_{ppd} + \frac{2R_c Y_{exp}^2}{G_0(1 + \beta_1 + \beta_2)}, \quad a_3 = iQ_L \quad (6.92)$$

The general geometric role of these three coefficients on the Q-circle become a little more clear here. S_{ppd} can be solved to be a_1/a_3 . a_2 is the reflection coefficient at the resonant

frequency, $S_{pp}(\omega_0)$. Most importantly, $Q_L = \text{Im}(a_3)$ in both transmission and reflection cases.

Two effects that aren't included in the transmission analysis are an explicit phase-delay in transmission lines and any cross-talk that may exist between the coupling structures. Any non-zero length cable will introduce some phase delay. This is a frequency dependent phase shift of the form, $\phi(f) = \phi_{ref}f/f_{ref}$, where ϕ_{ref} is the phase introduced by the delay at an arbitrary f_{ref} . This will have the effect of offsetting the phase angle as the scan approaches the resonance, resulting in a rotation of the Q-circle by an angle, $\phi \propto L_{line}$, where L_{line} is the length of the transmission lines. Because of the frequency dependence of this shift, a higher frequency regime will result in a larger angle rotation. One also expects this to distort the shape of the Q-circle slightly to a balloon-shape, with the strength of the effect proportional to the frequency regime of the resonance; this is harder to quantify and will be ignored for now. Cross-talk simply has the effect of a complex translation, re-centering the Q-circle to some point $X = (x_0, y_0)$. These two effects can be combined into a corrected S_{21} parameter:

$$S_{21} = (S_{21r} + X)e^{i\phi} \quad (6.93)$$

This doesn't change much other than changing the location of the Q circle. However, the order of this translation is unique; cross-talk shifts the circle first, and then it is rotated about the origin. The general geometric form of Equation 6.89 still holds, but the explicit coefficients in Equation 6.91 do not. The geometric interpretation of a_1, a_2, a_3 does hold; S_{21d} won't be at the origin, but will be at a_1/a_3 and $S_{21}(\omega_0) = a_2$. Since these two points will always be on opposite sides of the circle, despite its center in the complex plane, one can derive a useful expression for the radius of the Q-circle:

$$R_{Q-circle} = \frac{1}{2} \left| a_2 - \frac{a_1}{a_3} \right| \quad (6.94)$$

At this point one has everything needed to give a baseline process for estimating the Q-circle parameters from a VNA scan. One can determine the point closest to $S_{ij}(\omega_0)$ relatively easily from just finding the point with maximum or minimum $|S_{ij}|$ depending on if it is a

reflection or transmission parameter; a_2 is simply this complex point. The detuned point can be estimated by taking the two edge frequency points from the scan (the ones closest to $\omega \rightarrow \pm\infty$) and averaging them together:

$$S_{ij}(\omega \rightarrow \pm\infty) \approx \frac{1}{2}(S_{ij}(\omega_{max}) + S_{ij}(\omega_{min})) \quad (6.95)$$

This gives you a_1/a_3 . From there, the center of the Q-circle should be the average of these two points coordinates in the complex plane:

$$X_{center} \approx \frac{1}{2}(S_{ij}(\omega_0) + S_{ij}(\omega \rightarrow \pm\infty)) \quad (6.96)$$

The radius can then be calculated from Equation 6.94. There isn't an intuitive geometric way of getting a_3 alone from the circle, but if one takes a Lorentzian fit of the magnitude data, they now have an estimate. This forms a nice set of initial guesses to feed into a complex curve fitting process.

In terms of using this for a fitting process, there are many routes one can take to implement this geometric interpretation. If the goal is simply to measure Q_0 , there's a relatively simple way: perform a complex fit of an S_{21} scan to Equation 6.89, solving for the 3 coefficients (using the estimates as initial parameter values). Once done, simply take the imaginary part of a_3 to get a revised Q_L . Next, you can either determine β_1 and β_2 from a reflection Lorentzian measurement as described earlier, or you can also perform a complex fit to Equation 6.89 on the reflection Q-circles through the same process. In this case, one uses their a_{2i} from each reflection fit on ports $i = 1, 2$ to get β_i :

$$\beta_i = \frac{1 - |a_{2ii}|}{|a_{211}| + |a_{222}|} \quad (6.97)$$

and from there Equation 6.72 still applies for finding Q_0 . Combining all of this together in terms of fit parameters:

$$Q_0 = Im(a_{321}) \left(1 + \frac{2 - |a_{211}| - |a_{222}|}{|a_{211}| + |a_{222}|} \right) \quad (6.98)$$

One can take this much further if they wish to quantify more specifics about the system. Further analysis in ref. [77] gives you exact more exact expressions for the diameters of

each S-parameter's Q-circle; this is useful for quantifying the losses in the measurement setup. Inverse mapping techniques outlined in Ref. [91] can give you good values for ϕ and X to quantify phase-delay and cross-talk; this reference also contains several other fitting techniques of note.

Short note on Impedance Mismatches

Impedance mismatches are also often common in measurement setups between cables, connectors, and antennae. This can usually be diagnosed by just looking at the magnitude of a reflection measurement, $|S_{11}|$ or $|S_{22}|$: A well-matched setup will have a constant response off-resonance, whereas the presence of standing waves indicates an impedance mismatch. In the complex plane, this will manifest itself in a second Q-circle in the off-resonance region, near the detuned point, whose center does not coincide with the center of the Smith circle; this is often called a coupling loss circle. See Ref. [76, 77] for more details.




Figure 6.9: A Q-circle with an impedance mismatch [76].

6.8 High frequency axion Searches

6.8.1 Motivation: the problems at high frequencies

Upon further examination of Equation 3.31, it is apparent that many of the experimental variables depend on the search frequency. Most notably, since $f_{TM_{010}} \propto 1/R$, one will have to create a smaller cavity for higher frequency searches. Since $\frac{df}{dt} \propto V^2$, this would imply that volume contributes a factor $f^{N_V} = f^{-4}$ slow down in scan rate! But this isn't the only factor that will degrade at higher frequencies; Equation 6.37 implies that the conductivity of copper cavities will degrade the Q as the anomalous skin depth grows at higher frequencies. One can show that $Q \propto f^{-2/3}$ for a classic copper cavity, and since Q is linear with the scan rate, this means a power of $N_Q = -2/3$ added to the scan rate slow-down. Finally, although currently ADMX operates above the standard quantum limit for system noise temperature, this noise level grows with frequency such that $T_{sys} \propto f^1$. This results in a $N_{T_{sys}} = -2$, because $\frac{df}{dt} \propto T_{sys}^{-2}$. One is saved a tiny bit by the fact that scan rate explicitly depends on frequency, $\frac{df}{dt} \propto f^2$, but all in all this isn't enough and the total frequency dependence can be written:

$$\frac{df}{dt} \propto f^{N_f + N_{T_{sys}} + N_Q + N_V} = f^{2 - 2 - (2/3) - 4} = f^{-14/3} \quad (6.99)$$

Clearly, the biggest problem here is the degradation from the volume drop off, and any high frequency search needs to resolve or mitigate the $N_V = -4$ factor. Future ADMX systems, such as run 2 and ADMX EFR, do this by having multiple smaller cavities whose total volume is comparable to the current ADMX main volume. A naive estimate for the number of cavities needed can be made with the expression, $N(f) = (f/f_0)^2$, where f_0 is the frequency of the original single cavity system. However this assumes cavities of the same length and ignores the wasted space in the magnet bore due to packing multiple cavities together. Using Table 6.1, one can see that going from the ADMX-main single cavity to the run 2A frequency range implies $N \approx 6$ is necessary to keep the same volume; however this doesn't account for the tuning rod perturbation. Run 2 will end up being a 4 cavity system.

Beyond the geometry, the trick with these systems is how and when one combines these

separate cavity signals. If one does not care about matching these cavities signals coherently, they then constitute N independent haloscope experiments, and the total signal power of each cavity must be combined in quadrature such that $P_{sig,N} = \sqrt{N}P_{sig,0}$. However, if the cavity signals can be combined in phase, one can simply sum the powers $P_{phase-matched} = NP_{sig,0}$, offering the potential of a \sqrt{N} improvement. Original prototype multi-cavity systems were to power combine cavity signals before the primary amplifier to minimize the number of RF receivers for the system. This obviously also requires resonant frequencies of these cavities to be matched when combining. A mismatch in summing each cavity's Lorentzian curve will have the effect of spreading it out, lowering the net quality factor of the system. Tuning each cavity adds to the complexity of this matching process, requiring some sort of cavity frequency "locking" protocol to make sure everything is stable before a net digitization would take place.

For these reasons and more, it was settled on having each cavity have its own RF receiver. Data for each cavity is stored separately, and they are combined digitally in the analysis. This gives each cavity the flexibility of tuning through its range at its own pace, not being slowed down by the other cavities, or additional calibration processes. There's still the possibility of phase-combining their separate signals digitally, trying to get at the \sqrt{N} improvement for at least some of the run. The only down-side to this route is the number of amplifiers, cables, and hardware also increases with N , requiring more fridge space and generally increasing the number of things to be maintained on site.

6.8.2 ADMX Run 2

ADMX run 2 is a 4-cavity array designed to cover the 1.3 – 2 GHz range. Similar to the ADMX-main, each tube has a knife-edge on its end that is clamped to a recessed area on the end cap plate via a stainless steel collar as pictured in Fig 6.10. Each is copper-plated in the same manner as ADMX-main. As shown in Figure 6.10, the original prototype had all four cavities attached to the same end cap plate, with a joint tuning rod system. This system would have one coarse tuner that turned 4 tuning rods all at once, as well as 4 individual




Figure 6.10: The run 2 cavity system assembled at LLNL.

fine tuners that would allow for adjustment of the cavity frequencies within 1 kHz for locking them together. However, upon switching from an analog to digital combining system, this was changed to 4 different cavity end plates and tuners that join together. This cuts down on the tolerance requirements for the end plate, and allows for each cavity to be assembled and tuned individually. It has an overall packing footprint such that it can operate in the current ADMX magnet bore at UW, although it will still require many changes to the site and insert to accommodate the necessary electronics. This includes converting the current




Figure 6.11: The run 2A cavity array original coarse tuner design.

stepper motors and gearboxes into an 8 piezoelectric motor system, covering each tuning rod and strong port antenna. Additionally, 4 separate RF cryogenic receiver chains, instead of 2, need to be able to fit within the system. Because of these space requirements, ADMX Sidecar will also be retired before this system is commissioned. The overall system has a volume of 85 L, which is a slight reduction (62.5%) from ADMX run 1 volume of 136 L. Table 6.1 lists the empty cavity frequency characteristics.

The system was first assembled at LLNL by myself and former LLNL staff scientist Nathan Woollett, where we also measured the empty cavities. The previous plots, Figure 6.4 and 6.5, were generated from a single cavity of this system. Figure 6.12 shows the empty reflection measurement of each cavity, with the frequency drift due to the difference in machining between them. The appreciable difference in Cavity C's frequency partially motivated the switch to separate tuning systems, because it was a larger frequency drift than the original fine tuning system could handle. It has since been shipped to Fermilab

where it resides currently for cryogenic testing. It will be shipped to UW for commissioning in the insert after the conclusion of run 1D.




Figure 6.12: Comparison of the TM_{010} mode weakly coupled reflection measurement for each cavity of the run 2 system. Drifts in the dip are due slightly different couplings. Only cavity C had an appreciably different frequency > 1 kHz.

6.8.3 ADMX EFR

ADMX Extended Frequency Range is an 18-cavity array that will cover the 2 – 4 GHz range to DFSZ sensitivity. The goal is to have a scan speed that is 5x greater than the current ADMX, which makes siting it at UW in the current magnet bore not possible. Instead, ADMX-EFR will be located in its own building at Fermilab with a completely new site design, outlined in Figure 6.13.

A 9.4 T research MRI magnet manufactured by GE Medical Systems has been requisitioned from University of Illinois Chicago for the project. With a 80 cm bore allowing for more volume, this means a 47% increase in $B_0 V_{bore}$ over the current ADMX magnet.




Figure 6.13: Diagram of the new EFR site at Fermilab. The resonator system and amplifier systems are stored and cooled separately. This allows the cold electronics to be cooled further and minimize the magnetic shielding requirements from the 9.4 T magnet. The system is horizontal to accommodate the MRI magnet design.

If this volume increase is fully utilized for cavities, one expects an increase in scan rate by $\approx 2.65 - 4.1\times$ depending on what operating field value one chooses for the new magnet, $\approx 8.4 - 9.4$ T. Because this an MRI magnet, the entire system must be horizontal as pictured in Figure 6.13.

The other major change to the site construction is the separation of the cold electronics and resonator cooling system. One dilution refrigerator will cool the cavity system to 100 mK, while the electronics package will be cooled separately to 25 mK. This allows the 18 JPA SQUID amplifiers to be stored out of the magnetic field, negating the need for an active bucking coil system; a passive magnetic shield will be used instead. Quantum squeezing techniques will hopefully allow this electronics package to operate below the standard quantum limit. Low-loss coaxial cabling will be critical for connecting this system to the cavities.




Figure 6.14: The proposed ADMX-EFR 18 cavity array. Each cavity is a hexagonal clam-shell to maximize packing efficiency and allow for superconducting coatings. Every cavity is topped with a rotor piezoelectric motor and geared tuning arm, as well as a linear piezoelectric motor for individual adjustment of the antenna coupling.

The actual cavity system itself will be 18 "clam-shell" cavities with hexagonal exteriors to maximize the packing volume, pictured in Figure 6.14. The nominal cavity diameter and height are 128 mm and 1 m respectively, with empty cavity characteristics listed in Table 6.1. This gives each cavity a loaded volume of 12.1 L for the 2-3 GHz range, and 10.1 L for the 3-4 GHz range; larger tuning rods will be implemented half-way through. This improves the total volume to 182 – 218 L, a 34-60% improvement over the current ADMX cavity. Currently several reduced length aluminum prototypes, pictured in Figure 6.15 are being tested at LLNL before the design is finalized. Considerations are being made for this system to be superconducting; the clam-shell design lends itself well to the coating processes. Additionally, cost-benefit analyses are being done on the efficacy of niobium substrate; the results of this dissertation are informing many EFR decisions! (see Chapter 9). A projected scan speed of the EFR system is shown in Figure 6.16. Stay tuned for a publication soon on the design and projection of this future ADMX system.



Figure 6.15: An open aluminum reduced length EFR prototype cavity being tested at LLNL.

6.8.4 ADMX Orpheus

ADMX Orpheus is a path-finder experiment sited at the UW CENPA site that implements an alternative approach to the volume degradation problem; it is a tunable, dielectrically-loaded cavity operating at higher order mode, allowing the detection volume to remain the same. Although I did not work on Orpheus during my time at UW, I believe its important to include in order to emphasize that multi-cavity systems are not the only approach ADMX is taking to tackle high frequency searches. Dielectrics are strategically placed at every $1/4$ th of a half wavelength of the operating mode to suppress the electric field in those regions as shown in Figure 6.17. This results in a $\vec{E} \cdot \vec{B} \neq 0$ for the higher order mode. One can think




Figure 6.16: The instantaneous scan rate estimated for full ADMX-EFR run. The orange line represents our conservative estimate (*threshold*). The green line represents our most optimistic estimate (*objective*) assuming optimal performance of all components, superconducting cavities and squeezing.




Figure 6.17: Diagram of a dielectrically-loaded cavity operating at a multi-wavelength mode. The dielectrics (grey boxes) are placed at every anti-node such that they suppress parts of the mode wavelength (\vec{E}) so that the axion form factor is non-zero ($\propto \vec{E} \cdot \vec{A}$) [35].

of this as producing a non-zero form factor. However, because Orpheus is a Fabry-Perot cavity, with an open-walled design, typically the form factor is combined with cavity volume and referenced as an effective volume, V_{eff} :

$$V_{eff} = \frac{|\int dV \vec{B}_0 \cdot \vec{E}|^2}{B_0^2 \int dV \epsilon_r |E|^2} \quad (6.100)$$

The model of this Fabry-Perot cavity is pictured in Figure 6.18. Mirrors on either end form the resonator, reflecting the electromagnetic waves back and forth between the two. A small aperture on the top mirror acts as the strong port, and the bottom aperture is the weak port. The operating mode is the TEM_{00-18} . The cavity is tuned by adjusting the position of the bottom mirror, changing the length between the mirrors. Scissor jacks evenly adjust the spacing of the dielectric plates as it tunes so that they remain in the right relative locations to create the non-zero V_{eff} .




Figure 6.18: CAD model diagram for Orpheus Cavity [35].

Orpheus has a tuning range from 15.8 GHz to 16.8 GHz. In the summer of 2021, ADMX Orpheus performed a search for dark photons in this range, the results of which can be found in Ref. [35]. A 1.5 T dipole magnet is being constructed at UW for an axion search in this same range to a nominal $g_{a\gamma\gamma} \approx 3 \times 10^{-12} GeV^{-1}$, twice the sensitivity as the CAST limit. The magnet design can be found in ref. [36]. The proposed limits are shown in Figure 6.19. Stay tuned for results!




Figure 6.19: The project Orpheus axion limit assuming a 1.5 T dipole magnet[35]. Figure adapted from [87].

Chapter 7

SRF CAVITY THEORY

Superconducting Radio Frequency Cavities, although challenging to implement on axion haloscopes, are an established tool in the physics community, and the standard for achieving the highest Q cavities in the laboratory. They are essential in the accelerator community for efficiently transferring the large RF power along particle beams [101]. In recent years, they have become a crucial tool for housing and high fidelity read out of qubits for the Quantum Information Science (QIS) community [51]. Within the dark matter community, superconducting cavities have been used to look for dark photons at an incredibly high sensitivity [34]. The axion community has done several feasibility studies of SRF cavities for a variety of materials over the past few years, particularly NbTi [30, 7, 41], Nb₃Sn [94], and YBCO [6, 4, 5]. The CAPP axion institute will soon be publishing their first axion search results using a YBCO cavity and tuning rod [122].

Superconductivity refers to the property certain solid materials exhibit where internal electrical resistance vanishes. Normal conductors, like copper, exhibit an internal resistance that decreases with temperature gradually until absolute zero. However, superconductors each have a characteristic critical temperature, T_c , below which their resistance drops abruptly to zero. This phenomenon was first observed in 1911 by Heike Kamerlingh Onnes, who immersed a solid mercury wire in liquid helium that went superconducting at 4.2 K (he was also the first physicist to liquefy helium in 1908, for which he won the 1913 Nobel Prize in Physics). Since then, a myriad of materials have been shown to be superconductors with critical temperatures ranging from just thousandths of a degree above 0 K to record holding yttrium and lanthanum hydride compounds that range from $T_c \approx 224 - 288$ K [112, 45]. However, this upper limit is routinely pushed to higher levels with the hope of one day having

a room temperature superconductor. Figure 7.1 shows the critical temperature for elements of the periodic table, although this does not do justice to the vast variety of compounds that exhibit the property [53]. Throughout this section I will be drawing from two introductory textbooks on superconductivity, Ref. [89, 111].

Periodic table of superconductivity

H																	He
Li 0.0004	Be 0.026											B 11	C	N	O 0.6	F	Ne
Na	Mg											Al 1.18	Si 8.2	P 13	S 17.3	Cl	Ar
K	Ca 29	Sc 19.6	Ti 0.5	V 5.4	Cr	Mn	Fe 2.1	Co	Ni	Cu	Zn 0.87	Ga 1.1	Ge 5.35	As 2.4	Se 8	Br 1.4	Kr
Rb	Sr 7	Y 19.5	Zr 0.85	Nb 9.25	Mo 0.92	Tc 8.2	Ru 0.5	Rh 0.0003	Pd	Ag	Cd 0.5	In 3.4	Sn 3.7	Sb 3.9	Te 7.5	I 1.2	Xe
Cs 1.3	Ba 5		Hf 0.38	Ta 4.5	W 0.01	Re 1.7	Os 0.7	Ir 0.1	Pt	Au	Hg 4.15	Tl 2.4	Pb 7.2	Bi 8.5	Po	At	Rn
Fr	Ra		Rf	Db	Sg	Bh	Hs	Mt	Ds	Rg	Cn	Nh	Fl	Mc	Lv	Ts	Og
Lanthanides		La 6	Ce 1.7	Pr	Nd	Pm	Sm	Eu 2.7	Gd	Tb	Dy	Ho	Er	Tm	Yb	Lu 0.1	
Actinides		Ac	Th 1.4	Pa 1.4	U 1.3	Np	Pu	Am 1.0	Cm	Bk	Cf	Es	Fm	Md	No	Lr	

T_c (K) Ambient pressure superconductor
T_c (K) High pressure superconductor

Figure 7.1: The periodic table of elements with the critical temperature of known superconductors listed. Some elements only exhibit superconductivity above atmospheric pressure, known as high-pressure superconductors [53].

7.1 The Meissner effect

An additional effect of superconductivity is the expulsion of magnetic fields from a superconductor during its superconducting transition. This was first noticed in 1933 by German physicists Walther Meissner and Robert Ochsenfeld when measuring the magnetic field distribution outside samples of superconducting lead and tin, and is now referred to as the Meissner effect [83]. An external magnetic field was applied to samples that were subsequently cooled below their critical temperature, causing the samples' internal magnetic field

to drop to zero. They detected this effect through the conservation of magnetic flux; when the internal field decreased, the exterior field around the superconductor increased. This gave a defining feature to superconductors versus a classical perfect conductor for the first time. What is clear, even without a theory of superconductivity, is that the superconductor itself must be generating a magnetization, \vec{M} , in response to the external field, \vec{H} . As \vec{H} is increased, eventually the superconductor reaches a limit where \vec{M} can no longer cancel out the external field, known as the critical field, \vec{H}_c , forcing the material back into a normal state. Being that currents generate magnetic fields, often this is also expressed as a critical current in the case of superconducting wires. This critical field was found empirically to be well approximated by the parabolic relation:

$$H_c(T) \approx H_c(0 \text{ K}) \left(1 - \left(\frac{T}{T_c} \right)^2 \right) \quad (7.1)$$

where $H_c(0 \text{ K})$ is the projected critical field at 0 K.

This effect is the heart of the problem for using superconducting cavities in axion haloscopes. Axion haloscopes can only operate effectively in high fields, in excess of several Tesla, as discussed in Chapter 3. For reasons that will be clearer by the end of this chapter, there are no known materials that exhibit perfect superconductivity in excess of around 1 T. This means most materials will be forced back into a normal state within a haloscope, where their resistance is higher than copper. This is why ADMX has so far settled on using a high conductivity normal conductor, copper.

7.2 London theory

Although superconductivity can only be explained in full with quantum mechanics, the first theory of superconductivity was purely classical in nature. Only two years after the discovery of the Meissner effect, in 1935, Fritz and Heinz London worked out a phenomenological theory of superconductors now known as London theory [78]. The crux of superconductors is that Ohm's law, $\vec{j} = \sigma \vec{E}$, clearly cannot apply because $\sigma = 0$. The brothers imagined the superconducting current as free electrons within an external uniform electric field, which

would follow the Lorentz force law, $\vec{F} = m\dot{v} = -e\vec{E} + e\vec{v}_s \times \vec{B}$, with $\vec{j}_s = -n_s e\vec{v}_s$, a superconducting current density. Assuming the electrons are just being driven by the electric field, one can derive the first London equation:

$$\frac{\partial \vec{j}_s}{\partial t} = \frac{n_s e^2}{m_e} \vec{E} \quad (7.2)$$

where e and m_e are the electron charge and mass, and n_s is a number density of charges. The 2nd equation can be derived by applying Faraday's law to the first to obtain:

$$\frac{\partial}{\partial t} \left(\vec{\nabla} \times \vec{j}_s + \frac{n_s e^2}{m_e} \vec{B} \right) = 0 \quad (7.3)$$

They recognized any solutions that were non-zero constants would be non-physical due to the recent discovery of the Meissner effect, and required the time derivative be exactly zero, resulting in the 2nd London equation:

$$\vec{\nabla} \times \vec{j}_s = -\frac{n_s e^2}{m_e} \vec{B} \quad (7.4)$$

Notice that this is the key step that differentiates it from a perfect conductor; it can only be said that a perfect conductor has $\dot{B} = 0$ within its bulk, whereas the field itself is zero, $\vec{B} = 0$, within a superconductor. This means a perfect conductor can hold trapped magnetic flux whereas a superconductor cannot. Often, these two equations can be expressed as one equation in terms of a superconducting vector potential, \vec{A}_s :

$$\vec{j}_s = -\frac{n_s e^2}{m_e} \vec{A}_s \quad (7.5)$$

of which they defined the 'London gauge' for \vec{A}_s :

$$\vec{\nabla} \cdot \vec{A}_s = 0, \quad \vec{A}_s \cdot \hat{n} = 0 \quad (7.6)$$

where \hat{n} is the surface normal vector of the superconductor. Additionally, they required $\vec{A}_s = 0$ within the superconducting bulk. This becomes clearer by taking the curl of Equation 7.4 resulting in a Helmholtz equation:

$$\nabla^2 \vec{B} = \frac{1}{\lambda_s^2} \vec{B} \quad (7.7)$$

where λ_s is the London penetration depth,

$$\lambda_L = \sqrt{\frac{m_e}{\mu_0 n_s e^2}} = \sqrt{\frac{m_e c^2}{4\pi n_s e^2}} \quad (7.8)$$

where the second definition in terms of c is more common. This name makes more sense when one solves Equation 7.7 for the case of an infinite plane superconducting bulk sitting in a constant external magnetic field, B_0 , oriented parallel to the surface of the superconductor along the \hat{z} . Defining \hat{x} to be the perpendicular direction where $x < 0$ is free space and $x > 0$ is within the superconductor, one can solve for B_z within the superconductor to be:

$$B_z(x) = B_0 e^{-x/\lambda_L} \quad (7.9)$$

Hence, this is the length scale by which an external field penetrates within a superconductor before it is exponentially suppressed. Typical values are 50 to 500 nm with Nb_3Sn having a value of about 160 nm. One can reason from this that the super-currents must also flow near the surface of the superconductor in order for this magnetic field distribution to hold. This theory triumphs at giving a clear classical picture of the Meissner effect and surface super-currents, yet it still fails at a microscopic explanation of superconductivity. Even more so, the charge carriers in this theory, electrons, being fermions and bound to the Pauli exclusion principle, pose problems when developing a quantum theory of superconductivity.

7.3 Ginzberg-Landau theory

The 2nd phenomenological theory of superconductivity was developed by Soviet physicists Lev Landau and Vitaly Ginzburg, first published in 1950 [74], who approached the problem with quantum mechanics in mind. Lev Landau made use of his theory of 2nd order phase transitions, published in 1937 [73], describing the free energy density of the superconductor, f_s , in terms of a complex order parameter field, $\vec{\psi}$, analogous to a wave function; the quantity $|\psi|^2$ is a measure of the local density of superconducting electrons, n_s . As Landau's explanation of super-fluidity in liquid helium was only years earlier, in 1941, their thought was the electrons in a superconductor form a super-fluid, and this would represent the fraction

of electrons that had condensed into this state. They assumed $|\psi|$ would be non-zero below the transition temperature, T_c , but also small with small fluctuations ($|\nabla\psi|$ is small), which actually holds well for most low-temperature, conventional superconductors we deal with in this dissertation, but not the more exotic modern high-temperature (HTC) superconductors. This assumption allows one to assume the form of the free energy density to have the form of a field theory exhibiting a U(1) gauge symmetry:

$$f_s = f_n + \alpha|\psi|^2 + \frac{1}{2}\beta|\psi|^4 + \frac{1}{2m^*} \left| \left(-i\hbar\vec{\nabla} - \frac{e^*}{c}\vec{A} \right) \psi \right|^2 + \frac{\vec{B}^2}{8\pi} \quad (7.10)$$

where α and β are the phenomenological parameters of the theory, f_n is the free energy density of the normal conductor state, m^* and e^* are the effective charge and mass of the superconducting current carrier, and \vec{B} and \vec{A} correspond to the external magnetic field and its vector potential. One wants to minimize the total free energy, $F = \int f_s d^3r$, with respect to variations in the order parameter, ψ , which gives the first Ginzburg-Landau (GL) equation:

$$\alpha\psi + \beta|\psi|^2\psi + \frac{1}{2m^*} \left(-i\hbar\vec{\nabla} - \frac{e^*}{c}\vec{A} \right)^2 \psi = 0 \quad (7.11)$$

One can then use Ampere's law to derive the superconducting current in the presence of the external field to be:

$$\vec{\nabla} \times \vec{B} = \frac{4\pi}{c} \vec{J}; \quad \vec{J} = \frac{e^*}{m^*} \text{Re} \left\{ \psi^* \left(-i\hbar\vec{\nabla} - \frac{e^*}{c}\vec{A} \right) \psi \right\} \quad (7.12)$$

This is referred to as the 2nd GL equation. A student of quantum mechanics will notice how this is remarkably similar to the time-independent Schrodinger equation with the exception of the non-linear term, $\beta|\psi|^2\psi$. In the absence of an external field this simplifies greatly:

$$\alpha\psi + \beta|\psi|^2\psi = 0 \quad (7.13)$$

This equation has the trivial solution $\psi = 0$, which would correspond to the normal conducting state, $T > T_c$, but below the transition temperature one expects a non-trivial solution of the form $|\psi|^2 = -\alpha/\beta$. This means in order for a real non-trivial solution to exist below T_c , but not above, one can assume the functional form of α to be $\alpha(T) = \alpha_0(T - T_c)$ where α_0

must be positive. The linear dependence can be made here under the assumption that one is near the transition temperature. Note this also requires β to be positive and temperature independent. Hence the order parameter is zero close to the transition, which is ideal for second order phase transitions. Therefore in the superconducting state absent of magnetic fields, $|\psi|^2$ is :

$$|\psi|^2 = -\frac{\alpha_0(T - T_c)}{\beta} \quad (7.14)$$

Considering the 1-dimensional case, where we use a normalized wave function, $s(x) = \psi(x)/|\psi|^2$, on the first GL equation, 7.11, for the case of zero field, $\vec{A} = 0$, one gets a differential equation of the form:

$$\xi^2 \frac{d^2 s(x)}{dx^2} + s(x) - s(x)^3 = 0 \quad (7.15)$$

where,

$$\xi = \sqrt{\frac{\hbar^2}{2m^*|\alpha|}} \quad (7.16)$$

This equation shows how ξ is a length scale over which the superconducting order parameter will fluctuate; this is referred to as the coherence length of the superconductor. Note that it will have a temperature dependence of $(1 - T/T_c)^{-1/2}$.

If we return to the 2nd GL equation, Equation 7.12, considering the non-zero field case, one can perform the same analysis as in the London theory: apply the curl to the equation to get the form $\nabla^2 \vec{B} = \lambda^{-2} \vec{B}$. This gives us the penetration depth for GL theory to be:

$$\lambda_{GL} = \sqrt{\frac{m^*c^2}{4\pi e^{*2}|\psi|^2}} = \sqrt{\frac{m^*c^2\beta}{4\pi e^{*2}|\alpha|}} \quad (7.17)$$

This similarly has the $(1 - T/T_c)^{-1/2}$ temperature dependence. In principle, this should be equal to the London penetration depth, Equation 7.8, and is commonly just referred to as λ . The ratio of these two length scales, λ and ξ , are critical for understanding the behavior of the superconductor in fields, and one can see that it must be temperature independent:

$$\kappa = \frac{\lambda}{\xi} = \frac{m^*c}{e^*\hbar} \sqrt{\frac{\beta}{2\pi}} \quad (7.18)$$

This is called the Ginzburg-Landau parameter, or GL parameter. Ginzburg and Landau were able to show that for $\kappa \gtrsim 1$ the superconductor had a positive free energy at the superconducting boundary; this was the case for pretty much all superconductors known at the time. Intuitively, one can see that if $\kappa > 1$, meaning the penetration depth is greater than the coherence length of the superconductor, the superconducting state will run into instability problems at its boundaries as an external field is raised (more on this later).

If one considers again the zero field transition of superconductor, the energy density gap between the normal and superconducting states, $f_s - f_N$, can be found from Equation 7.10 to be $-\alpha^2/2\beta$. This can be thought of as the transition energy between the two states, which would be indicative of a drastic change in electron behavior since such a large spike in conductivity occurs below the phase transition. One can imagine the hypothetical scenario where a superconductor is cooled through its superconducting transition all the way to absolute zero. Once at absolute zero, an external field is applied to the superconductor, and slowly raised. The superconductor would respond according to the Meissner effect, repelling the field according to its field strength, until the magnetic field exceeds a critical value, H_{c0} , denoting the 0 K value. One might then think that this field value corresponds with a magnetic energy, $H_{c0}^2/8\pi$, equal to that energy gap, $\alpha^2/2\beta$. However, this drastically underestimates the critical fields that are observed in the laboratory for superconductors, only giving values on the order of 10^{-2} T.

Finally, notice that this energy gap, $-\alpha^2/2\beta$, has a temperature dependence $(1-T/T_c)^{-1}$. This seems to suggest that if the superconducting state has some sort of binding energy, it is gradually weakened as it approaches the critical temperature; perhaps electrons are entering a bound state that allows them to form a superfluid. Remember electrons are spin-1/2 fermions, and cannot form a Bose condensate.

7.4 BCS theory

The first microscopic theory of superconductivity was published in 1957 by John Bardeen, Leon Cooper, and Robert Schrieffer, now called by the last initials of the three, BCS theory

[17]. They would win the Nobel prize for this work in 1972. The fundamental problem with electrons condensing into a charged super-fluid is that they do not obey the Bose-Einstein statistics necessary for its formation. BCS theory solves this by suggesting the formation of composite bosons made from two electrons called Cooper pairs. Cooper showed that even an arbitrarily small attraction between electrons in the crystalline lattice of a metal can cause a paired state with an energy lower than the Fermi energy. This is conventionally explained as a electron-phonon interaction; in the classical picture of a metal, an electron distorts the local lattice of positive ions towards itself due to its electrical attraction, increasing the local positive charge density, which can attract electrons (with opposite spin) that would otherwise be repelled by the electron. The collective motion of this positive lattice acts as the phonon.

This pairing distance must be long enough to shield the electron-electron repulsion, to cause the paired state. This is a fundamentally quantum mechanical phenomenon, highlighted by the fact that this minimum shielding distance between these bound electrons is usually 100-1000 nm, which is much larger than the spacing of conduction electrons themselves; this means cooper pairs overlap significantly, enough so that there may be up to 10^7 electrons between the two electrons in a cooper pair! This is also why some may frown at my description of a composite boson being formed from two individual electrons; the paired state is energetically preferential to each electron and electrons will go in and out of it accordingly with many different partners.

This binding energy is very weak, of the order 10^{-3} eV, thus thermal energy easily breaks these pairs up unless significantly cooled. This means upon cooling, electrons start to form cooper pairs within the metal, which then form a composite charged bosonic gas, allowing the condensation of a super-fluid below its T_c . Once pairs are condensing together, the energy required to break any given pair depends on all the other pairs breaking as well, making the energy barrier to break the state macroscopically high. Any collisions/friction of the positive ions with the cooper pairs has insignificant energy in comparison, thus all the pair states are preserved and electron flow is unabated with zero resistance!

Another cool piece of evidence for this theory is that it supports the isotope effect.

The phonon frequency is inversely proportional to square root of the lattice ions mass; the critical temperature of many conventional superconductors follow the same inverse square root relation. In the classical picture, heavier lattice ions are displaced less by electrons, lowering the binding energy, and requiring more cooling before a condensate can form.

Another important prediction of BCS theory is an energy gap, Δ , half the energy to break a cooper pair, being dependent on the critical temperature. The value was predicted to increase from zero at T_c to a limiting value,

$$\Delta(0) = 3.528k_B T_c \quad (7.19)$$

This result agreed very well with experiments and gave the theory credence. This energy gap can also be used to approximate a surface resistance of superconductor:

$$R_{BCS} = A \frac{f^2}{T} e^{-\frac{\Delta}{k_b T}} \quad (7.20)$$

which will be useful for predicting surface resistances of SRF cavities. All in all the BCS theory was capable of explaining all the phenomenon of superconductivity predicted by London theory and GL theory. Where it starts to fail is in the explanation of HTC superconductors, such as YBCO, but that is not relevant to the conventional superconductors discussed in this dissertation.

7.5 Type-I and type-II superconductors

In the same year BCS theory was published, 1957, Alexei Abrikosov published work on superconductors with high GL parameters, κ [2]. The experimental discovery of high κ superconductors in Lead-Bismuth alloys was during 1935-36 [98], but hardly noticed at the time. Much of this work was inspired by the theory work of Richard Feynman and Lars Onsager on quantum vortices in superfluids [52, 88]. Fritz London also had showed the quantization of magnetic flux in superconducting solenoids, which was also very relevant to the work [79]. Ginzburg and Landau had shown in their work that above some value of κ , superconductors would exhibit a negative surface energy, as mentioned earlier. Nonetheless,

Abrikosov was able to take GL-theory to precisely show that for $\kappa > 1/\sqrt{2}$, superconductors would exhibit this negative surface energy and generalize the ideas around this distinction between superconductors. What are now called Type-I superconductors, $\kappa < 1/\sqrt{2}$, behave as described so far; upon super-conduction, they repel any external magnetic fields below their critical field completely, H_c . However, above this singular critical field, the Meissner effect breaks down entirely, and the material returns to the normal state. In accordance with Equation 7.1, the phase diagram for Type-I superconductors is shown in Figure 7.2.




Figure 7.2: The phase diagram of a type-I superconductor. The transition line represents the parabolic relation of B_c , the critical field, with temperature. If $B < B_c$ and $T < T_c$, the critical temperature, the object is superconducting. Note that B_c and H_c are often used interchangeably [25].

For $\kappa > 1/\sqrt{2}$, the superconductor would exhibit a transitional phase where the external magnetic field would partially penetrate the superconductor via vortices in the normal state, while the remainder would remain superconducting (see Figure 7.3). The magnetic flux passing through each vortex is quantized by the magnetic flux quantum $\Phi_0 = h/2e$, giving them the name fluxons. Because of this, the number of vortices is directly proportional to the flux passing through the superconductor, $N_{vortex} \approx \Phi_{ext}/\Phi_0$. Furthermore, the relative size of each vortex is proportional to the penetration depth, $\approx \pi\lambda^2$, meaning there is a sharp cutoff between the superconducting and normal state about the vortex core. Abrikosov




Figure 7.3: Type-II superconducting film in the intermediate vortex phase, gray arrows represent the magnetic field penetrating through the quantized flux vortices, which generate their own circular surface currents (green arrows). The red arrow represents the Meissner surface current [109].

found that these vortices try to arrange themselves in a lattice structure, now known as an Abrikosov Lattice. This lattice of vortices can be imaged with SQUID scanning microscopy, as pictured in Figure 7.4. Being an intermediate phase, this implies there are two critical fields: one field below which the superconductor is fully in the Meissner state, H_{c1} , and one field above which the material is fully in the normal state, H_{c2} . The phase diagram for the type-II superconductor is pictured in Figure 7.5.

An important practical feature of type-II superconductors is the typical difference between H_{c1} and H_{c2} values. Both type-I superconductors' H_c and type-II superconductors' H_{c1} are pretty much never in excess of a single Tesla, and are more commonly on the order of 100 mT [100]. However, common type-II superconductors' H_{c2} are often multiple Teslas and can be much higher than the highest laboratory magnetic fields generated; this isn't surprising because the magnets themselves are made out of these type-II superconductors! This why type-II superconductors are most commonly used in applications: Magnet wire, Josephson devices, and SRF cavities. In the case of SRF cavities for axion haloscopes, with $B > 1$ T at a minimum, type-II superconductors operating in this mixed vortex state are the only




Figure 7.4: A scanning SQUID microscopy image of a 200 nm thick YBCO film (Type-II superconductor) after field cooling at $6.93 \mu\text{T}$ to 4 K. The black dots are flux vortices or fluxons. [117].

option; it is therefore important to understand what drives the resistance in this mixed state.

7.6 Fluxons in SRF cavities

At the high frequencies of RF cavities, similar to Chapter 6, the response of the superconductor must be modelled with a complex surface impedance $Z_s = R_s + iX_s$, where R_s and X_s are the surface resistance and reactance respectively. Often it is the case to consider this the result of a complex resistivity or complex conductivity, depending on your preference. In the case of resistivity, $Z_s = \sqrt{i\omega\mu_0\tilde{\rho}} = i\omega\mu_0\tilde{\lambda}$, where $\tilde{\lambda}$ refers to a complex shielding length, this length is a function of H and T , and should capture the quasi-particle scattering and vortex motion that governs the impedance of a type-II superconductor. This means being a function of the complex conductivity of the superconducting area, $\sigma_f = \sigma_1 + i\sigma_2$, as well as the resistivity of the vortices and their motion, ρ_{vm} [111]. If one is not close to a transition, near T_c or H_{c2} , one can assume that the imaginary part of the complex conductivity




Figure 7.5: The phase diagram of a type-II superconductor. Both critical fields follow a parabolic relation with temperature. In zone a, $B < B_{c1}$, the object is completely superconducting. In zone b, $B_{c1} < B < B_{c2}$, the object is the mixed superconducting state, where the external field penetrates through quantized flux vortices, called fluxons. In zone c, $B > B_{c2}$, the field penetrates fully and the object is in the normal state [26].

dominates, $\sigma_1 \ll \sigma_2$. This simplifies the impedance to:

$$Z_s = i\omega\mu_0 \sqrt{\lambda^2 - \frac{i\rho_{vm}}{\omega\mu_0}} \quad (7.21)$$

where λ is the penetration depth. In the zero-field case, there will be no fluxons in the superconductor and therefore $\rho_{vm} = 0$, which simplifies this equation to $R_{s,0} \approx 0$ and $X_{s,0} = \omega\mu_0\lambda$.

The fluxon motion is what is driving the resistivity; the cavity mode fields displace the fluxons on the cavity surface, which dissipates power in the system. In order to minimize this motion, superconductors are often made intentionally with defects causing smaller grain boundaries; these "pin" down the fluxons to a smaller area, therefore limiting the power absorption from their motion. This process is called flux pinning. One can imagine at a high frequency mode, the induced surface currents of the mode will displace the fluxons from their equilibrium position in the lattice by just a small amount. Therefore one can ignore any

mutual interactions between fluxons during their displacement. Therefore we will consider the gross vortex motion an average effect of single fluxon model; the Gittleman-Roseblum (GR) model in this case [54] can be used to write out the vortex motion resistivity [93]:

$$\rho_{vm,GR} = \rho'_{vm} + i\rho''_{vm} = \rho_{ff} \frac{1}{1 - i\frac{\omega_p}{\omega}} \quad (7.22)$$

This introduces two new parameters: The flux-flow resistivity, ρ_{ff} , and the de-pinning frequency, ω_p . The flux-flow resistivity is associated with the flux-flow viscosity, η , by the relation:

$$\rho_{ff} = \frac{\Phi_0 B}{\eta} \quad (7.23)$$

where Φ_0 is the magnetic flux quantum. This highlights the fact that these fluxons behave very much like a fluid within the crystal lattice. In principle, η , should be independent of field, therefore it is commonly set by its value near the 2nd critical field:

$$\eta = \frac{\Phi_0 B_{c2}}{\rho_0} \quad (7.24)$$

The value of ρ_0 is typically a constant set by the manufacturing; in the case of Nb_3Sn , $\rho_0 = 0.71 - 9.7 \times 10^{-7} \Omega \text{ m}$ is used depending on the observed grain size; a smaller grain size means a higher resistivity [95].

The pinning frequency, ω_p , is very important to our purposes because it marks the boundary between fluxons behaving elastically, $\omega < \omega_p$, and vortex oscillations being purely dissipative, $\omega > \omega_p$; the operating mode frequency must be less than the pinning frequency of the superconductor. It is sometimes also called the de-pinning frequency. It is related to the flux-flow viscosity by the equation:

$$\omega_p = \frac{k_p}{\eta} \quad (7.25)$$

where k_p is called the pinning constant. One can see that from Equation 7.22, that ω_p can be obtained directly from the real and imaginary parts of the resistivity:

$$\omega_p = \omega \frac{\rho''_{vm}}{\rho'_{vm}} \quad (7.26)$$

This is useful for measuring the pinning frequency experimentally.

A more comprehensive solution to this problem would be to consider a full equation of motion for the fluxons that not only takes into account their harmonic motion from the resonance, but also the Lorentz force acting on them and a pinning force from the grain boundaries. The Lorentz force arises because each vortex is a loop of current, and the magnetic field of the mode as well as any external field will push them around. One can write out the magnetic field of each vortex outside its core as:

$$B(r) = \frac{\Phi_0}{2\pi\lambda^2} K_0\left(\frac{r}{\lambda}\right) \quad (7.27)$$

where $K_0(z)$ is the first Bessel function. The function diverges for $r < \lambda$, when in reality the entire inside of the fluxon can be approximated by:

$$B(0) \approx \frac{\Phi_0}{2\pi\lambda^2} \ln(\kappa) \quad (7.28)$$

This field will be normal to the surface of fluxon, whereas the current will subsequently be circling within the surface plane. This current can be approximated as constant from the core field:

$$I_{fluxon} \approx \frac{\Phi_0}{\mu_0\lambda} \ln(\kappa) \quad (7.29)$$

The actual Lorentz force will then depend on the angle between the normal vector of the fluxon and the direction of the net external field as it overlaps with the fluxon. For our purposes, the external field of the magnet will be significantly greater in magnitude than the resonant mode field, and therefore H_{mode} can be ignored when considering the strength of the force. In this case, because the external field is axial and static, the Lorentz force motion arises from the harmonic motion of the fluxons induced by the magnetic field of the resonant mode.

If we consider a right cylindrical cavity, fluxons on the end caps will have a normal vector parallel to the external magnetic field; a small displacement of the current loop in any direction will have maximal Lorentz force in the direction of the displacement. These fluxons will be in an unstable equilibrium, where their motion will be increased, subsequently

increasing the power dissipated in the cavity. Therefore, surface perpendicular to the external field will exhibit maximal losses.

On the other hand, a fluxon on the wall of the cavity, when displaced a small amount, will actually be in a semi-stable arrangement; displacements along the axis of the field will exhibit no Lorentz force, and displacements along the circumference of the cavity would result in procession about its original center if undisturbed. This makes superconducting surfaces parallel to the magnetic field have minimal losses! This is the motivation for using hybrid SRF cavities in very high magnetic fields for axion detection. By keeping surfaces perpendicular to the field made of copper, one avoids the maximal losses of the superconductor, while parallel surfaces should hypothetically still be superconducting with a lower resistance than copper.

Putting this together, one can write the 1-dimensional Lorentz force acting on the fluxon for parallel surfaces, with strength γ :

$$F_{Lorentz}(t, z) \approx \frac{\Phi_0 B_0}{\mu_0 \lambda} \ln(\kappa) e^{-\frac{z}{\lambda}} \cos \omega t = \gamma e^{-\frac{z}{\lambda}} \cos \omega t \quad (7.30)$$

where B_0 is the peak field of the magnet. The factor of $e^{-\frac{z}{\lambda}} \cos \omega t$ takes into account the field penetration into the surface, and harmonic motion of the fluxons from the mode.

The pinning force is a bit of a mess to approximate, as it depends on the position function of the fluxon from its origin in the surface plane, $\vec{x}(t, z)$; z represents the depth of the fluxon into the surface in this case. This force can be written:

$$F_{pinning}(t, z) = -\frac{2e\kappa^2 U_0}{\lambda} \vec{x}(t, z) \quad (7.31)$$

where e is the electron charge and U_0 is called the fundamental pinning energy. This energy can be found by finding $|\vec{x}_{max}|$, the distance from the fluxon origin that maximizes the pinning force, which should be correlated to the grain size of crystals in the superconducting material. This can be done by finding the zero of the first derivative of force, which can actually be written out for the 1-dimensional case:

$$F'_{pinning}(x) = \frac{8x^2 \xi^2}{(x^2 + \xi^2)^3} - \frac{2\xi^2}{(x^2 + \xi^2)^2} \quad (7.32)$$

The fundamental pinning energy can then be calculated:

$$U_0 = \frac{f_p \xi^2}{2ex_{max}} \left(\left(\frac{x_{max}}{\xi} \right)^2 + 1 \right)^2 \quad (7.33)$$

f_p is a fundamental pinning force constant roughly equal to 8×10^{-6} N/m in the case of Nb₃Sn.

The vortex equation of motion can then be written out in full:

$$\eta\lambda^2 \frac{dx(t, z)}{dt} = \epsilon \frac{d^2x(t, z)}{dz^2} + F_{Lorentz}(t, z) + F_{pinning}(t, z) \quad (7.34)$$

where all the variables have been defined except for ϵ ; this is called the vortex line elasticity.

It can be calculated from the expression:

$$\epsilon \approx \frac{\Phi_0^2}{4\pi\mu_0\lambda^2} \left(\ln(\kappa) + \frac{1}{2} + e^{-0.4-0.8\ln(\kappa)-0.1(\ln(\kappa))^2} \right) \quad (7.35)$$

This equation can then be solved under three boundary conditions:

1. The displacement is zero at $t = 0$, $x(0, z) = 0$
2. There exists a maximum depth where the displacement drops to zero, $x(t, Z_{max}) = 0$.
This will be $Z_{max} \approx 10\lambda$ where the field is roughly zero, or even better, $Z_{max} = t_{coating}$ the thickness of the superconducting coating.
3. The derivative of displacement with respect to depth must be zero at the surface of the superconductor, $\frac{dx(t, 0)}{dz} = 0$.

If the solution to this equation converges, one expects the plot of x vs $\frac{dx}{dt}$ to be an ellipse, representative of the harmonic motion. The sensitivity of the surface resistance to magnetic fields (Ω/T) can then be calculated:

$$S = \frac{\mu_0\omega^2}{2\pi^2} \int_0^{2\pi/\omega} \int_0^{Z_{max}} \frac{dx(t, z)}{dt} e^{-z/\lambda} \cos(\omega t) dz dt \quad (7.36)$$

This sensitivity is irrespective of the cavity wall geometry, which if we are to calculate a total surface resistance, must be integrated over. In this case, we will assume an axial symmetry

along the cavity, such that the angle of the cavity surface with respect to the axial direction is just a function of axial direction, $\theta(z)$ (this is not the surface thickness direction, which we already integrated over). Additionally, we will ignore the end caps and assume an empty cavity, such that the radius of the surface from the center-line can be defined, $r(z)$; this will allow us to integrate over circumferences of the wall in the azimuthal direction of the cavity ϕ . We only consider the axial component of the magnetic field, $B_z(z)$, along the length of the solenoid, however, one can define a cavity tilt parameter, $\Delta\theta_{cav}$, which is the angle of the cavity center-line with respect to the magnetic field center-line. In this way one can integrate over the surface of the cavity to get the power dissipated:

$$P_d = \frac{1}{2\mu_0^2 B_0} \int_0^{L_{cav}} \int_0^{2\pi} S B_z(z)^2 |\sin \Delta\theta_{cav} \cos \phi \cos \theta(z) + \cos \Delta\theta_{cav} \sin \theta(z)| r(z) d\phi dz \quad (7.37)$$

This can then be paired with the energy stored in cavity, U_{cav} , which is independent of any superconducting physics to get the quality factor according to equation 3.4.

Typically, the best way to perform these calculations is numerically using Python or Mathematica and importing a data file that contains a table of the cavity geometry values of interest; $\theta(z)$, $r(z)$, and $B(z)$.

Chapter 8

EXPERIMENTAL SRF CAVITY WORK AT LLNL

This chapter covers the experimental work on superconductors and SRF cavities I performed at LLNL. This includes DC resistivity measurements of superconductors, as well as some prototype cavity measurements. The highlight of this work is my first published paper that I was first and corresponding author for; it makes use of the multi-mode decomposition technique, outlined in Chapter 6, to characterize a NbTi SRF cavity.

The work done on superconducting cavities at Lawrence Livermore National Laboratory (LLNL) had two goals: Qualify potential superconducting materials for use in ADMX operating conditions, and demonstrate the efficacy of future cavity geometries such as the hybrid SRF cavity and clam-shell cavity design. This meant being able to test materials under cryogenic and high magnetic fields in order to measure materials' critical temperature and fields. On the RF side, being able to measure the RF impedance properties, such as the pinning frequency would also be useful. Finally, an actual direct measurement of the increased losses from a cavity end cap perpendicular to field, versus a parallel wall, would give credence to the hybrid cavity concept.

8.1 *PPMS system*

The bulk of this work was done in a Quantum Designs Physical Properties Measurement System (PPMS). Pictured in Figure 8.1 and Figure 8.2, the PPMS is an exchange gas cryostat with a 26.5 mm bore, equipped with a superconducting magnet capable of fields up to 16 T applied axially along the bore; the minimum temperature is 1.8 K. Because all of the potential superconductors had critical temperatures well above 1.8 K, this cryostat was fine for our purposes; field performance at milli-Kelvin temperatures could be extrapolated using

Equation 7.1.



Figure 8.1: Quantum Design's Physical Properties Measurement System (PPMS) from the outside.

The uniform field region height of the magnet was approximately 40 mm. The field deviation of magnet was directly measured with a hall probe as a function of distance from the field center, as shown in Figure 8.3 and Figure 8.4; past 60 mm, the field deviates beyond 20% of its center value. This deviation tends to go down for higher magnet strength settings, at least for distances less than 60 mm from the field center.

The proprietary use of the PPMS system is material characterization under various field and temperature conditions via DC resistance measurements. Samples are mounted on "pucks" that are inserted into the PPMS. As shown in Figure 8.2, the sample of the puck's surface is actually located about 40 mm below the field center, deviating the field value by roughly 5-15% depending on the set strength. Depending on the measurement setup, care was taken to account for these deviations accordingly.




Figure 8.2: Diagram of the PPMS bore.




Figure 8.3: Absolute field deviation of the 16 T PPMS magnet as function of distance from field center. Data was taken at three different magnet field strengths: 0.5, 1, and 2 T.

8.2 DC measurements

Proprietary to the PPMS's uses, DC resistance measurements of superconductor samples provided a fast and effective way of measuring their critical field and temperature. A typical




Figure 8.4: Percent field deviation of the PPMS magnet as function of distance from field center.

superconductor exhibits a finite DC resistance above its critical temperature, which abruptly transitions to zero at the critical temperature, T_c , so by cooling the sample and continuously taking resistance measurements, this transition can be observed, and T_c measured, as pictured in Fig 8.5.

By ramping the PPMS magnet, one can force the sample, if a Type II superconductor, into the intermediate vortex state, above the first critical field, B_{c1} . Because one is measuring a DC resistance, one doesn't actually expect to see a change at this first transition, because the frequency of the applied current source is so low, it won't move the fluxons very much, and the "swiss-cheese" structure of the superconductor will let the current through unimpeded. It is only when the density of fluxons is great enough that large portions of the superconductor are normal and very few superconducting DC "paths" remain that the DC resistance will start to increase. Eventually, when the sample has gone fully normal, at B_{c2} , this DC resistance plateaus; this resistance should also be relatively equal to the resistance of the sample measured just above T_c . This process is outlined in Fig 8.6. Since we are qualifying materials that will be operating in the intermediate phase no matter what, measurement of

this second critical field is sufficient.




Figure 8.5: A typical DC resistance versus temperature graph for determining T_c . Resistance measurements are periodically taken as the sample is cooled through its critical temperature, At the critical temperature, the superconductor's resistance abruptly drops, making it easy to read off T_c . The residual resistance is from the leads connecting to the sample, in this case, being only a two-wire measurement. This particular sample is a coating of NbTiN made at LLNL that had an abnormally low T_c of ≈ 2.8 K; typically it is between 10-12 K. This transition temperature was also not particularly stable; upon cycling the sample back through critical temperature, several resistance measurements imply it was superconducting up to 3.1 K. All of this indicated that the coating was compromised by impurities in the deposition process.

8.2.1 4-wire Measurements

A 4-wire measurement is a much more precise way of measuring resistance across a particular load, such as a superconducting sample, rather than a classical, 2-wire digital multi-meter (DMM) measurement. In the two-wire measurement, the DMM supplies a known current to two wire leads that the sample is connected in series to, and then measures the voltage drop between these two leads with a voltmeter in parallel; the resistance, by Ohm's law, is then the ratio of this voltage to the supplied current (the equivalent circuit is shown in Figure 8.7). However, being connected in series, this includes the resistance of the leads




Figure 8.6: A typical DC resistance versus magnetic field strength graph for determining B_{c2} . Resistance measurements are periodically taken as the magnet is ramped through the samples' critical fields. For type II superconductors, such as NbTiN in this case, the resistance will start to increase and plateau into a fully normal state above B_{c2} . The residual resistance is from this being only a two-wire measurement. This particular sample is a coating of NbTiN, the same used in Figure 8.5, made at LLNL. It had an abnormally low B_{c2} of ≈ 5.3 T; typically it is in excess of 15 T. The sample was cycled back through its critical field when the magnet was ramped down, generating the hysteresis loop structure; this is from trapped flux from the initial ramp shifting and increasing the measured resistance on the ramp back down. Both the low B_c and large hysteresis loop, indicate that the coating was compromised by impurities in the deposition process.

themselves, rather than just the sample. This can be fine for the purpose of determining critical temperatures and fields, as long as the lead resistance is relatively small compared to the change in resistance of the sample upon superconduction; this is the case for Figure 8.5 and 8.6. But if one's goal is precisely measuring the sample resistivity as a function of temperature and field, the lead resistance must be calibrated out. This is also useful for measuring near-zero resistances, like superconductors, where the leads resistance will dominate.

The solution to this is relatively simple: Separate the fixed current source leads from the voltmeter measurement leads that connect to the sample, giving it 4 connected wires. This circuit is pictured in Figure 8.8. Because a voltmeter is designed to have an exceedingly




Figure 8.7: The circuit diagram for a 2-wire resistance measurement. A fixed current source and voltmeter are connected in parallel internally within the DMM. Two wire leads are then connected in series with the sample load of interest.

high impedance, very little current passes through the leads connected to the voltmeter, hence their individual voltage drop is negligible. This means the voltage measured across the voltmeter is practically the same as voltage drop across the sample, whereas it wasn't in Figure 8.7.




Figure 8.8: The circuit diagram for a 4-wire resistance measurement. One set of leads connects the sample to a fixed current source, while another connects to a voltmeter reading the voltage drop solely across the sample.

8.2.2 Mounting Samples

Samples were mounted to proprietary PPMS pucks that had 2-4 sets of 4-wire solder contacts (see Figure 8.9); This meant a single puck could hypothetically take 2-4 samples at a time if space permitted. These pucks had a sample space just over 1 cm^2 , so samples had to be chosen and sometimes cut to size to accommodate. It was best for a sample, whether bulk or a coating, to have an elongated rectangular shape, creating a wire-like current path through it. Four very thin (20 micron diameter) gold wires would be epoxied to the sample with two leads on either short-end of the sample, one short-end being the positive voltage and current side, and the other end the negative lead side. These wires would be placed perpendicular to the long side of the sample, with the inner two leads being the voltage connection. This was a silver conductive epoxy, called H20E EPO-TEK from Ted Pella Inc, that cured at roughly 150 degrees Celsius for 5 minutes. The use of silver and gold is to still minimize the lead resistance as much as possible. Once cured, the gold wire leads would be soldered to the 4 puck contacts appropriately. A small amount of N-grease was used to hold the sample to the puck surface. Throughout this process a microscope was used to see the very thin wires and where their tips were actually being placed. Because they are very delicate, it was common for a gold wire to break at either the epoxy or solder stage of mounting, or even thermally upon cooling in the PPMS. If this happened late in the mounting or cooling process, we could still run a two-wire measurement; Often the 4 wires were just redundant connections rather than critical to a 4-wire measurement. Once a sample was properly mounted and the puck connectivity was confirmed, it would be inserted into the PPMS for cooling and subsequent measurement.

8.2.3 Operating Procedure

Once samples are mounted and loaded into the PPMS, the PPMS is cooled overnight. This is done in two stages typically; a rapid cooling, 10 K/min, to just above the critical temperature, typically around 10-20 K, and a slow cool, 1 K/min, from there to base fridge temperature,




Figure 8.9: A PPMS puck for 4-wire resistance measurements of samples. This particular puck has 3 sets of 4 solder contacts ($V+$, $V-$, $I+$, and $I-$) to handle 3 simultaneous sample characterizations. The central square where samples would be placed is about 1 cm^2 .

2 K. The sample would then be given some extra time to ensure full thermalization. Usually a procedure starts at the end of a work day, is set to cool overnight, and by morning it has been sitting for several hours at the base fridge temperature (see Figure 8.10). Depending on the smoothness of the cool down through the critical temperature and the data quality, a subsequent warming might be done to get a clearer transition temperature measurement. Then magnet ramps are performed accordingly to measure the critical field. In the case of Figure 8.10, the bulk NbTi sample went through 3 ramps; the first two of these we failed to measure the critical field because a wire broke in situ, giving intermittent readings of the resistance. After warming, and re-mounting the sample, the third ramp produced the needed data.

8.2.4 Results

In Table 8.1, I highlight three main DC measurement results: NbTiN coatings made at LLNL, a bulk chip of NbTi metal used to machine the test cavity in following sections, and Nb_3Sn coatings on bronze substrate given to us by FSU. Unfortunately, we were unable to make measurements of Fermilab's Nb_3Sn coating on pure Niobium substrate, because we




Figure 8.10: A plot of the PPMS operating temperature and magnetic field versus time during a 33 hour measurement of a bulk NbTi sample. Samples are typically cooled overnight, and magnet operations begin the following day. Sample may lose connectivity due to thermal or magnetic strains, at which point the system is warmed and sample is removed for re-mounting (This was the case after the 2nd magnet ramp in the plot). Warming above the critical temperature in between magnet ramps can also be useful for expelling any trapped flux in the sample.

couldn't get any samples of the right size made; this was the coating process used to make the Sidecar SRF tuning rod described in Chapter 9. Similarly, the YBCO samples acquired were too large to fit within the PPMS, and remain untested for the next LLNL student.

At the beginning of my stay at LLNL, a local deposition shop on site made several NbTiN samples in their system. Unfortunately, we discovered that the base pressure of their system was too high 10^{-5} Torr than what is needed to produce clean coatings ($\approx 10^{-9}$ Torr). This resulted in very low T_c and B_{c2} measurements, as shown in Figure 8.5 and 8.6. Without the proper deposition system and cheaper alternatives available, these nitride coatings were ultimately abandoned for our purposes.

During the manufacturing of a bulk NbTi test cavity, we were able to ask the machinist to chip off a small portion of the square rod stock for DC testing. This cavity will be discussed in the following two sections in detail, but the DC results of this chip are included here:

Material (Source)	T_c (K)	B_{c2} (T)
NbTiN coating on silicon (LLNL)	2.8	5.3
Bulk NbTi (American Elements Alloy Nb-55% Ti-45%)	8.5	> 14 T
Nb ₃ Sn on Bronze (FSU)	10.3	> 14 T

Table 8.1: Critical temperature and fields of various materials measured at LLNL. The NbTiN coating was produced at LLNL. The NbTi was acquired from American Elements. The Nb₃Sn was supplied by Lance Cooley’s group at Florida State University (FSU).

$T_C \approx 8.5$ K and $H_{c2} > 14$ T. We were unable to determine the full critical field value because it was not permitted to ramp the magnet above 14 T to its maximum at 16 T. As shown in Figure 8.11, the NbTi sample had begun transitioning to a normal state, but had yet to plateau; typically Nb-45% Ti-55% alloys such as this have critical fields between 14-15 T, so this agrees well with that figure.




Figure 8.11: Measuring the B_{c2} of a bulk NbTi sample. $R1$ and $R2$ denote two different resistance bridges in the PPMS, that each were connected to a NbTi sample. Because the magnet was limited to 14 T, it can only be said that $H_{c2} > 14$ T. This is in agreement with similar alloys that exhibit an $H_{c2} \approx 14 - 15$ T.

The Nb₃Sn sample highlighted in Table 8.1 was made by Dr. Wenura Withanage, working under Dr. Lance Cooley at the time, and uses an innovative coating process allowing for a bronze substrate (see Figure 8.12). This was of interest to us at ADMX, because a bronze substrate would have the heat capacity properties of copper, as well as a the lower price tag when compared to a pure Niobium substrate. As expected for a Nb₃Sn coating, the critical field was above 14 T, and therefore could not be precisely measured; Typically it is around 18 T.



Figure 8.12: Nb₃Sn coating deposited on bronze substrate. This deposition work was done at Florida State University by Lance Cooley's group, see Ref. [121] for details.

This sample was also tested with a variable angle rotating sample stage made by Quantum Designs (see Figure 8.13). This allowed an operator to rotate the sample surface with respect to the magnet field, rather than just a perpendicular arrangement. The operating procedure for this becomes a bit longer; one has to perform a full ramp up and down, ensure there's no trapped flux, then rotate, and then repeat. For this reason, we only tested at 5 angles: 0, 30, 45, 60, and 90 degrees. As shown in Figure 8.14, the sample actually remained

superconducting in the parallel arrangement (90 degrees in the figure) all the way to the 14 T limit. The other angles accordingly start to transition at lower field values with the perpendicular arrangement having the lowest transition field strength at ≈ 12.5 T; this is exactly the angular dependence one hopes to see.




Figure 8.13: The Quantum Designs PPMS rotating puck stage. This allows for samples to be rotated with respect to the magnetic field axis from its typical perpendicular arrangement.

8.3 PPMS cavity design

In order to perform RF measurements within the system, a cavity and readout system had to be designed to fit within the PPMS bore. This design was used to machine cavities out of a variety of materials; cavities made out of copper, aluminum, and bulk NbTi have been made, as pictured in Figure 8.15A-C. This enables a direct comparison of material performance. The clam-shell design lends itself well to being coated, although no cavity has been coated yet as of the time of writing. The clam-shell halves are joined by 4 4-40 screws, and are aligned by 3 alignment pins pictured in Figure 8.15B-C.

The bore, being very narrow, limited the internal diameter of any cavity that could fit




Figure 8.14: The resistance versus magnetic field of Nb_3Sn on bronze substrate for multiple angles with respect to the field axis; 0 degrees corresponds to the sample being perpendicular to the field axis, while at 90 degrees the sample surface is parallel. Successive Magnet ramps at 5 different angles were made. Note that the parallel arrangement does not begin its superconducting transition, while the other angles with perpendicular components begin to transition above ≈ 12.5 T.

to 23 mm, which put the frequency of the TM_{010} mode at ≈ 10 GHz. Additionally, the internal height was limited to be 53 mm because of the magnet coil height; this would put the bottom end cap within ≈ 25 mm below the field center, whereas the top end cap would be ≈ 28 mm above. Since most measurements would be taken at high field strengths, this was deemed within the uniform field range, and close to symmetric about the field center such that field deviations could be ignored. As shown in Table 8.2, limiting the cavity length had the added benefit of keeping the end cap geometric factor within an order of magnitude of the wall geometric factor; any longer and the walls would dominate over any end cap contributions (this was important for making multi-mode decomposition measurements).

The geometric factors shown in Table 8.2 were simulated using the Ansys HFSS software package. The simulated cavity modes' field structure and surface currents (see Figure 8.16) were used to calculate the surface and volume integrals for geometric factor in Equation




Figure 8.15: PPMS clamshell cavity machined from a bulk NbTi square stock. A) The NbTi cavity assembled next to aluminum prototype. B) A profile of both NbTi halves side by side. C) One half of the copper clamshell close-up. D) A close-up of the cavity top without and with fixed pin antennae attached; additional holes are for mounting to probe.

6.29. The set of nine geometric factors correspond to the three modes of interest, TM_{010} , TM_{011} , and TM_{012} , as well as the three sub-surfaces of interest, top end cap, bottom end cap, and walls. One may notice that the top and bottom end caps do not differ very much in values; nonetheless it was important to quantify this difference as the top end cap contained the antenna port holes. Python software enabled one to combine these geometric factors in various combinations, most commonly the total end cap geometric factor, as well as their associated error.

The error in these G-factors were calculated through a parametric analysis according to the cited machine tolerance, $\Delta X = \pm 0.005''$. The height and radius were varied by three points, $X + \Delta X$, X , and $X - \Delta X$, generating 9 configurations of the cavity design; the nine geometric factors were subsequently calculating for each configuration. The final uncertainty

	TM_{010}	TM_{011}	TM_{012}
Walls	457.0 ± 2.2	494.6 ± 3.4	549.9 ± 5.5
Top End Cap	3846 ± 41	2045 ± 15	2214 ± 6
Bottom End Cap	4127 ± 55	2029 ± 26	2192 ± 21
Total Endcaps	1991 ± 38	1019 ± 17	1101 ± 12
Total	371.7 ± 9.3	332.9 ± 7.0	366.8 ± 6.2

Table 8.2: Geometric factors for clamshell cavity design for first three TM modes by interior surfaces. These were calculated from simulation in the HFSS software package and are expressed in Ohms.




Figure 8.16: A field overlay of the cavity modes' surface currents in AnsysTM HFSS software. A) The TM_{010} mode. B) The TM_{011} mode. C) The TM_{012} mode. Scale bar for the magnitude of the surface currents picture left in A/m.

was calculated according to Equations 6.64, 6.65, and 6.66. The fractional uncertainty was kept to less than 5% as shown in Figure 8.17. Unfortunately, subsequent calculation of the weight matrix and its determinant for the three possible 2-mode decompositions showed that

the TM_{011}/TM_{012} mode combination had a statistically zero determinant and couldn't be used.




Figure 8.17: The estimated fractional error in geometric factor by mode and subsurface. Note that the end caps factor is the sum of the top and bottom end cap contributions, and the total geometric factor is the sum of the end caps and walls. A) The error in the TM_{010} mode geometric factors. B) The error in the TM_{011} mode geometric factors. C) The error in the TM_{012} mode geometric factors.

Two antenna ports on the top of the cavity enabled one to perform two-port measurements and characterization (see Figure 8.15D). Fixed pin antennae attached with low-profile 4-40 screws; these antennae were cut such that $\beta < 0.01$, making the cavity weakly coupled and direct measurement of Q_0 could be made. The cavity was mounted to a custom multi-use probe, shown in Figure 8.18C. The fixed pin antennae were attached by SMA connections

with an SMA to SMP junctions several inches above, with stainless steel semi-rigid coaxial cable up to the top of the probe where SMA vacuum feedthroughs enabled connection to the warm space. A 28-pin DC connector next to the SMA feed-throughs (Figure 8.18B) connected to a temperature sensor located on the top of cavity. A puck screwed into a center-line hole on the bottom of the cavity, which thermally sunk the cavity to the PPMS cryostat as well as connecting it to the system's internal temperature sensor. A vacuum flange at the top of the probe sealed it to the PPMS bore with an ISO-K clamp.

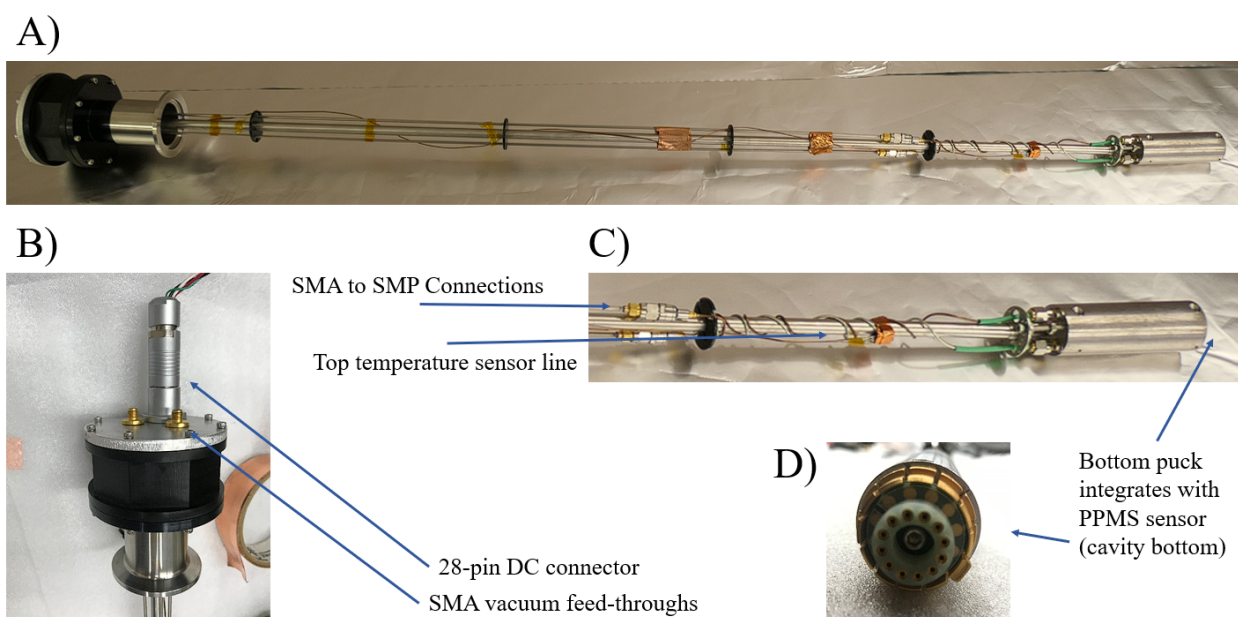


Figure 8.18: The PPMS readout probe with an aluminum clamshell cavity mounted after a test. A) The entire probe with NbTi cavity mounted on bottom. B) The probe top with DC and SMA vacuum feed-through connections. C) The probe bottom with cavity, SMA to SMP connections, and temperature sensor line. D) Cavity bottom puck that ensures thermal connection to PPMS internal temperature sensor.

8.4 Results of a bulk NbTi clamshell surface decomposition

Using the PPMS clamshell design outlined in the last section, as well as the novel multi-mode decomposition technique outlined in Chapter 6, we were able to directly show the

losses in an SRF cavity were dominated by the perpendicular field surfaces, the end-caps, while the surfaces parallel to the field remained superconducting, with minimal loss. The superconducting NbTi cavity was machined from a Nb-45% Ti-55% square stock with 99% purity purchased from American Elements by LLNL. An annealed copper clamshell was used as the control in this measurement.

Each clamshell cavity was cooled in the PPMS to 2 K in zero field by cooling to 10 K at 10 K/min, then to 2 K at 1 K/min and then held for ≈ 10 minutes to allow the cavity to fully thermalize. This cool down shifts the resonant frequencies from room temperature by about 25 kHz upward. The magnet was then ramped in steps (0.1 T steps from 0 to 1 T and 1 T steps from 1 to 10 T), and held in between until the mode frequencies and Q factors settled to constant values (≈ 1 minute), at which point the resonant peaks were recorded from the VNA for the three *TM* modes of interest (TM_{010} , TM_{011} , and TM_{012}).

The zero field Q factor for the NbTi cavity was measured to be $\approx 70,000$ at 2 K directly after the cavity was machined. The cavity was then sent to Jefferson Laboratory for a cleaning and annealing process. Post-annealing, the cavity had a zero field quality factor at 2 K of nearly 1 million; annealing and cleaning has a significant effect on the zero field Q. For comparison, the copper cavity had a Q of about 180,000 at 2 K. The critical temperature of the NbTi was observed to be approximately ≈ 8 K based on the beginning of the superconducting transition for both the pre-annealed and post-annealed cavity; this is actually slightly under the typical 10 K usually cited for NbTi [23]. The mode quality factor as a function of applied magnetic field is shown in Figure 8.19. As expected, the NbTi cavity Q degrades exponentially as a function of field. Eventually these quality factors were then converted to total mode resistances via Equation 6.49 shown in Figure 8.20.

In order to test the consistency of the multi-mode decomposition, a two-mode decomposition to the walls and total end caps contributions was performed with combinations of the three modes measured. The TM_{010}/TM_{011} and TM_{010}/TM_{012} decomposition surface resistances agreed very well and confirmed the hypothesis that the end caps were the primary contribution to the surface resistance in field for the NbTi cavity (Figure 8.21).




Figure 8.19: NbTi versus copper clamshell in PPMS system: The cavity quality factor for the first three TM modes as function of applied magnetic field. Data was taken at a temperature of 2 K.

The TM_{011}/TM_{012} combination unfortunately did not produce a physical result; this is because the geometric factors of the two modes were nearly identical, making the determinant of the weights matrix very nearly zero as discussed in Chapter 6 (almost two orders of magnitude smaller than the other two combinations), inflating errors when taking the inverse (see Table 8.3). Having determined what caused this mode-combination to fail, it was excluded from the final data set.

These results can be compared to the results of the QUAX hybrid NbTi cavity in Ref. [7] and Ref. [41]; the NbTi film is thick enough (3-4 Microns) to be in the bulk limit and was tested in a similar temperature and field regime. The first analysis in Ref. [7] estimates a similar resistance for their copper cone end-caps at 4 K of 5 m Ω . However they must calculate




Figure 8.20: NbTi versus copper clamshell in the PPMS system: The total mode surface resistance for the first three TM modes as function of applied magnetic field.

their NbTi resistance using this copper resistance value with their methodology; they get a value of 20 ± 20 mΩ for NbTi in zero field, much higher than our measurement at zero field, and closer to the resistance we saw at 5 T (Figure 8.21). As for higher field values, the quality factor drops off at a similar rate, starting at 1.2 million at zero field to around 35,000 at 5 T. It is a bit curious and discouraging considering our cavity is entirely NbTi while the QUAX is a hybrid design. This can be explained partially by operating at 4 K versus 2 K, a significantly different cavity geometry, and lower frequency regime. The results of the analysis found in [41] are in better agreement with ours; they cite a Q of 1.2 million at 4 K and 0 T versus our 1.4 million at 2 K, and this dropped to 200,000 at 5 T versus our 100,000. This implies their cavity must have had a lower total resistance at 5 T because of the hybrid construction.

Mode Combination	$\text{Det}(C_{\text{ms}})$	$\Delta(\text{Det}(C_{\text{ms}}))$
TM_{010}/TM_{011}	0.140	0.016
TM_{010}/TM_{012}	0.146	0.016
TM_{011}/TM_{012}	0.006	0.013

Table 8.3: PPMS clamshell cavity multi-mode decomposition: determinant of weight matrix and associated error for the three possible mode combinations.

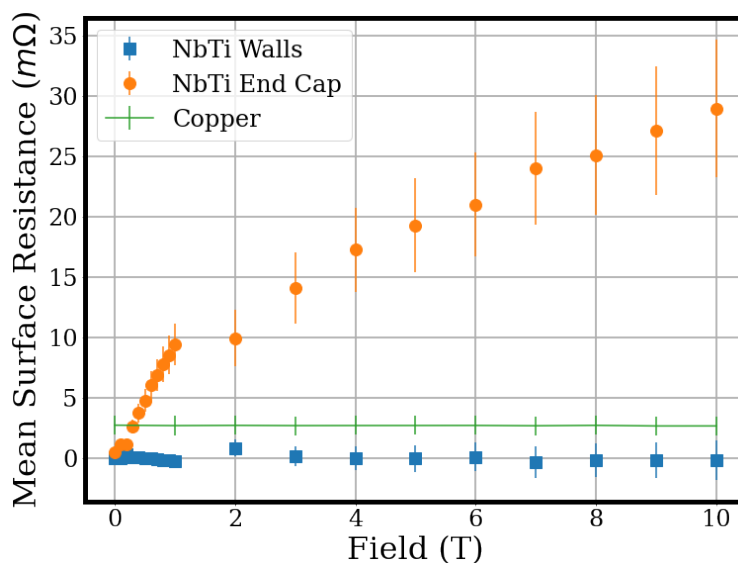


Figure 8.21: The mean surface resistance contributions of the walls and end caps of the NbTi cavity based on two different two-mode decompositions: TM_{010}/TM_{011} and TM_{010}/TM_{012} . This is compared to the mean copper cavity resistance, averaged over surface type for the same mode combinations.

As discussed in Chapter 6, the two sources of error in this measurement are the quality factor and geometric factor. The geometric factor errors in the PPMS cavity design were discussed earlier and reported in Table 8.2. The quality factor uncertainty was determined from the estimated covariance matrix of the Lorentzian magnitude fit; this is an output of the `Scipy.optimize.curve_fit` package. Typically this error was constant with ΔQ on the

Parameter	Variable	Maximum Uncertainty
Quality Factor	ΔQ	4.2%
Geometric Factor	ΔG_{ms}	5.0%
Total Mode Resistance	ΔR_m	6.6%
Weight matrix	ΔC_{ms}	4.9%
Inverse weight matrix	ΔC_{ms}^{-1}	12.6%
Surface Resistance	ΔR_s	5.7m Ω

Table 8.4: A summary of the maximum estimated uncertainties for the surface decomposition measurement setup. They are expressed as percentages of their nominal values with the exception of surface resistance, since many values measured were effectively zero while the cavity was superconducting. Note that ΔC_{ms}^{-1} and ΔR_s does not include the TM_{011}/TM_{012} decomposition because of near-zero determinant.

order of 10^2 . Any measurements with $\frac{\delta Q}{Q} > 0.10$ were considered outliers, and the fit was subsequently reviewed by plotting the data; this occurred only seven times in the data set, all for the superconducting NbTi cavity at low field. Upon further inspection, it was clear that the fit function failed because of the large window size compared to the very narrow resonant peak width; this could be corrected for by reducing the window size and by removing some of the off-resonant data from the fit in post-processing.

The maximum estimated uncertainty for each parameter along the analysis chain for the decomposition is then reported in Table 8.4. As you can see, the inverse weight matrix compounds a very high increase in fractional uncertainty from the amount of multiplications needed in its construction. Since the surface resistance was statistically zero for the NbTi wall measurements, it was better to report the maximum absolute uncertainty, which was in the NbTi endcaps, $\pm 5.7 \text{ m}\Omega$, rather than a misleading high fractional uncertainty in the NbTi walls.

A final check on the efficacy of this method was to look closely at the control copper cavity, and its variance for each mode combination decomposition. Being a non-magnetic material, one expects to not see a difference between the end-cap and wall resistivity (resistivity referring to resistance per unit surface area). The resistance may be slightly different based on the differing surface areas, but overall this should be captured by the geometric factor. As you can see in Figure 8.22, the TM_{010}/TM_{011} combination does follow this pattern; however, the TM_{010}/TM_{012} does not, indicating it might not be a stable or effective mode pairing. Nonetheless, this observable difference between the two surfaces in this mode pairing is on the order of the error bars themselves, and is far smaller than the difference observed for the NbTi cavity surfaces in Figure 8.21. Overall, this gives confidence in the efficacy of the decomposition method for its future use.

The final take-away is that this result supports the construction of hybrid SRF cavities. The NbTi walls had a lower resistance than copper up to 10 T, and its overall degradation seems to only be from the increased resistance in the end caps. If one takes the wall resistance with maximum positive error, giving the highest possible resistance, one sees that it is still lower than copper by a factor of ≈ 2.3 . Using this minimum improvement factor to project a Q for this cavity if it were to be a hybrid, with copper end caps and NbTi walls, one expects a minimum Q of 315,000 up to 10 T, which is a 67% improvement over the all copper PPMS clamshell. The details of this measurement were reported in Review of Scientific Instruments in March 2023 [30].

8.5 Design and status of a Hakki-Coleman PPMS Cavity

Another cavity design compatible with the PPMS was pursued for precise RF surface impedance measurements of planar superconducting samples. Pictured in Figure 8.23, this cavity is loaded with a dielectric sapphire crystal cylinder, which focuses the electromagnetic field of the mode away from the wall region and onto the end cap regions above and below the cylinder. Additionally, the increased permittivity of the dielectric ($\epsilon \approx 10$) shifts the mode frequency downward, allowing one to probe lower frequency regimes with smaller physical

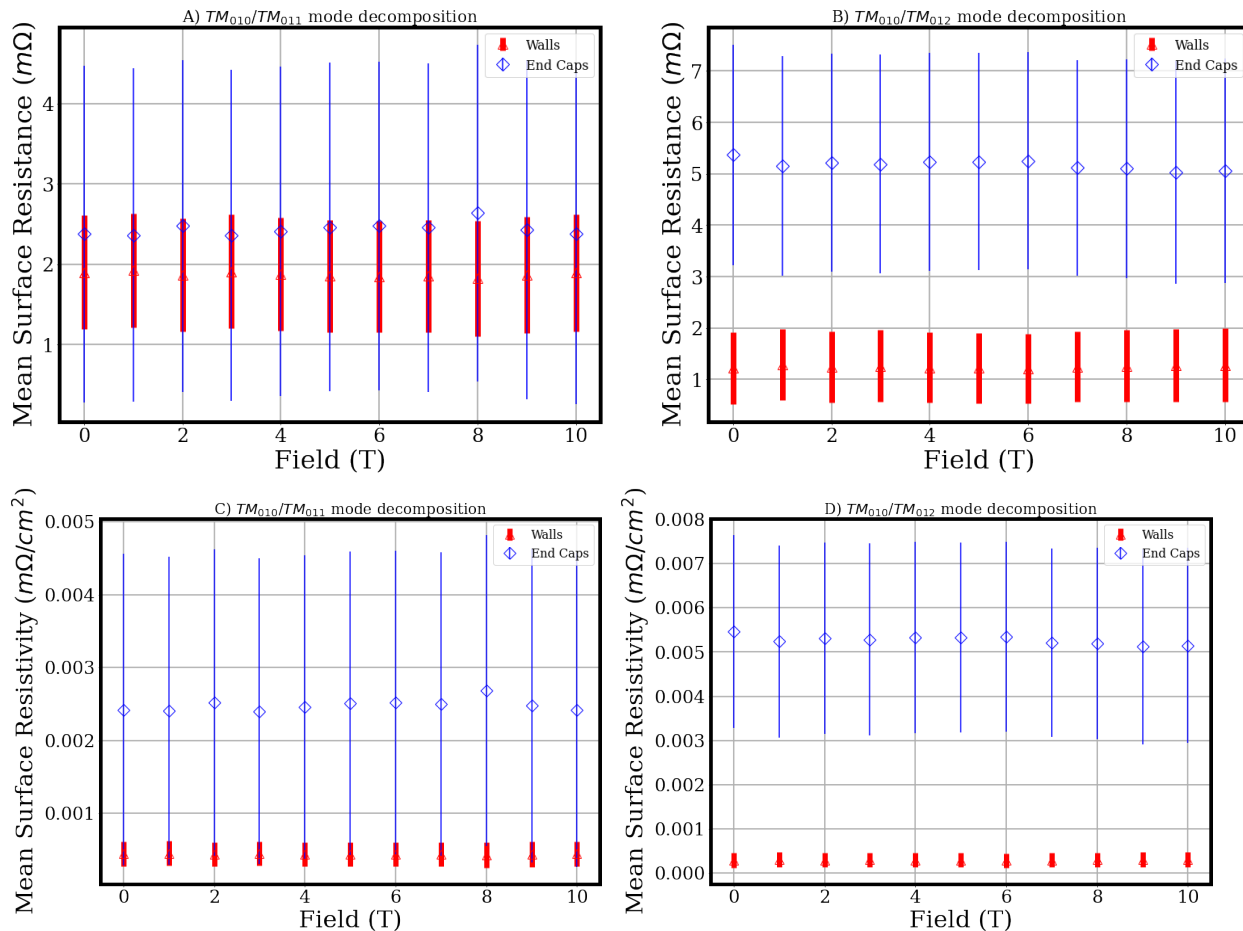


Figure 8.22: The two two-mode surface resistance decompositions for the pure copper clamshell cavity vs field. A) Surface resistance of the cavity sub-surfaces for the TM_{010}/TM_{011} combination. B) The same plot for the TM_{010}/TM_{012} combination. C) The corresponding resistance per unit surface area for the TM_{010}/TM_{011} combination. D) The same plot for the TM_{010}/TM_{012} combination. These bottom plots account for the differing surface areas of the walls and end caps. Note the differing error bars are due to the higher uncertainty in the end cap geometric factors.

cavity dimensions. This is not a new technique, and is used for high sensitivity impedance measurements [8, 37]. This is similar to a Hakki-Coleman type resonator [62]; however, it was made into a clam-shell to accommodate the small bore size of the PPMS and did not contain any springs or vernier drive that typically clamp the cylinder to the sample surface (more on this later).




Figure 8.23: An aluminum clamshell dielectric loaded cavity for precise surface impedance measurements of samples. A) The two halves have circular insets for the sapphire cylinder and rectangular inset below for inserting sample blocks. The bottom circular inset acts as a mask to remove edge defects from the exposed surface. B) The cavity loaded with sapphire rod and aluminum sample block. Superconducting coatings would be sputtered to a substrate block of similar size and can be exchanged easily. C) The cavity contains the same exterior features for compatibility with the PPMS. D) A rising cylinder lifts the shorter cavity to center the sample surface in the PPMS magnet.

The TE_{011} mode is the principle mode of operation because the circular currents are easily focused away from the walls, and they also will generate maximal vortex motion on the sample surface. Since this sample surface is also oriented perpendicular to the external field of the PPMS, this gives the maximum expected surface resistance for a sample. A sapphire crystal of 4.5 mm diameter was chosen to give an expected mode frequency of 8.5 GHz. This was chosen over an 8 mm cylinder design, similar to the resonator Ref. [8], which had an operating frequency of 14.5 GHz. The height of the cavity was reduced from 53 mm

to 13 mm, with two 1 mm circular insets (Figure 8.23A-B) that cradled the 15 mm length rod; this was intended to reduce the wall surface area, and subsequently increase the geometric factor of the walls, making it a negligible contribution to the quality factor. The bottom inset additionally exposed a uniform circular region from the center an inserted sample block; this masked any edge defects on the sample coating, and controlled the geometry of the exposed sample region.

To make this compatible with the PPMS system, all of the exterior features were kept the same as the original clam-shell design (Figure 8.23C). Since this cavity had the reduced length, a spacing cylinder was attached below the cavity, which lifted the sample surface to the exact center of the magnet by 40 mm (Figure 8.23D). This also kept the overall length of the assembled system the same, so that the PPMS puck and probe seated correctly within the bore.

The surface impedance measurement worked as follows. The Quality factor and resonant frequency of the mode are monitored by measuring and fitting to the S parameters periodically, with T and H_{ext} being varied. The coupling of the antennae again are kept weak such that $Q_L \approx Q_0$. Additionally, running the network analyzer with low power input signals minimized the RF amplitude of the mode magnetic field, so Z_s was unaffected by the measurement itself, which Niobium and Nb₃Sn can be very sensitive to [63, 82]. The total geometric factor of this cavity, in principle, is dominated by the sample surface, such that all other sub-surfaces are negligible. This then enables us to write down the surface impedance of the sample by the relation:

$$Z_s = R_s + i\Delta X_s = \frac{G_s}{Q} - i2G_s \frac{\Delta f_0}{f_0} - C \quad (8.1)$$

where Δ represents the variation with respect to the reference zero-field reactance, $X_{s,0} = \omega\mu_0\lambda$, and frequency, and C represents a background constant. The background constant is determined by calibration with a matched metal sample with the cavity in field. After this, a superconducting sample can be loaded, and zero-field calibration can be made: $R_s(H = 0, T \rightarrow 0) \approx 0$ as well as $R_s = X_s$ for $T > T_c$ are used as fixed point relations to center the

absolute impedance values (this is a real quasi-particle conductivity, also called the Hagens-Rubens limit). The data is then taken by fixing the magnet strength at several values and then varying the temperature for each field strength at a fixed rate from the fridge base temperature to slightly above zero field critical temperature. Once the absolute complex impedance data is in hand, Equation 7.21 can be used to find the complex resistivity values, ρ'_{vm} and ρ''_{vm} , the ratio of which, $r = \rho''_{vm} / \rho'_{vm}$, is directly related to the pinning frequency, $r = f_p / f = \omega_p$. As stated in Chapter 7, this pinning frequency provides the upper frequency limit for elastic vortex motion in the superconductor.

This project unfortunately hit a standstill during my PhD tenure due to several design flaws. The principle design flaw was that the rod was too narrow and the walls were too long to achieve the desired geometric factor arrangement. As you can see from Table 8.5, the cavity is dominated by the mask region of the bottom end cap, rather than the sample surface as intended. Furthermore, the top end cap and walls themselves, having a smaller geometric factor, contribute more loss to the Q than the sample. The cavity used in Ref. [8], which was 8 mm in diameter and 5 mm in height, had an incredibly high wall geometric factor of 420,00 Ω ; this is due to the shorter wall height as well as the increased focusing power from a larger diameter crystal. This design also had a top end cap with a lower geometric factor than the sample surface, but the ratio between the two in our design is much closer to equality. The solution to this is relatively simple: Build a new shorter cavity with a wider crystal cylinder. However, this does drastically change the frequency regime of the measurement; our original research goal was to measure the impedance in the frequency regime below 10 GHz.

In practice, this clam-shell cavity design had several other flaws. As stated earlier, because the design lacked a spring/compression system to ensure the crystal was pressed against the sample surface, thermal contractions likely caused a gap to develop upon cooling. This manifested in unpredictable mode structures due to the gap. It also would further increase the sample geometric factor, further decreasing its contribution to the quality factor. Additionally, the sample block had clear issues with thermalizing at cryogenic temperatures

TE_{011} Geometric Factor (Ω)	Our design (4.5x15 mm Rod)	8x5 mm Rod from Ref. [8]
Walls	1160	420,000
Mask	265	64
Top End Cap	1260	42
Total Non-sample surface	166	25
Sample Surface	2940	58
Total	173	18

Table 8.5: A comparison of geometric factors expressed in Ohms between the LLNL PPMS-compatible clamshell Hakki-Coleman Resonator versus a previous, more traditionally designed Hakki-Coleman resonator used in Ref. [8]. This highlights a design flaw in our resonator; the sample surface is much less significant than the other surfaces in the resonator, therefore making accurate surface impedance measurements of the sample very difficult.

because it lacked a tight connection with the cavity, causing it to be much hotter than the PPMS temperature. Attempts were made to mitigate this by pressing small amounts of indium metal into the gaps between the block and rectangular insert, but were unreliable. Furthermore, because the surface current structure of the TE_{011} mode crosses the seam of the clam-shell, the surface resistance of the cavity surfaces are increased, decreasing their sub-surface Q and drowning out the sample surface Q signal in measurement. For all these reasons, a real surface impedance measurement of a superconductor was never made with this cavity during my time at LLNL.

I bring this project up because it is a great example of the lessons learned through practical failures in grad school. This was a great idea in principle, and had been done before, but it was only by doing it and dissecting the design details that we were able to see the deeper requirements that make this cavity design work. It is my hope that someone in the future can take note of this if pursuing the same measurement scheme.

Chapter 9

ADMX SIDECAR RUN 1D PRELIMINARY RESULTS

This chapter will go over the ADMX Sidecar preparations, operations, and results so far in run 1D. The highlight of this Sidecar run is the first implementation of a Nb_3Sn superconducting tuning rod. It is late April 2024 at the time of writing. ADMX run 1D began data-taking operations in earnest in mid-December 2023 after multiple delays. The main experiment is expected to run for about a year for its first half of data-taking with the next expected insert extraction taking place around December 2024. The ADMX Sidecar is as always "along for the ride" so to speak and is unable to be extracted for repairs or testing. Everything must be done "in-situ" and must also be done in accordance with the main experimental operations. Unlike the LLNL PPMS system discussed in section 8.1, the Sidecar is in a much less controllable environment in terms of adjusting temperatures and magnetic fields; this makes characterizing the cavity challenging and sometimes impossible. The goal of this chapter is to not only highlight the work that has been done so far in characterizing the superconducting rod, but the axion search analysis that could be done given less than ideal operating conditions.

9.1 Previous Sidecar runs overview

The ADMX Sidecar was introduced by former ADMX UW graduate student Christian Boutan, as a high-frequency (4-7 GHz) side axion experiment for not only testing emerging technologies that could help solve the problems with high frequency axion searches, but for taking real axion search data at these frequencies, albeit at a much lower sensitivity than the main experiment. This was only possible because of a bit of free space on the top of the main cavity and below the 1 Kelvin plate above. This region is in the fringe field of the

main magnet, giving the Sidecar only an average magnetic field of 3.58 T at full current, 220 A, significantly reducing its sensitivity to axions; this is the main reason why it does not achieve of KSVZ or DFSZ sensitivity. This fringe field region has been mapped, with the magnitude of the field plotted through each center cross section in Figure 9.1. Although the field is slightly curved in this region, the off-axis components are still relatively small, with an average $\frac{B_z}{|B|} \approx 0.981$ over the region. In addition to the cavity sitting in this reduced field region, the Sidecar must have all the other necessary systems for a haloscope: Tunable rod and antennae, a cold, low-noise RF receiver chain, and the warm electronics to support these systems.

Sidecar’s premier run was during run 1A, where it’s emerging technology was the use of the piezoelectric motors for tuning the rod and antennae. This was one of the first successful demonstrations of cavity resonant tuning with such a motor. The stick-slip motors were manufactured by Attocube [110] to operate at 10 mK, 31 T, and high-vacuum. The Attocube ANR240/RES rotary motor and ANPz101eXT12/RES linear motor were mounted to the cavity’s top end cap and were used for tuning rod control and antenna coupling adjustment respectively. These motors can produce a significant heat load, so extra care had to be taken not to heat the Sidecar cavity or overall insert too much. The linear piezoelectric motor also did not function for a majority of this run, rendering the antenna coupling stationary. This run only implemented a Low Noise Factory HFET amplifier, LNF-LNC1_12A, thermally sunk to the top of the helium reservoir, as its primary amplifier. They measured a noise temperature of $4.0 \pm K$ using the Y-factor method 1 highlighted in section 5.2. The cavity itself was a knife-edge construction similar to the main cavity, with two end caps bolted to a 0.121 m tall, 0.064 m inner diameter cylindrical barrel with very sharp edges that ‘bite’ into the end caps; the tuning rod was a 0.013 m diameter, hollow copper tube with fused end caps. This cavity had an unloaded Q of 22,000 empty, and 11,000 with the tuning rod. The system was able to exclude axions in 3 regions, 4202-4249 MHz, 5086-5799 MHz, and 7173-7203 MHz, to varying sensitivities beyond the CAST limit; the last region was done using the higher TM_{020} mode, which was a novel technique at the time. For more details of

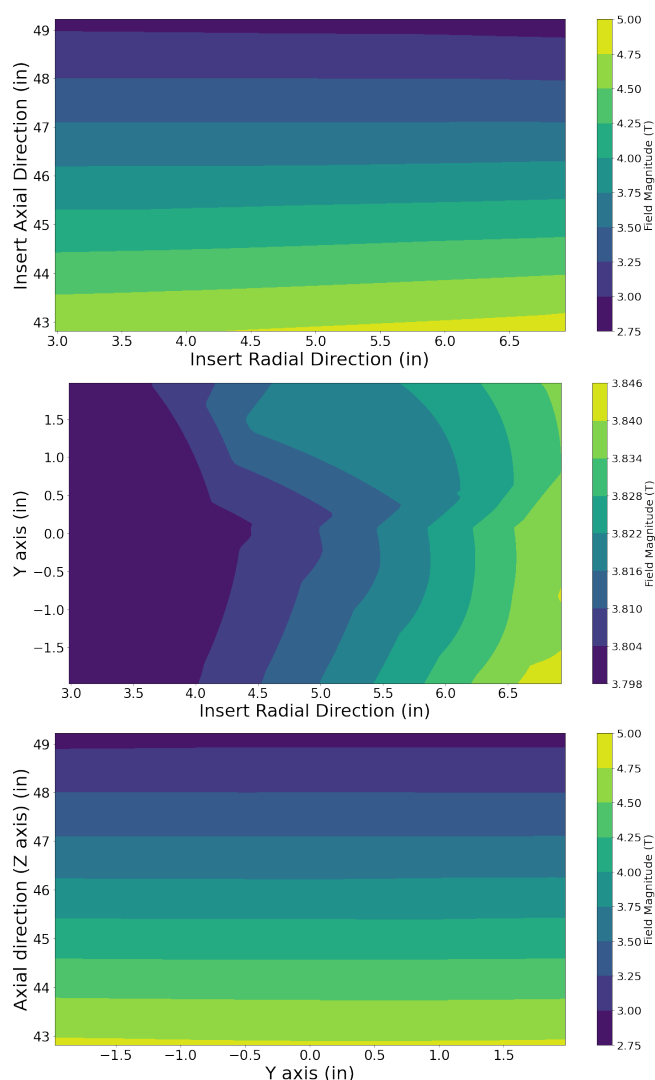


Figure 9.1: The ADMX Sidecar Magnetic Field Region. The magnitude of the magnetic field is plotted for each cross-section passing through the center of the Sidecar cavity. The radial coordinate uses the ADMX insert center as zero. The axial 'Z' coordinate uses the center of the magnet coils as its zero. The 'Y' coordinate is the only coordinate that uses the Sidecar center as its zero. Note that this plot doesn't display the direction of the field vector, which is predominantly in the axial direction; off-axis components are quite small with the average $B_z/|B|$ over the region being ≈ 0.981 .

this run, see ref. [27].

The first major set of upgrades to Sidecar came together in time for ADMX run 1C. The

premier technology for this run was the introduction of a quantum noise-limited amplifier, a Josephson Traveling Wave Parametric Amplifier (JTWPA), as the primary amplifier. JTWPA's are known for delivering high gain over a wide (3 GHz) bandwidth, which is ideal for more broadband axion searches. This amplifier, pictured in Figure 9.2 was fabricated at and acquired from MIT Lincoln Laboratory [86]. It is made of niobium with a critical temperature of roughly 9.3 K. Being a field-sensitive component, this amplifier was mounted within the SQUIDadel, so that it would be protected by the bucking coil magnet. This amplifier, because of its broadband nature, has an operational gain typically limited to around 10 dB. When the JTWPA was paired with the Sidecar HFET as a secondary amplifier, which had a noise temperature of $3.7 \pm 0.2 K$ using Y factor method 2 (outlined in Section 5.3) the system had an overall system noise of $925 \pm 80 mK$. That HFET amplifier towards the end of this run failed, and was sent for repair after the insert was extracted. The RF schematic of this amplifier chain is pictured in Figure 9.3.

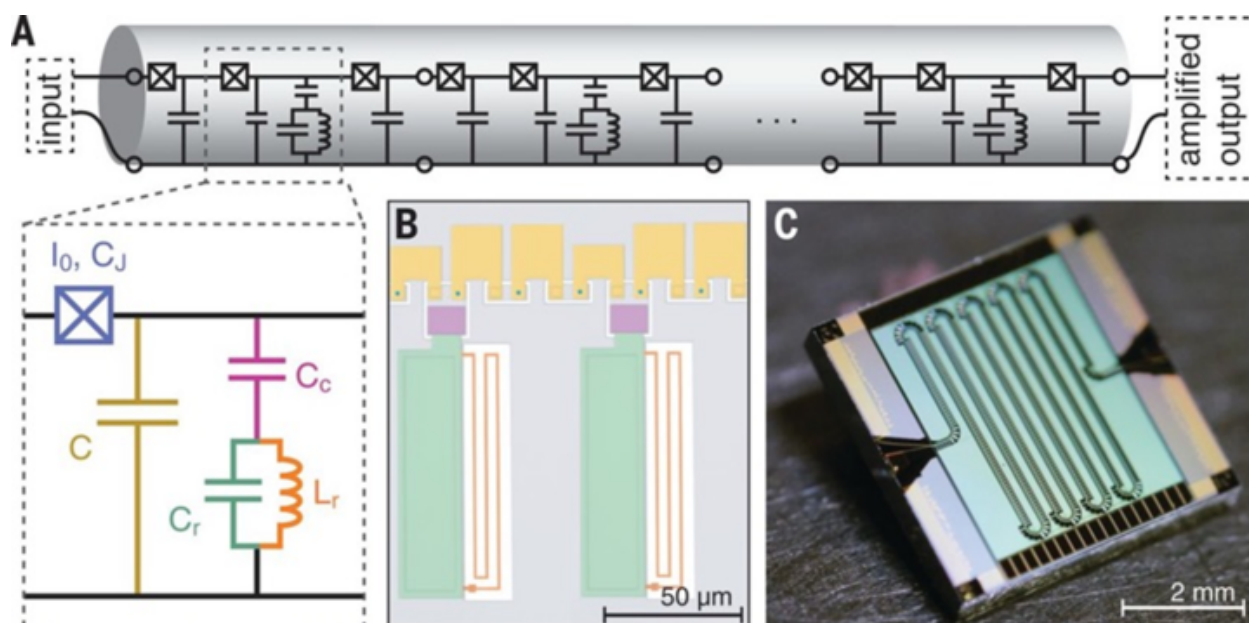


Figure 9.2: The Sidecar JTWPA. A) The circuit diagram of the amplifier. B) A diagram of the amplifier micro structure. C) A picture of the entire amplifier chip.

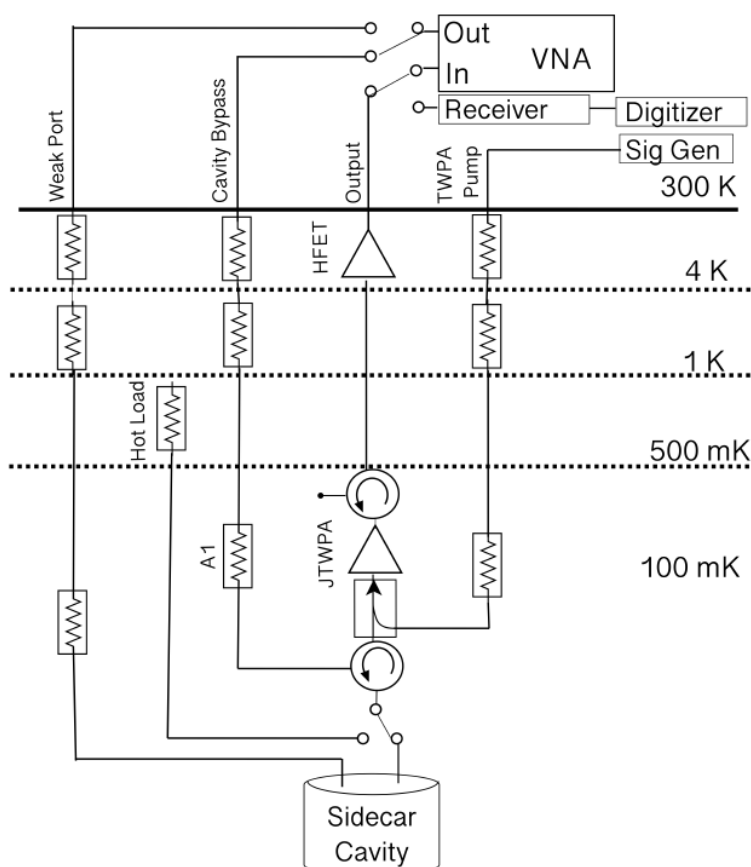


Figure 9.3: An RF schematic of the run 1C and 1D Sidecar receiver. The JTWPA is placed between two circulators to prevent power being reflected back into the cavity and to allow for periodic RF calibrations. Although a completely separate chain, these magnetic-sensitive cold electronics resided in the SQUIDadel along side the main cavities' sensitive electronics protected by the bucking coil magnet.

Concurrent with the upgrade to the RF chain, a whole new Sidecar cavity was fabricated at LLNL, and pictured in Figure 9.4 disassembled. This cavity was a clam-shell construction, with two cylindrical halves machined from a solid copper block and bolted together with machine screws around the flat face of either half. The clam-shell seam, in theory, being along the axial direction, should not disturb the currents of the TM_{010} mode or any higher axion-search modes, therefore maintaining Q . The 2" diameter copper tuning rod was hollowed out

and the end caps welded back on to minimize the mass, and therefore the torque required by the piezoelectric motor to rotate. The alumina axles were epoxied to copper tabs that bolted into an inset on the top and bottom of the rod. These axles would pass through the apertures that are seen slightly off center axis of either half of the cavity in Figure 9.4. These tuning port holes were made diagonal with varying steps to accommodate different size ball-bearings to optimize the smoothness of the rod turning; however, it was later suspected that this may have been a source of RF leakage, and would inform the cavity design in run 1D. During run 1C operations, the cavity unfortunately had a very low Q of about 700. This was thought to be because of RF leakage, most likely through the tuning rod port, but also possibly through the clam-shell seam if it had loosened up significantly under cryogenic conditions.




Figure 9.4: The run 1C sidecar clam-shell cavity disassembled with 2" copper tuning rod and alumina axles.

Finally, the piezoelectric motors were moved from being mounted directly to the top of

the Sidecar cavity to the 1 Kelvin plate above. The exposed top alumina axle was clamped by a flexible bellow stainless steel rod coupler, which then clamped to a stainless steel rod that was directly attached to the rotor piezoelectric motor above through a port hole in the 1 Kelvin plate. Because the alumina was an insulator, there shouldn't be a significant thermal connection between the 1 kelvin plate and the cavity. The linear piezoelectric motor was mounted to the underside of the 1 Kelvin plate, and a NbTi superconducting coaxial cable antenna was clamped to it; the antenna would be superconducting during operations, and would not have the thermal conductivity to create a link between the cavity and 1 kelvin plate. The whole idea behind this change was to minimize the heating of the cavity and mixing chamber plate during tuning operations; tuning rates in run 1A were limited by the heat generated by the motors. The piezoelectric motors had mixed success during this run with both motors failing at certain times. The rotor motor, although working initially, failed before the JTWPA could be run for axion search data, and as a result the axion limits reported in Ref. [20] only covered 2 MHz of frequency space, roughly the bandwidth of the digitization scans taken.

The successes and failures of these past two runs would inform the preparations and operations of run 1D Sidecar. As one will see, run 1D Sidecar seems to still suffer from some of the problems of its past: A low Q possibly from RF leakage, and failing piezoelectric motors.

9.2 Preparations and design changes for run 1D

This section will go over the design changes to the Sidecar system, warm testing, and installation into the ADMX insert.

9.2.1 Superconducting Tuning Rod

The highlighted novel technology of this run is a Nb₃Sn coated tuning rod manufactured at Fermilab and overseen by Fermilab senior scientist, Dr. Sam Posen. The coating process comes from the accelerator cavity world, using pure niobium metal as the substrate for the

coating, and is outlined in Ref. [95]. It is meant to be a thick film, on the order of about 3 microns, thick enough to prevent the field from penetrating into the bulk niobium. At their request, the rod diameter was reduced from 2" to 1.75", and the axle tab inset was redesigned to accommodate a center line hole; this would allow for hanging the rod by a wire strung through the center line hole when it was placed in the furnace for coating. The completed rod is pictured in Figure 9.5. An inadvertent benefit of this was that the detection volume of the cavity increased slightly, from 0.58 L to 0.61 L. This also shifted the tuning frequency range downward to 3.6-6.2 GHz. Two versions were made: One solid Nb Rod, and one that was hollowed out and the end-caps welded back in place. The hollow version was preferred by us because it would have a reduced torque for the piezoelectric motor. The solid version, however, was much easier for the Fermilab engineers to manufacture initially. In the end, the hollow rod was the rod installed in the cavity from which data was taken. This additional solid rod, however, is planned to be installed in a cavity at LLNL for further testing soon.

9.2.2 Copper Clamshell Cavity

The copper clamshell cavity was re-made for this run by collaboration colleagues at the University of Sheffield. The design remained almost entirely the same with the exception of the tuning rod axle port; the port was extended out from the top of the cavity exterior and the number of ball-bearing inset steps was reduced from 4 to 2, as pictured in Figure 9.6. The idea here was to reduce any RF leakage through the port coming from the free space around unused steps, by creating a much longer port. Additionally, the tuning rod axles were changed from alumina to sapphire, in order to increase the thermal conductivity and dielectric constant of the axles. This was all motivated by the low Q seen in run 1C Sidecar operations.

9.2.3 Simulation Work

The cavity assembly was simulated using the COMSOL emw software package. The eigenmode solver was used to extract simulated values for the mode frequency, axion form factors,



Figure 9.5: The Nb_3Sn coated tuning rod from a side and top viewpoint. The rod is 1.75" in diameter, and features an inset for mounting the copper sapphire axle tabs. The center and right holes are tapped with 8-32 threads for mounting these tabs. The center hole goes through the length of the rod so that it could be hung by a wire in the coating furnace.

as well as an array of geometric factors for performing multi-mode measurements later.

The real axion form factor was calculated using the magnetic field map shown in Figure 9.1, as well as an ideal form factor that assumed the magnetic field had perfect axial structure ($\vec{B} = B_0\hat{z}$). This was necessary for the TM_{010} mode, because these values would be required for feeding into the axion search analysis using that mode later. For the other modes, it proved useful for identifying TM modes of interest. The TM_{010} form factors, real and ideal, over the entire tuning range are plotted in Figure 9.7.

A total of 12 sub-surface geometric factors were tracked in simulation: The exterior copper cavity surface, the tuning rod axle assembly, the entire tuning rod, the tube wall of the rod, and the end caps including the curved fillet surfaces that curved towards the wall. The rod wall region was further subdivided into 5 sub-surfaces along its axis based on height




Figure 9.6: A close up of the run 1D Sidecar cavity tuning rod port profile as seen from one clam-shell half. The extension increases the length of port, minimizing RF power leakage out.

within the magnetic field: rod bottom, rod middle/bottom, rod middle, rod top/middle, and rod top. This was in anticipation of the magnetic field changing from the bottom to the top of the cavity interior, and the hope that one might be able to see a change in surface resistance along the length of the rod. The tuning rod angle was the primary geometric parameter varied, where 0 degrees represented a centered rod and 180 degrees represented a rod at the wall of the cavity, which as it was varied would give us a simulated mode map for the modes of interest. The first 5 TM modes were identified at first using axion form factor, but were eventually followed up by plotting the simulated electric fields and identified by eye especially near mode crossings. This initial simulation pass used the base design dimensions and took half degree steps in the tuning rod angle. Because of the large amount of data in this set, it has been chosen just to show the TM_{010} mode geometric factor for the tuning rod and exterior cavity over the whole tuning range, pictured in Figure 9.8. The gaps in this plot are due to the geometric factor dropping to near-zero values near mode crossings, and are cut from the data set. One would not want to operate the cavity in these frequency regions anyway. This initial pass didn't take into account error in the geometric factor because of




Figure 9.7: The simulated real and ideal axion form factor for the TM_{010} mode as a function of the mode resonant frequency from COMSOL simulations. The ideal form factor assumes a perfectly straight magnetic field, whereas the real form factor data interpolates the magnetic field map data shown in Figure 9.1. Because the real field isn't straight and constant, it is less than the ideal form factor. Both form factors drop to zero in mode crossing regions, of which there are 5 crossings visible in this plot. The real form factor data is used later in the axion search analysis for calculating expected axion signal power for each digitization scan.

the computational time.

Once the target frequency range for the run was chosen to be 5.2-5.6 GHz follow up simulations were made just in that run range to lower computational time. This allowed us to get error estimates of the geometric factors by adding a nested parametric loop within the rod angle loop that varied the major dimensions of the cavity according to its machine tolerance for each value of tuning rod angle. Five major dimensions were chosen for this analysis: The diameter of the cavity, the height of the cavity, the diameter of the tuning rod, the height of the tuning rod, and the radius of the fillet on the rod and cavity end cap regions. Each variable received 3 possible values based on the base design value, x_0 , and the machine tolerance, Δx : $x_0 + \Delta x/2$, x_0 , $x_0 - \Delta x/2$. This created a set of $3^5 = 243$ possible geometries of the cavity at each possible tuning angle of which standard deviations in the geometric factor could be taken, as outlined in section 6.5.2. The tuning rod steps




Figure 9.8: The simulated TM_{010} mode geometric factors for the rod and exterior cavity as a function of mode resonant frequency over the entire tuning range of the cavity. The geometric factor tends to drop to zero near mode crossings, which makes it indistinguishable with TE modes, and therefore these regions are cut from the data set, producing the gaps in the plot data. When making a run plan, it was best to choose a region between mode crossings with relatively constant, well-separated geometric factors, such as the region between 5.2-5.6 GHz.

had to be increased for this simulation to 2.5 degree steps because computational time was significantly increased otherwise. These were then plotted over the data from the initial pass with error bars, as pictured in Figure 9.9, specifically showing the exterior cavity geometric factor for each mode. The uncertainty in geometric factor, unsurprisingly, got quite large near mode crossing regions, and even if one was not in a mode crossing, but close to it, the error in the geometric factor would be higher. This work is still on-going as how to best optimize where in the tuning rod angle space is it best to perform a decomposition. The cavity characterization section of this chapter will go into the challenges faced performing the decomposition on the Sidecar system. Table 9.1 will show the values of all the geometric factors and their uncertainties for the starting TM_{010} frequency of Run 1D, 5521 MHz, which corresponds to a tuning rod angle of roughly 26 degrees.

	TM_{010}	TM_{011}	TM_{012}	TM_{020}
Frequency (MHz)	5521.0 ± 33.6	5655.8 ± 20.0	5842.5 ± 26.3	6060.7 ± 26.0
Exterior cavity	150.05 ± 10.05	116.89 ± 26.67	311.44 ± 1.70	158.60 ± 15.41
Tuning rod (whole)	694.62 ± 66.67	568.02 ± 132.28	1299.9 ± 7.94	786.53 ± 78.16
Rod axle assembly	54160 ± 9966	38018 ± 21386	3971305 ± 670343	67609 ± 20900
Rod walls	1386 ± 49	2032 ± 387	1322 ± 8	1691 ± 72
Rod end caps	1405 ± 222	793 ± 199	78789 ± 10196	3047 ± 576
Top end Cap	2814 ± 471	1585 ± 398	157656 ± 20445	3047 ± 576
Bottom end Cap	2808 ± 427	1585 ± 398	157508 ± 20366	3047 ± 576
Rod wall top	9628 ± 338	10640 ± 66697	26992 ± 232	11555 ± 252
Rod wall top/middle	5497 ± 265	9419 ± 904	5102 ± 44	10329 ± 4604
Rod wall middle	6649 ± 315	38101 ± 17901	3435 ± 32	6912 ± 2525
Rod wall middle/bottom	5514 ± 265	9412 ± 892	5122 ± 43	10365 ± 4623
Rod wall bottom	9631 ± 355	10634 ± 6701	27049 ± 267	11534 ± 261
Total	123.1 ± 17.4	96.7 ± 36.9	251.2 ± 2.4	131.73 ± 21.4

Table 9.1: The geometric factors in Ohms of the Sidecar cavity for a tuning rod angle of 26.74 degrees. This angle corresponds to the TM_{010} mode frequency that the Sidecar cavity was left at for the start of run 1D, and therefore they are the relevant values used for characterizing the cavity in cool down and magnet ramp. Unfortunately, this was near a mode-crossing for the TM_{020} mode so some errors are quite large for that mode. There is anomalous behavior in the values for the TM_{012} mode as well, which has much larger values and smaller errors than the other modes. TM_{012} mode also had a measured mode frequency that did not correspond to this rod angle.





Figure 9.9: The simulated Sidecar exterior cavity geometric factors with error bars for each mode of interest in the target run range. The error was calculated by simulating $3^5 = 243$ possible geometries based on the machine tolerances of the design at each tun rod angle; because of the large computational time necessary, this was only done by moving the tuning rod in steps of 2.5 degrees from 10 to 50 degrees. This median value set of G factors is then plotted against the initial base geometry that used half degree steps. It agrees well in between mode crossings, but the uncertainty increases greatly near mode crossings.


9.2.4 Warm cavity assembly and testing at LLNL

The cavity and tuning rod were received at LLNL in the summer 2023. Minor alterations to the cavity and tabs were made to ensure proper fitting and rotation of the tuning rod at room temperature. Very special care had to be taken to prevent rubbing of the tuning rod against the interior end-cap surfaces; the design accommodates for 0.005" gaps on either side between the rod and interior end cap surface, as pictured in Figure 9.10. The sapphire axles were epoxied to their copper tabs with stycast; several sets had to be made to ensure the tab was flush with the rod surface in its inset, and the axle was perpendicular with the tab (pictured in Figure 9.11). Any deviations from this produced friction in the rod

rotation. There were several issues with the silver-plated screws cold-welding themselves to the Nb substrate, so moly-powder dry lubricant was used, as well as much shorter screws than initially designed for. Additionally, the axle ball bearings had to sit in their insets on the cavity exterior with very tight tolerances; when the clam-shell halves were fully tightened together, any misalignment could cause the bearing to clamp on the axle and cause friction. After the rod with axles had been placed and aligned within one cavity half, and ball-bearings placed in their insets, a stainless steel coupler and copper clamp were placed on the cavity top and bottom respectively (pictured in Figure 9.12) and tightened flush to the ball-bearing bushing. This locked in the position of the rod and its gap. The other cavity half would then be installed on top and tightened in a star-pattern formation.



Top



Bottom

Figure 9.10: A close up of the run 1D Sidecar cavity tuning port. The cavity is opened with tuning rod and axle resting inside and ball bearing inserted in place above. Special care had to be taken to ensure there was no rubbing of the rod with the cavity interior, and no rubbing of the axle assembly within the port.

A test installation of the cavity was done on a mock-up test stand of the ADMX insert T-plate and 1 kelvin plate, including the piezoelectric motor assembly. At the time of testing,



Figure 9.11: The Sidecar tuning rod with copper axle tab installed. The sapphire axle was epoxied to the copper tab with stycast. Vented, silver-plated 8-32 screws attached the tab to the Nb-substrate tuning rod.

only the solid Nb rod had arrived at Livermore, which was used for these assemblies and tests; the hollow rod would only arrive in time to be immediately sent to the University of Washington for installation in the insert. Nonetheless, the warm temperature tests with the piezoelectric motors informed us on how inconsistent torque required to turn the rod was in certain angles of its rotation path. Initial tests required using 60 V from the power supply to get consistent stepping. After re-assembling the cavity many times, this number was brought down to about 30 V to reliably tune the cavity through its range; however, certain smoother turning regions could be operated with a voltage as low as 18 V. For context, the copper rod used in run 1C that had gotten stuck at cryogenic temperatures, had operated in warm testing at 28 V; the copper rod being 36% heavier than the solid niobium rod indicated that there were frictional issues in the new assembly and that it wasn't just the weight of the solid rod that was causing them.

Nonetheless, warm RF testing proceeded with taking a mode map of the cavity pictured in Figure 9.13. Because of the high resistance of room temperature Nb_3Sn , Q_0 was only




Figure 9.12: The Sidecar cavity with rod, axles, ball-bearings and stainless steel axle coupler all installed. The copper clamp on the cavity bottom was used for rod thermalization, whereas the axle coupler on top connected to the rotor motor. Both of these locked the gaps on either side of the rod in place, ensuring no rubbing, before the other cavity half was added.

around 500-2500 throughout the tuning range for the TM_{010} mode. The mode frequencies lined up with simulation values, as well as the expected mode crossings. It was at this point that a target frequency range of 5.2 to 5.6 GHz for the run was settled on, because it offered the longest track of frequency space without any mode crossings, and had only been probed by Sidecar run 1A in 2016 to a low sensitivity value. With hints at the time that the CAPP institute was going to release axion limits up to 5.3 GHz, the cavity would be set to start at the upper end of this range, at about 5.5 GHz when installed at UW.

9.2.5 HFET repairs

The Sidecar HFET amplifier was repaired by Low Noise Factory in the interim, after failing towards the end of run 1C. It was sent back to UW in August 2023, and passed all room temperature and liquid nitrogen dunk testing, following in line with the factory specs given




Figure 9.13: A room temperature mode map of the run 1D Sidecar taken at LLNL. Because of the high resistance of room temperature Nb_3Sn , Q_0 was only around 500-2500 throughout the tuning range.

by LNF. It was installed in the insert without helium testing or a full noise temperature measurement.

9.2.6 Installation of Cavity on ADMX insert

The cavity was finally assembled and installed at UW in September 2023 with the freshly received hollow tuning rod, as pictured in Figure 9.14. The rod assembly turned very smoothly in comparison with previous assemblies done at LLNL. A copper strap was fabricated on site to give a thermal link to the tuning rod and cool it through its bottom axle, pictured in Figure 9.15. The strap has an estimated cooling power of $77\mu\text{W}$, and the cooling would be limited by the single sapphire axle, which could only transfer heat at $0.52\mu\text{W}$. Following the specific heat value models given in Ref. [65], a cooling time constant of 26.3 hours was roughly estimated for the rod.



Figure 9.14: The run 1D Sidecar cavity installed in the ADMX insert with piezoelectric motors mounted.




Figure 9.15: The copper tuning rod cooling strap attached to the bottom axle. This strap is then directly bolted to the mixing chamber plate. The sapphire rod keeps the rod electrically isolated, while still providing thermal conductivity for cooling.

Because of the high resistance of the insert wiring at room temperature used for the piezoelectric motors, motor testing was limited and unreliable. Voltages from the power

supply had to be set very high as a result, always in the 55-70 V range to actuate the motors. The linear motor successfully actuated the antenna at room temperature, but the rotor showed clear signs of missed actuation or "skipping". An oscilloscope was used to check the wave form of the saw-tooth wave originating from the power supply, and it was discovered that the wave-form was being deformed in the warm space wiring. The wiring started at the power supply as a BNC cable that lead to a breakout board where the leads were split into 5 parallel paths to reduce the overall resistance; these 10 leads then went into a dB-50 cable, which stepped down to a micro-dB cable that attached to a vacuum feed-through on the insert. The 5 parallel lines would then be recombined on the 1-K plate, right before the connection to the motor itself. Oscilloscope measurements showed that this degradation was occurring at the point where the signal was split into 5 parallel paths. We were able to bypass this problem temporarily by wiring with other cables at the time, and test the tuning rod motor successfully at 70 V. The decision was made that the wiring within the insert was operating correctly, and that the warm wiring would be replaced after the insert went into the magnet, when on-site personnel were not preoccupied with other insert-specific work. It was still not clear if the tuning rod would move once things got cold. The strong port antenna was left in a weakly coupled state with a coupling constant of 0.1, with the expectation that upon cool down this would drop to roughly 0.01. This was to keep the cavity in a state for characterizing it during the cool down and magnet ramp. After the Q characterization data was taken, it would be adjusted to critical coupling for an axion search starting at 5.5 GHz.

9.3 Hybrid cavity characterization

This section will cover the quality factor measurements taken on the hybrid Sidecar cavity during the cool down and magnet ramp. As stated earlier, this was a challenging task because the insert is not designed for controlled cooling and warming. Additionally, there was no way of measuring the superconducting transition of the tuning rod, because its critical temperature should be higher than the NbTi SQUID loops within the JTWPA; the JTWPA

would be opaque until the system had cooled below at least 9 Kelvin. Simulations of the geometric factors were done beforehand with the hope of using the multi-mode decomposition technique outlined in chapters 6 and demonstrated in 8 with the NbTi PPMS cavity. The plan was to measure the first 5 TM modes: TM_{010} , TM_{011} , TM_{012} , TM_{013} , and TM_{020} . It was then noticed that the TM_{013} was too close to a mode crossing with the TM_{020} in the current tuning configuration, and could not be observed well. What further complicated this was the rotor piezoelectric motor failed to move the tuning rod, even after the warm wiring had an attempted repair.

Each multi-mode measurement was taken using the same process on the Sidecar DAQ. Because there is only one vector network analyzer for Sidecar, modes couldn't be tracked simultaneously. Furthermore, the DAQ software couldn't be easily adapted to automate this measurement process, and the center frequency of the VNA had to be adjusted via a remote operator on our DAQ website. The TM_{010} was the default mode that the system would be left in after a measurement, so it was only necessary to periodically cycle through the higher modes. The process went as follows:

1. Turn off axion search data cycles such as digitizations and TWPA re-biasing, but keep the network analyzer scans running each data cycle: transmission, reflection, and transmission wide scan.
2. Set the network analyzer to the mode target frequency and ensure good windowing parameters for Q fitting. Note the starting time stamp for the first scan.
3. Wait at least 15 minutes while the network analyzer does many Q fittings. These values of frequency, coupling, and loaded Q are stored in the database to be pulled later. By waiting 15 minutes and getting many fits, a standard deviation and uncertainty in these values can be set.
4. Set the network analyzer to the next mode frequency and record the time.

5. Repeat until all modes of interest have associated start and stop time stamps. Set the frequency back to the TM_{010} mode, and record its start time stamp to be at least 15 minutes before the first time stamp recorded.

A script then took the recorded time stamps from each measurement, and pulled the relevant sensor values during those time periods from the database, including values for relevant temperature sensors, and the magnet current. A weakness in this method was that, especially during the cool down, temperatures were not guaranteed to be the same for each mode in a given measurement; if the temperature changed significantly over the course of the hour it took to measure, then the last mode measured would not be equivalent to the first. This meant that performing a surface resistance decomposition between those two modes wouldn't produce necessarily an accurate result. This was less relevant when ramping the magnet because it is done in steps, and so the magnet current was pretty much constant over the measurement; however, the magnet ramp does heat the insert by eddy currents, meaning differences in temperatures arose based on how long it had been since the field last changed. Additionally, these had to be made with a static antenna, so the coupling could not be changed mode to mode; in principle this shouldn't be an issue as long as one measures beta well. This all accumulated into a data set that has too many variables to really know what is going on with the SRF tuning rod. Nonetheless, I will now present the results of it so far.

9.3.1 Cool Down Studies

The first cool down measurement was taken until the first signs of cavity structure appeared in the VNA, which occurred when the Sidecar cavity temperature sensor was at about 5.02 K. Three more multi-mode measurements were made over the course of the cool down roughly at 3.5 K, 2.5 K, and 0.67 K. One strange aspect was that although an appreciable frequency shift had occurred as expected upon cool down for all the modes, the coupling constants for each mode were close to their room temperature values: The TM_{010} was about 0.15, TM_{011} about 0.9, TM_{012} was 0.12, and TM_{020} was 0.7. This may suggest that the antenna

didn't fully thermalize yet when the measurements were taken. Another important check on the system was to see if the measured mode frequencies all aligned on the simulated mode map in correspondence with a single tuning rod angle as show in figure 9.16; this would be necessary for the multi-mode decomposition because one needs to pick one set of G factors corresponding to one tuning rod angle. All but the TM_{012} mode seem to correspond well to a single angle, with an average of about 27.1 ± 0.4 degrees. The TM_{012} is anomalously lower, and when included in the average, shifts things to 28.7 ± 2.9 degrees. This was not noticed in room temperature testing. It is therefore hard to say that the system is well aligned with simulations, and a multi-mode decomposition result shouldn't be taken with confidence.




Figure 9.16: The measured mode frequencies during cool down plotted onto the simulated mode map of which simulated geometric factors are derived. One hopes that all mode frequencies would align closely with a single tuning rod angle, but that is not the case here; the TM_{012} is at an anomalously lower frequency than the other 3, giving it a corresponding tuning angle of about 33.7 degrees, whereas the other lay much closer to the TM_{010} average angle of 26.7 degrees. This θ_{010} angle is plotted as a vertical line, whereas the θ_{avg} for all four modes is plotted with the upper and lower 1-sigma bounds plotted as well.

Most importantly, the unloaded quality factor measurements were much lower than ex-

pected, even much lower than what is expected for an all copper cavity. Figure 9.17 shows the experimental unloaded quality factors versus temperature for each of the four TM modes. The horizontal lines in this plot show the expected cryogenic quality factor for various surface resistances of the tuning rod; the $R_{rod} = R_{Cu}$ corresponds to an all copper cavity, whereas the $R_{rod} = 0$ corresponds to the best possible hybrid Q where the superconducting rod doesn't contribute any surface resistance. In this way, one expects as the system cools, the experimental data should cross these lines, and ideally settle somewhere above the $R_{rod} = R_{Cu}$, but below the $R_{rod} = 0$ line. As one can see, this did not happen, and the quality factor, although increasing as the system cooled, was lower than the expected value for a rod exhibiting 10x the surface resistance of cryogenic copper.

The expected quality factors in Figure 9.17, however, do not account for any possible RF leakage in the cavity. As mentioned when describing Sidecar run 1C, it was presumed that there must have been RF leakage causing the very low Q of 700 during that run. One can calculate based off the expected quality factor for that cavity that it would've had a Q_{leak} of about 720 to have such a low result. There were no tests between run 1C and 1D that confirmed that this leakage issue had been resolved by the efforts made adapting the cavity tuning rod ports to their new configuration in run 1D. Furthermore, it was noticed in room temperature testing that the torque applied to the screws holding the clam-shells together did have an effect on Q if they were too loose, although this was never quantified. As the copper contracts in cool down, it is possible that gaps may have opened up, increasing the Q_{leak} and spoiling the unloaded Q, however, this is very hard to quantify or prove. Figure 9.18 explores this idea by introducing an associated Q_{leak} to the expected Q value lines for the TM_{010} mode: A) shows the same data from Figure 9.17 for TM_{010} , B) shows the values assuming a $Q_{leak} = 720$ in accordance with Sidecar in run 1C, and C) shows a $Q_{leak} = 10000$. One can see that Q_{leak} can't be as bad as Sidecar run 1C, because the implied intrinsic quality factors would be far higher than even the ideal hybrid cavity, which isn't physically possible. However, if one sets a higher value, $Q_{leak} = 10000$, the experimental data does start to follow the expected values; in this way, $Q_{leak} = 10000$ represents the rough lower

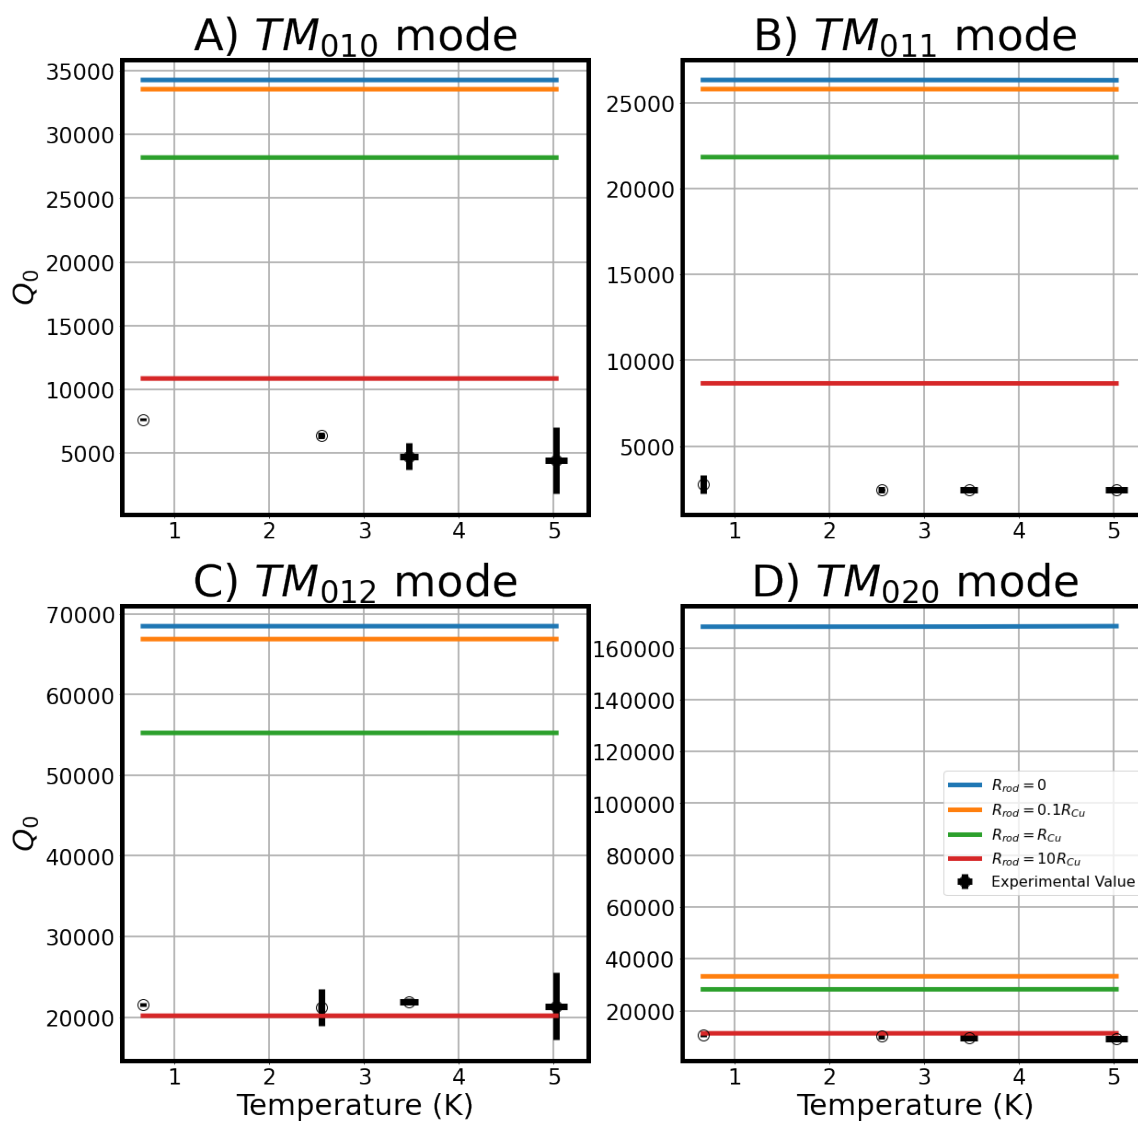


Figure 9.17: The unloaded quality factors versus the Sidecar temperature during the insert cool down for each of the relevant modes: A) TM_{010} , B) TM_{011} , C) TM_{012} , and D) TM_{020} . In addition, horizontal lines of the expected cryogenic quality factor for each mode with varying tuning rod surface resistance is shown. The $R_{rod} = R_{Cu}$ line corresponds to the expected cryogenic Q for an all copper cavity, whereas the $R_{rod} = 0$ would be the upper bound where the rod contributes zero surface resistance to the Q .

bound on the Q_{leak} , because any lower value would imply intrinsic Q values that are in excess of what is physically possible. However, this is a totally prescribed value; one still doesn't

have any upper bound on the Q_{leak} , which is the critical value to determining the actual surface resistance of the tuning rod. Without an upper bound, one could just pick a Q_{leak} that makes the behavior of the rod look agreeable, as is shown in 9.18C.

Another simple way of getting at the behavior of the rod in the cool down, without a full multi-mode decomposition, is to follow the analysis done in section 6.4, and solve for the ratio of the rod resistance to R_{Cu} using Equation 6.55. From the ratio of the rod and cavity geometric factors, one can see that the ideal hybrid Q will only be ≈ 1.20 that of the all copper cavity. Substituting the experimental values of quality factor to solve for κ as is done in Figure 9.19, one can see the rod resistance is still much higher than copper. Its important to note here that this will produce slightly different ratios than 9.17 because it doesn't account for the geometric factor of the axle assemblies and only looks at G_{cavity} and G_{rod} . This also assumes that the copper resistance is exactly as expected from theory and that there is no leakage. Nonetheless it tells the same story; the resistances are much higher than expected.

Finally, we can translate this quality factor information into total mode resistances using the total geometric factor for each mode, as shown in Figure 9.20. This would be done in preparation of performing a multi-mode decomposition, but several factors prevented that analysis from going further. As outlined in section 6.5, once a set of mode resistances had been constructed, the C_{ms} matrices would be constructed and their determinants calculated. For simplicity, we stuck with just 2x2 matrix decompositions. With 4 modes, there are 6 possible 2-mode combinations. In this case, there were 10 possible 2-surface combinations from the 11 geometric factors that were simulated. The base decomposition would be just the tuning rod versus the exterior cavity. The subsequent ones would combine the other geometric factors with the exterior cavity versus a single surface of interest: the rod walls, the rod end caps, top end cap, bottom end cap, and the 5 wall segments top to bottom. The geometric factors would be selected based on a chosen tuning rod angle that best corresponded to the system state. As outlined with Figure 9.16, we were not very confident in this value because of the anomaly in the TM_{012} mode; for simplicity, the angle correspond-

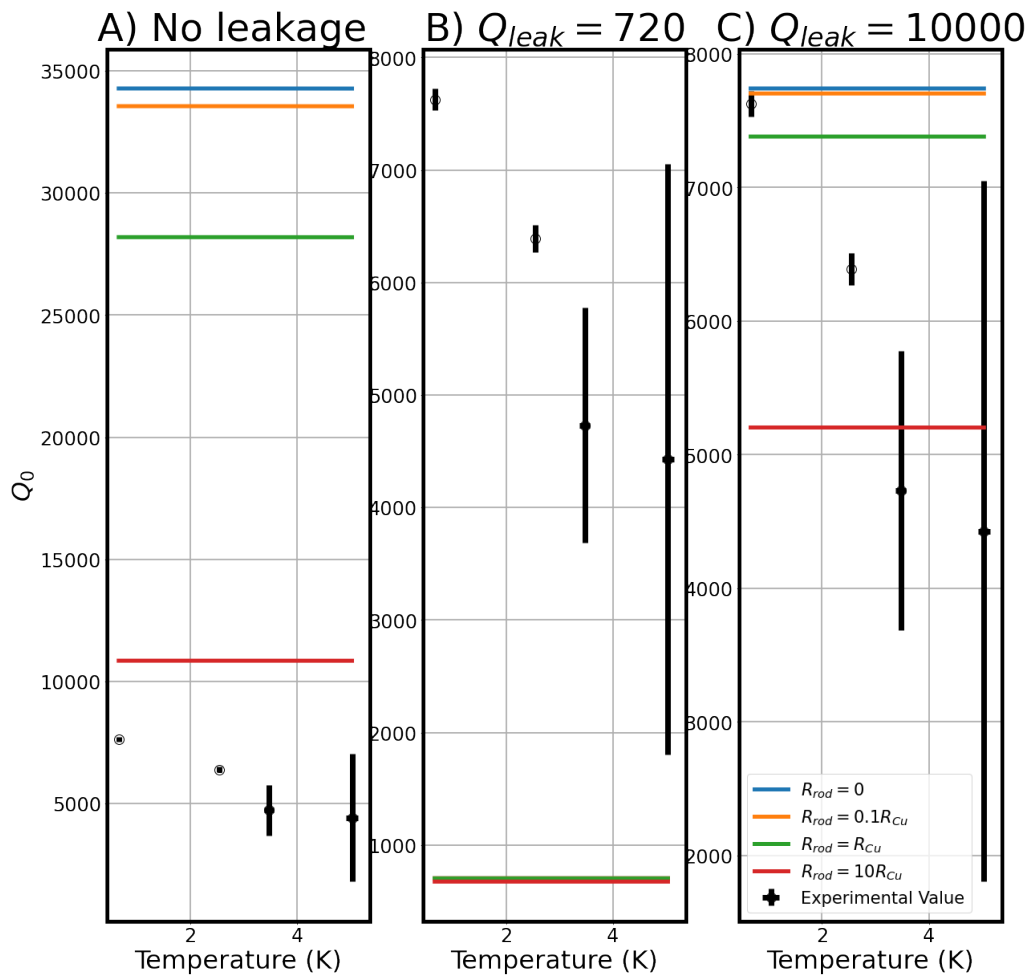


Figure 9.18: The unloaded quality factors of the TM_{010} mode versus Sidecar cavity temperature. The horizontal lines are the same as Figure 9.17, however they introduce an associated Q_{leak} value. A) shows the values with no associated leakage, $Q_{leak} \rightarrow \infty$. B) $Q_{leak} = 720$, which was the expected leakage Q for Sidecar during its previous run. One can see that a leakage Q this low is not possible because the experimental data greatly exceeds the expectation for an ideal hybrid cavity. C) $Q_{leak} = 10000$ represents a prescribed lower bound on the leakage Q where the experimental Q settles into a good hybrid cavity value between an all copper cavity and ideal superconducting hybrid. However, nothing prevents the actual Q_{leak} from being much higher than this, and the cavity surface resistances actually behaving less than ideally.

ing to the TM_{010} mode was chosen, 26.7 degrees. Once these 60 matrices were constructed, their determinants were calculated, and only 18 of them were statistically non-zero, meaning




Figure 9.19: The ratio of the tuning rod resistance to copper during the Sidecar cool down. This is calculated using Equation 6.55, with the experimentally measured quality factors, and the expected quality factors for an all copper cavity. It doesn't account for any RF leakage, or sub-quality copper resistance within the cavity. This method will also produce different values than the implied resistance values in Figure 9.17.

they qualified for being inverted and use in a decomposition. Of the remaining 18, only 1 of them was a mode combination that didn't include the anomalous TM_{012} mode that we

were not confident about. All of them produced non-physical, negative values for the surface resistances, which were not usable. The best solution to this would be to tune the cavity to another frequency, further away from mode crossings for all modes of interest, and converge on a better value for the tuning rod angle; one wants the cavity modes to be in region where they are distinct from one another to minimize the number of matrices with zero determinants. Unfortunately, at the time of writing, the Sidecar cavity has been unable to be tuned from its starting frequency using the piezoelectric motors, so this is not possible. I'll also note that it won't be possible to get new cool down data until the system is warmed up, which is most likely not to occur until December 2024.

9.3.2 Magnet ramp

Although the Quality factor measurements during the cool down were far less than anticipated, the magnet ramp might still shed some light on an important unanswered question: Is the rod actually superconducting? If it was still in a normal state, one would expect there to be no degradation in the quality factor as the magnetic field was increased. If it was superconducting, but there was some unaccounted for RF leakage, impurities, etc. that were degrading the overall base Q, then one still expects the quality factor to degrade in field.

The plots in this section will pretty much be the same as the last section, but replace the X-axis with magnetic field instead of temperature. Measurements were made at 0 T, 0.5 T, 2 T, 3.6 T, and 3.8 T, where the field strength is the average magnetic field within the Sidecar cavity for the given magnet current. As alluded to earlier, when the magnet current is being increased, the insert is heated significantly, and it takes significant time for things to cool off. The magnet was initially ramped at a starting temperature of about 700 mK to 200 amps of current with two short pauses at the 0.5 T and 2 T points; the system then cooled completely before the 3.6 T data point was taken. Similarly, when the magnet was finally ramped to 225 amps several weeks later, the system was cooled until the 3.8 T data was taken. Unfortunately, this created a dataset that had different temperatures and fields, which makes deciphering the behavior difficult. The first three data points were at about

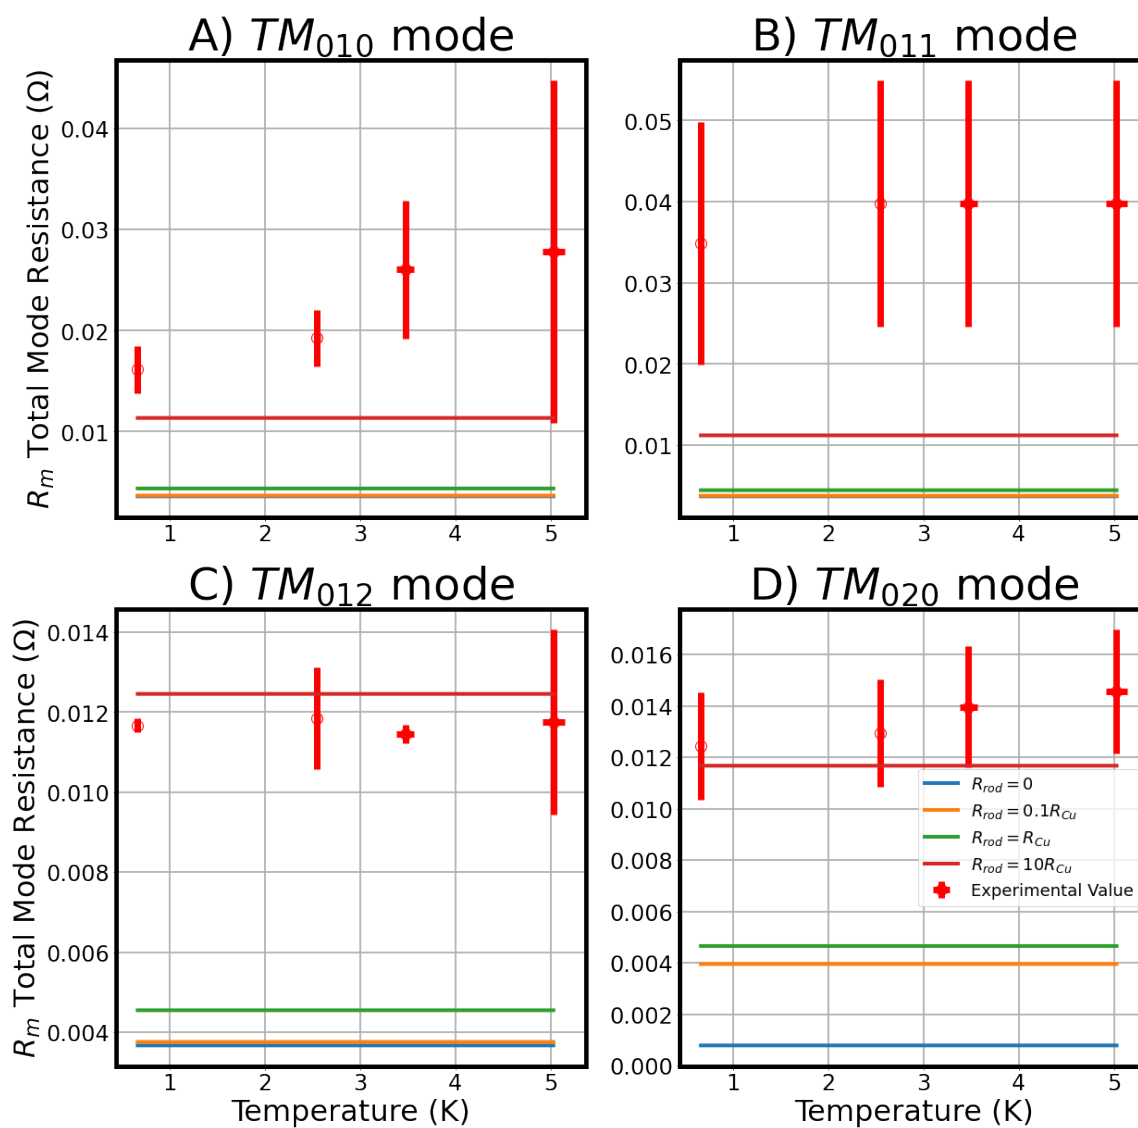


Figure 9.20: The total mode resistance versus Sidecar cavity temperature for each mode of interest. The horizontal lines represent the expected total mode resistance for possible resistance values of the tuning rod; the total mode resistance is a weighted average based on the geometric factors of the superconducting surfaces and copper surfaces.

700 mK whereas the last two were at 115 mK. Figure 9.21 shows this relation, and how it was well correlated with a shift in resonant frequency for all the modes.

This was slightly corrected by adjusting the resonant frequency of all the points to be

the same for all Q values within a given mode, but this didn't capture the decrease in Δf as the cavity cooled to 115 mK. This information is entangled with the possible increase in Δf due to Q -degradation from the magnetic field. Even after correcting for the resonant frequency shift, some of these points still showed an increase in Q , as shown in Figure 9.22, which implies the cooling dominated over any degradation due to the field. To add to this, the last two points were also taken after the Sidecar antenna had been critically coupled for axion data-taking; this should be fine as it is accounted for in β during the calculation of Q_0 , but the relation does break down for $\beta \gg 1$, and the TM_{020} did register the highest β of 10.8 and 3.7 for its 3.6 and 3.8 T values respectively.




Figure 9.21: The Sidecar cavity temperature versus the average Sidecar magnetic field during the ramp. The initial ramp was done before the system had fully cooled to base temperature, whereas the later two points the insert had fully cooled to its base value at around 115 mK. This temperature shift was well correlated with a shift in the Sidecar mode frequencies with the TM_{010} frequency displayed on a separate y-axis here. This is also seen as an increase in Q within most of the subsequent ramp plots.

When one overlays the expected quality factor lines to the ramp data, as shown in Figure

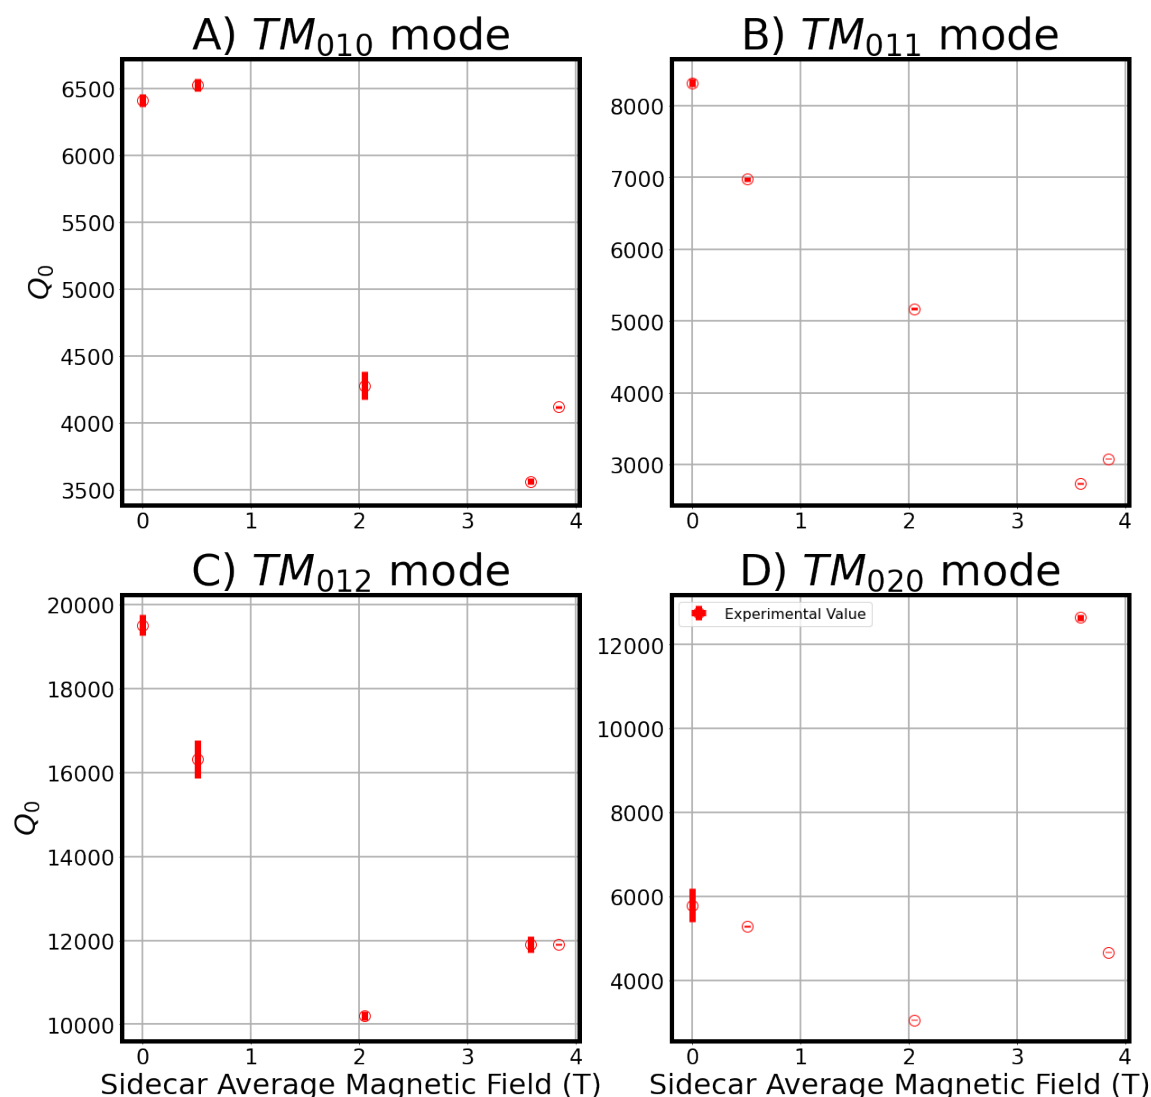


Figure 9.22: The unloaded quality factors for the modes during the magnet ramp. All of them seem to show signs of Q -degradation, indicative of superconducting behavior. The small rises in Q for the 3.6 T and 3.8 T values could be because of the much lower temperature of the insert when they were taken.

9.23 it is still clear that the cavity is behaving very poorly. Nonetheless, following the same Q_{leak} analysis done on the cool down data, as shown in Figure 9.24, one can see that the cavity behavior, if the cavity had a $Q_{leak} = 10000$, is in line for what one would expect of a hybrid SRF cavity of this construction; the Q degrades with field strength, just not from

a high zero field Q . It is important to note the lower zero field Q here as compared to the coldest temperature Q in the cool down which was about 1000 higher; its not entirely clear what changed here, as the temperature is only slightly higher, $\approx 40mK$, at 700 mK. Another likely reason is the coupling constant increased from about 0.17 in the cool down to 0.45 at the beginning of the ramp for unknown reasons, as the antenna had not been adjusted at all. It is important to note however, that one expects the leakage quality factor to be constant at this temperature because the copper exterior should already be fully thermally contracted.

The rod resistance to copper resistance ratio during the ramp is shown in Figure 9.25. The values are similarly high as they were in Figure 9.19, but there is still a clear increase in resistance from the zero to 3.8 T field points, as would be expected if the rod is superconducting. Figure 9.26 shows a similar story for the total mode resistances during the ramp; the total mode resistance increases with field. Similarly to the cool down data, the 60 possible surface decompositions all had statistically zero determinants, included the anomalous TM_{012} mode, or just produced non-physical surface resistance values. One would need to tune to a different tuning rod angle and wait until the magnet is ramped down again, which will not be possible until November 2024 at the earliest. In conclusion, the magnet ramp shows signs that the rod is superconducting, and could just be suffering from some sort of unknown radiative Q .

9.4 Noise temperature calibrations

There were two opportunities so far in run 1D to characterize the HFET noise temperature, and therefore the overall system noise temperature of the Sidecar system. Both of these implemented the y-factor method 2, outlined in section 5.3, that uses the mK attenuator as a variable temperature noise source. This was because the Radial coaxial switch was unable to be actuated successfully to point towards the dedicated hot load; plans are in order to bypass the pulsed switching system and attempt a direct DC connection to actuate, but this hasn't been approved or scheduled at the time of writing. Nonetheless, two opportunities arose where the mK attenuator was heated close to 1 Kelvin, enabling the Y-factor method

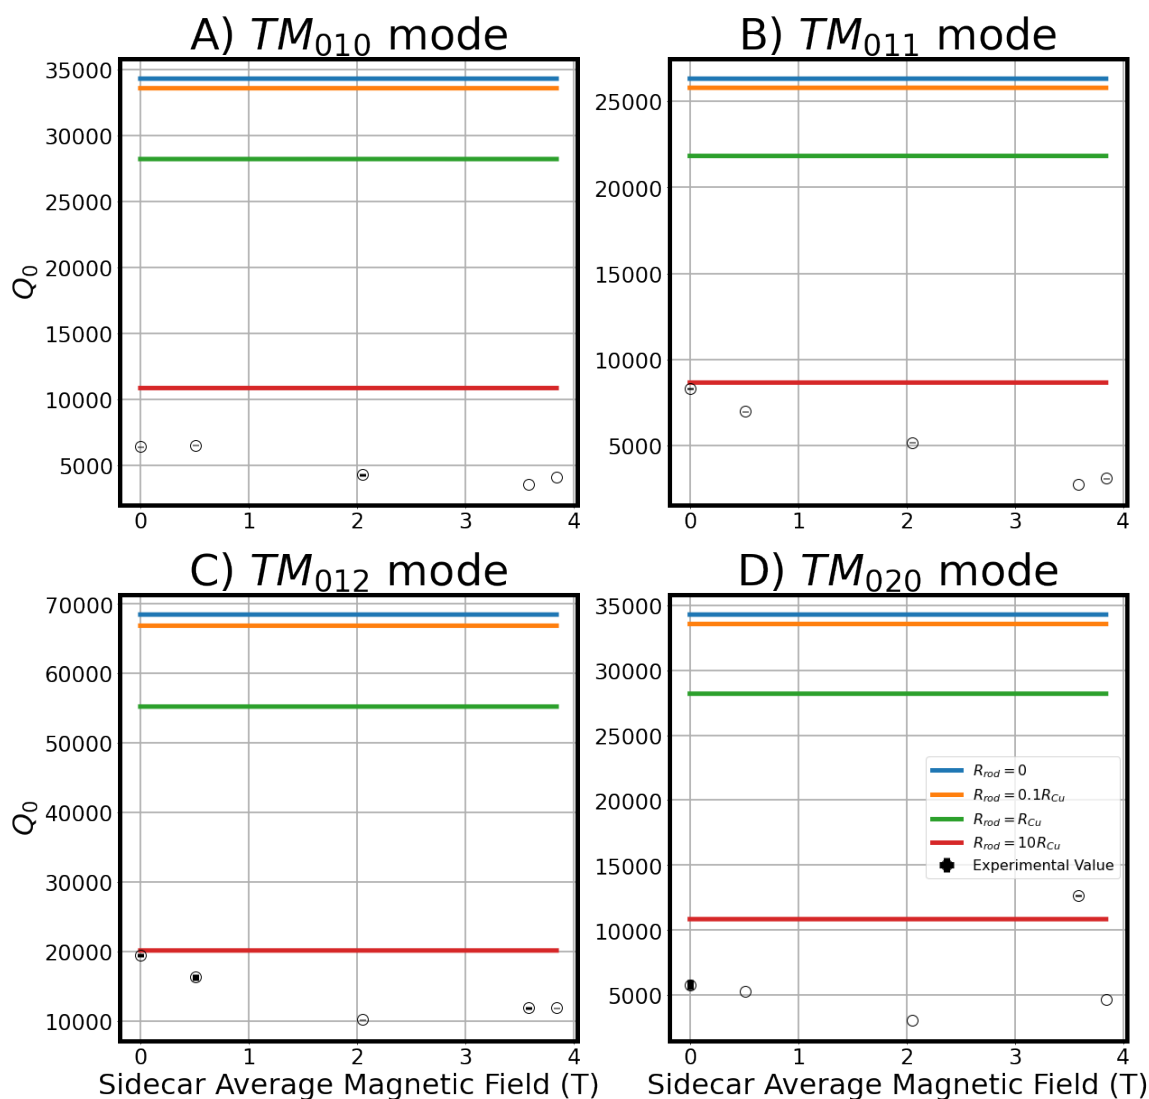


Figure 9.23: The unloaded quality factors for the modes during the magnet ramp compared to expected hybrid cavity Q values.

2 to be performed as it cooled back down: The main magnet ramp in December 2023 (see Figure 9.27), and an anomalous heating of the dilution refrigerator that took place at the end of March 2024 (see Figure 9.28).

During these periods, regular data taking on-resonance digitizations were turned off, and a wide-power measurement was performed over the entire tuning range of the cavity; a

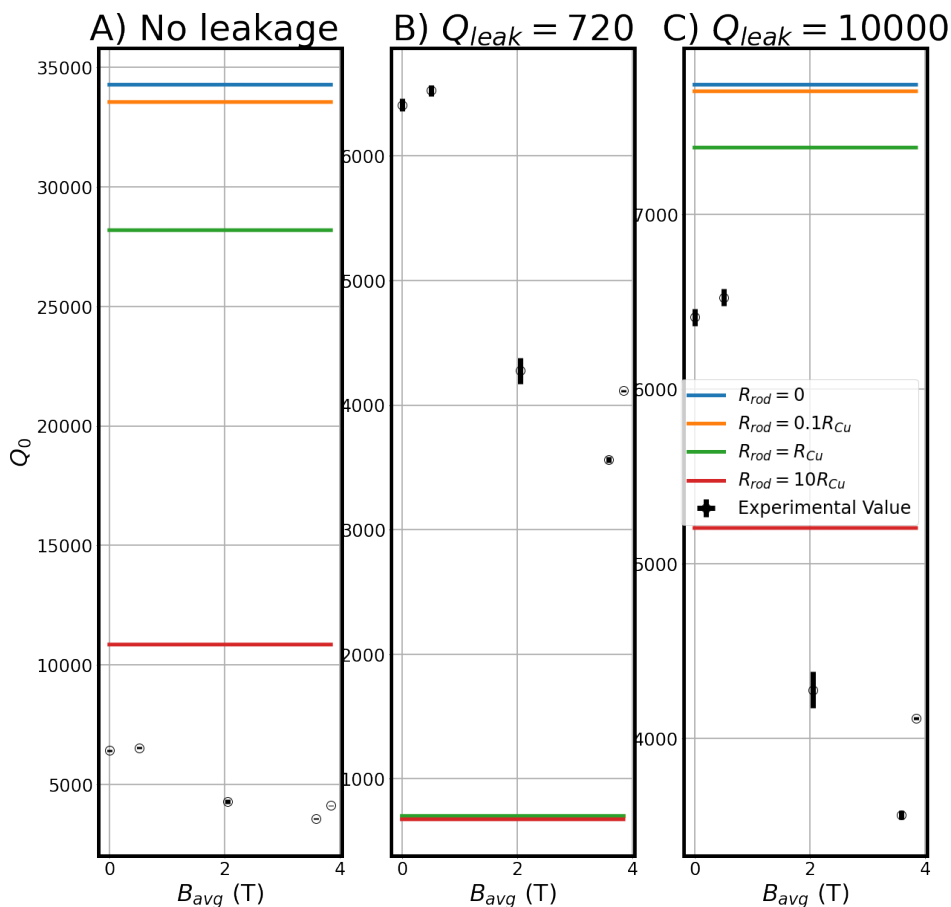


Figure 9.24: The TM_{010} mode unloaded quality factor versus magnetic field strength compared to the expected quality factor for varying rod resistances and different leakage Q s. A) shows the expected values for a cavity with no leakage ($Q_{leak} \rightarrow \infty$). B) shows a very low $Q_{leak} = 720$ that was estimated in the previous Sidecar run 1C. As one can see, this value is too low, as it implies the experimental data exceeds the ideal hybrid cavity case, which is nonphysical. C) shows a $Q_{leak} = 10000$, this was a prescribed lower bound estimate for Q_{leak} in the cool down data (see Figure 9.18) because it set the lowest temperature experimental value near the ideal hybrid cavity expected Q . In this case, the zero field data point does not fall near the $R_{rod} = 0$ value but below the $R_{rod} = R_{Cu}$ this might be because the cavity temperature was slightly warmer (700 mK versus 666 mK), changing the resistance of the rod, but one doesn't expect Q_{leak} to change with temperature at this regime because copper has already fully thermally contracted.

100 second power digitization (same 3 MHz bandwidth as normal data-taking) was taken every 200 MHz from 4-6 GHz. The average power was then calculated for each of these




Figure 9.25: The expected ratio of tuning rod resistance to copper resistance for each mode during the magnet ramp. This value is derived using Equation 6.55 and the other analysis outlined in section 6.4.

digitizations. Due to complications with the control scripts during the magnet ramp, these predefined frequency steps changed about half-way through the December cool down, and the latter half took steps every 60 MHz from 5223-5823 MHz. Because of how the analysis

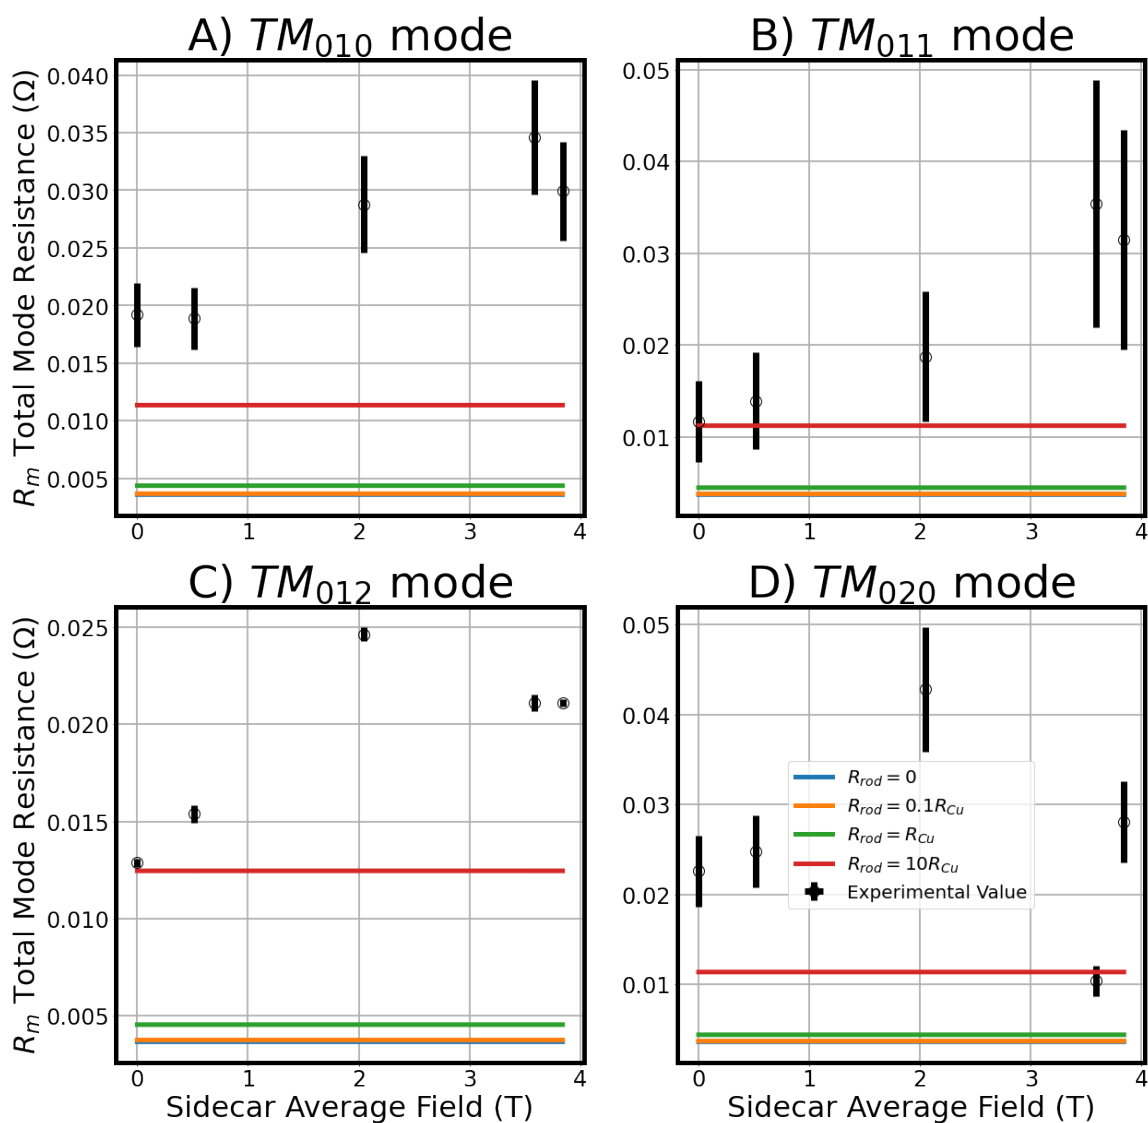


Figure 9.26: The Sidecar total mode resistances versus the magnetic field strength during the run 1D magnet ramp. The expected total mode resistance for a hybrid cavity with varying tuning rod resistance is also shown as colored horizontal lines.

was set up, it was easier to divide this magnet ramp cool down period into two separate measurements for each set of frequencies. The limitation to doing this, however, is that each individual measurement has roughly half the sweep in temperatures; the first half only took data from 0.75 K to 0.45 K, and the second half from 0.45 K to 0.1 K. Additionally, since

the 2nd half of data was taken at colder temperatures, but the uncertainty in temperature stayed constant, it has overall higher fractional uncertainties in the systematic error, which is evident at the lower fit temperatures in Figure 9.29. In contrast, the fits from the first half of the December cool down (see Figure 9.30) do not have this effect, and are more dominated by the error in the fit itself. The fits from the march cool down (see Figure 9.31), covering the entire temperature swing, does show an increase in systematic error at the low temperatures, but it also had a higher temperature difference registered between the cavity and attenuator, increasing the systematic error.




Figure 9.27: The relevant temperature sensor values versus time during the December magnet ramp cool down noise characterization period. This period was divided into to separate measurements because the frequency steps taken during the wide-power measurement changed part way through. The points that occurred after the gap in data at about 30,000 seconds constitute the temperatures used for the second measurement.

This is a good time to discuss how the error bars were set in Figs. 9.30, 9.29, and 9.31. A key assumption when using Equation 5.5 for these fits, is that the attenuator and cavity have the same temperature. In the previous run 1C Sidecar run, the temperature sensors used




Figure 9.28: The relevant temperature sensor values versus time during the March anomalous fridge crash cool down noise characterization period. This crash heated the system above 1 kelvin, whereas the previous magnet ramp cool down only reached about 0.75 kelvin. This gave this measurement a much larger swing in temperature, which is better for the fit integrity.

were not fully calibrated from the manufacturer, Lakeshore, and as a result they measured statistically the same temperature within their mediocre uncertainties. In this run, the sensors had been fully calibrated, and had an improved uncertainty ± 2 mK, which meant they now registered an observable temperature difference between them. This meant that the temperature uncertainty had to be set not only on the calibration uncertainty of the individual sensors, but the measured difference between the two sensors, which was an order of magnitude higher, at 40-60 mK. Figs. 9.32 and 9.33 show this measured difference versus time for each cool down period; one can see that the March period registered a noticeably higher difference. These temperature differences set the blue error bars seen in Figs. 9.30, 9.29, and 9.31 directly as the error in the x-axis, and indirectly for the y-axis error. The systematic uncertainty in the power was set by the equation:

$$\Delta P = |P| \sqrt{\left(\frac{\Delta T}{T}\right)^2 + \left(\frac{\Delta \epsilon}{\epsilon}\right)^2} \quad (9.1)$$

where P is the measured average power for the scan, T is the attenuator temperature during that scan, ΔT is predominantly the temperature difference between the attenuator and cavity

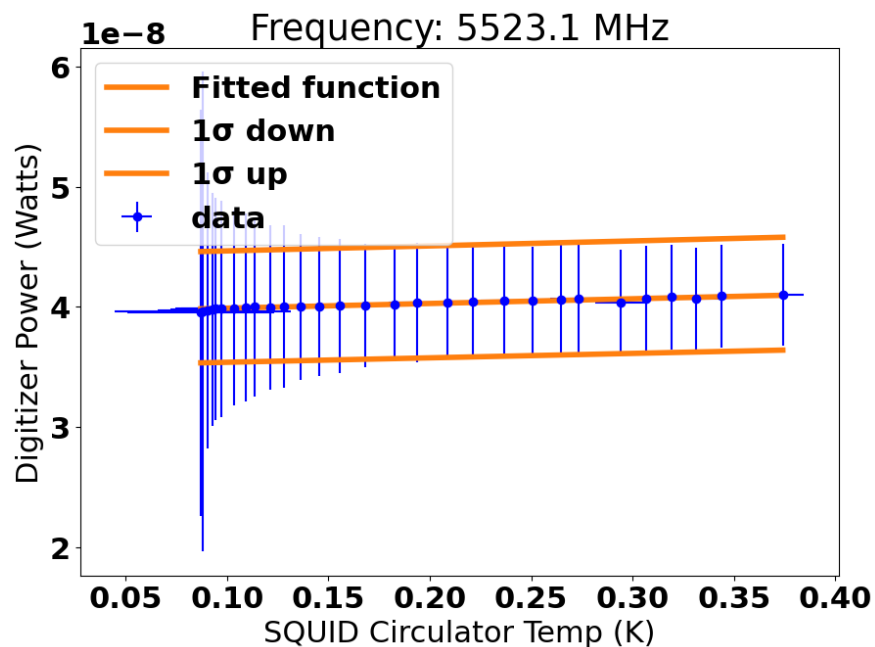


Figure 9.29: An example fit from the 2nd half of the December magnet ramp cool down. The blue error bars represent the systematic error due to the temperature difference between the attenuator and cavity, which is carried over as an uncertainty in the measured power. Because this temperature uncertainty is relatively constant, the fractional uncertainty increases at lower temperatures, causing the increase in the power uncertainty near 0.1 K. The orange 1σ lines represent the error in the fit itself. These two errors are then combined when setting the error in the HFET noise temperature.

during the scan (but it also includes the calibration error in the sensors themselves), and ϵ and $\Delta\epsilon$ is the emissivity, or transmission efficiency, between the cavity and JTWPA, which is measured ex-situ. Since this part of the receiver remained pretty much unchanged, it was fair to use the previous ex-situ measurement of ϵ during run 1C, which was 0.4 ± 0.04 . It is planned to remeasure this value the next time the insert is removed from the magnet. Both are comparable sources of error to one another, $\approx 10\%$, depending on the temperature at the time of the digitization; warmer digitizations ϵ error dominates and vice-versa. The fit is then performed which gives an uncertainty in the fitted values; this error was then combined with the systematic error when setting the uncertainty in the receiver temperature.




Figure 9.30: An example fit from the 1st half of the December magnet ramp cool down, which has significantly reduced errors compared to the 2nd half. The blue error bars represent the systematic error due to the temperature difference between the attenuator and cavity, which is carried over as an uncertainty in the measured power. The orange 1σ lines represent the error in the fit itself. These two errors are then combined when setting the uncertainty in the HFET noise temperature.

The HFET receiver noise temperature plot for the 3 measurements is shown in Figure 9.34. The first half of the December cool down gave the lowest noise temperature of $4.1 \pm 0.25K$ at 5521 MHz, the frequency the cavity is currently tuned to. This measurement agrees the best with the run 1C measured temperature, $3.7 \pm 0.2K$ at 4798 MHz, and the run 1A temperature, $4.0 \pm 0.3K$, albeit it predicts a warmer temperature at that lower frequency of $5.2 \pm 0.2K$. The second half of the December cool down has the most bizarre, inconsistent curve, and predicts the highest noise temperature of $10.1 \pm 0.3K$ at 5521 MHz; because of the constrained set of frequencies, it is also only, roughly, valid from 5.2-5.8 MHz. The March cool down data produced a result much closer in shape to the first half of the December cool down, but that had an overall noise temperature in between the two December results,




Figure 9.31: An example fit from the March fridge cool down. The blue error bars represent the systematic error due to the temperature difference between the attenuator and cavity, which is carried over as an uncertainty in the measured power. This measured temperature difference was roughly 20-30 mK higher than in the December cool down, increasing the overall error. The orange 1σ lines represent the error in the fit itself. These two errors are then combined when setting the error in the HFET noise temperature.

$7.9 \pm 0.2K$ at 5521 MHz. Nonetheless, none of these measurements statistically agreed with one another, and a very conservative, combined result had to be set using data from all 3 measurements, as shown in Figure 9.34. This increased the uncertainty in the measurement quite a bit. The final combined noise temperature was $8.7 \pm 1.7K$ at 5521 MHz. This combined result would be used in the subsequent axion search analysis for calculating the overall system noise temperature.

9.5 Axion search analysis

The Sidecar's ability to search for axions has been very limited so far in run 1D. The piezo-electric rotary motor has been unable to tune the cavity, making the search limited to the




Figure 9.32: The registered temperature difference between the attenuator and cavity during the December magnet ramp cool down measurement.




Figure 9.33: The registered temperature difference between the attenuator and cavity during the March cool down measurement.

bandwidth of the digitizer. It was for this reason that the digitizer bandwidth was expanded to a maximum width of 3 MHz from its 2 MHz value during run 1C, which had also been expanded for the same reason. A bandwidth larger 3 MHz was not possible because of the




Figure 9.34: The HFET noise temperature versus frequency for each of the three noise calibration measurements, as well as the combined result for all three. Note that the December cool down part 2 measurement was taken at a different set of frequencies than the other two. Since all three measurements seem to disagree with one another, a conservative combined result was made to be used in the subsequent axion search analysis.

presence of a 3 MHz pass-band filter in the warm electronics. Due to complications with programming the digitizer, another quirk of this change was that the width of a single frequency bin was kept at 100 Hz, and the number of bins expanded to 30,000, whereas it is typically only 500-2000 in the main cavity analysis. This significant increase in the number of bins and the fact that the same frequencies were being probed in hundreds of digitizations posed some unique problems to running the analysis. To add to this, a very low loaded quality

factor, ≈ 1900 , and higher than normal HFET temperature, $8.7 \pm 1.7K$, as discussed earlier, lowered the sensitivity of any given individual scan.

The Sidecar cavity antenna was critically coupled on December 21st 2023, and the JTWPA enabled and biased for data-taking. Due to complications with the JTWPA re-biasing script not working very well when the cavity isn't tuning, the initial data did not keep a consistent TWPA gain; the script was fixed by January 11th and data-taking commenced in earnest with an average TWPA gain of about 10 dB. It was noticed later in the run that the JTWPA was being compressed when measuring its gain, because of the network analyzer power being set too high; somewhat fortunately, this resulted in the gain being underestimated, and therefore using this data would be fine as it would only produce a more conservative limit.

It also became apparent that, when running the analysis on hundreds of digitization scans with the relatively the same frequencies, that co-adding spectra after individually fitting them, compounded otherwise small errors into much larger errors in the grand spectrum; normally there are much fewer overlapping spectra that get added together so this error is mitigated. This resulted in a non-Gaussian SNR spectrum, even for a data set of less than one day. Solving this issue is still in progress at the time of writing. For the purposes of this dissertation, the subsequent sections go through the analysis performed on a single random day of data-taking, January 15th, even though there are months of data at this point. It is a goal of the future Sidecar analysis to resolve this issue and determine the maximum length dataset one could use before further systematic errors arise and limit further axion sensitivity.

9.5.1 Extracting the Background Warm Receiver Shape

The first step in the analysis is to fit the background in the digitizer due to the warm receiver components. This is done by disconnecting the coaxial connections from the insert, terminating them, and performing a wide power measurement for at least 3 cycles. This is the same wide power measurement used in characterizing the HFET noise temperature.

These digitizations' timestamps are called out to be used in the background determination, and are cut from the axion dataset. In this case, a Savitisky-Golay filter was used to fit the background with a window size of 7700 points and of 3rd order; the main experiment currently uses the same but with a window size of just 17 points. The window size was increased significantly because the number of bins had been increased by a factor of 60. In theory, keeping this value at 17 should be fine, because, since it's a warm receiver measurement, there is no chance of a narrow axion signal being hidden in this background. Nonetheless, if kept small, it would introduce fine structure when divided out of the individual spectra later in the analysis; this exacerbated the compounding fit error issues we were seeing at the stage of this analysis (more on this later). Figure 9.35 shows the warm receiver shape and its fit, and Figure 9.36 shows the resultant power residue after subtracting the fit; the residue is a normalized noise spectrum.




Figure 9.35: The Sidecar warm receiver shape and its fit for background extraction.

Parameter	Average Value for 14 hour dataset
Reflection frequency	5521.052 ± 0.035 MHz
Transmission frequency	5521.311 ± 0.024 MHz
Reflection Q	1919 ± 209
Transmission Q	1941 ± 40
β coupling	1.01 ± 0.03
transmission Δf FWHM	2.84 ± 0.06 MHz
JTWPA gain	12.40 ± 0.06 dB
JTWPA SNRI	9.78 ± 0.02 dB
T_{HFET} noise temperature	9.34 ± 1.78 K
T_{sys} noise temperature	1.37 ± 0.42 K
$V_{eff} = CV_{cav}$ effective volume	0.165 ± 0.005 L
B_0 average field	3.585 ± 0.000003 T

Table 9.2: The average value of run parameters and their uncertainties for a random 14 hour data-set used for initial Sidecar run 1D results.

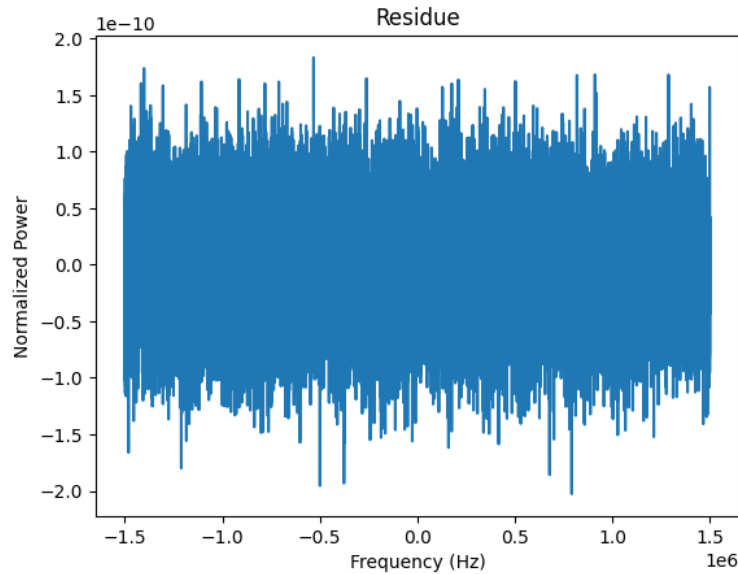


Figure 9.36: The Sidecar warm receiver residue after the background fit is removed.

9.5.2 Run Parameter Extraction

The next major step in the analysis pulls relevant sensor and RF calibration values during the time each raw spectra used in the run dataset was taken. This includes all the variables that go into calculating Equations 3.25 and 3.26: f_0 , Q_0 , β , C , V , B_0 , and T_{sys} . The T_{sys} is calculated using Equation 5.11, with values of the T_{HFET} derived in the previous section, ϵ from ex-situ measurements, and the SNRI of the TWPA measured during the last calibration period of the spectra; the SNRI is calculated from Equation 5.9 based on the TWPA gain. The average B_0 for Sidecar is calculated based on the current in the main magnet and assuming a constant ratio of field strength to current (3.58 T/220 A). The C and V are coded in a form factor configuration file based on the COMSOL simulations done for the TM_{010} mode; they are actually expressed in the analysis as $V_{eff} = V_{cav}C(f_0)$, an effective volume based on the form factor at the given resonant frequency. The f_0 , β , and Q_L are pulled from the periodic network analyzer measurements in the data-taking cycle and associated with the timestamp of each individual spectra; both transmission and reflection

values are pulled, but the transmission values of f_0 and Q_L are what are used for weighting the spectra later. Table 9.2 lists the average values of these various parameters for the 14 hour data set that was used to set initial results.

A weakness in the Sidecar data-taking process is that it doesn't stop taking data during helium reservoir fills. These happen approximately once a day, and last only a few minutes. This is a problem because it causes several of these parameters to shift temporarily for about 1-3 hours, predominantly the resonant frequency, coupling, and loaded quality factors; this temporary shift is visible in a time plot of the quality factor for the specific day that this initial analysis used (Figure 9.37). As one can see its a relatively small shift compared to the normal variance in the data. No cuts were made on this data, but it is worth noting, as it may play into the systematic errors of a more precise search.




Figure 9.37: The Sidecar loaded quality factor (transmission) over the course of the 14 hour data set used for analysis. Evident in this plot is a small shift in Q (≈ 150) due to a helium reservoir fill that took place roughly at hour 4. The value recovers after about 2 hours, but it is shifted by ≈ 75 ; This is probably within the uncertainty in Q for any individual fit, so it is not considered a problem at the moment.

Finally, the JTWPA gain performance over the course of the day used for analysis is plotted in Figure 9.38; it is more resilient to the helium fills, but also has a larger variance

to begin with throughout the day.




Figure 9.38: The Sidecar TWPA gain performance over the course of a 14 hour data set.

9.5.3 Raw Spectrum Preparation

The next step in a typical axion analysis would be to prepare all of the raw spectra taken in the data-taking time period. The typical steps in this process for each spectra would be:

1. Check to see if the spectra should be cut from the dataset for several reasons (time cuts, parameter cuts, etc.). This isn't very relevant in this set because there were no cuts during the 14-hour period used for this analysis.
2. Subtract/divide out the warm receiver fit shape from the raw spectra that was made earlier. This is shown in Figure 9.40.
3. Perform a "cold-receiver" filter fit to the individual spectra and divide it out from the individual spectra. Figure 9.39 shows a raw spectrum and its Padé fit at this stage of the process. Special care has to be taken here not to filter out potential axion signals.

A Padé polynomial fit was used in this analysis. This produces a normalized filtered spectra that is shown in Figure 9.41.

4. Apply a Lorentzian weighting to each bin of the spectra according to the quality factor and resonant frequency associated with the individual spectra. This ensures that when co-adding spectra later that bins further off-resonance do not get counted equal to bins closer to resonance. It is expressed in yocto-watts/Kelvin because a system noise temperature has not been applied to the spectrum yet; this will be done in the grand assembly. This is shown in Figure 9.42.




Figure 9.39: An example raw spectrum shown in blue with its cold receiver Padé fit shown in orange.

This process was done for the 408 spectra that were in the random 14-hour data set chosen. The final statistics of the process agreed well with the ideal. The distribution of the standard deviations was Gaussian about the expected value with the expected variance. This variance should be 0.01 about a mean value of 1. The standard deviation increase, which is the ratio of the individual spectrum standard deviation to the expected value, is plotted




Figure 9.40: An example raw spectrum with the warm receiver shape removed, but no cold receiver fit applied; there is still residual structure unaccounted for.




Figure 9.41: The raw spectrum after the cold receiver fit being removed. It is now normalized, and should be flat and centered about zero power.

as a function of digitization index in Figure 9.43. The deviation of individual bins about the mean is plotted in Figure 9.44. This was also very Gaussian, indicating this process

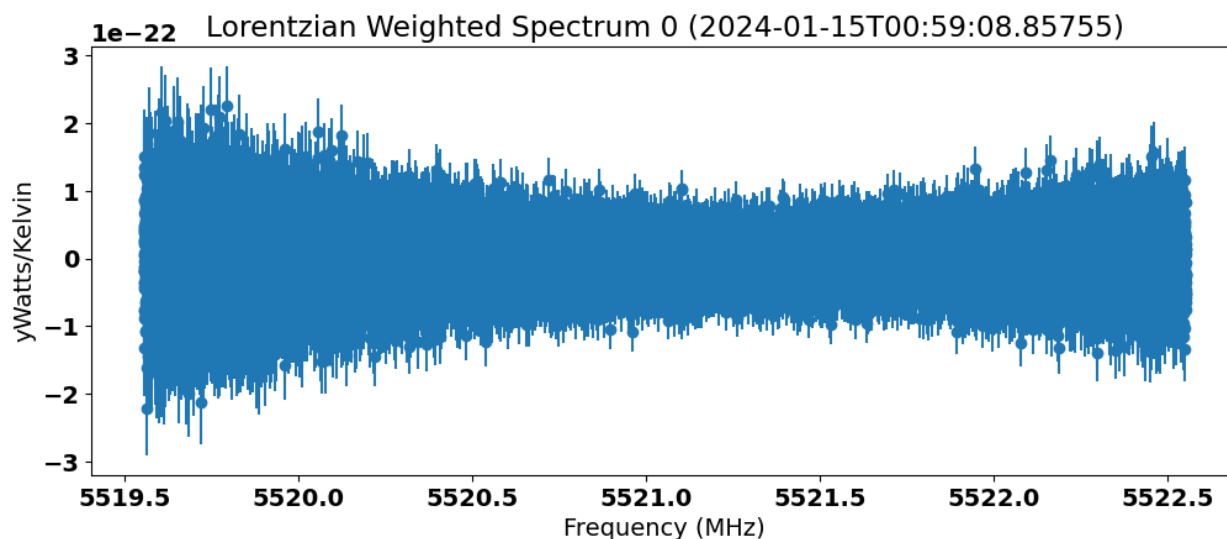


Figure 9.42: The final Lorentzian weighted spectrum sent for co-adding into a grand spectrum. The Lorentzian weighting applied matches the associated resonant frequency and quality factor at the time of the digitization.

successfully normalized the noise for each individual spectra.

We will run into problems with what we did here in the next step however; the cold receiver fit isn't perfect and there is some receiver structure in each bin that wasn't captured. This means when we add many of these spectra together that error will start to compound into a larger error. This isn't an issue in a normal analysis with cavity tuning because only on the order of 10 spectra get added into a single grand spectrum assembly bin, and the fit error is uncorrelated with the grand spectrum bins across spectra because there is tuning between each spectra. In this case, because all the spectra cover the relatively same frequency range, this error will compound on the order of the number of spectra included in the set, 408 scans in this case. More on this next.

9.5.4 Grand Spectrum Assembly

The next step in the analysis is to co-add all of these prepared spectra into a grand spectrum as outlined in section 5.5.3. Fig 9.45 shows the number of spectra being added in any given




Figure 9.43: The standard deviation increase versus the spectrum index of the 408 spectra included in the data set. The data is centered about the expected mean with a ± 0.01 variance as expected.




Figure 9.44: The individual bin deviations of raw spectra about the mean fitted to a Gaussian distribution.

frequency bin of what will become the grand spectrum. Because there is no tuning, almost every spectra in the sample set, 408 total, contributes to every bin in the grand spectrum,

with the exception of the fringe frequencies, because there are slight differences in the center frequency for each spectrum.




Figure 9.45: The number of raw spectra that contributed to each frequency bin in the grand spectrum.

The DFSZ SNR for each frequency in the grand spectrum is shown in Figure 9.46; remember that Sidecar is not meant to be sensitive to DFSZ axions, hence the very low SNR. These DFSZ SNR sensitivity values are in agreement with the expected sensitivity of Sidecar for a set of 408 spectra. Again, because this data set has no tuning, a unique expected feature of this 'static' spectrum is that it should produce a Lorentzian shape about the resonant frequency as seen in the plot. The reason it isn't centered in the plot is because of a slight mismatch between the target frequency that the DAQ instructed the digitizer to center spectra to and the measured resonant frequency of the cavity; it is a difference of about 200 kHz.

The final spectrum is expressed as power excess in each bin, and its associated standard deviation. In this way, any spikes in power above 3σ in the grand spectrum would be flagged as potential axion candidates; this wasn't the case in this Sidecar dataset. Furthermore, one




Figure 9.46: The DFSZ SNR of the grand spectrum versus frequency. Its Lorentzian shape is a product of the Lorentzian weighting of the individual raw spectra. The off-center peak is due to a mismatch between the spectra centers and the measured resonant frequency of the cavity.

expects a Gaussian distribution about the center frequency for a set of data all taken about the same center frequency; in the case of a data set where there was tuning at a constant rate, one expects this shape to be averaged out and the grand spectrum flattened, as is the case for normal ADMX grand spectra. Figure 9.47 shows the grand spectrum for the sidecar dataset; it is filled with excess structure and not ideal. Figure 9.48 shows the histogram of SNR excesses, which does not fit well to its expected Gaussian, indicating unknown structure in the combined dataset.

This excess structure in the grand spectrum is presumed to be because of the error alluded to earlier; the individual error in each cold receiver fit is being compounded as the raw spectra are being added together. Normally this error would not grow very large if there is tuning, because the raw spectra bins would not all align with each other in the same grand spectra bin very well; each spectra would be offset by a tuning step. This error only grows for the number of spectra included in the dataset, which is why only a 14 hour data set

is used here. The solution to this would be to restructure the analysis to perform the raw spectra adding procedure before the cold receiver fit, and then perform a single cold receiver fit on the combined spectrum to normalize it. This, unfortunately, requires some significant changes to the analysis code to get working and is outside the scope of this dissertation. However, the plan is to solve this eventually and be able to expand the size of the dataset to improve the search limits. Nonetheless, with the less than ideal grand spectrum shown here, some axion limits can be set.




Figure 9.47: The grand power excess spectrum versus frequency for the Sidecar run 1D sample dataset. It is apparent that there is significant noise from the receiver leaking through due to the compounding error from individual spectra fits. Ideally, this spectrum should be Gaussian for a static frequency dataset.

9.5.5 Axion limits

Although the traditional analysis process would require one to temper the grand spectrum to make sure there are no axion candidates hidden within the residual structure of Figure 9.47. Considering Sidecar is not even close to KSVZ or DFSZ sensitivity, one can go forward and set a limit under the assumption that all candidates were eliminated within the sensitivity of Sidecar. This is done by setting a minimum value of $g_{a\gamma\gamma}$ based on a 90% confidence




Figure 9.48: A histogram of the SNR excess in the grand spectrum.

limit. As shown in Figs. 9.49, 9.50, for the 14 hour sample data set used throughout this section, the average minimum $g_{a\gamma\gamma}$ set was $2.71 \pm 0.70 \times 10^{-13} GeV^{-1}$ between 5519.5 and 5522.6 MHz (22.823-22.836 μeV). The previous best limit was Sidecar run 1A with $g_{a\gamma\gamma} = 2.46 \pm 0.06 \times 10^{-12} GeV^{-1}$, which tuned through this 3 MHz region in 2016. By just increasing the number of spectra at this very specific frequency to the order of hours, one is able to get a limit 10x better. Assuming this analysis could be run for a longer data set with a leveled grand spectrum, one could theoretically use the 2+ months of Sidecar data to drill down further; this should follow a \sqrt{N} relation with the number of spectra. Expanding from a 14 hour data set to 30 days would then theoretically improve the limit by a factor 7.7. However one will eventually run up against further systematic issues as the data set is expanded to be longer and longer. It is a goal of the future Sidecar paper to find what this maximum set length would be, given that Sidecar is potentially stuck at this frequency for many months more before an extraction can occur.




Figure 9.49: Preliminary axion exclusion limits for Sidecar run 1D, based on a 14 hour dataset. This work is shown in green against the previous Sidecar run 1A limits in this region (red), and the CAST broadband limit in this region (blue).

9.6 Future prospects for Sidecar

As should be apparent from the previous sections, there are still quite a few mysteries and problems to solve on the current ADMX Sidecar; even though new, better limits could be set in the small frequency range the cavity was stuck taking data on.

The performance of the superconducting tuning rod remains unclear as it stands; there maybe RF power leaking through the apertures in the clamshell that are causing the very low baseline Q , but one can't know for sure. This degraded performance from external factors is entangled with the performance of the tuning rod. Plans are being made to test the second, solid Nb_3Sn tuning rod in another copy of the Sidecar cavity within a dilution refrigerator at




Figure 9.50: A zoomed in exclusion limit plot of the Sidecar run 1D preliminary dataset shown in green. The previous Sidecar run 1A limit is shown in red, and the CAST experimental limit shown in blue.

LLNL. This will give a new baseline zero field Q. Additionally, because the amplification will only be done through LNF HFETs, cavity data can hopefully be taken at higher temperatures during the cool down, opening the possibility of observing the superconducting transition. This fridge can be opened on a more regular basis, and therefore a variety of tests can be performed quantifying any potential leakage from the clam shell. In addition there is some work to be done on optimizing the geometric factor calculations and determining optimal tuning rod angles for performing a surface decomposition on Sidecar successfully.

In terms of data-taking and analysis, Sidecar will continue to take axion data until approximately December 2024. In that time, it may be possible to attempt to tune the rotor piezoelectric motor again; even a small change could move the TM_{020} mode away from its mode crossing, enabling axion data taking on that mode as well. There is still the possibility

of successfully actuating the coaxial switch by directly connecting it to its power supply; this would enable a dedicated hot load measurement, which could improve the HFET noise calibration. As it stands, the HFET noise temperature is significantly degraded compared to previous runs, potentially because its recent repairs were not successful; any further noise measurements would only help. Finally, the filtering error issues present in the power excess spectrum must be resolved to get a publishable limit. After this is resolved, increasing the size of the dataset further will allow the limits to be improved by an order of magnitude potentially.

All of this work will eventually inform the future direction of ADMX. Superconducting cavities have gained a lot of interest in my time studying them as an ADMX graduate. The CAPP institute has continued to push the limits of their YBCO HTS superconducting cavities, with many anticipated results coming soon. With more axion experiments appearing in the field, competition for the immediate higher frequencies above 1 GHz is possible. Being able to tune at a competitive rate is crucial for ADMX to remain relevant; superconducting cavities may be a necessity to do so. To implement this technology successfully will involve taking this graduate prototype experiment into an engineering development phase for ADMX-EFR, ensuring the cavities will meet the quality factor standard to operate. There is still much work to be done for anyone interested in either axions or superconducting cavities!

BIBLIOGRAPHY

- [1] 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. Hélaine, 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. Lemière, 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. Quéméner, M. Rawlik, D. Rebreyend, I. Rienäcker, D. Ries, S. Roccia, G. Rogel, D. Rozpedzik, A. Schnabel, P. Schmidt-Wellenburg, N. Severijns, D. Shiers, R. Tavakoli Dinani, J. A. Thorne, R. Virost, J. Voigt, A. Weis, E. Wursten, G. Wyszynski, J. Zejma, J. Zenner, and G. Zsigmond. Measurement of the permanent electric dipole moment of the neutron. *Phys. Rev. Lett.*, 124:081803, Feb 2020.
- [2] A.A. Abrikosov. The magnetic properties of superconducting alloys. *Journal of Physics and Chemistry of Solids*, 2(3):199–208, 1957.
- [3] Adwaele. Dilution Refrigerator Wiki.
- [4] Danho Ahn, Ohjoon Kwon, Woohyun Chung, Wonjun Jang, Doyu Lee, Jinhwan Lee, Sung Woo Youn, HeeSu Byun, Dojun Youm, and Yannis K. Semertzidis. First prototype of a biaxially textured $\text{YBa}_2\text{Cu}_3\text{O}_{7-x}$ microwave cavity in a high magnetic field for dark matter axion search, 2021.
- [5] Danho Ahn, Ohjoon Kwon, Woohyun Chung, Wonjun Jang, Doyu Lee, Jinhwan Lee, Sung Woo Youn, HeeSu Byun, Dojun Youm, and Yannis K. Semertzidis. Biaxially textured $\text{YBa}_2\text{Cu}_3\text{O}_{7-x}$ microwave cavity in a high magnetic field for a dark-matter axion search. *Phys. Rev. Appl.*, 17:L061005, Jun 2022.
- [6] Danho Ahn, Dojun Youm, Ohjoon Kwon, Woohyun Chung, and Yannis K. Semertzidis. High quality factor high-temperature superconducting microwave cavity development for the dark matter axion search in a strong magnetic field. *arXiv*, 2019. arXiv:1902.04551.

- [7] D. Alesini, C. Braggio, G. Carugno, N. Crescini, D. D'Agostino, D. Di Gioacchino, and S. Tocci. Galactic axions search with a superconducting resonant cavity. *Physics Review D*, 99:101101, May 2019.
- [8] Andrea Alimenti, Nicola Pompeo, Kostiantyn Torokhtii, Tiziana Spina, René Flükiger, Luigi Muzzi, and Enrico Silva. Surface impedance measurements on nb₃sn in high magnetic fields. *IEEE Transactions on Applied Superconductivity*, 29(5):1–4, 2019.
- [9] V. Anastassopoulos, S. Aune, K. Barth, A. Belov, H. Bräuninger, G. Cantatore, J. M. Carmona, J. F. Castel, S. A. Cetin, F. Christensen, J. I. Collar, T. Dafni, M. Davenport, T. A. Decker, A. Dermenev, K. Desch, C. Eleftheriadis, G. Fanourakis, E. Ferrer-Ribas, H. Fischer, J. A. García, A. Gardikiotis, J. G. Garza, E. N. Gazis, T. Gerasis, I. Giomataris, S. Gninenko, C. J. Hailey, M. D. Hasinoff, D. H. H. Hoffmann, F. J. Iguaz, I. G. Irastorza, A. Jakobsen, J. Jacoby, K. Jakovčić, J. Kaminski, M. Karuza, N. Kralj, M. Krčmar, S. Kostoglou, Ch. Krieger, B. Lakić, J. M. Laurent, A. Liolios, A. Ljubičić, G. Luzón, M. Maroudas, L. Miceli, S. Neff, I. Ortega, T. Papaevangelou, K. Paraschou, M. J. Pivovarov, G. Raffelt, M. Rosu, J. Ruz, E. Ruiz Chóliz, I. Savvidis, S. Schmidt, Y. K. Semertzidis, S. K. Solanki, L. Stewart, T. Vafeiadis, J. K. Vogel, S. C. Yildiz, K. Zioutas, and C. A. S. T. Collaboration. New cast limit on the axion–photon interaction. *Nature Physics*, 13(6):584–590, Jun 2017.
- [10] A. Arbey and F. Mahmoudi. Dark matter and the early universe: A review. *Progress in Particle and Nuclear Physics*, 119:103865, 2021.
- [11] Asimina Arvanitaki, Masha Baryakhtar, and Xinlu Huang. Discovering the qcd axion with black holes and gravitational waves. *Phys. Rev. D*, 91:084011, Apr 2015.
- [12] S. Asztalos, E. Daw, H. Peng, L. J Rosenberg, C. Hagmann, D. Kinion, W. Stoeffl, K. van Bibber, P. Sikivie, N. S. Sullivan, D. B. Tanner, F. Nezrick, M. S. Turner, D. M. Moltz, J. Powell, M.-O. André, J. Clarke, M. Mück, and Richard F. Bradley. Large-scale microwave cavity search for dark-matter axions. *Phys. Rev. D*, 64:092003, Oct 2001.
- [13] 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.
- [14] Kyu Jung Bae, Ji-Haeng Huh, and Jihn E Kim. Updating the axion cold dark matter energy density. *Journal of Cosmology and Astroparticle Physics*, 2008(09):005, sep 2008.

- [15] Guillermo Ballesteros, Javier Redondo, Andreas Ringwald, and Carlos Tamarit. Standard model—axion—seesaw—higgs portal inflation. five problems of particle physics and cosmology solved in one stroke. *Journal of Cosmology and Astroparticle Physics*, 2017(08):001, aug 2017.
- [16] R. Ballou, G. Deferne, M. Finger, M. Finger, L. Flekova, J. Hosek, S. Kunc, K. Macuchova, K. A. Meissner, P. Pugnatt, M. Schott, A. Siemko, M. Slunecka, M. Sulc, C. Weinsheimer, and J. Zicha. New exclusion limits on scalar and pseudoscalar axion-like particles from light shining through a wall. *Phys. Rev. D*, 92:092002, Nov 2015.
- [17] J. Bardeen, L. N. Cooper, and J. R. Schrieffer. Microscopic theory of superconductivity. *Phys. Rev.*, 106:162–164, Apr 1957.
- [18] S. M. Barr. Solving the strong CP problem without the peccei-quinn symmetry. *Phys. Rev. Lett.*, 53:329–332, Jul 1984.
- [19] 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.
- [20] C. Bartram, T. Braine, R. Cervantes, N. Crisosto, N. Du, 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, C. Cisneros, 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. Murch, M. Goryachev, B. T. McAllister, A. Quiskamp, C. Thomson, M. E. Tobar, V. Bolkhovsky, G. Calusine, W. Oliver, and K. Serniak. Dark matter axion search using a Josephson Traveling wave parametric amplifier. *Review of Scientific Instruments*, 94(4):044703, 04 2023.
- [21] 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.
- [22] Gianfranco Bertone and Tim M. P. Tait. A new era in the search for dark matter. *Nature*, 562(7725):51–56, Oct 2018.
- [23] Frank J. Blatt. *Modern Physics*. McGraw-Hill, 1992. Type II superconductors critical fields and temperatures.
- [24] Sz. Borsanyi, Z. Fodor, K. H. Kampert, S. D. Katz, T. Kawanai, T. G. Kovacs, S. W. Mages, A. Pasztor, F. Pittler, J. Redondo, A. Ringwald, and K. K. Szabo. Lattice qcd for cosmology, 2016.
- [25] Frederick Bouquet and Julien Bobroff. Phase diagram (b, t) of a type i superconductor, 2011.
- [26] Frederick Bouquet and Julien Bobroff. Phase diagram (b, t) of a type ii superconductor, 2011.
- [27] C. Boutan, M. Jones, B. H. LaRoque, N. S. Oblath, R. Cervantes, N. Du, N. Force, S. Kimes, R. Ottens, L. J. Rosenberg, G. Rybka, J. Yang, G. Carosi, N. Woollett, D. Bowring, A. S. Chou, R. Khatiwada, A. Sonnenschein, W. Wester, R. Bradley, E. J. Daw, A. Agrawal, A. V. Dixit, J. Clarke, S. R. O'Kelley, N. Crisosto, J. R. Gleason, S. Jois, P. Sikivie, I. Stern, N. S. Sullivan, D. B. Tanner, P. M. Harrington, and E. Lentz. Piezoelectrically tuned multimode cavity search for axion dark matter. *Phys. Rev. Lett.*, 121:261302, Dec 2018.
- [28] Christian Boutan. *A Piezoelectrically Tuned RF-Cavity Search for Dark Matter Axions*. PhD thesis, University of Washington, 2017.
- [29] T. Braine, R. Cervantes, N. Crisosto, N. Du, S. Kimes, L. J. Rosenberg, G. Rybka, J. Yang, D. Bowring, A. S. Chou, R. Khatiwada, A. Sonnenschein, W. Wester, G. Carosi, N. Woollett, L. D. Duffy, R. Bradley, C. Boutan, M. Jones, B. H. LaRoque, N. S. Oblath, M. S. Taubman, J. 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, J. H. Buckley, P. M. Harrington, E. A. Henriksen, and K. W. Murch. Extended search for the invisible axion with the axion dark matter experiment. *Phys. Rev. Lett.*, 124:101303, Mar 2020.

- [30] T. Braine, G. Rybka, A. A. Baker, J. Brodsky, G. Carosi, N. Du, N. Woollett, S. Knirck, M. Jones, and ADMX Collaboration. Multi-mode analysis of surface losses in a superconducting microwave resonator in high magnetic fields. *Review of Scientific Instruments*, 94(3):033102, 03 2023.
- [31] B. M. Brubaker, L. Zhong, S. K. Lamoreaux, K. W. Lehnert, and K. A. van Bibber. Haystac axion search analysis procedure. *Phys. Rev. D*, 96:123008, Dec 2017.
- [32] Benjamin M. Brubaker. *First results from the HAYSTAC axion search*. PhD thesis, Yale University), 2018.
- [33] Adam Burrows, M. Ted Ressel, and Michael S. Turner. Axions and sn 1987a: Axion trapping. *Phys. Rev. D*, 42:3297–3309, Nov 1990.
- [34] R. Cervantes, C. Braggio, B. Gioccone, D. Frolov, A. Grasselino, R. Harnik, O. Melnychuk, R. Pilipenko, S. Posen, and A. Romanenko. Deepest sensitivity to wavelike dark photon dark matter with srf cavities. *arXiv*, August 2022. arXiv:2208.03183.
- [35] 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 μeV dark photon dark matter: Detailed design, operations, and analysis. *Phys. Rev. D*, 106:102002, Nov 2022.
- [36] Raphael Cervantes. *A Search for Wavelike Dark Matter with Dielectrically-loaded Multimode Cavities*. PhD thesis, Universtiy of Washington, 2021.
- [37] L.F. Chen, C. Ong, C. P. Neo, V.V. Varadan, and V.K. Varadan. *Microwave Electronics: Measurement and Materials Characterization*. John Wiley and Sons Ltd., 2004.
- [38] Mark W. Coffey and John R. Clem. Unified theory of effects of vortex pinning and flux creep upon the rf surface impedance of type-ii superconductors. *Phys. Rev. Lett.*, 67:386–389, Jul 1991.
- [39] N. Crisosto, P. Sikivie, N. S. Sullivan, D. B. Tanner, J. Yang, and G. Rybka. Admx slic: Results from a superconducting lc circuit investigating cold axions. *Phys. Rev. Lett.*, 124:241101, Jun 2020.
- [40] Edward Daw. *A Search for Halo Axions*. PhD thesis, Massachusetts Institute of Technology, 1998.

- [41] Daniele Di Gioacchino, Claudio Gatti, David Alesini, Carlo Ligi, Simone Tocci, Alessio Rettaroli, Giovanni Carugno, Nicolo Crescini, Giuseppe Ruoso, Cateriana Braggio, Paolo Falferi, C. Sebastiano Gallo, Umberto Gambardella, Gerardo Iannone, Gianluca Lamanna, Augusto Lombardi, Renato Mezzena, Antonello Ortolan, Ruggero Pengo, Enrico Silva, and Nicola Pompeo. Microwave losses in a dc magnetic field in superconducting cavities for axion studies. *IEEE Transactions on Applied Superconductivity*, 29(5):1–5, 2019.
- [42] R. H. Dicke. The Measurement of Thermal Radiation at Microwave Frequencies. *Review of Scientific Instruments*, 17(7):268–275, 12 1946.
- [43] Duane A. Dicus, Edward W. Kolb, Vigdor L. Teplitz, and Robert V. Wagoner. Astrophysical bounds on the masses of axions and higgs particles. *Phys. Rev. D*, 18:1829–1834, Sep 1978.
- [44] 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, 1981.
- [45] A. P. Drozdov, P. P. Kong, V. S. Minkov, S. P. Besedin, M. A. Kuzovnikov, S. Mozaffari, L. Balicas, F. F. Balakirev, D. E. Graf, V. B. Prakapenka, E. Greenberg, D. A. Knyazev, M. Tkacz, and M. I. Eremets. Superconductivity at 250 k in lanthanum hydride under high pressures. *Nature*, 569(7757):528–531, May 2019.
- [46] N. Du, N. Force, R. Khatiwada, E. Lentz, R. Ottens, L. J Rosenberg, G. Rybka, G. Carosi, N. Woollett, D. Bowring, A. S. Chou, A. Sonnenschein, W. Wester, C. Boutan, N. S. Oblath, R. Bradley, E. J. Daw, A. V. Dixit, J. Clarke, S. R. O’Kelley, N. Crisosto, J. R. Gleason, S. Jois, P. Sikivie, I. Stern, N. S. Sullivan, D. B Tanner, and G. C. Hilton. Search for invisible axion dark matter with the axion dark matter experiment. *Phys. Rev. Lett.*, 120:151301, Apr 2018.
- [47] Nick Du. *Extended Search for Axion Dark Matter with ADMX*. PhD thesis, University of Washington, 2021.
- [48] L. D. Duffy and P. Sikivie. Caustic ring model of the milky way halo. *Phys. Rev. D*, 78:063508, Sep 2008.
- [49] 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.

- [50] 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.
- [51] Fermilab. Superconducting quantum materials and systems, 2020.
- [52] R.P. Feynman. Chapter ii application of quantum mechanics to liquid helium. In C.J. Gorter, editor, *Progress in Low Temperature Physics*, volume 1, pages 17–53. Elsevier, 1955.
- [53] José A. Flores-Livas, Lilia Boeri, Antonio Sanna, Gianni Profeta, Ryotaro Arita, and Mikhail Erements. A perspective on conventional high-temperature superconductors at high pressure: Methods and materials. *Physics Reports*, 856:1–78, 2020. A perspective on conventional high-temperature superconductors at high pressure: Methods and materials.
- [54] Jonathan I. Gittleman and Bruce Rosenblum. Radio-frequency resistance in the mixed state for subcritical currents. *Phys. Rev. Lett.*, 16:734–736, Apr 1966.
- [55] Paolo Gondolo and Georg G. Raffelt. Solar neutrino limit on axions and keV-mass bosons. *Phys. Rev. D*, 79:107301, May 2009.
- [56] Alexander V. Gramolin, Deniz Aybas, Dorian Johnson, Janos Adam, and Alexander O. Sushkov. Search for axion-like dark matter with ferromagnets. *Nature Physics*, 17(1):79–84, August 2020.
- [57] Daniel Grin, Giovanni Covone, Jean-Paul Kneib, Marc Kamionkowski, Andrew Blain, and Eric Jullo. Telescope search for decaying relic axions. *Phys. Rev. D*, 75:105018, May 2007.
- [58] Alan H. Guth. Inflationary universe: A possible solution to the horizon and flatness problems. *Phys. Rev. D*, 23:347–356, Jan 1981.
- [59] C. Hagmann, S. Chang, and P. Sikivie. Axions from string decay. *Nuclear Physics B - Proceedings Supplements*, 72:81–86, 1999. Proceedings of the 5th IFT Workshop on Axions.
- [60] 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.

- [61] Christian A. Hagmann. *A Search for Cosmic Axions*. PhD thesis, University of Florida, 1990.
- [62] B.W. Hakki and P.D. Coleman. A dielectric resonator method of measuring inductive capacities in the millimeter range. *IRE Transactions on Microwave Theory and Techniques*, 8(4):402–410, 1960.
- [63] D.L. Hall, J.J. Kaufman, M. Liepe, J.T. Maniscalco, and R.D. Porter. High Performance Next-Generation Nb₃Sn Cavities for Future High Efficiency SRF Linacs. In *Proc. of Linear Accelerator Conference (LINAC'16), East Lansing, MI, USA, 25-30 September 2016*, number 28 in Linear Accelerator Conference, pages 398–401, Geneva, Switzerland, May 2017. JACoW. <https://doi.org/10.18429/JACoW-LINAC2016-TUOP07>.
- [64] Keisuke Harigaya and Masahiro Kawasaki. Qcd axion dark matter from long-lived domain walls during matter domination. *Physics Letters B*, 782:1–5, 2018.
- [65] A. T. Hirshfeld, H. A. Leupold, and H. A. Boorse. Superconducting and normal specific heats of niobium. *Phys. Rev.*, 127:1501–1507, Sep 1962.
- [66] Michael Hotz. *A SQUID-Based RF Cavity Search for Dark Matter Axions*. PhD thesis, University of Washington, 2013.
- [67] John David Jackson. *Classical electrodynamics*. Wiley, New York, NY, 3rd ed. edition, 1999.
- [68] Ryan Janish and Elena Pinetti. Hunting dark matter lines in the infrared background with the james webb space telescope, 2023.
- [69] Darko Kajfez. *Q-Factor*. John Wiley & Sons, Ltd, 1994.
- [70] Jihn E. Kim. Weak-interaction singlet and strong CP invariance. *Phys. Rev. Lett.*, 43:103–107, Jul 1979.
- [71] Darin Kinion. *First Results from a Multiple-Microwave-Cavity Search for Dark-matter Axions*. PhD thesis, University of California Davis, 2001.
- [72] Lawrence Krauss, John Moody, Frank Wilczek, and Donald E. Morris. Calculations for cosmic axion detection. *Phys. Rev. Lett.*, 55:1797–1800, Oct 1985.
- [73] Lev Davidovich Landau. On the theory of phase transitions. i. *Phys. Z. Sowjet.*, 11:26, 1937.

- [74] Lev Davidovich Landau and V L Ginzburg. On the theory of superconductivity. *Zh. Eksp. Teor. Fiz.*, 20:1064, 1950.
- [75] Erik W. Lentz, Thomas R. Quinn, Leslie J. Rosenberg, and Michael J. Tremmel. A new signal model for axion cavity searches from n-body simulations. *The Astrophysical Journal*, 845(2):121, aug 2017.
- [76] K. Leong and J. Mazierska. Precise measurements of the q factor of dielectric resonators in the transmission mode-accounting for noise, crosstalk, delay of uncalibrated lines, coupling loss, and coupling reactance. *IEEE Transactions on Microwave Theory and Techniques*, 50(9):2115–2127, 2002.
- [77] K. Leong, J. Mazierska, and J. Krupka. Measurements of unloaded q-factor of transmission mode dielectric resonators. In *1997 IEEE MTT-S International Microwave Symposium Digest*, volume 3, pages 1639–1642 vol.3, 1997.
- [78] F. London, H. London, and Frederick Alexander Lindemann. The electromagnetic equations of the supraconductor. *Proceedings of the Royal Society of London. Series A - Mathematical and Physical Sciences*, 149(866):71–88, 1935.
- [79] Fritz London. *Superfluids: Macroscopic Theory of Superconductivity*. John Wiley & Sons, 1950.
- [80] Dmitry Lyapustin. *An improved low-temperature RF-cavity search for dark-matter axions*. PhD thesis, University of Washington, 2015.
- [81] Ana T. Malagon. *Search for 140 micro-eV Pseudoscalar and Vector Dark Matter Using Microwave*. PhD thesis, Yale University, 2014.
- [82] M. Martinello, A. Grassellino, M. Checchin, A. Romanenko, O. Melnychuk, D. A. Sergatskov, S. Posen, and J. F. Zasadzinski. Effect of interstitial impurities on the field dependent microwave surface resistance of niobium. *Applied Physics Letters*, 109(6):062601, 08 2016.
- [83] W. Meissner and R. Ochsenfeld. Ein neuer effekt bei eintritt der supraleitfähigkeit. *Naturwissenschaften*, 21(44):787–788, Nov 1933.
- [84] John Miller. The nhmfl 45-t hybrid magnet system: Past, present, and future. *Applied Superconductivity, IEEE Transactions on*, 13:1385 – 1390, 07 2003.
- [85] Ann Nelson. Naturally weak cp violation. *Physics Letters B*, 136(5):387–391, 1984.

- [86] Kevin O'Brien, Chris Macklin, Irfan Siddiqi, and Xiang Zhang. Resonant phase matching of josephson junction traveling wave parametric amplifiers. *Phys. Rev. Lett.*, 113:157001, Oct 2014.
- [87] Ciaran O'HARE. cajohare/axionlimits: Axionlimits, Jul 2020.
- [88] L. Onsager. Statistical hydrodynamics. *Il Nuovo Cimento (1943-1954)*, 6(2):279–287, Mar 1949.
- [89] R. Parks. *Superconductivity in two volumes*, volume 2. Marcel Dekker, Inc, 1969.
- [90] R. D. Peccei and Helen R. Quinn. CP conservation in the presence of pseudoparticles. *Phys. Rev. Lett.*, 38:1440–1443, Jun 1977.
- [91] Paul J. Petersan and Steven M. Anlage. Measurement of resonant frequency and quality factor of microwave resonators: Comparison of methods. *Journal of Applied Physics*, 84(6):3392–3402, 09 1998.
- [92] A.M. Polyakov. Quark confinement and topology of gauge theories. *Nuclear Physics B*, 120(3):429–458, 1977.
- [93] N. Pompeo and E. Silva. Reliable determination of vortex parameters from measurements of the microwave complex resistivity. *Phys. Rev. B*, 78:094503, Sep 2008.
- [94] S. Posen, M. Checchin, O.S. Melnychuk, T. Ring, and I. Gonin. Measurement of high quality factor superconducting cavities in tesla-scale magnetic fields for dark matter searches. *arXiv*, January 2022. arXiv:2201.10733.
- [95] S Posen and D L Hall. Nb3sn superconducting radiofrequency cavities: fabrication, results, properties, and prospects. *Superconductor Science and Technology*, 30(3):033004, jan 2017.
- [96] David M. Pozar. *Microwave Engineering, 4th Edition*. Wiley, 2011.
- [97] Chris W. Purcell, James S. Bullock, and Manoj Kaplinghat. The dark disk of the milky way. *The Astrophysical Journal*, 703(2):2275, sep 2009.
- [98] J. N. RJABININ and L. W. SHUBNIKOW. Magnetic properties and critical currents of supra-conducting alloys. *Nature*, 135(3415):581–582, Apr 1935.
- [99] Joseph Thomas Rogers. *Limits on the Electromagnetic Coupling and Density of Galactic Axions*. PhD thesis, University of Rochester, 1987.

- [100] James William Rohlf. *Modern Physics from a to Z0*. Wiley, 1994. Type I superconductors critical fields and temperatures.
- [101] S. Russenschuck and G. Vandoni. Cern accelerator school: Superconductivity and cryogenics for accelerators and detectors, May 2002.
- [102] 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. *Physical Review Letters*, 127(8), August 2021.
- [103] A. Savitzky and M. J. E. Golay. Smoothing and differentiation of data by simplified least squares procedures. *Analytical Chemistry*, 36:1627–1639, January 1964.
- [104] 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, 1980.
- [105] Peter S. Shternin, Dmitry G. Yakovlev, Craig O. Heinke, Wynn C. G. Ho, and Daniel J. Patnaude. Cooling neutron star in the Cassiopeia A supernova remnant: evidence for superfluidity in the core. *Monthly Notices of the Royal Astronomical Society: Letters*, 412(1):L108–L112, 03 2011.
- [106] P. Sikivie. Experimental tests of the "invisible" axion. *Phys. Rev. Lett.*, 51:1415–1417, Oct 1983.
- [107] P. Sikivie. Detection rates for “invisible”-axion searches. *Phys. Rev. D*, 32:2988–2991, Dec 1985.
- [108] 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. Blonar, 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.
- [109] Can-Li Song and Qi-Kun Xue. Cuprate superconductors may be conventional after all. *American Physical Society Viewpoints*, 2017.
- [110] Attocube systems AG. Eglfinger Weg 2, 85540 Haar, Germany.
- [111] M. Tinkham. *Introduction to Superconductivity*. Mcgraw-Hill, 1996.

- [112] Ivan A. Troyan, Dmitrii V. Semenov, Alexander G. Kvashnin, Andrey V. Sadakov, Oleg A. Sobolevskiy, Vladimir M. Pudalov, Anna G. Ivanova, Vitali B. Prakapenka, Eran Greenberg, Alexander G. Gavriliuk, Igor S. Lyubutin, Viktor V. Struzhkin, Aitor Bergara, Ion Errea, Raffaello Bianco, Matteo Calandra, Francesco Mauri, Lorenzo Monacelli, Ryosuke Akashi, and Artem R. Oganov. Anomalous high-temperature superconductivity in yb6. *Advanced Materials*, 33(15):2006832, 2021.
- [113] Michael S. Turner. Axions from sn1987a. *Phys. Rev. Lett.*, 60:1797–1800, May 1988.
- [114] Michael S. Turner. Periodic signatures for the detection of cosmic axions. *Phys. Rev. D*, 42:3572–3575, Nov 1990.
- [115] N. Viaux, M. Catelan, P. B. Stetson, G. G. Raffelt, J. Redondo, A. A. R. Valcarce, and A. Weiss. Neutrino and axion bounds from the globular cluster m5 (ngc 5904). *Phys. Rev. Lett.*, 111:231301, Dec 2013.
- [116] Steven Weinberg. A new light boson? *Phys. Rev. Lett.*, 40:223–226, Jan 1978.
- [117] Frederick S. Wells, Alexey V. Pan, X. Renshaw Wang, Sergey A. Fedoseev, and Hans Hilgenkamp. Analysis of low-field isotropic vortex glass containing vortex groups in yba2cu3o7x thin films visualized by scanning squid microscopy. *Scientific Reports*, 5(1):8677, Mar 2015.
- [118] S. Wiegers, Peter Christianen, H. Engelkamp, Andries den Ouden, Jos A.A.J. Perenboom, U. Zeitler, and J. Maan. The high field magnet laboratory at radboud university nijmegen. *Journal of Low Temperature Physics*, 159:389–393, 04 2009.
- [119] F. Wilczek. Problem of strong p and t invariance in the presence of instantons. *Phys. Rev. Lett.*, 40:279–282, Jan 1978.
- [120] T. L. Wilson. *Techniques of radio astronomy*, 2011.
- [121] Wenura K Withanage, Andre Juliao, and Lance D Cooley. Rapid nb3sn film growth by sputtering nb on hot bronze. *Superconductor Science and Technology*, 34(6):06LT01, may 2021.
- [122] Dojun Youm, HeeSu Byun, Hyunkyung Kim, Jinsu Kim, Jiwon Lee, Ohjoon Kwon, Seongtae Park, Woohyun Chung, Andrei Matlashov, Boris Ivonov, Caglar Kutlu, Jinmyeong Kim, Seonjeong Oh, Sergey Uchaikin, Yannis K. Semertzidis, Arajana F. Van Loo, and Yasunobu Nakamura. High temperature superconducting cavities : Current progress and future plans for axion searches at capp. 18th Patras Workshop, 2023.

- [123] D.B. Yu. *An Improved RF Cavity Search for halo Axions*. PhD thesis, Massachusetts Institute of Technology, 2004.
- [124] A. R. Zhitnitsky. On Possible Suppression of the Axion Hadron Interactions. (In Russian). *Sov. J. Nucl. Phys.*, 31:260, 1980.

VITA

Tom Braine is now a *Postdoctoral Researcher* located at *Pacific Northwest National Laboratory (PNNL)* in Richland, WA as of *July 2024*. He welcomes any comments or inquiries at `tbraine@uw.edu`.