UCLA
UCLA Electronic Theses and Dissertations

Title
Traceable and Precise Displacement Measurements with Microwave Cavities

Permalink
https://escholarship.org/uc/item/71x57775

Author
Koulakis, John

Publication Date
2014

Peer reviewed|Thesis/dissertation

University of California

Los Angeles

Traceable and Precise Displacement Measurements with Microwave Cavities

A dissertation submitted in partial satisfaction

of the requirements for the degree

Doctor of Philosophy in Physics

by

John Pandelis Koulakis

2014

ABSTRACT OF THE DISSERTATION

Traceable and Precise Displacement Measurements with Microwave Cavities

by

John Pandelis Koulakis

Doctor of Philosophy in Physics

University of California, Los Angeles, 2014

Professor Károly Holczer, Chair

The difficulty of making accurate, repeatable, sub-nanometer displacement measurements has limited the progress of nanotechnology and surface science. Scanning probe microscopy, an important set of tools for characterizing nanoscale structures, is capable of atomic resolution imaging, but has yet to realize its full, metrological potential. This work evaluates the feasibility of using microwave cavities to address this need. Accurate and stable RF frequency references have become ubiquitous and are an attractive option for realizing traceable distance measurements through the resonant frequency of microwave cavities operating in TEM modes. A method of measuring the resonant frequency of such cavities capable of sensing picometer displacements is developed. The concept is demonstrated with a variable-length, 10 GHz coaxial cavity, and proves to have a resolution of $60 \text{ fm Hz}^{-1/2}$, and a range of 10 µm. Independent measurements with an interferometer verify that the device is capable of displacement measurements accurate to 1% without external calibration, and with non-linearity $< 5 \times 10^{-4}$ of the measured range. Appropriate mechanical design can extend the range and improve the accuracy. Incorporating this system into scanning probe microscopes would allow them to measure sub-atomic distances confidently.

The dissertation of John Pandelis Koulakis is approved.

Tatsuo Itoh

Eric R. Hudson

Károly Holczer, Committee Chair

University of California, Los Angeles

2014

iii

. . . to my parents,
who nurtured my curiosity.

Contents

vi

LIST OF FIGURES

LIST OF TABLES

Vita

2006	B.A. (Physics), Pomona College.
2007	Mellon Post-Baccalaureate Fellow.
2007-2008	Graduate Assistance in Areas of National Need (GAANN) Fellow.
2007–2014	Teaching/Research Assistant, Physics and Astronomy Department, University of California, Los Angeles.
2009	M.S. (Physics), University of California, Los Angeles.

Publications and Presentations

John P. Koulakis, Catalin D. Mitescu, Franqise Brochard-Wyart, Pierre-Gille De Gennes and Etienne Guyon (2008). The viscous catenary revisited: experiments and theory. *Journal of Fluid Mechanics*, **609**, 87-110.
doi: 10.1017/S0022112008002413.

John P. Koulakis and Catalin D. Mitescu (2007). The viscous catenary. *Physics of Fluids*, **19**, 091103. doi: 10.1063/1.2775166.

Prize-winning entry in the 24th Annual Gallery of Fluid Motion at the American Physical Society Division of Fluid Dynamics 2006 conference.

CHAPTER 1

Introduction - The Frequency-Distance Encoder

As technology shrinks there is an increasing need for systems that can accurately and reproducibly measure nanoscale distances [1, 2]. The growing importance and wealth of surface science has brought about changes in conventional instruments, and the development of new tools, to deal with its special needs — reduced diffraction volume, loss of periodicity, weakened cohesion, increased contamination, etc. About 30 years ago, an important group of instruments emerged — scanning probe microscopes (SPM) — that is particularly well adapted to characterize surface properties at nanoscale distances. However, its metrological accuracy lagged far behind what was customary for bulk lattice measurements, typically $\sim 1\%$ of interatomic distances, or a few picometer.

Historically, early SPMs used piezo actuators that did not incorporate position sensors. General practice was to rely on known characteristics of the sample to roughly calibrate the axes [3]. Often, bulk crystallography data was used to assign dimensions to images, despite the fact that the reason to study surfaces is to infer the difference from the bulk[1]. This practice has occasionally led to the mislabeling of surface atoms and the misinterpretation of images. Improved SPMs incorporate sensors capable of sub-nanometer resolution into closed-loop positioning stages to minimize the effects of hysteresis, piezo creep, and drift [5, 6]. Independent calibration with interferometers, certified reference artefacts[2], or other means is

---

[1]One of the most striking examples is the 7x7 reconstruction of the Si(111) surface, resolved by Binnig [4].

[2]A comprehensive list is available at http://www.nanoscale.ptb.de/nanoscale/nanoscale-standards.html.

required. Closed-loop positioning is, however, commonly turned off when atomic resolution is desired, due to the relatively large sensor noise.

More recently, there has been significant effort at national metrology institutes (NIST, NPL, PTB) to develop the next generation of high-stability SPM capable of traceable, high-resolution, displacement measurement [7, 8]. Attempts to realize truly metrological SPM have traditionally employed integrated laser interferometers [3, 9–11]. Unfortunately, due to a combination of positioning noise and/or nonlinearities associated with fringe interpolation [12], metrological SPMs have not yet been able to reliably measure surface lattice parameters. Some have suggested mapping out non-linearities with x-ray interferometry [13, 14], while others [15, 16] have taken different approaches. Progress towards achieving a target uncertainty of 10 pm in the next generation of optical interferometers is described in [17].

This work investigates the possibility of using a microwave cavity as a self-calibrated distance sensor - coined the "frequency-distance encoder." With the proper choice of mode, the resonant frequency will depend only on the length of the cavity, with the speed of light as the only conversion factor. A measurement of the cavity resonant frequency will then uniquely determine the distance between the cavity ends, to a precision equal to that of the frequency measurement. RF frequency standards disciplined to atomic transitions provide stability and traceability [18] to the definition of the meter. It will be shown in this dissertation that the technique is capable of achieving the desired accuracy.

Part I is a description and evaluation of a method of measuring the cavity resonant frequency relative to the frequency of an oscillator, providing a feedback signal for closed-loop control of the frequency-distance encoder. Chapter 2 is a general overview of frequency control methods, one of which is chosen for use with the frequency-distance encoder. Chapters 3 and 4 develop the theory of the lock while chapter 5 covers the hardware and its performance. The technique is capable

of detecting length/frequency fluctuations of the cavity down to the level of the frequency noise of the oscillator. A frequency noise-floor of 0.04 Hz $\text{Hz}^{-1/2}$ has been demonstrated using common sources, corresponding to a distance noise-floor of 60 fm $\text{Hz}^{-1/2}$ for a 10 GHz cavity. The bandwidth of the measurement is only limited by the cavity linewidth, and can be megahertz if desired. In closed-loop, the cavity length can be set by changing the frequency of the source. The length inherits the stability and accuracy of the frequency reference, provided proper cavity design. With recent advances in phase-lock loop technology, the frequency can be stepped with ~ 10 µHz precision (length steps of 15 am) while maintaining the stability of the frequency reference.

In part II, the focus shifts to the design and performance of the realized frequency-distance encoder. To best serve as a frequency-to-distance converter, a microwave cavity should provide two parallel reference planes, preferably coinciding with the location of two cavity walls, whose separation defines the length to be controlled. It is also important that a mode exists where the resonant frequency is related to the length through only the speed of light. A common example of a cavity with the desired properties is the coaxial resonator, whose theory is presented in chapter 6.

Coupling to microwave cavities is typically done through an iris joining a waveguide to the cavity or via a loop of wire entering the cavity. For this application, it is important that the coupling structure has a minimal effect on the cavity resonant frequency. Details of the coupling structure employed are given in chapter 7, and its effect on the resonant frequency is discussed in section 11.3.

To make the cavity length tunable, it must be cut in a way that allows moving the two reference planes apart, while keeping them parallel and minimizing any perturbation to the resonant mode. In a coaxial cavity operating in the fundamental TEM mode, there is no current that crosses the plane exactly half way between the cavity ends. This permits slicing it in that plane without disturbing

the mode. Moving the cavity halves apart will leave a gap between the pieces and perturb the mode. But, if the gap is small, the perturbation can be ignored. The mechanical design of the frequency-distance encoder is described in chapter 8, and the effect of the gap is discussed in detail in section 11.1.

The setup and general performance of the frequency-distance encoder is presented in chapter 9 and environmental influences are discussed in chapter 10. Chapter 11 is an evaluation of the largest effects that alter the frequency-to-distance conversion, and may contribute systematic errors to length measurements. The frequency-distance encoder is shown to be self-calibrating to an accuracy[3] of 1%, with non-linearity less than 5×10^{-4} of the measured range. Its resolution is 60 fm $\text{Hz}^{-1/2}$, and its mechanical bandwidth can be as large as 8 kHz. Such a device would allow quantitative study of surface atomic structure if incorporated into an SPM.

---

[3]Calibrating with an independent source can improve the accuracy.

Part I

The Microwave

Pound-Drever-Hall Lock

CHAPTER 2

Introduction to Automatic Frequency Control

Automatic frequency control (AFC) is needed to generate a feedback signal for the frequency-distance encoder. Various methods of AFC are presented here, one of which is chosen for further development in the subsequent chapters.

The concept of AFC has been around since the early 20th century [19]. A frequency discriminator, typically a relatively stable resonator or cavity, is employed to measure a microwave frequency. The signal produced by the various AFC schemes described below is a voltage proportional to the frequency difference between the oscillator and cavity. Traditionally, this signal is fed back to lock the oscillator to the cavity, thereby stabilizing it. However the inverse, locking the cavity to the oscillator, may also be worthwhile if a stable oscillator is available.

Among the earliest methods of AFC was the traditional frequency modulation technique [20] shown in figure 2.1a. Microwave is directed to a frequency-matched cavity resonator while a frequency modulation (FM) is applied to the source with a modulation frequency (typically ~ 70 kHz) much smaller than the resonance width of the cavity. The response to this modulation is separated from the back-reflected microwave signal with the help of a lock-in amplifier referenced to the FM frequency, yielding a signal that is proportional to the first derivative of the cavity absorption as a function of frequency. The integrated error signal is fed back to control the frequency of the microwave source to complete the AFC loop, in a typical bandwidth of ~ 1 kHz. This technique has been widely used in fields as diverse as Electron Paramagnetic Resonance (EPR) and FM radio receivers.

(a) Traditional AFC   (b) Pound Lock

(c) DC-AFC   (d) Pound-Stein-Turneaure
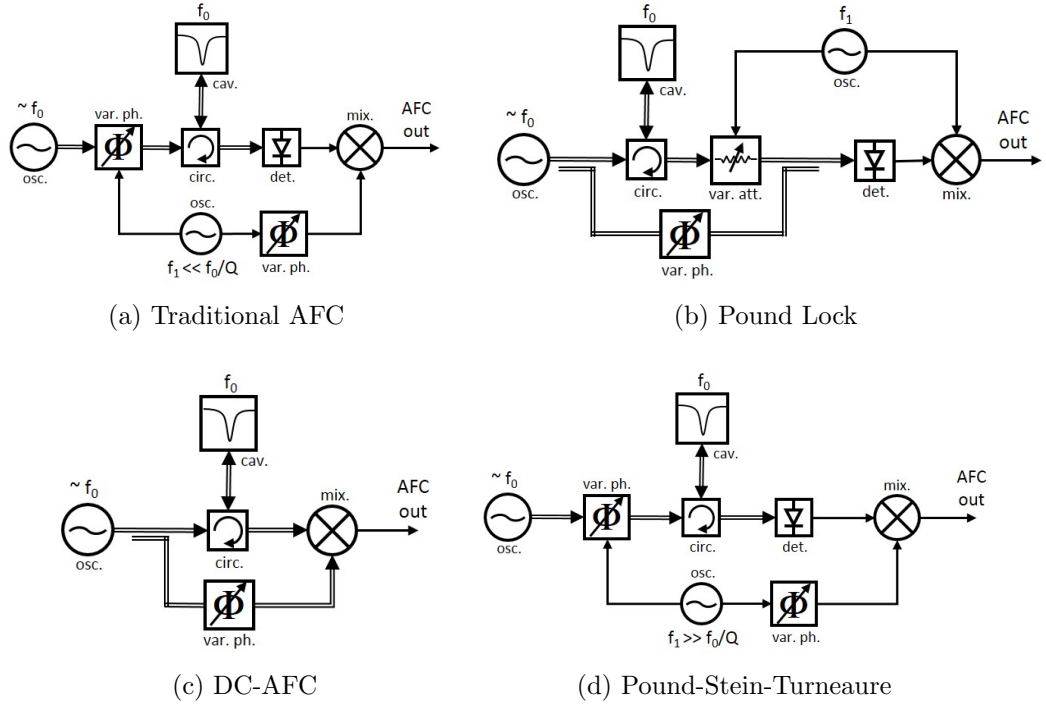
Figure 2.1: Simplified block diagrams of the most commonly encountered automatic frequency control (AFC) circuits. Abbreviations above: Oscillator (osc.), variable phase-shifter (var. ph.), circulator (circ.), cavity (cav.), detector (det.), mixer (mix), variable attenuator (var. att.).

With this technique, the FM modulation must be slow compared to the linewidth, which limits the bandwidth that is achievable. A variation of this method is to modulate the cavity frequency instead of the oscillator frequency [21] to achieve the same result.

Another classic method is the so-called "Pound" lock [22] and its variations [23]. Figure 2.1b is a block diagram of the classic Pound lock[1] adapted for modern components. The oscillator output is split. Part is used as a reference, and the other part is directed to a frequency-matched cavity resonator. The key realization of Pound was that the phase of the back-reflected microwave contains information about the location of the cavity resonance. This phase is measured by comparing

---

[1]Pound describes two version of AFC, the DC lock and the AC lock. The AC version is described here.

the back-reflected microwave to the reference. Pound originally included a form of amplitude modulation in order to increase the signal-to-noise ratio of his detection scheme, but with modern components, this is no longer necessary. Removing the amplitude modulation and replacing the detector with a mixer yields the "DC-AFC" circuit of figure 2.1c. For proper operation, the two arms of the Pound lock need to be set in quadrature. As the two arms are physically different transmission lines, differential variations in length can adversely affect the stability of the lock.

This concern is eliminated in the form of Pound stabilization developed to a high degree by Stein and Turneaure [24, 25] among others [26, 27], herein referred to as the Pound-Stein-Turneaure (PST) lock, figure 2.1d. Stein and Turneaure frequency modulate the source at a frequency much larger than the cavity linewidth. The generated sidebands travel the same path to the cavity as the carrier. As the modulation is much greater than the cavity linewidth, the sidebands are immediately reflected from the cavity coupler and serve as a reference to measure the phase change of the back-reflected carrier.

Soon after, this technique was adapted to optical frequencies, and became well-known as the "Pound-Drever-Hall (PDH) laser stabilization technique." Drever and Hall [28, 29] used the PST scheme to lock/stabilize lasers to Fabry-Perot resonators. The PDH frequency-stabilizing technique is fundamental in the construction of ultra-stable lasers and became an indispensable tool of current atomic-physics research. It is widely discussed and understood in the scientific literature (see for example [30]), taught in undergraduate laboratory classes, and used at top level research institutes.

It is somewhat surprising that the microwave version, the original PST lock, never got traction, wide recognition, or much use in its native field. One of the reasons may be that frequency modulation faster than the linewidth of room temperature cavities was difficult to realize. A cavity with a Q of 1,000 at 10 GHz would have a linewidth of 10 MHz requiring frequency modulation greater than $\sim$

100 MHz. The community that did use the PST lock was limited to those working with superconducting cavities, having exceedingly narrow linewidths ~ 100 Hz where the speed requirement was not so stringent. They also fully understood and took advantage of the fact that the reference and carrier of the PST lock traveled the same path.

Chapters 3-5 describe our modern realization of the PST lock and contain a detailed analysis of its performance. As the term "PDH" lock has become ubiquitous, the modern implementation of the PST lock will herein be referred to as the "Microwave Pound-Drever-Hall lock" (µPDH). Current realizations of the PST lock use devices that did not exist during its original development, and provide capabilities and features that were not originally possible. Vector modulators, for instance, provide an alternative to the pure phase/frequency modulation originally employed. With advances in Phase-Locked-Loops (PLL) and Direct Digital Synthesis (DDS), microwave sources can be phase-locked to stable, frequency references while maintaining the ability to step the frequency with high resolution. The subsequent chapters will describe how to convert the synthesized frequency to a stable and precise length, through a microwave cavity.

CHAPTER 3

Microwave Reflection off of a Cavity

In this chapter, a summary of cavity theory is presented, and notation employed throughout the text is introduced. Readers familiar with the topic are encouraged to skip ahead and refer back as needed for clarification on symbols used.

3.1 The Reflection Coefficient

The heart of the PDH signal is the cavity reflection coefficient, Γ. A wave traveling down a waveguide is incident onto a cavity as depicted in figure 3.1. Most of the wave is immediately reflected from the coupler (iris shown), and a small fraction enters the cavity. The field builds up in the cavity, and part of it leaks back out into the waveguide. In steady state, the reflected signal is the sum of the immediately reflected field and the cavity leakage field, and is described by the complex reflection coefficient Γ.

Altman, [31], gives Γ as[1],

$$\Gamma = -\frac{r - e^{-\nu + i\mu}}{1 - re^{-\nu + i\mu}},\tag{3.1}$$

where r^2 is the reflectivity of the coupling structure, $\mu = 2k_1 L$ is the phase contributed by a round trip pass through the cavity of length L, k_1 is the propagation

---

[1]The notation has been changed from [31], where $r = \sqrt{1 - k^2}$. The real number k represents the coupling through the iris. The real number r is always greater than zero. The π phase change upon reflection, which would result in a negative r in some notation, has been included separately. See also appendix A for a derivation of $|\Gamma|$ in terms of the scattering parameters of the coupling structure.
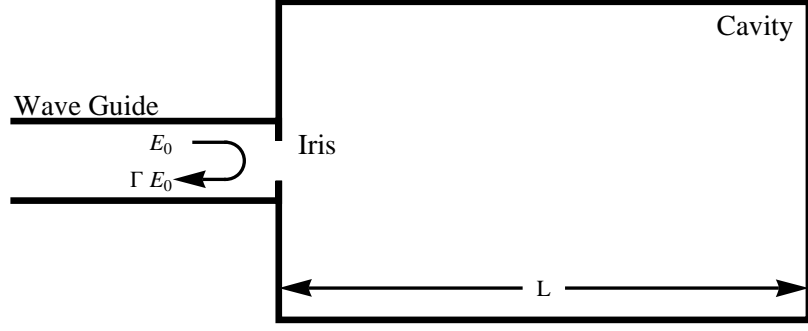
Figure 3.1: Reflection off a microwave cavity is described by the reflection coefficient Γ, which is a function of the difference between the source frequency and the cavity resonant frequency, the cavity Q, as well as the coupling between the waveguide and the cavity.

constant in the cavity, $\nu = 2k_2 L$ is the round trip loss within the cavity, and k_2 is the attenuation constant in the cavity. For later convenience, various forms of Γ and its magnitude are given here,

$$\Gamma = \frac{-r(1 + e^{-2\nu}) + e^{-\nu}(1 + r^2)\cos\mu + ie^{-\nu}(1 - r^2)\sin\mu}{(1 - re^{-\nu})^2 + 4re^{-\nu}\sin^2\frac{\mu}{2}}, \tag{3.2a}$$

$$|\Gamma|^2 = 1 - \frac{(1 - r^2)(1 - e^{-2\nu})}{(1 - re^{-\nu})^2 + 4re^{-\nu}\sin^2\frac{\mu}{2}}. \tag{3.2b}$$

These equations can be written more elegantly in terms of $\rho = -\ln r$,

$$\Gamma = \frac{-\cosh\nu + \cosh\rho\cos\mu + i\sinh\rho\sin\mu}{\cosh(\nu + \rho) - \cos\mu}, \tag{3.3a}$$

$$|\Gamma|^2 = \frac{\cosh(\nu - \rho) - \cos\mu}{\cosh(\nu + \rho) - \cos\mu}, \tag{3.3b}$$

$$= 1 - \frac{\sinh\nu\sinh\rho}{\sinh^2\frac{\nu+\rho}{2} + \sin^2\frac{\mu}{2}}. \tag{3.3c}$$

One might recognize $\frac{4re^{-\nu}}{(1-re^{-\nu})^2} = \frac{1}{\sinh^2\frac{\nu+\rho}{2}}$ as being the equivalent of the coefficient of finesse for optical cavities.

11
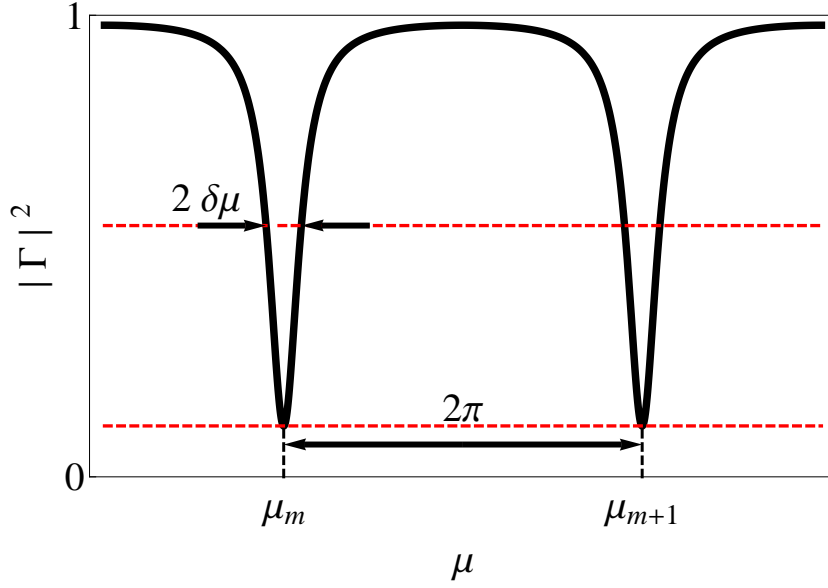
Figure 3.2: The cavity resonances appear as dips in the reflected power as a function of frequency. In terms of the frequency-like parameter μ, the linewidth is $2\delta\mu$ and the resonances are spaced by 2π.

3.2 Conditions for No Reflection and the Coupling Coefficient

Setting equation (3.3b) equal to zero gives the conditions for no reflection,

$$\cos\mu = \cosh(\nu - \rho). \tag{3.4}$$

The only way this equation can be satisfied is if $\nu = \rho$ (round trip cavity loss is equal to the losses through the coupler), and $\mu = 2\pi m$, with m an integer. For cavities of interest, m will be small, typically 1.

The cavity coupling coefficient, β is commonly defined in terms of the mini-

mum[2] power reflected, $|\Gamma_{min}|^2$, [32, 33] through,

$$\Gamma_{min} = \frac{\beta - 1}{\beta + 1}, \tag{3.5a}$$

$$\Rightarrow \quad \beta = \frac{1 + \Gamma_{min}}{1 - \Gamma_{min}}. \tag{3.5b}$$

The cavity is undercoupled for $\beta < 1$, overcoupled for $\beta > 1$, and critically coupled at $\beta = 1$. Notice that the magnitude $|\Gamma_{min}|$ is the same for both $\beta = x$ and $\beta = 1/x$. The minimum value of Γ is at $\mu = 2\pi m$ regardless of ν and ρ, and is

$$\Gamma_{min} = -\frac{\sinh \frac{\nu - \rho}{2}}{\sinh \frac{\nu + \rho}{2}}. \tag{3.6}$$

Solving for the coupling coefficient gives,

$$\beta = \frac{\tanh \frac{\rho}{2}}{\tanh \frac{\nu}{2}} \approx \frac{\rho}{\nu}, \tag{3.7}$$

where the approximation holds for small ρ and ν, valid for most practical cavities.

3.3 Cavity Linewidth and Q

The cavity linewidth $2\delta\mu$, shown in figure 3.2, is found from equation (3.3c). The half-width $\delta\mu$ that sets (3.3c) equal to the average of its extreme values yields,

$$\delta\mu = \arccos \left[\frac{1}{\cosh(\nu + \rho)} \right], \tag{3.8a}$$

$$= \mathrm{gd}(\nu + \rho) \approx \nu + \rho. \tag{3.8b}$$

In proceeding from (3.8a) to (3.8b), we have taken advantage of the identity $\cos[\mathrm{gd}(x)] = 1/\cosh x$, where $\mathrm{gd}(x) = \int_0^x \frac{d\zeta}{\cosh \zeta}$ is the Gudermannian function [34].

---

[2]Γ_{min} is negative for undercoupled cavities. There is also a subtlety here regarding where the reference plane for Γ is chosen - the difference between the detuned-open and detuned-short planes discussed in [32]. The form of the equation presented here is chosen to be consistent with the convention employed throughout.
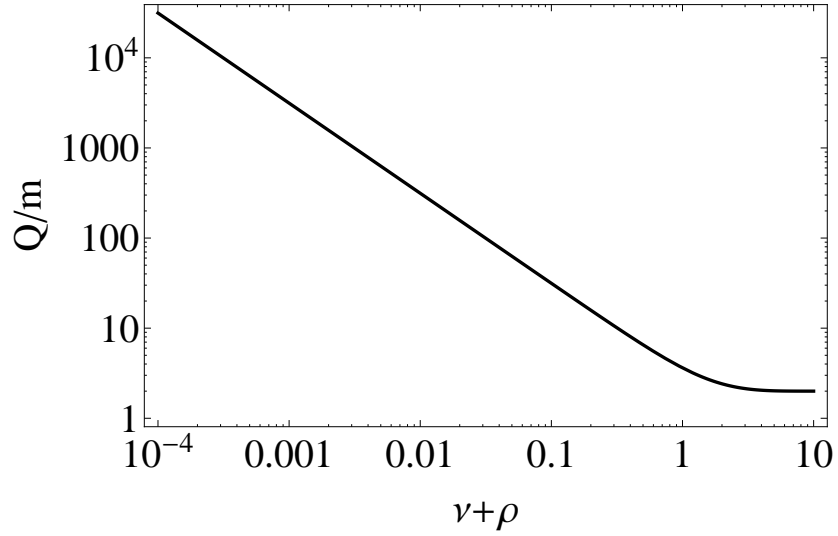
Figure 3.3: A log-log plot of the inverse halfwidth as a function of $\nu + \rho$, the cavity round-trip losses.

Let $\mu_m = 2\pi m$ be the center of the resonance of interest. In the limit that $\delta\mu \ll 2\pi$ (always true for practical cavities), the cavity Q can be calculated,

$$Q = \frac{\mu_m}{2\delta\mu} = \frac{\pi m}{\nu + \rho} \qquad \text{for } \delta\mu \ll 2\pi. \tag{3.9}$$

The Q calculated here is what is commonly referred to as the "loaded Q," Q_L, and can be easily decomposed into the "unloaded Q," $Q_U = \pi m/\nu$, and the "external Q," $Q_E = \pi m/\rho$. From the form of equation (3.9), it is clear that the Q's obey the relation $1/Q_L = 1/Q_U + 1/Q_E$. Q/m as a function of $\nu + \rho$ is plotted in figure 3.3.

In terms of $\delta\mu$ and β, Γ can be written,

$$|\Gamma|^2 = 1 - \frac{4\beta}{(1+\beta)^2} \frac{1}{1 + \sin^2\frac{\mu}{2}(\cot^2\frac{\delta\mu}{2} - 1)}, \qquad (3.10\text{a})$$

$$\approx 1 - \frac{4\beta}{(1+\beta)^2} \frac{1}{1 + \frac{(\mu-\mu_m)^2}{\delta\mu^2}}, \qquad (3.10\text{b})$$

$$= 1 - \frac{4\beta}{(1+\beta)^2 + \frac{(\mu-\mu_m)^2}{\delta\mu_0^2}} \qquad (3.10\text{c})$$

where the second and third lines holds for $\delta\mu \ll 2\pi$ and $(\mu - \mu_m) \ll 2\pi$. In the third line, the loaded halfwidth $\delta\mu$ has been replaced with the halfwidth of the unloaded resonator $\delta\mu_0 = \nu$. In this form, it is clear that the cavity mode curves can be described by the universal curves shown in figure 3.4a.

(a) Reflected Power

(b) Phase of Reflection
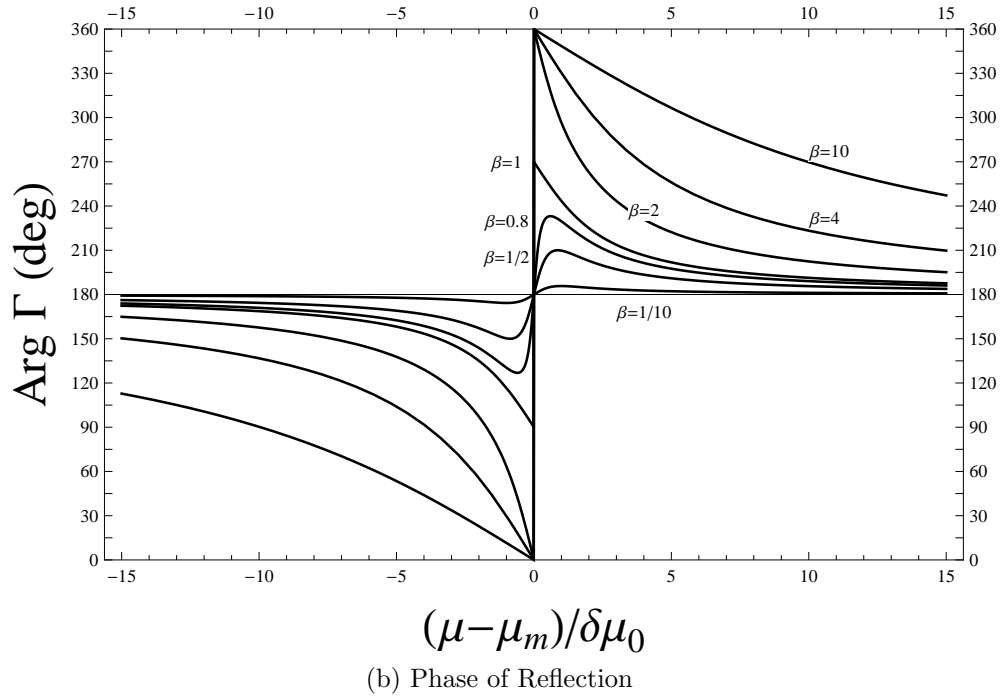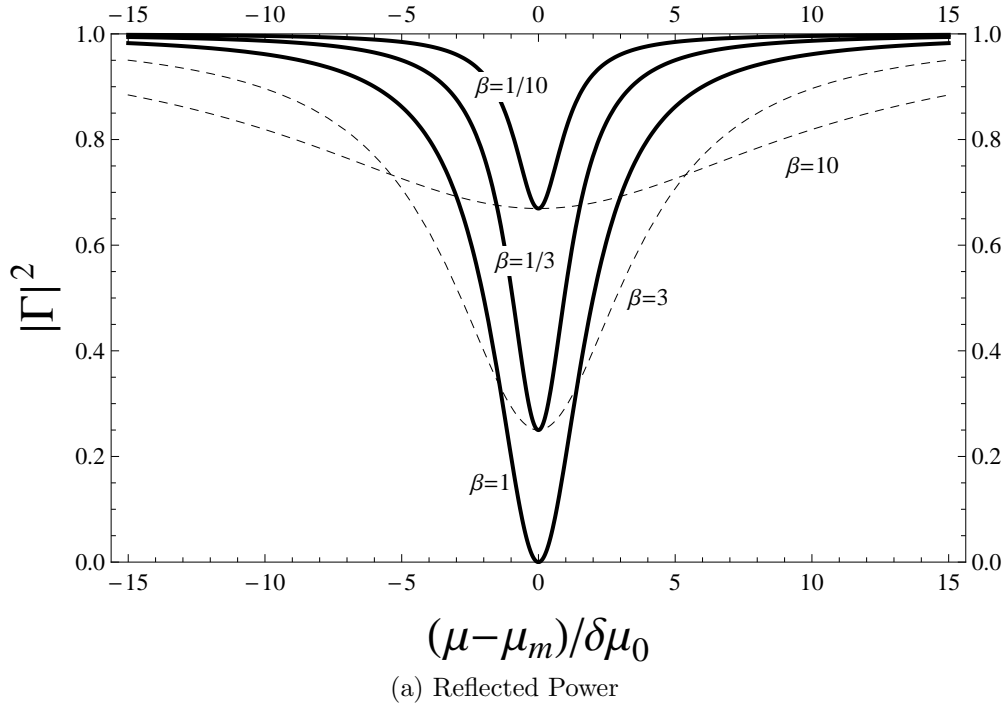
Figure 3.4: Universal curves describing the reflection coefficient of high Q resonators. The family of curves defined by equation (3.10c) is shown in panel (a) for various coupling coefficients. The phase of the reflection coefficient is plotted in panel (b). The abscissa of both panels is scaled by the halfwidth of the unloaded resonator.

16

3.4 Geometric Understanding of Γ

Γ can be written in the form of a circle in the complex plane,

$$|\Gamma - Z|^2 = R^2, \tag{3.11}$$

where Z is the location of the center of the circle and R is its radius, as shown in figure 3.5. Z and R are both real, and given by,

$$Z = -\frac{\sinh \nu}{\sinh(\nu + \rho)}, \tag{3.12a}$$

$$R = \frac{\sinh \rho}{\sinh(\nu + \rho)}. \tag{3.12b}$$

When $\beta = 1$, $Z = R$ and the circle intersects the origin. For $\beta < 1$, the circle excludes the origin, and for $\beta > 1$ the circle includes the origin. Polar plots of Γ are shown in figure 3.6 for various β in the limit of high Q. A Smith chart is superimposed over the background for reference.

The Γ circle is parametrized by μ. As μ varies over a period of 2π, the entire circle is traced, going slowly when away from resonance and quickly when near it. Figure 3.7 shows the position of Γ along the circle for 0.25 increments of $(\mu - \mu_m)/\delta\mu_0$. It is clear that $d\Gamma/d\mu$ is largest near the resonance. In terms of the angle spanning the circle, θ, Γ can be written,

$$\Gamma = Re^{i\theta} + Z, \tag{3.13}$$

where θ is measured counterclockwise from the positive real axis and is given by,

$$\sin \theta = \frac{\sinh(\nu + \rho) \sin \mu}{\cosh(\nu + \rho) - \cos \mu}, \tag{3.14a}$$

$$\cos \theta = \frac{-1 + \cosh(\nu + \rho) \cos \mu}{\cosh(\nu + \rho) - \cos \mu}. \tag{3.14b}$$

17
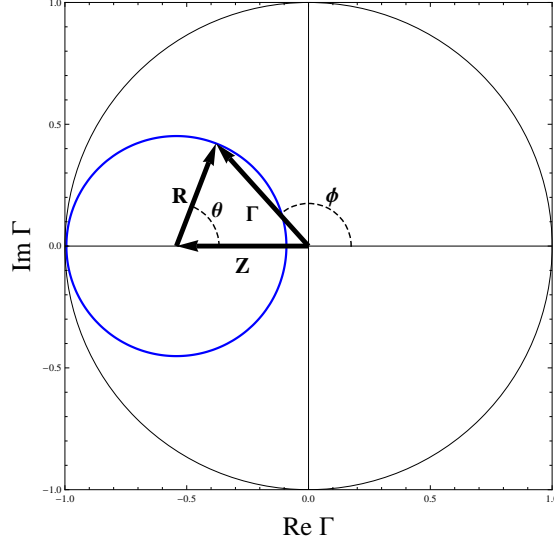
Figure 3.5: The geometric relationship between Γ, Z, and R is shown. As μ goes through a full 2π period, Γ traces a circle in the complex plane, whose center and radius are Z and R respectively. The figure is drawn for $\nu = 0.12$ and $\rho = 0.1$.
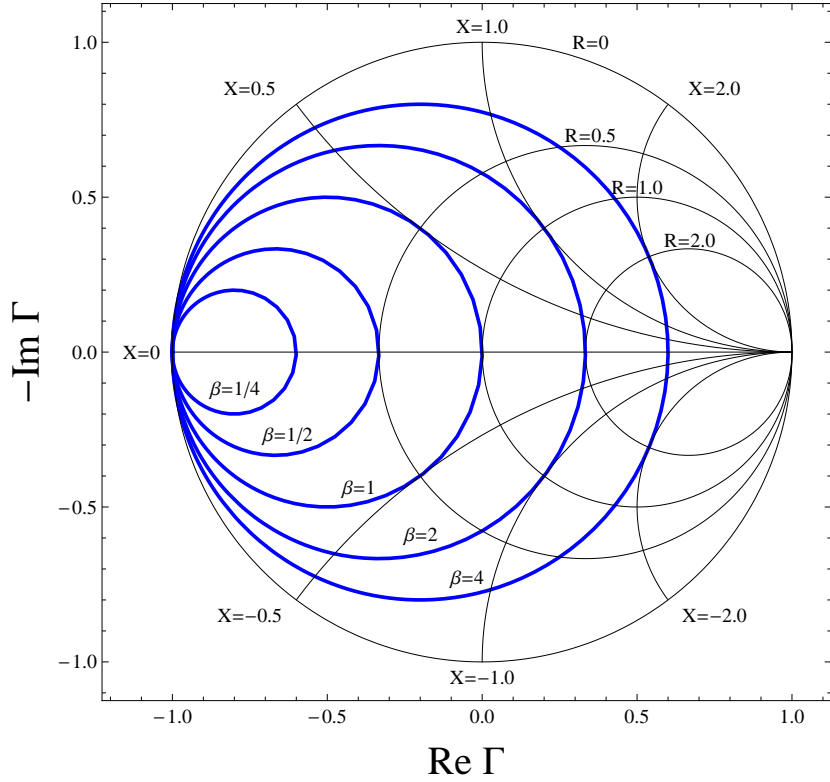
Figure 3.6: A Smith plot of the complex reflection coefficient Γ is shown for various β. When the cavity is undercoupled ($\beta < 1$), the circle does not enclose the origin. When the cavity is overcoupled ($\beta > 1$) the circle includes the origin. R and X refer to the normalized constant resistance and reactance circles respectively.

18
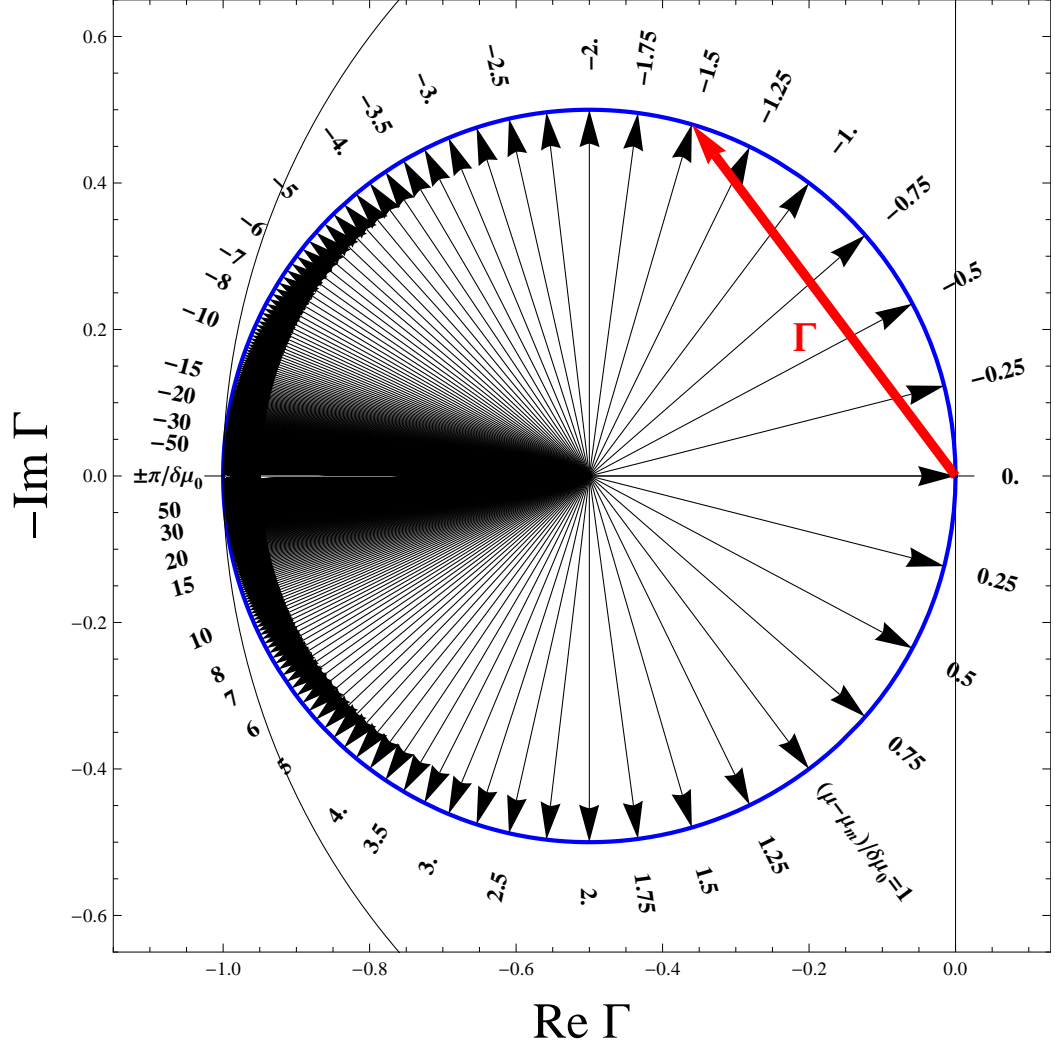
Figure 3.7: The reflection coefficient Γ traces a circle in the complex plane. As the frequency-like parameter μ varies over a 2π range, the entire circle is traced. The arrows drawn from the center of the circle to its circumference are drawn for increments of $(\mu - \mu_m)/\delta\mu_0$ equal to 0.25, showing that $\frac{d\Gamma}{d\mu}$ is largest near the resonance and is negligible away from it. The figure is drawn for $\beta = 1$ in the high Q limit.

3.5 Changes in Γ

Many types of cavity perturbation measurements are closely related to how Γ changes near a resonance. Changes in μ due to a change in the source frequency, the cavity length, or the dielectric constant of a sample inserted in the cavity result in a change in the imaginary part of Γ. Changes in ν due to a change in losses within the cavity (the presence of an electron paramagnetic absorption line for example), result in a change of the real part of Γ. The relevant change in Γ is determined by the intended application. In general, the change in Γ near a resonance due to a change in μ or ν, $\Delta\mu$ or $\Delta\nu$ respectively, is given by[3],

$$\Delta\Gamma|_{\mu_m} = \frac{\sinh\rho}{\cosh(\rho+\nu)-1}(-\Delta\nu + i\Delta\mu), \tag{3.15a}$$

$$= \frac{2\beta}{(1+\beta)^2 \tan\delta\mu_0}(-\Delta\nu + i\Delta\mu). \tag{3.15b}$$

Critical coupling and high Q maximizes $\delta\Gamma|_{\mu_m}$ yielding,

$$\Delta\Gamma|_{\mu_m,\beta=1} = \frac{-\Delta\nu + i\Delta\mu}{2\delta\mu_0}. \tag{3.16}$$

The sensitivity to changes in both ν and μ is proportional to the inverse linewidth.

---

[3]In proceeding from the first to the second line, the identity $\sinh \mathrm{gd}^{-1} x = \tan x$ is used.

CHAPTER 4

μPDH Theory

The theory of an optical Pound-Drever-Hall lock is well-known and understood [28–30] and most of the theory applies just as well for the microwave version. This chapter is not about rehashing old results, but rather about 1) understanding the differences that may arise due to the wildly different devices used, 2) understanding how the PDH signal may be altered when its setup is not ideal, and 3) taking advantage of the flexibility of the components available at microwave frequencies to develop techniques that may have been far more difficult at optical frequencies.

The primary purpose of the re-derivation of the PDH signal is to understand the effects that come in due to imperfections in the setup. For example, how does the (non)-orthogonality of the carrier and sum of the sidebands affect the signal? What is the expected effect of any amplitude imbalance between the sidebands? What can be gained from being able to fine-tune the cavity coupling, and how good does the coupling have to be? How much larger than the cavity linewidth does the modulation frequency have to be? To answer these questions, a more general derivation of the PDH signal is needed than is usually encountered.

The chapter concludes with introducing a single-sideband lock and quadrature detection technique that is easily implemented with the μPDH bridge. Both the real and imaginary components of the cavity reflection coefficient can be simultaneously measured, which has applications beyond a simple cavity lock.

4.1 Modulation: Phase and Vector

4.1.1 Phase modulation

The mathematical form of a wave undergoing phase modulation is,

$$E = E_0 e^{i(\omega t + z \cos \Omega t)}, \tag{4.1a}$$

$$= E_0 e^{i\omega t} \sum_{p=-\infty}^{\infty} i^p J_p(z) e^{ip\Omega t}, \tag{4.1b}$$

where z is the modulation depth and Ω is the modulation frequency. In the second line the exponential has been expanded in terms of Bessel functions, J_n [35]. Typically, z is chosen so that most of the power is in the carrier and the two nearest sidebands, and the higher order terms ignored. Graphically, phase modulation can be represented as in figure 4.1. In a phasor diagram of the field in the frame rotating at the carrier frequency ωt, the (unmodulated) carrier appears as a constant vector along the real (in-phase) axis. A pure phase modulation is represented by sweeping this constant-magnitude vector back and forth through an angle of $\pm z$, or sweeping the black vector over the blue region in the figure.

4.1.2 Vector modulation

Vector modulation is a Cartesian form of modulation. The incoming wave is passed through a 90° splitter, to form "in-phase" ($\boldsymbol{I}$) and "quadrature" ($\boldsymbol{Q}$) components. The amplitude of each component is independently adjusted proportionally to its baseband input, resulting in amplitudes "A" and "B" in each arm respectively. The waves are added back together, and can sum to produce any vector in the phasor plot. Written mathematically, with $\tan \phi = B/A$, the wave
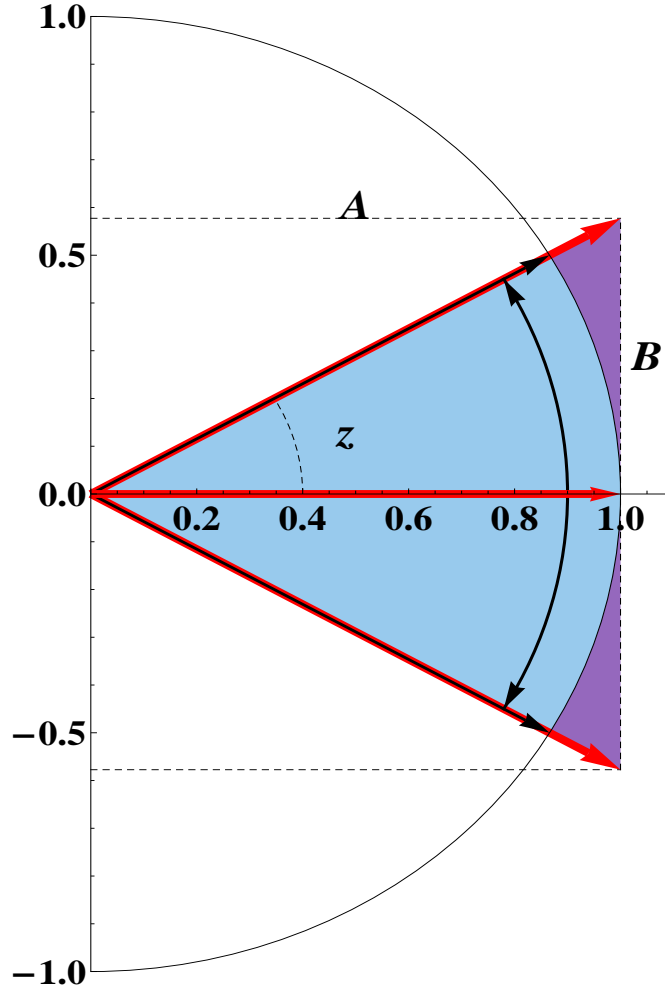
Figure 4.1: Comparison of phase and vector modulation for the purposes of a µPDH lock. The red vector along the real axis represents the instantaneous signal. If it is phase modulated, its length remains fixed, as it sweeps over the blue region. The Fourier decomposition of this motion contains every harmonic of the modulation frequency, with relative amplitudes determined by the modulation depth. If the signal is vector modulated by keeping A constant and varying B sinusoidally, its projection along the real axis remains constant, as its projection on the imaginary axis varies sinusoidally. The resultant vector sweeps over the combined blue and purple areas. The Fourier decomposition of the vector modulated motion contains only the carrier and two sidebands at plus and minus the modulation frequency.
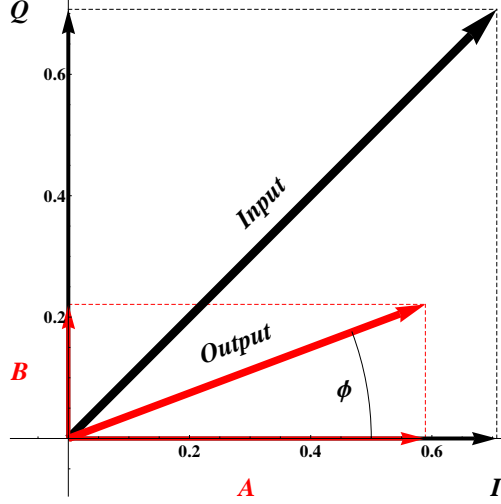
Figure 4.2: A vector modulator splits the input signal into quadrature components (black vectors). The amplitude of both components is independently adjusted, resulting in amplitudes A and B respectively. The two components are then recombined to produce the output (red vectors).

exiting the vector modulator is,

$$E_0 e^{-i\pi/4} e^{i\omega t} \rightarrow E_0 e^{i\omega t}(A - iB) \tag{4.2a}$$

$$= E_0 \sqrt{A^2 + B^2}\; e^{i(\omega t - \phi)}. \tag{4.2b}$$

The arbitrary phase $e^{-i\pi/4}$ is added to the input so that the $\boldsymbol{I}$ and $\boldsymbol{Q}$ channels lie on the real and imaginary axes. This is shown graphically in figure 4.2.

4.1.3 Creating a carrier and two sidebands for μPDH locking

To achieve optical PDH locking, lasers are usually phase modulated. As shown above, the result of a pure phase modulation contains the carrier, and sidebands at plus and minus all multiples of the modulation frequency whose amplitude ratio is determined by the modulation depth. With a vector modulator, a similar sideband structure can be created, with no higher order sidebands, and the additional freedom of being able to independently adjust the carrier and sideband powers. This can be done by fixing one channel, say the $\boldsymbol{I}$ channel ($A \rightarrow A_0$ is constant),
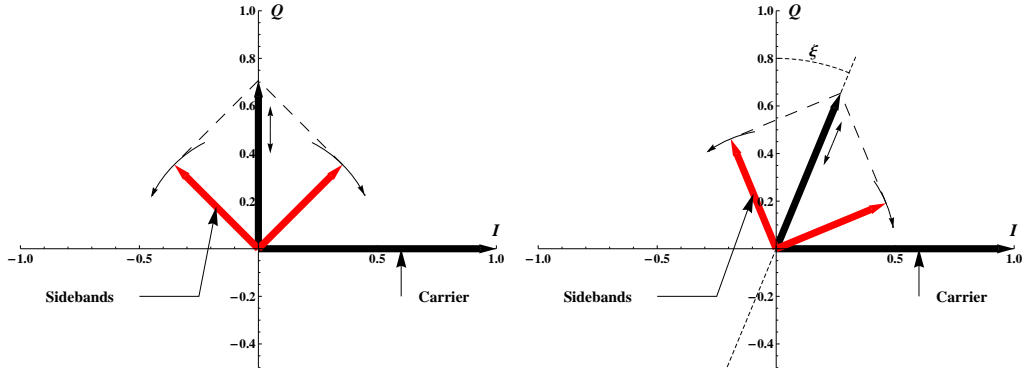
(a) Carrier orthogonal to sum of sidebands (b) Carrier at an arbitrary angle with the sum of the sidebands

Figure 4.3: Panel a) is a phasor diagram showing the carrier (black vector along I axis) and two counter-rotating sidebands (red vectors) as created by the vector modulator. The sum of the sidebands (black vector along Q axis) is always orthogonal to the carrier. Panel b) shows an angle ξ between the carrier and the sum of the sidebands created by applying a DC level to both channels. This angle should be zero for optimal performance of the µPDH signal, and would not intentionally be created. However, group delay effects can lead to it being non-zero, and the DC levels could be adjusted to cancel it out.

and sinusoidally varying the other channel ($B \rightarrow B_1 \sin \Omega t$). Mathematically,

$$E_0 \, e^{-i\pi/4} \, e^{i\omega t} \rightarrow E_0 \, e^{i\omega t}(A_0 - iB_1 \sin \Omega t) \tag{4.3a}$$

$$= E_0 \, e^{i\omega t}(A_0 - \frac{B_1}{2} e^{i\Omega t} + \frac{B_1}{2} e^{-i\Omega t}). \tag{4.3b}$$

The Q component oscillating along the imaginary axis can be written as a sum of two sidebands at frequencies of $\omega - \Omega$, and $\omega + \Omega$, and represented graphically as two oppositely-rotating vectors, as in figure 4.3a.

The difference between the vector modulated version and the pure phase modulation done in optics is depicted graphically in figure 4.1. The phase modulated signal sweeps only the blue area, whereas the vector modulated signal sweeps the combined blue and purple regions.

In the example just given, the carrier and sum of the two sidebands are always orthogonal, figure 4.3a. If, however, DC voltage is applied instead to the same

channel as the modulation, the carrier and the sum of the sidebands will be parallel. Further, by applying a DC voltage to both channels (with modulation only on one), any angle between the carrier and sum of the sidebands can be produced, as in figure 4.3b. Explicitly, $A \to A_0$ and $B \to B_0 + B_1 \sin(\Omega t + \psi)$, and

$$E_0 \, e^{-i\pi/4} \, e^{i\omega t} \to E_0 \, e^{i\omega t}(A_0 - i(B_0 + B_1 \sin(\Omega t + \psi))) \tag{4.4a}$$

$$= E_0 \, e^{i\omega t} \left(\sqrt{A_0^2 + B_0^2} e^{-i\xi} - \frac{B_1}{2} \, e^{i(\Omega t + \psi)} + \frac{B_1}{2} \, e^{-i(\Omega t + \psi)} \right), \text{ or} \tag{4.4b}$$

$$E_0 \, e^{-i\pi/4} e^{i\xi} \, e^{i\omega t} \to E_0 \, e^{i\omega t} \left(\sqrt{A_0^2 + B_0^2} - \frac{B_1}{2} \, e^{i\xi} \, e^{i(\Omega t + \psi)} + \frac{B_1}{2} \, e^{i\xi} \, e^{-i(\Omega t + \psi)} \right), \tag{4.4c}$$

where the last line follows by pulling an extra phase $\xi = \arctan B_0/A_0$ into the input so that the output is rotated to keep the carrier lying on the real axis. The phase shift ψ defines the relative "turn on time" of the modulation and the carrier. The sum of the sidebands is orthogonal to the carrier for all values of ψ. The phases ξ and ψ are helpful in understanding the effects of group delay, and how to combat them, discussed in section 4.3.

4.1.4 Creating a carrier and a single sideband

For single sideband locking, a carrier and a single sideband is needed. This can be generated by applying an out-of-phase sinusoidal signal to the **I** and **Q** channels, and applying a DC offset to one or both. The simplest example is $A \to A_0 + B_1 \cos \Omega t$ and $B \to \pm B_1 \sin \Omega t$. Then,

$$E_0 \, e^{-i\pi/4} \, e^{i\omega t} \to E_0 \, e^{i\omega t}(A_0 + B_1 \cos \Omega t \mp i B_1 \sin \Omega t) \tag{4.5a}$$

$$= E_0 \, e^{i\omega t}(A_0 + B_1 \, e^{\mp i\Omega t}), \tag{4.5b}$$

so the upper or lower sideband can be selected by setting the I and Q channels $\pm 90°$ out of phase.

4.2   The Microwave Pound-Drever-Hall Signal

After modulation of the type described in section 4.1.3, the microwave incident on the cavity has the form given by (4.4c),

$$E = \Re[E_c e^{i\omega t} + E_s e^{i\xi} e^{i(\omega+\Omega)t} - E_s e^{i\xi} e^{i(\omega-\Omega)t}]. \tag{4.6}$$

Here, the amplitude constants in (4.4c) have been renamed to E_c, the amplitude of the carrier, and E_s, the amplitude of the two sidebands. The arbitrary phase ψ has been set to π. As before, ω is the carrier frequency, Ω is the modulation frequency, and ξ is the angle between the carrier and the oscillating-sum of the sidebands, as depicted in figure 4.3b. It will be shown that ξ determines the amount that the real part of the reflection coefficient influences the µPDH signal. After reflection off the cavity, each term is multiplied by the corresponding reflection coefficient, given by (3.3a), and the field incident on the diode has the form,

$$E = \Re[E_c \Gamma_0 e^{i\omega t} + E_s \Gamma_+ e^{i\xi} e^{i(\omega+\Omega)t} - E_s \Gamma_- e^{i\xi} e^{i(\omega-\Omega)t}], \tag{4.7}$$

where Γ_0, Γ_+, and Γ_- are shorthand for $\Gamma(\omega)$, $\Gamma(\omega+\Omega)$, and $\Gamma(\omega-\Omega)$ respectively. The time-averaged (over period $2\pi/\omega$) power incident on the diode is,

$$P \sim \frac{1}{2}(E_c^2|\Gamma_0|^2 + E_s^2|\Gamma_+|^2 + E_s^2|\Gamma_-|^2)$$

$$-E_s^2[(\Im[\Gamma_+]\Im[\Gamma_-] + \Re[\Gamma_+]\Re[\Gamma_-])\cos(2\Omega t)$$

$$+(\Re[\Gamma_+]\Im[\Gamma_-] - \Re[\Gamma_-]\Im[\Gamma_+])\sin(2\Omega t)]$$

$$+E_c E_s[(\Re[\Gamma_0]\Re[\Gamma_+] + \Im[\Gamma_0]\Im[\Gamma_+])\cos(\Omega t + \xi) \quad (4.8)$$

$$-(\Re[\Gamma_0]\Re[\Gamma_-] + \Im[\Gamma_0]\Im[\Gamma_-])\cos(\Omega t - \xi)$$

$$+(\Re[\Gamma_+]\Im[\Gamma_0] - \Re[\Gamma_0]\Im[\Gamma_+])\sin(\Omega t + \xi)$$

$$+(\Re[\Gamma_-]\Im[\Gamma_0] - \Re[\Gamma_0]\Im[\Gamma_-])\sin(\Omega t - \xi)].$$

The constant terms serve as a bias. The 2Ω terms are the result of the beating of the two sidebands, and contain no information about the reflection coefficient of the carrier. The terms that oscillate at Ω are the result of the sidebands beating with the carrier, and contain the relevant information. The diode operates in the square-law regime and produces a signal proportional to the incident power. This signal is mixed with a local oscillator of the form $\sin(\Omega t + \theta)$ and low-pass filtered to produce the µPDH signal ϵ,

$$\frac{2\epsilon}{E_c E_s} \sim \cos\theta\cos\xi\Big[-\Re[\Gamma_0](\Im[\Gamma_+] + \Im[\Gamma_-]) + \Im[\Gamma_0](\Re[\Gamma_+] + \Re[\Gamma_-])\Big] \quad (4.9a)$$

$$-\cos\theta\sin\xi\Big[\Re[\Gamma_0](\Re[\Gamma_+] + \Re[\Gamma_-]) + \Im[\Gamma_0](\Im[\Gamma_+] + \Im[\Gamma_-])\Big] \quad (4.9b)$$

$$+\sin\theta\cos\xi\Big[\Re[\Gamma_0](\Re[\Gamma_+] - \Re[\Gamma_-]) + \Im[\Gamma_0](\Im[\Gamma_+] - \Im[\Gamma_-])\Big] \quad (4.9c)$$

$$+\sin\theta\sin\xi\Big[-\Re[\Gamma_0](\Im[\Gamma_+] - \Im[\Gamma_-]) + \Im[\Gamma_0](\Re[\Gamma_+] - \Re[\Gamma_-])\Big]. \quad (4.9d)$$

The µPDH signal is contained in the second term of line 4.9a. In a perfect setup, θ and ξ are both be zero, and the modulation is much larger than the cavity width, so that $\Im[\Gamma_+] = \Im[\Gamma_-] = 0$, $\Re[\Gamma_+] = \Re[\Gamma_-] = -1$. The ideal signal is,

$$\frac{\epsilon}{E_c E_s} \sim \Im[\Gamma_0] = 4Q\frac{\beta}{1+\beta}\frac{f - f_0}{f_0}, \quad (4.10)$$

where β is the coupling coefficient defined in section 3.2.

All of the remaining terms are relatively small, and can be arranged by order of smallness in terms of perturbations from the ideal setup. Some are worth pointing out, and their magnitude is estimated in section 11.5.

The first term in line (4.9b) should be a concern. It puts a restriction on the (non)-orthogonality of the carrier and sum of the sidebands (angle ξ) that depends on the degree of coupling,

$$\Re[\Gamma_0] \sin \xi \ll \Im[\Gamma_0] \cos \xi. \tag{4.11}$$

The presence of this term would very clearly lead to locking off resonance, even if θ is properly set.

The first term in line (4.9c) shows how the cavity coupling and sideband imbalance (different amplitude) would combine to produce an out-of-phase signal in θ, which would result in locking off resonance if θ is not set properly. Although the sideband imbalance is not written explicitly, it is straightforward to see that it would result in different coefficients on the $\Re[\Gamma_+]$ and $\Re[\Gamma_-]$ terms, whose difference would then be non-zero. This effect is clearly visible and is discussed in section 5.3.

Similarly, the first term in line (4.9a) should be zero if the sidebands are balanced because $\Im[\Gamma_+] = -\Im[\Gamma_-]$, however any imbalance would lead to a finite value. Fortunately, $\Im[\Gamma_\pm]$ can be made smaller by increasing the modulation frequency. As the amplitude of this term depends on three parameters not being proper (sideband imbalance, cavity coupling, and modulation frequency), it can safely be ignored.

4.3 The Effects of Group Delay

In practice, after the microwave has been properly modulated, it passes through one or more filters, discussed in section 5.2. However, filtering may have the unintentional effect of changing the relative phase of the sidebands and carrier. Other devices, such as waveguides, circulators, or amplifiers, may do the same. In this section, group delay effects on the µPDH signal are investigated.

Any linear, time-invariant system can be characterized by its transfer function, defined to be the ratio of the output to the input,

$$H(i\omega) = \frac{V_{out}(i\omega)}{V_{in}(i\omega)} = |H(i\omega)|e^{i\phi(\omega)}, \tag{4.12}$$

where $\phi(\omega) = \arg H$. The group delay is then defined,

$$\tau = -\frac{d\phi(\omega)}{d\omega}. \tag{4.13}$$

Let τ_0, τ_+, τ_- be the group delays for the carrier and the plus and minus sidebands respectively. Equation (4.6), in the more general form given by (4.4c), is then modified to,

$$E = \Re[E_c e^{i\omega(t+\tau_0)} + E_s e^{i\xi} e^{i\psi} e^{i(\omega+\Omega)(t+\tau_+)} - E_s e^{i\xi} e^{-i\psi} e^{i(\omega-\Omega)(t+\tau_-)}]. \tag{4.14}$$

Recall that ξ describes the angle between the carrier and the sum of the two sidebands, defined in figure 4.3b, which should as close to zero as possible. The arbitrary phase ψ describes the "turn on time" of the modulation relative to the carrier, as discussed in section 4.1.3. Note that with the transformations,

$$\xi' = \xi - \omega\tau_0 + \frac{\omega}{2}(\tau_+ + \tau_-) + \frac{\Omega}{2}(\tau_+ - \tau_-), \quad \text{and} \tag{4.15a}$$

$$\psi' = \psi - \Omega\tau_0 + \frac{\omega}{2}(\tau_+ - \tau_-) + \frac{\Omega}{2}(\tau_+ + \tau_-), \tag{4.15b}$$

equation (4.14) can be re-written as,

$$E = \Re[E_c e^{i\omega(t+\tau_0)} + E_s e^{i\xi'} e^{i\psi'} e^{i(\omega+\Omega)(t+\tau_0)} - E_s e^{i\xi'} e^{-i\psi'} e^{i(\omega-\Omega)(t+\tau_0)}]. \qquad (4.16)$$

In other words, the presence of group delay cause three effects, two which do not affect the signal, and one that does. First, $t \to t + \tau_0$, is an arbitrary overall time shift. Second, there is a shift in the apparent time the modulation was turned on relative to the carrier. These two effects are inconsequential. The third effect is to change the angle between the sum of the sidebands and the carrier. This is more serious because the µPDH signal becomes sensitive to the real part of the reflection coefficient. When considering the effects of group delay, the quantity that must be minimized to preserve a clean µPDH signal is given by (4.15a).

4.4 Single Sideband Lock

The ability to easily create a single sideband with a vector modulator opens the possibility of a quadrature detection scheme that measures both the real and imaginary components of the cavity reflection coefficient. As the sideband rotates around the phasor diagram, its projection along the real and imaginary axes samples the two components of the reflection coefficient. With the proper two-channel lock-in detection, both signals can be recovered.

Modulation is performed as described in section 4.1.4. The microwave incident on the cavity has the form given by equation (4.5b),

$$E = \Re\left[e^{i\omega t}\left(E_c + E_s e^{\mp i\Omega t}\right)\right]. \qquad (4.17)$$

After reflection off the cavity, each term is multiplied by the corresponding am-

plitude reflection coefficient, and the field incident on the diode has the form,

$$E = \Re\left[e^{i\omega t}\left(E_c\Gamma_0 + E_s\Gamma_{\mp}e^{\mp i\Omega t}\right)\right]. \tag{4.18}$$

For simplicity, set $\Gamma_{\mp} = -1$. The low frequency components of the power incident on the diode are,

$$P \sim \frac{1}{2}[E_c^2|\Gamma_0|^2 + E_s^2 - 2E_cE_s(\Re[\Gamma_0]\cos\Omega t \mp \Im[\Gamma_0]\sin\Omega t)]. \tag{4.19}$$

The diode signal is taken to a quadrature mixer. All that remains is to set the phase of the local oscillator so that the real and imaginary parts of the reflection coefficient are separated. This is done in several steps. Recall that the single sideband modulation was done by driving the I baseband input with signal of the form $A_0 + A_1\cos\Omega t$ and the Q baseband input with a signal of the form $\pm B_1\sin\Omega t$, with $A_1 = B_1$. To choose the proper LO phase, A_1 is temporarily set to zero. The modulation is now the same as for the standard double sideband PDH. The LO phase is chosen so that one of the quadrature mixer outputs is maximized, and the other is zero. The channel with the maximum signal is now sensitive to $\Im[\Gamma_0]$. When A_1 is set equal to B_1 again, this channel will remain sensitive to $\Im[\Gamma_0]$, and the second channel will be sensitive to $\Re[\Gamma_0]$. The $\Im[\Gamma_0]$ signal can be used the same as before, for locking or other purposes. The $\Re[\Gamma_0]$ signal contains information which was not available in the double sideband version.

This can be used, for example, to detect the quadrature EPR signal - the dispersion and absorption signals will be cleanly separated. In the frequency-encoder application, the quadrature signal would detect changes in the cavity coupling, which could be used in a second feedback loop to keep the coupling fixed if needed.

CHAPTER 5

µPDH Hardware and Performance

The previous two chapters dealt with the theoretical understanding of generating the µPDH signal. In this chapter the hardware needed to implement it is described in detail. The performance is evaluated and a comprehensive analysis of the system signal-to-noise ratio is given.

5.1 Simplified Hardware Description

When properly set up, the µPDH error signal is given by equation (4.10),

$$\epsilon \propto \sqrt{P_c P_s} Q \frac{\delta f}{f}, \tag{5.1}$$

where ϵ is the response to a frequency difference of δf, Q is the cavity quality factor, and P_c and P_s are the carrier and sideband power respectively. As fluctuations in the length of the cavity are indistinguishable from frequency fluctuations of the source, the cavity cannot be stabilized better than the frequency noise of the source. Thus, the best-case noise spectrum of the µPDH signal in $Hz\ Hz^{-1/2}$ is the frequency noise spectrum of the source. Optimal microwave circuit design comes down to minimizing the close-in frequency noise of the transmitter, and setting the carrier and sideband power levels high enough to make the receiver noise negligible in comparison.

33
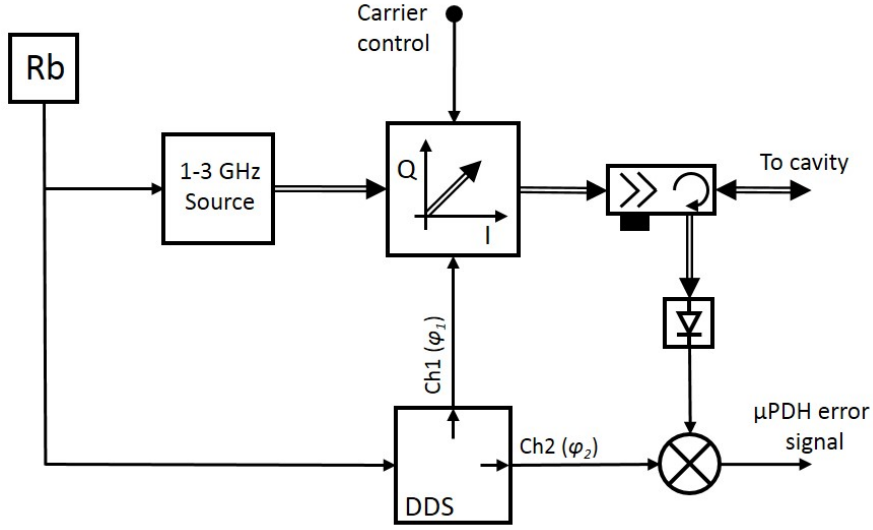
Figure 5.1: Block diagram of a 1-3 GHz µPDH lock.

5.1.1 Low frequency µPDH lock

A block diagram of a low-frequency (1-3 GHz) µPDH bridge is shown in figure
5.1. For clarity, only functionally relevant parts are shown. Other units, such as
amplifiers, attenuators, isolators, filters etc. whose only purpose is to set signal
levels and conditions assuring the optimal functioning of the interconnected, are
not shown. In general, these elements need to be set to match the characteristics
of the actually selected components during the construction process.

All signal sources are phase-locked directly, or via multiple stages, to a 10 MHz
Rubidium clock reference assuring 11 digit absolute precision. The working fre-
quency of the source can be chosen in the 1-3 GHz frequency range by selecting
a low phase-noise VCO from the many models available on the market. Avail-
able single chip DDS/PLL units (such as the AD9540 or AD9956) can be used
to turn the selected VCO into a digitally controllable (with 48 bit precision) mi-
crowave source that is phase locked to the Rubidium clock and has better than
1 Hz stability.

The selected source serves as a local oscillator for an IQ (or vector) modula-

tor. Proper modulation of the source is critical in implementing the μPDH lock. Thanks to advances in the cell-phone industry, recently available IQ modulators have become excellent devices to modulate 1-3 GHz microwaves. Chips such as the AD8340/1 or AD5373, for example, can produce arbitrary modulation at up to 230 MHz. For a μPDH lock, the carrier modulation needs to be as described in section 4.1.3 - the sum of the sidebands should be orthogonal to the carrier so that the error signal is sensitive to the carrier's phase and not its amplitude. This can be accomplished by applying a DC voltage to one of the baseband inputs (the I for instance) and driving the other (Q) at the modulation frequency f_1. The Fourier spectrum of such a signal contains only three lines, at f_0 and $f_0 \pm f_1$, with the required phase relationship. In addition, the power level of the carrier and sidebands is independently adjustable. This is in contrast to a phase-modulated spectrum, which contains all the harmonics of the modulation frequency with relative amplitude set by the modulation depth.

A dual channel DDS (see Analog Devices AD9958 for example) is used to generate the modulation frequency f_1 on both outputs, with independently adjustable phase (14 bit/ 0.02° resolution) and amplitude. The first output of the DDS is connected to the Q input of the IQ modulator through a bias-tee (not pictured) to generate the sidebands.

A four-port circulator directs the modulated microwave to the resonant cavity and diverts the back-reflected signal to a broadband, high-bandwidth detector diode. The signal-of-interest is carried in the diode output at frequency f_1. It is filtered (not pictured) and connected to the RF input of a mixer. The mixer LO signal is supplied by the second channel of the DDS, whose programmable phase-offset feature allows finding and precisely adjusting the LO phase for optimal detection of the f_1 component. The low-pass-filtered IF output is the μPDH error signal.
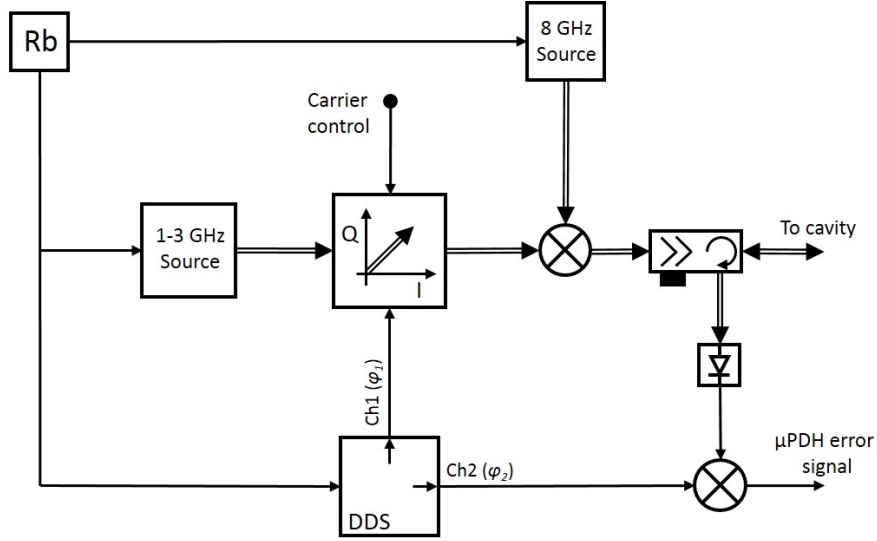
Figure 5.2: Block diagram of a higher-frequency (9-11 GHz pictured) µPDH lock. A fixed-frequency source is used to upconvert the modulated microwave generated by the transmitter portion of the low-frequency bridge of figure 5.1 to a higher frequency.

5.1.2 High frequency µPDH lock

10 GHz and higher frequency µPDH locks can be constructed with little modification to the basic architecture of the above-described low frequency bridge as shown in figure 5.2. A fixed-frequency source that is 1.5-2 GHz away from the intended working frequency is chosen based on its close-in phase noise characteristics. It is phase locked to the Rb reference and used as a local oscillator for an up-converting mixer, whose intermediate frequency input is connected to the signal generator part of the low-frequency bridge. The RF output of the mixer is amplified and filtered by a 1 GHz band-pass filter (not shown), removing everything but the working frequency and sidebands. The receiver arrangement is very much the same as for the low-frequency version.

Devices that can vector (or phase) modulate a 10 GHz signal directly (and at the required speed) exist, but amplitude and phase imbalances lead to asymmetric sidebands, resulting in an additional component to the diode signal that is at the

modulation frequency, but out-of-phase with the signal-of-interest. This leads to more stringent coupling requirements (see equation (4.9)) and complicates setting the optimal phase of the LO. At higher frequencies, the devices simply do not exist and up-conversion is the only option.

5.2 Detailed Description of a 10 GHz µPDH Bridge

Figure 5.3 shows a detailed block diagram of the realized 10 GHz µPDH bridge. The general layout is the same as that shown in figure 5.2, but figure 5.3 shows all of the components used in practice.

The transmitter intermediate frequency is provided by a ZCommunications ZRO1820A1LF 1.8 GHz VCO, chosen for its low free-running phase-noise[1]. An AD9956 evaluation board transforms this VCO into a digitally controllable frequency source, with its on-board DDS and PLL. Part of the VCO output (divided by 8) is used to clock the DDS, which generates a frequency that depends on its 48-bit tuning word. The DDS output is compared to a 100 MHz reference by the PLL, whose output voltage drives the VCO frequency. In this mode of operation, the DDS clock frequency is adjusted so that the DDS output is always 100 MHz. As the DDS output frequency can be set with 48-bit precision, the VCO frequency can be digitally set to a resolution of $\sim 10^{-5}$ Hz. For all practical purposes, the VCO can be considered continuously sweepable. The 100 MHz reference is generated by a Wenzel SC-cut crystal oscillator that is phase-locked to a Stanford Research Systems PRS10 Rubidium oscillator, providing long-term stability. Its output is taken to a distribution amplifier, and disseminates to several devices. As the VCO is phase-locked to the reference, the close-in phase noise of the reference should be as low as possible[2], or at least lower than the noise of the PLL

---

[1]The VCO free-running phase noise is specified as -88 dBc Hz^{-1} at 1 kHz offset, -118 dBc Hz^{-1} at 10 kHz offset, and -140 dBc Hz^{-1} at 100 kHz offset. Its tuning sensitivity is 5 MHz V^{-1}.

[2]The 100 MHz Wenzel oscillator is specified to have a free-runing phase noise of -

components. In this setup, the noise of the PLL is the limiting factor, discussed in section 5.5.

The VCO output is amplified and drives the RF input of an AD8341 vector modulator. 0-0.5 V DC is applied to the $\boldsymbol{I}$ baseband input to set the carrier level between -17 dBm and 14 dBm. The $\boldsymbol{Q}$ baseband input is driven by channel 1 of a two-channel DDS (AD9958), with both channels running at 77 MHz. The AD9958 is clocked by a 400 MHz signal generated by quadrupling one of the distribution amplifier outputs. The amplitude of each sideband incident on the cavity can be set from -18 dBm to 0 dBm in 6 dB steps, by setting the appropriate bits within the chip. In typical operation, the carrier is set to 12 dBm and the sidebands to -12 dBm.

A Wenzel 8 GHz multiplied crystal oscillator[3] (MXO) serves as an LO for the up-converting mixer whose IF input is driven by the output of the vector modulator. An isolator prevents back-reflections from the mixer from pulling the frequency of MXO. The RF output of the mixer contains the carrier at 9.8 GHz with the vector-modulation sidebands at 9.8 ± 0.077 GHz, as well as the analog at 6.2 GHz. A 9.5 ± 0.5 GHz bandpass filter passes rejects the set of lines around 6.2 GHz. A second isolator prevents back-reflections from the bandpass filter from reaching the mixer and creating unwanted mixing products. High mixer LO-RF isolation is important because the LO leakage generates a large spur in the microwave spectrum that mixes with broadband noise and can degrade the signal-to-noise ratio of the µPDH signal. Another approach is to choose a narrower bandpass filter that can remove the LO frequency as well.

---

125 dBc Hz^{-1} at 100 Hz offset, -155 dBc Hz^{-1} at 1 kHz offset, and -170 dBc Hz^{-1} at 10 kHz offset. It is phase locked to the Rubidium reference in a 100 Hz bandwidth.

[3]The 8 GHz Wenzel MXO is specified to have a free-running phase noise of -57 dBc Hz^{-1} at 10 Hz offset, -87 dBc Hz^{-1} at 100 Hz offset, -107 dBc Hz^{-1} at 1 kHz offset, and -129 dBc Hz^{-1} at 10 and 100 kHz offsets. It is phase locked to the 100 MHz Wenzel oscillator in a bandwidth of 60 Hz.

Figure 5.3: Detailed block diagram of the realized 10 GHz µPDH bridge.

The 9.8 GHz signal is then amplified and filtered by two band-reject (notch) filters at $9.8\pm2\times0.77$ GHz. The notch filters remove noise at the image frequencies of the carrier and are crucial in maximizing the signal-to-noise ratio. These should be narrow enough not to cause significant group delay effects at the sideband frequencies. This completes the description of the transmitter portion of the bridge.

A 4-port circulator, with high port-to-port isolation, directs the microwave to the cavity. Since better than 50 dB coupling is required (see below, figure 5.12), poor circulator port-to-port isolation would imply that by "critically coupling," the cavity reflection is actually adjusted to cancel out the circulator leakage. Changing the frequency would upset this balance, as would length changes in the path to the cavity.

A standard X-band rectangular TE_{102} EPR measuring cavity with a loaded Q of 2585, whose mode curve is shown in figure 5.4, is modified to make it tunable. The front-plate optical window was replaced with a structure where a rectangular piezo-block bends the cavity wall inward upon extension. Such modification allows about 150 kHz frequency tunability with a 5-10 kHz bandwidth - sufficient to hold the cavity in lock under normal laboratory conditions for several hours. A somewhat more elaborate solution takes advantage of the fact that the sample access of almost every commercial EPR cavity is a through hole of about 10-11 mm diameter. There is plenty of space for a metal cylinder, mounted on a mechanical stage, to be inserted into the cavity. Note that the metal cylinder will not penetrate into the cavity more than 1-2 mm; depending on the cavity mode, the corresponding frequency change is several hundred MHz. This metal cylinder is used as a rough, manual adjustment to set the cavity frequency near the source frequency, whereas the piezo is used to keep the cavity under lock.

The circulator directs the back-reflected microwave to a Herotek DDS218 biased Schottky diode detector operating in the square law regime, with a sensi-
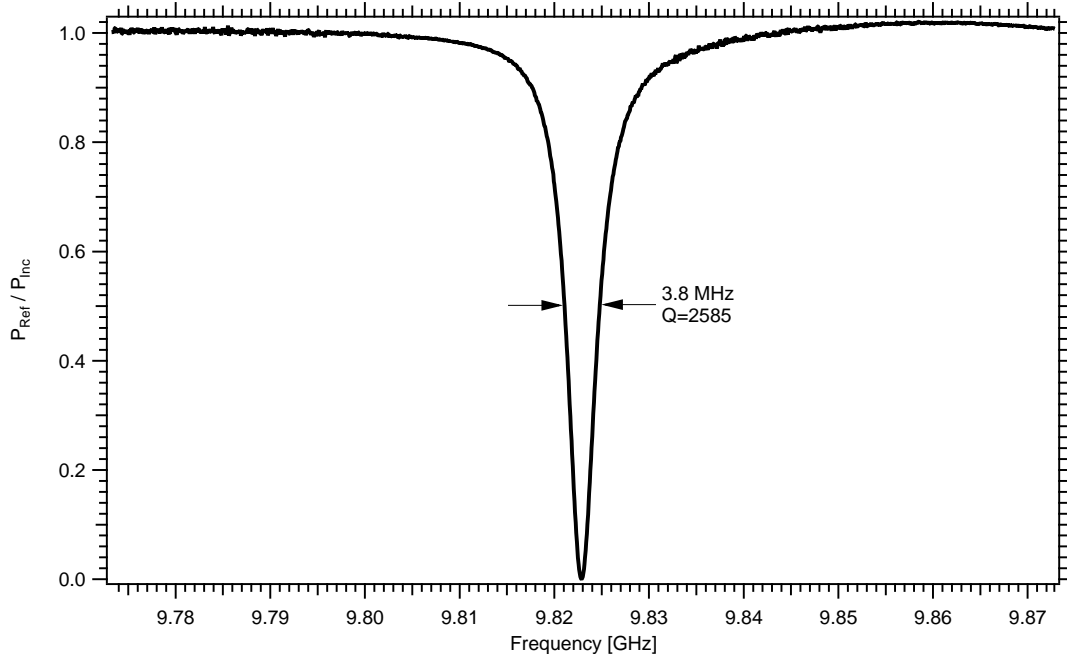
Figure 5.4: A standard TE_{102} EPR cavity mode curve. The cavity loaded Q is 2585 when critically coupled.

tivity of 2200 mV mW^{-1} and a Tangential Signal Sensitivity (TSS) of -53 dBm at 100 µA bias. The diode signal goes to a preamplifier and is filtered to remove the $2f_1$ beating of the sidebands before being taken to the RF input of the down-converting mixer. The LO is provided by the second channel of the AD9958, whose programmable-phase feature allows finding the proper phase setting. The low-pass filtered IF output is the µPDH signal which is optionally taken to a digital feedback controller that drives the tuning piezo of the cavity to lock it to the source.

5.3 Tuning the System

To set up the µPDH lock, the source and cavity are coarse adjusted to roughly the same frequency to within a couple hundred kilohertz. The cavity coupling is adjusted with the variable size iris between the waveguide and cavity that is built into most EPR cavities to achieve better than 50 dB coupling. Critical coupling

(a) Phase-Frequency Plot (b) µPDH signal sensitivity
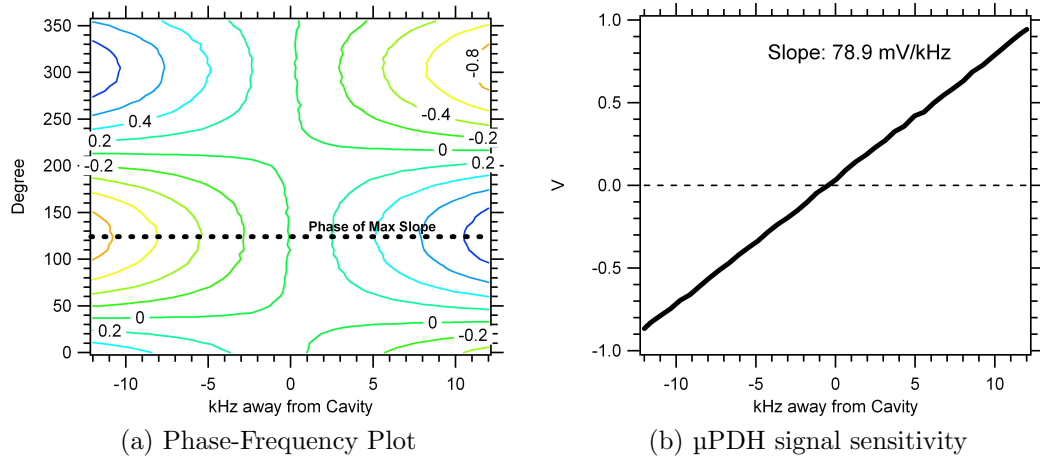
Figure 5.5: The optimal sensitivity of the µPDH signal requires that the phase of the local oscillator of the down-converting mixer is set properly. The error signal is shown in panel (a) as a function of microwave frequency and LO phase. The optimal phase is selected based on the largest change in error signal when the frequency is swept through the cavity. Panel (b) shows the cross section of the phase-frequency plot at the optimal phase. The slope of this line is the sensitivity of the µPDH signal.

has three purposes: 1) A large carrier power can be used without saturating the detector diode, 2) it cancels many noise sources to first order, and 3) it minimizes the effects of asymmetric sidebands.

The sideband power level is chosen to be as large as possible while keeping the diode operating in its square-law regime and without saturating the preamplifier with the large $2f_1$ beat signal. As the $2f_1$ beat signal is proportional to P_s and the f_1 signal is proportional to $\sqrt{P_c P_s}$, increasing P_s lowers the dynamic range of the detector. The carrier power is chosen high-enough so that the signal-to-noise ratio is limited by the source frequency noise.

For optimal performance of the system, the relative phase of the 2 channels of the DDS (the phase of the local oscillator of the down converting mixer) must be properly set. The optimal phase is found by repeatedly sweeping the VCO frequency over the cavity while stepping the phase of the LO. Plotting the error signal against the phase and VCO frequency results in the phase-frequency plot
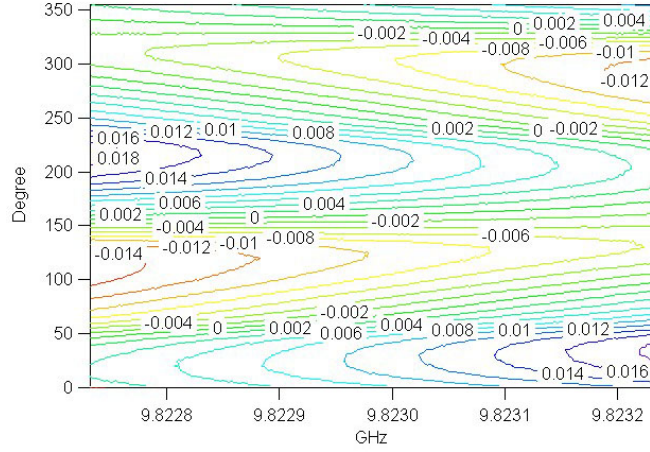
Figure 5.6: An erroneous signal is produced when the large-amplitude $2f_1$ component of the diode signal is mixed by the local oscillator at f_1, shown in the phase-frequency plot above. In generating the plot, the carrier power is purposely set low so that the artificial signal can be seen clearly. Adding a filter after the preamplifier removes this artificial signal.

in figure 5.5a. For every phase setting, the slope of the signal is calculated at its zero crossing. The phase resulting in the maximum slope is the optimal phase. The cross section of the phase-frequency plot at the optimal phase (dotted line in figure 5.5a) is shown in figure 5.5b. The slope of this line is the sensitivity of the instrument.

The filter after the preamplifier is necessary, because if the $2f_1$ component is not removed from the signal, an erroneous mixing signal is produced. Removing the filter and setting a low carrier power (so that the $2f_1$ component is the only harmonic reaching the mixer) results in the phase-frequency plot in figure 5.6. That signal is added to the wanted one when the filter is not in place.

The cavity coupling can also alter the appearance of the phase-frequency plot as shown in figure 5.7. As predicted by equation 4.9, this is caused by the sidebands having slightly different amplitude (about 0.5 dB asymmetry in this case). Panel 5.7a displays the complex cavity reflection coefficient plotted in various ways, showing that the cavity is overcoupled. Panel 5.7b is the phase-frequency plot measured while the cavity was in this overcoupled position. Figures 5.7c and

43

5.7d are the analog when the cavity is critically coupled, and figures 5.7e and 5.7f are the analog when the cavity is undercoupled. The red curves on the right side of figures 5.7b, 5.7d, and 5.7f are the cross sections of the phase-frequency plots along the red dotted lines in the figures, located at the cavity resonant frequency. The amplitude of this sin curve is proportional to the real part of the reflection coefficient, and changes sign when the cavity goes from over to under coupled.

To lock the cavity to the source, the error signal is amplified, filtered, and taken to the cavity piezo. The frequency fluctuations of the cavity are suppressed within the bandwidth of the feedback loop, and any fluctuations of the source are imposed on the cavity. The open and closed loop noise spectra of the error signal are presented in figure 5.8 for various gain levels in the loop. The onset of oscillation is visible in the high gain curve.
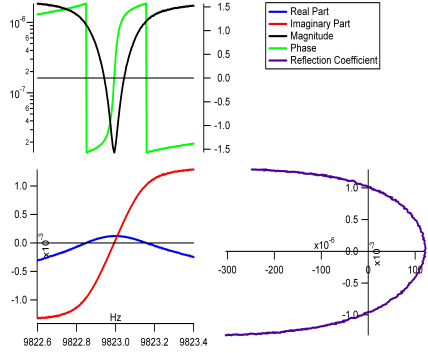
(a) Overcoupled Reflection Coefficient
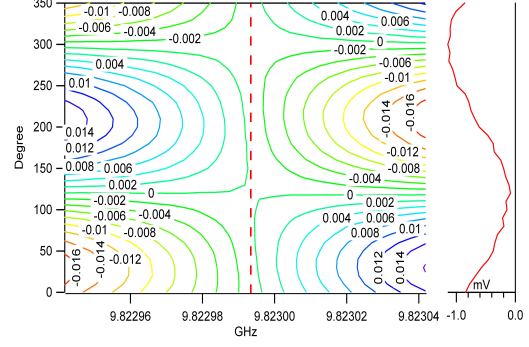
(b) Overcoupled Phase-Frequency Plot
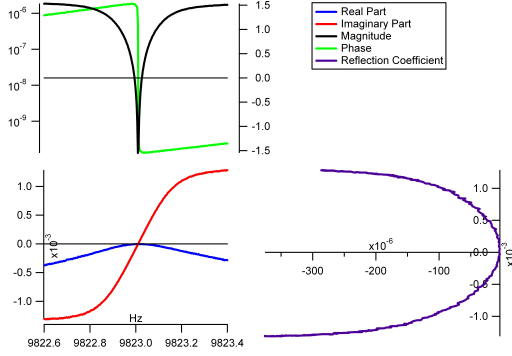
(c) Critically Coupled Reflection Coefficient
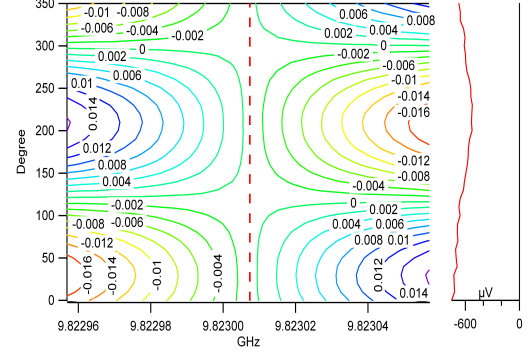
(d) Critically Coupled Phase-Frequency Plot
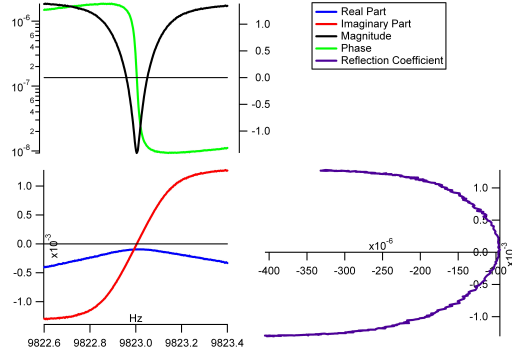
(e) Undercoupled Reflection Coefficient
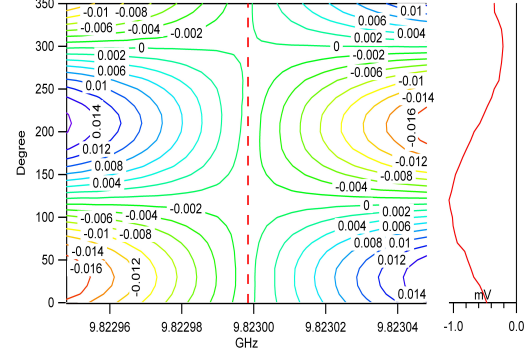
(f) Undercoupled Phase-Frequency Plot

Figure 5.7: The cavity coupling influences the phase-frequency plot as shown above. A small difference in the sideband amplitude makes the µPDH signal sensitive to the real part of the reflection coefficient if the phase of the down-converting mixer LO is not set properly.

5.4   Signal-to-Noise Optimization

The µPDH technique measures the frequency difference between the microwave source and the cavity. Any frequency fluctuations of the source cannot be distinguished from those of the cavity and therefore the best-case cavity stability that can be hoped for is equal to that of the source. The µPDH bridge should be designed with equal or better quality than the source, so that it does not degrade the best-case performance.

The best-case stability of the cavity can be calculated from the phase noise of the source, typically specified by $L_\phi(f_m)$, the single-sideband phase-noise power spectral density relative to the carrier per hertz, where f_m is the frequency offset from the carrier. The frequency-noise power spectral density $S_f(f_m)$ is related to the (double sideband) phase-noise power spectral density $S_\phi(f_m) = 2L_\phi(f_m)$
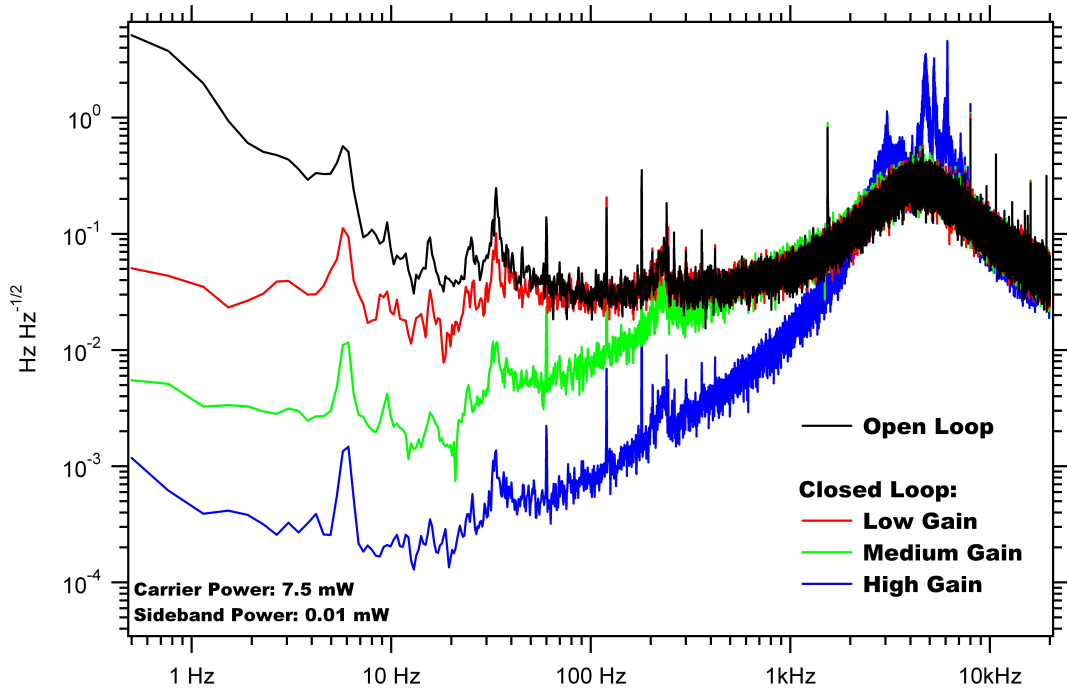
Figure 5.8: The open and closed loop noise spectra of the µPDH signal. In closed loop, the error signal is driven to zero within the bandwidth of the loop. The bandwidth increases with loop gain until the onset of oscillation, visible in the high gain curve.

through $S_f(f_m) = f_m^2 S_\phi(f_m)$ [36]. The µPDH noise spectrum in Hz Hz$^{-1/2}$ is equal to $\sqrt{S_f(f_m)} = f_m\sqrt{2L_\phi(f_m)}$ when it is source frequency-noise limited. For example, our source has $L_\phi(f_m) =$ -91 dBc Hz^{-1} at $f_m=$ 1 kHz offset from the carrier. The corresponding cavity frequency fluctuation is 0.04 Hz Hz$^{-1/2}$.

Other sources of noise may prevent one from reaching the frequency-noise limit. These can be categorized broadly into receiver or transmitter noise. The receiver noise level is due mainly to the diode detector and the preamplifier. These can be made insignificant by increasing the $Q\sqrt{P_cP_s}$ product so that the detected frequency-noise power is larger than the receiver noise. Then, if the transmitter contribution is negligible, the ultimate signal-to-noise limit for the given source has been achieved.

5.4.1 Transmitter noise, and the noise spectra

In the fabricated system, the overall signal-to-noise limit is determined by the transmitter; the receiver is negligible. To understand how the transmitter contributes noise, it is necessary to look at the microwave spectrum in detail. Figure 5.9a shows the spectrum of microwave power that reflects off the cavity when the carrier and cavity are aligned. The carrier is visible in the middle, but has been reduced in amplitude due to cavity absorption. Two sidebands are symmetric about the carrier at ±77 MHz. Two dips in the broadband noise spectrum are visible at $\pm2 \times 77$ MHz and are due to the notch filters removing noise at those frequencies. Several spurs are also visible. A closeup of the reflection around the cavity is shown in figure 5.9b. The dip is due to the cavity absorbing broadband microwave noise and the spike in the middle is the carrier reflection.

Figure 5.10a shows the signal-to-noise spectrum, and figure 5.10b shows the noise spectrum, of the open-loop error signal. The spectra are the same, except for division by the sensitivity. Nonetheless, both figures are useful in identifying
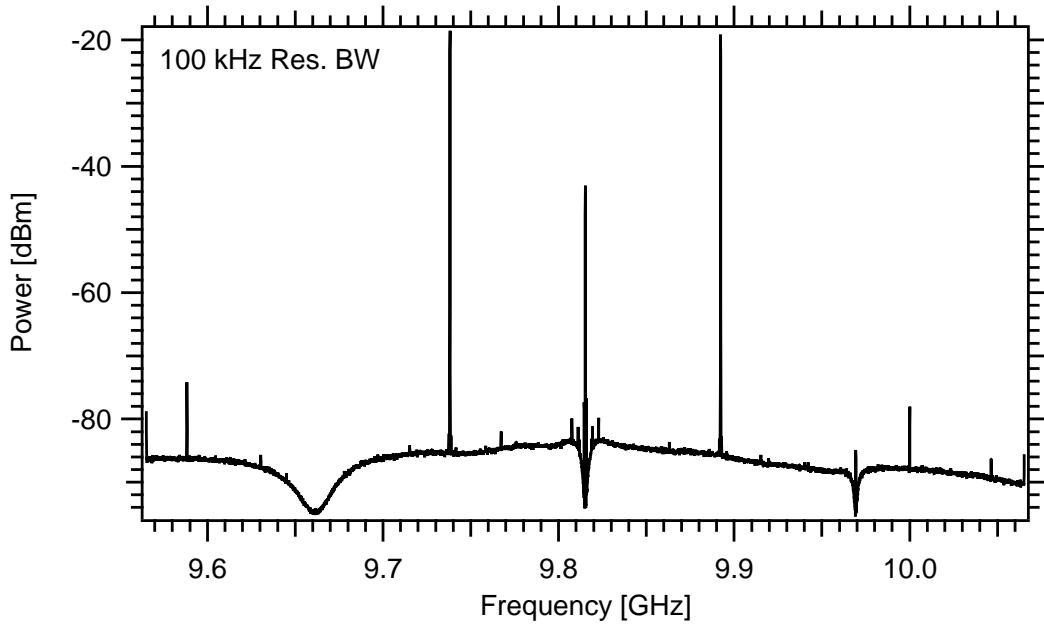
(a) Wideband Reflection off the Cavity
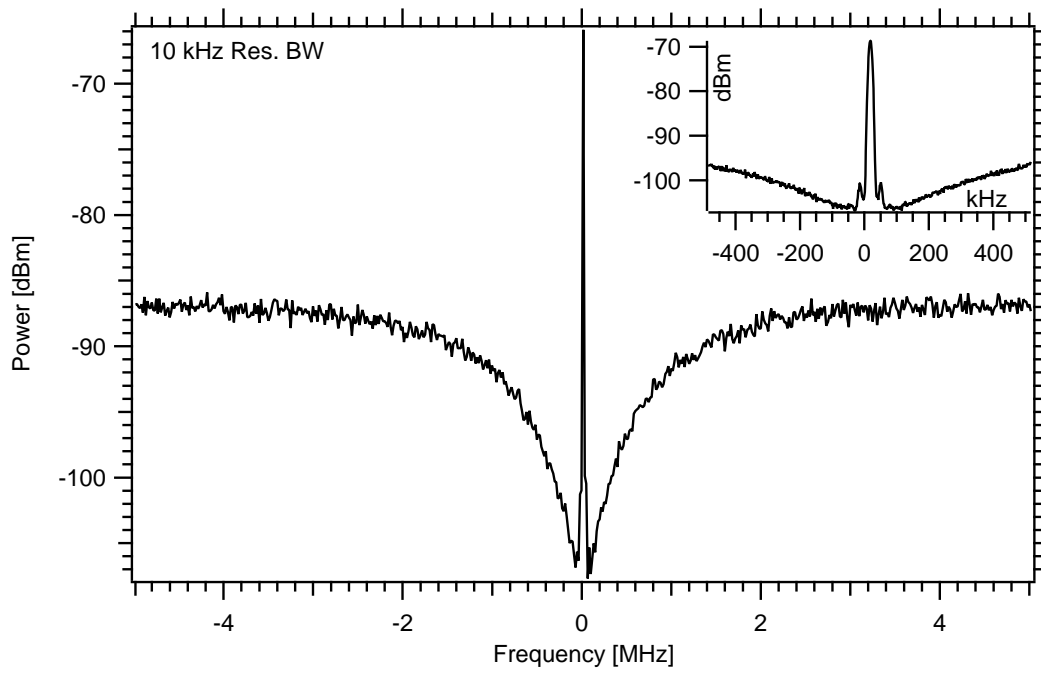
(b) Narrowband Reflection off the Cavity

Figure 5.9: Spectrum of microwave power reflecting off the cavity. The sidebands and notch filters can be seen in panel a), which shows the reflection over a 500 MHz bandwidth. A close-up around the cavity is shown in panel b).
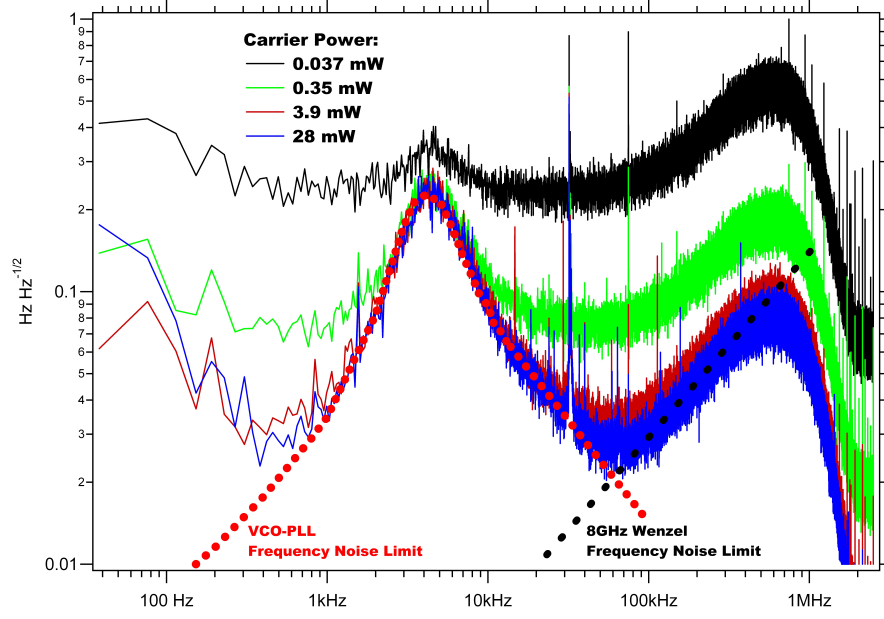
(a) Signal-to-noise spectrum
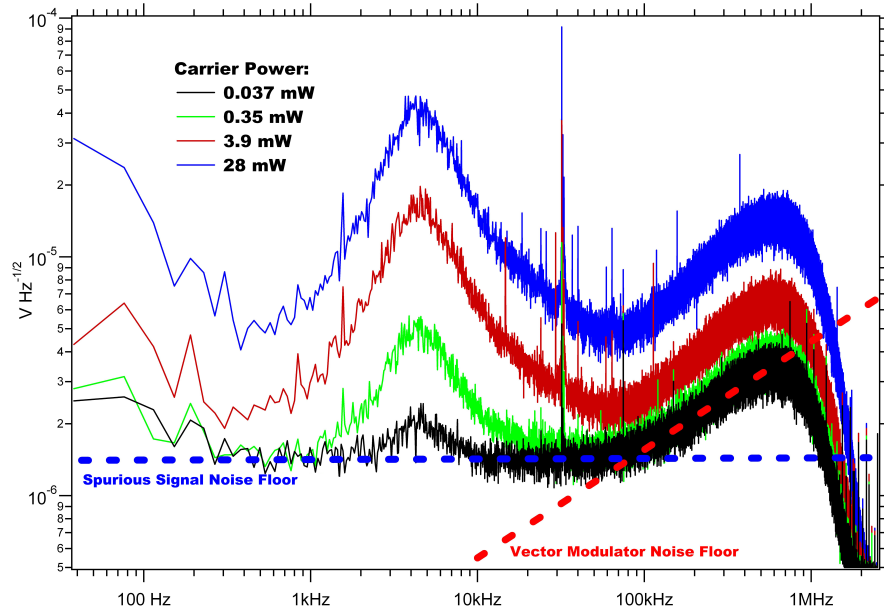
(b) Noise spectrum

Figure 5.10: The signal-to-noise spectrum (a) and the noise spectrum (b) of the open-loop µPDH error signal. The different colored curves correspond to different carrier power levels, all measured with a sideband power of 0.01 mW. Both spectra are useful in identifying the various noise sources, as the dependence on the carrier power is visible in the noise spectrum, but drops out of the signal-to-noise spectrum. At high carrier powers the frequency noise limit of the source is reached (blue and red curves). At low carrier powers, other noise limits are reached, shown by the labeled dotted lines in the figures (green and black curves).

49

the noise sources contributing at different power levels (discussed below). For example, the frequency-noise limit, having the same proportionality to P_c as the signal, appears as a power-independent curve in the signal-to-noise spectrum, figure 5.10a. The vector modulator and spurious signal noise floors, however, are not proportional to P_c and would shift with power if displayed in figure 5.10a. Instead, they are displayed in figure 5.10b, where they are power independent. To understand these spectra, it is necessary to consider the different processes that can result in a diode signal at frequency f_1 - particularly all the mixing products that can result in a frequency f_1. These are shown graphically in figure 5.11.

Figure 5.11 is a cartoon of the microwave spectrum before (dashed curve) and after (solid curve) filtration and reflection off the cavity. Both the cavity and the carrier are at frequency $f_0 = 9.823$ GHz and the sidebands at frequencies $f_0 \pm f_1$. The cavity absorbs most of the power in the carrier; any reflection is due to the coupling not being quite critical. The sidebands, being far from the cavity, are completely reflected. The overall microwave bandwidth is set by a bandpass filter, 1 GHz wide in our system. Close to the signals, the noise power is due to the phase noise of the source. The broadband noise is due to the vector modulator noise floor being up-converted to 10 GHz. Two notch filters remove noise at frequencies $f_0 \pm 2f_1$. After reflecting off the cavity the microwave is detected, generating a signal at the modulation frequency f_1, among others. The mixing products that can generate a diode signal at f_1 are depicted graphically in the figure.

5.4.2   Source 1: The signal

The useful μPDH signal is the result of the sidebands mixing with the back-reflected carrier. As discussed above, this contribution is proportional to $Q\sqrt{P_c P_s}$ and is zero when the carrier is the same frequency as the cavity.
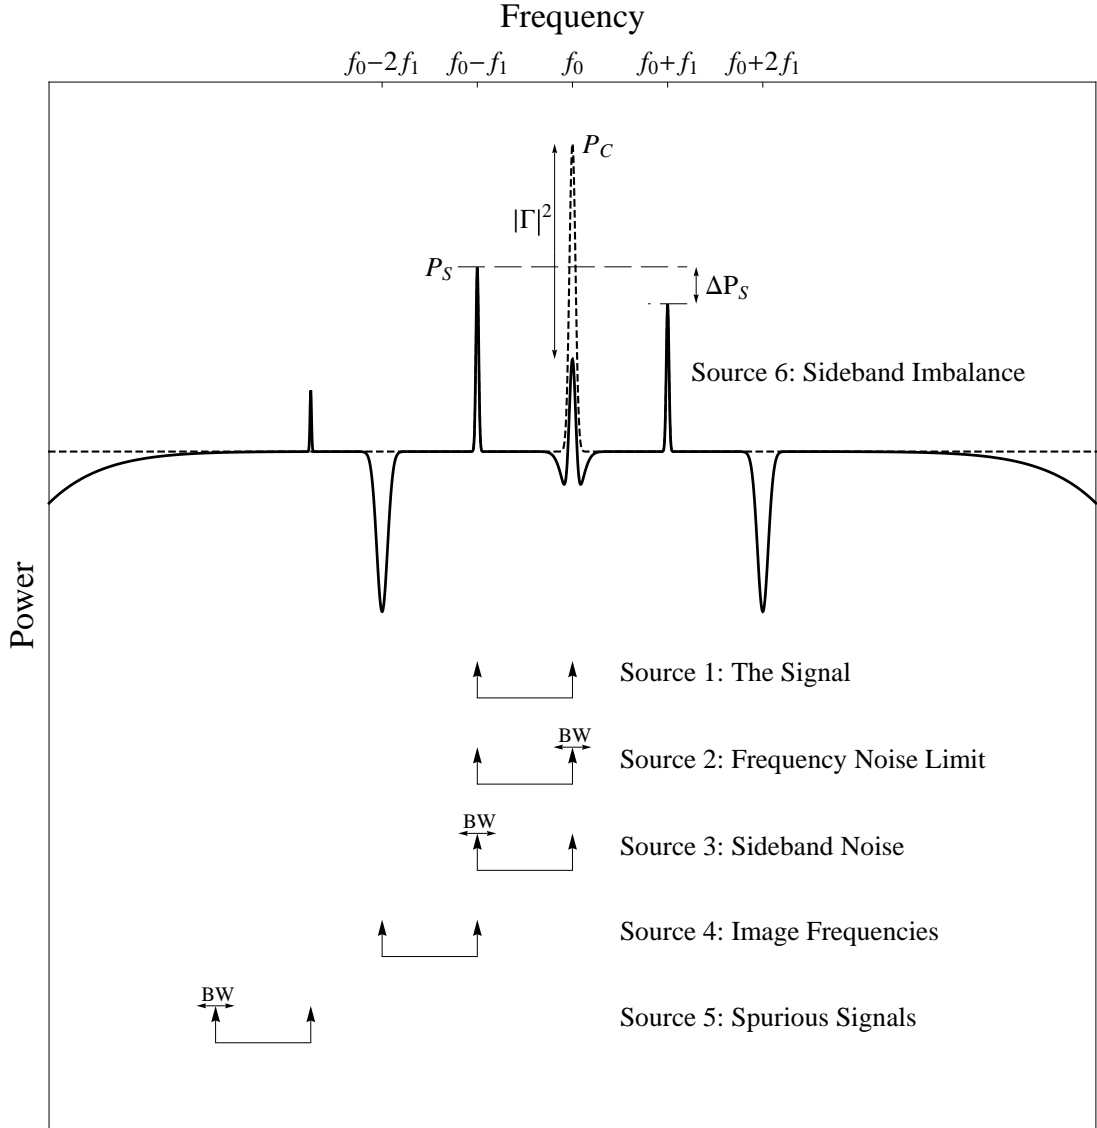
Figure 5.11: A cartoon depicting the microwave spectrum of figure 5.9a as well as various mixing products that can result in a diode signal at f_1. All these sources contribute to the signal-to-noise ratio of the system.

5.4.3 Source 2: The frequency-noise limit

The sidebands mix with the cavity-filtered noise around the carrier, which is primarily due to its phase noise. The carrier phase noise is a combination of the VCO (reduced by the PLL loop) and the 8 GHz Wenzel MXO phase noise, with an added contribution from the vector modulator noise floor. As the carrier

power is reduced, the combined phase noise power of the VCO and MXO is also suppressed in proportion to the carrier power. The noise power added by the vector modulator, however, is constant and becomes the dominant contributor to the carrier phase noise as its power is reduced.

This process is visible in figure 5.10. At high carrier power levels (greater than about 4 mW), the signal-to-noise ratio is completely determined by the phase noise of the VCO and MXO (blue curve in figures 5.10a and 5.10b). For these high powers, the shape of the spectrum between 1 and 60 kHz is due to the phase-locked VCO (red, dotted line in figure 5.10a). Above 60 kHz, the shape of the spectrum is determined by the MXO (black, dotted line in figure 5.10a). The down turn that begins around 600 kHz is due to the half width of the cavity (~ 1 MHz) as well as low-pass filtering at 1 MHz. Lower than 1 kHz, the spectrum reflects actual cavity drift.

As the power of the carrier is reduced, the vector modulator noise floor becomes relevant. This is visible in the black and green curves in figures 5.10a and 5.10b from about 80 kHz and higher, and is traced by the red dashed line in figure 5.10b. This presumably extends to lower frequencies as well, but it is not visible because of the spurious signal noise floor, discussed below. The relative contribution to the noise spectrum of the vector modulator and the MXO can be estimated by comparing the maximum of the peak at 4 kHz (VCO noise) to the maximum of the peak at 600 kHz (combined MXO and vector modulator noise). As long as the MXO is the dominant contributor, the ratio of the peak values will be power independent. As the carrier power is lowered and the vector modulator contribution becomes significant, the ratio will change with the power.

The noise power due to VCO and MXO frequency noise is proportional to $Q\sqrt{P_c P_s}$. Since this is the same proportionality as the signal, increasing that product does not improve the signal-to-noise ratio. The contribution of the vector modulator, however, is neither proportional to P_c nor P_s, so increasing either does
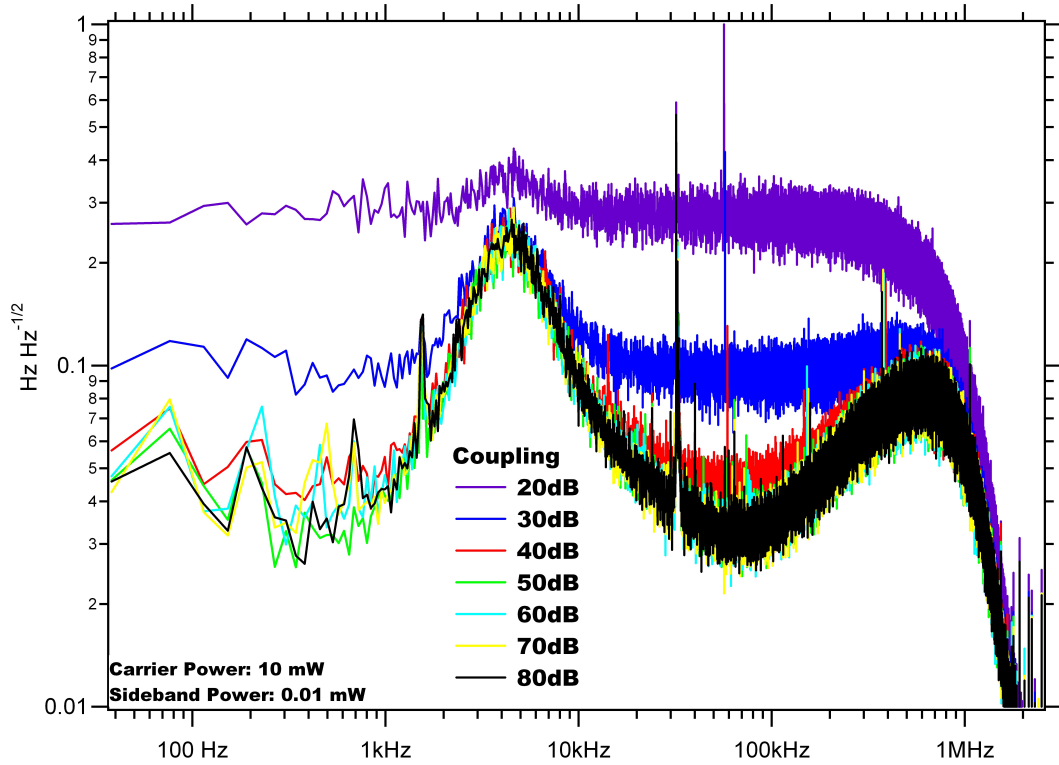
Figure 5.12: The signal-to-noise spectrum for various levels of coupling. Increasing the coupling improves the signal-to-noise level until the frequency noise limit is reached. Any further increase in coupling does not help. This occurs around 50 dB coupling for our system.

help.

5.4.4 Source 3: Sideband phase and amplitude noise

The back-reflected carrier mixes with close-in noise around the sidebands, which is mostly due to the vector modulator noise floor (because the sideband power is relatively low). Since the power of the back-reflected carrier drops sharply with better coupling, this noise source can be suppressed to negligible levels. The signal-to-noise spectrum as a function of coupling is shown in figure 5.12. The noise drops as the coupling is increased until the frequency noise limit is reached at about 50 dB coupling. Increasing the coupling no longer helps after that.
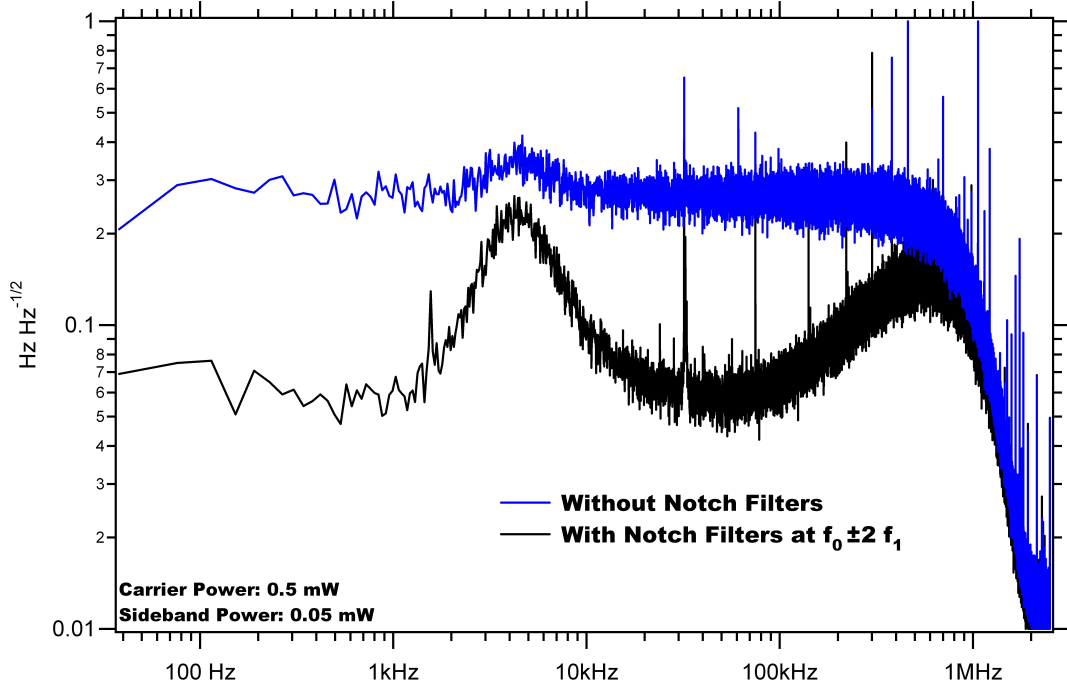
Figure 5.13: Notch filters can improve the signal-to-noise ratio in the low power limit.

5.4.5 Source 4: Image frequencies

Any noise present at frequencies $f_0 \pm 2f_1$ is mixed to f_1 by the sidebands. This can be a large contribution to the noise, especially for low carrier powers. It is proportional to $\sqrt{P_s}$, but not $Q\sqrt{P_c}$, so increasing Q or P_c improves signal-to-noise, but increasing the sideband power does not. Adding notch filters at $f_0 \pm 2f_1$ can reduce this contribution to negligible levels. Figure 5.13 shows the signal-to-noise spectrum measured both with and without using notch filters. At this low power, the filters reduce the high noise level (blue curve) by enough to reach the frequency noise limit (black curve).

5.4.6 Source 5: Spurious signals

This noise is a result of spurious signals mixing with broadband microwave noise that is offset from it by f_1. These spurs are unavoidable with the use of digital

synthesizers, but can be minimized with proper filtering and/or choosing the right digitally-synthesized frequency. In this case, the spur that causes the most havoc is not from the digital synthesis, but rather from insufficient LO-RF isolation in the up-converting mixer. The amplitude of the 8 GHz line, even after the broadband filter, dwarfs all other spurs by far. For low carrier powers, this sets the noise limit for frequencies below 80 kHz (blue, dashed line in figure 5.10b). In this case, signal-to-noise can be improved by increasing the $Q\sqrt{P_cP_s}$ product, because the amplitude of the 8 GHz line is independent of it. However, if the limiting spur is created by non-linearity in some device (mixer, or vector modulator for instance), the amplitude of the spur may grow with power, and each case must be considered individually.

5.4.7 Source 6: Sideband imbalance

An imbalance in the sideband power adds a contribution to the diode signal at f_1, but out-of-phase with the signal. This is not a source of noise, but it can complicate picking the optimal LO phase, as well as increasing the degree of coupling required to reach the frequency-noise limit. This signal is proportional to the real part of the cavity coupling coefficient Γ and the sideband asymmetry ΔP_s, and can be reduced by minimizing either one. It is symmetric about the cavity frequency (as opposed to the anti-symmetric µPDH signal) and makes the phase-frequency plot look skewed (see figure 5.7). Since it is out-of-phase with the signal of interest, it is irrelevant once the proper phase is chosen. However, its presence does require that the phase is chosen more accurately (see section 11.5.3).

5.5 Composition of the Frequency-Noise Limit

The source frequency noise is a combination of the frequency noise of the 1.8 GHz VCO - reduced by its PLL - and the 8 GHz MXO. A detailed breakdown of the sources of the transmitter phase noise is given in this section.

A 1.8 GHz signal is generated by a fractional-N PLL internal to the AD9956 chip as shown in figure 5.14. The VCO frequency is divided by $R = 8$ to generate a clock for the DDS. The DDS and digital-to-analog converter (DAC) generate a 100 MHz signal that is compared to the 100 MHz reference in a charge-pump phase detector. The low-pass filtered charge-pump output drives the VCO frequency. The closed-loop, single-sideband phase noise of the VCO is measured with an Agilent E4440A spectrum analyzer[4], shown by the blue dots in figure 5.15[5]. Limits imposed by the residual phase noise of the DAC and RF divider, and charge-pump phase detector (taken from the AD9956 datasheet [37]) are plotted in black and red respectively. Additionally, the phase-noise of the reference oscillator with respect to (wrt) 1820 MHz[6] is shown in purple. It is clear that the closed-loop VCO phase noise is limited by the internal circuitry of the AD9956, and not the cleanliness of the reference. The free-running VCO phase noise is also plotted in green in the figure. Its intersection with the phase-detector noise level indicates that the optimal loop bandwidth would be roughly 1.5 kHz. The actual loop bandwidth is closer to 4 kHz, as indicated by the local maximum in the closed-loop VCO phase noise at that frequency.

---

[4]Due to the non-linear processes that naturally limit the oscillation amplitude from growing indefinitely, AM noise is generally suppressed in oscillators, especially close-in to the carrier. Consequently, the noise observed on a spectrum analyzer close-in to a signal is predominantly due to phase-noise. This is doubly true in the case of multiplied sources, as in the MXO, where the multiplication process enhances phase noise, but not amplitude noise [38].

[5]Following [39], 1.99 dB is added to the spectrum analyzer measurement to correct for its effective noise bandwidth as well as its log-scale response.

[6]The phase noise of an oscillator grows by a factor of N^2 when its frequency is multiplied by N [38]. To take this into account, $20 \log(1820/100)$ dB have been added to the reference oscillator's specifications.

The 8 GHz MXO follows its specifications very closely. Its phase-noise is plotted in figure 5.16 (red-dashed curve) along side the closed-loop VCO phase noise (blue-dashed curve). When the two signals are mixed to produce the 9.8 GHz signal, their phase-noise adds to produce the black curve. To generate the VCO and MXO curves over the 5 decades shown in the figure, data from the specification sheet of the sources was relied upon when spectrum analyzer measurements were not available[7]. For frequencies below 100 Hz and above 70 kHz, the MXO is the limiting component of the transmitter. In between, the components of the PLL are limiting.

The frequency-noise limit of the µPDH signal-to-noise ratio is given by $\sqrt{S_f} = f_m\sqrt{2L_\phi(f_m)}$. Inserting the black curve of figure 5.16 for $L_\phi(f_m)$ gives the red curve in figure 5.17, the frequency-noise limit of our source. It matches the observed spectrum between 1-500 kHz. Below 1 kHz, the spectrum is believed to be due to actual motion of the cavity (a real signal, not noise). Above 500 kHz, the deviation is due to low-pass filtering at 1 MHz, as well as the offset frequency becoming comparable to the cavity half-width.

---

[7]The high-quality of the sources made it difficult to measure their phase noise with a spectrum analyzer. In some frequency ranges, the spectrum analyzer noise overwhelmed the measurement and the results were not reliable. Data from the source specification sheet was used to fill in the gaps where needed.
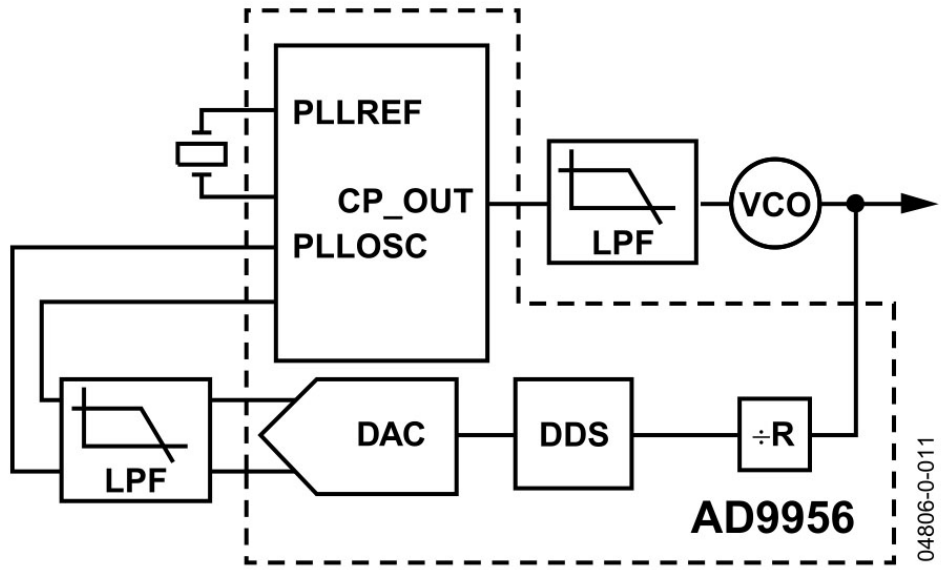
Figure 5.14: Block diagram of the fractional-N PLL that generates the 1.8 GHz signal. Reproduced from [37].
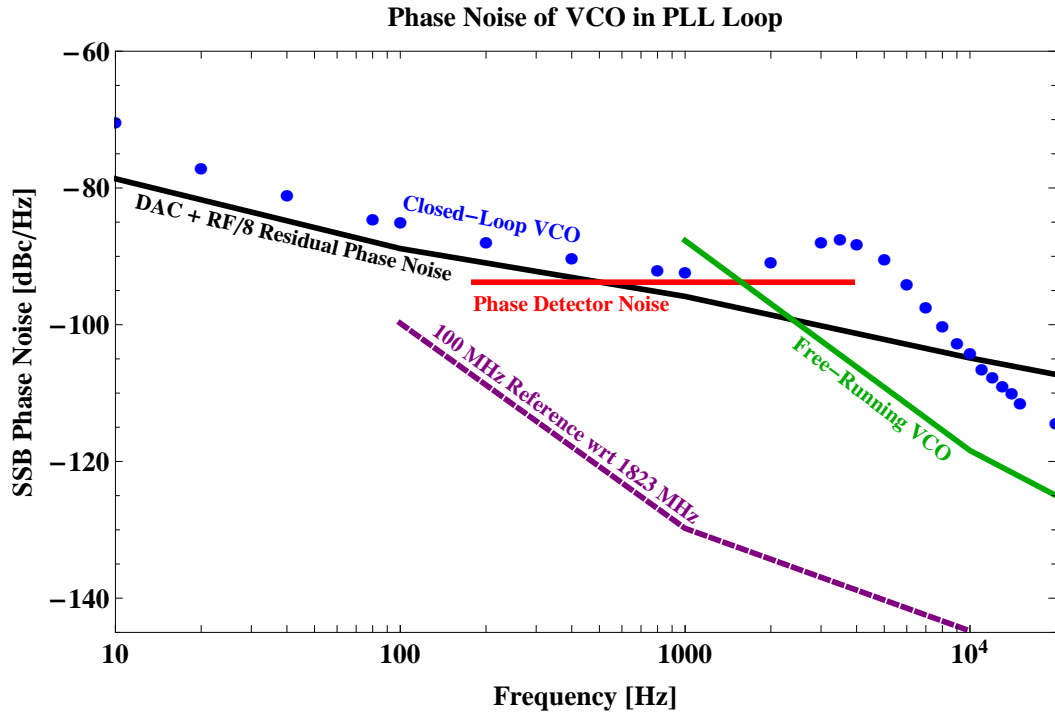
Figure 5.15: The various contributions to the closed-loop VCO phase noise. Within the loop bandwidth ($\sim$ 4 kHz), the PLL performance is limited by the DAC residual phase noise and the phase detector noise floor. The 100 MHz Wenzel crystal reference is far cleaner than the components within the loop.

58

Combined Phase Noise
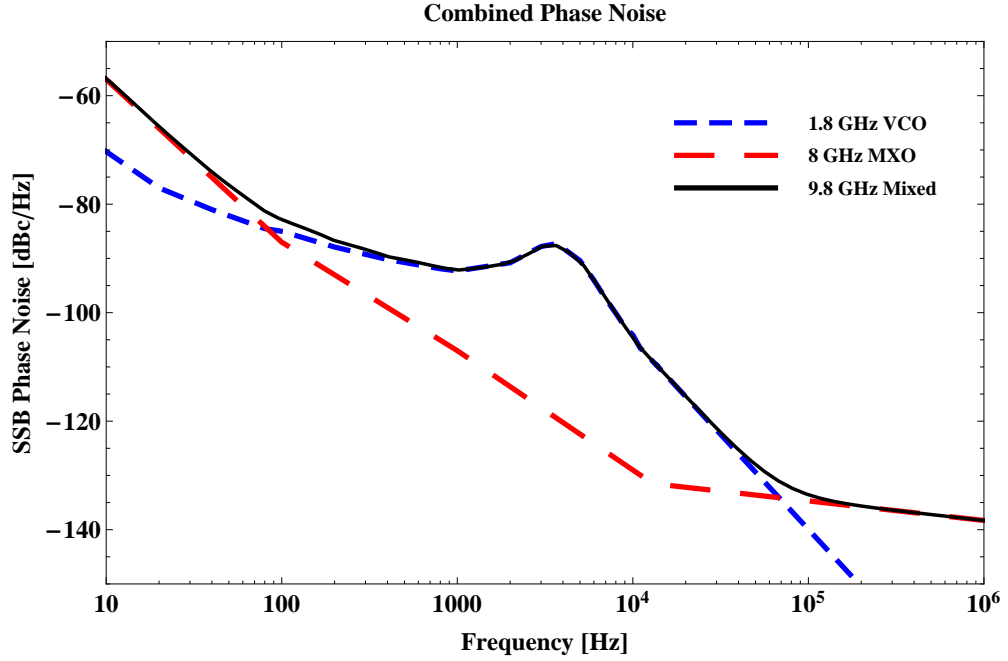
Figure 5.16: The phase noise of the VCO (blue, dashed curve) and MXO (red, dashed curve) add when they are mixed to produce the overall phase noise spectrum of the source (black curve). The frequency noise limit of the source is calculated from this curve and is shown by the red curve in figure 5.17.

μPDH Noise Spectral Density
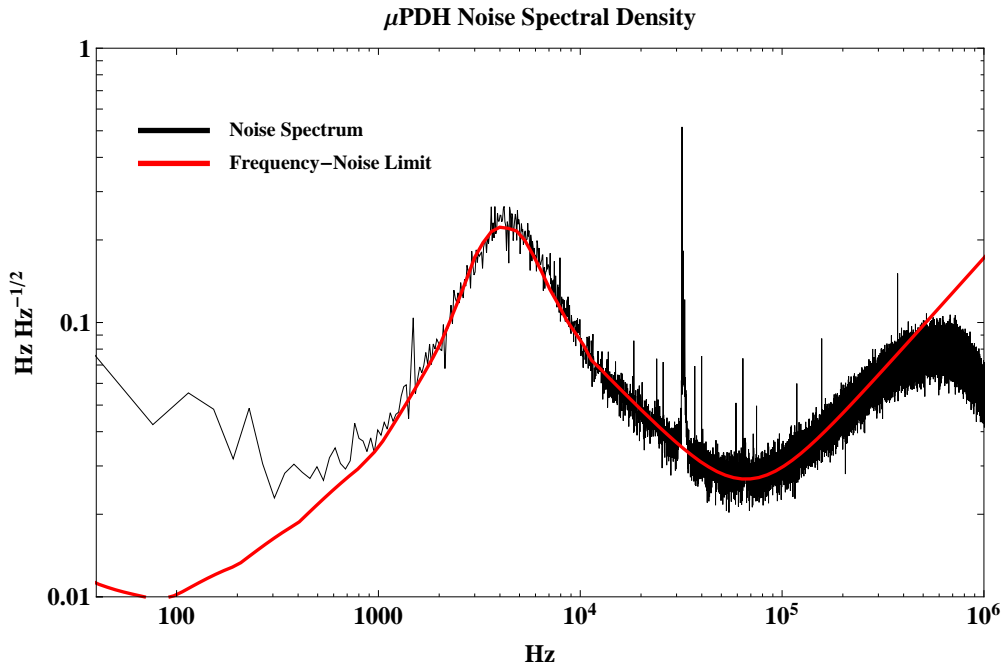
Figure 5.17: The frequency-noise limit calculated from the phase-noise of the sources (red curve) agrees with the observed noise spectrum (black curve) between 1-500 kHz.

5.6   Single Sideband Lock

In the traditional PDH and µPDH schemes, the vector sum of the two sidebands is always orthogonal to the carrier, and therefore the PDH signal is only sensitive to the imaginary part of the cavity reflection coefficient, $\Im[\Gamma]$. If $\Im[\Gamma]$ is the only relevant information for the application at hand, this is great, because 1) the detected signal is relatively immune from changes in the cavity coupling and 2) choosing the phase of the down-converting mixer is straightforward, as described in section 5.3. However, knowing the real part of the cavity reflection coefficient, $\Re[\Gamma]$, may be useful for some applications.

Thanks to flexibility of the vector modulator, the µPDH scheme can be modified with minimal hardware changes to measure both $\Im[\Gamma]$ and $\Re[\Gamma]$. A block diagram of such a bridge is shown in figure 5.18. It is largely the same as the system in figure 5.3, so for simplicity, only changes will be described here. The 2-channel DDS board that generates the modulation frequency is replaced by a 4-channel version of the same board, the AD9959. Channels one and two are used to generate a single sideband by driving the $\boldsymbol{I}$ and $\boldsymbol{Q}$ baseband inputs of the vector modulator 90° out-of-phase, as described in section 4.1.4. The DC levels on the the baseband inputs are set so that the carrier lies along the $\boldsymbol{I}$ axis (this simplifies tuning the detector). The detector is modified by the addition of a second channel. The amplified and filtered diode signal is split and taken to two mixers, whose local oscillators are provided by channels three and four of the 4-channel DDS. Setting the local oscillator phases properly is critical in separating $\Im[\Gamma]$ from $\Re[\Gamma]$.

Proper phasing is done in multiple steps. First, modulation of the $\boldsymbol{I}$ baseband input is turned off, reverting to double-sideband detection. A phase-frequency plot, as in figure 5.5a, is generated for both channels, and the optimal phase for $\Im[\Gamma]$ detection is chosen for both channels as described in section 5.3. The optimal

phase for $\Re[\Gamma]$ detection will be 90° off from this, so 90° is added to the phase of one of the LOs. Turning back on modulation on the $\boldsymbol{I}$ baseband input will result in a full quadrature detection of the cavity reflection coefficient.

As a proof of principle, single-sideband detection was used to measure the quadrature EPR spectra shown in figure 5.19. The bridge of a Bruker e680 EPR spectrometer was replaced with the quadrature bridge of figure 5.18. A rectangular TE_{102} EPR cavity was modified as described in 5.2 to make it tunable and allow locking it to the source. The $\Im[\Gamma]$ output of the bridge served two purposes. The low frequency component of the signal locked the cavity to the source in a bandwidth less than 5 kHz, while the 100 kHz component (the field-modulation was 100 kHz) was taken to the spectrometer signal channel to give the EPR dispersion signal. The $\Re[\Gamma]$ output of the bridge was taken to a second signal channel to give the EPR absorption signal. The conditions at which the spectra were measured are given in table 5.1.
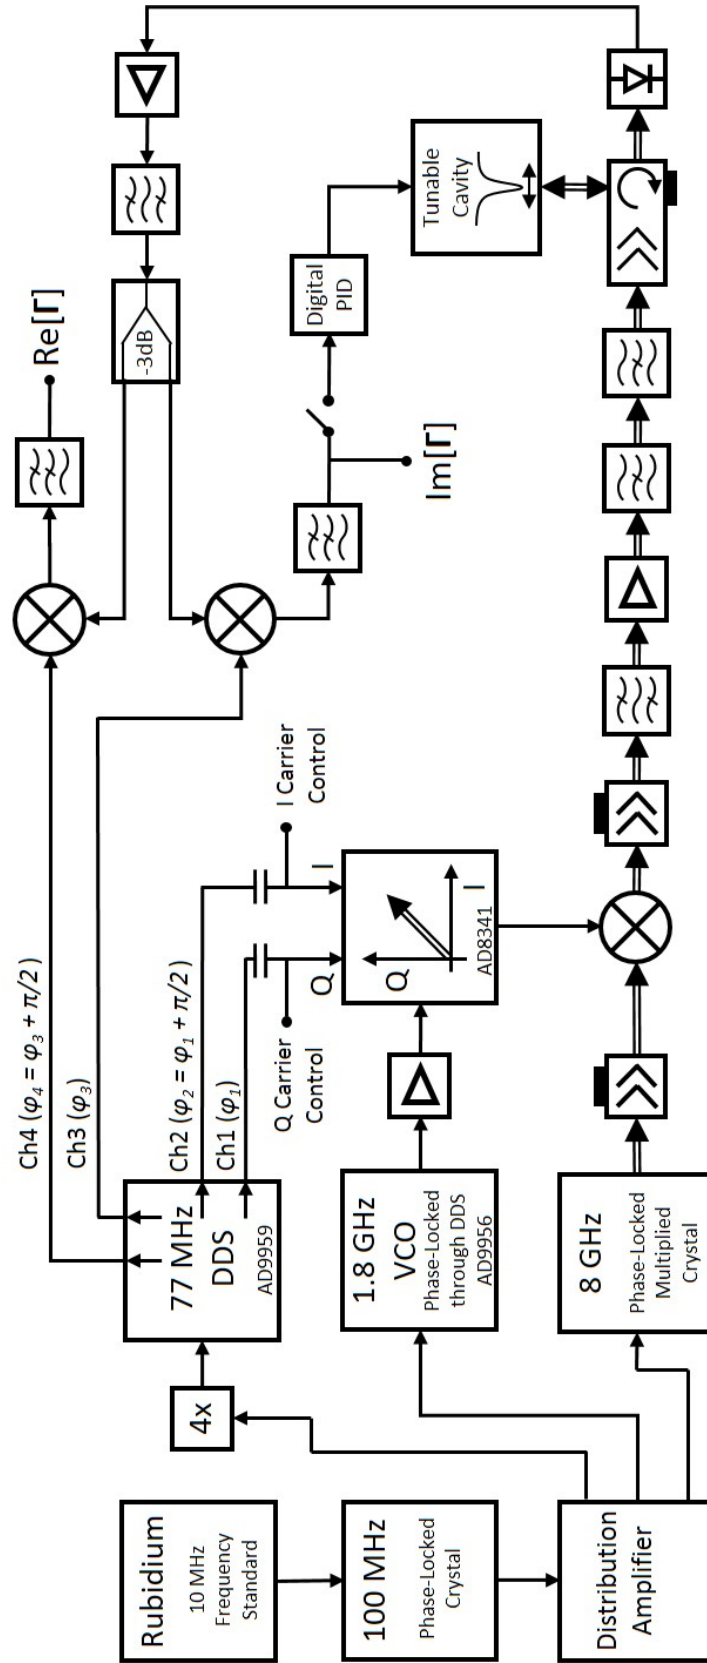
Figure 5.18: Block diagram of a single-sideband quadrature bridge that is able to measure both the real and imaginary parts of the cavity reflection coefficient.

	DPPH	PNT	Weak Pitch
Microwave Power [mW]	4.5	0.6	12
Field Modulation [G]	0.2	0.2	6
Time Constant [ms]	1.28	10.24	1311
Sampling Time [ms]	5.12	40.96	163.8

Table 5.1: The conditions during measurement of the EPR spectra in figure 5.19

(a) α,α'-diphenyl-β-picryl hydrazyl (DPPH)

(b) Perinaphthenyl Radical (PNT)

(c) 0.0003% pitch in KCl (Weak Pitch)
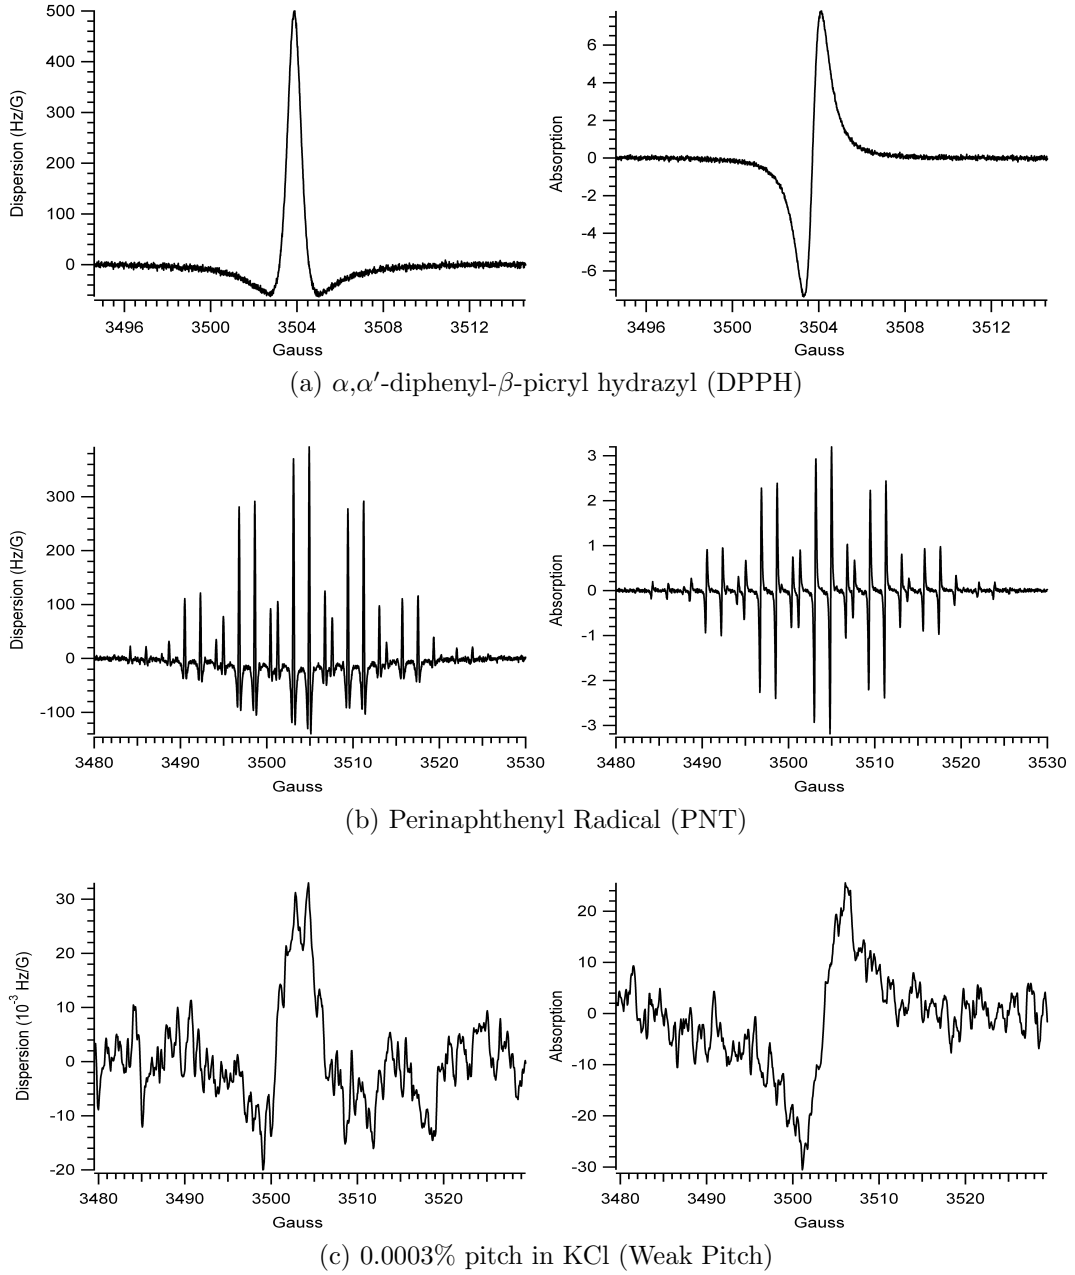
Figure 5.19: The EPR spectra of three common calibration sources provided by Bruker, measured with the quadrature bridge of figure 5.18.

Part II

The Frequency-Distance Encoder

CHAPTER 6

Coaxial Resonator Theory

For frequency-to-distance conversion with microwave cavities, it is important that only the cavity length, and not any transverse dimension, influences its resonant frequency. It is well-known that two-conductor transmission lines can support transverse electromagnetic (TEM) modes, and that the propagation constant of these waves depends only on the speed of light in the homogeneous and isotropic dielectric between them [40]. Shorting the ends of such a transmission line, by inserting two conducting planes perpendicular to its axis, produces a resonant cavity supporting modes with the desired properties. A coaxial cavity is the simplest cavity supporting TEM modes, and a derivation of its Q and the resonant frequency of the lowest non-TEM mode are given in this chapter.

6.1   Q of the Fundamental TEM Mode

The electric, $\vec{E}$, and magnetic, $\vec{H}$, fields inside a half-wavelength, coaxial cavity are,

$$\vec{E} = \frac{V_0}{\ln b/a} \frac{1}{\rho} \sin kz \ \hat{\rho},$$

(6.1a)

$$\vec{H} = \frac{-i}{\eta} \frac{V_0}{\ln b/a} \frac{1}{\rho} \cos kz \ \hat{\phi},$$

(6.1b)

where ρ is the cylindrical radial coordinate, z is the axial coordinate, b and a are the radii of the outer and inner conductors, V_0 is the sinusoidal voltage amplitude

(a) Electric Field   (b) Magnetic Field
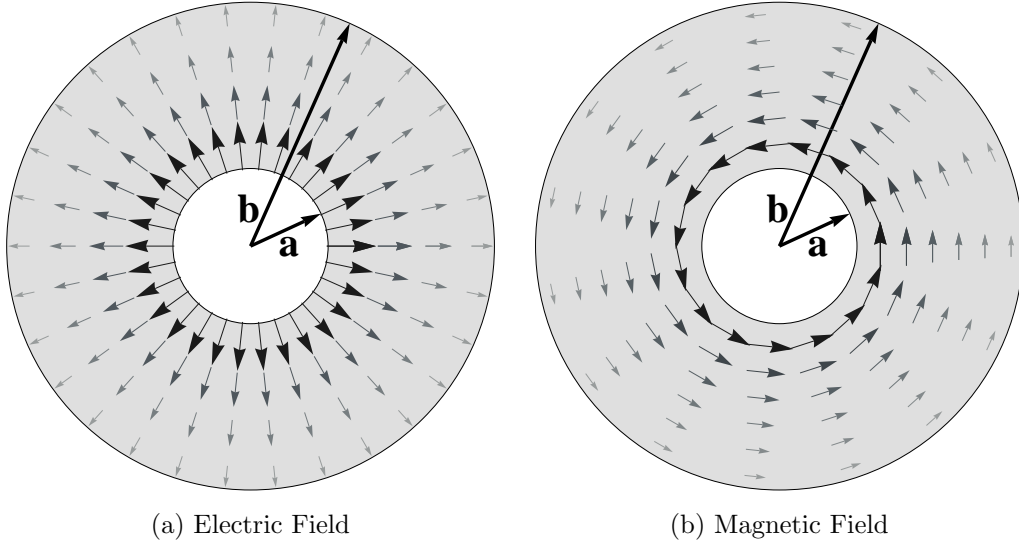
Figure 6.1: The electric and magnetic fields in a cross-section of the coaxial cavity.

between them, $\eta = \sqrt{\frac{\mu}{\epsilon}}$ is the impedance of the medium filling the cavity, and $k = \pi/L$. The time dependence $e^{-i\omega t}$ is implied. The fields are shown in figure 6.1. The time-averaged energy in the cavity is,

$$W = \frac{\epsilon}{4}\int |E|^2 dV + \frac{\mu}{4}\int |H|^2 dV, \tag{6.2a}$$

$$= \frac{\pi \epsilon V_0^2 L}{2 \ln b/a}. \tag{6.2b}$$

The time-averaged power absorbed in the conductor is,

$$P_{loss} = \frac{1}{2\sigma\delta}\int |H_{||}|^2 ds, \tag{6.3a}$$

$$= \frac{2\pi V_0^2}{\sigma\delta\eta^2}\frac{1}{\ln b/a}\left[1 + (\frac{1}{a} + \frac{1}{b})\frac{L}{4\ln b/a}\right], \tag{6.3b}$$
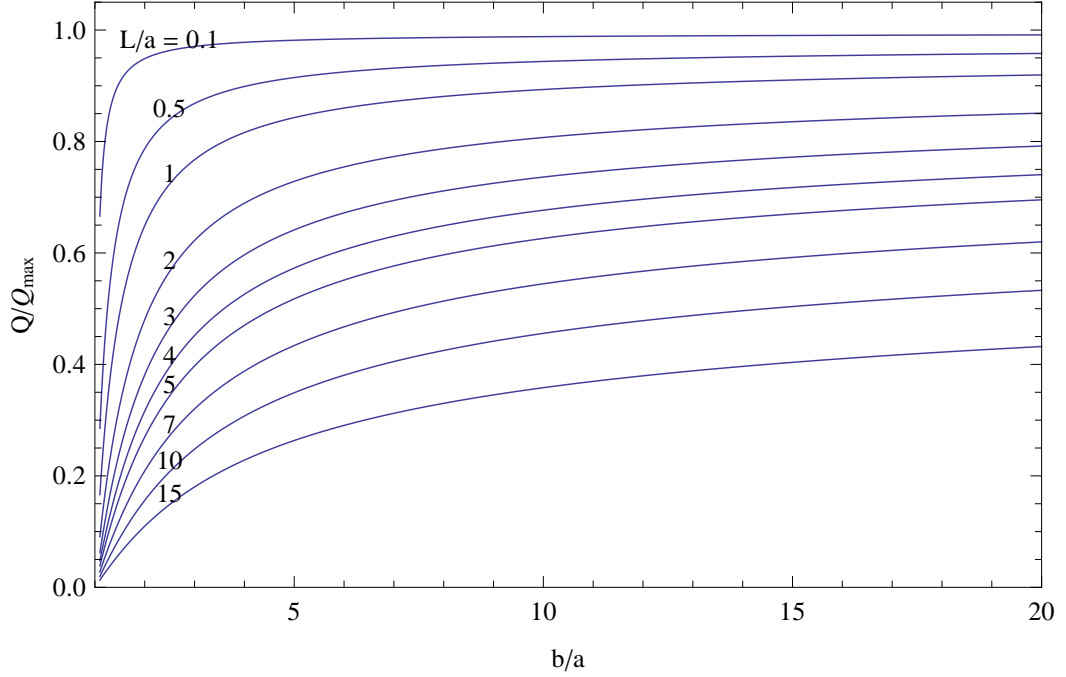
Figure 6.2: Q as a function of the cavity dimensions. The outer and inner radii are b and a respectively, and L is the cavity length.

where σ is the conductivity of the conductor, $\delta = \sqrt{\frac{2}{\mu_c \omega \sigma}}$ is the skin depth, and μ_c is the magnetic permeability of the conductor. Finally, the cavity Q is,

$$Q = \frac{\omega W}{P_{loss}}, \tag{6.4a}$$

$$= \frac{Q_{max}}{1 + \frac{L}{4 \ln b/a}(1/a + 1/b)}, \tag{6.4b}$$

where $Q_{max} = \frac{L\mu}{2\delta\mu_c}$. Q/Q_{max} as a function of b/a is plotted in figure 6.2 for various values of L/a. $Q_{max} = 9540$ for a gold, coaxial cavity of 15 mm length. Such a cavity with an inner and outer diameter of 3.5 mm and 11 mm respectively would have an unloaded Q of 2750. The unloaded Qs of the lowest 10 modes, numerically calculated in Ansoft HFSS, are given in table 6.1.

Mode	Frequency (GHz)	Q
1	9.991 + i 0.00185	2705
2	16.95 + i 0.00171	4973
3	16.95 + i 0.00171	4973
4	19.98 + i 0.00261	3825
5	24.22 + i 0.00233	5200
6	24.22 + i 0.00233	5200
7	27.75 + i 0.00214	6475
8	27.75 + i 0.00214	6475
9	30.00 + i 0.00320	4685
10	32.70 + i 0.00240	6825

Table 6.1: Numerical calculations of the ten lowest-frequency resonant modes and unloaded Qs of a gold coaxial cavity of 15 mm length and 3.5 mm and 11 mm inner and outer diameter respectively.

6.2 TE$_{111}$ Mode

The TEM waveguide mode has no cutoff frequency; it propagates down to DC frequencies with a propagation constant that does not depend on the dimensions of the cavity. It is desirable to operate at a frequency where there are no resonant TE or TM modes nearby. In a coaxial cavity, after the fundamental TEM resonance, the next highest mode is the TE$_{111}$ mode. The resonant frequency of the TE$_{111}$ mode as a function of the cavity dimensions is derived here.

The propagation constant of the TE$_{11}$ waveguide mode is $\beta = \sqrt{k^2 + k_c^2}$, where $k = \omega/c$ is the standard plane wave propagation constant, and k_c is the cutoff frequency [41]. Imposing the boundary condition at the $z = 0, L$ planes gives $\beta L = \pi$. The cutoff frequency is given by the solution to the transcendental equation,

$$J_1'(k_c a)Y_1'(k_c b) - J_1'(k_c b)Y_1'(k_c a) = 0, \tag{6.5}$$

where J_1' and Y_1' are the derivatives of the first order Bessel functions of the first and second kind respectively. Solving for the resonant frequency gives,

$$\omega_{111}^2 = c^2(k_c^2 + \beta^2), \tag{6.6}$$

68
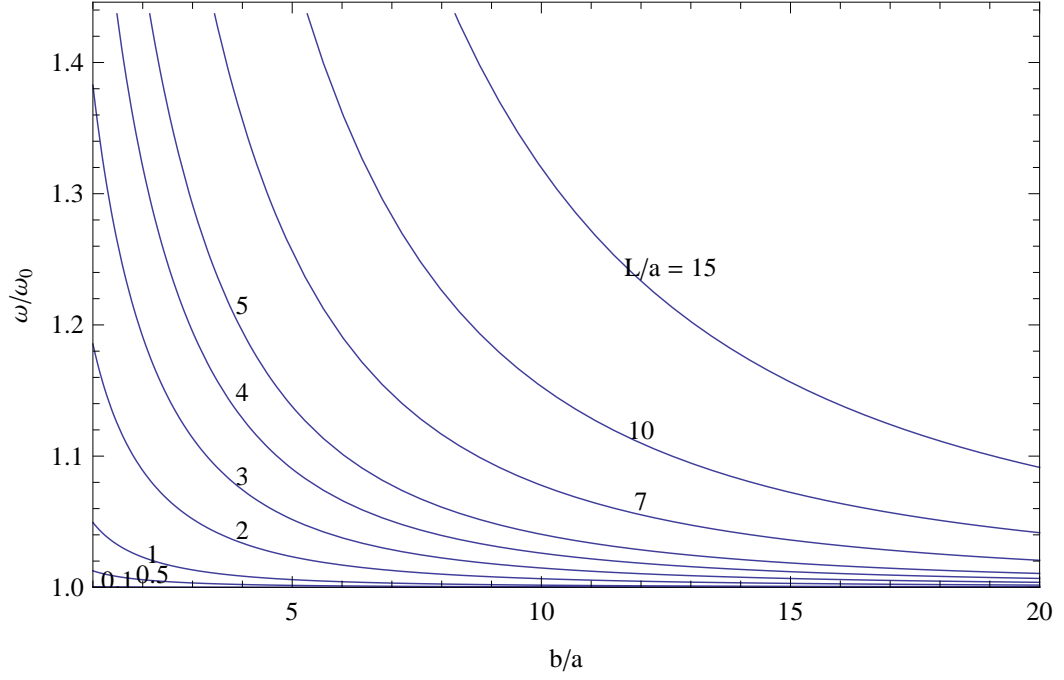
Figure 6.3: The resonant frequency of the TE$_{111}$ mode of a coaxial cavity normalized to the frequency of the fundamental TEM mode. The outer and inner radii are b and a, and L is the cavity length.

and normalizing to the resonant frequency of the TEM mode, $\omega_0 = c\pi/L$,

$$\frac{\omega_{111}^2}{\omega_0^2} = 1 + \frac{k_c^2 L^2}{\pi^2}. \tag{6.7}$$

Figure 6.3 shows the normalized resonant frequency of the next highest mode as a function of the cavity dimensions. For the example cavity with dimensions given above, the next highest resonance is predicted to be at a frequency 1.683 times higher than the TEM resonance.

CHAPTER 7

Cavity Coupling Design

Having chosen the cavity and mode to be used, a coupling structure must be designed that allows electromagnetic energy to enter the right mode. Additionally, the coupling must be variable in order to achieve the > 50 dB return loss necessary to optimize signal-to-noise, as discussed in section 5.4.4.

The best coupling structure is to be judged on the basis of 1) its influence on the resonant frequency of the cavity, 2) the range of coupling tunability, and 3) the practicality of the design for an eventual application (for instance, the center may need to be left clear for a laser or fiber to pass). The ideal structure would be smoothly tunable between say ~ 10 dB under to ~ 20 dB over coupled and display no frequency change over this range. Months were spent modeling and experimenting with various coupling structures without finding one that met all the above requirements as well as was hoped. It turns out that it is very difficult to achieve all three at the same time. There were important lessons learned along the way that will be discussed in section 7.1. The final structure chosen, presented in section 7.2, has an excellent tunability range, and is mechanically robust and repeatable. In addition, it has negligible influence on the cavity resonant frequency (see section 11.3). There is, however, no clear path down the axis of the cavity, which may be important for some applications.
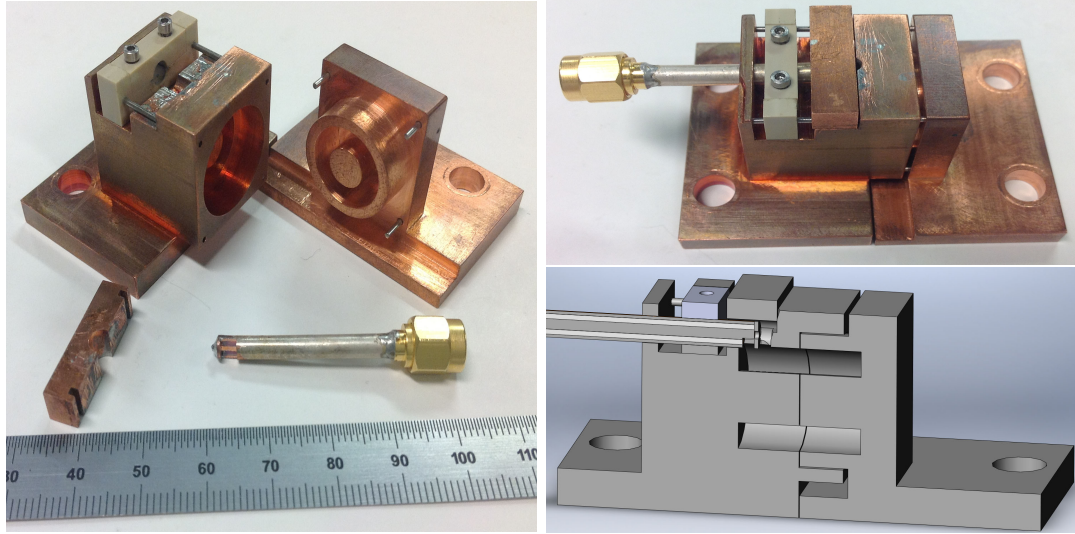
Figure 7.1: One of the early attempts at coupling into a coaxial cavity. The various pieces are shown individually on the left and assembled on the top right. The bottom right is a cross section of a model showing the internal structure of the figure above it. A cut near the shorted end of a semirigid cable is placed above a cut in the coaxial cavity wall. The amount of overlap of the cuts determines the degree of coupling.

7.1 Early Attempts

On the order of a dozen different type of coupling mechanisms were modeled in Ansoft High Frequency Structure Simulator (HFSS), but only a couple made it to fabrication. The first is shown in figure 7.1. The purpose of building this cavity was two-fold: 1) test the coupling mechanism and 2) experimentally verify that the cavity mode is not altered significantly, as indicated by the Q, when the two halves are moved apart by a bit. There is an incoming semirigid coaxial cable that is shorted at one end. At the shorted end, a cut is made in the outer conductor, allowing some magnetic field to leak out. A cut of comparable size is also made on the outer conductor of the coaxial cavity, near the end. The shorted, semirigid cable is placed so that the cut in the cable lies above the cut in the cavity so that the fields couple. By adjusting the overlap of the cuts, the coupling can be varied.

Apart from the fact that the mechanical stability of the design was terrible, the

71
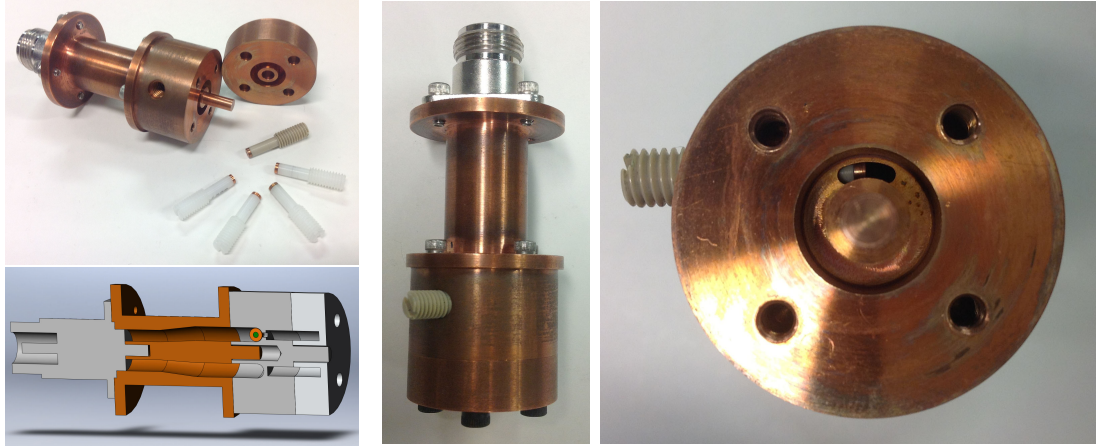
Figure 7.2: A more refined attempt at coupling into a coaxial cavity. The various pieces are shown individually in the top left, and assembled in the center. The bottom left is a cross section of the model showing the internal structure of the assembly. The hole that couples the waveguide to the cavity is shown on the right, which is a face-on view of half the cavity. A copper pill, whose position can be varied with the visible dielectric screw, can be seen behind the hole.

downfall of the coupling mechanism was that it was too weak. It could never be brought to an overcoupled state with these semirigid and cavity dimensions. The mismatch between the semirigid and cavity diameters played a roll (low overlap of the fields), as did the fact that the hole was cut in the outer wall of the cavity, where the fields are weaker. This cavity was useful, however, in that it experimentally verified that a coaxial cavity can be sliced down the middle, and the two halved moved apart, without a significant degradation in Q.

A second attempt, see figure 7.2, proved much more successful. Here the incoming coaxial cable is along the same axis as the cavity, and the length scales of the incoming fields are better matched to the ones of the cavity mode. The structure starts with a standard N-type connector that ends in a free space coaxial structure. Then there is a slight diameter transition to get to a size that matches the cavity. This "oversized coax" region ends at a shorting wall in which a wedge-shaped hole is cut, joining the coax to the cavity. A copper ring, the "pill" is placed on the coax side of the hole, which helps concentrate the field, and makes

the hole appear to be larger. The position of the pill along the hole can be varied to provide tunability. The highest coupling is when the pill is centered on the wedge, and decreases as it is moved to the side. Therefore, the size of the hole is chosen to be as small as possible while still maintaining a small region of overcoupling.

This coupling structure performs extremely well. It has a smooth, repeatable, and mechanically stable coupling adjustment that can achieve better than 70 dB matching. The tuning range can be from -6 dB under to greater than -10 dB over, and there is a 5-6 MHz frequency shift in going from -6 dB under coupled to critically coupled. Regrettably, there is no clear path down the axis of the cavity, which might be advantageous in some applications.

Fabricating some of the pieces in this structure is not trivial. The half-cavity with the coupling hole is especially difficult. Once it has been machined, it is not easy to go back and change the size of the coupling hole if needed. Therefore, it is of interest to know the range of coupling change one can expect by altering a different piece, the diameter of the pill for instance. To this end, a set of pills with slightly different diameters was fabricated. With no pill in front of the coupling hole, the cavity is 6 dB undercoupled. As each pill is inserted in front of the hole, the coupling increases until at some point the cavity is critically coupled. Continuing past this critically-coupled position overcouples the cavity, and the coupling increases until a maximum is reached when the pill is centered in front of the hole. The coupling here is the "maximum coupling" shown in figure 7.3. As the pill diameter is increased, the coupling increases as well. A change of over 20 dB in the coupling range was achieved by increasing the pill diameter by only 0.0118 inch. The resonant frequency change in going from pill removed to critically coupled was 5-6 MHz for all pills tested, without any measurable trend as a function of the pill diameter.
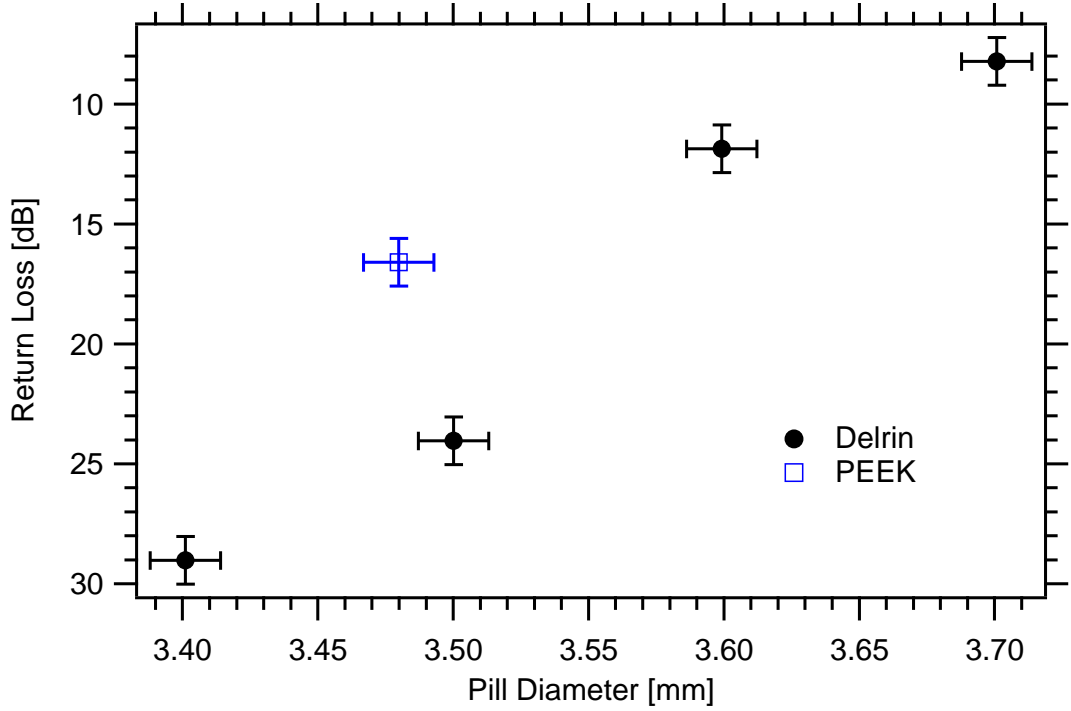
Figure 7.3: The range of coupling of the structure in figure 7.2 can be changed by altering the diameter of the copper pill placed in front of the coupling hole. Without a pill in front of the hole, the cavity is 6 dB undercoupled. The points plotted are the cavity return loss when the pill is centered on the hole, the position of highest coupling. Changing the pill diameter provides an easy means of adjusting the coupling range without having to alter the coupling hole, a delicate operation. The dielectric that holds the copper pill appears to play a small role as well, as there is a visible difference in using PEEK instead of delrin.

7.2 Coupling of the Realized Frequency-Distance Encoder

A variation of the coupling structure shown in figure 7.2 was chosen for the frequency-distance encoder. The final coupling structure is shown in figure 7.4, with figures 7.5-7.9 filling in the details. Numerical modeling in HFSS helped determine the final dimensions. The various materials in the model are shown explicitly in figure 7.5. The dimensions are defined in figure 7.6 and are given in table 7.1. As shown in figure 7.4e, the coupling structure can be divided into four regions based on their function: Region A) the semirigid cable, region B) the diameter transition, region C) the variable coupling, and region D) the cavity.

(a) Top View

(b) Side View

(c) Top View - Transparent

(d) Side View - Transparent

(e) Important Regions in Coupling Structure
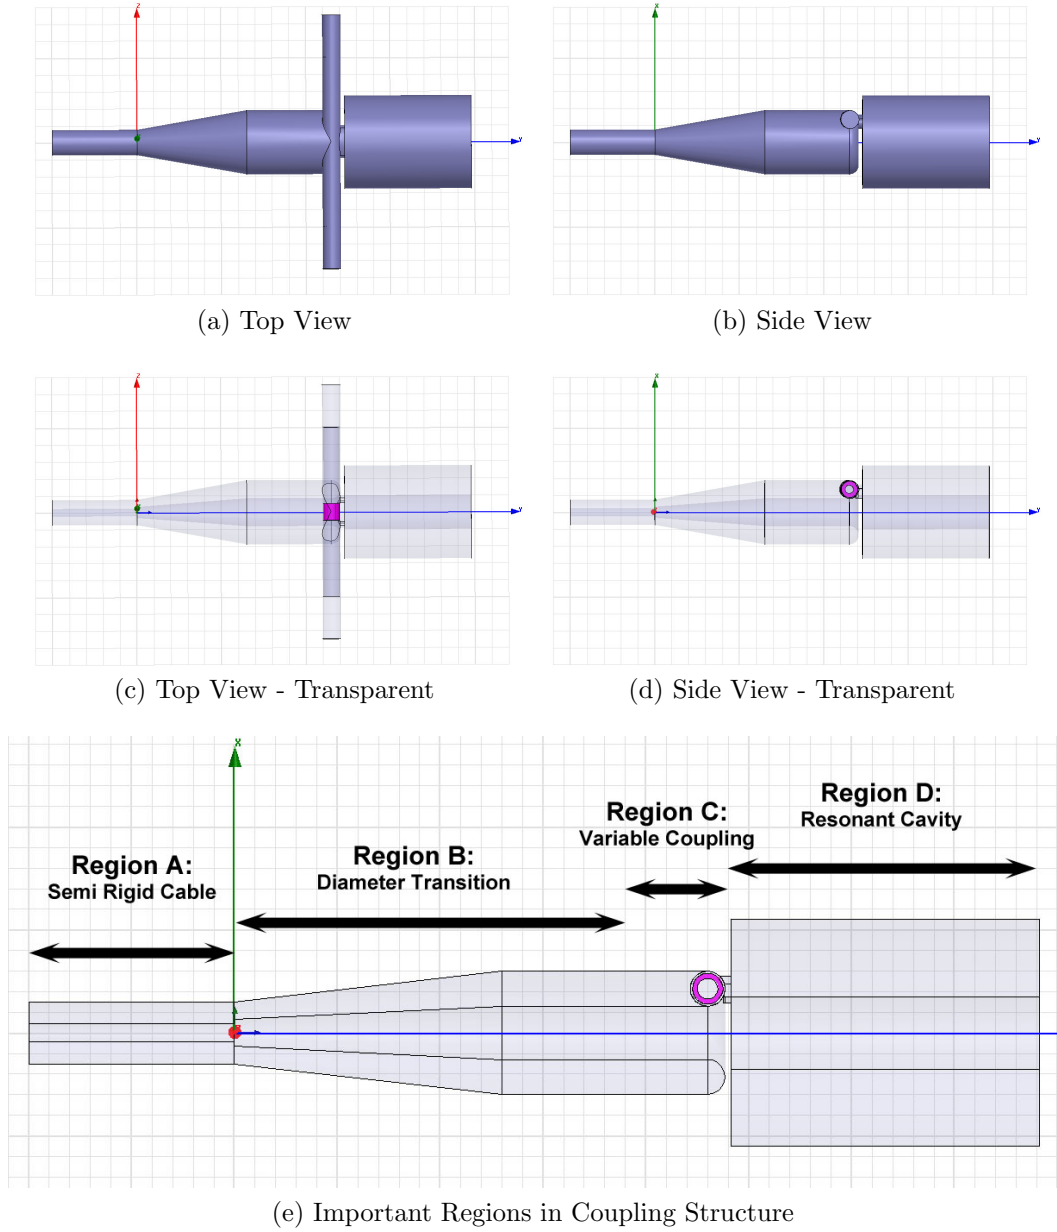
Figure 7.4: Various views of the cavity coupling structure modeled in Ansoft HFSS. The white space around the model is a conductor (gold) in which volumes of vacuum and dielectric (gray, shaded regions), as well as copper (pink), have been inserted. Figure 7.5 shows a breakdown of the various materials within the coupling structure. As labeled in panel e), the structure is divided into four regions based on their purpose. The dimensions of the model are defined in figure 7.6 and are given in Table 7.1. A detail of region C, the variable coupling region, is provided in figures 7.8 and 7.9.

(a) Top View - Vacuum

(b) Side View - Vacuum

(c) Top View - Copper

(d) Side View - Copper

(e) Top View - PEEK

(f) Side View - PEEK

(g) Top View - PTFE

(h) Side View - PTFE
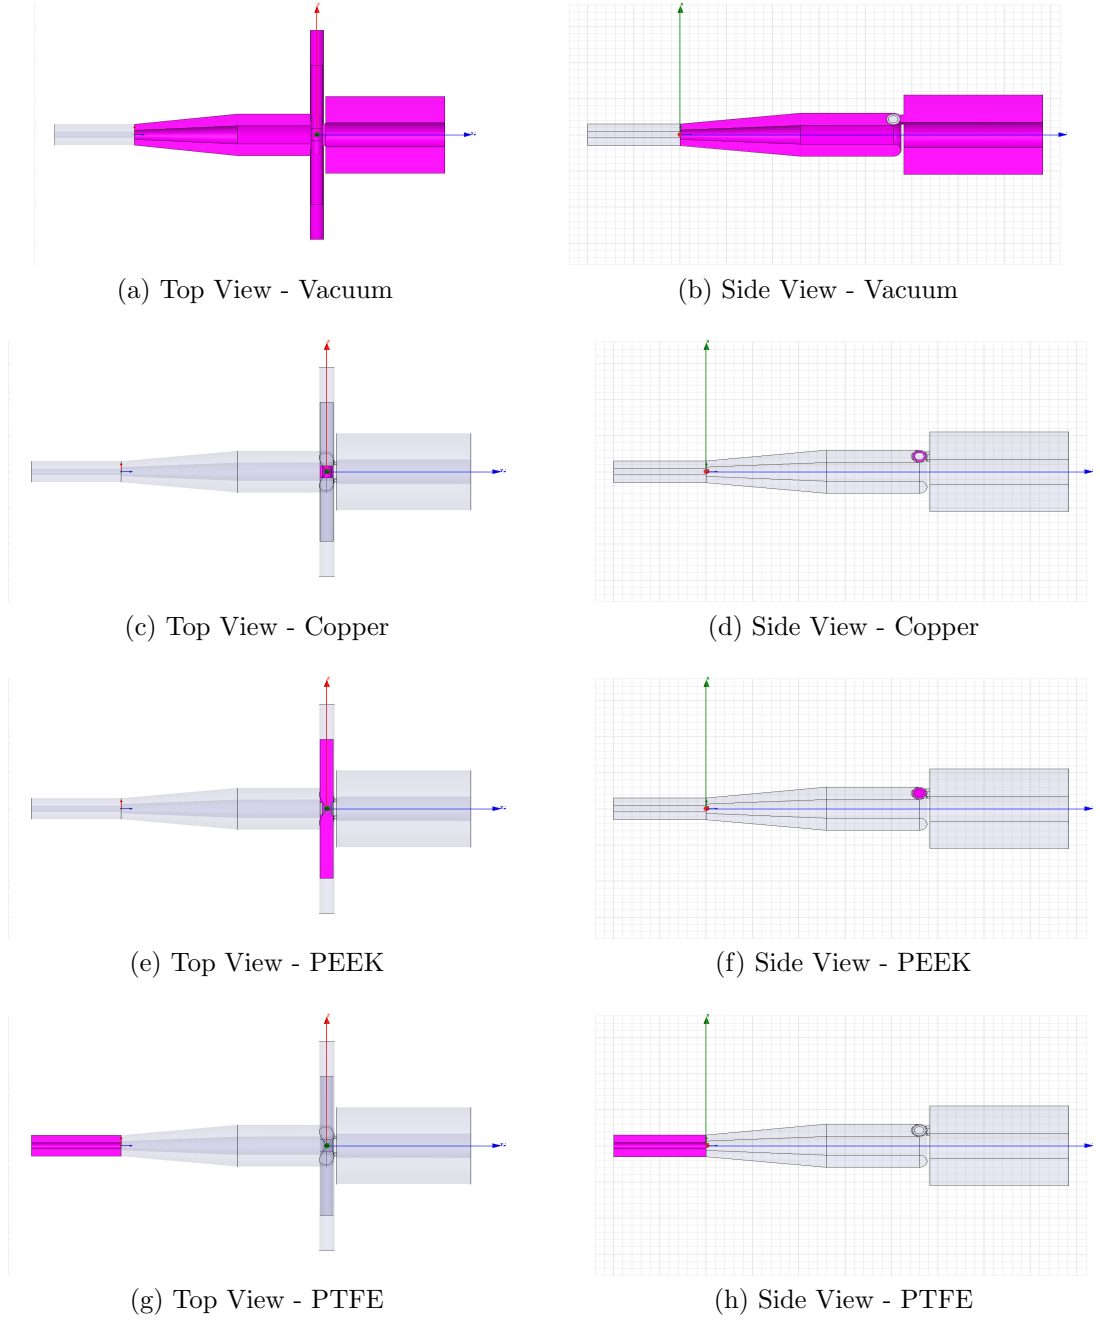
Figure 7.5: The different materials in the HFSS model. Any empty space surrounding the model is modeled as gold (not pictured). The structure is predominantly coaxial, with inner and outer conductors formed by the empty space in the model. The pink, shaded regions correspond to the volumes filled with the material in the panel label. The incoming coaxial cable (left side of images) is PTFE filled (panels g, h). The rest of the coaxial structure is vacuum filled (panels a, b), except for the coupling pill that is copper (panels c, d), and the PEEK rod (panels e, f) holding it in front of the cavity coupling hole.
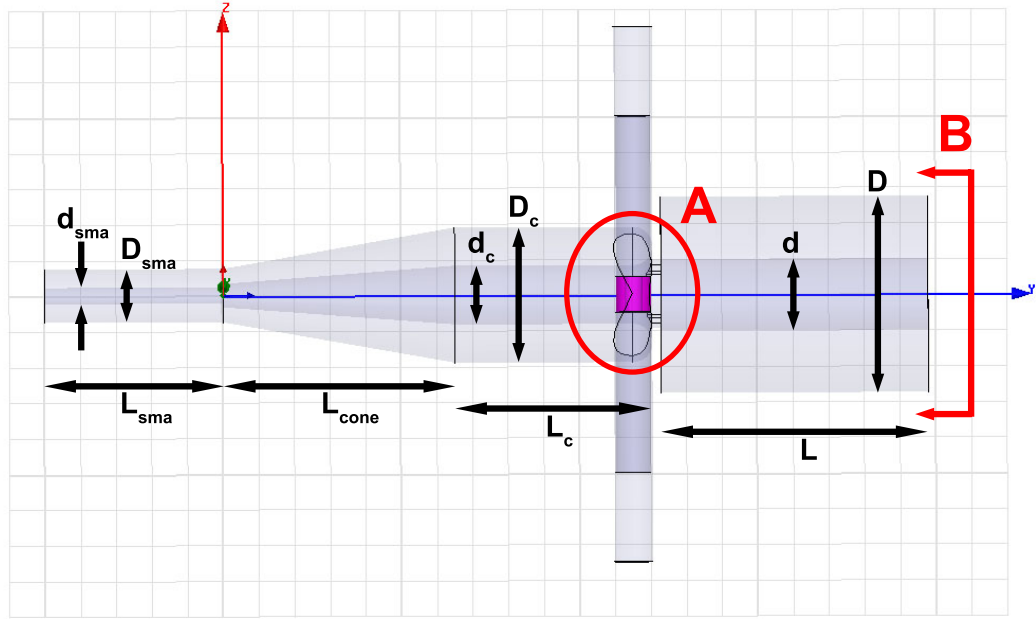
(a) Definition of model dimensions
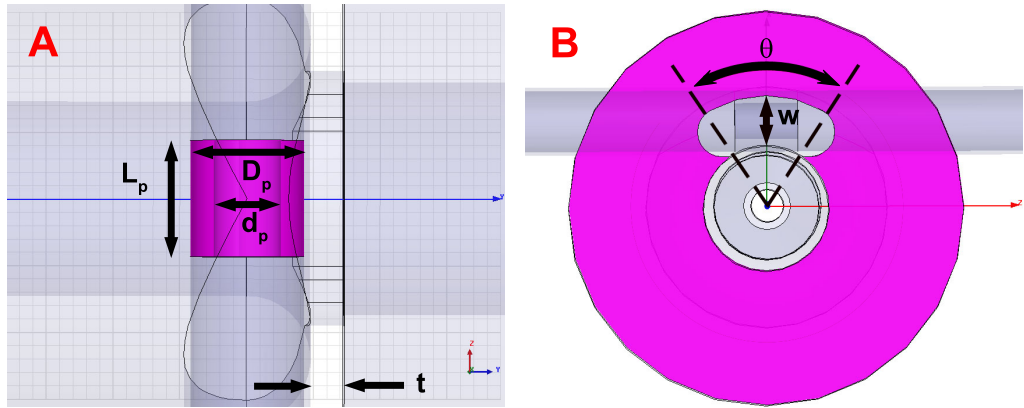
(b) Top View - Close-up of coupling structure

(c) Front View - Close-up of coupling structure

Figure 7.6: The dimensions of the cavity and coupling structure shown in figures 7.4 and 7.5. The values for the dimensions are listed in table 7.1.

Variable	4 mm design (mm)	3.5 mm design (mm)
d	4	3.5
D	11	-
L	15	-
d_c	3.3	2.92
D_c $(2.3d_c)$	7.59	6.72
L_c	10	-
w	1.25	1.38
θ	60°	-
t	0.5	-
d_p	1	-
D_p	1.95	1.7
L_p	2	1.75
d_{sma}	0.9	-
D_{sma} $(3.33d_{sma})$	3	-
L_{sma}	10	-
L_{cone}	13	-

Table 7.1: Table of dimensions defined in figure 7.6. The coupling structure was designed for two values of the cavity inner diameter, 3.5 mm and 4 mm, with the values for each design shown in the corresponding column. A dash in the value column indicates the value is the same for both designs.

Region A represents a standard, PTFE-filled, semiridged cable that brings the microwave to the system. The left end of the semirigid cable in the model is the input and output port; this is the only place energy can enter or leave the system, other than by being dissipated by ohmic or dielectric losses. HFSS models the fields at this port as being purely in the TEM mode. The length of region A does not effect the function of the rest of the structure, and it can be imagined as being infinitely long.

Region B is a transition region which serves to increase the diameter of the coaxial cable to facilitate the eventual coupling to the much larger cavity. It is vacuum filled, and the ratio of the outer to inner diameter is kept fixed as the overall diameter is increased in order to keep the coaxial cable impedance constant. At the boundary of regions A and B there is a step in the diameter of the inner conductor because of the dielectric change from PTFE to vacuum. The ratio of

outer to inner diameter jumps from 3.33 in PTFE to 2.3 in vacuum so that the impedance is maintained at 50 ohm in both regions. The diameter change is done slowly on the scale of the wavelength in order to minimize any reflection. Once the final diameter is reached, it is held fixed for some length in order to allow any evanescent modes to decay and purify the TEM mode.

Since the coupling structure would necessarily come after the cavity circulator in a μPDH setup, it is critical to minimize any reflections that do not come from the cavity itself. The scattering of the transition from region A through to the end of region B is tested by modeling it independently of regions C and D, as shown in figure 7.7a. The scattering coefficients between the input and output are displayed in figures 7.7b, 7.7c, and 7.7d, where the reflection is seen to be about −25 dB in the frequency range of interest.

Region C is where the incoming wave couples to the cavity. Due to the complex geometry here, its structure will be described step-by-step following the images in figure 7.8 for clarity. The oversized coax comes to an end in a circular grove (semi-torus) that is tangent to both the inner and outer conductors, and is separated from the cavity by a thin metallic wall - figures 7.8a, 7.8e, and 7.8i. The semi-torus serves to concentrate the magnetic field nearest to the cavity and helps minimize the size of the necessary coupling hole.

An arc-shaped wedge is cut into the thin wall, joining the coax to the cavity - figures 7.8b, 7.8f, and 7.8j. The bottom edge of the wedge is tangent to the inner conductor of the cavity, where the cavity magnetic field is highest, and the diameter of the incoming coax is chosen so that the semi-torus is centered on the wedge in order to maximize the overlap of the leakage field with the cavity field.

A cylindrical thru-hole is cut perpendicular, but offset, to the axis of the coax in such a way that it is symmetric about the wedge and concentric with the circle formed by the intersection of the torus and a plane containing the coax axis - figures 7.8c, 7.8g, and 7.8k. A short, copper tube, the "pill" (shown in pink in
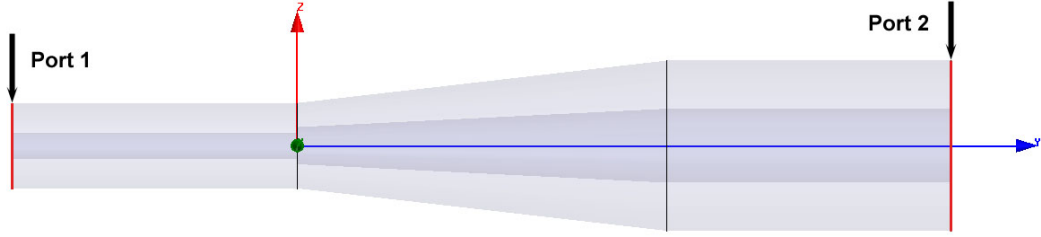
(a) Model of Transition, Regions A and B Only
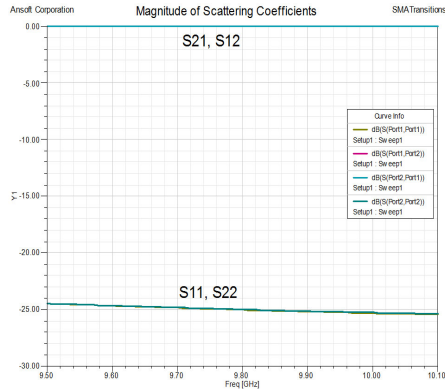
(b) Magnitude of Scattering Coefficients
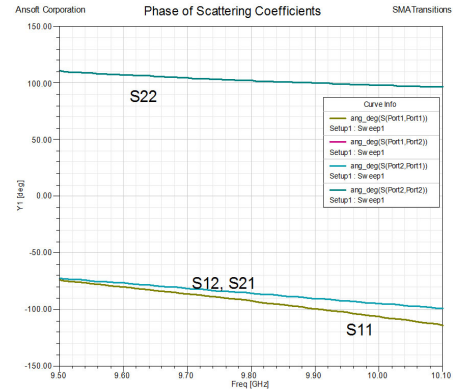
(c) Argument of Scattering Coefficients
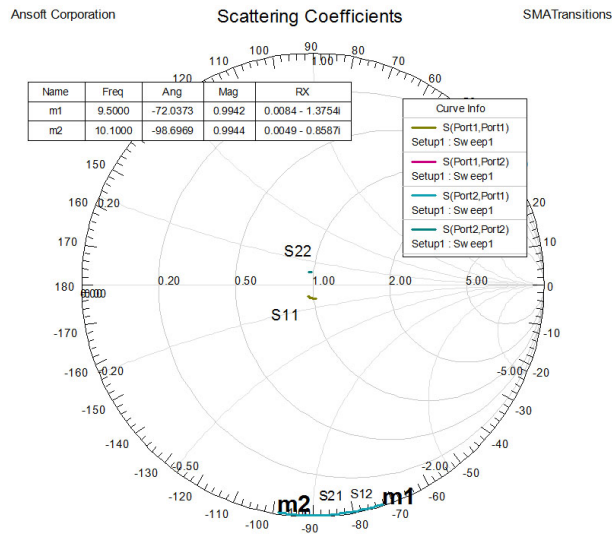
(d) Smith Chart with Complex Scattering Coefficients

Figure 7.7: The scattering from the transition from standard, semirigid, coaxial cable to an oversized, vacuum-filled coax. The reflection from the entire transition is about -25 dB. The geometry is shown in panel a), where ports 1 and 2 are defined. The curves in panels b), c) and d), which sometimes lie on top of each other, are labeled to avoid ambiguity.

80

(a) Top 1   (b) Top 2   (c) Top 3   (d) Top 4

(e) Side 1   (f) Side 2   (g) Side 3   (h) Side 4

(i) Front 1   (j) Front 2   (k) Front 3   (l) Front 4
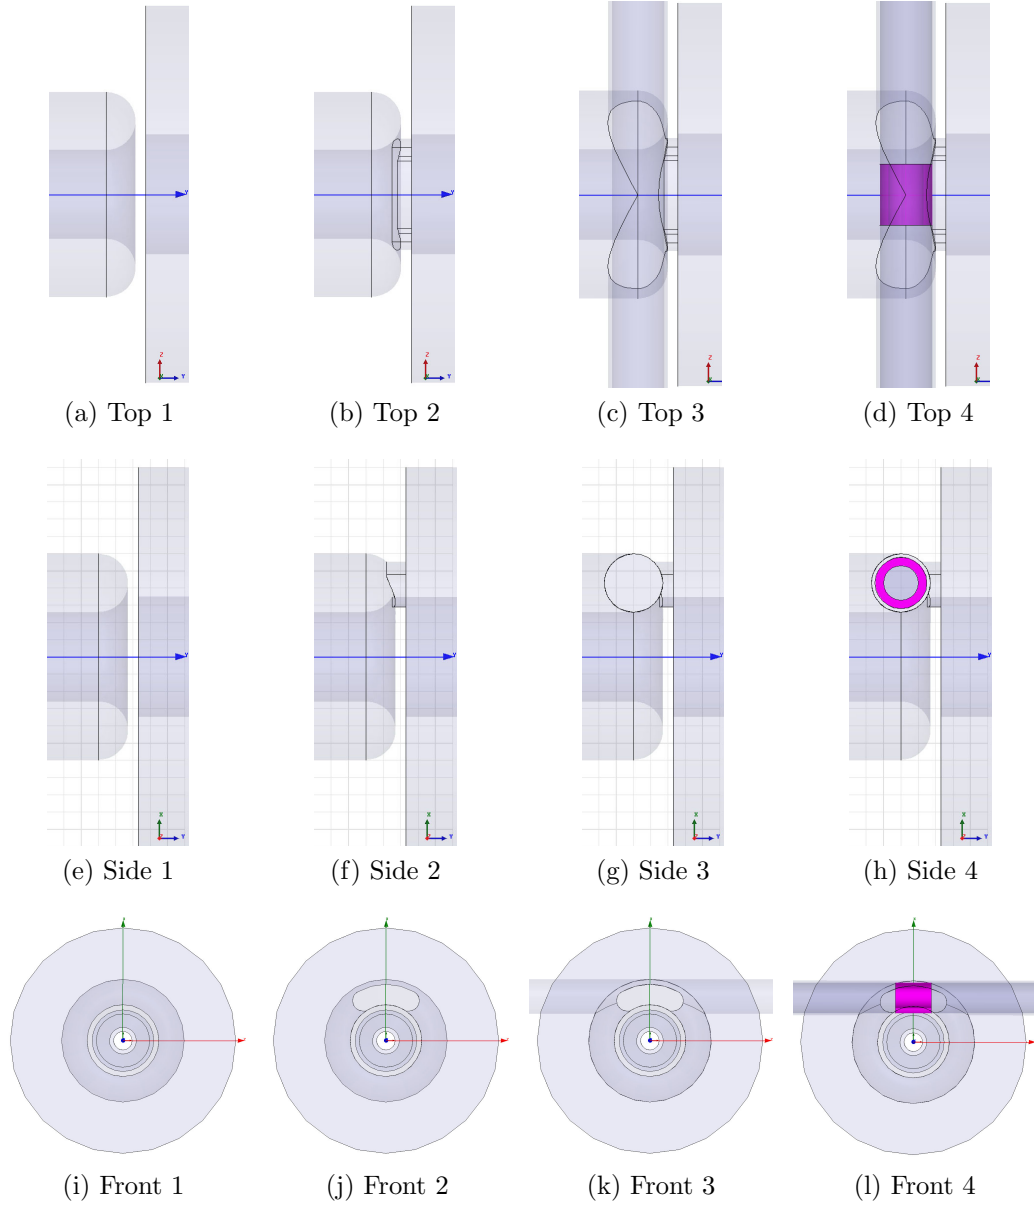
Figure 7.8: Various views of the coupling structure as its complex geometry is built up step-by-step. Panels a-d are a top view, e-h are a side view, and i-l are a front view.

(a) Pill Centered
Most Over Coupled

(b) Pill Offset
Intermediate Coupling

(c) Pill Removed -
Most Under Coupled
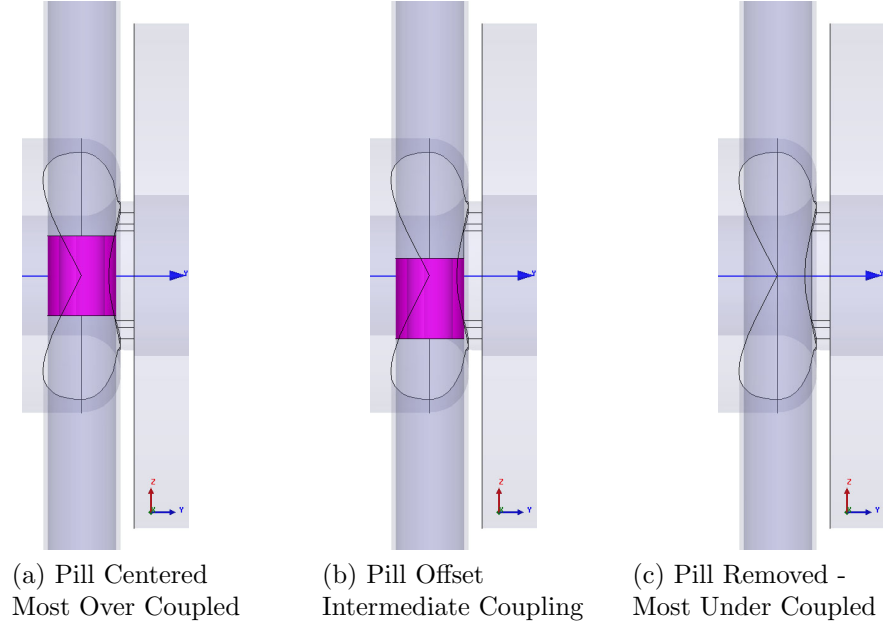
Figure 7.9: The position of the metallic pill in front of the wedge-shaped hole joining the coax with the cavity determines the degree of coupling between the two. As the pill concentrates the magnetic field in front of the hole, the cavity is the most overcoupled when the pill is centered on the hole, and becomes progressively undercoupled as the pill is moved to either side.

the figures), is supported by a dielectric rod inserted into the cylindrical hole and is positioned in front of the wedge - figures 7.8d, 7.8h, and 7.8l. The rod is made of polyetheretherketone (PEEK), chosen for its mechanical and stability properties[1]. The pill is of slightly smaller diameter than the cylindrical hole and does not touch any metal surfaces. The position of the pill along the axis of the cylindrical thru-hole can be smoothly adjusted and the overlap of the pill with the hole determines the degree of coupling between the coax and the cavity - figure 7.9. The pill serves to concentrate the magnetic field in front of the wedge-shaped hole, and consequently makes the hole appear bigger. This may be a bit counterintuitive, as one might naively expect that a piece of metal placed in front of a hole makes the hole appear smaller to an incoming microwave. As a result, the coupling is greatest when the pill is centered on the gap and decreases as the

---

[1]The dielectric constant and loss tangent are taken to be 3.3 and 0.004 at X-band [42].

pill is moved out in either direction. Thus, there are two positions where the cavity is critically coupled.

Figure 7.10 shows the scattering parameter S_{11} of the entire coupling structure, including the cavity, for various positions of the pill. As seen in figures 7.10e and 7.10f, when the pill is removed, the resonant frequency is 9.9508 GHz and the cavity is 6 dB undercoupled. The cavity is near critical coupling when the pill is about 0.5 mm off center from the gap - figures 7.10c and 7.10d. Here the coupling is better than 35 dB and the resonant frequency is 9.9352 GHz, a 15 MHz shift from the 6 dB undercoupled position. When the pill is centered, figures 7.10a and 7.10b, the cavity is 20 dB overcoupled and the resonant frequency is 9.93445 GHz. The 3 dB width of the critically coupled line is about 7 MHz, corresponding to a Q of 1420. This is a factor of 1.9 down from the unloaded Q calculated in table 6.1, very close to the expected factor of 2 difference. The shift of the resonant frequency from the unloaded cavity to the critically coupled cavity is 54 MHz.
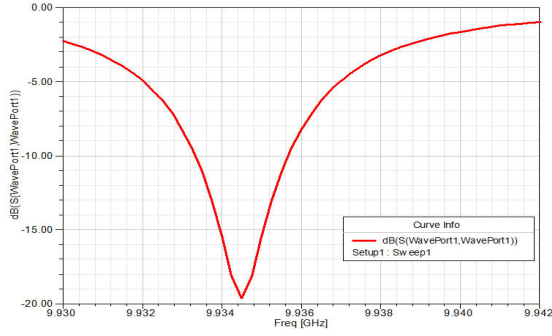
(a) $|S_{11}|^2$ - Pill Centered
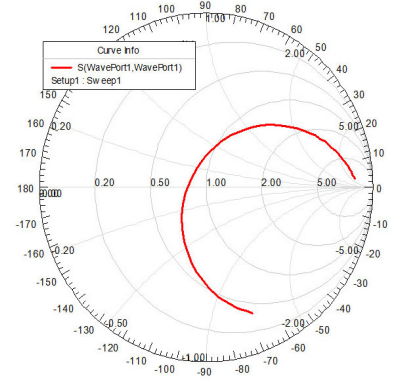
(b) S_{11} - Pill Centered
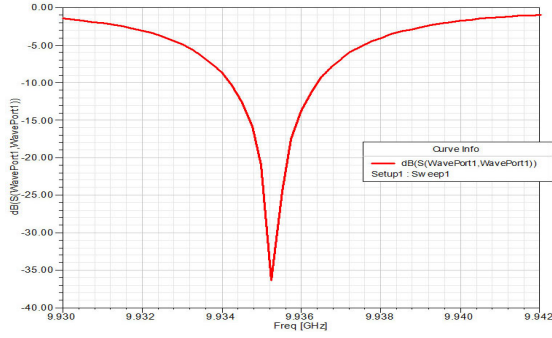
(c) $|S_{11}|^2$ - Pill Offset by 0.5 mm
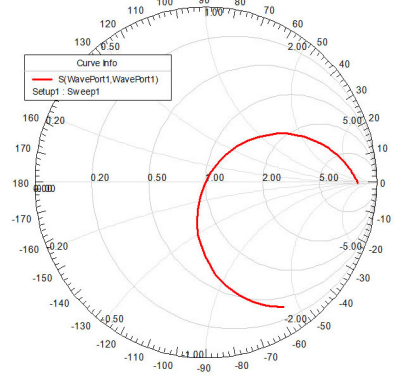
(d) S_{11} - Pill Offset by 0.5 mm
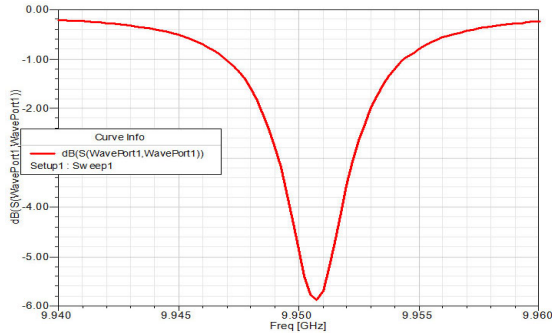
(e) $|S_{11}|^2$ - Pill Removed
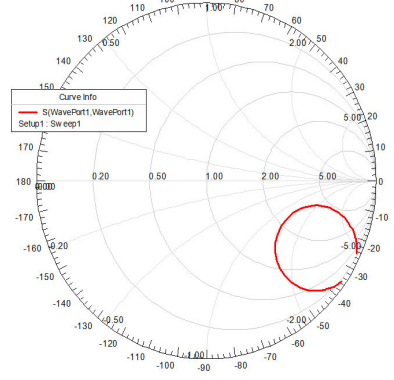
(f) S_{11} - Pill Removed

Figure 7.10: The scattering parameters of the cavity as the overlap of the pill and coupling hole is varied. In panels a and b, the pill is centered on the gap, the most overcoupled position. At $\sim$ 0.5 mm to either side of center, the cavity is critically coupled - panels c and d. When the pill is removed completely, the cavity is the least coupled - panels e and f.

CHAPTER 8

Mechanical Design of the Frequency-Distance Encoder
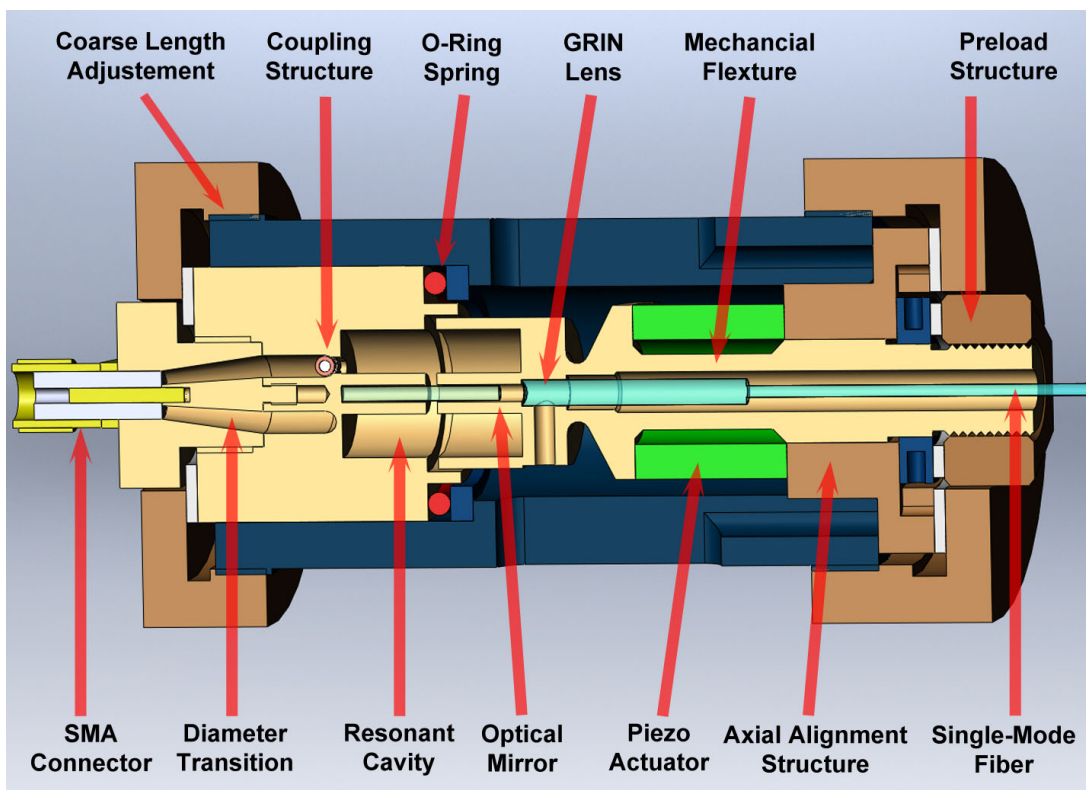
Figure 8.1: The internal structure of the frequency-distance encoder.

There are several important considerations to take into account in the mechanical design of the frequency-distance encoder. 1) The volume of space that contains any appreciable microwave field must be a faithful reproduction of the simulated model, described in section 7.2. The biggest deviation from the model is that the cavity must be cut somewhere in order to tune its length. To mini-

mize the perturbation to the resonant mode, the cut must be made where there are no currents in the walls - at the antinode of the electric field, right down the middle of the cavity. 2) A coarse (0.05-0.3 mm) adjustment of the cavity length is necessary in order to match it to the frequency of the source. The source is tunable over a range of ~ 25 MHz, corresponding to a length of about 40 µm, so the cavity length needs to be set within this range. 3) A smooth and repeatable fine-scale ($\lesssim 10$ µm) length adjustment is necessary to lock the cavity to the microwave frequency. The moving mass should be minimized in order to achieve a high-bandwidth feedback loop. 4) The two half-cavities need to be aligned to the required tolerance. Finally, 5) the length and/or motion of the microwave cavity needs to be transferred to some application to serve a useful purpose. To this end, it is advantageous to provide mounts on the microwave cavity to make the relevant points in the application lie within the planes of the microwave cavity walls. In this way, the precise and stable distance between them can be transfered to the application.

Figure 8.1 is a model of a frequency-distance encoder designed with an internal optical interferometer to measure its motion. The physical realization of the device is shown in figure 8.2. The mechanical design of the device can be separated into two functionally different components. On one side is the microwave portion of the device - the microwave input, variable coupling, and the "fixed" half of the resonant cavity. On the other is the "moving" half of the resonant cavity, incorporating a mechanical flexture, piezo actuator, and preload structure. The two halves are passively aligned with a cylindrical tube and are held in place with end caps that screw onto the outside of the tube.

8.1   The Fixed Half-Cavity

The microwave input is a standard SMA connector that transitions to a physical realization of the structure described in section 7.2. The input structure was made in several pieces, where possible, to simplify the machining. As the shorted end of the coax input as well as the corners of the resonant cavity hold high current, they were made out of a single piece of metal to minimize ohmic losses. All pieces that contain microwave current were machined out of aluminum 7075 and were subsequently gold plated. The fixed half-cavity assembly is shown in figures 8.2b and 8.2c.

The half-cavity assembly slides into a hard-anodized aluminum tube so that its axis is well defined. Dowel pins located on its outer face slide along groves in the tube, preventing any rotation. The tube has a step in its inner diameter against which a spacer ring of adjustable thickness is placed, followed by a rubber o-ring. The half-cavity slides into the tube until it presses against the o-ring, and is held in place by an endcap that screws onto the outside of the tube. Tightening the endcap compresses the o-ring, and provides a coarse adjustment of the microwave cavity length. A PEEK washer between the endcap and the half-cavity protects the metal surfaces from direct, metal-to-metal rubbing and wear.

(a)
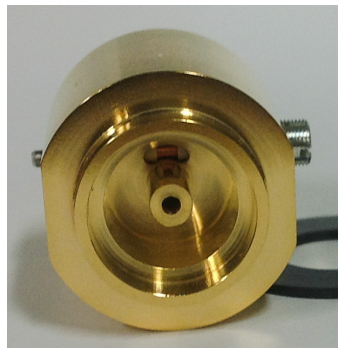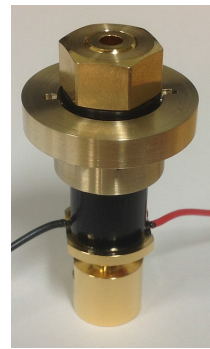
(b)   (c)   (d)
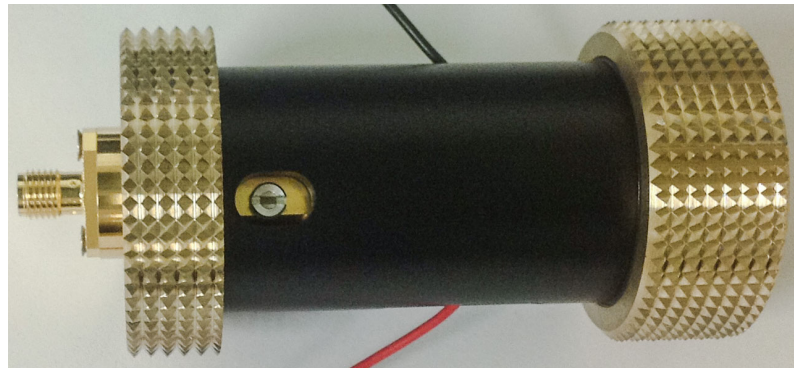
(e)

Figure 8.2: The individual components (a), intermediate assemblies (b,c,d), and completed assembly of the frequency-distance encoder. Figure 8.1 shows the internal structure.

8.2 The Moving Half-Cavity

The other half of the microwave cavity is designed for fast, small-scale motion. Its position within the device is visible in figure 8.1 and can be viewed alone in figures 8.3a and 8.3b. The half-cavity is machined onto the end of a mechanical flexture designed to elongate under the force of the piezo. The diameter of the piece varies along its length because the different regions serve different purposes. After the cavity, the piece narrows considerably, then taperes back out to the piezo diameter. This narrow point isolates the cavity from the strain of the flexture portion, minimizing deformation of the cavity wall. After reaching the piezo diameter, there is a step change to a much smaller diameter that stays constant for a few millimeter. This narrow, thin-walled region is designed to elongate under axial stress, discussed below. Then the diameter increases again and stays constant for another few millimeter. This region serves to align the half-cavity with the tube axis via an intermediate piece. A region with a square cross section follows, and then some threads to hold a nut.

The entire moving half-cavity assembly is visible in figure 8.2d. First a cylindrical piezo tube is slid over the flexture until it comes to rest on the flat surface. A piece follows that presses against the piezo, and aligns the half-cavity axially within the tube. Then comes a metal washer with a square inner hole that fits over the square-cross section region of the flexture. A standard-shaped PEEK washer

Property	Value	Units
Elastic Modulus	$7.2 \cdot 10^{10}$	N/m^2
Poisson's ratio	0.33	-
Shear Modulus	$2.69 \cdot 10^{10}$	N/m^2
Mass Density	2810	kg/m^3
Tensile Strength	$5.7 \cdot 10^{8}$	N/m^2
Yield Strength	$5.05 \cdot 10^{8}$	N/m^2

Table 8.1: The mechanical properties of aluminum 7075 as used in the Solidworks finite element study.

and a bronze nut complete the preload structure. Tightening the nut preloads the piezo so that it is always under compression. The piezo is not designed to support any shear stress, and consequently any torque that may be transmitted to it from tightening the nut could damage it. Because the washer with the square inner hole cannot rotate relative to the piezo, it cannot transmit any torque, and protects it. The PEEK washer prevents the two metal surfaces from rubbing when the nut is tightened and reduces wear. The assembly is inserted into the tube to align it with the other half-cavity and is held in place by another endcap.

Finite element analysis in Solidworks is used to determine the deflection of the flexture under load, as shown in figures 8.3b, 8.3c and 8.3d. Aluminum 7075 is chosen as the flexture material due to its light weight, stiffness, and yield strength, and is gold plated for passivation and electrical conductivity. A list of its mechanical properties is given in table 8.1. Under a 100 N load, the cavity deflection is 1.2 µm and the highest stress in the piece is 12.5×10^6 N/m^2, a factor of 40 less than the yield strength. The spring constant is therefore 83 N/µm and the flexture has an elastic range of 48 µm.

The stiffness of the PEEK washer is much higher than the stiffness of the flexture, and can be ignored. Its stiffness can be approximated using Hooke's Law,

$$\frac{F}{A} = E\,\frac{\Delta t}{t},\tag{8.1}$$

where F is the force acting over the area of the washer $A \approx 106$ mm^2, $E = 3.6$ GPa is Young's Modulus for PEEK, and $t = 1$ mm is the thickness of the washer. The effective spring constant relating force to the change in thickness is $EA/t = 380$ N/µm, more than a factor of 4 greater than the spring constant of the flexture.

8.3   The Optical Interferometer

In order to measure the motion and performance of the frequency-distance encoder, an optical interferometer is placed down the axis of the cavity. A single mode fiber brings a laser beam to a 3 mm focal length lens assembly, shown in figures 8.4a and 8.4b, that focuses it onto a metal mirror, figure 8.4c. The assembly consists of a gradient index (GRIN) lens optically bonded to a concave lens, whose curvature matches a surface of constant phase when the lens assembly is at the proper working distance. The mirror is formed by evaporating copper or aluminum onto the polished end of a 1 mm diameter quartz rod. As the glass-air interface does not have an anti-reflective coating, there is about a 4% reflection at the surface of the lens that serves as a reference for the interferometer. Some of the light that reflects from the mirror gets coupled back into the fiber and interferes with the reference beam. As the distance between the lens and the mirror is swept, there is a sinusoidal variation in the reflected power with a peridicity of $\lambda/2$.

The lens and quartz rod are held within the encoder assembly as shown in figure 8.1. The rod is attached to the fixed half-cavity in such a way that the uncoated end of the rod lies close to the plane of the cavity wall. The fiber and lens are held by a set screw in the moving half-cavity, again trying to set the lens in the plane of the cavity wall.
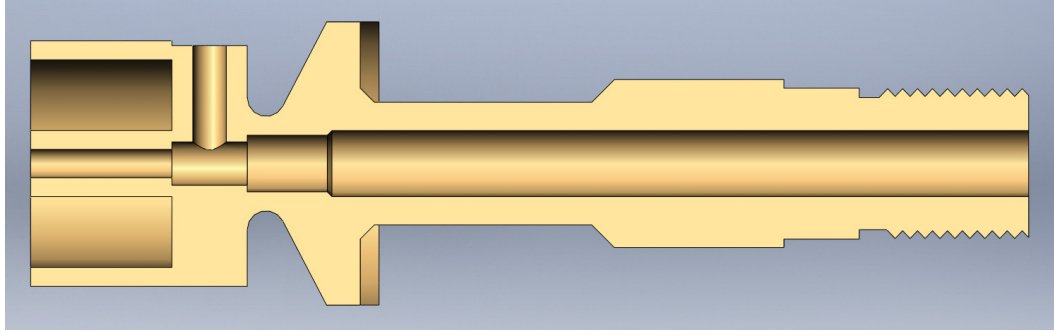
(a) Cross section of the moving half-cavity
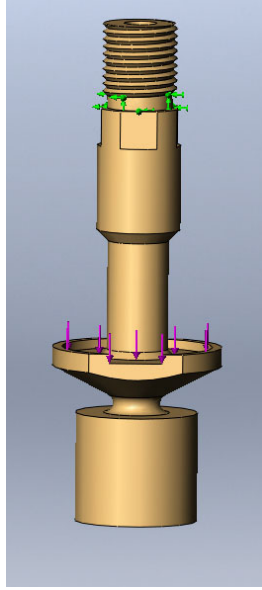
(b) Fixed points
and loads
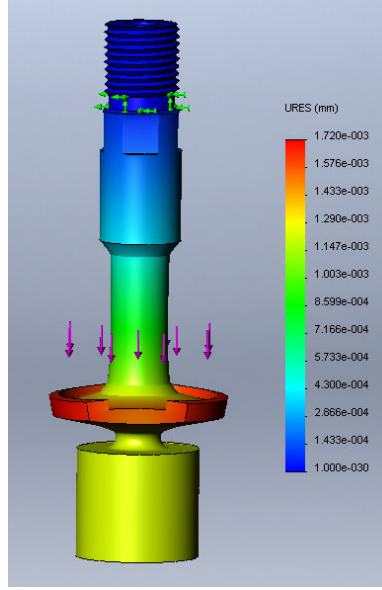
(c) Deformation under load
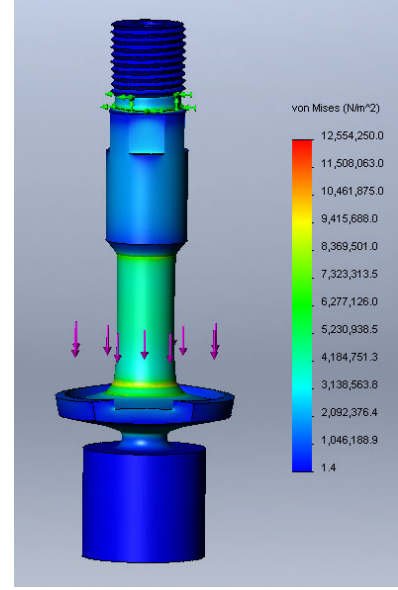
(d) Stress under load

Figure 8.3: The mechanical deformation of the flexture under load as determined by finite element analysis in Solidworks. The thin wall of the piece is designed to elongate under the force of the piezo, and provide the $\sim$ 10 µm range of motion. In the figures, the green arrows point to the face that is fixed and not allowed to move while the purple arrows represent the applied load. The material modeled is aluminum alloy 7075, and there is a 100 N force applied uniformily to the face perpendicular to the purple arrows. Panels a) and b) show the unloaded shape of the piece, while panels c) and d) display the loaded shape, with the deformation scaled by a factor of 3100. The magnitude of the deformation is shown by the color scale in panel c) while the stress in the material is shown by the color scale in panel d). Under the 100 N load, the cavity moves by 1.2 µm and the highest stress in the piece is a factor of 40 less than the yield strength. This implies a spring constant of 83 N/µm and an elastic range of 48 µm.
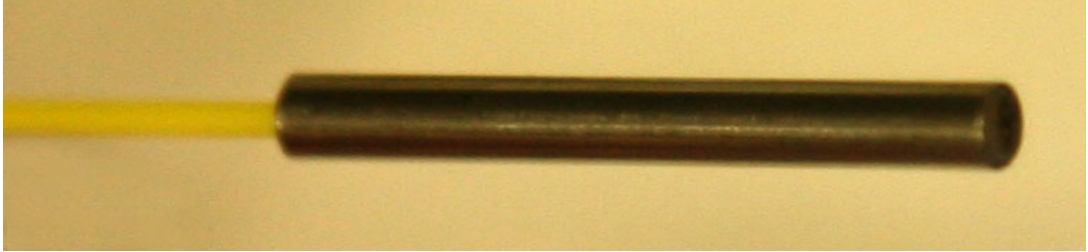
(a) Input fiber and lens assembly
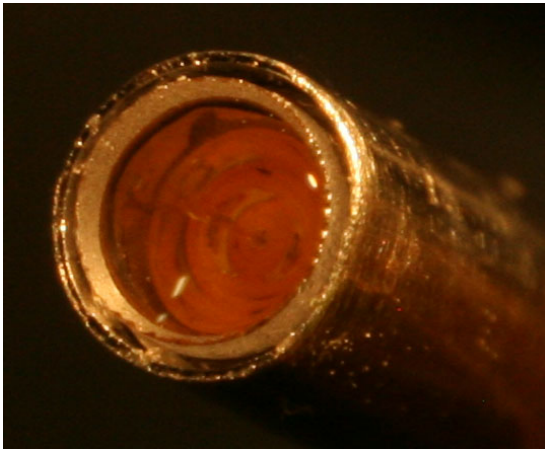
(b) Close-up of lens
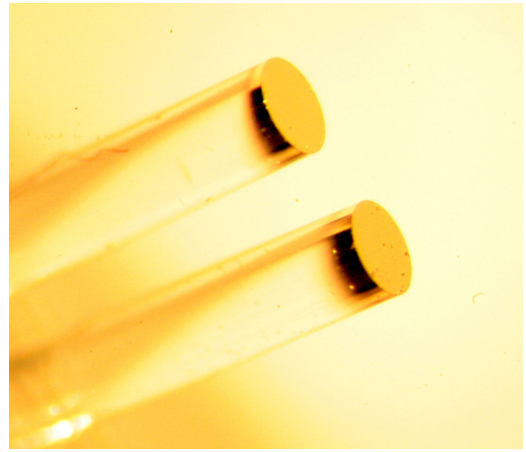
(c) Copper coated quartz rods

Figure 8.4: An optical interferometer is used to observe the motion of the frequency-distance encoder. A single mode fiber brings laser light to a lens assembly (a and b), that focuses it onto a copper or aluminum mirror on the end of a quartz rod (c). There is about a 4% reflection from the glass-air interface that serves as a reference beam for the interferometer. The transmitted light reflects off the metal mirror and is partially coupled back into the fiber by the lens, where it interferes with the reference. The distance between the lens and the mirror determines whether there is constructive or destructive interference.

CHAPTER 9

The Frequency-Distance Encoder System

In this section the performance of the frequency-distance encoder is described. Its range, resolution, and mechanical bandwidth are measured.

9.1 The Microwave Cavity

The first order of business is to characterize the cavity. Its Q is optimized, and the coupling structure is examined.

9.1.1 Cavity Q

The cavity was connected to a calibrated Agilent 8720ES Network Analyzer and the coupling screw was adjusted until critical coupling was achieved. Figure 9.1a shows the mode curve of the critically-coupled cavity on a linear scale. In the first iteration of the gold plating, the linewidth was measured to be 22 MHz, implying a Q of 450. From the HFSS eigenvalue calculations in table 6.1, the unloaded Q is expected to be 2685. When critically coupled, the Q is expected to drop by a factor of 2 to 1340 - so the measured Q is about 3 times too low.

This factor of 3 discrepancy can be due to the gold plating not being sufficiently thick, or to the presence of impurities in the gold. In gold plating aluminum, a layer of nickel is first deposited, followed by the gold. If the gold is too thin then there are significant losses in the nickel and the Q would be low. An estimate of the variation in Q that could be expected from such an effect is obtained by

| (a) Low-purity gold | (b) High-purity gold |
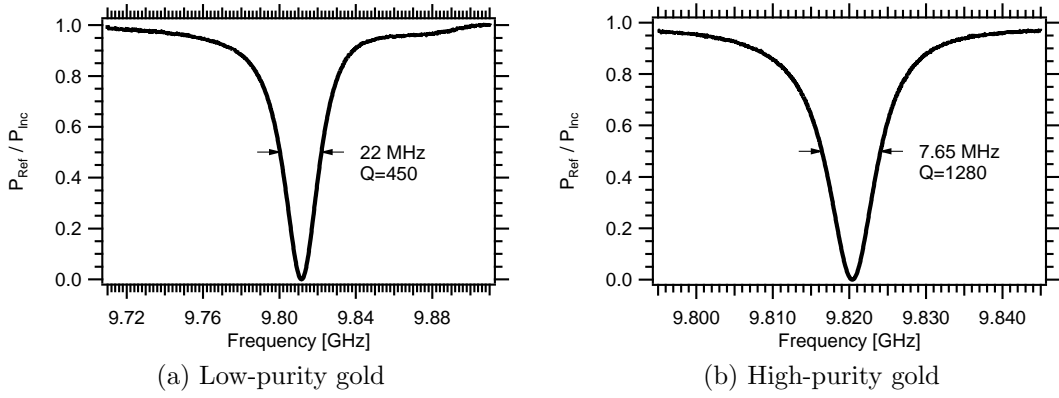
Figure 9.1: The cavity resonance shown on a linear scale. Panel a) shows the mode curve with the original gold plating. The linewidth is 22 MHz, implying a Q of 450. Panel b) shows the mode curve after plating with higher-purity gold. The cavity linewidth becomes 7.65 MHz, implying a Q of 1280. This is closer to the expected Q of 1340.

calculating the best and worst case scenarios. A pure gold cavity should have a Q of 1340, but a pure nickel cavity should have a Q of 33, a huge range of variation.[1]

To estimate the thickness of gold required, note that the Poynting vector of the electromagnetic field at the surface of a real conductor has a component pointing into the conductor, representing the energy flow into it. As both the E and H fields fall off exponentially with the skin depth, the Poynting vector is proportional to $e^{-2x/\delta}$, where x is the distance measured into the conductor and δ is the skin depth. This corresponds to -8.7 dB per skin depth. After four skin depth, the energy flow into the conductor has dropped by 3.3×10^{-4}, and after 5 skin depth, by 4.5×10^{-5}. Thus 4-5 skin depth should be sufficient to preclude any significant perturbation to the Q. The skin depth in gold at 10 GHz is 0.75 µm, so a gold thickness of 3-4 µm is required.

To verify the thickness of the plating, one of the gold plated pieces is inserted into a Nova 600 Scanning Electron Microscope (SEM) and Focused Ion Beam

---

[1]The cavity Q of nickel should be a factor of $\sqrt{\sigma_{Ni}\mu_{Au}/\sigma_{Au}\mu_{Ni}}$ less than the Q of the gold cavity, where σ_x is the conductivity of material x, and μ_x is its magnetic permeability. $\sigma_{Au} = 4.1 \times 10^7 \ \Omega^{-1}\mathrm{m}^{-1}$ while $\sigma_{Ni} = 1.45 \times 10^7 \ \Omega^{-1}\mathrm{m}^{-1}$. $\mu_{Au} = 1$, however $\mu_{Ni} = 600$. From this, a factor of 41 difference is expected, and verified by HFSS eigenvalue calculations.

(a) Cut in gold-plated surface   (b) Close-up of gold layer
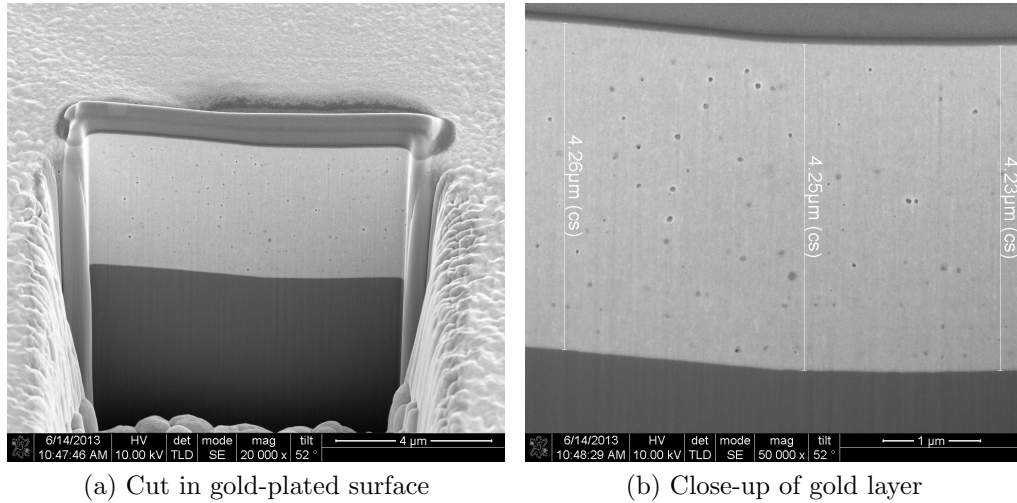
Figure 9.2: Focused ion beam images of the original gold plating. Panel a) shows a plane perpendicular to the surface of the piece that was exposed by ion beam milling. The lighter colored top layer is the gold, followed by a darker nickel layer. Panel b) is a close up of the gold layer with thickness measurements shown. As this is greater than 5 skin depth, it is sufficiently thick.

(FIB) system. An ion beam is used to cut a wedge into the piece, exposing a plane perpendicular to the surface as shown in figure 9.2a. The lighter colored top layer is the gold, followed by the darker nickel layer. The nickel is thicker than the depth of the cut, and therefore the aluminum is not visible in the image. Figure 9.2b is a close up of the gold layer with thickness measurements shown. The gold is 4.25 µm thick, or 5.67 skin depth, well above the requirement.

As the gold layer was sufficiently thick, the presence of impurities was suspected. To run an elemental analysis on the sample, it was placed in a Hitachi S4700 SEM with an Energy-Dispersive X-ray (EDX) probe. The EDX spectrum is used to determine the elemental composition of the material. Figure 9.3 shows the x-ray spectrum of the gold layer, at the point labeled "1" in the top left image. The machine identifies the peaks in the spectrum as being due to both gold and nickel content at point 1. By comparing the amplitude of the peaks, it is determined that there is $\sim 20\%$ nickel content by atom in the gold layer. Clearly this is way too high.
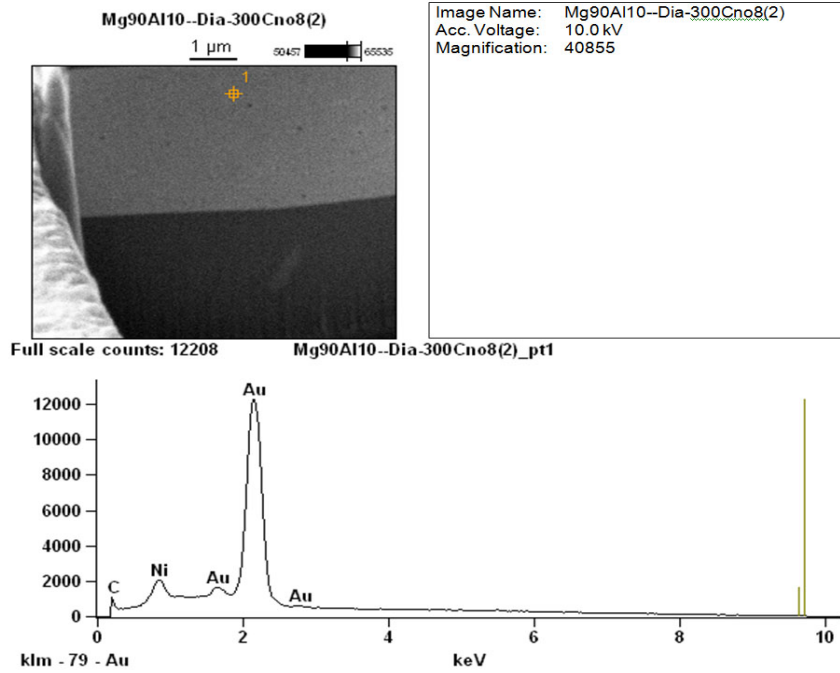
Figure 9.3: SEM image (top left) and EDX spectrum (bottom) of the original gold plating. Peaks corresponding to gold and nickel are visible in the spectrum. From the ratio of the peak amplitude, it is estimated that there is 20% nickel by atom in the plating. This is likely the cause of the Q being a factor of 3 lower than expected.

The cavity was sent to Metal Surfaces Inc for re-plating. The specification on the plating was nickel sulfamate plate 0.00005 inch (1.3 µm) minimum thickness followed by gold 0.0002 inch (5 µm) minimum thickness per MIL-G-45204C Type III (99.9%), Grade A hardness (soft). The higher purity plating is visually different than the lower purity plating as shown in figure 9.4. It appears as a dull, matte, yellow finish as opposed to a shiny metallic finish. FIB SEM images are taken to verify the thickness of the new layer, shown in figure 9.5. EDX data confirm that the re-plated surface has a higher purity of gold. The cavity linewidth, shown in figure 9.1b, is 7.65 MHz, and gives a Q of 1280, within 5% of simulation. The signal-to-noise ratio of the µPDH signal measured with the cavity was improved by the higher Q, shown in figure 9.6.

Figure 9.4: Some gold-plated aluminum pieces from two different batches. The shiny pieces on the left are from a run that left $\sim 20\%$ nickel in the gold. The dull pieces on the right are a 24 carat plating, specified to be 99.9% gold.

9.1.2 The next-highest mode

The cavity was designed so that the lowest frequency mode is the TEM mode, and that it is well separated from the next highest mode in order to prevent mode mixing. As discussed in section 6.2, the second mode is the TE111 mode. Figure 9.7 shows the cavity reflection as the frequency is scanned over a range of 10 GHz. The TEM mode is visible at 9.82 GHz, and the TE111 mode is visible at 16.75 GHz, as expected from the theory of section 6.2.

(a) Cut in re-plated surface   (b) Close-up of cut
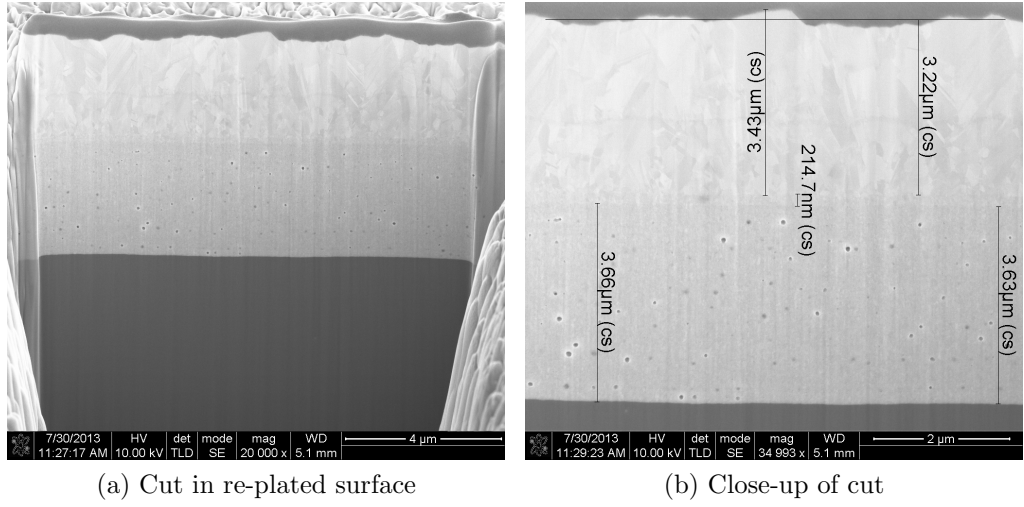
Figure 9.5: Focused ion beam images of the re-plated gold surface. Panel a) shows a plane perpendicular to the surface that was exposed by ion beam milling. The dark layer towards the bottom is the original nickel layer. Right above it is the original gold layer, followed by a slightly lighter new layer - the re-plated gold. Panel b) is a close up of the gold with thickness measurements shown. The crystalline structure of the purer gold layer contrasts with the porous structure of the impure layer below.
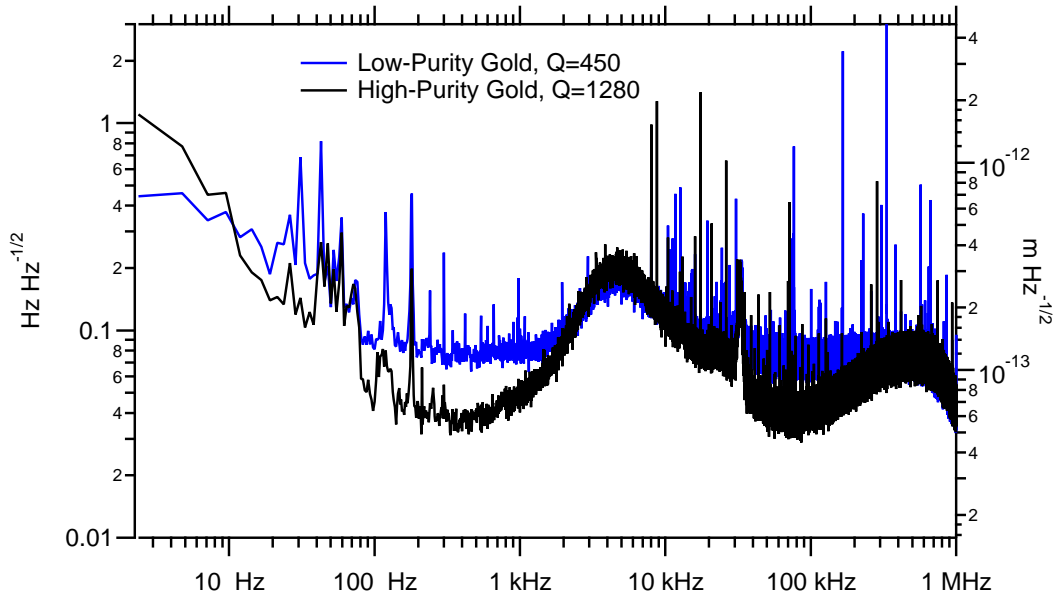
Figure 9.6: The μPDH noise spectrum measured with cavities with different gold plating. The higher purity gold triples the cavity Q, and helps lower the noise spectrum to the frequency-noise limit.

99
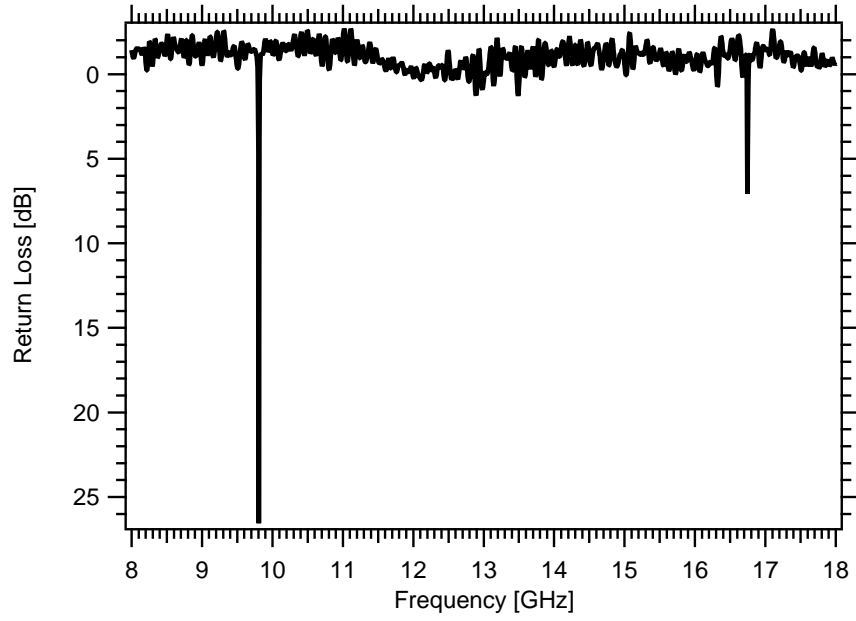
Figure 9.7: The two lowest-frequency resonant modes of the cavity. The TEM mode is at 9.82 GHz, and the TE111 mode is at 16.75 GHz, in agreement with the predictions of section 6.2.
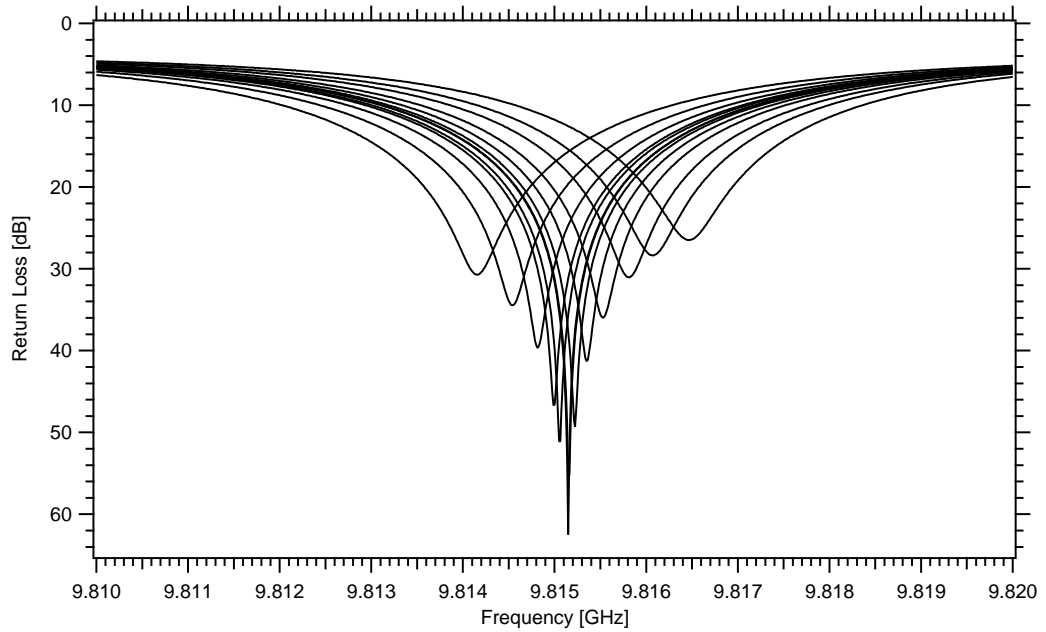
Figure 9.8: The cavity mode curve as the position of the coupling screw is changed.

9.1.3 Cavity coupling

The first test of the coupling structure was to critically couple the cavity. Since the range of coupling is extremely dependent on the dimensions of the coupling pill, a set of pills with diameters in the range 0.064-0.076 inch was made. It was experimentally determined that the pill with 0.071 inch diameter worked best[2], with larger diameters coupling too much, and smaller diameters not enough. The length of the cavity was set so that the resonance frequency was around 9.82 GHz. Coupling to 50 dB was straightforward and stable. At about the 60 dB level, it became sensitive to vibrations and pressure on the cable. 80 dB coupling was achieved, but only for short periods of time due to drift.

Figure 9.8 shows the cavity mode curve as the position of the coupling screw is changed. The curves on the right are undercoupled, and shift towards the left as the coupling is increased. A Lorentzian of the form,

$$|S_{11}|^2 = 1 - \frac{1 - |\Gamma_{min}|^2}{1 + 4\frac{(f-f_0)^2}{\delta f^2}},$$
(9.1)

is fit to the cavity mode curves. Here, f is the frequency, f_0 is the cavity frequency, δf is the Full-Width-at-Half-Max (FWHM), and $|\Gamma_{min}|^2$ is the reflection minimum. Recall that Γ_{min} is related to the coupling coefficient β through equation (3.5b). The variation of the fit parameters with the position of the coupling screw is shown in figure 9.9. The zero position of the screw is chosen to be at the location of critical coupling. The oscillations with screw position have a periodicity matching a full rotation of the screw (80 turns per inch), and indicate that the coupling pill was not axially symmetric with the screw.

---

[2]The 0.071 inch pill was used for the high-Q version of the cavity. For the low-Q version, the 0.0745 inch pill was needed.
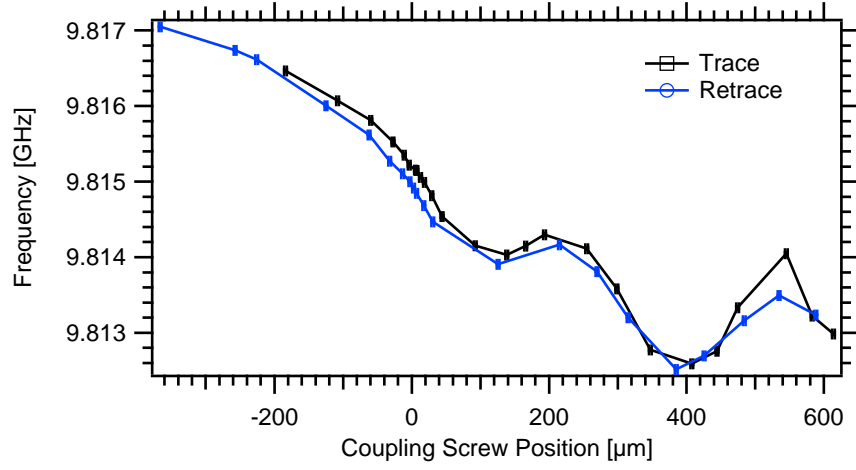
(a) Resonant frequency
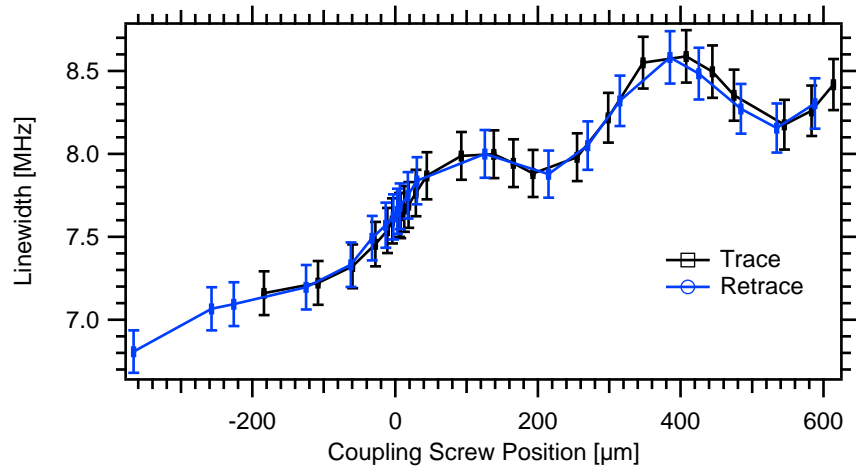
(b) Linewidth
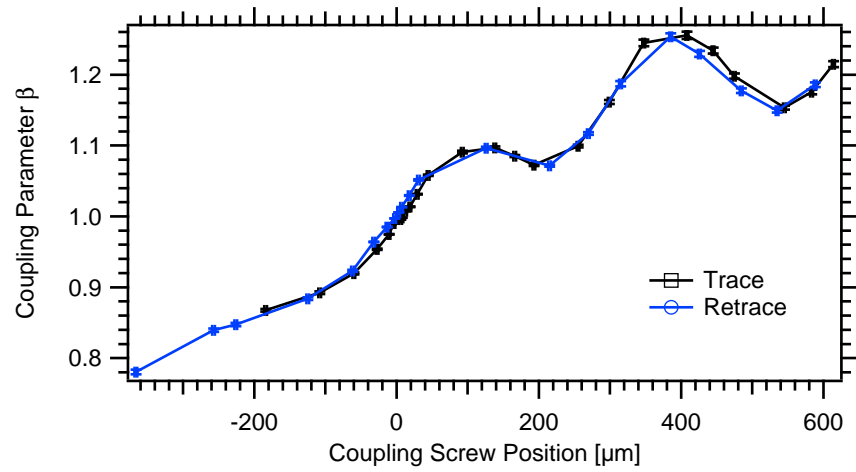
(c) Coupling Parameter

Figure 9.9: Various cavity parameters as the position of the coupling pill is changed. The curves are measured with the Q=1280 cavity.

9.2   The μPDH Signal Measured with the Coaxial Cavity

The μPDH signal is set up as described in section 5.3. The procedure is the same as it was with the EPR cavity. Since the Q of the coaxial cavity is lower than the Q of the EPR cavity (1280 compared to 2585), a reduction of signal for the same microwave power is expected. Increasing the microwave power partially makes up for this, but there are practical limitations to this. Increasing the carrier power by changing the DC level on the vector modulator baseband inputs is limited by the amplitude of spurious signals that grow faster than the carrier power. Increasing the sideband power is limited by the amplitude of the sideband-sideband beat signal, which grows faster than the sideband-carrier beat signal, saturating the preamplifier. In practice the carrier and sideband powers are set as high as possible without causing noticeable signal-to-noise degradation (about 12 dBm carrier and -12 dBm sideband power).

The noise spectral density of the μPDH signal measured with the coaxial cavity is shown by the black curve in figure 9.10. The broad peak at 5 kHz is recognized as the VCO-PLL phase noise peak. A few 60 Hz harmonics are visible. The peaks at 20 - 80 Hz are due to a fan in a nearby fume hood and vibrations of some system components. The μPDH noise spectrum measured using the EPR cavity is also shown (blue curve) as a best-case comparison. Since it is (mostly) frequency-noise limited, it is Q and power independent. The noise measured with the coaxial cavity below 2 kHz is slightly higher, due to the lower $Q\sqrt{P_c P_s}$ product, vibrations of system components, and low-frequency acoustic noise. The difference between 10-70 kHz is due to mechanical resonances of the coaxial cavity components.

In order to speed up the settling time of the VCO when changing frequency, the bandwidth of the VCO phase-lock loop was increased as described in appendix B. After doing so, the μPDH noise spectral density changed considerably, to that shown in figure 9.11. The noise density at frequencies up to 10 kHz was decreased
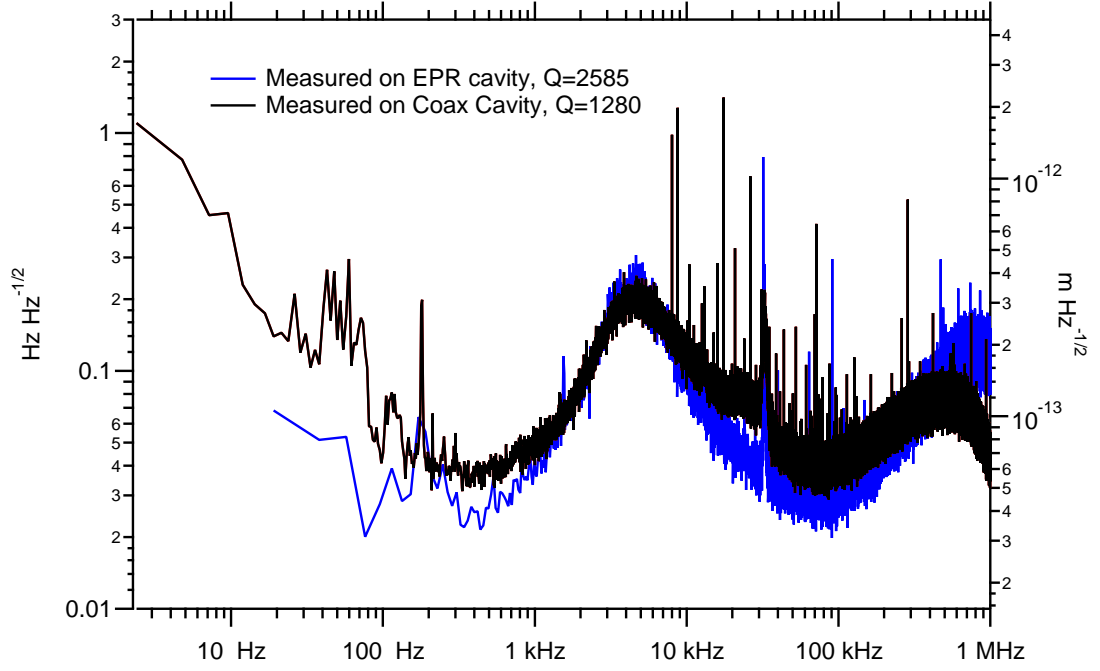
Figure 9.10: The noise spectral density of the µPDH signal measured with the coaxial cavity (black curve). The spectrum is measured at a carrier power of 15 mW and sideband power of 0.063 mW. The µPDH noise spectrum measured when using the EPR cavity is also shown for comparison (blue curve).

at a cost of increasing the noise density at frequencies greater than 10 kHz. Since the bandwidth of the cavity lock will be less than 10 kHz, this is a beneficial tradeoff for the frequency-distance encoder application. Figure 9.12 shows the noise in the time domain over a 5 ms period for 1, 3, and 10 kHz bandwidths. There is 1.6 Hz RMS noise, corresponding to 2.5 pm, in an 0.20-1 kHz bandwidth.
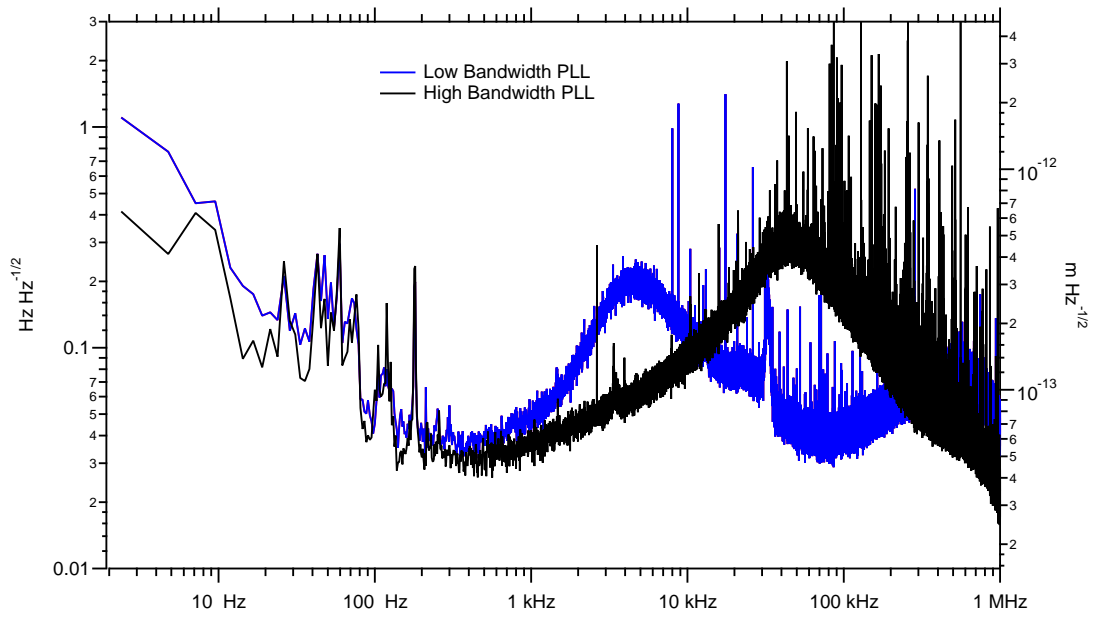
Figure 9.11: Increasing the bandwidth of the VCO phase-lock loop lowered the noise spectral density at frequencies below 10 kHz at a cost of increasing the noise spectral density at frequencies above 10 kHz.
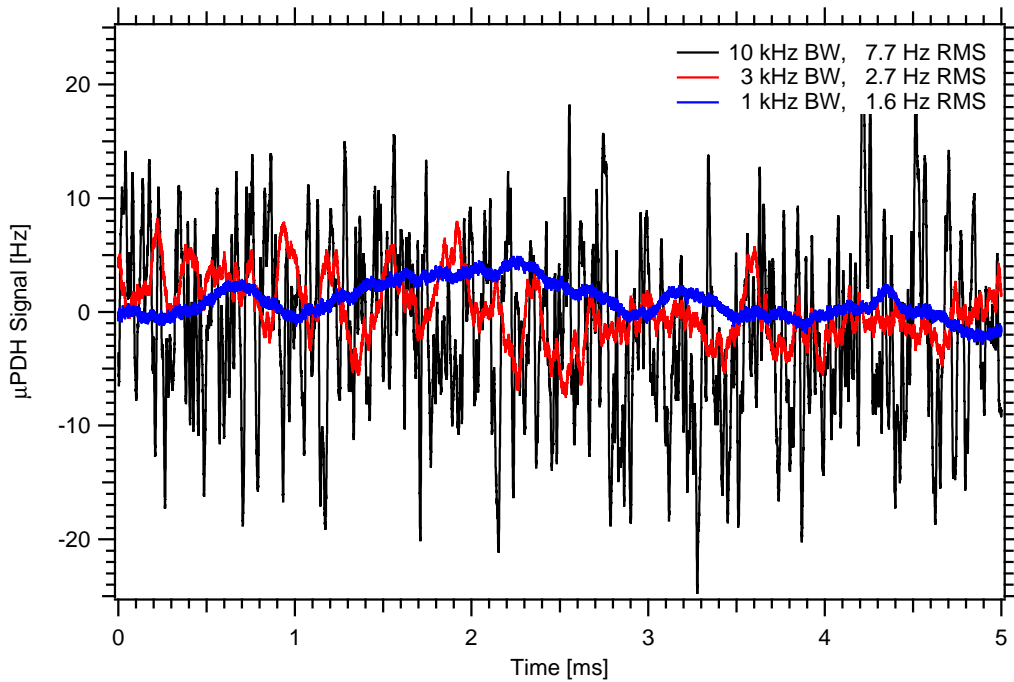
Figure 9.12: The noise of the µPDH signal in the time domain for three different bandwidths after increasing the PLL bandwidth. There is 1.6 Hz RMS noise in an 0.2-1 kHz bandwidth.

9.3 The Locked Cavity

The µPDH signal is fed back through a digital feedback controller and high voltage amplifier to the piezo to keep the cavity resonant frequency locked to the frequency of the source. As the source frequency is changed, the cavity follows. The interferometer is used to independently measure the motion of the cavity.

The feedback is done on an Asylum Research MFP 3D Atomic-Force Microscope controller. The µPDH signal is digitized using a 16-bit, 5 MHz ADC and taken to a digital signal processor that performs the filtering and proportional-integral gain feedback control. The signal goes to a 24-bit, 100 kHz DAC, followed by a high-voltage (-10-150 V) amplifier and taken to the piezo (Piezomechanic GmbH model HPSt 150/14-10/12).

9.3.1 Tuning range and resolution

The VCO has a tuning range of 20 MHz, corresponding to a distance change of 30 µm. The piezo, however, has an unloaded stroke of roughly 12 µm. In addition, when the piezo is properly preloaded, its range drops considerably. The stiffness of the piezo is specified as $k_1 = 250$ N/µm, and as discussed in section 8.2, the spring constant of the flexture about $k_2 = 83$ N/µm. Therefore, the range is expected to be reduced by a factor of $k_1/(k_1 + k_2) \approx 0.75$ to 9 µm.

Figure 9.13 displays the interferometer signal as the frequency is scanned over a range of 7.5 MHz, as well as the voltage applied to the piezo in closed-loop. The signal from both the forward (trace) and reverse (retraces) scans are visible. The piezo hysteresis curve is visible in the bottom panel of the figure. The range of motion of the cavity can be deduced from the $\lambda/2$ periodicity of the interference pattern and the known laser wavelength of 1310 nm. There are 18.7 periods visible in the figure, yielding a range of 12.2 µm. The period is 400 kHz, giving a frequency-to-distance conversion factor of 612 Hz nm^{-1} (measured with the low

106
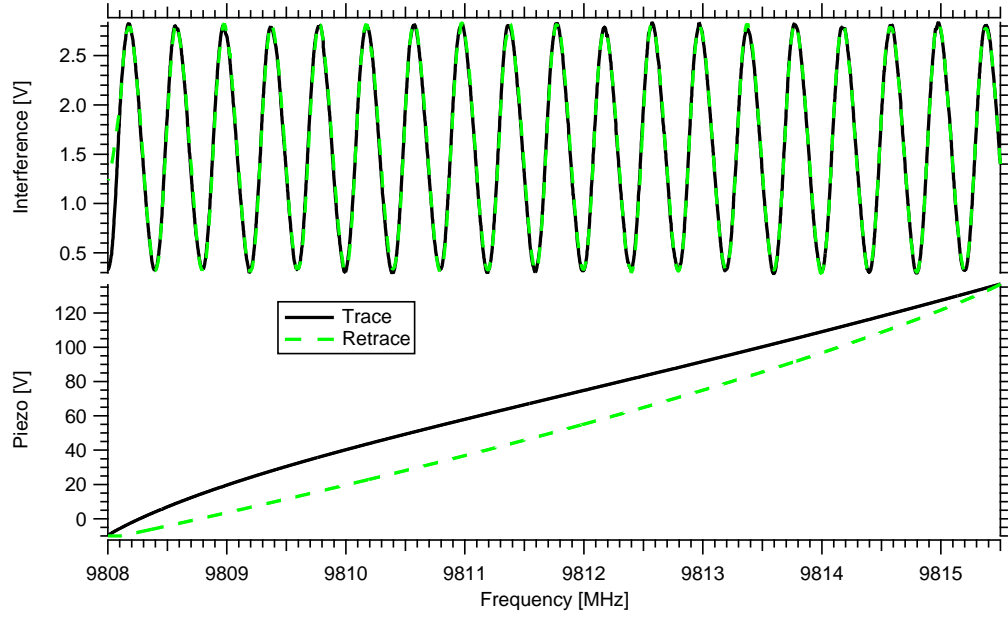
Figure 9.13: The interference pattern (above) and piezo voltage (below) during a full-range, forward (trace) and backward (retrace) sweep of the cavity. The $\lambda/2$ periodicity of the interference pattern gives the range of motion as being 12.2 µm. The closed-loop piezo voltage gives the piezo hysteresis curve.

Q cavity). This is to be compared to the theoretical value $2f^2/c = 642$ Hz nm^{-1}, to be discussed in detail in chapter 11. After further tightening the preload nut the range of motion dropped by a factor of 2 to about 6.2 µm. The interpretation is that the preload in the first case was not due to the flexure but to something softer. After tightening the nut, the piezo was properly seated and loaded against the flexure and the drop in the range of motion was observed. The factor of two difference in range implies that $k_2 \approx k_1$. This is likely a more accurate estimate for the stiffness of the flexure than the one given by Solidworks.

Assuming that the feedback controller is noise free, the resolution of the system is set by the µPDH noise level discussed in section 9.2, 1.6 Hz or 2.5 pm in a kHz bandwidth. In principle, the VCO frequency can only be set to discrete frequencies determined by the fractional-N PLL. However, as the AD9956 chip is used to lock the VCO to the reference through a DDS, the frequency can be set with 48 bit resolution. This corresponds to ~ 10 µHz steps, far better than the µPDH RMS
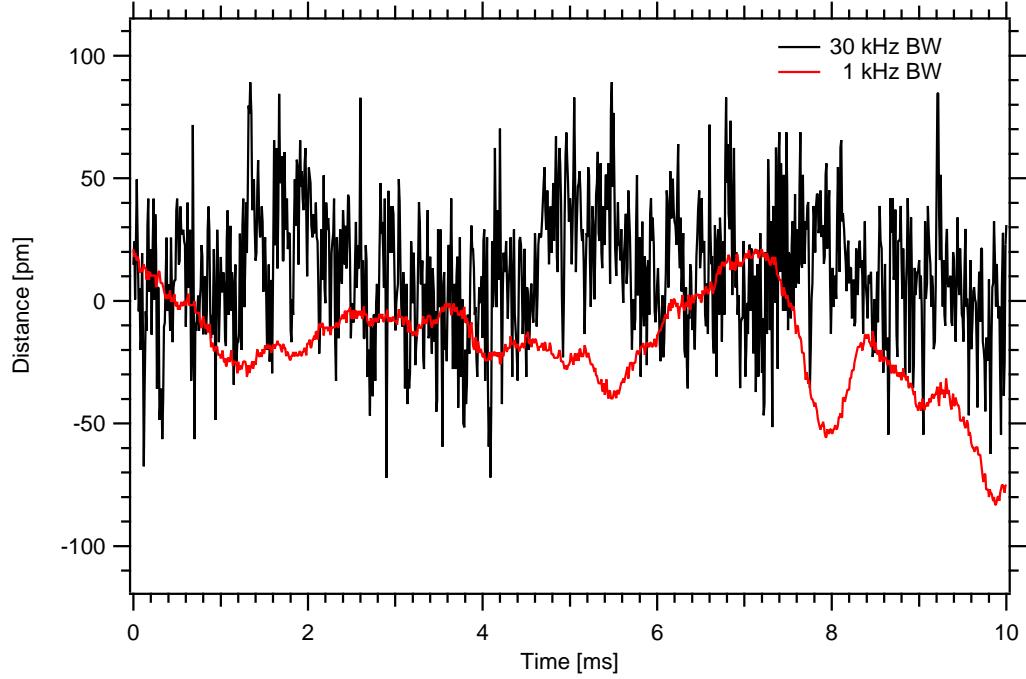
Figure 9.14: A sample of the interferometer noise in two different bandwidths over a 10 ms time interval.

noise level in any practical bandwidth. The range-to-resolution ratio is a few times 10^6.

9.3.2 Interferometer performance

The laser used in the interferometer is a QPhotonics QFBGLD-1300-5 single-mode-fiber-pigtailed laser diode at 1310 nm. To measure small-scale distance changes (less than $\lambda/4$), the cavity frequency is chosen so that the interferometer sits on a slope. Multiplying the slope of the curve by the frequency-distance conversion factor gives the volts-to-distance calibration. A sample of the interferometer noise is shown in figure 9.14. The RMS noise in 1 kHz bandwidth is about a factor of 4 larger than the µPDH noise.

Figure 9.15 shows the interferometer noise spectral density compared to the µPDH noise spectral density. The much poorer low frequency performance of the interferometer is due to its instability. It constantly drifts on the second to minute
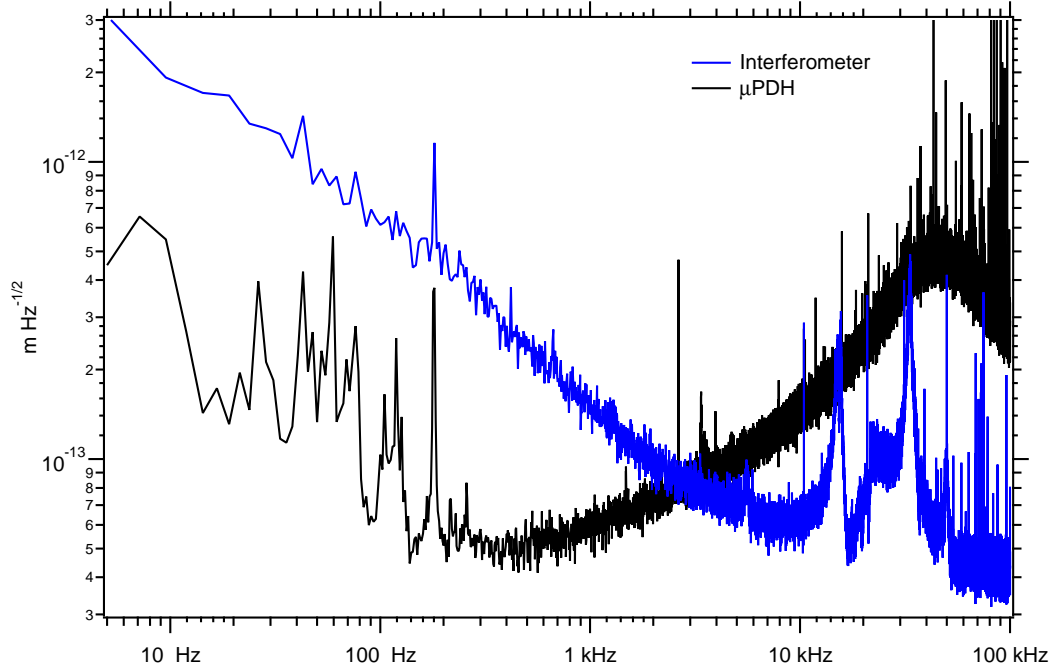
Figure 9.15: The noise spectral density of the interferometer signal compared to that of the μPDH signal.

timescale, and mode hopes occasionally. This is likely due to the 7.6×10^{-6} K^{-1} wavelength stability of the laser discussed in section 10.2.4, and can be reduced by placing the laser diode in a better controlled environment or stabilizing it by some other method. The large peaks between 5-50 kHz are mechanical resonances. They were also visible in the μPDH noise spectrum before the PLL bandwidth was increased.

Figure 9.16 shows the motion of the cavity as its frequency is stepped back and forth by 5 kHz. In open loop, there is a 5 kHz step in the μPDH signal as the setpoint is changed. The settling time of the μPDH signal is discussed in section 9.3.3. The interferometer signal stays flat as there is no motion of the cavity. The initial spike in the open-loop μPDH signal before it reaches the 5 kHz value reflects the VCO adjusting to the new set point of the PLL. In closed loop, the piezo adjusts the cavity length to bring the μPDH signal back to zero, as shown by its exponential decay. The resultant 8 nm step in the cavity length is visible

109
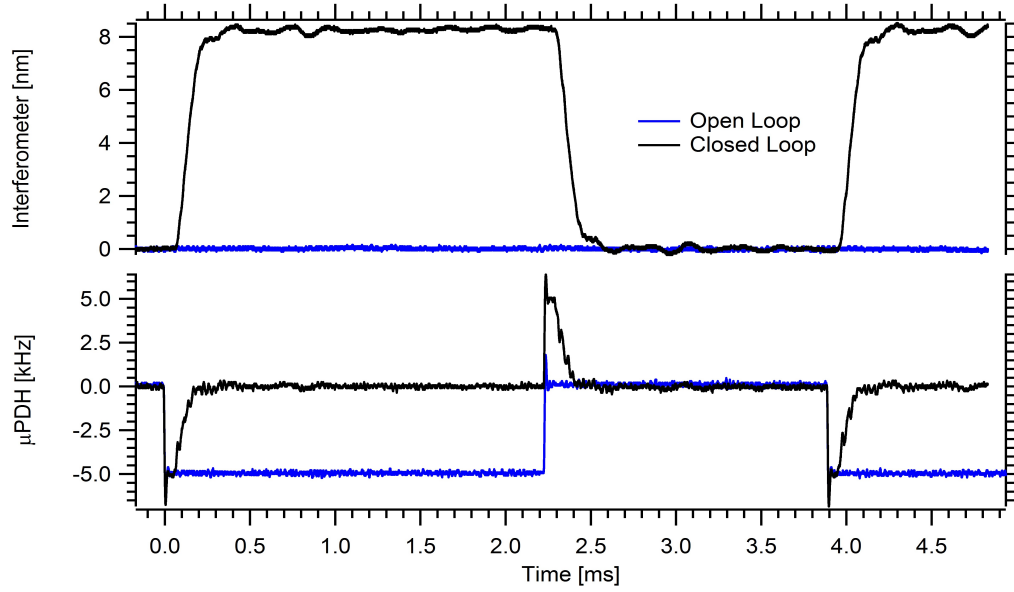
Figure 9.16: The interferometer and µPDH signals as the source frequency is stepped back and forth by 5 kHz. In open loop, the µPDH signal moves up to 5 kHz and the interferometer signal stays flat as there is no motion of the cavity. In closed loop, there is a step in the µPDH signal followed by an exponential decay back to zero as the piezo adjusts the cavity length to match the frequency. The interferometer signal shows the resultant motion of the cavity, a step of 8 nm. The interferometer signal was measured though a 10 kHz low-pass filter, while the µPDH signal was measured through a 100 kHz low-pass filter.

in the interferometer closed-loop curve.

From the measured the µPDH noise, the system should be able to achieve 2.5 pm steps in 1 kHz bandwidth, but unfortunately the interferometer is not capable of seeing them given its instability and noise. In partial consolation, 40 pm steps are visible in figure 9.17, created by stepping the cavity frequency by 25 Hz.
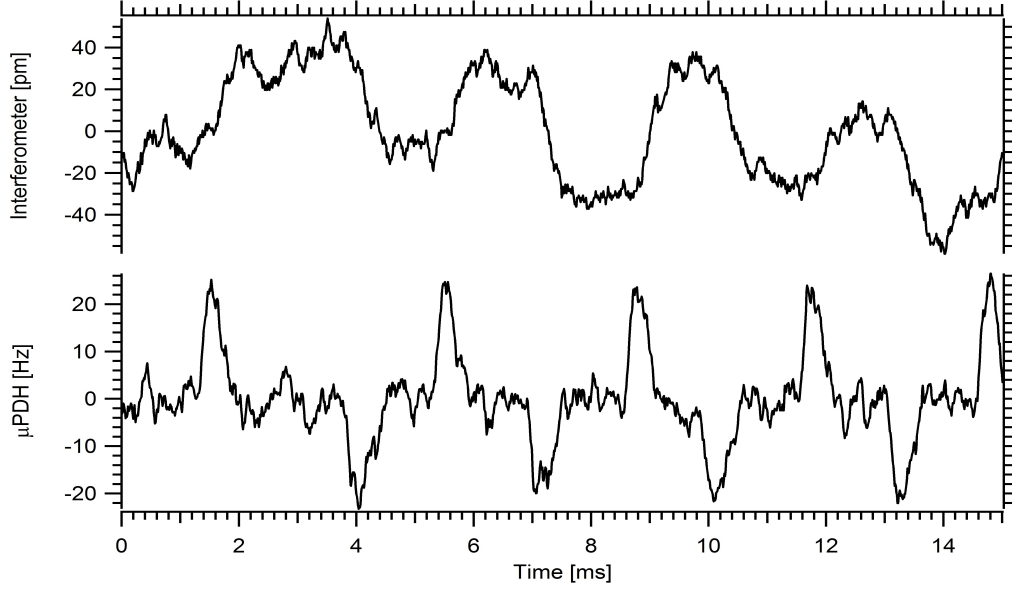
Figure 9.17: The closed-loop interferometer and µPDH signals as the source frequency is stepped back and forth by 25 Hz, resulting in 40 pm step in the cavity length. The signals were measured though 3 kHz low-pass filters.

9.3.3 Locking bandwidth

The frequency-distance encoder cannot change its length faster than the VCO frequency settles. After optimization of the PLL loop as described in appendix B, the settling time of the VCO is much faster then the expected mechanical bandwidth. Figure 9.18 depicts the open loop µPDH signal as the setpoint of the PLL is stepped by 5 kHz. The VCO frequency overshoots and oscillates a bit around its final value, similar to the behavior of an underdamped oscillator. The time constant of the exponential envelope of the oscillations is 8.8 µs. This is a bandwidth of 18 kHz.

To measure the speed of the mechanical response, the frequency is stepped in closed loop, and the rise-time of interferometer signal is measured. The response of the bare-metal flexture is limited to ~ 2 kHz because of its high Q mechanical resonant modes. Even at that bandwidth, ringing of ~ 1 nm amplitude is visible that decays over several milliseconds. To dampen the ringing, a layer of spray-on

111
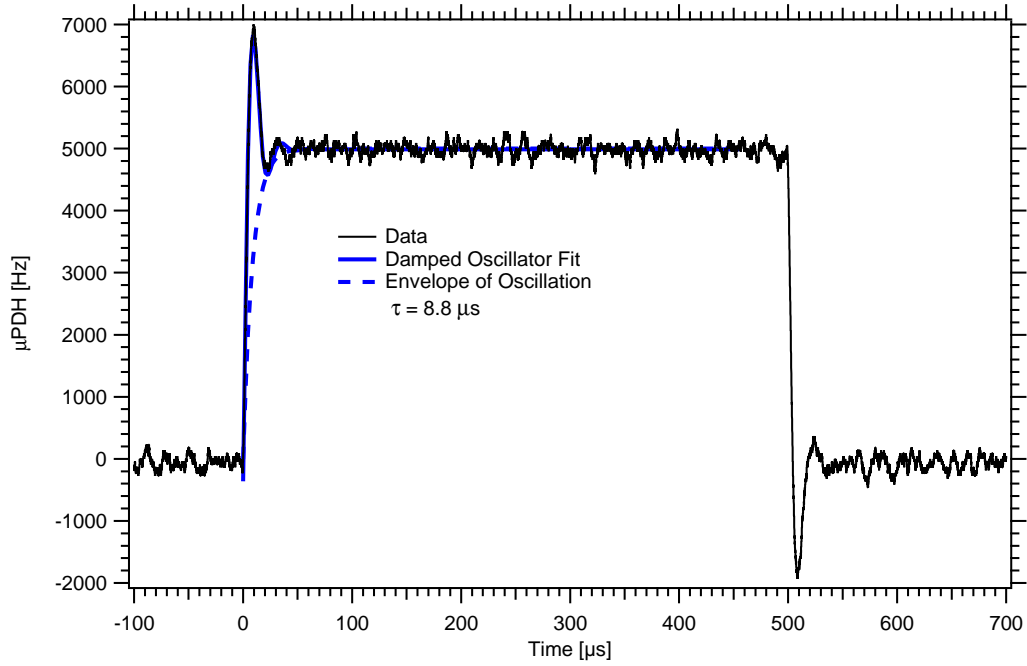
Figure 9.18: The response of the VCO frequency to a 5 kHz step in the setpoint. The response is similar to an underdamped oscillator with a natural frequency of 43 kHz and damping constant of 0.42. The time constant of the exponential envelope is $\tau = 8.8$ µs.

Figure 9.19: Rubber is added around the cavity flexture to increase damping and minimize ringing.

rubberized sealant is added around the flexture as shown in figure 9.19. With the damping, the mechanical step response rise time was 45 µs (3.5 kHz bandwidth), with significantly suppressed ringing. The bandwidth, however, was not limited by the mechanics, but by delay in the feedback controller.

Replacing the digital feedback controller with an SRS SIM960 100 kHz analog PID controller improves the response tremendously, shown in figure 9.20. The time
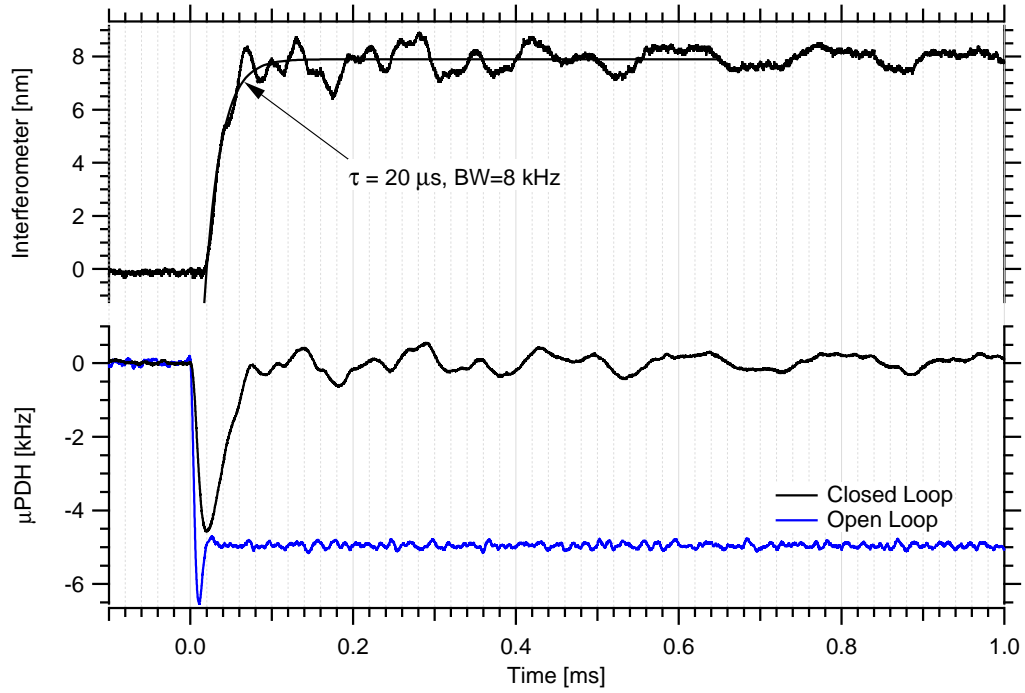
Figure 9.20: An analog PID controller improves the time constant of the response to 20±0.5 µs, or a bandwidth of 8 ± 0.2 kHz. Ringing of about 2 nm amplitude is visible, indicating additional damping is necessary. The interferometer signal is viewed through a 100 kHz filter while the µPDH signal is viewed through a 30 kHz filter.

constant is reduced to 20.0 ± 0.5 µs, corresponding to an 8.0 ± 0.2 kHz bandwidth.

Ringing is visible in the response, indicating additional damping may be helpful.

CHAPTER 10

Drift Budget

Here is a tally of some of the effects that may cause a drift in the frequency-distance encoder, despite it being locked to a stable frequency. The drift of the interferometer measurement is also considered. A summary of the various sources of drift is given in table 10.1. Experimental measurement of the three largest effects is subsequently presented.

---

[1]Evaluated for lab conditions of $p = 100.5$ kPa, $T = 25$ °C, and a relative humidity[2] of 40%.

Drift Source	Frequency		Distance	
Temp. Coefficient of Resistivity	-1.7	kHz K^{-1}	-2.8	nm K^{-1}
Index at 10 GHz due to Temp.[1]	12	kHz K^{-1}	19	nm K^{-1}
Index at 10 GHz due to tot. Pres.[1]	-31	Hz Pa^{-1}	-0.047	nm Pa^{-1}
Index at 10 GHz due to Humidity[1]	-412	Hz Pa$^{-1}_{\text{H}_2\text{O}}$	-0.63	nm Pa$^{-1}_{\text{H}_2\text{O}}$
Expansion of Quartz Rod	-		-8.4	nm K^{-1}
GRIN Lens Offset (Al Expansion)	-		46	nm K^{-1}
Index at 1310 nm due to Temp.[1]	-		-2.6	nm K^{-1}
Index at 1310 nm due to tot. Pres.[1]	-		.0078	nm Pa^{-1}
Index at 1310 nm due to Humidity[1]	-		-.001	nm Pa$^{-1}_{\text{H}_2\text{O}}$
Laser Wavelength Stability	-		-23	nm K^{-1}

Table 10.1: Various sources of drift and their expected magnitude. The frequency change is for an open loop cavity. The distance change is for a closed loop cavity, where a positive change means the cavity gets longer.

Metal	$\rho \cdot 10^{-8}$ (Ω m)	$\frac{1}{\rho}\frac{\partial\rho}{\partial T}$ (K^{-1})	δ (μm)	$\frac{\partial\delta}{\partial T}$ (nm K^{-1})	Drift (nm K^{-1})
Silver	1.587	.00375	0.640	1.20	2.40
Copper	1.678	.00401	0.658	1.32	2.64
Gold	2.214	.00367	0.756	1.39	2.78

Table 10.2: Drift due to the skin depth temperature dependence for various conductors. The values for ρ are taken from [43] at $T = 293$ K. The values for $\frac{\partial\rho}{\partial T}$ are calculated by linearly extrapolating between the values of ρ at $T = 293$ K and $T = 298$ K in [43] and taking the slope.

10.1 Microwave Cavity Drift

10.1.1 Skin depth temperature dependence

The skin depth in metals is given by

$$\delta = \sqrt{\frac{2\rho}{\omega\mu}}, \tag{10.1}$$

where ω, ρ, and μ are the microwave angular frequency, the metal resistivity, and metal permeability respectively. At room temperature, the resistivity typically varies by a few parts per thousand per degree, which leads to an effective change in cavity length of twice the change in skin depth (there are two walls). Estimates for the size of this type of drift are given in table 10.2.

10.1.2 Dielectric constant at 10 GHz

Changes of the index of refraction, n, of the gas within the microwave cavity affect its resonance frequency though,

$$\frac{df}{f} = -\frac{dn}{n}. \tag{10.2}$$

Measurements of the index of refraction of dry air and water vapor are given in [44], [45], and [46]. A convenient formula in [44] gives the refractive index up to

30 GHz as,

$$(n-1) \times 10^6 = \frac{0.776}{T} \left(p + 48.1 \, \frac{e}{T}\right), \tag{10.3}$$

where T is the temperature in Kelvin, p is the total pressure (dry air plus water vapor) in Pascal, and e is the partial pressure of water vapor in Pascal.

From here, several quantities of interest can be calculated, namely the change of index with respect to total pressure with the temperature and ratio of dry to moist air held fixed, $\frac{\partial n}{\partial p}\big|_{T,e/p}$; the change of index with respect to temperature with the pressures held fixed, $\frac{\partial n}{\partial T}\big|_{p,e}$; and the change of index with respect to the partial pressure of water vapor with the temperature and total pressure held fixed, $\frac{\partial n}{\partial e}\big|_{p,T}$,

$$\frac{\partial n}{\partial p}\bigg|_{T,e/p} = 0.776 \times 10^{-6} \, \frac{1}{T} \left(1 + 4.81 \times 10^3 \, \frac{e}{pT}\right), \tag{10.4}$$

$$\frac{\partial n}{\partial T}\bigg|_{p,e} = -0.776 \times 10^{-6} \, \frac{p}{T^2} \left(1 + 9.62 \times 10^3 \, \frac{e}{pT}\right), \tag{10.5}$$

$$\frac{\partial n}{\partial e}\bigg|_{p,T} = 3.73 \times 10^{-3} \, \frac{1}{T^2}. \tag{10.6}$$

Evaluating these derivatives for typical lab conditions of $p = 100.5$ kPa, $T = 25\,°C$, and a relative humidity[2] of 40% and combining with equation 10.2 gives the fractional change in frequency as,

$$\frac{1}{f} \frac{\partial f}{\partial p}\bigg|_{T,e/p} = -3.14 \times 10^{-9} \ \text{Pa}^{-1}, \tag{10.7}$$

$$\frac{1}{f} \frac{\partial f}{\partial T}\bigg|_{p,e} = 1.24 \times 10^{-6} \ \text{K}^{-1}, \tag{10.8}$$

$$\frac{1}{f} \frac{\partial f}{\partial e}\bigg|_{p,T} = -4.2 \times 10^{-8} \ \text{Pa}^{-1} \, _{\text{H}_2\text{O}}. \tag{10.9}$$

The magnitude of the expected frequency or length change is given in table 10.1.

---

[2]The vapor pressure of water is given in [47] as 3.2 kPa at 25 °C.

10.2 Interferometer Drift

10.2.1 Quartz rod

The fused quartz rod is fixed onto a surface coplanar with one wall of the microwave cavity. Aluminum was evaporated on the opposite end of the quartz rod, forming one mirror of the interferometer. The rod is 14 mm long and has a thermal expansion of 0.6×10^{-6} K^{-1}, corresponding to a drift of 8.4 nm K^{-1}.

10.2.2 Offset between the GRIN lens and microwave cavity wall

The surface of the GRIN lens is not coplanar with the microwave cavity wall. It is estimated that the surfaces are offset by ~ 2 mm because the quartz rod was fabricated about 2 mm longer than needed. Consequently the GRIN lens had to be moved back in order to keep the quartz mirror near the focal plane of the GRIN lens. The piece holding the GRIN lens is made of aluminum, with a thermal expansion of 23×10^{-6} K^{-1}, corresponding to a drift of 46 nm K^{-1}. This may be partially compensated by the lens expanding in the opposite direction.

10.2.3 Dielectric constant at infrared wavelengths

NIST provides a refractive index of air calculator as part of the Engineering Metrology Toolbox available on their website[3]. Based on this data, the sensitivity of the index of refraction at 1310 nm to changes in temperature, humidity,

---

[3]http://emtoolbox.nist.gov/Wavelength/Ciddor.asp

and pressure is estimated as,

$$\frac{\partial n}{\partial T}\bigg|_{p,e} = -8.79 \times 10^{-7} \ \ \text{K}^{-1}, \tag{10.10}$$

$$\frac{\partial n}{\partial p}\bigg|_{T,e/p} = 2.61 \times 10^{-9} \ \ \text{Pa}^{-1}, \tag{10.11}$$

$$\frac{\partial n}{\partial e}\bigg|_{p,T} = -0.36 \times 10^{-9} \ \ \text{Pa}^{-1} \ _{\text{H}_2\text{O}}. \tag{10.12}$$

10.2.4 Laser wavelength stability

With a 3 mm long interferometer, the laser wavelength needs to be stable to 3×10^{-7} in order for the measurement to be stable at the 1 nm level. This corresponds to a wavelength stability of 0.4 pm (68 MHz) for a 1310 nm laser. The current QPhotonics fiber pigtailed laser diode, model number QFBGLD-1300-5, operates at a wavelength of 1310 nm and is specified to have a wavelength temperature coefficient of 0.01 nm K^{-1}. It requires temperature stabilizing to 0.05 K for measurements to be stable at the 1 nm level.

10.3 Measured Sensitivity to the Environment

The microwave cavity and interferometer are sitting in atmosphere and are exposed to the changing lab environment. It is observed that the largest factors influencing the frequency-distance encoder are humidity, temperature, and air pressure. In an eventual application, a temperature-stabilized vacuum environment is almost certainly necessary.

To measure the effects of the environment on the system, the cavity is locked to a frequency and the interferometer signal, piezo voltage, temperature, air pressure, and humidity are recorded over several hours. Figure 10.1 displays the various curves over a 15 hour time-span. It appears that the piezo voltage follows the temperature most closely, and that the interferometer signal follows the partial
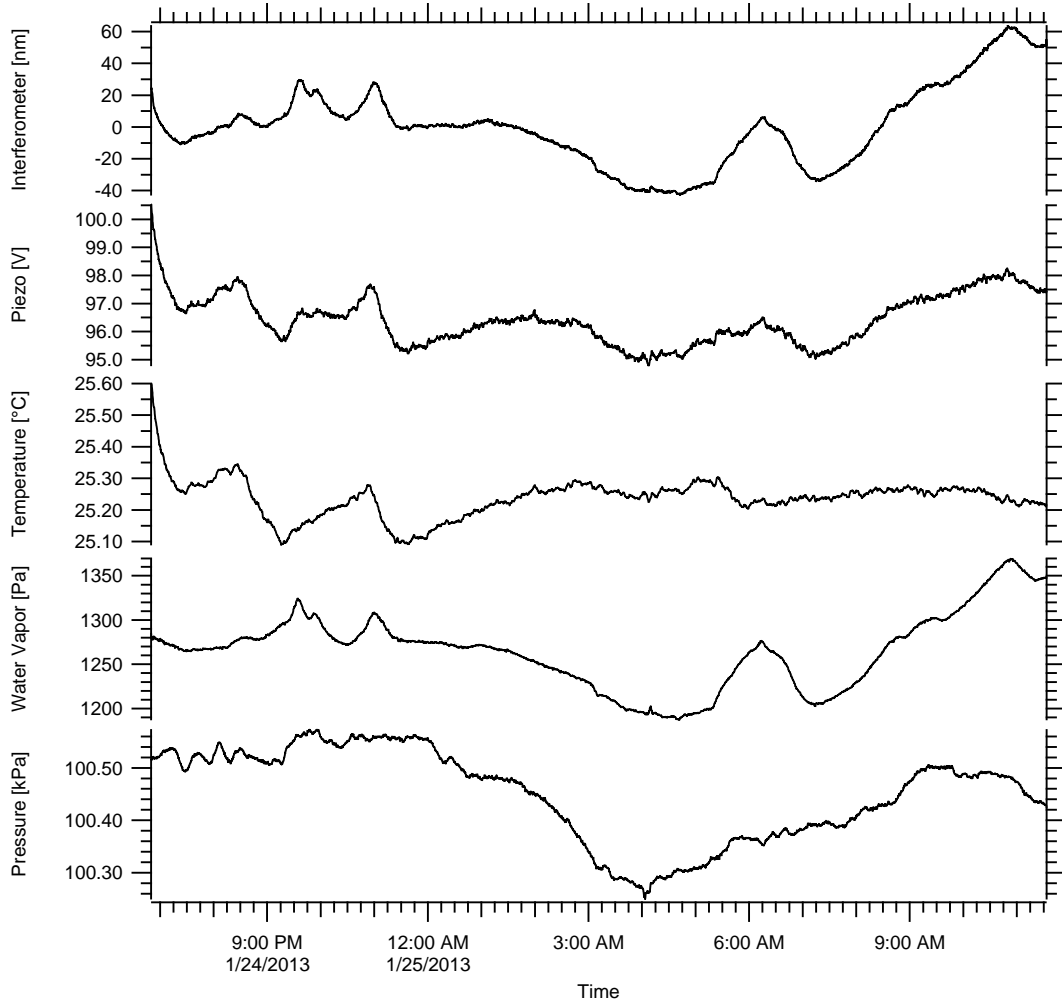
Figure 10.1: The cavity drift over almost 17 hours along with measurements of the piezo voltage, temperature, air pressure, and humidity. The interferometer signal remained on the linear portion of one fringe during the entire time span.

pressure of water vapor the closest.

A quantitative analysis is done by finding the linear combination of changes in the temperature, water vapor, and pressure curves that best fits changes in the interferometer and piezo voltage traces. The procedure is described in detail in appendix C. Figure 10.2a shows the best-fit linear combination to the piezo voltage, which seems to be a good fit. Figure 10.2b shows the relative contribution of the various sources, with temperature and humidity contributing roughly the same amount, and pressure contributing the least. The coefficients are given in

119
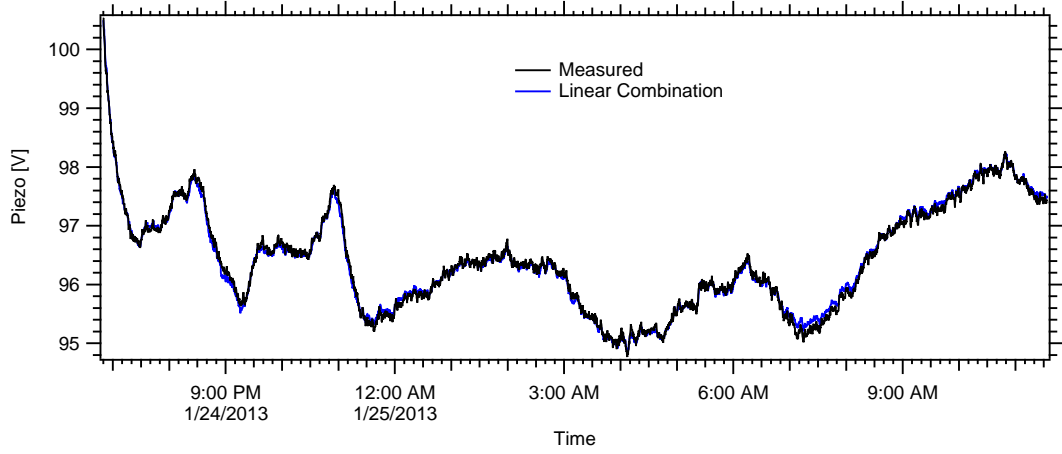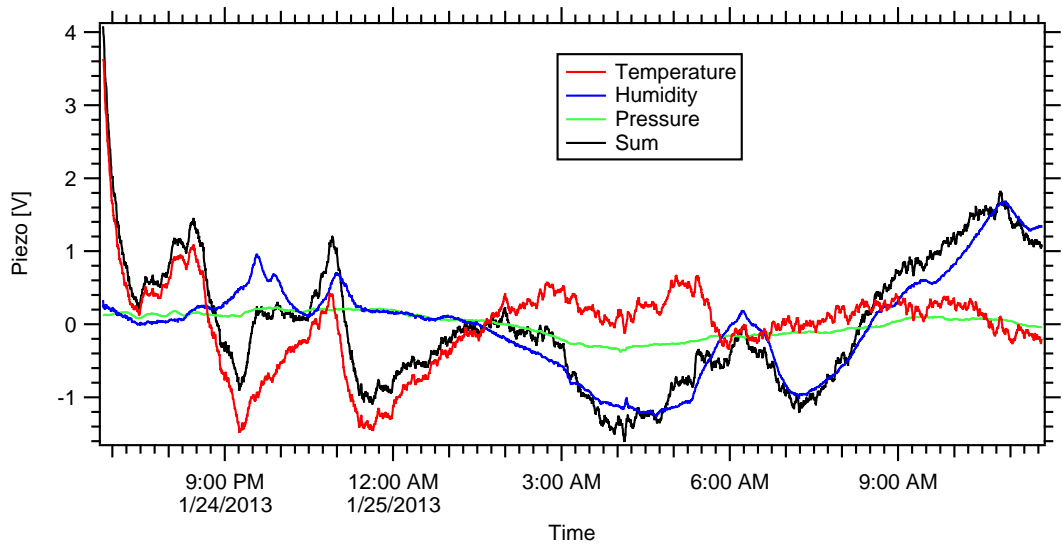
(a) Best Fit Curve

(b) Contributions of Drift Sources

Figure 10.2: The piezo voltage can be written as a linear combination of the temperature, water vapor pressure, and total pressure with coefficients given in table 10.3. Panel a) shows the best fit curve compared to the measured curve. The contribution of the various effects is shown in panel b).

table 10.3.

Figure 10.3a shows the best fit linear combination to the interferometer curve. The fit matches reasonably well, but there are noticeable discrepancies that indicate there is an effect that was not included. The wavelength variation of the laser is suspected to account for the discrepancies, but without a means of measuring the wavelength to the necessary precision, this remains to be verified. Figure
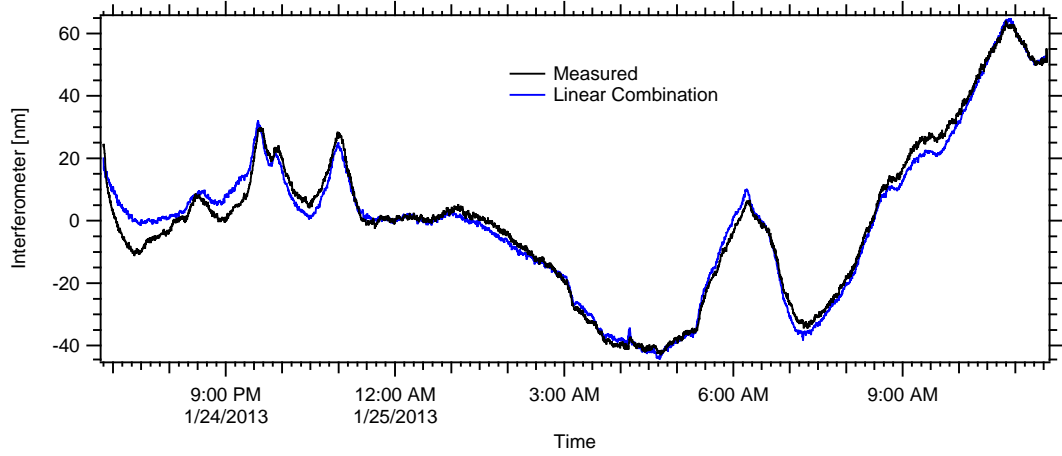
(a) Best Fit Curve
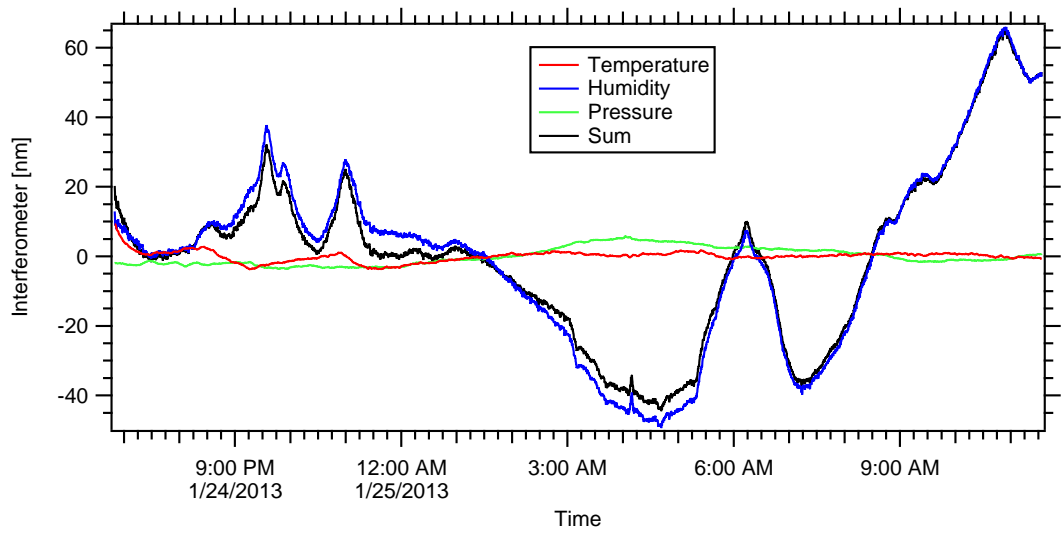
(b) Contributions of Drift Sources

Figure 10.3: The interferometer curve can be approximately written as a linear combination of the temperature, water vapor pressure, and total pressure curves with coefficients given in table 10.3. Panel a) shows the best-fit curve compared to the measured curve, while panel b) shows the contribution of the various effects.

10.3b shows the relative contribution of the various sources, with humidity being the largest contributor by far, and temperature and total pressure contributing roughly the same amount. The coefficients are also given in table 10.3.

The interferometer curve temperature coefficient of 25.4 nm K^{-1} is in agreement with what is expected from adding up the temperature-related drift sources

Variable	Piezo Voltage		Interferometer Motion	
Temperature	9.99	V K^{-1}	25.4	nm K^{-1}
Humidity	0.0161	V Pa$^{-1}_{H_2O}$	0.631	nm Pa$^{-1}_{H_2O}$
Total Pressure	0.00188	V Pa^{-1}	-0.029	nm Pa^{-1}

Table 10.3: The coefficients of the temperature, humidity, and total pressure curves that give the best fit linear combination to the piezo voltage and interferometer curves.
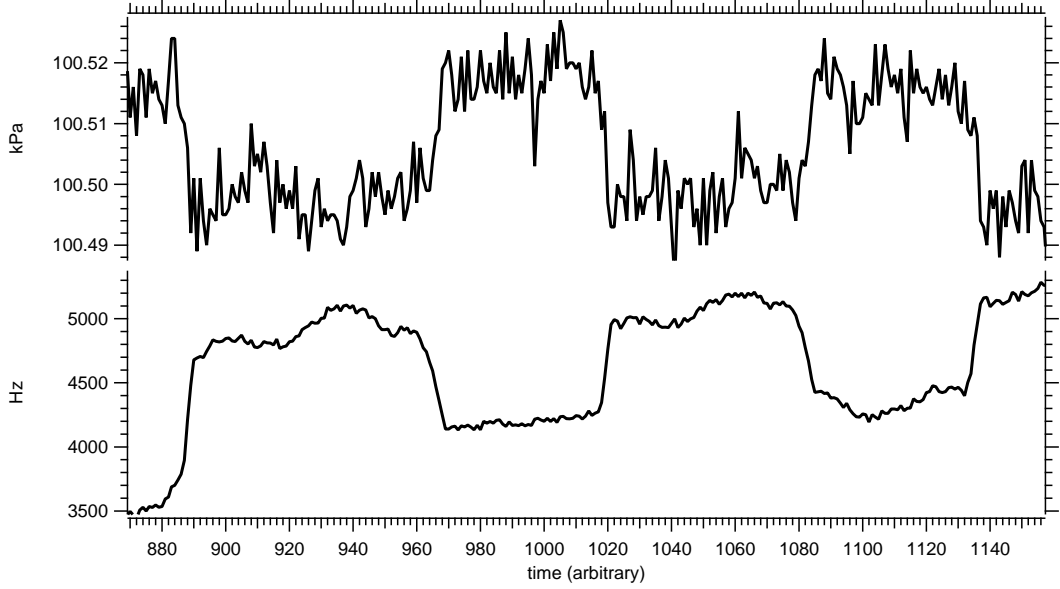
Figure 10.4: The µPDH signal as the pressure in the room is stepped by about 20 Pa. The frequency shift of the cavity is determined to be 36 ± 5 Hz Pa^{-1}.

in table 10.1 (28 nm K^{-1})[4]. The humidity coefficient of 0.631 nm Pa$^{-1}_{H_2O}$ matches the calculated value in table 10.1 to better than a percent. The total pressure coefficient is off by 25% (-0.029 compared to -0.039 nm Pa^{-1}), likely due to inaccurate estimates of the composition of air.

The total pressure coefficient was also measured by an independent method. Opening and closing the door to the lab changes the pressure by about 20 Pa. The pressure and µPDH signals are recorded in figure 10.4 as the door to the lab is repeatedly opened and closed. Dividing the frequency step by the pressure step,

---

[4]The uncertainty in some of the values in table 10.1 is larger than the discrepancy. The biggest contributors being the uncertainty in the position of the GRIN lens and the laser wavelength stability coefficient.

the frequency shift is determined to be 36 ± 5 Hz Pa^{-1}. This is within uncertainties of the value expected from the known pressure and temperature dependence of gas permittivity discussed in section 10.1.2.

CHAPTER 11

The Frequency-Distance Relationship

The frequency-distance relationship for an ideal, unperturbed cavity is,

$$f = \frac{c}{2L}. \tag{11.1}$$

To use the frequency-distance encoder as a metrological tool that does not require calibration, it is necessary that this relationship is obeyed as closely as possible. In this section, the frequency-distance conversion is investigated and compared to the optimal.

The interferometer is used to independently measure the length change of the
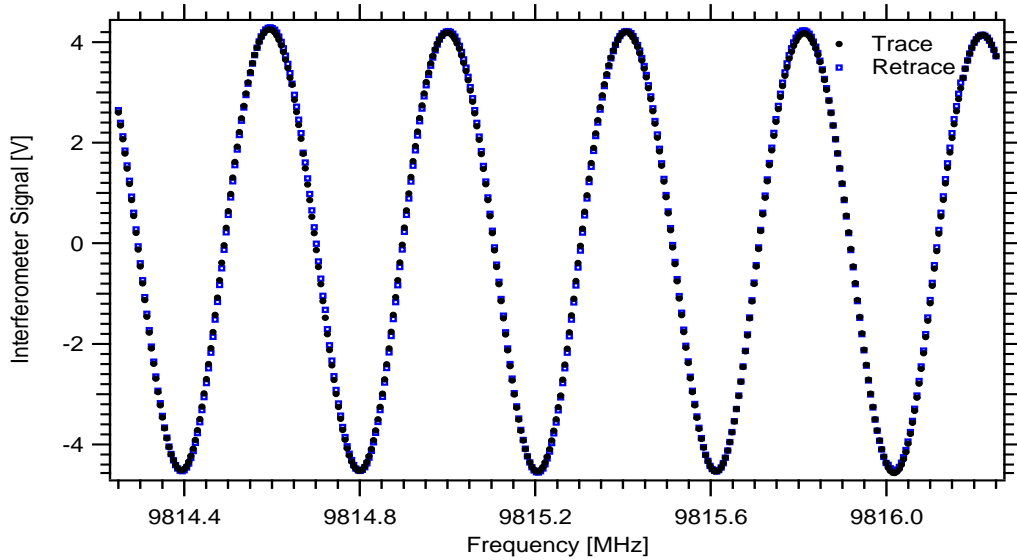
Figure 11.1: The interference measured as the cavity frequency is scanned. The frequencies at which the peaks occur are extracted to make the frequency-distance plot of figure 11.2.
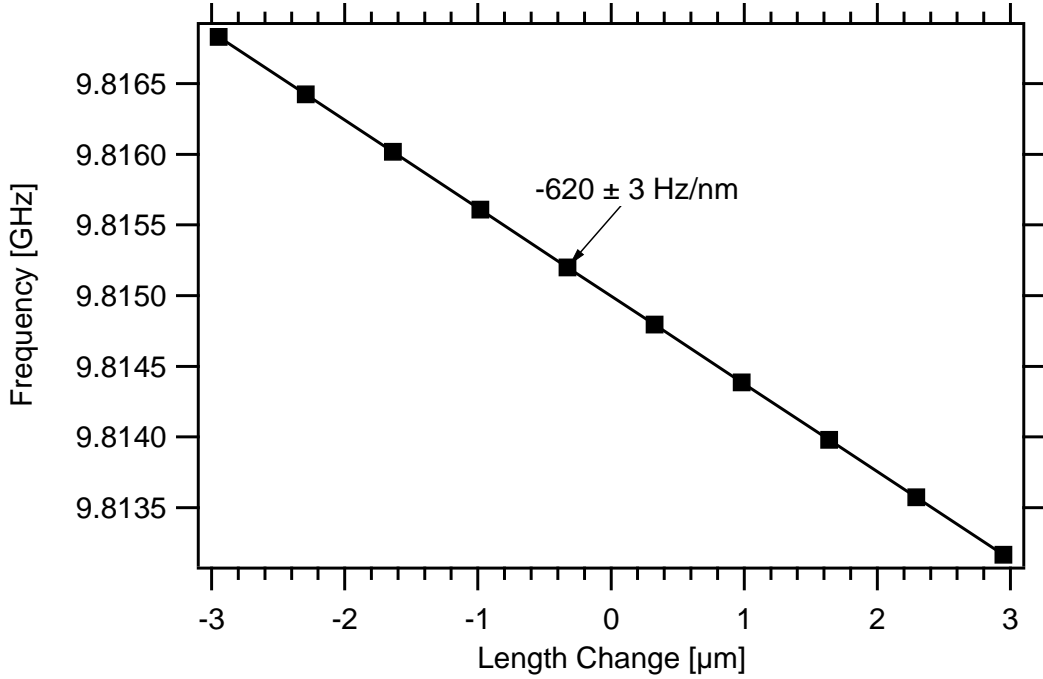
Figure 11.2: The frequencies at which maxima occur in the interferometer pattern are plotted against the length change of the cavity. For an ideal cavity, the frequency-distance curve would have a linear coefficient of -642.7 Hz nm^{-1} at 9.815 GHz.

cavity as its frequency is swept. Many interferometer curves are recorded over the range of motion of the cavity, as in figure 11.1. The peaks of the interferometer pattern (spaced $\lambda/2 = 655$ nm apart) are selected, and the frequencies at which they occur are plotted against the distance traveled. Many (10-20) such curves are averaged to produce a figure like 11.2. Fitting a line to the data gives a slope of -620 ± 3 Hz nm^{-1}. At the center frequency of the figure, 9.815 GHz, df/dL would be -642.7 Hz nm^{-1} for an ideal cavity, giving a deviation of $3.5\% \pm 0.5\%$.

To estimate the uncertainty in the slope measurement, the sensitivity of the slope to measurement technique was investigated. Repeating the scan and running the algorithm to extract the peaks and fit a line results in a standard deviation of the measurements on the order of $\sim 10^{-4}$. Changing the speed of the scan does not affect the result. In addition, rebooting the bridge and rerunning the phase-frequency plot algorithm to choose the optimal LO phase does not make
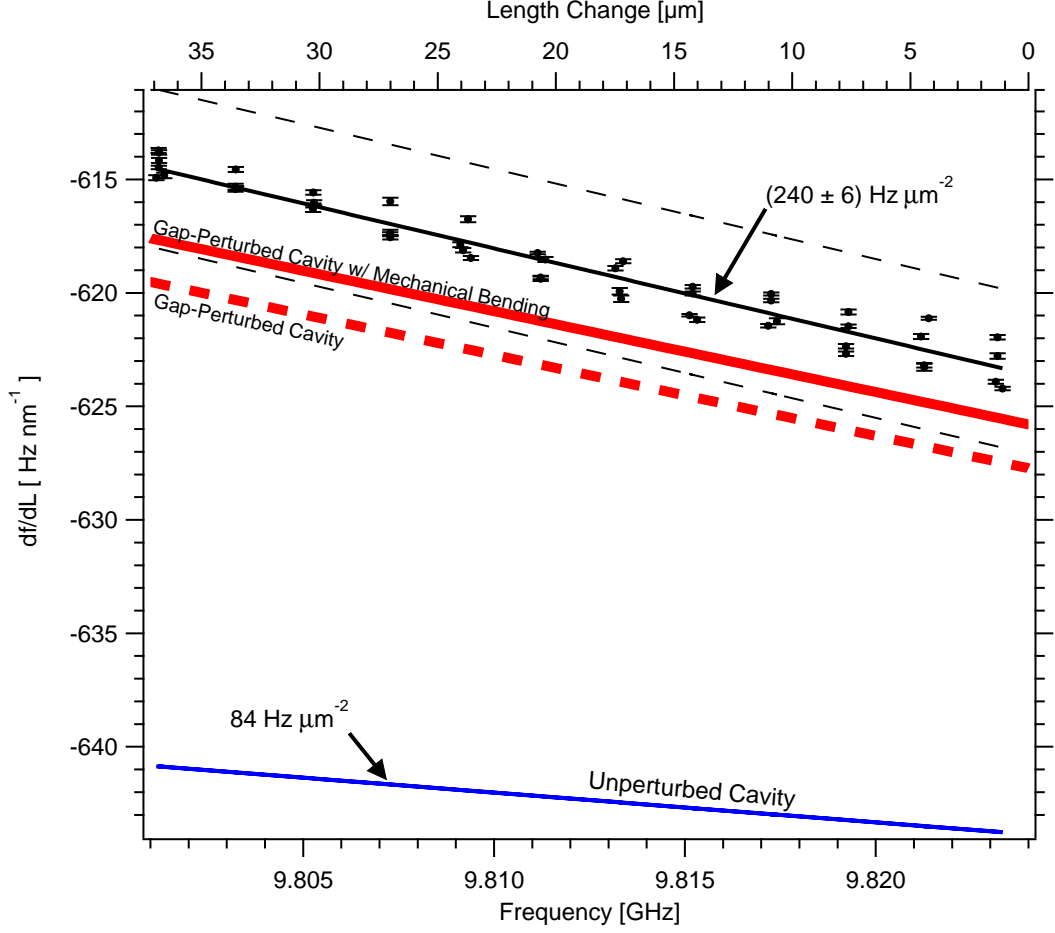
Figure 11.3: The slope of the frequency-distance relationship for an unperturbed cavity (blue line) compared to the measured slope (black points) as a function of frequency for the Q=1280 cavity. The dashed black lines indicate the range of mechanical repeatability of the measurements. Fitting a line to the data points gives the second derivative of the frequency-distance relationship as $\frac{d^2 f}{dL^2} = 240 \pm 6$ Hz µm^{-2}. The dashed red line indicates the expected slope as a function of frequency for a gap-perturbed cavity, using the theory developed in section 11.1. The solid red line indicates the expected slope when the gap and mechanical bending, section 11.2, are taken into account.

a difference. Removing and re-mounting the fiber GRIN lens changes the measurement on the order of 5×10^{-4}. Backlash in the coarse length-adjust knob makes changes on the 2×10^{-3} level. The largest effect comes from disassembling and rebuilding the entire cavity mechanics, which makes changes on the order of 5×10^{-3}.

Measuring df/dL over the range of the VCO frequency results in figure 11.3. Each point in the figure is measured by taking 2 MHz sweeps centered around the frequency plotted. The coarse length adjust knob was used to vary the length of the cavity over the range needed, as the piezo was limited to only a 4 MHz range. The scatter visible in the points is due to small mechanical changes from moving the knob, and is much larger than the statistical error of each point, indicated by the error bars. A curve corresponding to an ideal cavity is also shown for comparison. Explanations for the deviation are explored in the sections that follow.

11.1 Effect of the Gap

To make the cavity length adjustable, it was cut down the center. The location of the cut was chosen based on the fact that there is no current at that location, and therefore slicing the cavity in half would not perturb the mode as long as the gap is infinitesimally thin. In this section the effect of a finite gap width on the cavity resonant frequency is investigated. First, the results of numerical simulations are presented, followed by an analytical estimate of the perturbation.

11.1.1 Numerical simulations

The discontinuity of the cavity inner and outer conductors is simulated in Ansoft HFSS in order to quantify its effect on the cavity resonance frequency. Two models are simulated, one where the center conductor ends at the gap, figure 11.4a, and one where the center conductor bridges the gap, but is step-discontinuous, figure 11.4b. The outer conductor has a step discontinuity at the gap in both models, going from 11 mm diameter to 13.5 mm diameter. The cavity is cut out of copper, and modeled as two distinct halves separated by a thin vacuum layer, as would be the case in any actual mechanical design. In order to close the structure,

(a) Discontinuous Center Conductor   (b) Step-Discontinuous Center Conductor
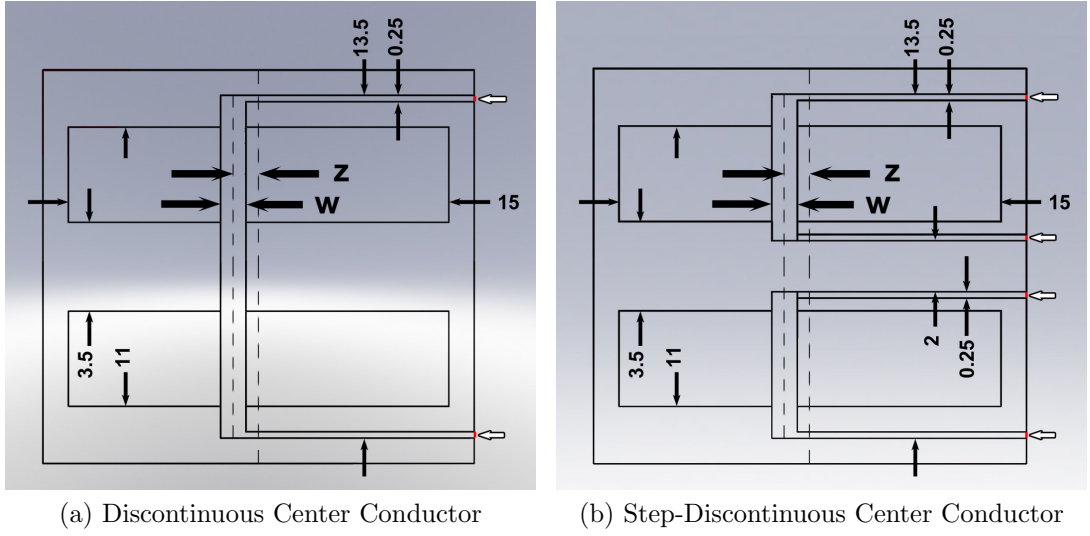
Figure 11.4: The cavity models simulated in HFSS to quantify the effect of the central gap on the resonance frequency. Panel a) shows the model without a center conductor bridging the gap while panel b) shows the model with a center conductor bridging the gap. The red edges indicated by the outlined, white arrows indicate the location of perfectly matched layers (PML) which prevent any reflection of an incident wave. The gap width w, and its position relative to the center of the cavity z, are adjustable parameters. The units of the displayed dimensions are mm.

perfectly matched layers (PML) are added to terminate the space between the halves, indicated by the red lines, and white arrows in the figures. The gap, of width w, is not assumed to be centered on the cavity. Its location is given by the gap-center to cavity-center distance z.

The cavity eigenfrequencies are calculated for $w = 0.5, 1,$ and 1.5 mm and z between ± 1.5 mm in 0.25 mm steps. In actuality, $w \approx 100$ µm and $z \lesssim 100$ µm, but due to the numerical accuracy of the simulation it is not possible to get meaningful results with such small perturbations and reasonable calculation time. The approach taken is to find the resonant frequency's functional dependence on w and z, and extrapolate to small values.

The results are shown by the markers in figure 11.5, with the unperturbed resonant frequency indicated by the red line. For a fixed w, the frequency varies
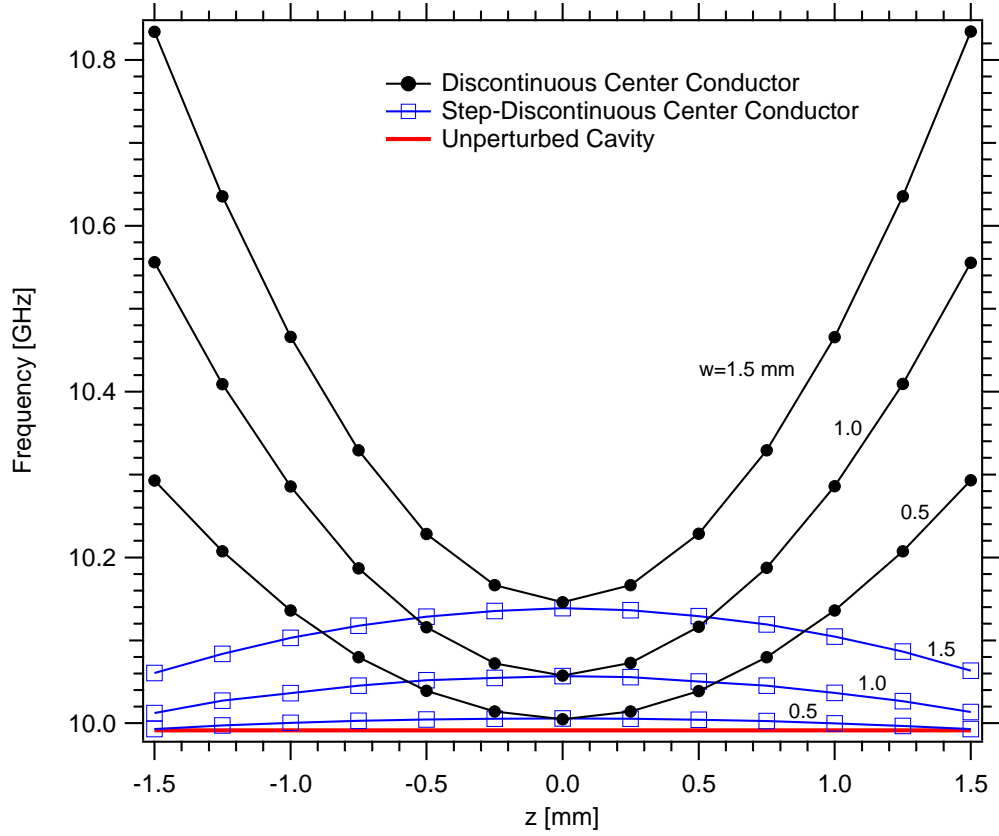
Figure 11.5: The cavity eigenfrequencies calculated by HFSS for a range of z and w are shown by the markers above. The cavity resonance frequency can be expressed in the form of equation (11.2). The parameters A and B are extracted by curve fitting and are given in table 11.1.

parabolically with z, with positive curvature when the center conductor is discontinuous, and negative curvature with the step-discontinuous center conductor. The perturbation to the resonant frequency is overall smaller when the center conductor bridges the gap, but for small z, there is no difference. The $z = 0$ points are plotted as a function of w in figure 11.6. The dependence of the resonant frequency on w and z can be expressed as,

$$f = f_0 + Aw^2 + Bwz^2, \tag{11.2}$$

where $f_0 = c/2L$ is the resonant frequency of the gap-less cavity, and values for the coefficients A and B are found by curve fitting and given in table 11.1. The
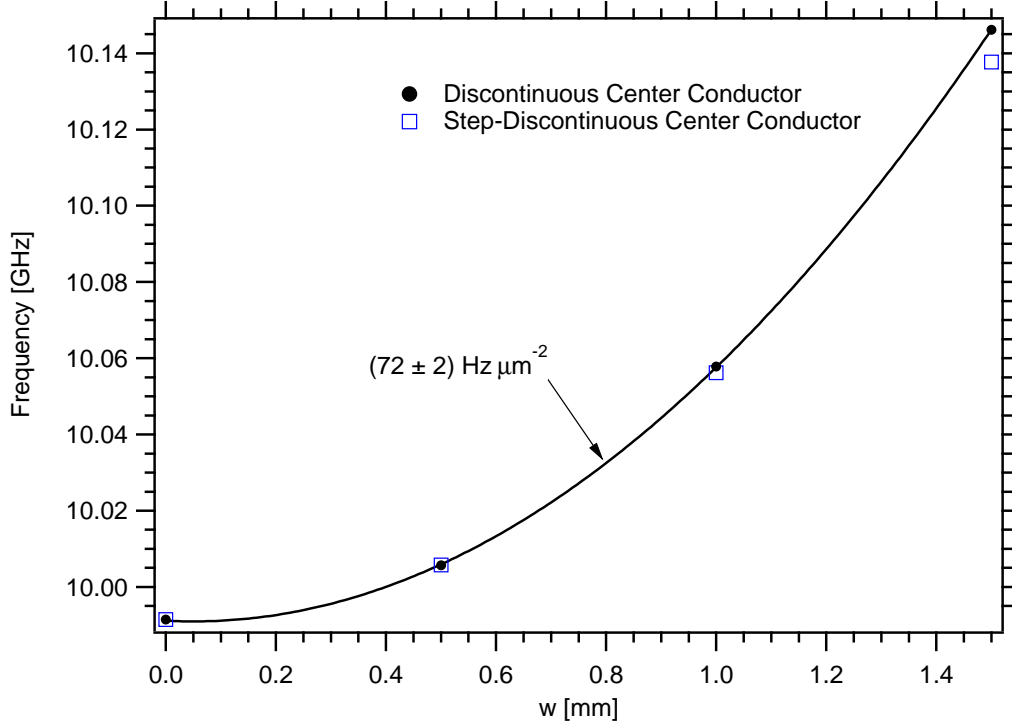
Figure 11.6: The cavity eigenfrequencies calculated by HFSS for $z = 0$ as a function of w. The correction to the cavity frequency is quadratic in w, with coefficient 72 ± 2 Hz µm^{-2}.

Model	A [Hz µm^{-2}]	B [MHz mm^{-3}]
Discontinuous Center Conductor	72 ± 2	215 ± 10
Step-Discontinuous Center Conductor	65 ± 2	-25 ± 3

Table 11.1: The values of the coefficients of equation (11.2), describing the perturbed resonant frequency calculated by HFSS and shown in figure 11.5.

uncertainty in the values of A and B is due to their sensitivity to the details of how the fitting is done. For example, forcing the fits to agree with known limiting cases changes the result.

To get to the correction to the slope of the frequency-distance curve, note that $dw/dL = 1$ and $dz/dL = 0$. The corrected slope is,

$$\frac{df}{dL} = \frac{df_0}{dL} + 2Aw + Bz^2. \tag{11.3}$$

and second derivative,

$$\frac{d^2 f}{dL^2} = \frac{d^2 f_0}{dL^2} + 2A. \tag{11.4}$$

Interestingly, the correction to the second derivative does not depend on the gap width, nor on how accurately it is centered. The measurement of the second derivative in figure 11.3, gives 240 ± 6 Hz µm^{-2}. Subtracting the unperturbed second derivative, 84 Hz µm^{-2}, and dividing by 2 gives an experimental value of A as 78 ± 3, compared to the numerical simulation value of 72 ± 2 Hz µm^{-2}.

The correction to the first derivative depends on both the gap width and the centering accuracy. The gap of the high Q cavity was measured by using the cavity piezo as a sensor and turning the coarse length adjust knob until the two halves touched. A voltage jump across the piezo was observed when the cavity resonant frequency was 9893 ± 1 MHz. This implies that at a resonant frequency of 9815 MHz, the gap width is 121 ± 2 µm. By design, it was intended to be about 150 µm. The second term in equation (11.3) is then $2Aw = 17.4 \pm 0.3$ Hz nm^{-1} at 9815 MHz.

To estimate z, the depth of the cavity on the moving and fixed halves was measured with a caliper. The depth in the moving half is 7.65 mm, and the depth on the fixed side is 7.51 mm. However, the coupling hole makes the fixed half-cavity effectively deeper. This effect can be estimated from the shift of the resonant frequency in going from the uncoupled cavity to the critically coupled cavity[1], -9.4 ± 0.2 MHz. Using the frequency-distance conversion of 643 Hz nm^{-1}, this corresponds to moving the cavity wall back by 14.6 ± 0.3 µm. In effect, the fixed half-cavity is 7.525 mm deep, and therefore z is estimated to be 60 ± 10 µm. The magnitude of the third term in equation (11.3) is then 0.75 ± 0.25 Hz nm^{-1}.

The prediction of the gap-perturbed frequency-distance slope is shown by the

---

[1]The frequency shift of the cavity in going from uncoupled to undercoupled cannot be measured directly because of the permanently present coupling hole. The value is obtained by extrapolating the curve in figure 11.9 below to $\beta = 0$.

dashed red line in figure 11.3 as a function of frequency, and is less than 1% away from the measured values.

11.1.2 Cavity perturbation theory

An approximate formula for the resonant frequency of the gap-perturbed cavity is derived. Standard microwave cavity perturbation theory [48, 49] gives the change in frequency of a resonant cavity upon deformation of its boundary as,

$$\frac{\omega - \omega_0}{\omega_0} \simeq \frac{\int_{\Delta V} (\mu|H_0|^2 - \epsilon|E_0|^2) dv}{\int_V (\mu|H_0|^2 + \epsilon|E_0|^2) dv}, \tag{11.5a}$$

$$= \frac{\Delta W_m - \Delta W_e}{W_m + W_e}, \tag{11.5b}$$

where E_0 and H_0 are the unperturbed fields, V is the volume of the unperturbed resonator, and ΔV is the volume removed from the original mode. In the second line, ΔW_m and ΔW_e are the changes in the stored magnetic and electric energy respectively and $W_m + W_e$ is the total stored energy in the cavity. The unperturbed fields are given by equation (6.1).

For the problem at hand, the volume within the gap did not exist as part of the unperturbed mode, and therefore the unperturbed fields cannot be used to calculate ΔW_m and ΔW_e. In addition to the volume ΔV scaling with the gap size, the fields within ΔV must scale with it too. Fortunately, as the gap width $w \sim 100$ µm is much smaller than the wavelength, electrostatic theory can be used to estimate the electric field within the gap. As it is located at the electric field maximum, there is no magnetic field, so $\Delta W_m = 0$. ΔW_e is then calculated from the electrostatic fields and plugged into equation (11.5b) to give the perturbed frequency.

An approximate solution for the electric field within the gap is achieved by separation of variables in the cylindrical coordinates ρ and z (no ϕ dependance),

(a) Inner Gap, $\rho < a$ (b) Outer Gap, $\rho > b$
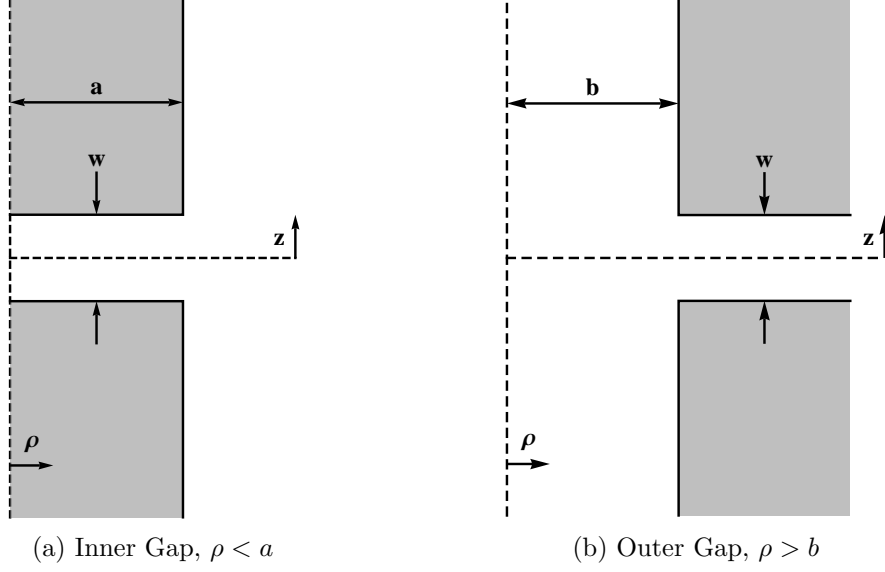
Figure 11.7: Geometry defining the dimensions for solving the electrostatic field within the gap regions. Dimensions are not to scale, in actuality $w \ll a \sim b$.

with the gap assumed to be centered in the cavity. Figure 11.7 shows the geometry for both $\rho < a$ and $\rho > b$. The z coordinate is measured from the center of the gap.

The boundary conditions imposed are that 1) the potential must be constant for $z = \pm w/2$, 2) the fields must decay for $\rho \gg b$ and $\rho \ll a$, and 3) the $\hat{\rho}$ component of the fields must match the unperturbed mode at $\rho = b$ and $\rho = a$. Consideration of conditions 1) and 2) gives a potential of the form,

$$\Phi(\rho, z) = \Phi_0 + \begin{cases} \sum_{n \geq 0} A_n \cos(k_n z) K_0(k_n \rho), & \text{for } \rho > b \\ \sum_{n \geq 0} B_n \cos(k_n z) I_0(k_n \rho), & \text{for } \rho < a \end{cases} \tag{11.6}$$

where Φ_0, A_n, and B_n are constants, $k_n = (2n+1)\pi/w$, and $K_n(x)$ and $I_n(x)$ are modified Bessel functions of the second and first kind respectively. The electric

fields are[2],

$$\overrightarrow{E}(\rho, z) = \begin{cases} \sum_n A_n k_n \left[\cos(k_n z) K_1(k_n \rho) \, \hat{\rho} + \sin(k_n z) K_0(k_n \rho) \, \hat{z}\right], & \rho > b \\ \sum_n B_n k_n \left[-\cos(k_n z) I_1(k_n \rho) \, \hat{\rho} + \sin(k_n z) I_0(k_n \rho) \, \hat{z}\right] & \rho < a. \end{cases}$$

$$(11.7)$$

Applying boundary condition 3) gives the constants,

$$A_m k_m = (-1)^m \frac{4V_0}{w k_m K_1(k_m b) b \ln \frac{b}{a}}, \tag{11.8a}$$

$$B_m k_m = -(-1)^m \frac{4V_0}{w k_m I_1(k_m a) a \ln \frac{b}{a}}. \tag{11.8b}$$

The integral over $|E|^2$ becomes,

$$\int_{\Delta V} |E|^2 dV = \frac{16 V_0^2}{\pi^2 \ln^2 \frac{b}{a}} \sum_m \frac{w^2}{k_m (2m+1)^3} \left[\frac{1}{b^2 \, K_1^2(k_m b)} \int_{k_m b}^{\infty} x \left[K_1^2(x) + K_0^2(x) \right] dx \right.$$

$$\left. + \frac{1}{a^2 \, I_1^2(k_m a)} \int_0^{k_m a} x \left[I_1^2(x) + I_0^2(x) \right] dx \right]. \tag{11.9}$$

The integrals over the modified Bessel functions can be performed by replacing them with their series expansion[3] for $x \gg 1$. This is justified as $k_m b \sim k_m a \gg 1$. The zero limit on the integral from 0 to $k_m a$ can be replaced with a finite value ϵ satisfying $k_m a \gg \epsilon \gg 1$ with negligible error. This reflects the physical reality that the electric field in the $\rho < a$ gap will decay over a scale comparable with w, and there is negligible field near $\rho = 0$. The integrals are solved analytically,

$$\int_{\Delta V} |E|^2 dV \simeq \frac{16 \, V_0^2 \, w^2}{\pi^2 \ln^2 \frac{b}{a}} \left(\frac{1}{b} + \frac{1}{a} \right) \sum_m \frac{1}{(2m+1)^3} \tag{11.10a}$$

$$\approx 1.7 \times \frac{V_0^2 \, w^2}{\ln^2 \frac{b}{a}} \left(\frac{1}{b} + \frac{1}{a} \right), \tag{11.10b}$$

where the second line follows by evaluating the sum, $\sum_m 1/(2m+1)^3 = 7\zeta(3)/8 \approx$

---

[2]The identities $\frac{dK_0}{dx} = -K_1$ and $\frac{dI_0}{dx} = I_1$ have been used.

[3]$K_n(x) \simeq \sqrt{\frac{\pi}{2x}} e^{-x}$ and $I_n(x) \simeq \frac{1}{\sqrt{2\pi x}} e^x$ for $x \gg 1$

1.052, and $\zeta(x)$ is the zeta function.

Finally, plugging into equation (11.5b), with the use of 6.2b gives the perturbed resonant frequency as,

$$f = f_0 + 1.7 \times \frac{cw^2}{4\pi L^2 \ln \frac{b}{a}} \left(\frac{1}{b} + \frac{1}{a} \right), \tag{11.11a}$$

$$= f_0 + \kappa w^2. \tag{11.11b}$$

The correction to the unperturbed frequency is parabolic in the gap width. Inserting values for a and b corresponding to the fabricated cavity, and $L = 15.28$ mm ($f_0 = 9.81$ GHz), the coefficient of w^2 is $\kappa = 114$ Hz µm^{-2}.

The coefficient of w^2 has now been obtained three different ways. The preceding derivation puts it at 114 Hz µm^{-2}, numerical simulations at 72 ± 2 Hz µm^{-2}, and experiment at 78 ± 3 Hz µm^{-2} if the deviation of the measured second derivative is all blamed on the gap. The 50% deviation of the perturbation theory result can be attributed to the theory not taking into account perturbations of the field within the cavity. Lowering the field within the cavity would lower the time-averaged energy in the electric field, pushing the value of κ down.

11.2 Stress-Induced Bending in Moving Half-Cavity

By design, the moving half-cavity is in constant tension to provide the preload for the piezo. Its shape was partially motivated by the desire to have the cavity section be stress-free to avoid deformation; however, stresses in the flexure part of it still cause some minor deflection. In this section the amount of bending in the moving half-cavity is quantified.

Solidworks Finite-Element modeling is used to "measure" the deformation of the cavity. Under a 100 N load, the flexure deforms as in figure 8.3. To look at small-scale differences in the motion of the cavity, "sensors" are set up in
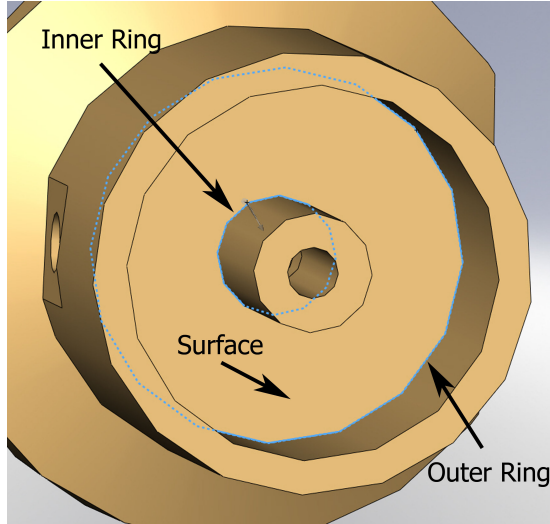
Figure 11.8: The cavity deforms slightly as it is strained to move. The inner ring moves 0.27% more and the outer ring moves 0.16% less than the average deflection. Under a load of 100 N, this corresponds to a difference of 5 nm between the inner and outer rings, compared to the average motion of 1230 nm.

Solidworks, which average the deflection over a path or area. Figure 11.8 shows how three sensors are defined. The average deflection along rings defined by the intersection of the flat cavity wall with the outer and inner cylindrical walls is 1228.3 nm and 1233.61 nm respectively. The average over the entire surface of the cavity end is 1230.28. Thus it is found that the the inside of the cavity moves 0.27% more, and the outside of the cavity moves 0.16% less, than the average motion. Considering that the fiber GRIN lens is held on the inner surface of the cavity, it can be expected to move about 0.3% further than the average deflection. This implies that $\frac{df}{dL_{avg}} = 1.003 \times \frac{df}{dL_{fiber}}$, pushing the measured value slightly closer to the expected value. The solid red line in figure 11.3 shows the gap-corrected, and stress-induced-bending-corrected frequency-distance slope. The result is within the range of mechanical reproducibility.
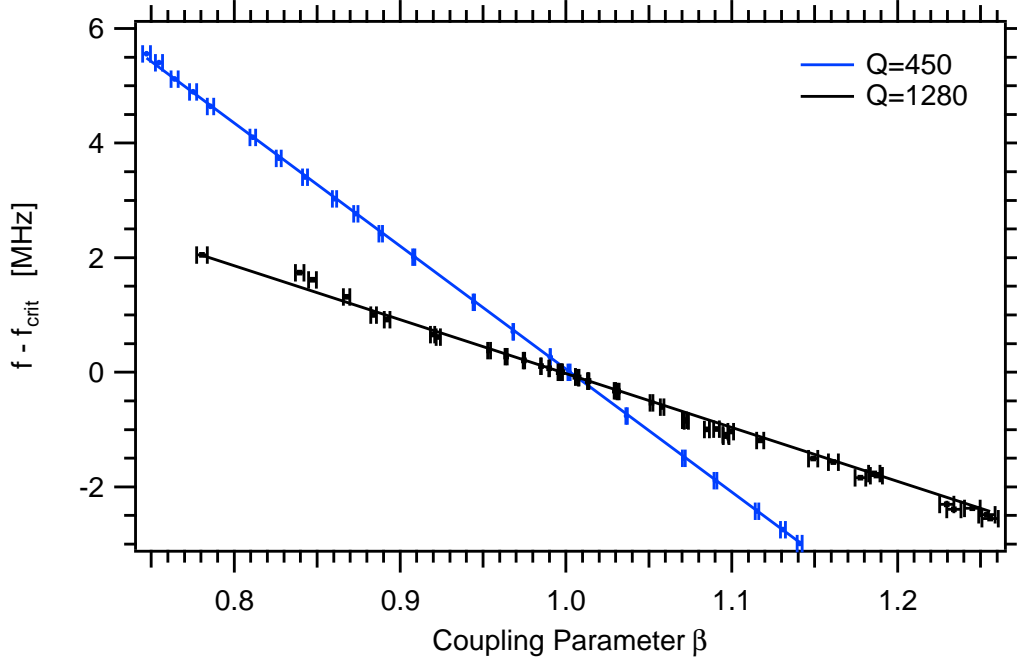
Figure 11.9: The cavity resonant frequency, f, as a function of β, when the cavity length is fixed and the coupling screw is adjusted. f_{crit} is the cavity frequency when it is critically coupled. The slope of the curves are $\left.\frac{\partial f}{\partial \beta}\right|_L = -21.5 \pm 0.5$ MHz and -9.4 ± 0.2 MHz for the cavities of Q=450 and Q=1280 respectively.

11.3   Effect of the Coupling Structure

Having taken into account the corrections to the frequency-distance relationship due to the gap within the cavity, and mechanical bending, there remains an $0.4\% \pm 0.6\%$ deviation. Coupling effects are herein investigated. In this section, the effects of the gap are ignored; the cavity is treated as gap-free.

If the cavity coupling structure adds a phase shift to the microwave reflecting in the cavity, the effective cavity length, and the frequency-distance relationship is modified. Let $S(x, \omega)$ be the scattering parameter experienced by the wave internal to the cavity when it reflects off the coupler. In this section, S is understood to be $-S_{22}$ of the scattering matrix defined in appendix A. The variable x represents the cavity geometry, specifically the position of the coupling screw.

The cavity resonance condition (see appendix A) is,

$$2\frac{\omega}{c}L + \arg S(x, \omega) = 2\pi, \tag{11.12}$$

which implies a relationship between the differentials of the variables,

$$\frac{2}{c}\omega dL + \left.\frac{\partial \arg S}{\partial x}\right|_\omega dx + \left(\frac{2}{c}L + \left.\frac{\partial \arg S}{\partial \omega}\right|_x\right) d\omega = 0. \tag{11.13}$$

The slope of frequency-distance relationship (at constant x) is,

$$\left.\frac{\partial \omega}{\partial L}\right|_x = -\frac{\omega}{L}\left[\frac{1}{1 + \frac{c}{2L}\left.\frac{\partial \arg S}{\partial \omega}\right|_x}\right]. \tag{11.14}$$

The unperturbed slope, $-\omega/L$, is changed by the presence of $\left.\frac{\partial \arg S}{\partial \omega}\right|_x$. If, however, the position of the coupling screw is adjusted as the frequency is scanned in such a way to keep the coupling constant[4], the slope of the frequency-distance relationship is modified as,

$$\left.\frac{\partial \omega}{\partial L}\right|_{|S|} = -\frac{\omega}{L}\left[\frac{1}{1 + \frac{c}{2L}\left.\frac{\partial \arg S}{\partial \omega}\right|_x\left(1 - \left.\frac{\partial \arg S}{\partial |S|}\right|_\omega\left.\frac{\partial |S|}{\partial \arg S}\right|_x\right)}\right]. \tag{11.15}$$

Interestingly, if the coupling structure is such that

$$\left.\frac{\partial \arg S}{\partial |S|}\right|_\omega\left.\frac{\partial |S|}{\partial \arg S}\right|_x = 1, \tag{11.16}$$

then measuring at constant coupling removes all the ill effects of the presence of $\left.\frac{\partial \arg S}{\partial \omega}\right|_x$. This condition holds for some simple coupling structures but is not true in general[5]. In this case, an estimate given below, puts $\left.\frac{\partial \arg S}{\partial |S|}\right|_\omega\left.\frac{\partial |S|}{\partial \arg S}\right|_x > 1.6$, or

---

[4]This can be expressed by the condition $d|S| = \left.\frac{\partial |S|}{\partial \omega}\right|_x d\omega + \left.\frac{\partial |S|}{\partial x}\right|_\omega dx = 0$.

[5]One example where it holds is a simple, aperture-coupled cavity. An equivalent circuit for a small aperture in a transverse waveguide wall is a shunt inductance. The scattering parameters for such a circuit depend only on the ratio of ω and ω_0, where ω_0 depends on the value of the
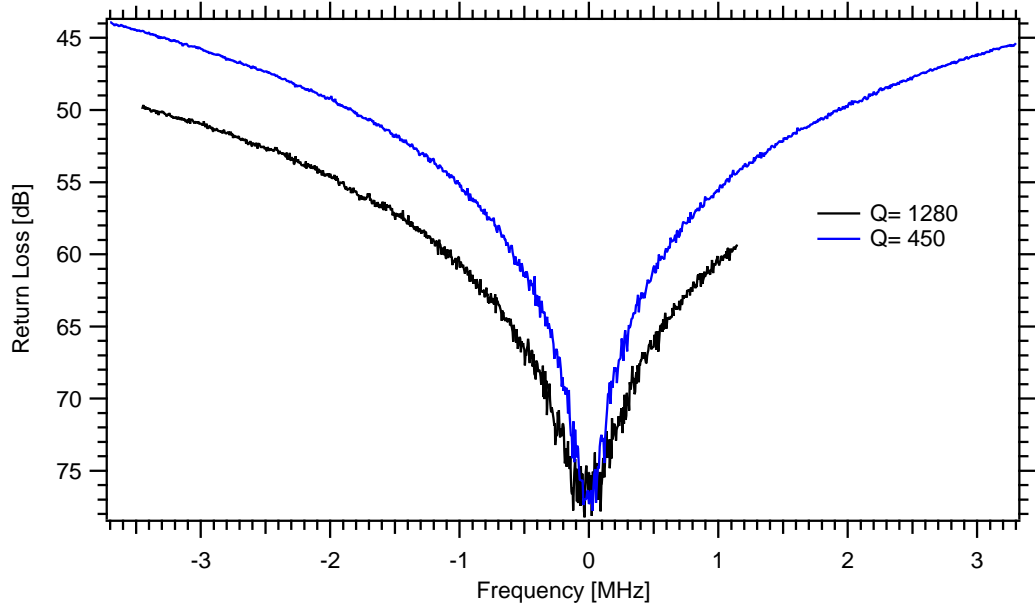
Figure 11.10: The cavity return loss as the frequency is swept, and the locked cavity follows. The zero of the frequency axis is adjusted to be at the best-coupled location. The higher Q cavity (black curve) remains better matched over the same frequency range. At small offsets from critical coupling, the measurement is limited by the spectrum analyzer noise floor.

< -8. In terms of easily measured derivatives, equation 11.14 can be expressed as[6],

$$\left.\frac{\partial \omega}{\partial L}\right|_x = -\frac{\left.\frac{\partial |S|}{\partial L}\right|_\omega}{\left.\frac{\partial |S|}{\partial \omega}\right|_L - \left.\frac{\partial |S|}{\partial \omega}\right|_x}, \tag{11.17a}$$

$$= \left.\frac{\partial \omega}{\partial L}\right|_{|S|} \left[\frac{1}{1 - \left.\frac{\partial |S|}{\partial \omega}\right|_x \left.\frac{\partial \omega}{\partial |S|}\right|_L}\right]. \tag{11.17b}$$

The derivatives in equation 11.17 can be measured as follows. As the position of the coupling screw is changed while the cavity length is fixed, the resonant

---

inductance (or the size of the aperture). In this case, S can be written as a function of one parameter only.

[6]In going from equation (11.17a) to equation (11.17b), the identity $\left.\frac{\partial |S|}{\partial L}\right|_\omega \left.\frac{\partial \omega}{\partial |S|}\right|_L \left.\frac{\partial L}{\partial \omega}\right|_{|S|} = -1$ has been used.
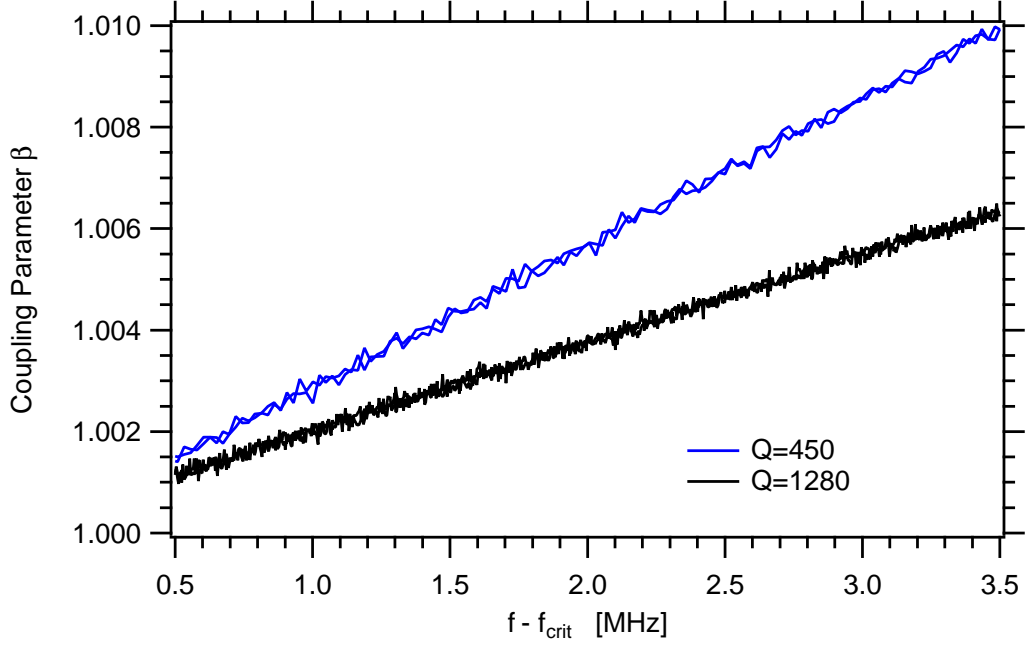
Figure 11.11: The coupling coefficient increases linearly as the frequency is offset from the best-coupled one. The slope of the curves are $\left.\frac{\partial \beta}{\partial f}\right|_x = (2.8 \pm 0.2) \times 10^{-3}$ MHz^{-1} and $(1.7 \pm 0.1) \times 10^{-3}$ MHz^{-1} for the cavities of Q=450 and Q=1280 respectively.

frequency shifts, as was displayed in figure 9.9. Plotting the frequency shift against the coupling parameter β gives figure 11.9. The slope of this curve is[7] $\left.\frac{\partial f}{\partial \beta}\right|_L = -\frac{1}{2Q_u} \left.\frac{\partial \omega}{\partial |S|}\right|_L$, and is equal to -21.5 ± 0.5 MHz or -9.4 ± 0.2 MHz for the cavities of Q=450 and Q=1280 respectively.

It is also observed that when the cavity is locked to the source, and the frequency is swept, the coupling (return loss) changes as shown in figure 11.10. The corresponding coupling parameter β is plotted in figure 11.11 as a function of the frequency offset from critical. The slope of this curve is $\left.\frac{\partial \beta}{\partial f}\right|_x = -2Q_u \left.\frac{\partial |S|}{\partial f}\right|_x$ and is equal to $(2.8 \pm 0.2) \times 10^{-3}$ MHz^{-1} or $(1.7 \pm 0.1) \times 10^{-3}$ MHz^{-1} for the cavities of Q=450 and Q=1280 respectively.

Plugging these values into equation 11.17b gives $\left.\frac{\partial \omega}{\partial L}\right|_x = \left.\frac{\partial \omega}{\partial L}\right|_{|S|} (1 + \epsilon)$, where $\epsilon = -.06 \pm .004$ and $-.016 \pm .001$ for the low and high Q cavities respectively.

---

[7]In the notation of chapter 3, $\beta = \frac{\rho}{\nu}$, $\nu = \frac{\pi}{Q_u}$, and $|S| = e^{-\rho}$. So, $\beta = \frac{Q_u}{\pi}(1 - |S|)$.
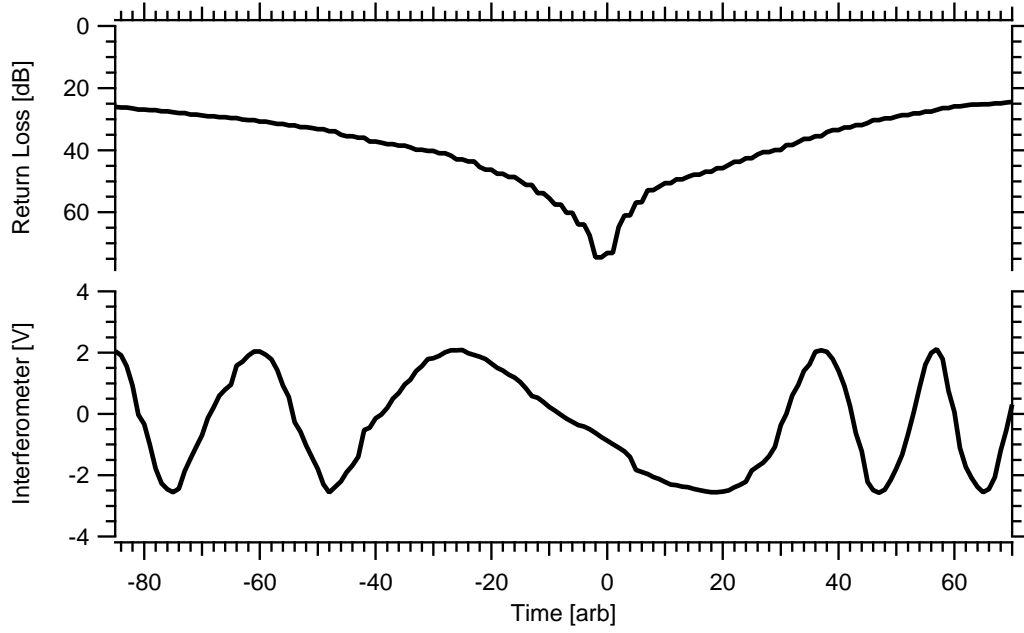
Figure 11.12: The cavity return loss and interferometer signal as the position of the coupling screw is changed. The cavity starts under coupled, and the cavity length becomes shorter as the coupling is increased. The positions of the maxima and minima in the interferometer signal indicate a length step of $\lambda/4 = 327.5$ nm. The coupling coefficient β (calculated from the return loss at those locations) is plotted against the length difference in figure 11.13. The flattening of the return loss curve near critical coupling is due to the measurement hitting the spectrum analyzer noise floor, and does not reflect the actual return loss.

If condition 11.16 were true, ϵ would be the correction to the frequency-distance slope due to the coupling structure. In general, however, an additional derivative, either $\left.\frac{\partial\omega}{\partial L}\right|_{|S|}$ or $\left.\frac{\partial|S|}{\partial L}\right|_{\omega}$ is needed (see equation 11.17a).

Evidence that $\left.\frac{\partial\arg S}{\partial|S|}\right|_{\omega} \left.\frac{\partial|S|}{\partial\arg S}\right|_{x}$ does not equal 1 for this coupling structure comes from consideration of the low-Q cavity (see section 11.4). In that case, the ratio of $\left.\frac{\partial\omega}{\partial L}\right|_{|S|}$ to $\left.\frac{\partial\omega}{\partial L}\right|_{x}$ is known to be $1.06 \pm .004$, but the frequency-distance slope only deviates from the ideal by $-4.7\% \pm 0.5\%$. The only way this can be reconciled is if $\left.\frac{\partial\arg S}{\partial|S|}\right|_{\omega} \left.\frac{\partial|S|}{\partial\arg S}\right|_{x}$ is greater than 1, or an unknown effect of comparable magnitude is balancing it out.

Measuring $\left.\frac{\partial\omega}{\partial L}\right|_{|S|}$ directly would require some modification of the hardware to allow automated tuning of the coupling structure. An error signal can be

141
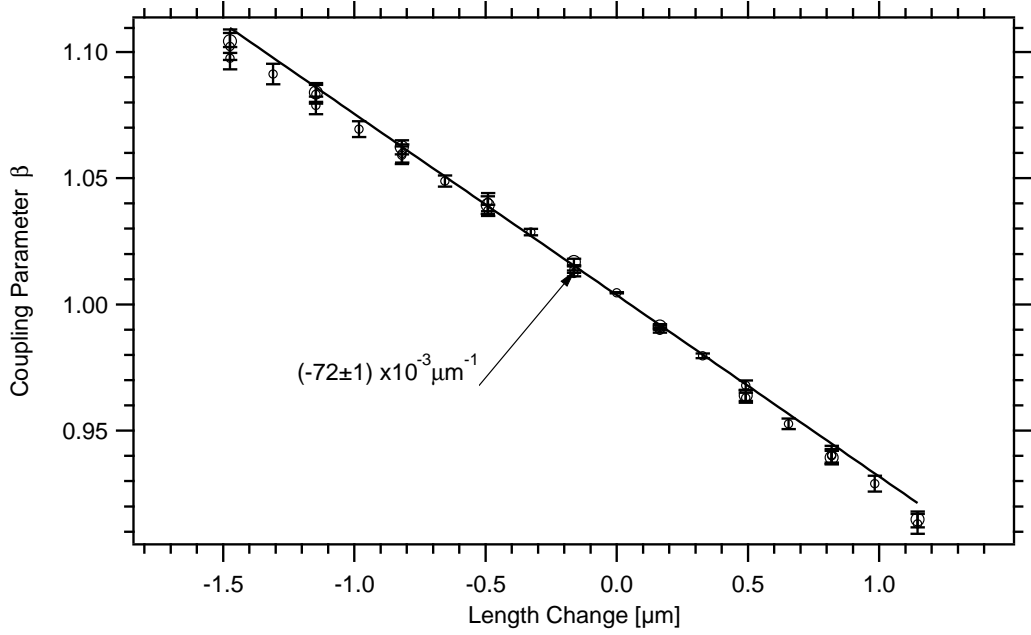
Figure 11.13: The coupling coefficient β is plotted against the change in cavity length (for the Q=1280 cavity) when the coupling screw is adjusted and the cavity resonant frequency is held fixed at 9815 MHz. The slope of the curve is $\frac{\partial\beta}{\partial L}\big|_\omega = (-72 \pm 1) \times 10^{-3} \ \mu\text{m}^{-1}$.

generated with the single-sideband detection technique discussed in sections 4.4 and 5.6, and fed back to keep the coupling constant. However, as this hardware is not in place, $\frac{\partial|S|}{\partial L}\big|_\omega$ is measured instead, and $\frac{\partial\omega}{\partial L}\big|_{|S|}$ is calculated using the identity $\frac{\partial|S|}{\partial L}\big|_\omega \ \frac{\partial\omega}{\partial|S|}\big|_L \ \frac{\partial L}{\partial\omega}\big|_{|S|} = -1$.

It is relatively easy to measure $\frac{\partial|S|}{\partial L}\big|_\omega$ thanks to the frequency lock and interferometer already implemented. The cavity is locked to a frequency of 9815 GHz, while varying the position of the screw. The interferometer signal and cavity reflection are recorded, and shown in figure 11.12. As the coupling is increased, the cavity becomes shorter. The return loss at the maxima and minima of the interferometer signal is extracted to calculate the coupling coefficient β, and is plotted against the cavity length change in figure 11.13. The slope of the curve is $\frac{\partial\beta}{\partial L}\big|_\omega = -\frac{Q_u}{\pi} \frac{\partial|S|}{\partial L}\big|_\omega = (-72 \pm 1) \times 10^{-3} \ \mu\text{m}^{-1}$, for the Q=1280 cavity.

This gives a value of $\frac{\partial f}{\partial L}\big|_{|S|} = 675 \pm 20 \ \text{Hz nm}^{-1}$. Unfortunately, the uncer-

tainty is too high for the measurement to be valuable when looking for an effect that is less than 1%. In fact, the remaining $-0.4\% \pm 0.6\%$ discrepancy and the value of ϵ above, constrain its value much better, $\frac{\partial f}{\partial L}\big|_{|S|} = 650.5 \pm 3.5$ Hz nm^{-1}. This, combined with equations 11.14, 11.15, and 11.17b, constrains the value of $\frac{\partial \arg S}{\partial |S|}\big|_{\omega} \frac{\partial |S|}{\partial \arg S}\big|_{x}$ to,

$$\frac{\partial \arg S}{\partial |S|}\bigg|_{\omega} \frac{\partial |S|}{\partial \arg S}\bigg|_{x} > 1.6, \quad \text{or,} \tag{11.18a}$$

$$\frac{\partial \arg S}{\partial |S|}\bigg|_{\omega} \frac{\partial |S|}{\partial \arg S}\bigg|_{x} < -8. \tag{11.18b}$$

If it remains positive and less than 2, then $\frac{\partial f}{\partial L}\big|_{|S|}$ will give a slope closer to the unperturbed slope than $\frac{\partial f}{\partial L}\big|_{x}$, and vice versa.

In summary, direct measurements of corrections to the frequency-distance relationship due to the coupling structure are not sufficiently precise to be valuable. However, considering that other effects explain the deviation down to $0.4\% \pm 0.6\%$, it can be safely concluded that the influence of the coupling structure is smaller than the mechanical reproducibility of the setup.

11.4   Corrections for the Low-Q Cavity.

Since the Q=450 cavity was re-plated with a higher purity gold to obtain the Q=1280 cavity, some of the data sets required for piecing together the frequency-distance corrections for it are unavailable - specifically the values of w and $\frac{\partial |S|}{\partial L}\big|_{\omega}$. The limited frequency-distance slope data available for the low-Q cavity are shown in figure 11.14, and are $-4.7\% \pm 0.5\%$ away from the unperturbed cavity expectation.

The gap width w was not measured directly, and was estimated by another means. From the shift of resonant frequency in going from uncoupled to critically coupled, -21.5 ± 0.5 MHz in this case, the coupling hole is estimated to make the
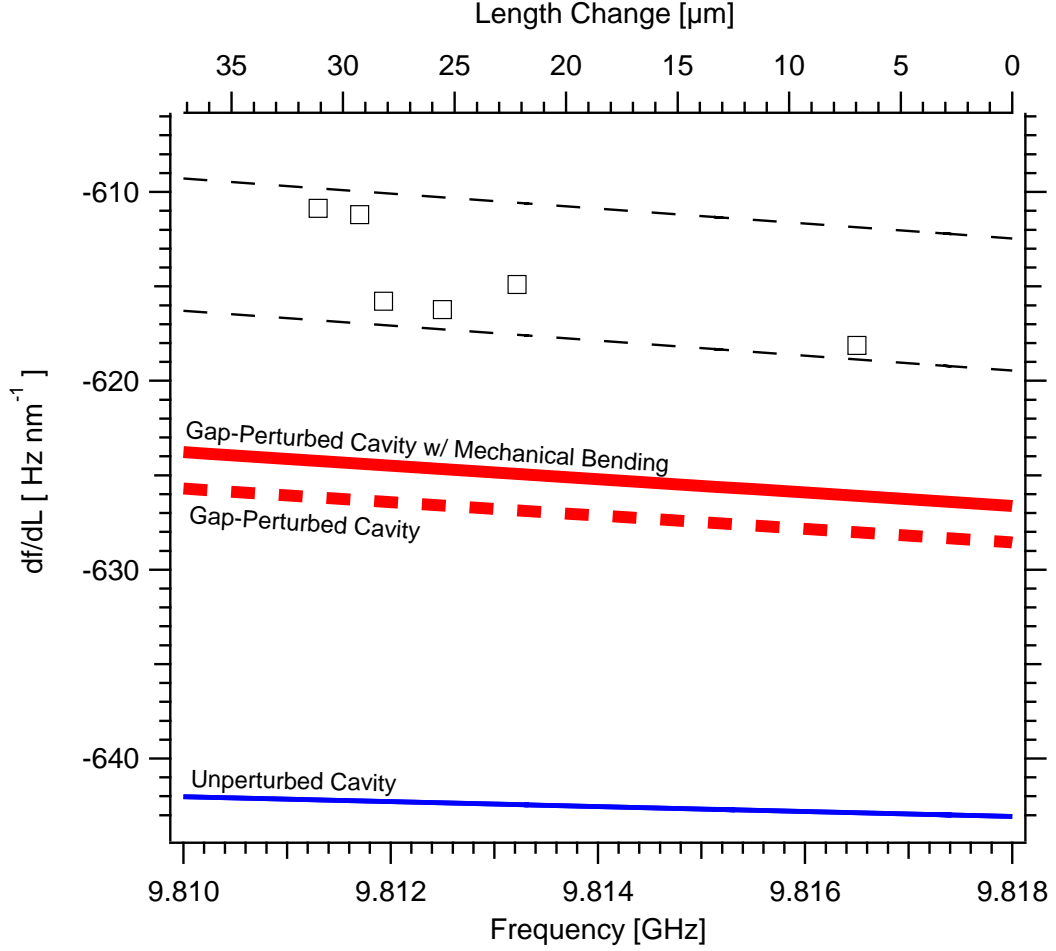
Figure 11.14: The slope of the frequency-distance relationship for an unperturbed cavity (blue line) compared to the measured slope (black points) as a function of frequency for the Q=450 cavity. The dashed black lines indicate the range of mechanical repeatability of the measurements. The dashed red line indicates the expected slope as a function of frequency for a gap-perturbed cavity, using the theory developed in section 11.1. The solid red line indicates the expected slope when the gap and mechanical bending, section 11.2, are taken into account.

fixed half-cavity 33.5 ± 1 µm longer. Since this is 19 µm more than with the high Q cavity, w is estimated to be 19 µm shorter to make up for it, $w = 102 \pm 3$ µm at 9815 MHz. The off-center distance z is also obtained as 50 ± 10 µm. Then the correction expected due to the presence of the gap is 15.25 ± 0.75 Hz nm^{-1} at 9815 MHz, and shown by the dashed, red curve in figure 11.3. The solid, red curve is obtained by adding the bending correction.

After those two corrections, there remains a $-1.5\% \pm 0.5\%$ deviation that can likely be attributed to the coupling structure. Although the ratio of $\frac{\partial \omega}{\partial L}\big|_{|S|}$ to $\frac{\partial \omega}{\partial L}\big|_x$ is known to be $1.06 \pm .004$ (see section 11.3), the absence of a measurement of $\frac{\partial |S|}{\partial L}\big|_\omega$ precludes comparison between the remaining deviation and coupling effects.

11.5 Improper Signal Generation

The μPDH signal is excellent at rejecting unwanted effects from varying coupling if it is properly set up. However imperfections in the signal generation process due to imbalance of sideband amplitude, insufficiently high modulation frequency, the non-orthogonality of the sum of the sidebands and the carrier, improperly-phased local oscillator, or combinations thereof, can cause the varying coupling to influence the error signal a small amount. The size of these effects can be estimated from the general form of the μPDH signal, equation 4.9, and the measured frequency dependence of the coupling in figure 11.10. They all turn out to be small, and they sum up to a worst-case $\sim 10^{-3}$ effect on the frequency-distance slope of the Q=450 cavity, and an order of magnitude smaller for the Q=1280 cavity.

11.5.1 Insufficiently high modulation frequency and imbalanced sideband amplitude

The first term in brackets in line 4.9a reflects the influence of varying coupling (a presence of $\Re[\Gamma_0]$) on the signal if the modulation frequency is insufficiently high and there is an imbalance in the sideband amplitude. In the presence of sideband imbalance, it is modified to,

$$\Re[\Gamma_0]\Im[\Gamma_+]\frac{\Delta P_s}{P_s} - \Im[\Gamma_0] = 0, \tag{11.19a}$$

$$\rightarrow \quad (f - f_0) = d_{ModFreq}\,(f - f_{crit}), \tag{11.19b}$$

145

where $2\Delta P_s$ and P_s are the imbalance and average of the sideband powers (power in the lower sideband assumed higher), and $\Im[\Gamma_-] = -\Im[\Gamma_+]$ has been used. This leads to an offset between the source and cavity frequencies $(f - f_0)$ that is proportional to the difference between the operating frequency and the frequency at which the cavity is critically coupled $(f - f_{crit})$. Fortunately this effect can be reduced by minimizing $\Im[\Gamma_+]$ by increasing the modulation frequency. The size of the constant $d_{ModFreq}$ is estimated for both the Q=450 and Q=1280 cavities at a modulation frequency of 77 MHz and 2 dB sideband imbalance, to be on the order of 5×10^{-4} and 4×10^{-5} respectively.

11.5.2 Carrier and sideband non-orthogonality

The first term in brackets in line 4.9b reflects the influence of varying coupling if the sum of the sidebands is not orthogonal to the carrier (by angle ξ defined in figure 4.3b), which can be due to imbalance in the IQ channels of the vector modulator or to the presence of group delay as discussed in section 4.3. The relevant terms are,

$$\Re[\Gamma_0]\sin\xi - \Im[\Gamma_0]\cos\xi = 0, \qquad (11.20\text{a})$$

$$\rightarrow \quad (f - f_0) = d_{Orth}\,(f - f_{crit}). \qquad (11.20\text{b})$$

This leads to an offset between the source and cavity frequencies that is proportional to the difference between the operating frequency and the frequency at which the cavity is critically coupled. The size of the constant d_{Orth} is estimated for both the Q=450 and Q=1280 cavities fabricated and $\xi = 1°$, to be on the order of 3×10^{-4} and 6×10^{-5} respectively.

11.5.3 Improperly-phased LO

The first term in brackets in line 4.9c reflects the influence of varying coupling if the down-converting mixer's local oscillator phase is not optimal (off by angle θ) and there is imbalance in the sideband amplitude. The procedure followed for choosing the optimal phase by examining the phase-frequency plot (see section 5.3) reproducibly picks the optimal phase in a range of $\pm 2°$ - so it is of interest to know the magnitude of the effect of being off by a couple of degrees. After modification due to the presence of sideband imbalance, the relevant terms are,

$$\frac{\Delta P_s}{P_s} \Re[\Gamma_0] \sin \theta - \Im[\Gamma_0] \cos \theta = 0, \tag{11.21a}$$

$$\rightarrow \quad (f - f_0) = d_{LO} \left(f - f_{crit} \right). \tag{11.21b}$$

This is a third effect that leads to an offset between the source and cavity frequencies that is proportional to the difference between the operating frequency and the frequency at which the cavity is critically coupled. The size of the constant d_{LO} is estimated for both the Q=450 and Q=1280 cavities fabricated, $\theta = 1°$ and a 2 dB sideband imbalance to be on the order of 7×10^{-5} and 2×10^{-5} respectively.

11.6 Other Considerations

A number of other miscellaneous corrections, all relatively small, were considered and summarized here. The diode laser wavelength was known to $\sim 5 \times 10^{-5}$, and fluctuated within that range due to mode hops. Changes of the index of refraction of air at both 10 GHz and 1310 nm are on the order of 10^{-6} (see chapter 10). The Gouy phase of a Gaussian beam[8] (the interferometer works with focused light) changes the spacing between the interferometer fringes by a factor $\sim 4 \times 10^{-5}$.

---

[8]The phase of a Gaussian beam is given by $kz + \arctan z/z_0$, where z is measured from the beam waist, and z_0 is the Rayleigh length [50, 51].

11.7 Summary of Frequency-Distance Corrections

Measurements have shown that in the case of the Q=1280 cavity, the slope of the frequency-distance relationship deviates from the ideal slope by -3.5%±0.5% at a frequency of 9815 MHz. Numerical and experimental evidence indicate that $-2.8\% \pm 0.1\%$ can be attributed to the presence of a 121 ± 2 µm gap within the cavity. If a 100 µm scan range is desired, then there is not much to win by reducing the gap width. On the other hand, with a scan range of 10 µm, this can be reduced by a factor of 10 by operating closer to 9890 MHz, and using an appropriate VCO.

Direct experimental measurement of coupling effects is not precise enough to be of value. However, considering that the remaining deviation is $-0.7\% \pm 0.6\%$, and about -0.3% can be attributed to stress-induced bending of the cavity, it can be safely concluded that coupling effects are less than 1%, and within the level of the mechanical reproducibility of the setup.

Deviations from the ideal frequency-distance relationship are currently all attributable to the physical design of the cavity. Imperfections in the setup of the feedback signal appear to be negligible. If the predictable effects of the gap and cavity bending are anticipated and corrected by an algorithm or some other means, the frequency-distance slope of the current cavity can be brought within the range of mechanical reproducibility, $\sim 0.5\%$, of the ideal slope. Further improvements will require tighter alignment tolerances and finer mechanics.

CHAPTER 12

Potential Applications

12.1 Calibrated Scanning Probe Microscopy

The main motivation for developing the frequency-distance encoder is for use in calibrated scanning probe microscopy, as discussed in chapter 1. Most approaches have used real-time interferometry based on iodine-stabilized He-Ne lasers to realize the meter [3, 10, 11]. These require complex and expensive laser systems, and the resulting noise-level is typically a few tenths of a nanometer. In addition, they begin to loose accuracy when measuring distances smaller than a fringe due to periodic non-linearity commonly seen in interferometers [12]. The frequency-distance encoder goes to the opposite limit - the entire scan range is within one fringe. Its resolution is determined by the source frequency noise [26] and can be as low as $60 \, \text{fm} \, \text{Hz}^{-1/2}$ using commonly available oscillators. Further, it is far more compact and inexpensive. Challenges involved in realizing metrological SPM are described in [52] and [53].

Incorporating the frequency-distance encoder into an SPM is not a trivial task. The x-y scanner will have to be redesigned to incorporate one encoder per axis, while minimizing cross-talk. In addition, system stability must be on par with the sensor noise, to fully take advantage of its capability. A resolution of a couple of picometer is useless if the image drifts several nanometer while acquiring an image. Putting the entire assembly in a thermally stabilized environment helps, but is not sufficient. Commercially available scanning probe microscopes typically

drift on the order of 200 nm K^{-1} so even millikelvin stability would correspond to hundreds of picometer of drift.

12.2 Laser Stabilization

Atomic-Molecular-Optics (AMO) physicists go through great lengths to build narrow and stable lasers. Cross-locking of lasers is a basic tool of modern optical research laboratories that enables them to build any-wavelength, high-quality, stable lasers starting from relatively broadband and unstable semiconductor light sources. In the multiple-step process to achieve a narrow wavelength distribution, the last one(s) is (are) to lock the light source to a high finesse Fabry-Perot resonator via PDH in a broader bandwidth than the incoming light distribution. However, the stability of the laser constructed in this way will be equal to the length stability of the last Fabry-Perot resonator. To assure long-term length stability, the resonator (built with mirrors of high reflectivity at both wavelengths) is also locked to a different wavelength light that is itself locked to a stable reference (typically an atomic transition), transferring stability from one wavelength to the other. As the cavity length must be an integer number of half wavelengths at both frequencies, the frequencies must differ by exactly a multiple of the free spectral range of the cavity. While a stability of 10^{-11} - 10^{-12} is routinely achieved using a Rubidium gas cell as a reference, the technique is fairly cumbersome and not finely tunable.

The state-of-the-art in stable, optical frequency synthesis makes use of mode-locked lasers and optical frequency combs to link lasers of any wavelength to a stable reference [55]. One line of the comb is locked to an atomic reference, while another is locked to a continuous-wave (CW) laser. Fine-tuning the repetition frequency of the mode-locked laser allows continuous tuning of the CW laser.

Stabilizing lasers with electronic oscillators is an alternative to optical cross-
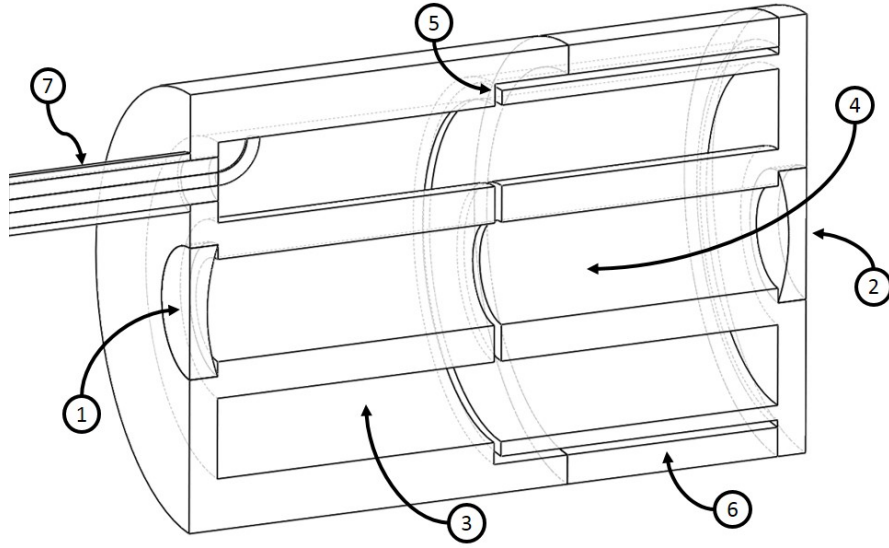
Figure 12.1: The "dual-cavity" - a structure containing both an optical and a microwave resonator that allows locking the optical and microwave frequencies [54]. The image shows a cross section of the structure, which is cylindrically symmetric, apart from the input cable. High-quality mirrors (1) and (2) are set coplanar with the ends of a coaxial microwave cavity (3), forming an optical Fabry-Perot etalon (4) of equal length with the microwave cavity. The cavity is split down the center (5) to allow its length to be fine-tuned with a piezo tube actuator (6). The microwave enters the cavity via a waveguide or coaxial cable (7).

locking. DeVoe and Brewer [56] developed a dual-frequency laser modulation scheme that both locks the laser to an optical cavity and the cavity in turn to an electronic reference. A single, continuous-wave laser is modulated at the cavity Free Spectral Range (FSR) to produce sidebands whose separation equals the spacing of the cavity axial modes. What follows in effect performs a Pound-Drever-Hall lock (PDH, see chapter 2) with all of the sidebands in parallel. The cavity is adjusted to keep its FSR, and hence the length, fixed. Sandford and Antill [57] measured the beating of two independent lasers, each PDH locked to adjacent axial cavity modes, to measure the cavity FSR. The beat note was phase locked to a reference, stabilizing the cavity length. Mode-locked lasers have also been referenced to stable electronic oscillators [58].

An alternative method of locking lasers to microwave oscillators can be envis-
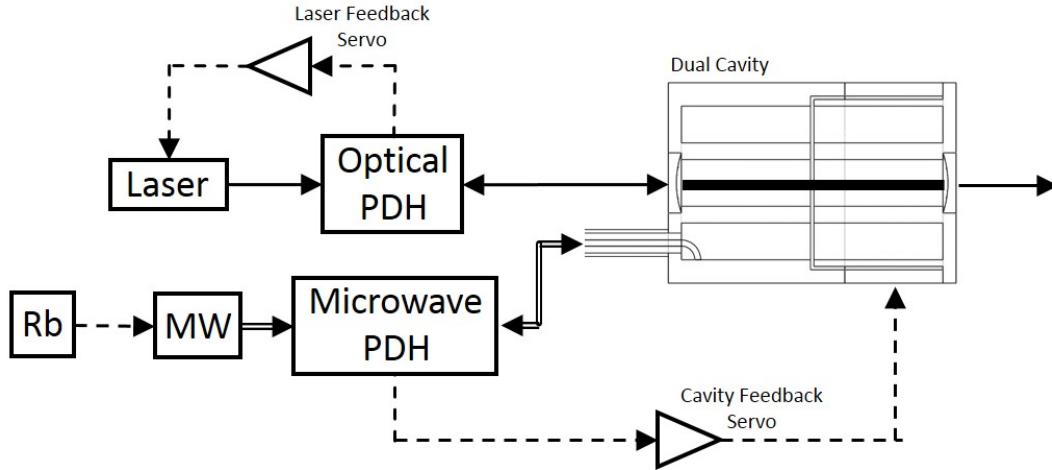
Figure 12.2: A system that uses the dual-cavity (figure 12.1) to transfer the frequency stability of a reference oscillator to a laser. The cavity length can be fine-tuned by altering the microwave frequency, in contrast to other schemes capable of achieving such a stability.

aged based on the work presented here. The frequency-distance encoder described in this dissertation is modified to house a high-quality optical Fabry-Perot resonator, adequate to be the last stage reference for a sub- Hz linewidth laser. The Fabry-Perot is placed in the center of a split coaxial microwave resonator, whose length is actively controlled and stabilized through the resonant frequency of the coaxial resonator [54]. This structure is referred to as the "dual-cavity".

The suggested architecture and key components of the dual-cavity are illustrated in figure 12.1. High-quality reflective mirrors (1) and (2) are set coplanar with the ends of a coaxial microwave cavity (3), forming an optical Fabry-Perot etalon (4) of equal length with the microwave cavity. The cavity is split down the center (5) to allow its length to be fine-tuned over several micron with a piezo tube actuator (6). The microwave enters the cavity via a waveguide or semirigid cable (7).

Figure 12.2 is a block diagram of a system utilizing the dual-cavity to lock the laser and microwave frequencies. A microwave source, whose frequency lies within the tuning range of the resonant cavity, is phase-locked to the Rubidium (or

otherwise) reference. A μPDH lock generates an error signal that is proportional to the frequency difference between the source and the coaxial cavity. The error signal is fed back via a servo to the dual-cavity piezo, which adjusts the length so that its resonant frequency matches the source. With proper mechanical design, the length stability of the microwave cavity is passively transferred to the optical Fabry-Perot etalon. Standard optical PDH generates an error signal that is fed back through a second servo to the laser, keeping its frequency locked to the optical resonator. Thus, the laser output has inherited the stability of the reference oscillator.

The cavity length is tunable over a several micron range with the precision and stability of the microwave source frequency. While the tuning range of the etalon exceeds the FSR of the optical cavity, its long-term stability (Allan variance) is comparable with the used Rubidium clock reference (2×10^{-12} over $100\,\mathrm{s}$) offering a convenient alternative to the currently used techniques. Although the 10^{-12} stability achievable in this way with commonly available equipment falls short of the stat-of-the-art in terms of highest stability achieved with rigid mirror distances, in this case the distance between mirrors is fine-tunable, allowing for continuous tuning of laser frequencies without compromising stability. In addition, it is less complex and less expensive, requiring fewer optical components, than other methods used, and is sufficient for many laboratory research projects.

12.3   Cavity-Stabilized Oscillators (CSO)

Microwave oscillators can be stabilized with the μPDH lock, in fact, this was the application that first led to the Pound-Stein-Turneaure (PST) lock [24–26]. Oscillators were locked to high-Q, superconducting resonators to produce the cleanest sources available. The quest for high-Q led to the development of the "whispering-gallery mode" sapphire-loaded resonator, which boasts unloaded Qs on the order

of 200,000 at room temperature or 6×10^9 at cryogenic temperatures [59]. Phase noise as low as -146 dBc Hz^{-1} at 1 kHz offset from a 10.24 GHz source can be achieved with room temperature cavities.

In recent CSO designs [60], sophisticated versions of the DC-AFC (figure 2.1c) have been employed to produce the feedback signal. The performance of the lock is tied to the noise figure of the feedback loop components. The µPDH lock is an alternative method of generating the feedback signal that uses different components and may provide an advantage. A detailed, side-by-side evaluation of the required hardware for both types of locks is warranted.

The fact that the "noise" of the µPDH lock used for the frequency-distance encoder is limited by the source frequency noise above 1 kHz, guarantees that the same signal could be used instead to clean up the source. Of course, the cavity would preferably have to be designed for high-Q and stability, rather than for distance control. Nevertheless, it appears that even standard, room-temperature microwave cavities could improve the phase-noise of some sources by 10-20 dBc Hz^{-1} or more.

12.4   Other Applications

The µPDH lock can be used to measure and control virtually any phenomenon that can influence a microwave resonant circuit. A microwave cavity made of some high-thermal-expansion material can be used as thermometer, as its resonance frequency will depend on temperature. A thermal expansion of 10^{-5} K^{-1} would give a frequency shift on the order of 100 kHz per degree for a 10 GHz cavity. Based on the < 0.1 Hz Hz$^{-1/2}$ sensitivity already demonstrated, temperature measurement with µK resolution is possible.

Another possible application is a magnetic field controller based on an Yttrium Iron Garnet (YIG). YIGs resonate at a microwave frequency dependent on the

154

applied DC magnetic field with a sensitivity of 2.8 MHz per Gauss [61]. Due to the relatively low quality factor of YIGs (a few hundred at 10 GHz), the µPDH signal-to-noise is expected to decrease somewhat. Despite this, field resolution better than 10^{-5} Gauss seem achievable, with the same stability as the frequency standard.

CHAPTER 13

Concluding Remarks

This work evaluates the performance of microwave cavities as calibrated, atomic-resolution, displacement measuring devices for scanning probe microscopes. Accurate and stable RF frequency references have become ubiquitous and are an attractive option for realizing traceable distance measurements through the resonant frequency of microwave cavities operating in TEM modes, a device referred to as the frequency-distance encoder.

The control electronics necessary for such a device include a precisely tunable source, a method of automatic frequency control, and a feedback servo. The error signal is generated by the µPDH technique, whose $60\,\text{fm}\,\text{Hz}^{-1/2}$ noise floor is limited by the frequency-noise of the sources chosen. The used Rubidium reference clock provides an accuracy of 5×10^{-11}, with comparable stability.

Required features of the cavity - the coupling structure, adjustable length, flexture-based kinematic positioning - perturb the frequency-to-length conversion, and compromise the accuracy of the displacement measurement, $\Delta L = -\frac{c}{2f^2}\,\Delta f$. The largest correction, a couple of percent in this case, is due to the presence of a 120 µm gap, and can be reduced proportionally with the scan range. Further, the effect of the gap can be predicted, and mostly accounted for. Coupling effects are $\lesssim 0.5\%$ and comparable to the mechanical reproducibility of the assembly, $\pm 0.5\%$.

The resolution of the realized device is equal to the µPDH noise, $60\,\text{fm}\,\text{Hz}^{-1/2}$. The range is 10 µm, limited by the mechanical stage and piezo actuator used. With

current design and known corrections, displacement measurements are accurate to 1% without independent calibration. Non-linearities due to the gap can be subtracted out, with remaining non-linearity $< 5 \times 10^{-4}$ of the scan range, ± 2.5 nm over 10 µm. The frequency-distance encoder would benefit most from improved mechanics, higher quality sources, and better Q. In the meantime, independent calibration (valid as long as the mechanics are not changed) can increase the accuracy by two orders of magnitude, where non-idealities in the µPDH signal become relevant.

These capabilities are compared to the results of atomic-resolution metrology obtained by the Molecular Measuring Machine (M^3) at NIST [3, 7]. The x and y axes of M^3 are measured with polarization-encoded heterodyne interferometers, whose noise level is a few 100 pm in a relatively low bandwidth. Over the 80 nm fringe period, there is a 2 nm non-linearity, resulting in as much as an 8% deviation of measured lattice parameters from bulk crystal x-ray diffraction measurements. The frequency-distance encoder described in this dissertation can resolve 2.5 pm in 1 kHz bandwidth, with non-linearity < 0.04 nm over a comparable range. A 1% measurement of lattice parameters could be made without independent calibration, and better if one is available. This accuracy is comparable to crystallographic data, and would allow establishing quantitative differences between surface and bulk structures, if incorporated in a scanning probe microscope.

APPENDIX A

Scattering Coefficients of the Coupling Structure and the Resonance Condition

The resonant cavity and coupling structure may be viewed as the network in figure A.1. An incoming transmission line (assumed matched) is connected to port 1 of a two-port coupling structure characterized by the generalized scattering matrix $[S]$. Port 2 is connected to a shorted transmission line of length L. The signal
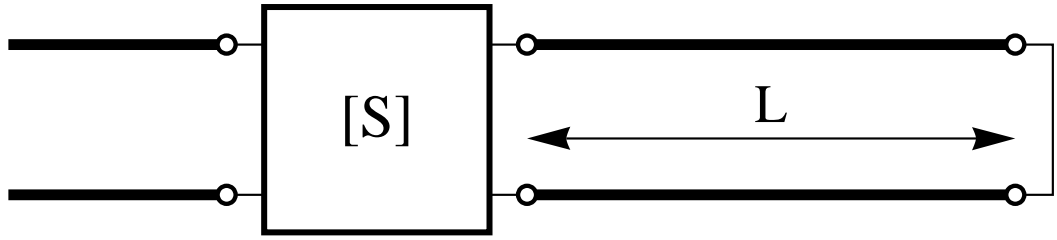

Figure A.1: The cavity may be represented by the network above. The two-port coupling structure is characterized by the scattering matrix $[S]$, and the cavity by a shorted transmission line of length L.
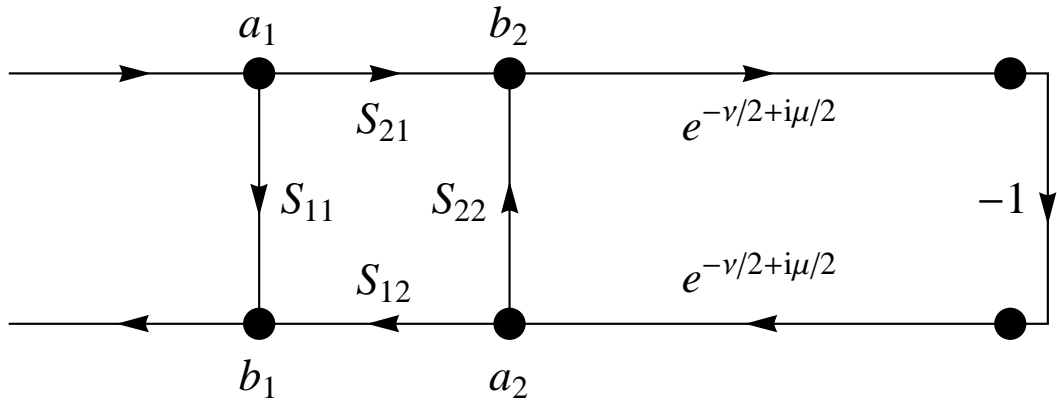

Figure A.2: The signal flow graph for the network in figure A.1.

flow graph of the network is shown in figure A.2. a_n and b_n are the amplitudes of the waves incident and reflected from port n respectively. The attenuation and phase contributed by the round-trip travel from port 2 to the short and back to port 2 is $e^{-\nu+i\mu}$.

Then, the reflection coefficient for the cavity is,

$$\frac{b_1}{a_1} = \frac{S_{11} + (S_{11}S_{22} - S_{12}S_{21})e^{-\nu+i\mu}}{1 + S_{22}e^{-\nu+i\mu}}, \tag{A.1}$$

and,

$$\frac{b_2}{a_1} = \frac{S_{21}}{1 + S_{22}e^{-\nu+i\mu}}. \tag{A.2}$$

The components of the scattering matrix are not independent. If the coupling structure is built from passive components, it will be reciprocal, and $[S] = [S]^t$. If the structure is also lossless, then the scattering matrix is unitary. With these two assumptions, the most general scattering matrix is,

$$[S] = \begin{bmatrix} |S_{11}|\,e^{i\delta_1} & \pm i\sqrt{1-|S_{11}|^2}\,e^{i(\delta_1+\delta_2)/2} \\ \pm i\sqrt{1-|S_{11}|^2}\,e^{i(\delta_1+\delta_2)/2} & |S_{11}|\,e^{i\delta_2} \end{bmatrix}. \tag{A.3}$$

With $S_{11} = S_{22} = -r$, the reflection coefficient defined in chapter 3 is recovered, and the cavity resonance condition is $\mu = 2\pi n$ for integer n.

If the phase δ_2 is different from an odd multiple of π, the resonance condition will be altered. To extract out the minus sign from S_{22} explicitly, let $\delta_2 = (2m+1)\pi + \phi$ for an integer m. The resonance condition is then, $\mu + \phi = 2n\pi$. This condition is used to understand the effects of a phase shift associated with changing the coupling in section 11.3.

APPENDIX B

Charge Pump Filter Optimization

The fastest cavity response that can be hoped for is one that equals the settling time of the VCO and Phase-Lock-Loop (PLL) on the AD9956 testboard. The bandwidth of the PLL is set by the VCO voltage-to-frequency sensitivity, the charge-pump (CP) phase detector current, and the filter after the charge-pump. The CP base current for the AD9956 chip is 0.5 mA. This base value can be multiplied by a programmable integer from 1-8, the CP scale. The CP filter on the testboard is shown in figure B.1. The values of the components as set by the factory are listed in the table under the column "Filter 1".

The open loop µPDH signal is the ideal signal to analyze in order to measure the settling time of the VCO PLL loop, as it is a measure of the instantaneous VCO frequency when the cavity is fixed. Figure B.2 shows the µPDH signal as the frequency of the VCO is stepped by 5 kHz for the lowest and highest CP current scaling. The response of a PLL with a second-order loop transfer function
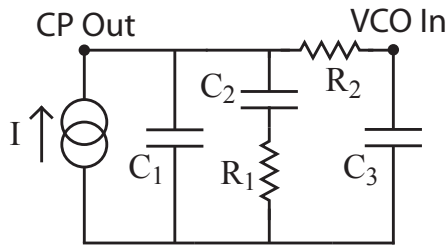
| Component	Filter 1	Filter 2
C_1	0.1 µF	3.3 nF
C_2	2.7 µF	27 nF
C_3	15 nF	1.5 nF
R_1	25 Ω	180 Ω
R_2	158 Ω	40 Ω

Figure B.1: The filter following the charge-pump (CP) in the VCO PLL loop. The component values set the bandwidth of the loop.
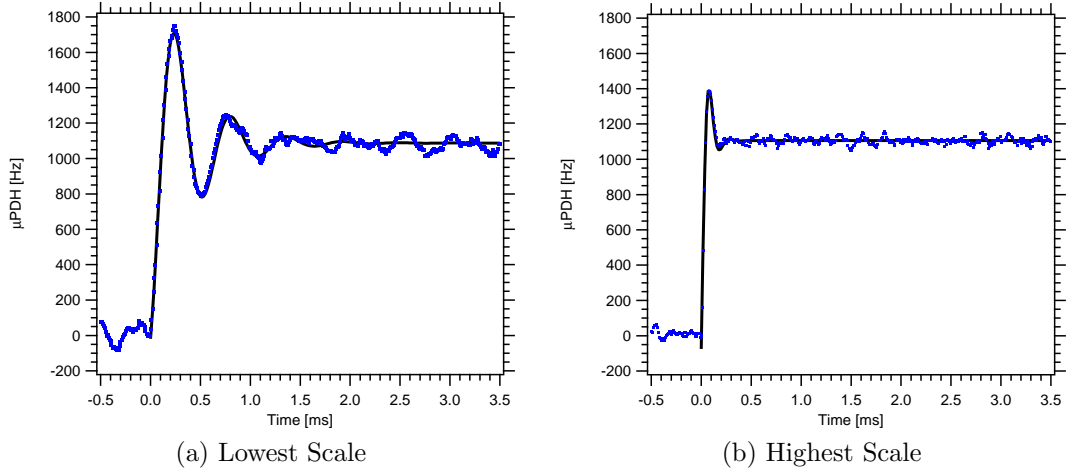
Figure B.2: The µPDH signal as the VCO setpoint is stepped by 5 kHz for the lowest and highest charge pump current scalings. The µPDH signal is the instantaneous frequency difference between the VCO and a fixed cavity. Curves shown were measured with filter 1.

is similar to a damped oscillator [62], and the observed response is indicative of an underdamped loop. While the loop transfer function with the filter shown in figure B.1 is not strictly second-order, components are chosen such that the third pole is much higher than the lower two, and in the frequency range of interest, the response can be approximated as being due to a second-order filter. Thus, the step response is fitted to a function of the form,

$$f(t) = y_0 + A\, e^{-\zeta\omega_0 t}\, \cos(\omega_d\, t - \phi), \tag{B.1}$$

with $\omega_d = \omega_0\sqrt{1 - \zeta^2}$. The damping constant, ζ, and natural frequency, ω_0, determined from the fits are shown in table B.1. The time constant, $\tau = 1/(\zeta\omega_0)$, in the table is the time constant of the exponential envelope decay. The bandwidth shown, BW$= 1/(2\pi\tau)$ is the bandwidth of a single-pole, low-pass filter that would give the same time constant.

The mechanical bandwidth of the frequency-distance encoder is expected to be several kHz, and consequently, the VCO-PLL bandwidth needs to be a few

161

CP Scale	Filter 1				Filter 2			
	f_n kHz	ζ	τ μs	BW kHz	f_n kHz	ζ	τ μs	BW kHz
1	1.82	.22	400	0.40	13.5	0.16	74	2.2
2	2.53	.29	220	0.72	19.3	0.22	37	4.3
3	3.14	.35	140	1.1	24.3	0.27	24	6.6
4	3.46	.46	100	1.6	27.6	0.31	19	8.6
5	4.39	.47	76	2.1	31.6	0.33	15	10
6	4.94	.45	71	2.2	35.1	0.37	12	13
7	5.32	.47	63	2.5	38.0	0.41	10	16
8	5.52	.48	60	2.6	42.9	0.42	8.8	18

Table B.1: The parameters of the VCO-PLL loop determined by fitting equation B.1 to the measured step response curves.

time larger in order not to limit its bandwidth. However, increasing the VCO-PLL bandwidth too much is not optimal, because spurs generated within the bandwidth are not suppressed. To set the PLL bandwidth properly, first it should be set large enough so that it definitely does not limit the mechanical bandwidth, then the mechanical bandwidth should be measured, and finally the PLL bandwidth should be reduced to the smallest possible without affecting the mechanical bandwidth.

To this end, the filter component values were changed to the ones listed in figure B.1 under "Filter 2". The VCO step-response was again measured for all CP current scalings, also shown in table B.1. The time constant is plotted in figure B.3 against the current scaling for both filters. It is seen that the PLL bandwidth using filter 2 is about a factor of 6 faster than when using filter 1. The step response at the highest CP current scaling is shown in figure 9.18, and has a rise time of 8.8 μs, implying a bandwidth of 18 kHz. This is expected to be higher than the mechanical bandwidth.
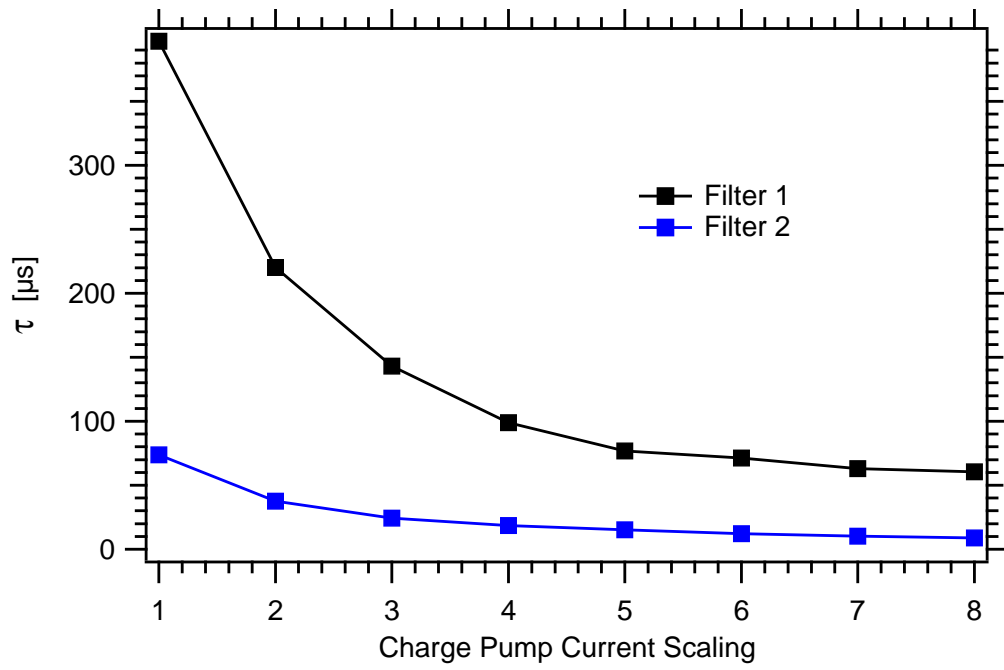
Figure B.3: Changing the charge-pump filter increased the bandwidth of the VCO PLL by about a factor of 6.

APPENDIX C

Best-Fit Linear Combination of Functions

After locking the cavity frequency and measuring the interferometer signal, piezo voltage, temperature, total pressure, and humidity over time, it is of interest to see if changes in the interferometer or piezo voltage curves can be written as a linear combination of changes in the others. More explicitly, if the interferometer or piezo voltage can be written as a function of P, T, H (pressure, temperature, humidity), it can be linearly expanded for small changes about a set of starting conditions P_0, T_0, H_0,

$$F(P,T,H)-F(P_0,T_0,H_0) \approx \left.\frac{\partial F}{\partial P}\right|_{T,H}(P-P_0)+\left.\frac{\partial F}{\partial T}\right|_{P,H}(T-T_0)+\left.\frac{\partial F}{\partial H}\right|_{T,P}(H-H_0).$$
(C.1)

Since $F(P,T,H) - F(P_0,T_0,H_0)$, $P - P_0$, $T - T_0$, and $H - H_0$ are known, it is possible to determine the partial derivatives.

Given a set of data, $D(t)$, that needs to be written as a linear combination of functions $g_i(t)$, how can a set of coefficients, a_i, be chosen to best approximate $D(t)$? The linear combination can be written as,

$$f(t) = a_1 g_1(t) + a_2 g_2(t) + ... = \sum_i a_i g_i(t).$$
(C.2)

The "error" in the fit between $D(t)$ and $f(t)$ can be defined as,

$$\epsilon = \int (D(t) - f(t))^2 dt.$$
(C.3)

164

The task is to minimize ϵ with respect to the set of a_i. Taking the partial derivatives with respect to a_i,

$$\frac{\partial \epsilon}{\partial a_i} = -2 \int \left(D(t) - \sum_j a_j g_j(t) \right) g_i(t) dt = -2 \left[D_i - \sum_j a_j G_{ij} \right], \qquad \text{(C.4)}$$

where,

$$D_j = \int D(t) g_j(t) dt \quad \text{and} \quad G_{ij} = \int g_i(t) g_j(t) dt. \qquad \text{(C.5)}$$

Setting equations C.4 equal to zero, we arrive at the matrix equation,

$$[G_{ij}][a_j] = [D_i], \qquad \text{(C.6)}$$

whose solution gives the best-fit coefficients.

References

[1] Richard K. Leach, Robert Boyd, Theresa Burke, Hans-Ulrich Danzebrink, Kai Dirscherl, Thorsten Dziomba, Mark Gee, Ludger Koenders, Valerie Morazzani, Allan Pidduck, Debdulal Roy, Wolfgang E. S. Unger, and Andrew Yacoot. The European nanometrology landscape. *Nanotechnology*, 22 (6), Feb 11 2011.

[2] Michael T. Postek and Kevin Lyons. Instrumentation, metrology, and standards: Key elements for the future of nano manufacturing. In *Instrumentation, Metrology, and Standards for Nanomanufacturing*, volume 6648, page 64802. SPIE, 2007.

[3] John A. Kramar, Ronald Dixson, and Ndubuisi G. Orji. Scanning probe microscope dimensional metrology at NIST. *Measurement Science & Technology*, 22(2), Feb 2011.

[4] G. Binnig, H. Rohrer, C. Gerber, and E. Weibel. 7x7 reconstruction on Si(111) resolved in real space. *Physical Review Letters*, 50(2):120–123, 1983.

[5] Neal B. Hubbard, Martin L. Culpepper, and Larry L. Howell. Actuators for micropositioners and nanopositioners. *Applied Mechanics Reviews*, 59(1-6): 324–334, 2006.

[6] Santosh Devasia, Evangelos Eleftheriou, and S. O. Reza Moheimani. A survey of control issues in nanopositioning. *IEEE Transactions on Control Systems Technology*, 15(5):802–823, Sep 2007.

[7] J. A. Kramar. Nanometre resolution metrology with the molecular measuring machine. *Measurement Science & Technology*, 16(11):2121–2128, Nov 2005.

[8] H.-U. Danzebrink, G. Dai, F. Pohlenz, T. Dziomba, S. Butefisch, J. Flugge, and H. Bosse. Dimensional nanometrology at PTB. In *2012 IEEE Inter-*

national *Instrumentation and Measurement Technology Conference (I2MTC 2012)*, page 4 pp. Instrum. Meas. Soc., 2012.

[9] Pongpun Rerkkumsup, Masato Aketagawa, Koji Takada, Tomonori Watanabe, and Shin Sadakata. Direct measurement instrument for lattice spacing on regular crystalline surfaces using a scanning tunneling microscope and laser interferometry. *Review of Scientific Instruments*, 74(3):1205–1210, 2003.

[10] S. Gonda, T. Doi, T. Kurosawa, Y. Tanimura, N. Hisata, T. Yamagishi, H. Fujimoto, and H. Yukawa. Real-time, interferometrically measuring atomic force microscope for direct calibration of standards. *Review of Scientific Instruments*, 70(8):3362–3368, Aug 1999.

[11] L. Howard, J. Stone, and J. Fu. Real-time displacement measurements with a Fabry-Perot cavity and a diode laser. *Precision Engineering-Journal of the International Societies for Precision Engineering and Nanotechnology*, 25(4): 321–335, Oct 2001.

[12] J. A. Stone and L. P. Howard. A simple technique for observing periodic nonlinearities in Michelson interferometers. *Precision Engineering-Journal of the American Society for Precision Engineering*, 22(4):220–232, Oct 1998.

[13] Andrew Yacoot and Ulrich Kuetgens. Sub-atomic dimensional metrology: developments in the control of x-ray interferometers. *Measurement Science and Technology*, 23:074003, 2012.

[14] Mehmet Çelik, Ramiz Hamid, lrich Kuetgens, and Andrew Yacoot. Picometre displacement measurements using a differential Fabry-Perot optical interferometer and an x-ray interferometer. *Measurement Science and Technology*, 23(8):085901, 2012.

[15] Mathieu Durand, John Lawall, and Yicheng Wang. High-accuracy Fabry-

Perot displacement interferometry using fiber lasers. *Measurement Science and Technology*, 22(9):094025, 2011.

[16] Youichi Bitou. High-accuracy displacement metrology and control using a dual Fabry-Perot cavity with an optical frequency comb generator. *Precision Engineering*, 33(2):187 – 193, 2009.

[17] Marco Pisani, Andrew Yacoot, Petr Balling, Nicola Bancone, Cengiz Birlikseven, Mehmet Celik, Jens Fluegge, Ramiz Hamid, Paul Koechert, Petr Kren, Ulrich Kuetgens, Antti Lassila, Gian Bartolo Picotto, Ersoy Sahin, Jeremias Seppa, Matthew Tedaldi, and Christoph Weichert. Comparison of the performance of the next generation of optical interferometers. *Metrologia*, 49(4):455–467, Aug 2012.

[18] Michael A. Lombardi. Selecting a primary frequency standard for a calibration laboratory. *Cal Lab International Journal of Metrology*, pages 33–39, 2008.

[19] C. Travis. Automatic frequency control. *Proceedings of the Institute of Radio Engineers*, 23(10):1125–1141, Oct 1935.

[20] Eugene F. Grant. An analysis of the sensing method of automatic frequency control for microwave oscillators. *Proceedings of the I.R.E.*, 37(8):943–951, 1949.

[21] G. G. Gerlach. A microwave relay communication system. *RCA Review*, 7 (4):576–600, 1946.

[22] R. V. Pound. Electronic frequency stabilization of microwave oscillators. *The Review of Scientific Instruments*, 17(11):490–505, 1946.

[23] W. G. Tuller, W. C. Galloway, and F. P. Zaffarano. Recent developments in

frequency stabilization of microwave oscillators. *Proceedings of the I.R.E*, 36: 794–800, 1948.

[24] S. R. Stein and J. P. Turneaure. Superconducting-cavity-stabilised oscillator of high stability. *Electronics Letters*, 8(13):321–323, 1972.

[25] S. R. Stein and J. P. Turneaure. The development of the superconducting cavity stabilized oscillator. In *27th Annual Frequency Control Symposium*, pages 414–420, 1973.

[26] A. G. Mann and D. G. Blair. Ultra-low phase noise superconducting-cavity stabilised microwave oscillator with application to gravitational radiation detection. *Journal of Physics D: Applied Physics*, 16:105–113, 1983.

[27] Bokuji Komiyama. Experimental results on a 9.2-GHz superconducting cavity stabilized oscillator. *IEEE Transactions on Instrumentation and Measurement*, IM-16(1):2–8, 1987.

[28] R. W. P. Drever, G. M. Ford, J. Hough, I. M. Kerr, A. J. Munley, et al. A Gravity Wave Detector Using Optical Cavity Sensing. In *Proceedings of the Ninth International Conference on General Relativity and Gravitation*, 1980.

[29] 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. *Applied Physics B*, 31:97–105, 1983.

[30] Eric D. Black. An introduction to Pound-Drever-Hall laser frequency stabilization. *American Journal of Physics*, 69:79–87, 2001.

[31] Jerome L. Altman. *Microwave Circuits*, page 205. D. Van Nostrand Company, Inc., 1964.

[32] J. E. Aitken. Swept-frequency microwave Q-factor measurement. *Proceedings of the Institution of Electrical Engineers-London*, 123(9):855–862, 1976.

[33] Gerd Hammer, Stefan Wuensch, Markus Roesch, Konstantin Ilin, Erich Crocoll, and Michael Siegel. Coupling of microwave resonators to feed lines. *IEEE Transactions on Applied Superconductivity*, 19(3):565–569, 2009.

[34] I. S. Gradshteyn and I. M. Ryzhik. *Table of Integrals Series and Products*, page 51. Academic Press, San Diego, California, 5th edition, 1965.

[35] I. S. Gradshteyn and I. M. Ryzhik. *Table of Integrals Series and Products*, page 987. Academic Press, San Diego, California, 5th edition, 1965.

[36] William F. Egan. *Phase-Lock Basics*, page 235. John Wiley & Sons, Inc., 2nd edition, 2008.

[37] Analog Devices. 2.7GHz DDS-Based *Agile*RF$^{\text{tm}}$ Synthesizer, 2004. AD9956 datasheet.

[38] W. P. Robins. *Phase Noise in Signal Sources*, pages 75–81. IEE Telecommunications Series 9. Peter Peregrinus Ltd., 1984.

[39] Agilent Technologies. Measuring noise and noise-like digital communications signals with spectrum and signal analyzers, 2009. Application Note 1303.

[40] David M. Pozar. *Microwave Engineering*, page 94. John Wiley & Sons, Inc., 3rd edition, 2005.

[41] David M. Pozar. *Microwave Engineering*, page 128. John Wiley & Sons, Inc., 3rd edition, 2005.

[42] K. V. Rajani, S. Rajesh, K. P. Murali, P. Mohanan, and R. Ratheesh. Preparation and microwave characterization of PTFE/PEEK blends. *Polymer Composites*, 30(3):296–300, 2009.

[43] W. M. Haynes, editor. *CRC Handbook of Chemistry and Physics*, pages 12–41. CRC Press, 93rd edition, 2012.

[44] Ernest K. Smith and Stanley Weintraub. The constants in the equation for atmospheric refractive index at radio frequencies. *Journal of the National Bureau of Standards*, 50(1), 1953.

[45] Gordon D. Thayer. An improved equation for the radio refractive index of air. *Radio Science*, 9(10), 1974.

[46] James C. Owens. Optical refractive index of air: Dependence on pressure, temperature and composition. *Applied Optics*, 6(1):51–59, 1967.

[47] Arnold Wexler. Vapor pressure formulation for water in range 0 to 100°C. a revision. *Journal of Research of the National Bureau of Standards*, 80A (5&6):775, 1976.

[48] David M. Pozar. *Microwave Engineering*, pages 300–303. John Wiley & Sons, Inc., 3rd edition, 2005.

[49] A. Berk. Variational principles for electromagnetic resonators and waveguides. *Antennas and Propagation, IRE Transactions on*, 4(2):104–111, 1956.

[50] Robert W. Boyd. Intuitive explanation of the phase anomaly of focused light beams. *J. Opt. Soc. Am.*, 70(7):877–880, Jul 1980.

[51] Joseph T. Verdeyen. *Laser Electronics*, page 71. Prentice-Hall, Inc., 2nd edition, 1981.

[52] Je Griffith and Da Grigg. Dimensional metrology with scanning probe microscopes. *Journal of Applied Physics*, 74(9):R83–R109, Nov 1 1993.

[53] Richard Leach. *Fundamental Principles of Engineering Nanometrology*. Elsevier Inc., 2010.

[54] Karoly Holczer, Emil Kirilov, John Koulakis, and Michael Stein. Actively controlled Fabry-Perot resonance cavities for laser stabilization, 2014. Submitted provisional patent.

[55] J. D. Jost, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward. Continuously tunable, precise, single frequency optical signal generators. *Optics Express*, 10:515–520, 2002.

[56] R. G. DeVoe and R. G. Brewer. Laser-frequency division and stabilization. *Physical Review A*, 30:2827–2829, 1984.

[57] Stephen P. Sandford and Charles W. Antill Jr. Laser frequency control using and optical resonator locked to an electronic oscillator. *IEEE Journal of Quantum Electronics*, 33:1991–1996, 1991.

[58] J. J. McFerran, S. T. Dawkins, P. L. Stanwix, M. E. Tobar, and A. N. Luiten. Optical frequency synthesis from cryogenic microwave sapphire oscillator. *Optics Express*, 14:4316–4327, 2006.

[59] C. McNeilage, J. H. Searls, E. N. Ivanov, P. R. Stockwell, D. M. Green, and M. Mossamaparast. A review of sapphire whispering gallery-mode oscillators including technical progress and future potential of the technology. In *Frequency Control Symposium and Exposition, 2004. Proceedings of the 2004 IEEE International*, pages 210–218, 2004.

[60] E. N. Ivanov, M. E. Tobar, and R. A. Woode. Applications of interferometric signal processing to phase-noise reduction in microwave oscillators. *IEEE Transactions on Microwave Theory and Techniques*, 46(10, 2):1537–1545, Oct 1998.

[61] J. Helszajn. *YIG Resonators and Filters*, page 219. John Wiley & Sons Ltd., 1985.

[62] William F. Egan. *Phase-Lock Basics*, page 59. John Wiley & Sons, Inc., 2nd edition, 2008.