U.S. DEPARTMENT OF ENERGY

PNNL-30639

Prepared for the U.S. Department of Energy
under Contract DE-AC05-76RL01830

Axion Dark Matter Experiment Cavity Locking; System Development and Testing

Christian R. Boutan
Matthew S. Taubman
Noah S. Oblath

April 2018

Pacific Northwest
NATIONAL LABORATORY

Proudly Operated by Battelle *Since 1965*

DISCLAIMER

PACIFIC NORTHWEST NATIONAL LABORATORY
operated by
BATTELLE
for the
UNITED STATES DEPARTMENT OF ENERGY
under Contract DE-AC05-76RL01830

Printed in the United States of America

Available to DOE and DOE contractors from
the Office of Scientific and Technical Information,
P.O. Box 62, Oak Ridge, TN 37831-0062
www.osti.gov
ph: (865) 576-8401
fox: (865) 576-5728
email: reports@osti.gov

Available to the public from the National Technical Information Service
5301 Shawnee Rd., Alexandria, VA 22312
ph: (800) 553-NTIS (6847)
or (703) 605-6000
email: info@ntis.gov
Online ordering: http://www.ntis.gov

PNNL-30639

Axion Dark Matter Experiment Cavity Locking; System Development and Testing

Christian R. Boutan
Matthew S. Taubman*
Noah S. Oblath

April 2018

Prepared for the ADMX Collaboration
under U.S. Department of Energy Contract DE-AC05-76RL01830

Pacific Northwest National Laboratory
Richland, Washington 99352

*Corresponding author: Matthew.Taubman@pnnl.gov

Executive Summary

The current configuration of the Axion Dark Matter eXperiment (ADMX) uses a single cavity to probe for axions in the mass-energy range of 2.4μeV – 3.7μeV, with corresponding frequency range 580 – 890 MHz. Subsequent phases will search for 6.2μeV – 40μeV (1.5 – 10 GHz) axions, requiring microwave cavities of smaller volume to interact at these higher frequencies. To compensate for the reduction in deposited axion power in a smaller cavity, the plan is to use multiple cavities that tune together in frequency. This strategy trades a small volume (insufficient sensitivity) for the added complexity of frequency-locking multiple cavities at each axion detection frequency step. The ADMX Collaboration has developed two locking systems for this purpose: the Sweep Locking system, based on a swept transmission measurement, and the Pound Locking system, which is the microwave predecessor to the well-known Pound-Drever-Hall technique for laser locking or stabilization. This report covers the requirements and technical challenges facing both of these techniques, a summary of tests performed at the University of Florida in order to compare them, and the final outcome and decision based on these tests and subsequent analysis.

The requirements for locking system performance are:

1. Locking of all cavities will be performed in under 10 seconds for each frequency step;
2. The accuracy of the locking will be within 10% of the full-width-half-maximum (FWHM) of the microwave cavity line width;
3. The locking process will produce less than 10 mW on average during the locking process.

Comparison tests between the two techniques were conducted at the University of Florida in late 2017/early 2018 under semi-realistic conditions. The results of the comparison tests showed that:

1. The Sweep Locking technique showed an accuracy of 2.5% of the cavity FWHM, while the Pound Locking technique showed an accuracy of 6.45% of the cavity FWHM. Both are well within the accuracy requirement, although there were some disturbing systematic variations in the lock frequencies achieved by the Pound Locking system that the authors feel would need explaining or mitigating.
2. The Sweep Locking system locked each cavity in about five seconds, indicating that it could lock four cavities in about 20 seconds in series. The Pound Locking system locked each cavity in about two seconds, indicating that it could lock four cavities in eight seconds in series.
3. Both techniques easily satisfied the heating criterion, being at 1.14 mW and 1.67 mW for the Sweep Locking and Pound Locking techniques respectively.

Two additional factors were brought into consideration after the test results were analyzed.

Firstly, the authors discovered an issue with the detection process in the locking systems involving power levels, which would cause massive signal-to-noise ratio (SNR) reductions in the locking process if not addressed. In a realistic scenario where experimental constraints fix the cavity locking SNR to be small, the detection diodes necessary for either Sweep Locking or Pound Locking would cause both systems to be saturated by noise. The authors found a solution to this issue for the Sweep Locking technique, involving a mixing stage and noise filtering that would improve the sensitivity by a factor of 30 dB. It is not clear to the authors how this problem could be solved in the Pound Locking case. Secondly, it was realized after the completion of the analysis from the comparison tests that the locking time of the Sweep Locking technique would be considerably reduced and the accuracy improved by replacing certain pieces of the system, including the above-mentioned mixing stage, with a vector network analyzer (VNA); the potential speed increase for the sweep process for the overall locking process is more than a factor of two.

Taking into account all of these considerations, authors select the VNA version of the Sweep Locking technique for the final cavity locking design. It produces the most accurate locks, it is best equipped to deal with low SNR situations, and the use of a VNA would satisfy the speed requirement. A VNA measurement was the standard by which both techniques were judged during the comparison tests. For the final locking system, by using a VNA to make the mode frequency measurements knowledge about the exact cavity frequencies will be available in the data analysis. Finally, the use of a VNA as part of the Sweep Locking feedback loop fits well with established collaboration practices of using VNAs as an integral part of ADMX operations.

Acronyms and Definitions

Q	Quality factor of a microwave or other cavity, or resonant electronic circuit
ADMX	Axion Dark Matter eXperiment
CENPA	Center for Experimental Nuclear Physics and Astrophysics
DFSZ	Dine-Fischler- Srednicki-Zhitnitsky axion model
FWHM	Full-Width-Half-Maximum
JPA	Josephson Parametric Amplifier
KSVZ	Kim-Shifman-Vainshtein-Zakharov axion model
PDH	Pound-Drever-Hall
PNNL	Pacific Northwest National Laboratory
RAM	Residual Amplitude Modulation
RF	Radio Frequency
SNR	Signal to Noise Ratio
TE	Transverse Electrical
TM	Transverse Magnetic
UF	University of Florida
VNA	Vector Network Analyzer

Contents

List of Figures

List of Tables

1 Introduction

The Axion is a well motivated hypothetical elementary particle. It's existence is a natural consequence of the strong CP problem of QCD being explained by the spontaneous breaking of a Peccei-Quinn symmetry. [1] Axions [2, 3, 4, 5, 6, 7, 8] are also excellent dark matter candidates because they are cold, stable, non-baryonic and can be produced in sufficient abundance to account for some or all of the missing matter energy density of the universe. The world's most sensitive search for these Dark-Matter Axions is being performed by the Axion Dark Matter eXperiment (ADMX) at the Center for Experimental Nuclear Physics and Astrophysics (CENPA) at the University of Washington in Seattle, Washington. [9, 10, 11, 12, 13, 14] Exploiting the inverse Primakoff effect where $a \rightarrow \gamma^*\gamma$ in a magnetic field, the experiment utilizes a high-Q, tunable resonant cavity, submerged in a large magnetic field and looks for the resonant conversion of axions into microwave photons. [8] The signature of an axion would appear as an exceptionally tiny excess of power in the cavity given by:

$$P_a = 8.83 \times 10^{-26}\,\mathrm{W} \left(\frac{g_\gamma}{0.36}\right)^2 \frac{\rho_a}{0.45\,\mathrm{GeV/cm^3}} \left(\frac{B}{7.5\mathrm{T}}\right)^2 \frac{V}{1\mathrm{L}} \cdot \frac{f}{1\,\mathrm{GHz}} \cdot \frac{Q_L}{10,000} \cdot \frac{C_{nlm}}{0.5} \tag{1.1}$$

where g_γ is a model dependent constant (-0.97 for KSVZ,*[4, 5] and $+0.36$ for DFSZ†[6, 7]), B is the magnetic field strength pointing along the z-axis of the cavity, V is the volume of the cavity, f is the resonant frequency of the mode, Q_L is the loaded quality factor of the resonant mode and C_{nlm} is a resonant mode dependent axion-photon conversion efficiency. For typical operating parameters, Equation 1.1 shows that the experiment must be sensitive to yocto-watt power fluctuations for an axion detection and illustrates the importance of maximizing experimental parameters such as B, V, etc.

Currently, ADMX is sensitive enough to detect or rule out weakly coupled DFSZ axions in the 460 MHz – 1.5 GHz (1.9 – 6.2 μeV). If the axion is not found within this range, ADMX will develop new techniques to search for higher mass axions. A current example of a higher mass search is the ADMX Sidecar experiment. [15] It uses a small, 0.38-liter cavity that performs an axion search roughly between 4 and 6 GHz. Piezo motors are used to rotate a tuning rod thereby tuning the cavity resonant frequency and adjust the antenna depth to maintain a critical coupling. The downside to scaling the cavity size down as a way of scaling up the frequency reach of an axion search is that the signal to noise ratio (SNR) suffers as a result of the reduction in cavity volume, to which the power is proportional as shown in Eqn. 1.1.

ADMX plans to search for 6.2 μeV – 40 μeV (1.5 – 10 GHz) axions using an array of multiple Sidecar-like cavities. With all cavities tuned to the same frequency, this strategy trades a small

*Kim-Shifman-Vainshtein-Zakharov axion model
†Dine-Fischler-Srednicki-Zhitnitsky axion model

volume (insufficient sensitivity) for the added complexity of frequency-locking multiple cavities each time the desired resonant frequency is stepped. It is this technical challenge that is the focus of this report. This report explains the requirements of such a cavity locking system, compares available techniques and describes our recommendation for the path forward.

2 Cavity Tuning

Here we review the basics of single cavity tuning and motivate the need for a technique that will quickly tune multiple ADMX cavities to the same frequency.

2.1 Single Cavity Axion Search

Currently, ADMX searches for the resonant conversion of axions to photon in a single resonant cavity. During data taking, over the course of several months, the cavity frequency is iteratively tuned over and over again and digitized spectra are taken at each frequency. During this time, experiment monitors the cavity mode structure, tracks the relevant resonant mode and measures the antenna coupling. Here we review these techniques.

2.1.1 Mode Form Factor

A bare resonant cavity supports a host of transverse magnetic (TM) and transverse electrical (TE) resonant modes,

$$f_{mnp}^{\mathrm{TM(TE)}} \equiv \frac{1}{2\pi\sqrt{\mu\epsilon}}\sqrt{\frac{x_{mn}^{(\prime)2}}{R^2}+\frac{p^2\pi^2}{d^2}} \tag{2.1}$$

where f is the frequency of the TM (or TE) mode, d is the cavity length, p is an integer and $x_{mn}^{(\prime)}$ stands for either x_{mn} or $x_{mn}^{\prime}$, being the nth roots of the Bessel functions $J_m(x)$ and derivatives $J_m^{\prime}(x)$ respectively. Keeping in mind that the axion to photon production rate is proportional to $\mathbf{E}\cdot\mathbf{B}$, a resonant cavity in externally applied magnetic field can be used to convert dark matter axions in the milky-way halo into excess energy stored in an electromagnetic mode. The efficiency of this conversion is given by the "form factor"

$$C_{nlm} \equiv \frac{\left(\int_V dV\,\mathbf{E}(\mathbf{x},t)\cdot\mathbf{B}(\mathbf{x})\right)^2}{VB^2\int_V dV\,\epsilon_r E^2} \tag{2.2}$$

where $\mathbf{E}(\mathbf{x},t)$ is the electric field of the nlm cavity mode, V is the cavity volume, ϵ_r is the relative permittivity within the cavity, and $\mathbf{B}$ is the externally applied magnetic field. If the magnetic field is applied parallel to the length axis of the cavity, only TM_{0n0} modes have a

non-zero form factor and can be used for searching for axions. The TM_{010} mode has the highest efficiency and is the only mode of interest in this report. Unless otherwise specified, "cavity frequency" refers the frequency of the TM_{010} mode.

2.1.2 Mode Tuning

Given the narrow bandwidth of a high-Q resonant cavity, the resonator must be made tunable in order to search for a wide range of axion masses. By inserting a movable metal or dielectric rod into the cavity, the electromagnetic boundary conditions of the cavity can be altered, thereby shifting the frequency of the TM_{010} mode. However, with the introduction a tuning rod (or rods) all symmetries needed to calculate an analytical solution for the form factor are broken and simulations must be performed to verify the axion-photon conversion efficiency. Furthermore, the mode structure of the cavity becomes more complicated and small perturbations to the cavity geometry can cause significant affects on the TM_{010} mode and how it interacts and mixes with other resonant modes in the cavity.

2.1.3 Mode Tracking

The art of ADMX data taking lies in the ability to track the TM_{010} mode as it tunes through and mixes with other interfering resonant modes. For a given tuning rod configuration, the cavity mode structure is characterized by an S_{21} transmission measurement by a vector network analyzer (VNA) . A swept signal is directed into a weakly coupled antenna on the cavity. Power that is on resonance with a resonant mode is absorbed into the cavity and read out by the "major port" which is another antenna that is typically maintained in critically coupled or slightly over-coupled state. The shape of the transfer function is a Lorentzian

$$h(f) = \frac{1}{1 + 4(f - f_0)^2/\Gamma^2} \tag{2.3}$$

where f is the swept frequency of the network analyzer, f_0 is the center frequency of the resonant mode and Γ is the full width half maximum (FWHM) of the response. The left panel of Figure 2.1 shows the S_{21} transfer function of the AMDX Sidecar TM_{010} mode for 70 slightly different tuning rod positions. At each tuning rod configuration, the experiment measures the TM_{010} frequency, uses that information to set the radio frequency (RF) receiver chain to mix that particular frequency down to 10.7 MHz and takes a digitized spectrum. The right panel of Figure 2.1 shows a map of the peak resonant frequencies of the cavity versus tuning rod position. In this mode map, it can be seen that the TM_{010} tunes over a large 2-GHz bandwidth but is periodically interrupted by a mode crossing with interfering TE modes. These problematic interactions with other modes creates regions where no data can

be taken. Taking data near a mode crossing can be difficult because the TM_{010} begins to take on characteristics of the TE mode, the slope of frequency versus tuning rod changes and the uncertainty in the calculated form factor increases.

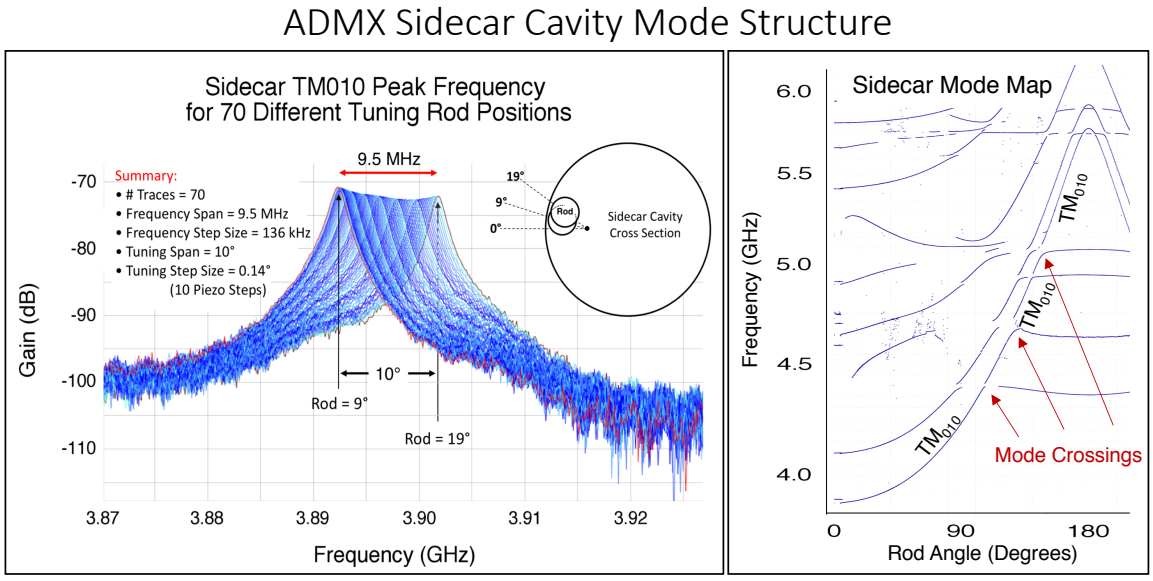

Figure 2.1: This figures shows the tuning of the sidecar experiment in the left panel, and sidecar cavity mode map in the right panel. 70 different Lorentzian resonance features are shown according to angular rod position in the cavity. Note the change in peak amplitude with frequency, which becomes important in section 4.1.

The "quality factor" (Q) of the mode can either be written in terms of f_0 and Γ, or of the geometry and skin depth (δ) of the cavity:

$$Q = \frac{\lambda}{\delta} \cdot \frac{(x_{mn}^2 + (p\pi r/d)^2)^{1/2}}{2\pi(1 + 2R/d)}, \tag{2.4}$$

where λ is the wavelength of the cavity mode in free space. The higher the Q, the more axion power that develops in the cavity (Eqn. 1.1) and the narrower the instantaneous bandwidth of the experiment.

2.1.4 Antenna Coupling

During data taking, the antenna of the major port is typically maintained at a critical coupling. The antenna coupling is measured with an S_{11} reflection measurement from a network analyzer. A swept signal is directed through a circulator and towards the antenna. Power that is off resonance is reflected off the antenna. If the antenna is critically coupled,

power that is on resonance will be fully absorbed into the cavity. This appears as a dip in the network analyzer measurement. The antenna depth is then adjusted until this absorption dip is minimized which is typically ∼30 dB below the reflected signal baseline.

2.2 Multi-Cavity Axion Search

With the introduction of multiple cavities, all of the operations detailed above need to be performed for each of these cavities. These additional measurements may have an impact on the duty cycle of the experiment. Depending on how the cavities are used together, there are other complications to be considered. During the data taking process, ideally an axion search tunes predictably and smoothly to achieve a SNR that is fairly flat in frequency. In an experiment that only uses one single microwave cavity, there is no need to tune it to particular precise frequencies, so long as its resonant frequency can be accurately measured. In a multi-cavity experiment, if the search is to operate with each cavity independently searching N different frequency regimes, then this statement is still true to a certain extent.* However, for a frequency-locked multi-cavity search, these assumptions are not true. We have to be able to move all the cavities to a specific frequency, and then be able to know to a high confidence level that they are, and will remain thus tuned for the duration of that measurement step.

2.2.1 The Frequency Comb vs the Frequency Lock

One potential strategy for a multi-cavity search would be to assign each of the N cavities a separate frequency region to scan, creating an N-toothed comb in frequency. While this approach would create N separate axion searches, the scan rate of the collective experiment would be reduced by $1/\sqrt{N}$ compared to the case where all cavities are tuned to the same frequency. This is because the the signal to noise (SNR) is given by the radiometer equation

$$\mathrm{SNR} = \frac{P_a(V/N)}{P_N}\sqrt{B_a t_{\mathrm{int}}} \tag{2.5}$$

where P_a is volume dependent axion power, P_N is the noise power, B_a is the axion bandwidth and t_{int} is the integration time. By solving for the integration time and noting that the total time would be reduced by the number of N axion searches, the scan rate for the multi-cavity comb can be found to be

$$\mathrm{Rate} \propto \frac{1}{t_{\mathrm{int}}/(N \text{ independent searches})} \propto \frac{1}{\sqrt{N}} \tag{2.6}$$

*It might not be prudent to have two such cavities search the same regime or overlap in frequency by chance.

as opposed to the N-independent scan rate achieved by tuning all cavities to the same frequency. The multi-cavity comb would also be difficult to scale since it would require N independent RF receiver chains. The downside to running an axion search with N frequency locked cavities is the tradeoff in the loss in sensitivity for an increase in risk and complexity. For a four-cavity array, the factor of two increase in scan rate compared to the frequency comb approach may be lost in the unintended down time associated with complicated experiments. On top of that, the need to carefully tune cavities to exact frequencies will take longer as each cavity over and undershoots the intended frequency in a fine-tuning feedback loop.

2.2.2 ADMX Run 2A

In 2020, ADMX run 2A will proceed with a frequency-locked four-cavity array (also referred to as a four-cavity system). The cavity system itself will be fabricated by the University of Florida (UF), and have a target frequency range of 1.5 to 2 GHz. This cavity system is shown in Figure 2.2, both CAD design and completed prototype, side by side. Linked coarse tuning rods will tune all cavities with a single motion. Small fine-tuning rods will be used to alter cavity mode frequencies slightly by being inserted or removed from the cavities using independent linear piezo motors. Two locking techniques are being considered for this four-cavity system, both of which aim to use these fine-tuning rods to frequency-lock the cavities together. One of these techniques is a computer-mediated search (Swept Locking), using cavity transmission (S_{21}), developed and advocated by The Pacific Northwest National Laboratory (PNNL). The other is the more traditional Pound technique [16], which is based on cavity reflection (S_{11}), being demonstrated and advocated by UF. Regardless which of these techniques is chosen, PNNL will be responsible for building and commissioning the final cavity locking system. This chapter discusses the requirements of the four-cavity system, which will be used as a basis upon which to evaluate and ultimately choose between the two proposed locking systems.

2.3 Fine-Tuning and Heating

It is expected that run 2A of ADMX will use Attocube motors in the baseline of a four-cavity array search. Attocube motors have four sources of heating: dissipated power from the exercising of the actuator, mechanical heating from the stick-slip motion, ohmic heating from leakage currents through the element itself and encoder heating. [17]

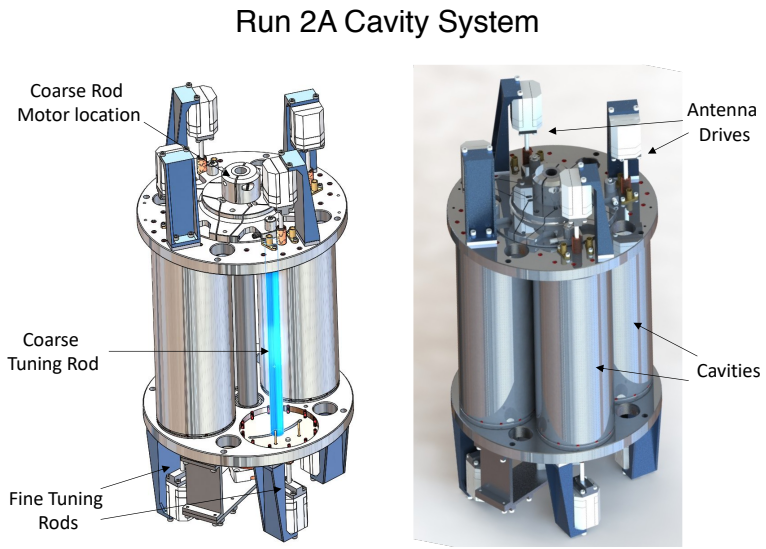

Figure 2.2: A CAD drawing of the ADMX run 2A cavity array prototype designed and built at the University of Florida

2.3.1 Dissipated Power

The power dissipated by the piezo motor is

$$P_{\text{dissipated}} = CV^2 \tan(\delta) f \tag{2.7}$$

where C is the capacitance of the actuator, V is the voltage applied to it, $\tan(\delta)$ is the loss tangent and f is the frequency of the sawtooth voltage applied to the motor. Since this scenario doesn't involve running the motor constantly over long periods of time, here it makes more sense to think in terms of the amount of energy deposited per "step"

$$E_{\text{dissipated}} = CV^2 \tan(\delta) \tag{2.8}$$

Here we assume a fixed voltage of 60V which was the working voltage used by the the ADMX Sidecar experiment. The other parameters in Eqn. 2.8 are summarized in Table 2.1.

Table 2.1: This table compares the capacitance, loss tangent and step size for room temperature and cryogenic operation.

Temperature	C	tan(δ)	Step Size
300K	2 μF	0.02 - 0.04	2 μm*
$<$ 4K	200 nF	slightly less than room temp val†	$\sim$1/10 room temp val‡

Table 2.1 reveals that while the amount of expected heating decreases by slightly more than a factor of 10 when cooled below 4 K, the step size also decreases by a factor of 10, requiring 10 times the number of steps to achieve the same frequency step size. This suggests that estimates of piezo heating at room temperature should be adequate for predicting the amount of heating at 150 mK temperatures. Since the definition of a "step" changes with temperature, it is more useful to know that the number of Joules per mm is

$$\begin{aligned} E_{\text{dissipated}}/\text{mm} &= CV^2 \tan(\delta)(\text{steps per mm}) \\ &= (2\mu\text{F})(60V)^2(0.03)(500 \text{ steps/mm}) \\ &= 100\,\text{mJ/mm} \end{aligned} \tag{2.9}$$

This is the largest contribution to the heating of the system.

2.3.2 Mechanical Power

There is an additional heating term from the work done by the stick-slip system

$$E_{\text{mechanical}} = F \cdot x \tag{2.10}$$

where F is a fixed clamping force on the moving fixture within the motor and x is the distance by which it travels. Attocube claims that a typical value for the amount of mechanical heating is on the order of 500 nJ over a 100 nm step. This corresponds to

$$E_{\text{mechanical}}/\text{mm} = 5\,\text{mJ/mm} \tag{2.11}$$

This mechanical source of heating is about a factor of 20 less than the dissipated power term.

*Measured by Christian Boutan on 8-11-2017.

†Attocube doesn't have a precise number but email correspondence between Christian Boutan and Attocube engineer Daniel Stroh on 1-14-16 suggests that the number is slightly less than the room temperature value.

‡By measuring the change in the frequency step sizes of the Sidecar experiment from room temperature to below 200mK, it was determined that the piezo step size decreases by about a factor of 10 (cite thesis).

2.3.3 Ohmic Power

This source of heating comes from the fact that the piezo is essentially a capacitor but unlike an ideal capacitor it has some finite resistance on the order of GΩ. This would result in leakage currents producing

$$P_{\text{ohmic}} = V^2/R \sim (60\text{V})^2/1\,\text{G}\Omega \sim 4\,\mu\text{W} \tag{2.12}$$

of heating if a constant voltage were left on the piezo element. While it is an option to leave a static DC voltage directly on the element to achieve nanometer positioning resolution, we would probably chose not to use this option. Since this source of heating is small and avoidable, it will be ignored in this calculation.

2.3.4 Encoder Heating

It was experimentally determined that the piezo encoders deposit quite a bit of heat into the system. ADMX will be fixing this with a switchbox that will allow that encoder cables to be switched in and out. Since this is an issue that will be solved in future versions of the experiment, it will be ignored in this calculation.

2.3.5 Attocube System Summary

Adding up all sources of heating, it is found that every millimeter of motion from a single motor will result in a little over 100 mJ of heat being deposited in the system. This is quite a bit of energy but the motion is not continuous. During regular operations of the experiment, the cavities will be tuned over a 10-second period of time and then the heating will cease for 100 seconds while signal from the cavities are digitized. Therefore, on average power deposited during this time is

$$P_{\text{total}} = 105\,\text{mJ}/110\,\text{s} \sim 1\,\text{mW} \tag{2.13}$$

This is an important conclusion because it has been decided that if the piezo system is thermally tied to the 1K plate, the heat load from the fine-tuning system must not exceed 10 mW. If the above calculations are correct, on average, each of the four piezo fine-tuning rods will be allowed to move 2.5 mm to achieve a 10% line width lock.

2.3.6 Impact of Tuning Rod Performance

Since the travel distance of the fine-tuning rods is fixed by the above argument to be less than 2.5mm per rod, it is important that a rod insertion or removal by this amount results in a frequency shift that is great enough to achieve a 10% line width lock, which is a requirement of system. Unfortunately, there are aspects of the final design that are unknown. For example, it is not known how well the coarse tuning rod system will keep all cavities well grouped in frequency. For the prototype cavity system described in section 5.1, cavities A, B and D remain within 10 MHz of each other (typically less than 7 MHz) but cavity C, which is theoretically identical to the others is sometimes 30 MHz off from the rest of the system. Any frequency misalignment will need to be corrected by the fine-tuning system. If the final coarse tuning mechanism works much better than the prototype and only produces a variance in cavity frequencies that is on the order of a few cavity line widths, the fine-tuning mechanism only needs to be able to tune a few Q widths for day to day operations. In other words, at 2 GHz with a Q of 30,000 to 40,000, at a bare minimum, the fine tuning system needs to be able to tune by 300 kHz to 600 kHz. Early plans for the fine-tuning rod design claimed that it would shift the cavity frequency by $\sim$ 1 MHz when extended fully. If we assume that a full extension of the rod is 12 mm, which is the full travel distance of the linear Attocube motor, then the constraint of only being allowed to move the rod 2.5 mm means that the fine-tuning system won't be able to tune the cavity more than 200 kHz* without overheating the system.

It is the authors opinion that piezoelectric heating and the potential frequency discrepancy between cavities be considered when designing the final tuning rod system. This did not become an issue with the comparison test (section 5) because the cavity prototype is 3 times smaller than the final system and fine-tuning rods where able to shift the cavity frequency by 10-20 MHz within the allowed piezo travel length.

*(2.5 mm/12 mm)(1 MHz) $\sim$ 200 kHz

3 Cavity Locking System Requirements

The four main concerns of the cavity locking system are accuracy, speed, heating and stability. On July 31, 2017, at the ADMX collaboration meeting, representatives from Lawrence Livermore National Laboratory (Nathan Woollett), PNNL (Matt Taubman, Noah Oblath, Christian Boutan), UF (Jihee Yang, Shriram Sadashivajois) and the University of Washington (Richard Ottens) met together and agreed upon the following requirements for the 1 – 2 GHz cavity system, comprising four cavities. (These may be different for later, higher-frequency versions comprising more cavities.)

3.1 Accuracy

So that the signals from all four cavities will add coherently, not only should they all be tuned to the same cavity mode, but they should be within 10% of the mode line width of one another. In other words, the tunings of all four cavities should be within a frequency region no bigger than 10% of the mode line width. Under these conditions, the phases of the cavities signals in response to axion decay should be within a 20° range.

3.2 Speed

In order to be able to complete the data sweep across the cavities' frequency range in a reasonable time frame, it was determined that in each frequency step, the complete locking process should take no more than 10 seconds for all four cavities.

3.3 Heating

One of the problems with the locking process is that the piezo motors inside the cryogenic cavity system generate heat each time they move. The dilution refrigerator that keeps the system at or below 150 mK can only handle a certain amount of heat. The Josephson parametric amplifier (JPA) is thermally tied to this 100 mK space controlled by the dil fridge. It has been determined that more than a 50 μW heat load into this space will unbias the JPA. This is a severe limitation on the piezo motors, making operation essentially unattainable. The solution was to move the heat-sinking of the piezo motors from the 100 mK space to the '1K plate', where

the requirements are much more relaxed. Here, a 10 mW heat load over the locking period of 10 seconds is acceptable. This will not unbias the JPA, or damage the dilution refrigerator.

4 Techniques

Currently, the ADMX cavity is tuned in a manual fashion, where cavity frequency measurements are not used iteratively to guide the tuning mechanism. For the ADMX run 2A, a four-cavity array will be tuned in a feedback loop where knowledge about the mode structure will be used to fine-tune each cavity to the same frequency. Two different techniques contend to be the fine tuning mechanism used by ADMX to frequency lock the array. The first technique uses the response of a swept signal through each of the cavities to identify the frequency of the TM_{010} mode. This technique will be known as Sweep Locking (formerly Hill Climbing) and has been investigated by the PNNL group. The second technique, investigated by UF, uses a phase-modulated carrier signal reflected off the critically-coupled antenna on the cavity to create a signal with the correct polarity and amplitude to drive a servo motor which tunes the cavity to the carrier frequency. This technique is known as Pound Locking, which is the microwave precursor to the laser locking or stabilization technique known as Pound-Drever-Hall (PDH). (See section 4.2 for details and history and relevant citations of Pound Locking.) The key differences between both of these techniques is that one is based on a transmission measurement (S_{21}) and the other is based on a reflection (S_{11}). In this chapter we review both systems and in the following chapter we compare both systems directly. Finally, in chapter 6 we discuss what each of these techniques would look like in a realistic axion search and look for the system that best matches the ADMX experiment.

4.1 Sweep Locking Technique

The Sweep Locking technique evolved from a simple Hill-Climbing technique, envisioned as a simplified transmission-based alternative to the Pound Locking technique. Figure 4.1 tells the story of the progression of the technique in four phases, labeled 0 to 3.

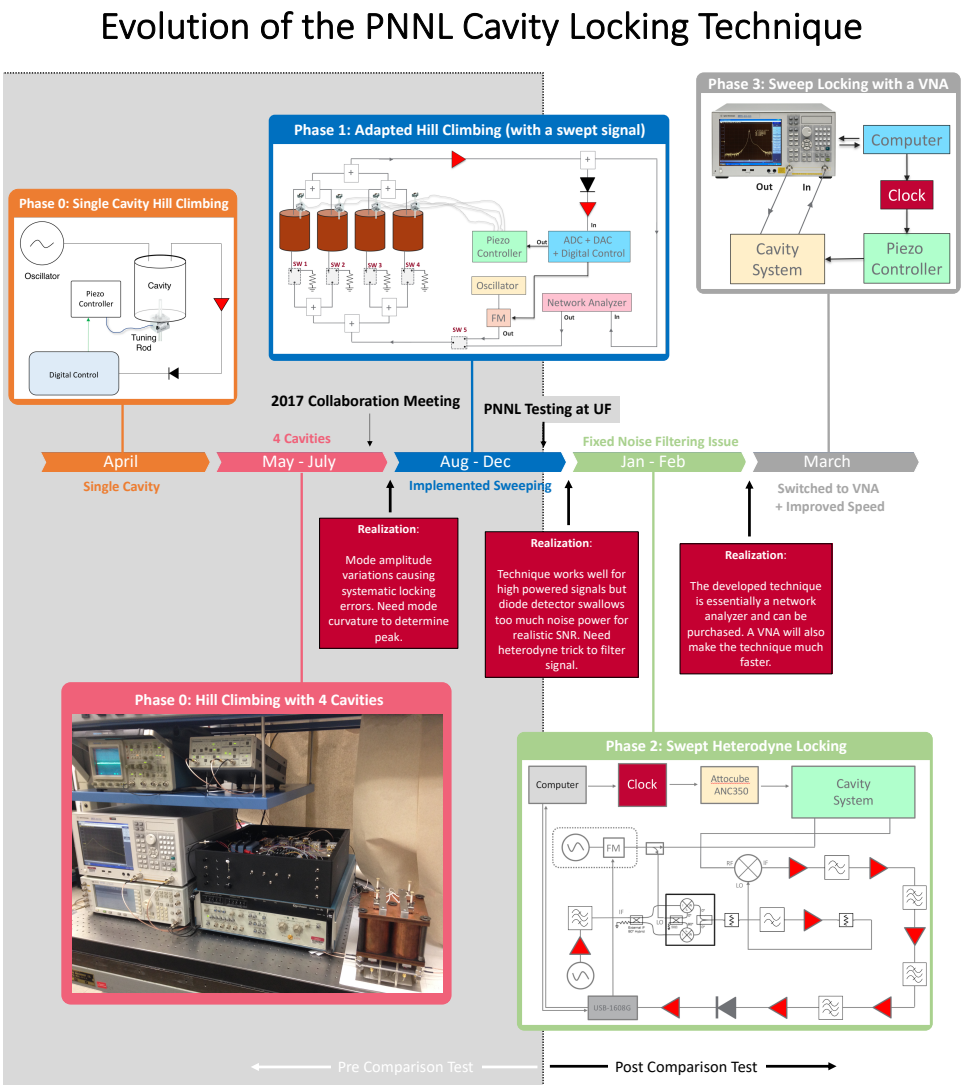

Figure 4.1: This figure shows the evolution of the PNNL Sweep Locking technique in four phases, 0 to 3.

4.1.1 Phase 0 (Hill Climbing)

The hill-climbing technique was the starting point for the present system, and used much of the same equipment. Shown on far left side of Figure 4.1, as single tone from a signal generator was injected into the cavity. It was demonstrated to be a simple, modulation-free algorithmic process, which would move the cavity tuning one way or the other to attain maximum microwave transmission as detected by an RF detector diode. A simple algorithm was written to move the tuning rods of the cavity until its transmission of a given microwave

tone from a signal generator was maximized. It was imagined that progressive tunings of cavities could be simply effected using this technique, because the distance to be moved in any one step would be small, or at least well known, and once a progressive set of tuning steps was established in an experimental run, the direction of tuning was assumed to be constant. It was also established that multiple cavities could be moved step-wise in a cycle, being relatively simple to do in software. This interleaved architecture had the advantage of removing any limitations due to the settling time of the Piezo actuators, which have a time constant of approximately 1 ms. This was envisioned to be particularly important in the penultimate version of the ADMX cavity system, in which 16 individual microwave cavities are required. This was seen as essential in this case, as lower limitations to the tuning time for the 16 cavities tuned in a serial manner may have been detrimental to the run time of the multi-GHz ADMX experiment. As the precision and speed of the hill climbing system was investigated, it was found that the interleaving process was limited by the switching time of the multiplexing unit that was being used. To avoid unnecessary dead time, the cavity locking speed, the interleaving option was deactivated for future testing.

As the technique was being fleshed out, it was found that there was a systematic error in locking precision caused by natural amplitude changes in the response of the resonant cavity mode. A subtle example of this effect can be seen in the slight gain changes of the Sidecar Lorentzian peaks in the left panel of Figure 2.1. As a result, maximizing the detected power as a function of piezo motor position did not yield a frequency coincident that of the peak of the mode, to sufficient accuracy. It was determined that the peak frequency of the mode must be determined not as a function of piezo motor position, but from the Lorentzian shape of the mode as a function of RF drive frequency. As a consequence, the hill climbing technique was modified to use a swept RF input signal.

4.1.2 Phase 1 (Modified Hill Climbing)

In this technique, a system was developed that scanned a signal generator tone through the peak of the microwave cavity. The difference between the peak frequency and target frequency was computed and then the piezo motor driven fine-tuning rod was moved accordingly using an assumption about the localized slope of the peak frequency versus motor step size. This process was then repeated until the difference between the peak and target frequencies was less than a specified tolerance. In this sense the technique became a sweeping technique, not a hill-climbing technique.

In December of 2017, this version of the system was shipped to UF and put through a series of tests (Section 5) conducted by Christian Boutan. The results of these tests showed that the Phase-1 system was very accurate but could probably be improved in terms of cavity locking speed.

4.1.3 Phase 2 (Swept Heterodyne Locking)

After completing the PNNL portion of the comparison testing, there was time to envision how a final locking system would be incorporated into ADMX. It was realized that the amount of allowed power sent through the cavity would be limited by the quantum electronics and that the SNR of the swept locking system might become an issue. It was also realized that the RF detector in the Phase-1 design was being exposed to a very wide bandwidth of unnecessary noise power. This background noise was not an issue in the comparison test since the locking signal could be made arbitrarily large in an isolated system. In the final system, the quantum front end amplifiers begin to compress with an input greater than −120 dBm and the amplitude of the signal injected by the locking system would need to be reduced considerably. In this scenario, the wide bandwidth of system noise being rectified by the detector would completely dwarf the locking signal. It was determined that a heterodyne filtering trick could be used to thread the locking signal through a narrow filter thereby cutting out almost all of the system noise power. This proof of concept was demonstrated and the results are shown in Figure 4.2.

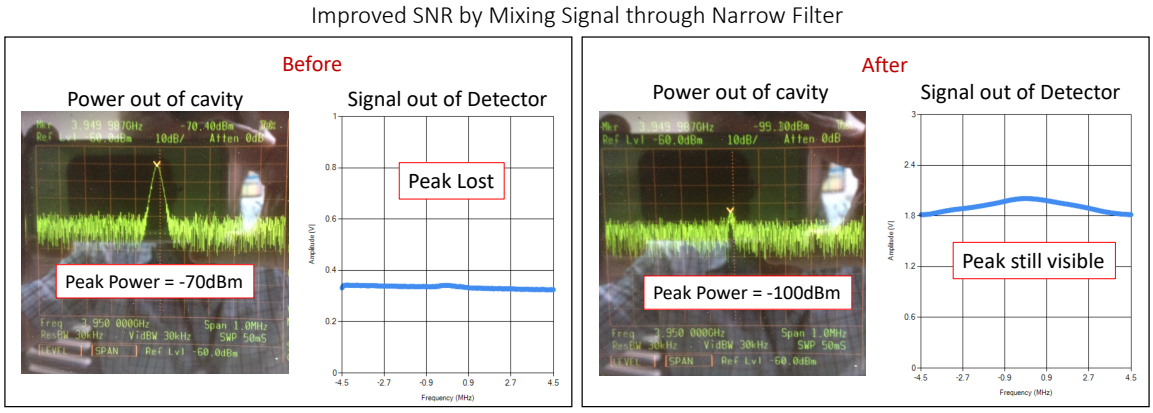

Figure 4.2: This figure shows the improvement of the SNR when the heterodyne technique is employed, which allows broadband noise to be removed with a narrow fixed filter.

This demonstration performed by mixing the swept cavity output down by a fixed local oscillator. While easy to implement for an in initial test, it limits the narrowness of the band pass filter to be the width of the swept range of several MHz. Furthermore, sweeping across the width of the filter imparts the filter shape to the Lorentzian response. Incorporating this heterodyne trick into the final system would involve the locking signal and the local oscillator signal sweep together in unison. The design for this (Phase-2) locking system is shown in the bottom right corner of Figure 4.1. This design utilizes a hybrid and a single sideband mixer to take a fraction of the locking signal and create another swept signal that is 70 MHz less than the initial signal. The resulting signal returning from the cavity system is mixed down to 70 MHz and through a narrow crystal filter. This technique essentially adapted the Phase-1 locking system for low-SNR scenarios.

4.1.4 Phase 3 (VNA Sweep Locking)

It was then realized that the design for the sweeping mechanism could be performed by a vector network analyzer. A speed tests was performed and it was discovered that a VNA could increase the measurement speed by an order of magnitude (acquiring a trace in 50 ms instead of 0.55 sec), improving the locking time of the technique by more than a factor of two. The VNA was also found to have a much higher frequency resolution. Replacing the preexisting Hill Climbing sweeping mechanism with a VNA lead to Phase-3 of the Sweep Locking technique.

4.2 Pound Locking Technique

4.2.1 Motivation, advantages and disadvantages

Pound Locking is a technique with considerable history behind it. Robert Vivian Pound was an American physicist who helped discover nuclear magnetic resonance (NMR), and who devised the famous Pound-Rebka experiment supporting general relativity [18]. In 1946, Pound published a paper [16] describing the frequency control of a microwave oscillator by an external high-Q microwave cavity, using an intermediate frequency facilitating better electronic control, to provide, higher resolution, more stable tools for the investigation of microwave absorption spectra. The same drive for finer, more stable tools has existed in the optical regime since the invention of optical spectroscopy. Since shortly after the invention of the laser to the present day, researchers sought to apply frequency stabilization to these new tools. Barger, Sorem and Hall achieved the first optical cavity stabilization of a dye laser in 1973 [19], using a technique locking to the side of the cavity fringe. This technique and its refinements [20, 21] became the basis for stabilization of commercial dye lasers of the day. The laser stabilization effort was supercharged by the desire to capture a laser beam between cavity mirrors separated by kilometers inside gravitational wave interferometers (feasibility studies in the early 1970s), in order to enable the detection of the murmurings of far-off massive objects [22, 23], and which have recently succeeded. [24] This led to the adoption of Pound's sideband modulation technique to that of optical cavity locking, and the now famous and standard Pound-Drever-Hall (PDH) stabilization technique. [25, 26, 27]

The principal advantage of Pound Locking for microwave systems, and indeed Pound-Drever-Hall locking for optical systems, is that the control bandwidth of the feedback servo system is not limited by the storage time of the cavity involved. This is because it is a heterodyne technique, responsive to phase shifts between the reflected carrier and reflected modulation sidebands. Relative movements between the drive frequency and that of the optical cavity cause such phase shifts immediately, without having to wait for the much slower cavity response. It is these phase shifts, which, after demodulation, produce the dispersion shaped error signal typical of reflected sideband locking systems. The response of the error signal to detuning

variations is sensitive to phase at response frequencies above the line width of the cavity because at these speeds the stored mode of the cavity remains relatively constant. Deviations at frequencies below the cavity line width produce signals responsive to frequency offsets, because the cavity mode has time to react to the drive frequency. The natural pole due to the distinction between these two regimes adds to the stability of the required servo systems. By comparison, techniques that rely on cavity transmission, such as the dither technique whereby slower larger frequency modulations are applied, their amplitude effects on the cavity mode being detected in transmission, are necessarily limited by the storage time of the optical cavity. Servo loops in such systems must accommodate this limitation. Clearly, this problem becomes worse as cavity Q values climb. This is a problem, since the required locking power, or servo gain is greater for narrower cavity resonances — it requires more intense control of the source to maintain lock to a narrower target. Transmission systems often fail in these cases, especially where the source has fast, or broad-band noise, as slow servo systems simply do not have the capability to move the noise power into a narrow band to couple with a narrow cavity. In short, the Pound and PDH techniques are the go-to techniques for locking involving high-Q cavity systems where the source may have considerable noise or broadband fluctuations to be tamed.

Another advantage of Pound and PDH locking is that the capture range of the lock is not limited by cavity line width or storage time. This is because the phase shifts between reflected carrier and sidebands that produce the error signal are persistent well outside the cavity line width, sustaining a non-zero error signal, often of significant amplitude for optimum ratios of modulation frequency and cavity line width. The appropriate phase of the error signal for locking is maintained within the frequency range defined by the first phase modulation sidebands, in other words, carrier frequency plus and minus modulation frequency. This, again, is a tremendous advantage for locking high-Q cavities quickly, reliably, and maintaining a robust system against unlocking conditions.

Still a further advantage of these phase-sensitive sideband locking techniques, and indeed one that Pound himself highlighted in his original paper [16], is that the technique is, to first order, immune to low frequency amplitude fluctuations of the source. This is because all involved sidebands are affected equally by such fluctuation. While the magnitude of the fast (phase regime) error signal my fluctuate as a result, or become noisy, this noise is only proportional to the error signal amplitude, and disappears at the zero crossing. Hence, the lock point of these techniques remains unaffected by amplitude noise of the source.

Disadvantages of Pound and PDH locking systems are relatively few, but they can be problematic. Depending on the system, these techniques can add significant complexity, which may or may not be tolerable, or practical. The necessity of producing a clean phase modulation can indeed be limiting for some systems, such as in the long-wave infrared optical regime, where good phase modulators are difficult to make. The use of laser drive current, or indeed source voltage for microwave systems, as a measure of moving the source frequency (or phase) is resorted to. However, this can have the deleterious effect of producing amplitude modulation at the phase modulation frequency. This *does* negatively affect the zero-crossing point of the error signal, because this residual amplitude modulation (RAM) produces

undesirable phase shifts in the modulation sidebands in a manner that always produces a signal upon detection, or an offset. To make matters worse, this offset *is* subject to low frequency power fluctuations of the source, and now *does* affect the frequency of the zero crossing point.

4.2.2 Operation

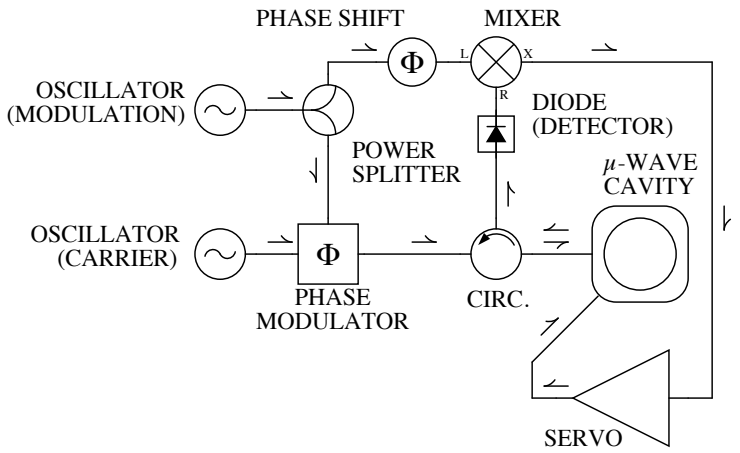

Figure 4.3: A schematic of a Pound-Drever-Hall cavity locking system.

Figure 4.3 shows a schematic of Pound Locking for microwave cavities, arranged in this case to lock the mode frequency of a microwave cavity to that of a fixed microwave oscillator,* sub-labeled 'carrier,' on the left of the figure. The output of this oscillator passes through a phase modulator, which imposes phase modulation sidebands on the carrier at some small fraction of the carrier frequency. The resulting phase-modulated carrier passes through a circulator, and is coupled into the microwave cavity. When the cavity is resonant with the carrier, it couples energy to the cavity, while the modulation sidebands are reflected. As mentioned previously, the technique operates by taking a reflection signal from the microwave cavity, and rectifying the RF (via the diode, or detector), in order to see the amplitude envelope. Phase modulation applied to the microwave source (derived from a second oscillator sub-labeled 'modulation') produces the phase modulation sidebands which are directly reflected by the microwave cavity, provided they are well outside the cavity line width, and away from other cavity resonances. The reflected power from both carrier and sidebands is diverted by the circulator through the diode (detector) and into a mixer, lock-in amplifier, or some other demodulation device, which produces a baseband signal. After some filtering (not shown) a DC signal is obtained, and passed to a servo. This DC signal is the Pound error signal, and has typical shapes similar to those shown in Figure 4.4. The width of

*...as distinct from moving the oscillator frequency to match that of the cavity.

the central features is that of the resonance of the microwave cavity; the separation between the central and side features of the signal is equal to the modulation frequency.

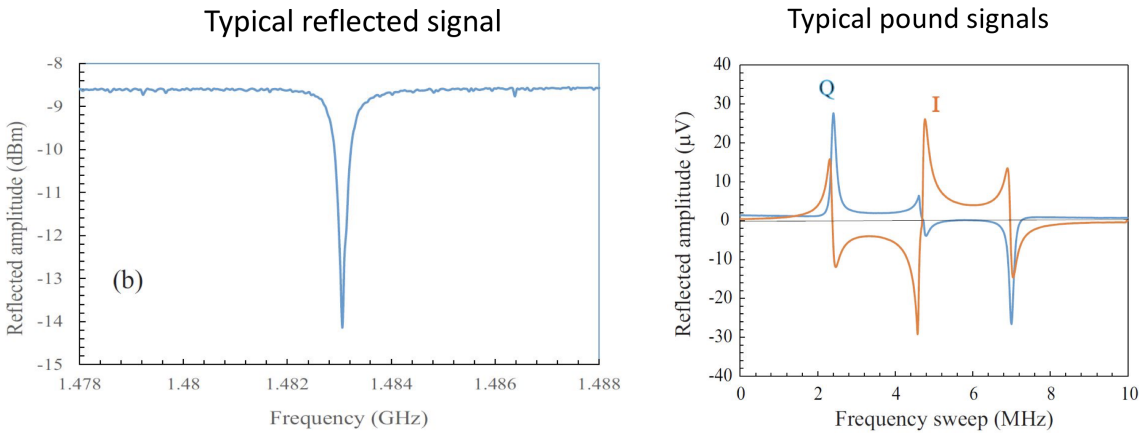

Figure 4.4: Left panel shows swept power reflected off of the cavity antenna. Right panel shows the I and Q quadrature Pound Locking signals.

5 Comparison Test

The purpose of the comparison test was to put both systems through the same set of semi-realistic conditions and see how each performed in the following categories: accuracy, speed, heating, stability. Comparison tests were carried out in Neil Sullivan's Laboratory at UF using their prototype cavity system.

5.1 Cavity System

Testing was performed with the four-cavity array shown in Figure 5.1. The prototype system was designed and built by Joseph Gleason at UF as a testbed for the first multi-cavity array for the ADMX experiment. The TM_{010} mode of the prototype tunes from 4.4 to 6.3 GHz compared to the 1-2 GHz target frequency range of the final system which will be roughly 3 times bigger. Each cavity of the array contains a large coarse tuning rod which shifts the resonant modes. The bottom of Figure 5.1 shows a mode map of cavity A measured by Jihee Yang. The map shows the TM_{010} mode (the lowest frequency tuning mode) along with other cavity modes versus tuning rod position.

Coarse tuning rods are mechanically connected to each other and are tuned simultaneously with a single rotary piezo motor. As the system tunes, tiny differences between cavity geometries become important and cavity frequencies which are required to remain within $f/(10Q)$ (Section 3) during operations drift apart. To counteract this frequency drift, each cavity is equipped with a smaller (metal or dielectric) fine tuning rod. Fine tuning is achieved with individual linear piezo actuators that insert or retract these rods through one of each cavities' end caps. It is this fine-tuning system that works in conjunction with and is a part of the planned cavity locking system. Figure 5.2 illustrates the importance of a fine tuning, cavity locking system. Error bars (present in the figure) are set by the $f/(10Q)$ of each cavity and the fact that they are typically smaller than the point size of the cavity frequency illustrates how far the unlocked cavities (top) drift out of specification. The bottom of the figure shows that effect of having the locking system activated is dramatic and that, in most cases, cavities sit well within the error tolerance away from the central target frequency. Only cavities A, B and D were used for the comparison testing. The mode structure of cavity C was too different to be fine-tuned into agreement with the other cavities.

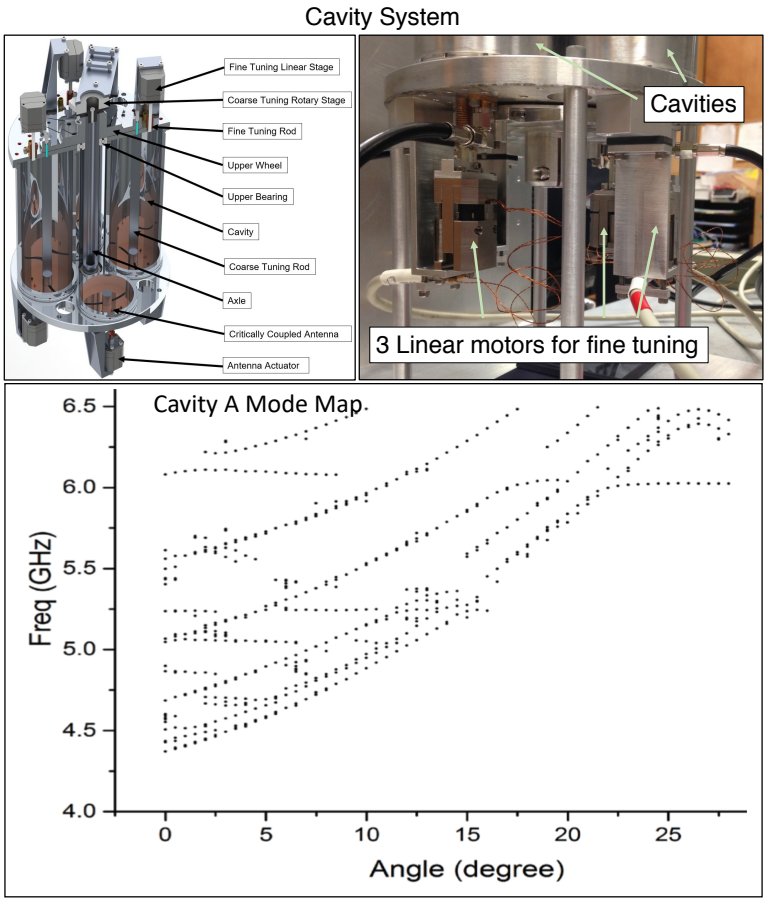

Figure 5.1: This figure shows the cavity system prototype used in the comparison tests at the University of Florida. A mode map of one of these cavities (cavity A) is shown in the bottom panel of this figure.

5.2 Procedure

This comparison test procedure was adapted from draft procedure written by Christian Boutan on October 30th 2017. Since the PNNL system was tested before the UF system, the particular implementation of the PNNL system testing was used to update and add detail to the procedure that the UF group would follow. This was done so that both sets of tests could be compared directly and so that both systems could be judged on an equal footing.

5.2.1 Test Protocol

Three cavities (A, B and D) will be locked 27 times as the coarse tuning rod tunes by one degree increments from the walls to the centers of the cavities. For each coarse tuning rod

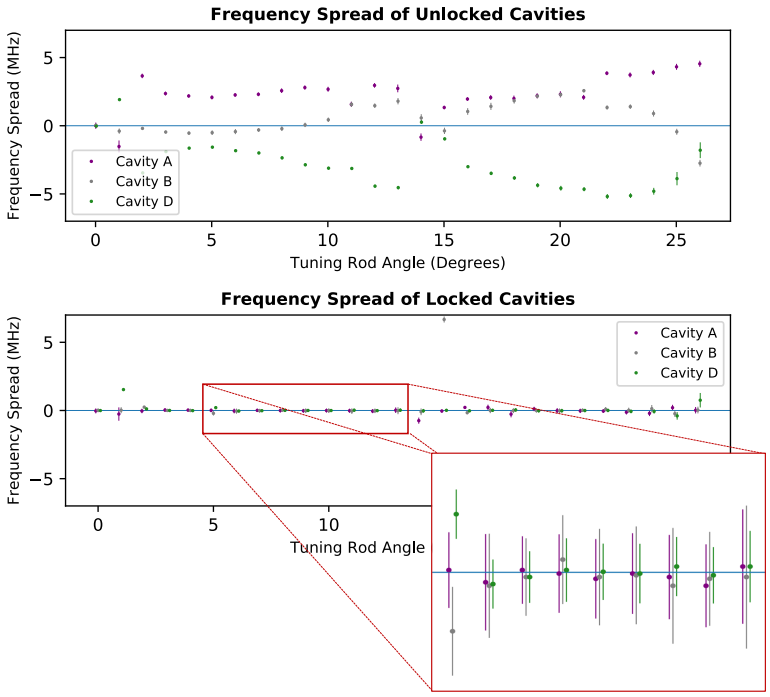

Figure 5.2: This figure shows improvement sustained by locking the microwave cavities versus unlocked operation. Error bars indicate the required locking tolerance for each cavity.

position, the cadence of the testing should follow the procedures listed below

1. Tune the coarse tuning rod by one degree and record the encoder value.
2. For each cavity/motor, record: encoder value, Q and Frequency.
3. Based on Cavity frequencies and motor positions, chose a target frequency, record that frequency and set the oscillator appropriately.
4. Initiate locking procedure
5. For each cavity: Measure final mode frequency with network analyzer
6. Make sure that your automated procedure successfully saved: cavity locking time, unavoidable steps and avoidable steps taken.
7. If you noticed anything out of the ordinary, make a note.
8. Repeat steps 1-7 over and over until you have tuned the coarse tuning rod from the wall to the center of the cavity.

This procedure is summarized in Figure 5.3 and represents a single test. This test comprising 27 different coarse tuning angles should be attempted twice for better statistics and may take 2 days because of the lack of automation.

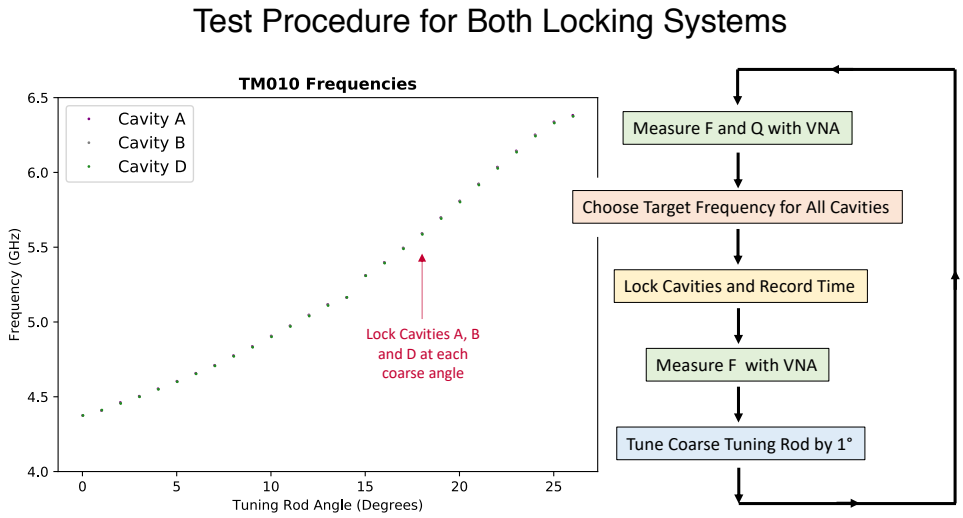

Figure 5.3: The left panel of this figure shows the successive frequencies in GHz of the coarse tuning rod positions at which cavities should be locked. The side of the figure is a block diagram of the procedure to be followed for each point appearing in the left hand side.

5.2.2 Saved Data

For each of the 27 cavity locking attempts, for each of the 3 cavities, for both of the identical attempted tests, the following values should be measured and recorded:

For each Coarse Rod Angle:

- Coarse Tuning Rod Angle
- Target (carrier) locking frequency chosen (see additional information below)
- Notes (see additional information below)

For each cavity:

- Pre-locked Fine-tuning Attocube encoder value
- Pre-locked Q
- Pre-locked Frequency (all significant figures available on the network analyzer)
- Post-locked Frequency (all significant figures available on the network analyzer)
- Time it took to lock cavity
- Unavoidable steps taken (see additional information below)
- Avoidable steps taken (see additional information below)
- Total steps taken (you can either record this directly or just add avoidable and unavoidable steps later)

5.2.3 Definitions/Clarifications

Target (carrier) locking frequency: Ideally, both locking systems would replicate the same target locking frequencies but realistically, the spread in viable target frequencies is very sensitive to coarse tuning rod and fine-tuning motor positions. The ideal target frequency will need to be set by the person performing the tests. Before deciding a target frequency, review the frequencies of each of the modes (already measured and recorded at this point) and observe where each of the fine-tuning motors are in their travel length. Pick a frequency that you think will cause minimal heating while at the same time keep the fine-tuning motors away from their extremes. Imagine you are actually taking axion data, and you want to keep the experiment running smoothly.

Notes: As you are locking the cavities for each of the 27 coarse rod positions, take note of mode structure features and any issues that come up. Later, you may need this information in order to explain a failed lock. Take note of mode crossings and double peaks. The double peaks (also observable in Jihee's 2017 mode maps) are a few MHz apart and may actually turn out to be the other cavities due to poor combiner isolation.

Unavoidable steps taken: If you can record the number of steps that you take, assign them a sign value depending on the direction. Then the unavoidable steps taken $= \sum$ (steps). By summing a vector representation of the number of steps gets rid of any instances when your system was inefficient. We would like you to record this data because it should be unrelated to the locking system being tested and your numbers should match ours.

Avoidable steps taken: The avoidable steps are the ones that separate an efficient locking system from an inefficient locking system and give you a sense of how much room there is for improvement. Avoidable steps $= \sum$ abs(steps) $- \sum$ steps $=$ total steps $-$ unavoidable steps.

Piezo Settings: Attocube step sizes will vary depending on the parameters set on the controller. The Attocube axes used to control the fine-tuning motors should be set to amplitude $=$ 60V, frequency $=$ 1kHz.

Fine-Tuning Movement: Fine-tuning motors should only be set to ideal starting positions at the beginning of the tests. They should NOT be reset or adjusted during the tests. The idea is to mimic real data taking. While resetting fine tuning motor positions between coarse tuning rod positions may improve cavity locking results, it would only result in more heating during an axion search. Pick your target frequency wisely to avoid the scenario where one of the motors runs out of travel length.

5.2.4 Post Processing

Saved data should be used to generate the following plots where Attocube coarse rod angles have been shifted to run from 0 to 26 degrees.

- A mode map of all cavity frequencies vs coarse tuning rod position. The scale of this plot is too big to note important features but it shows the context for the other plots.
- Frequency spread vs coarse rod position (like the mode map but with the target frequency removed) with dynamic error bars that show the 1/10th of a cavity FWHM tolerance for each cavity. Outliers and fine details should both be visible.
- A histogram of the frequency spread that has been normalized by the 1/10th of a cavity FWHM error bars such that successful cavity lock is represented by a number less than 1. Don't throw out outliers and failed attempts unless there is a clear explanation in your notes, such as a mode crossing.
- A plot showing the breakdowns of the locking times for each of the cavities for all trials.
- A histogram of the locking times.
- A plot showing the breakdown of the avoidable and unavoidable steps taken vs coarse rod angle
- A histogram of the avoidable and unavoidable steps taken per motor. They don't have to be on the same plot. It may make more sense to have them on separate plots so that the scales and bin sizes can be adjusted separately.
- A plot or text file containing the notes taken during the tests.

5.2.5 Summary

The goal of these tests is to put both systems though an identical set of semi-realistic scenarios to see where they fail and where they thrive. While 1 degree step sizes may seem arbitrarily large and unrealistic in the context of regular data taking the objective was to explore the entire mode map and subject the system to a full range of operating conditions: mode crossings, a wide range of frequencies and changing Qs. The following section details the specific implementation of both comparison tests, and Section 5.4 compares results of both systems.

5.3 Testing

5.3.1 Sweep Lock Testing

Testing of the PNNL Sweep Locking system was carried out December 11 – 14. by Christian Boutan at UF. Figure 5.4 shows pictures of the setup in Neil Sullivan's Lab as well as a

schematic of the Phase-1 system to be tested. Each set of tests took a day. During testing, the following issues were noted.

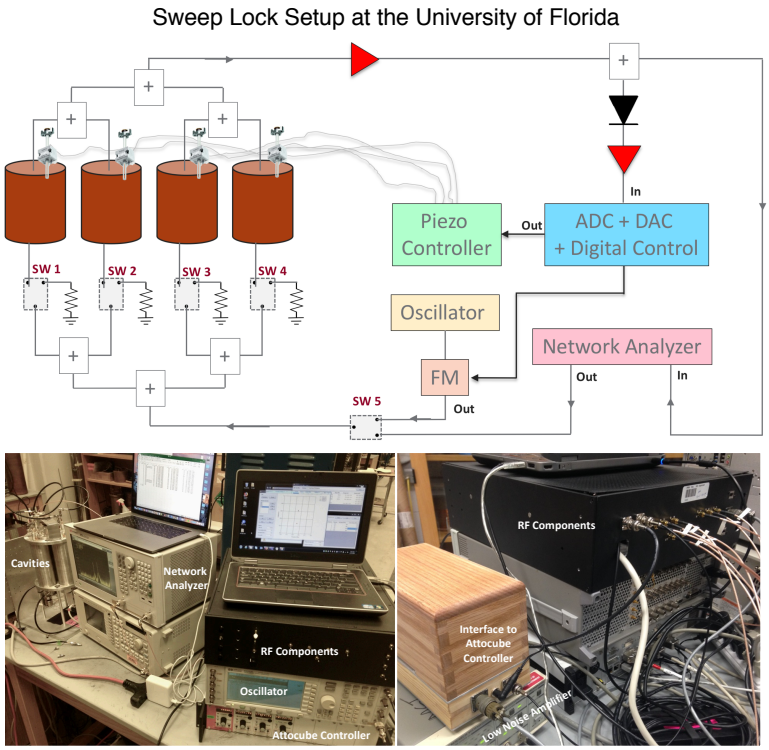

Figure 5.4: This figure shows a schematic of the Sweep Locking system developed at PNNL as well as pictures of the system being tested at the University of Florida.

Fine-Tuning Rods Run out of Travel Length in Mode Crossing: All of the A, B and D cavities are very similar but not identical. Typically one of the cavities would encounter a mode-crossing before the others, or the effect of the crossing would be more dramatic. In these scenarios, the slope of the localized mode map of one of the cavities will be greatly reduced, requiring increased tuning rod motion to produce the same frequency shift. Several times, in these situations, one of the motors would run out of travel length. No attempt to save the system or redo the measurement was made since we wanted the information about how these system would fail if commissioned. Mode crossings are a common occurrence in an axion cavity search and the fine-tuning system will need to help the experiment take data near (if not in) these regions of mixing. It is the authors' opinion that ADMX should consider either fatter rods or motors with a longer travel length.

Double-Peaks: During testing, there were a number of instances of there being an extra mode 2-10 MHz away from the TM_{010}. This happened so frequently that the occurrence didn't seem consistent with mode crossings and were dubbed "Double Peaks." After the fact, these double-peaks were observed as oval shaped points in Jihee

Yang's original mapping of the four-cavity array. This map and example network analyzer photos are shown in Figure 5.5).

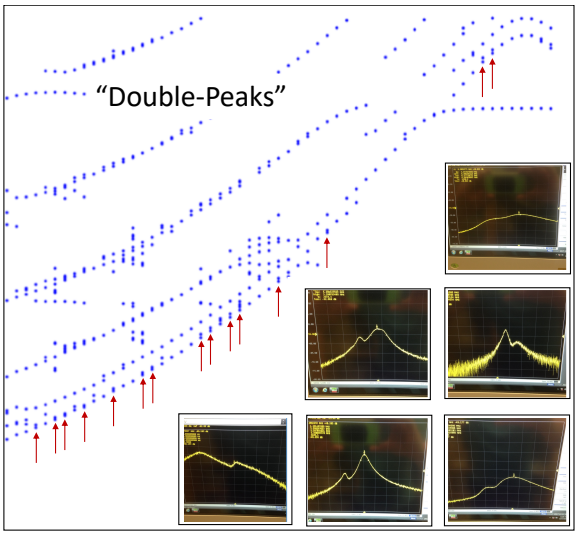

Figure 5.5: This figure shows the occurrence of double peaks in the mode maps of one of the test cavities used in the tests at the University of Florida. The pictures to the bottom right show examples of spectra also showing these double peaks.

Noted Issues vs Rod Angle: Problems and noteworthy issues were recorded.

Cavity *Q*s: Figure 5.6 shows how each of the cavity *Q*s change as a function of rod position. Error bars come from multiple measurements with different antenna couplings. This figure is present in the event that the reader has interest in the variance in cavity *Q*s in the context of multi-cavity power combining.

5.3.2 Pound Lock Testing

Testing of the UF Pound Locking system was carried out February 23 and 26 by Shriram Jois and Jihee Yang at UF. Figure 5.7 shows a schematic and a picture of the Pound Locking setup. Points worth noting are discussed below.

Mode crossings: Mode crossing regions were avoided (∼2, 15, 22 deg)

Error signal: Error signal and motor step were recorded in real time by a Labview program. Figure 5.8 shows examples of typical recorded error signals.

Antenna Coupling: At every coarse angle, strong port antennas were manually adjusted to achieve nearly critical coupling

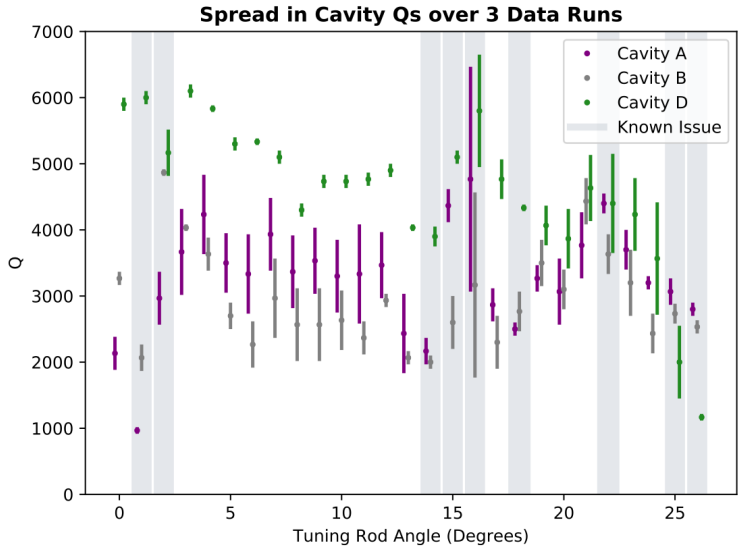

Figure 5.6: Qs of each of the cavities A, B and D versus tuning rod angle. Error bars come from multiple measurements for several tests with different antenna couplings.

Standing Waves: To minimize standing wave effect, a single-pole four-way switch was used to provide terminations for directional coupler. Figure 5.9 shows an example of standing wave interference that could potentially distort the TM modes therefore cause the inaccuracy in resonant frequency measurements.

Double-Peaks: Distorted peak shapes or double peaks were present in some frequency regions (see section 5.3.1)

Modulation Frequency Limit: The modulation frequency limit of the lock-in amplifier is 2.5 MHz, meaning that the capture range of the system is exceeded if cavity modes are separated by more than 2.5 MHz. In a realistic data taking scenario, this would rarely happen since the frequency step size of the experiment is a few to tens of kHz. However, for these tests, where the coarse tuning rod was tuned by degree increments, a technique called a "Multi-jump Lock" was implemented. Here, the Pound Locking system was initiated multiple times with different target frequencies in each iteration, thereby leading the mode to the right frequency in multiple steps. The results of each of these iterations were added together to give a total locking time, total motor steps etc.

5.4 Results

Here we present the results of both systems in the same section for easy comparison.

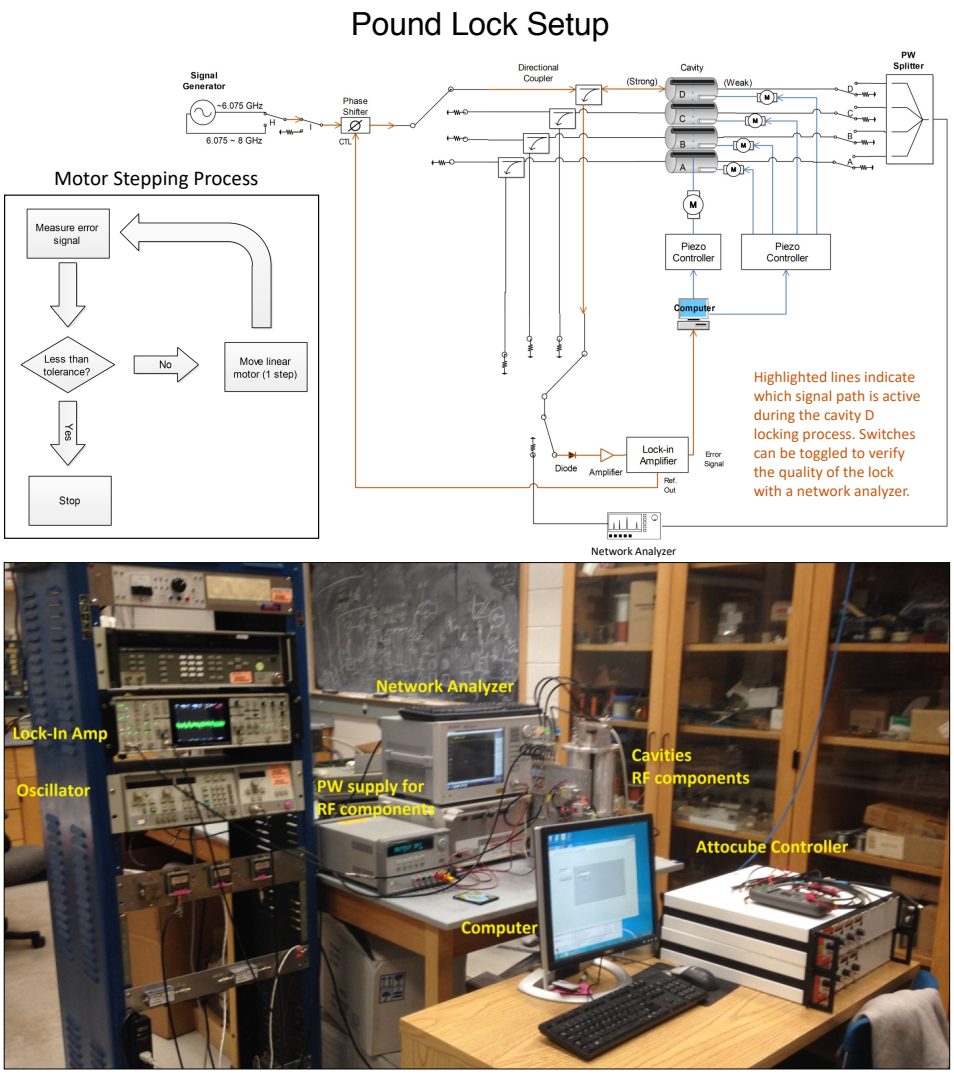

Figure 5.7: A schematic of the Pound Locking system built and tested at the University of Florida, as well as a picture of the setup in the laboratory.

5.4.1 Accuracy

The frequency spread in locked A, B and D cavities versus coarse tuning rod position is shown Figure 5.10. It is a measure of how far each cavity failed to reach the target frequency. Left to right, the PNNL Sweep Locking results run into the UF Pound Locking results. Top to bottom, results are organized by day. Error bars indicate how grouped cavity frequencies need to be in specification as dictated by the requirements (Section 3). Sweep locking error bars are different for each cavity since the tolerance is set by each cavity's Q. Pound Locking error bars are given by the average Q. Shaded regions indicate coarse tuning

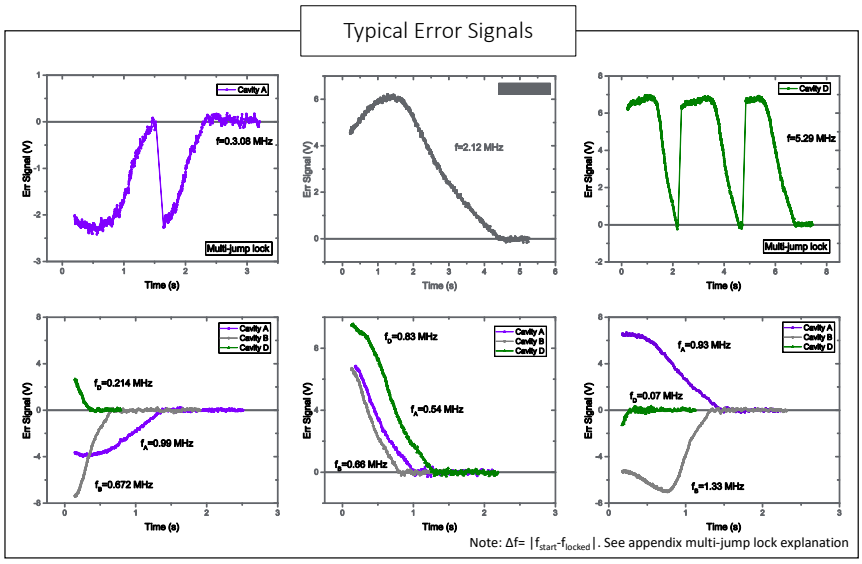

Figure 5.8: Examples of error signals taken during the Pound Locking tests conducted at the University of Florida.

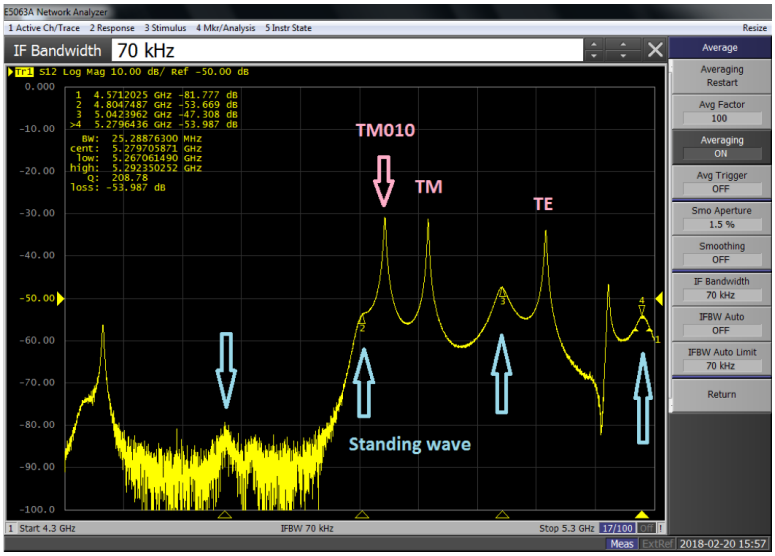

Figure 5.9: Standing waves present in the Pound Locking RF system, which could potentially distort the TM cavity modes, and hence interfere with the locking accuracy of the Pound Locking Technique.

rod positions where mode crossings are typically encountered.

The Sweep Locking results are very accurate and well within the requirement in almost all

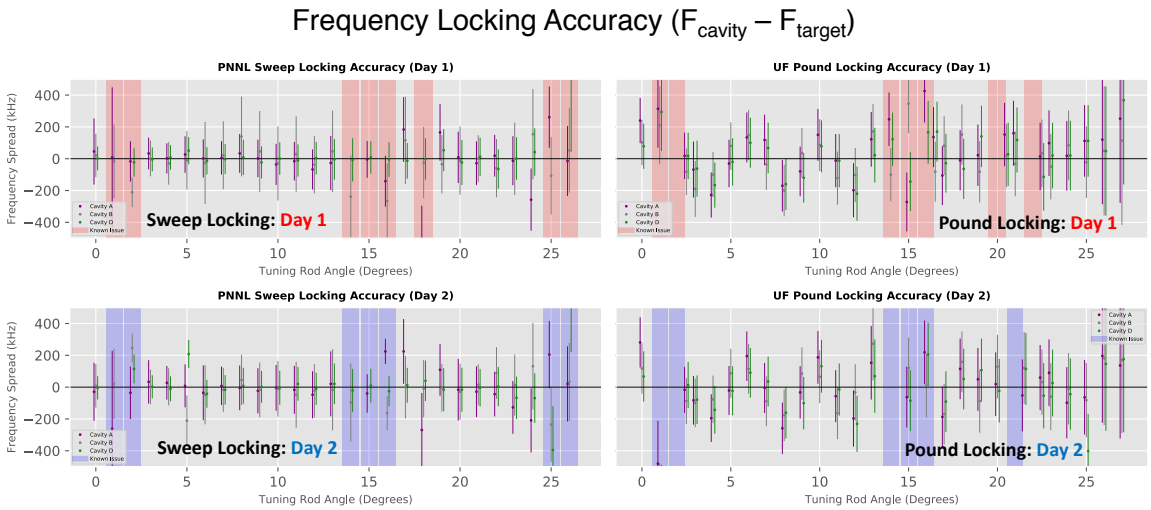

Figure 5.10: A comparison of the accuracy of the two systems tested at the University of Florida. Points plotted are the difference between mode frequencies and target frequencies. Sweep locking is to the left, and Pound Locking is to the right.

typical cases but cannot operate within spec in mode crossings. On average, the Pound Locking system operates within the requirements. The worrisome feature in the Pound Locking results is a clear systematic ripple that is repeated in the second day of testing and seems to be a function of rod position or frequency. If the cause of this systematic effect could be tracked down and fixed, the performance of the system could be greatly improved since the cavities themselves remain fairly well grouped. Unfortunately, if the frequencies are well locked but their center frequency is unknown, the experiment suffers a loss in sensitivity due to improper Lorentzian weighting in subsequent data analysis.

To avoid overshooting the peak and improve accuracy, the Pound Locking procedure was to switch the piezo motor drive voltage from 60V to 13V after equilibrium was reached. However, had the 60V piezo step size been a problem, the Sweep Locking would not have been able to achieve such precision. It is the opinion of the authors' that this evidence points to peak overshooting being caused by a measurement error as opposed to the piezo step size being inappropriately large.

Figure 5.11 shows these same results weighted by their error bars and histogrammed. Here the weighting scales the passing requirement to a value of one. On average, the Sweep Locking accuracy is 2.5% of a cavity FWHM, which is four times better than the requirement. The Pound Locking accuracy has a mean value of 6.45% of a cavity FWHM, which means that the system meets the requirement to be a viable technique. If the systematic ripple in the the locking accuracy were identified and fixed the vertical green mean accuracy line would be reduced from 6.45 to 3.37 which would be competitive with the Sweep Locking results.

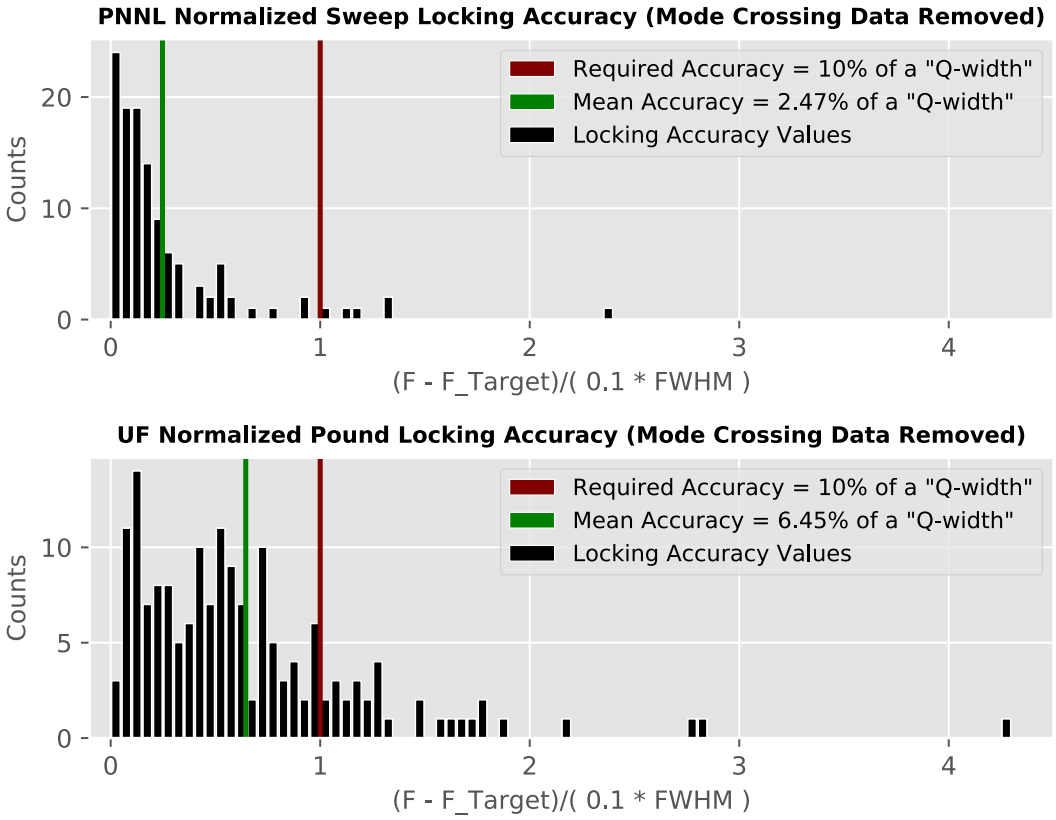

Figure 5.11: A comparison of the frequency accuracy of both systems weighted by the FWHM/10 requirement. Values less than 1 (to the left of the vertical burgundy line) represent results that pass the requirement.

5.4.2 Speed

All locking attempts for each cavity were timed during testing. A histogram of these individual cavity locking times is shown in Figure 5.12. The figure does not include data taken in mode crossings. The top figure shows Sweep Locking times and a green vertical line representing the 5.3 second mean value. In the event that the ADMX RF layout that has yet to be decided supports parallel locking (i.e. the ability to lock all cavities at the same time) this result meets the requirement specified in Section 3. However, the more likely outcome is that ADMX will choose a design that requires each cavity to be locked one at a time. In this scenario, the PNNL phase-1 Sweep Locking system would take just over 20 seconds to lock all cavities which is more than twice the required limit. By contrast, the Pound Locking system (bottom figure) set a 2-second individual cavity locking time which means that four cavities could easily be locked in under 10 seconds, meeting the 10 second requirement.

Furthermore, Figure 5.13 suggests that this speed could be improved. This example shows that while it takes 2.1 seconds to finish the locking process, were the algorithm refined, this time might be reduced close to the 1.2-second time needed to reach equilibrium. In terms of speed, the Pound Locking system is the clear winner.

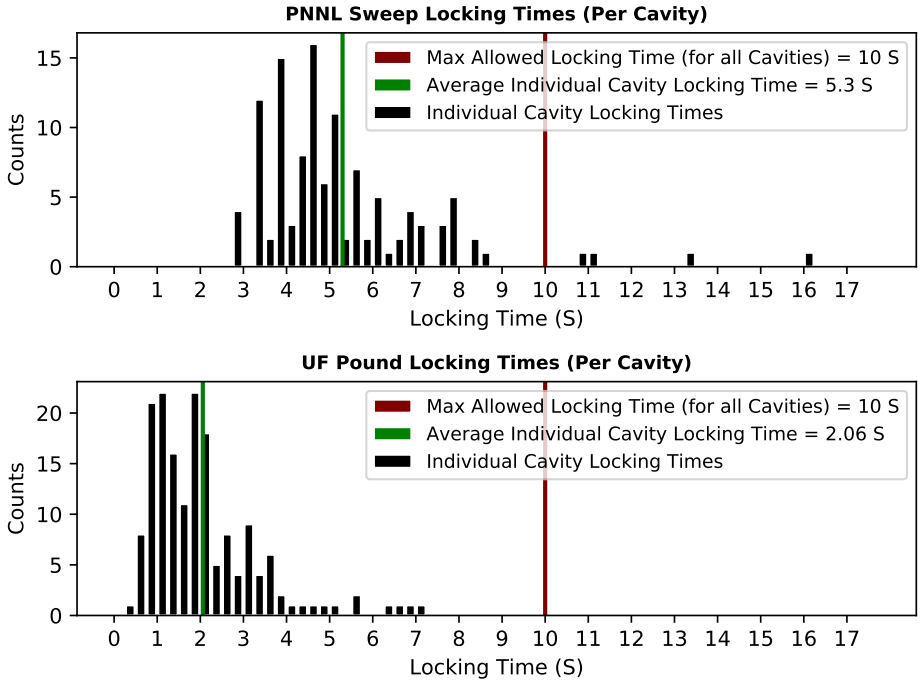

Figure 5.12: A histogram of the cavity locking times for both systems.

5.4.3 Heating

As described in Section 2.3, movement of the piezoelectric systems will cause heating in the cryogenic spaces of the experiment. The piezo motors used in the comparison tests and likely to be used in the final experiment are made by Attocube. Per step, these motors produce approximately

$$E_{\text{dissipated}} = CV^2 \tan(\delta) \tag{5.1}$$

Assuming typical cryogenic parameters (voltage = 60 V, C = 200 nF, $\tan(\delta) = 0.02$) the heating during normal operating experimental conditions would be $\sim$15 μJ per step. Using the knowledge learned from operating the Sidecar experiment, that a piezo step size bellow 4K is about a tenth of a piezo step size at room temperature, comparison test steps can be multiplied by 150 μJ to estimate the heating these systems would create in a cryogenic environment.

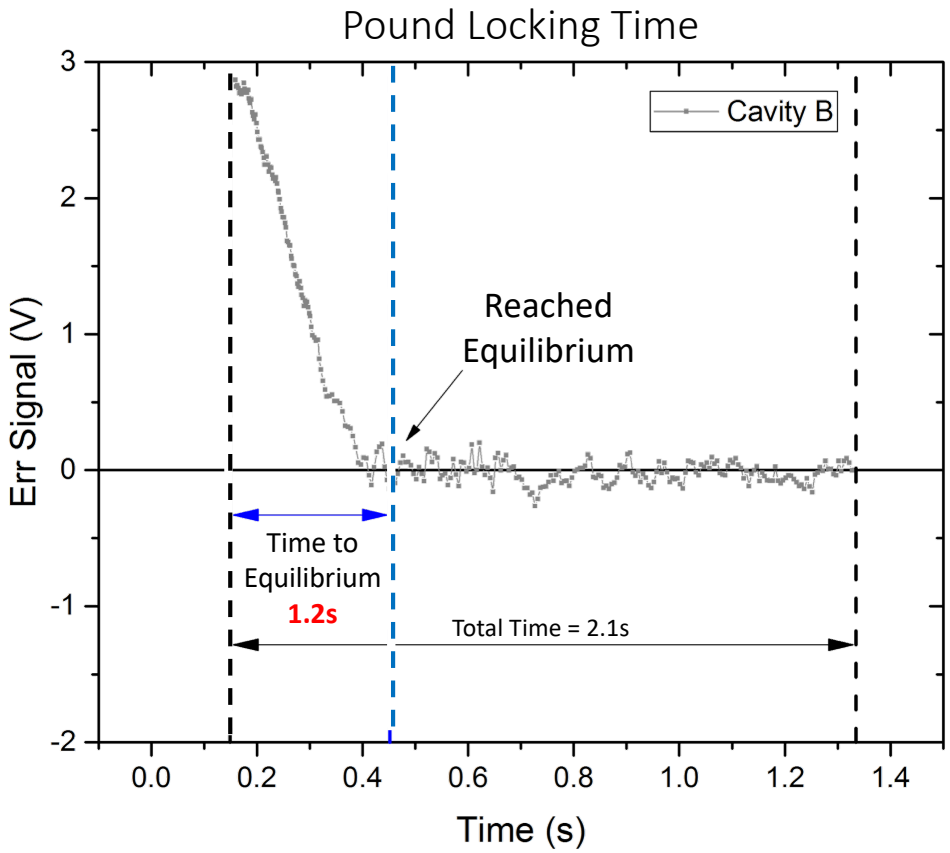

Figure 5.13: An example of the time taken by the Pound Locking system to reach equilibrium and then finish the locking process. The figure shows the potential to further speed up the Pound Locking process.

Since the heating requirement is specified in mW, this dissipated energy can be divided by the amount of time it takes the experiment to complete a single data taking cycle to achieve a heating in watts. Here we assume that the time associated with taking a digitized spectrum and performing other critical tasks is 120 seconds. In this case, the approximate heating is 1.25μW per step. By taking the 10mW heat load requirement, dividing by the number of cavities and solving for the number of steps, the outcome is that steps per motor $<$ 2000 to stay within the heat load budget. Figure 5.14 summarizes the measured number of steps for both systems. Steps are differentiated into "Unavoidable", "Avoidable" and "Total". Unavoidable steps are steps that must be taken to achieve a frequency lock. It represents the heating that is inevitable. Avoidable steps are the difference between the total and the unavoidable steps. They are a direct measurement of the inefficiency of the system and show if there is room for improving the technique.

When converting the results to mW per four-cavity system, the Sweep Locking technique produces 1.14mW and the Pound Locking system generates 1.67mW. The Sweep Locking technique comes out barely ahead but it's almost not worth mentioning. Both techniques

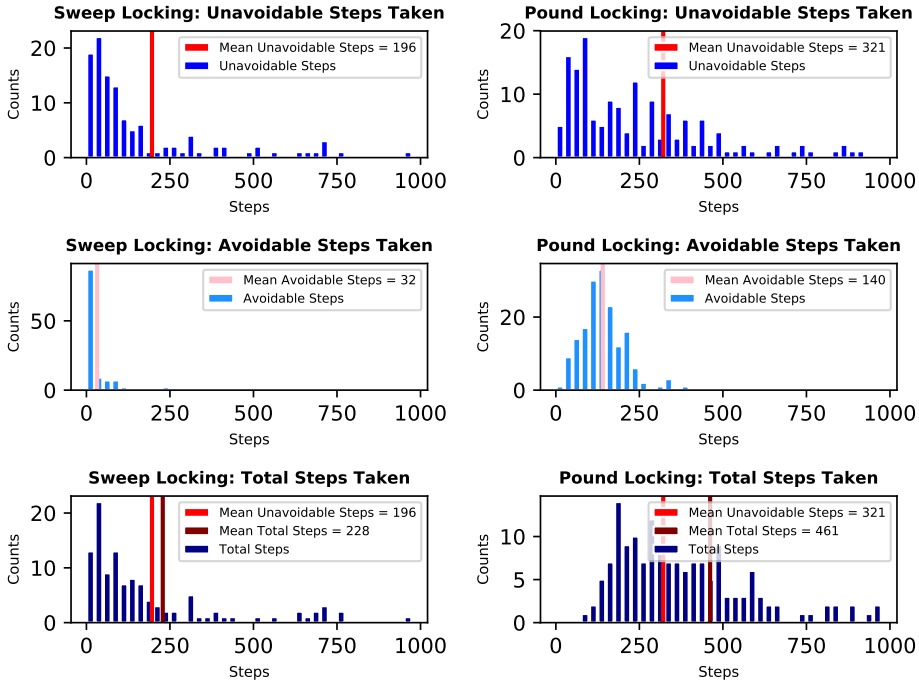

Figure 5.14: Unavoidable, Avoidable and Total Steps taken by both systems. Avoidable and total Pound Locking steps appear inflated and can't be mapped by the same $\sim 15\,\mu$J per step conversion because the motor voltage was switched from 60V to 13V.

should have nearly identical unavoidable steps, the Pound Locking system could probably be made more efficient and both systems meet the requirement by almost an order of magnitude.

5.5 Summary of Results

A summary of all of these results can be seen in Table 5.15. In short, the Sweep Locking system is roughly 2.5 times more accurate than the Pound Locking system, the Pound Locking system is roughly 2.5 times faster than the Sweep Locking system and both systems are well within the heating limit. The authors are worried about the systematic errors exhibited by the Pound Locking system and are concerned that the issue could get worse if the technique were taken out of a controlled environment. However, the Phase-1 Sweep Locking time is unacceptable in the event that parallel locking is not an option. In the following chapter we consider these results in the context of the rest of the experiment.

Summary of Comparison Tests

	Pound Lock (UF)	**Sweep Lock (PNNL)**	**Requirement**
Accuracy	6.45 % of FWHM	2.47 % of FWHM	< 10 % of FWHM
Time (per cavity) Time (4 cavities)	2.1 s 8.4 s	5.3 s 21.2 s	 < 10 s
Heating (per motor) Heating (4 motors)	417 µW 1.67 mW	286 µW 1.14 mW	 < 10 mW

Figure 5.15: A summary of comparison results

6 Further Analysis and Implementation

In the previous chapter we reviewed the results of the comparison tests. It was found that the PDH system was the faster system, that the PNNL technique was more accurate and that, in the context of an isolated set of tests, both systems appear to be viable. In this chapter we attempt to predict the final design of the ADMX RF in order to anticipate what it would take to implement each of these locking systems into the final design. Here we look for pitfalls associated with both locking systems that would not have been noticed in the comparison tests.

6.1 Unknowns and Prerequisites to a Final Design

The UF and PNNL systems both succeed at tuning multiple cavities to the same frequency. The question is which system will best compliment an incarnation of the ADMX experiment that is still in its designing phase. Here we review aspects of the experiment that are still unknown and may affect the outcome of the final design for the cavity locking system.

6.1.1 Choice of Piezo Motors

Attocube and JPE piezo motors are the two considered options for the tuning of the ADMX run 2A experiment. An Attocube system was used in the first three data runs of the Sidecar experiment, in the comparison tests and used for the heating calculations in this report. Tuning with the Attocube rotary motor has been very successful but there have been several failures with the linear (antenna) Attocube motors that haven't been diagnosed. Some issues were clearly the fault of the motors themselves and at other times, blame landed on mechanical systems unrelated to the motors. It has been successfully demonstrated that ethernet, USB and manual triggering from the back panel of the Attocube controller are all viable communication options. JPE systems have been tested at Lawrence Livermore National Laboratory (LLNL) and have been deployed in the current version of the Sidecar experiment. Currently, these motors are not working in the experiment and the issue has not been diagnosed. It has not been demonstrated that the direct motor triggering used by the Sweep Locking technique is possible.

Assuming that both locking techniques are able to communicate with either piezo motor system, in principal, the choice between Attocube and JPE shouldn't affect the choice between the Sweep Locking and Pound Locking systems. One exception would be if the JPE piezo motor system produced considerably more heating than the Attocube system. In this situation, cavity locking system heating efficiency would become more important.

6.1.2 Parallel vs Serial Locking

To this point, both cavity locking systems have been demonstrated with their own self contained cavity and switching systems. However, the final locking system must be integrated into an experimental architecture that is out of the hands of the cavity locking system designers. This ADMX RF architecture has not been specified and may have a dramatic impact on the design of the cavity locking system. For example, depending on the specifics of the RF cabling in the insert, the experiment could allow cavities to be locking simultaneously using four parallel locking systems. Such a system would cost approximately four times more than a single switched serial locking system. It would also run four times faster and would make the 10-second locking requirement a non-issue. However, a parallel RF design would require four independent signal paths out of the insert, in addition to the combined axion data output. Here, we propose two possible RF layouts to the collaboration. The layout in Figure 6.1 would be complemented by a single serial locking system whereas the layout in 6.2 would allow cavities to be independently locked, measured and diagnosed.

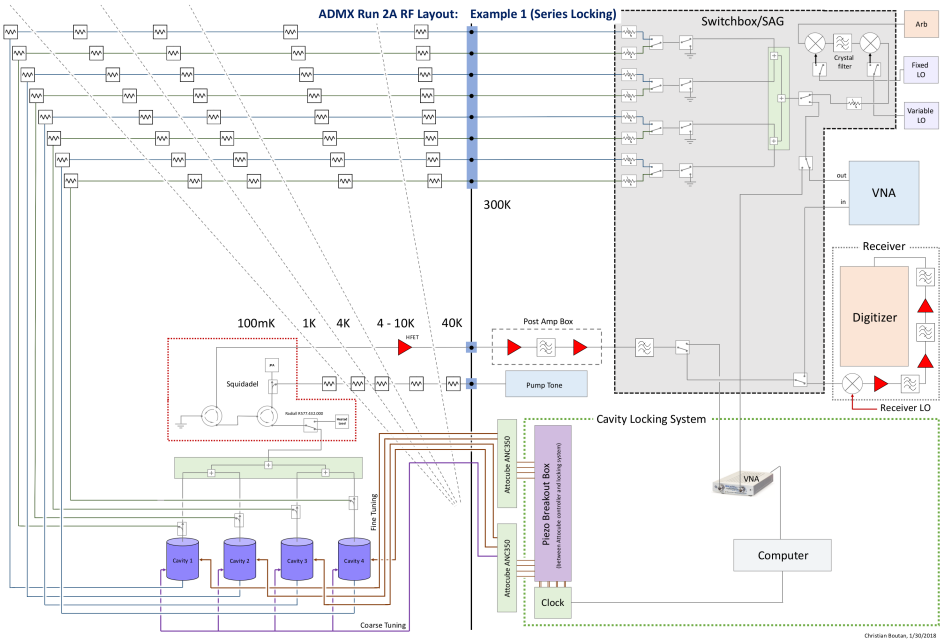

Figure 6.1: A single channel output example of a RF layout for run 2A of the ADMX experiment. This layout would require cavities be locked serially, one after the other. This figure shows a single channel implementation of the Phase-3 Sweep Locking system.

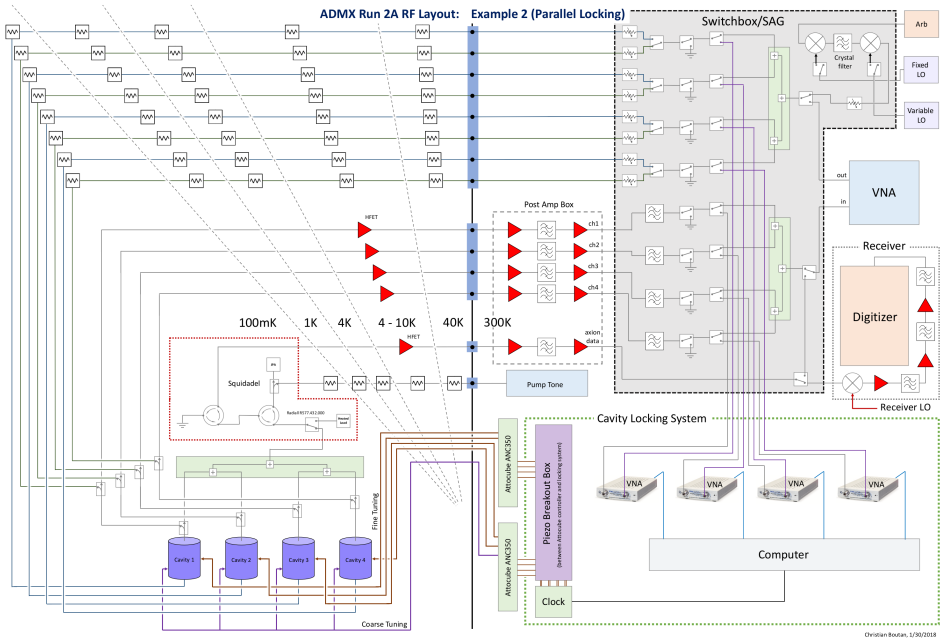

Figure 6.2: A multi-channel output example of a RF layout for run 2A of the ADMX experiment. This layout would allow each cavity to be independently measured permitting parallel (simultaneous) locking. This figure shows a multi-channel implementation of the Phase-3 Sweep Locking system.

6.1.3 Signal Power In/Out of the Insert

Treating the cryogenic insert as a black box, there will probably be restrictions on the amount of allowed injected power, and therefore limits on the locking SNR. This discussion is relevant because the locking systems in the comparison test were allowed to operate at arbitrarily high signal amplitudes and Sections 6.2.1.1 and 6.2.2.1 suggests the Phase-1 Sweep Locking system and the Pound Locking system will not operate at lower SNR.

If ADMX decides to build a single channel system similar to the one depicted in Figure 6.1, the locking SNR will be set by the compression point of the front end amplifier (approximately -120dBm) and the gain and system noise of the receiver chain. If, for example, the net gain in the insert is 50dB (20dB from JPA, 40dB from HFETs, 10dB loss), the locking signal exiting the insert will have a power of -70dBm. At this point in the chain, the SNR is essentially fixed even after further amplification and the chosen locking system will need to be able to operate with these input parameters.

If ADMX chooses a design that resembled the layout of Figure 6.2, compression in the JPA is no longer an issue since the signal is being taken out of an independent line. The question becomes, how much power can be applied to the JPA without damaging it? Email correspondence with Irfan Siddiqi and Sean O'Kelley suggest that -90dBm should be a safe

estimate but there is no clear answer. With an HEFT (40dB) on the bypass line to compensate for a 20dB loss through the weak port of the directional coupler, it would be possible to achieve a higher SNR especially if the JPA can be pushed to even higher powers. Both Siddiqi and O'Kelley proposed a switch that protects the quantum amplifier from seeing a large cavity locking signal intended for the additional signal bypass line. However, this switch would need to be operated every $\sim$ 2 minutes which would add risk and a lot of heating to the system.

In summary, the chosen locking system will need to be able to work with both piezo motor options over a wide possible range of signal powers. Furthermore, the design needs to be scalable/adaptable to be able to lock four cavities simultaneously.

6.2 Concerns about Implementation

Here we discuss possible pitfalls associated with implementing both techniques into ADMX run 2A.

6.2.1 Sweep Locking

6.2.1.1 Low SNR and RF Detectors

In the contest of the SNR discussion above, it was realized after the comparison test that the diode detector in the Phase-1 design would be saturated by noise in a low signal environment. In order to better filter the noise a stage of mixing was proposed. Figure 4.2 confirms the fear that the Phase-1 design would fail in low SNR situations by showing that a -70 dB signal is lost by the RF detector. The figure also shows that adding a stage of mixing improves the sensitivity of the technique by at least 30 dB. Thus, the Phase-1 design was recast to have a stage of mixing that passes the signal through a narrow filter. *This is essentially Phase 2.* This design was not tested beyond the proof of concept test shown in Figure 4.2. It was then realized that the swept measurement could be performed quicker and more accurately by a vector network analyzer. Phase 3, the newest version of the Sweep Locking technique uses a VNA in a feedback loop to perform the transmission measurements.

6.2.1.2 Locking Time

The Achilles heel of the technique from the perspective of the comparison test was the time needed to lock each cavity. This concern would apply to both the Phase-1 and Phase-2 designs. However, speed test comparison was performed and showed that the VNA could do the transmission measurement in 50 ms which is an order of magnitude faster than the 0.55 s

Phase-2 sweep time. It is estimated that had this system been used in the comparison tests, the average locking time would be 2 seconds per cavity which is competitive with the Pound Locking technique. While this exact system was not put through the comparison tests, ADMX has been using VNAs for a decade. In this sense, this new piece in the Sweep Locking technique has been extensively tested, can be purchased and is very fast.

6.2.2 Pound Locking

6.2.2.1 Low SNR and RF Detectors

Looking at the Pound Locking schematic in Figure 4.3 and setup in Figure 5.7 it appears as though this system would also suffer from the same loss of signal caused by the enormous noise bandwidth being taken in by the detector. In this case, it is not clear how the system would adapt. The Sweep Locking technique evolved by adding a stage of filtering and mixing. That approach doesn't appear to be an option in this situation.

6.2.2.2 Systematic Locking Error

The wiggles in the Pound Locking frequency spread shown in Figure 5.10 are worrisome. The effect was repeatable between the two days and the systematic effect has not been tracked down. The technique appears to be sensitive to cable length, phase shifts and standing waves. If this is the case, there is a fear that this effect could worsen if the technique is taken out of a controlled testing environment.

7 Chosen System

Based on all of the above considerations, we select the Phase-3 Sweep Locking technique for the final cavity locking design. It produces the most accurate locks, it is best equipped to deal with low SNR situations, and the use of a VNA will allow it to satisfy the speed-of-locking requirement. A VNA measurement was the standard by which both techniques were judged in the comparison test. In addition, by using a VNA to make the mode frequency measurements, we will know the exact cavity frequencies, which will be important for accurate data analysis. We propose the continued use of VNAs with the ADMX data taking process but instead of tuning via a manual operation (described in Sections 2.1.2 and 2.1.3) the VNA will be incorporated into a Sweep Locking feedback loop.

8 References

[1] R. D. Peccei and H. R. Quinn, "CP Conservation in the Presence of Pseudoparticles," *Phys. Phys. Lett.*, vol. 38, no. 25, pp. 1440–1443, 1977.

[2] S. Weinberg, "A New Light Boson?" *Phys. Phys. Lett.*, vol. 40, no. 4, pp. 223–226, January 1978.

[3] F. Wilczek, "Problem of Strong P and T Invariance in the Presence of Instantons," *Phys. Phys. Lett.*, vol. 40, no. 5, pp. 279–282, January 1978.

[4] J. E. Kim, "Weak-Interaction Singlet and Strong CP Invariance," *Phys. Phys. Lett.*, vol. 43, no. 2, pp. 103–107, July 1979.

[5] M. A. Shifman, A. I. Vainshtein, and V. I. Zakharov, "Can confinement ensure natural CP invariance of strong interactions?" *Nuclear Physics B*, vol. B166, pp. 493–506, July 1979.

[6] M. Dine, W. Fischler, and M. Srednicki, "A simple solution to the strong CP problem with a harmless axion." *Phys. Lett. B.*, vol. 104B, no. 3, pp. 199–202, August 1981.

[7] A. R. Zhitnitsky, "On Possible Suppression of the Axion Hadron Interactions," *Sov. J. Nucl. Phys.*, vol. 31, no. 260, pp. 497–504, 1980.

[8] P. Sikivie, "Experimental Tests of the 'Invisible' Axion," *Phys. Phys. Lett.*, vol. 51, no. 16, pp. 1415–1417, October 1983.

[9] 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. Andre´, J. Clarke, M. Mück, and R. F. Bradley, "Large-scale microwave cavity search for dark-matter axions," *Phys. Rev. D*, vol. 64, no. 092003, 2001.

[10] S. J. Asztalos, E. Daw, H. Peng, L. J. Rosenberg, D. B. Yu, C. Hagmann, D. Kinion, W. Stoeffl, K. van Bibber, J. LaVeigne, P. Sikivie, N. S. Sullivan, D. B. Tanner, F. Nezrick, and D. M. Moltz, "Experimental constraints on the axion dark matter halo density," *Astrophys. J. Lett.*, vol. 571, no. 1, pp. L27–L30, 2002.

[11] S. J. Asztalos, R. F. Bradley, L. Duffy, C. Hagmann, D. Kinion, D. M. Moltz, L. J. Rosenberg, P. Sikivie, W. Stoeffl, N. S. Sullivan, D. B. Tanner, K. van Bibber, and D. B. Yu, "Improved rf cavity search for halo axions," *Phys. Rev. D*, vol. 69, no. 011101, pp. 1–5, 2004.

[12] 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. Phys. Lett.*, vol. 104, no. 041301, pp. 1–4, 2010.

[13] J. H. *et al*, "Modulation sensitive search for nonvirialized dark-matter axions," *Phys. Rev. D*, vol. 94, no. 082001, pp. 1–13, 2016.

[14] J. S. *et al*, "Limits on axion-photon coupling or on local axion density: Dependence on models of the Milky Way's dark halo," *Phys Dark Universe*, vol. 14, pp. 95–102, 2016.

[15] C. R. Boutan, "A Piezoelectrically Tuned RF-Cavity Search for Dark Matter Axions," Ph.D. dissertation, U. Washington, Seattle (main), 2017.

[16] R. V. Pound, "Electronic Frequency Stabilization of Microwave Oscillators," *Rev. Sci. Instrum.*, vol. 17, no. 11, pp. 490–505, November 1946.

[17] T. Lindenberg and C. Bodefeld, "Technical Note: Power dissipation of a piezo," attocube systems AG, Munich, Germany, Tech. Rep., 2013.

[18] R. V. Pound, "Gravitational Red-Shift in Nuclear Resonance," *Phys. Phys. Lett.*, vol. 3, no. 9, pp. 439–441, 1959.

[19] R. L. Barger, M. S. Sorem, and J. L. Hall, "Frequency stabilization of a cw dye laser," *Appl. Phys. Lett.*, vol. 22, no. 11, pp. 573–575, June 1973.

[20] S. A. Helmcke and J. L. Hall, "Dye laser spectrometer for ultrahigh spectral resolution: Design and performance." *Appl. Opt.*, vol. 21, pp. 1686–1694, 1982.

[21] L. Hollberg, *Dye Laser Principles: With Applications*, ser. Quantum Electronics: Principles and Applications. ACADEMIC PRESS, INC., 1990, ch. 5. CW DYE LASERS.

[22] A. Abramovici, W. E. Althouse, R. W. P. Drever, Y. Gürsel, S. Kawamura, F. J. Raab, D. Shoemaker, L. Sievers, R. E. Spero, K. S. Thorne, R. E. Vogt, R. Weiss, S. E. Whitcomb, and M. E. Zucker, "LIGO: The Laser Interferometer Gravitational-Wave Observatory," *Science*, vol. 256, no. 5055, pp. 325–333, April 1992.

[23] B. P. Abbott *et al*, "LIGO: The Laser Interferometer Gravitational-Wave Observatory," *Reports on Progress in Physics*, vol. 72, no. 7, p. 076901, 2009.

[24] ——, "GW151226: Observation of Gravitational Waves from a 22-Solar-Mass Binary Black Hole Coalescence," *Phys. Phys. Lett.*, vol. 116, no. 241103, pp. 241 103–1, June 2016.

[25] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward, "Laser Phase and Frequency Stabilization Using an Optical Resonator," *Appl. Phys. B*, vol. 31, pp. 97–105, February 1983.

[26] J. Hough, D. Hils, M. D. Rayman, L. Ma, L. Hollberg, and J. L. Hall, "Dye-laser frequency stabilization using optical resonators," *Appl. Phys. E*, vol. 33, pp. 179–185, 1984.

[27] J. Helmcke, J. J. Snyder, A. Morinaga, F. Mensing, and M. Gläser, "New ultra-high resolution dye laser spectrometer utilizing a non-tunable reference resonator," *Appl. Phys. B*, vol. 43, no. 2, pp. 85–91, June 1987.

Pacific Northwest
NATIONAL LABORATORY

U.S. DEPARTMENT OF
ENERGY

902 Battelle Boulevard
P.O. Box 999
Richland, WA 99352
1-888-375-PNNL (7665)
www.pnl.gov