

Cryogenic Silicon Optical Reference Cavities

Thesis by
David Yeaton-Massey

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy

The logo for the California Institute of Technology (Caltech), consisting of the word "Caltech" in a bold, orange, sans-serif font.

California Institute of Technology
Pasadena, California

2016
(Defended November 12, 2015)

© 2016

David Yeaton-Massey

All Rights Reserved

Acknowledgments

First off, thanks go out to the NSF for funding this work under award NSF.IQIM. Secondly, thanks go to my advisor, Rana Adhikari, without whom I would not have the standards or knowledge I have today.

Thanks to many others - those who were grad students with me: Rob Ward, for letting me both him with questions while he was sitting next to me writing his thesis in my 2nd year (I now realize how annoying that sometimes was). Tara Chalermansongsak, for many hours of whiteboard discussion as we compared our (very similar) experimental topologies. Jenne Driggers, for not kicking me in the shins when she was frustrated by my pestering. Zach Korth, for patience shown as we play invisible tug of war with shared experimental equipment. Eric Quintero, for putting together the Cymac1 test stand. Xiaoyue Ni, for accepting "I just need it for a few more days" as an answer so many times. Evan Hall, for teaching me not to trust cloven hoofed animals. Brian Dawes, an abnormally helpful and competent first year as I was finishing (you'll do well). I would like to thank all the other grad students who spent time in the trenches with me.

The often under-appreciated post docs - Frank Seifert, Nicolas Smith, and Chris Wipf, who all put in amounts of work helping make the experiment possible which I did not truly appreciate until the end of this particular leg of the journey. I must also thank Jamie Rollins and Larry Price for the snark filled interludes when I barged into their office wanting to discuss some piece of signal processing.

Koji Arai, for sinking so much time into making me a better scientist; I can't count the times when I would come to you with something which confused me, and you would spend time after we parted re-familiarizing yourself with it, find me the next day, and explain it to me in depth. Aidan Brooks, for being not just a model scientist, but a model human. Rich Abbott, for taking so many hours teaching me about the difference between what LTSpice says I built and what I actually built. Stefan Ballmer, for that first year in lab. Jan Harms, for being his (ridiculous) illuminating self. Gabriele Vajente, for letting me barge in and ask him questions whenever I found him. Matt Abernathy, for helping us understand what we don't know about coatings.

Eric Gustafson, for listening to me complain about interpersonal problems in between telling stories. Steve Vass, for his implacable coffee standards. Bob Taylor, for helping us find time to clean

and bake our vacuum parts in the rush of all the demands LIGO placed on your plate. Mike Gerfen, for helping us fix the cryostat. A couple hard working SURFs who worked under me: Jennifer Schloss and Matthew Arran.

All the members of the LIGO Hanford Observatory, for my time up there.

To all my friends who were down here with me, I don't know if you helped keep me sane or go a bit crazy, but you helped make me...me. To all my friends not down here, you guys are definitely normal and helped show me that I was, in fact, crazy. To Io Kleiser, the adoption papers are in the mail. To everyone at RCJ Machado, you definitely helped keep me sane(r). To Jen Aubrecht, I owe you one. To my family, I don't even know where to begin; I wouldn't have even gotten to the foot of the mountain without you.

I offer a special thank you to Huan Tran, the postdoc at Berkeley who led me to this path by showing me what research is actually like. You left this world too early, but I am certain I am not the only one still carrying your memory.

Abstract

Thermodynamical fluctuations in temperature and position exist in every physical system, and show up as a fundamental noise limit whenever we choose to measure some quantity in a laboratory environment. Thermodynamical fluctuations in the position of the atoms in the dielectric coatings on the mirrors for optical cavities at the forefront of precision metrology (e.g., LIGO, the cavities which probe atomic transitions to define the second) are a current limiting noise source for these experiments, and anything which involves locking a laser to an optical cavity. These thermodynamic noise sources scale physical geometry of experiment, material properties (such as mechanical loss in our dielectric coatings), and temperature. The temperature scaling provides a natural motivation to move to lower temperatures, with a potential huge benefit for redesigning a room temperature experiment which is limited by thermal noise for cryogenic operation.

We design, build, and characterize a pair of linear Fabry-Perot cavities to explore limitations to ultra low noise laser stabilization experiments at cryogenic temperatures. We use silicon as the primary material for the cavity and mirrors, due to a zero crossing in its linear coefficient of thermal expansion (CTE) at 123 K, and other desirable material properties. We use silica tantala coatings, which are currently the best for making high finesse low noise cavities at room temperature. The material properties of these coating materials (which set the thermal noise levels) are relatively unknown at cryogenic temperatures, which motivates us to study them at these temperatures. We were not able to measure any thermal noise source with our experiment due to excess noise. In this work we analyze the design and performance of the cavities, and recommend a design shift from mid length cavities to short cavities in order to facilitate a direct measurement of cryogenic coating noise.

In addition, we measure the cavities (frequency dependent) photo-thermal response. This can help characterize thermo-optic noise in the coatings, which is poorly understood at cryogenic temperatures. We also explore the feasibility of using the cavity to do macroscopic quantum opto-mechanics such as ground state cooling.

Contents

Acknowledgments	iii
Abstract	v
1 Ultra-Stable Fabry-Perot Cavities	1
1.1 Intellectual Merit	1
1.2 Brownian Thermal Noise	2
1.3 Thermal Noise in Fabry-Perot Cavities	2
2 Silicon Reference Cavities	6
2.1 Mechanical Design	9
2.2 Cavity Assembly	11
2.3 Cavity Properties	13
3 The Experiment	17
3.1 The Supports	17
3.2 The Suspension	19
3.3 The Cryostat	20
3.4 Thermal Shielding and Control	21
3.5 Optical Layout	21
3.6 Lessons learned	29
4 Noise Budget	31
4.1 Thermal Noise	31
4.1.1 The Fluctuation Dissipation Theorem	31
4.1.2 Brownian Noise	33
4.1.3 Thermo-elastic Noise	34
4.1.4 Thermo-refractive Noise	35
4.1.5 Thermo-Optic Noise	36
4.1.6 Silicon Carrier Density Noise	36

4.2	Scatter	37
4.3	Seismic	38
4.4	Intensity	42
4.5	Temperature	48
4.6	Electronics	48
4.7	Quantum Fluctuations	58
4.8	RFAM	58
4.9	Photothermal Response	60
5	Towards Macroscopic Quantum Cavity Opto-Mechanics	67
6	Frequency Domain Subtraction	81
6.1	Wiener Filtering	81
6.2	Coherent Subtraction	82
7	Doubling Noise	88
7.1	Experimental Setup	89
7.2	Results	90
7.3	Noise Sources	92
7.4	Previous Bounds on Excess Frequency Noise	92
7.5	Conclusions	92
8	Conclusions and Future Work	94
A	Resonator Equations	96
B	Suspension Redesign	98
C	Window Scatter	103
D	Cavity Seismic Susceptibility	107
E	Electronics	109
F	Photothermal Transfer Functions	112
	Bibliography	116

List of Figures

1.1	Advanced LIGO Noise Budget	3
1.2	LIGO BLUE Noise Budget	4
2.1	Silicon Linear CTE vs Temperature	7
2.2	Silicon Q vs Temperature for Crystal Axes (100) and (101)	7
2.3	Silicon Thermal Conductivity vs Temperature	8
2.4	Silicon Heat Capacity vs Temperature	8
2.5	Cavity Geometry	10
2.6	Cavity Assembly Process	12
2.7	Higher Order Mode Spacing	14
2.8	Cavity Pole Measurements	15
3.1	Cavity Support Options	18
3.2	Cavity Support Structure	19
3.3	Cryogenic Suspension Platform	20
3.4	Cryostat	22
3.5	Radiation Shield Temperature Sensor	23
3.6	Cavity Shielding	23
3.7	Simplified Experimental Layout	26
3.8	Pictures of Experimental Setup	27
3.9	Pictures of Experimental Setup - Zoomed	28
3.10	PLL Schematic	29
4.1	Experimental Noise Budget	32
4.2	Evolution of Experimental Noise	32
4.3	Seismic Coupling Chain	40
4.4	Vertical Motion of Cryostat	41
4.5	Cavity Sensitivity to Vertical Motion	42
4.6	Cavity Noise from Vertical Motion	43
4.7	RIN Coupling	44

4.8	Frequency Noise to RIN Coupling from Net PDH Offset	45
4.9	Coherence of RIN and PDH Error Signals	46
4.10	Incident and Transmitted RIN	47
4.11	Temperature Noise	49
4.12	Electronics Noise Budget	50
4.13	PDH Loop Diagram	52
4.14	PDH Open Loop Transfer Function	53
4.15	PDH Loop Phase Budget	54
4.16	Laser Frequency Actuator Transfer Function	55
4.17	Beat Readout (PLL) Calibration	57
4.18	Quantum Vacuum Noisebudget	58
4.19	Frequency Noise due to RFAM	60
4.20	Temperature Dependent Model of Photothermal Transfer Function: Coherent Sum of Terms	61
4.21	Temperature Dependent Model of Photothermal Transfer Function: Substrate Expansion	62
4.22	300 K Model of Photothermal Transfer Function: All Terms	62
4.23	Diagram of Photothermal Measurement	64
4.24	Photothermal Transfer Function: Two Parameter Fit of Model to Measurement	65
4.25	Optical Cavity Absorption vs Temperature	66
5.1	Cavity Optomechanics Field Summary	68
5.2	Picture of Electrostatic Driver	69
5.3	Cavity Drum Mode Ring-Down Measurement Diagram	69
5.4	Cavity Drum Mode Ring-Down Measurements	70
5.5	Electro and Optomechanical Transfer Functions	71
5.6	PEEK Cavity Support COMSOL Simulation	72
5.7	PEEK Loss vs Temperature	73
5.8	New Cavity Support	73
5.9	Proposed Cavity Support Eigenmodes	76
5.10	Proposed Support Eigenmodes for Various Geometries	77
5.11	Strain in New Support	78
5.12	fQ Product vs Temperature	78
5.13	$f \times Q$ for Various Optomechanical Experiments	79
5.14	COMSOL Simulation of Short Cavity Eigenfrequency	80
6.1	Coherence Between Beat and Auxiliary Channels	83
6.2	Coherence Between Auxiliary Channels	84

6.3	Coherent Subtraction of PDH Residual from Beat Noise	85
7.1	Experimental Layout of Mach-Zehnder to Measure Doubling Noise	89
7.2	Time Series of Mach-Zehnder Output	90
7.3	Comparison of Experimental Results with Previous Bounds as Frequency Noise	91
7.4	Comparison of Experimental Results with Previous Bounds as Allan Deviation	93
B.1	Picture of New Cavity Suspension	98
B.2	Picture of Old Cavity Suspension	99
B.3	Spectra of New Suspension Platform Motion	100
B.4	Ring-Down Measurements of Platform Suspension Modes	101
B.5	Suspension Upgrade Noise Improvement	102
C.1	Cryostat Window Contamination	103
C.2	BRDF Measurement Setup	104
C.3	BRDF Measurement Images	105
D.1	Vertical Seismic Sensitivity vs Geometry	108
E.1	Optical Plant	109
E.2	PDH RF Photodetection and Demodulation	110
E.3	PDH Servo Box	110
E.4	Laser Diode Driver	111
F.1	Photothermal Fit 233 K	112
F.2	Photothermal Fit 255 K	113
F.3	Photothermal Fit 283 K	114
F.4	Photothermal Fit 290 K	115

List of Tables

2.1	Thermal and mechanical properties of potential cavity materials. α is the linear coefficient of thermal expansion, κ is the thermal conductivity, Q_{mech} is the mechanical quality factor, C is the heat capacity, and Y is the Young's modulus.	6
2.2	Cavity Optical Properties	13
5.1	Silicon Cavity Body Mode Quality Factor at 300 K and 123 K	68
5.2	References for legend labels in fig. 5.13.	80
C.1	Cryostat Window R T and BRDF	105

Chapter 1

Ultra-Stable Fabry-Perot Cavities

1.1 Intellectual Merit

Thermodynamical fluctuations in position and temperature, which occur in all physical systems at a finite temperature, are fundamentally limiting in several experimental areas at the forefront of precision metrology. The list of these fields includes torsion pendulums for measuring the equivalence principle and the gravitational inverse-square law, investigations into quantum behavior of macroscopic mechanical oscillators, cryogenic sapphire oscillators used in radio astronomy for pulsar timing, the suspensions and mirrors of interferometric gravitational wave detectors, and reference cavities used in atomic clocks and spectroscopy.

In the general landscape of frequency stability, suspended long baseline interferometers such as LIGO [1] are the best frequency standard at high frequencies, and are used with the purpose of detecting gravitational waves as small audio frequency perturbations to space-time. Atomic clocks such as the NIST F-1 Cs fountain clock [2], the NIST trapped Al⁺ single ion clock [3], and the JILA ⁸⁷Sr neutral lattice clock [4] are currently the best performers at low frequency. While the NIST F-1 clock is currently used as the international standard for the second, the Al⁺ and ⁸⁷Sr clocks have both been shown to outperform it in terms of long term fractional stability. Lasers locked to rigid Fabry-Perot cavities [5] are currently the best reference at intermediate frequencies, as such they are used to compare optical clocks [4] in a low noise way, and as the interrogating laser for the LIGO experiment [6]. All the above mentioned experiments are fundamentally limited by thermal noise, and benefit from a greater understanding of this noise. Improvements to Ultra-Stable Fabry-Perot Cavities would allow improved frequency stability of the lasers used to interrogate the optical clocks used to define the second, and opens the gate for improved clock stability by orders of magnitude, as shown in [4].

Additionally, the field of cavity optomechanics is pushing towards observing quantum behavior with increasingly large mechanical oscillators. Brownian thermal fluctuations are a classical fundamental noise source which will limit observing quantum behavior. It will be interesting to explore

using Ultra-Stable Fabry-Perot cavities to observe quantum behavior, as much engineering work has been done to minimize their thermal noise. Thermodynamical noise sources are fundamentally limiting many experiments at the forefront of precision metrology, and are thus of interest to a large scientific community.

1.2 Brownian Thermal Noise

Here we introduce what thermal noise is, and motivate the move towards cryogenic temperature to reduce this noise. Any mechanical system at a finite temperature is moving constantly, as temperature is the thermodynamic variable we use to describe how much the molecules of a system are moving. The fluctuation dissipation theorem relates the macroscopic motion of the system due to statistical fluctuations and the energy dissipated due to mechanical motion. It can be expressed as:

$$S_x(f) = \frac{k_B T}{\pi^2 f^2} \text{Re}\left[\frac{1}{Z(f)}\right], \quad (1.1)$$

where S_x is the single sided power spectral density of fluctuations in displacement, and $Z(f)$ is the complex impedance of the system. Away from any mechanical resonances in a system, $Z(f)$ scales as the mechanical quality factor, or the inverse mechanical loss $Z(f) \propto Q \propto \phi^{-1}$; lossy materials have higher thermal noise. For a given geometry, the only ways to improve the thermal noise of a system are to either decrease the mechanical loss angle, or the temperature of the system. We join in the relatively new effort to explore thermal noise in optical cavities at cryogenic temperatures [7, 8]. A direct measurement of the temperature scaling shown in eq. (1.1) for any thermal noise source in optical cavities would be of interest to the precision measurement communities mentioned above. See section 4.1 for a more in depth discussion of individual thermal noise sources, and equations which describe their coupling in this work.

1.3 Thermal Noise in Fabry-Perot Cavities

One of the most prevalent thermal noise sources in experimental topologies which lock a laser to a cavity is the thermal motion in the dielectric coating on the cavity mirrors. These coatings are generally made up of a stack of 1/4 wavelength thickness materials which have alternating (high and low) indices of refraction, low optical loss, and low mechanical losses. They are generally tens of layers thick. The fact that the optical coatings are volumetrically insignificant is offset by their anomalously high mechanical loss [9], which leads to them being a limiting noise source in the laser stabilization experiment discussed above.

To give a specific example: In the field of gravitational wave astronomy, after the first unambiguous detection, which is expected in the next few years, there will be a push towards another

generation of detectors with sensitivity beyond what is possible with current technology. In the Advanced LIGO detector, coating thermal noise is expected to be the largest of the thermal noise sources (see fig. 1.1) at their most sensitive detection frequencies. The first generation of gravitational wave detectors used alternating layers of SiO_2 ($n = 1.45$) and Ta_2O_5 ($n = 2.06$), while the current state of the art dopes these same coatings with titania (TiO_2) [10]. (This work uses non-titania doped $\text{SiO}_2:\text{Ta}_2\text{O}_5$ coatings).

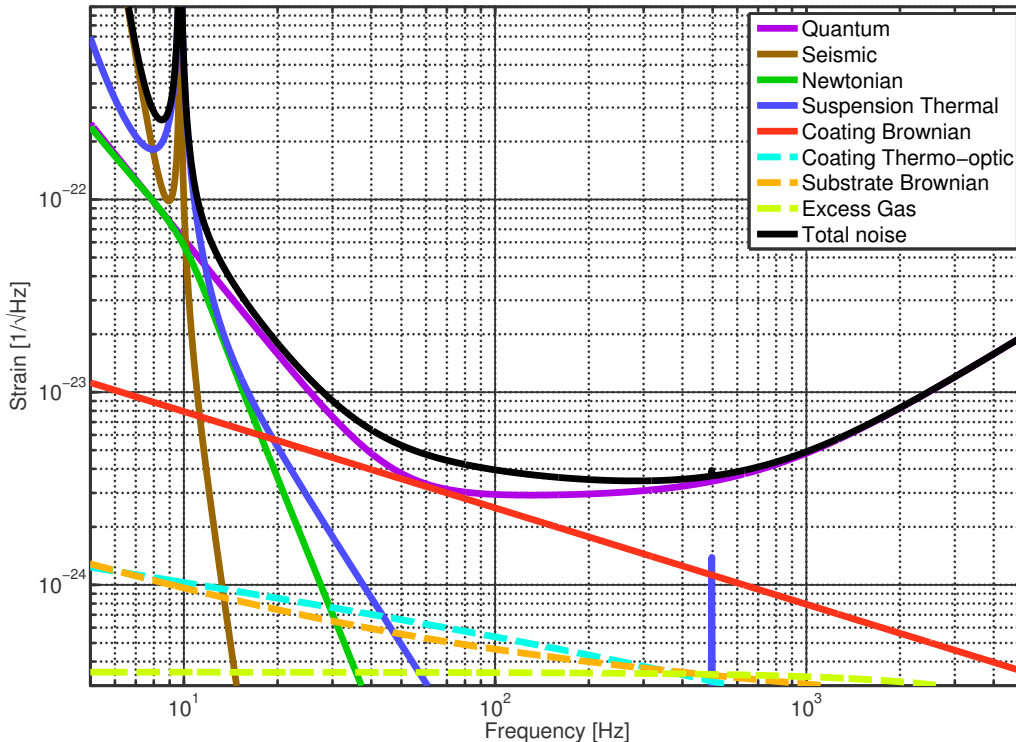


Figure 1.1: Theoretical strain sensitivity of the Advanced LIGO detector. A strain of $1 \times 10^{-23}/\sqrt{\text{Hz}}$ over the 4 km detector baseline translates to a mirror motion of $4 \times 10^{-20} \text{ m}/\sqrt{\text{Hz}}$, or laser frequency noise of $3 \times 10^{-9} \text{ Hz}/\sqrt{\text{Hz}}$. This plot was produced using the GWINC software package, maintained by members of the LIGO scientific collaboration.

While $\text{SiO}_2:\text{Ta}_2\text{O}_5$ coatings are being studied at cryogenic temperatures (both in this work and by others [11]) in order to understand their loss, the push towards using crystalline silicon optics (for reasons discussed in chapter 2) at cryogenic temperatures may allow for the use of lattice matched crystalline thin films for optical coatings. A future upgrade to the LIGO detectors dubbed "LIGO VOYAGER/BLUE" [8] involves switching to Silicon optics with crystalline coatings, with a 2000 nm light source, operating at 123 K. The noise budget for such a detector is shown in fig. 1.2. The (much improved) theoretical coating thermal noise shown in the noise budget is still limiting.

Once the process for generating low loss crystalline coatings is developed, crystalline coatings will likely be the best option to lower coating brownian noise in linear cavities from the cm scale (See

[12] the use of crystalline AlGaAs coatings with SiO₂ mirrors and cavities) to the km scale (LIGO). As the technique for depositing crystalline coatings was not yet mature enough, we decided to use the state of the art SiO₂:Ta₂O₅ coatings deposited on Si substrates in an attempt to measure the \sqrt{T} (when expressed as an amplitude spectral density, as in fig. 1.1 and fig. 1.2) scaling of coating thermal noise seen in eq. (1.1).




Figure 1.2: Amplitude spectral density of the proposed LIGO VOYAGER/BLEU upgrade using silicon optics at 123 K. Assuming the same 4 km baseline as Advanced LIGO and a 2000 nm laser, 1×10^{-23} in strain is 4×10^{-20} m/ $\sqrt{\text{Hz}}$ in displacement, and 2×10^{-9} Hz/ $\sqrt{\text{Hz}}$. The factor of four improvement seen in this silicon concept design over Advanced LIGO at 100 Hz is mostly due to lowering the coating thermal noise. A combination of moving to lower temperatures, lowering the coating loss by moving to a-Si (amorphous silicon) coatings, and increasing the beam size at the mirror contribute to this lower noise. This plot was produced using the GWINC software package, maintained by members of the LIGO scientific collaboration. See chapter 2 for a discussion of cryogenic silicon properties, and chapter 4 for discussion of coating noise.

As the field of gravitational wave astronomy and precision metrology move towards using silicon optics at cryogenic temperatures, coatings with low thermodynamical noise become of increasing interest. As the coating technologies which will likely be used at cryogenic temperatures in the future (crystalline coatings, amorphous silicon coatings) were not mature enough when we were designing the experiment, we chose to use SiO₂:Ta₂O₅ (silica tantala) coatings. These are the most prevalent coating at room temperature (it is the lowest optical loss coating currently available, and

thus the best choice for pushing the limits of high finesse cavities), so understanding their properties at cryogenic temperatures is still of merit.

Chapter 2

Silicon Reference Cavities

We saw in eq. (1.1) that going to cryogenic temperatures lets us decrease the thermal noise in an experiment. The thermal noise is also dependent (implicitly, in $Z(f)$) on the material properties of the optical components we choose to use in the experiment. Thus the choice of materials is driven by their contribution to the various thermal noise sources in the experiment, given in section 4.1, which we summarize here. To decrease thermal noise in the spacer and substrate, we want to use materials with higher Young's modulus, a lower mechanical loss angle (aka a higher quality factor), a lower thermal conductivity, a lower heat capacity, and a lower coefficient of thermal expansion. In the case of the Young's modulus, the spacer brownian noise scales as Y^{-1} , whereas the spacer thermoelastic noise scales as Y^2 . We also benefit from stiffer (higher Y) cavities in decreased susceptibility to seismic noise. The spacer thermoelastic noise is quite small in the parameter space of geometries we considered, so we do not decrease the cavity performance by choosing a high Y . Additionally, we want to explore the viability of using the cavities for opto-mechanics experiments which require a high Q , so we want this number to be as large as possible. The relevant material properties for silicon and sapphire (two leading candidates for optical material to make the mirrors and cavity spacer from) are shown in table 2.1. In addition to these material properties, we can only consider materials which are transmissive (very low optical loss) at some wavelength which is accessible to us (low noise lasers at this wavelength exist).

	α [K ⁻¹]	κ [W m ⁻¹ K ⁻¹]	Q_{mech}	C [J g ⁻¹]	Y [GPa]
Si(300 K)	2×10^{-6} [13]	1.5×10^2 [14]	7×10^7 [15]	6×10^{-1} [16]	130[17]
Si(123 K)	$2.2 \times 10^{-8}(T - 123)$ [13]	6×10^2 [14]	1.2×10^8 [15]	3.3×10^{-1} [16]	-
Si(18 K)	$-2 \times 10^{-9}(T - 18)$ [13]	4×10^3 [14]	1.4×10^8 [15]	3×10^{-3} [16]	-
Si(5 K)	7×10^{-11} [13]	5×10^2 [14]	3×10^8 [15]	3×10^{-5} [16]	-
Saph(30 K)	3×10^{-8} [18]	7.5×10^3 [19]	5×10^7 [20]	-	-

Table 2.1: Thermal and mechanical properties of potential cavity materials. α is the linear coefficient of thermal expansion, κ is the thermal conductivity, Q_{mech} is the mechanical quality factor, C is the heat capacity, and Y is the Young's modulus.




Figure 2.1: Silicon Linear CTE vs Temperature. This work was performed about the 123 K zero crossing, though there is another zero crossing at 18 K. The linear approximation of $\alpha_{Si}(T)$ in the neighborhood of these zeros are shown in eq. (2.1) and eq. (2.2), respectively. Data reproduced from [13].




Figure 2.2: Silicon quality factor at cryogenic temperatures for crystal axes (100) and (101). Plot reproduced from [15]. Actual Q varies with sample doping and polishing. Sensitivity to polishing scales as the surface to volume ratio.




Figure 2.3: Thermal conductivity as a function of temperature for a p-type silicon crystal grown in the (111) direction. Room temperature resistivity was roughly $2000 \Omega\text{-cm}$, which corresponds to a Boron dopant level of $6.6 \times 10^{12} \text{cm}^{-3}$. Data reproduced from [14]. Below 100 K, doping and orientation can cause order of magnitude variations in the conductivity. The variation in thermal conductivity from 100-300 K due to these variables is relatively stable.




Figure 2.4: Silicon heat capacity at cryogenic temperatures. Plot reproduced from [16].

We considered making the cavity spacer from silicon, sapphire, Zerodur, ULE, niobium and SiO₂ based on what we knew could be machined and believed could have a high Q. We ruled out ULE and Zerodur based on their very high mechanical loss. At room temperature, SiO₂ has mechanical loss $10^3 \times$ lower than ULE, though this degrades at cryogenic temperature, and the linear coefficient of thermal expansion (CTE) is not that low. Niobium was written off due to difficulty in machining reported by the cryogenic bar detector community (private communication with W. Johnson [21]).

Sapphire and silicon both have material properties we desire in a cavity spacer (high speed of sound, low CTE). The semiconductor industry has invested heavily in researching how to produce large, high purity silicon samples. This has led to the development and availability of inexpensive sources of silicon with ultra-low optical absorption (the lowest measured absorption of 1550 nm and 2000 nm light in silicon is 2.5 ppm/cm [22]). Silicon has several zero crossings in $\alpha_{\text{Si}}(T)$ [13]. We can expand $\alpha_{\text{Si}}(T)$ about these zero crossings, which yields:

$$\alpha_{\text{Si}}(T)|_{T=123\text{K}} \approx 2.2 \times 10^{-8}(T - 123\text{K})\text{K}^{-2}, \quad (2.1)$$

and

$$\alpha_{\text{Si}}(T)|_{T=18\text{K}} \approx -2 \times 10^{-9}(T - 18\text{K})\text{K}^{-2}. \quad (2.2)$$

Operating at these temperatures allows us to eliminate the first order coupling of ambient temperature fluctuations, as well as the thermal noise sources which scale as α . This means that we can use silicon at easily attainable temperatures (such as 123 K), whereas sapphire doesn't become a good option until very low temperatures.

A drawback of silicon comes in its already high thermal conductivity, which increases at 123 K (see fig. 2.3). In practice, this just sets an engineering requirement on how tightly we must hold the temperature about the zero crossing in order that silicon is still superior to sapphire in terms of thermal noise in the spacer.

Silicon (fig. 2.1, fig. 2.2, fig. 2.3, and fig. 2.4) makes an excellent choice for our cavity mirror substrates and spacer fig. 2.5. It has a high mechanical quality factor, a zero crossing in $\alpha_{\text{Si}}(T)$, and acceptable heat capacity and thermal conductivity (e.g., the thermal noise from the silicon detailed in section 4.1 is lower than the projected coating thermal noise). In addition, it is transmissive (2.5 ppm/cm loss) at 1550 nm, which is the wavelength of choice for the telecommunications industry and thus gives us many choices for low noise laser sources.

2.1 Mechanical Design

The mechanical dimensions of the cavity influences both the thermal noise levels, and some other practical engineering considerations. The main dimensions of interest are: **spacer length**, **spacer**

diameter, bore hole radius, mirror radius of curvature, mirror thickness, and optical coating.




Figure 2.5: The cavity is constructed by attaching two coated silicon mirrors to either end of a cylindrical silicon spacer with a bore hole through the center such that there is no material between the high reflective coatings of the inner mirror surfaces.

The *cavity pole* is one of the main design considerations we make when choosing **spacer length**, and **optical coating** (e.g., the mirror transmission and optical losses in the cavity). The cavity pole is set by the finesse, which, for a symmetrical cavity where the mirror transmission is much greater than the optical losses, is:

$$\mathcal{F} = \frac{2\pi}{\mathcal{L}_{tot}} = \frac{f_{FSR}}{2f_{cav}}, \quad (2.3)$$

where \mathcal{F} is the cavity finesse, \mathcal{L}_{tot} is the total roundtrip loss (including mirror transmission), $f_{FSR} = c/2L$ is the free spectral range of the cavity, f_{cav} is the cavity pole, L the cavity length, and c is the speed of light. Higher length and lower mirror transmission gives us a lower cavity pole, which gives us a more stable cavity. When we lock a laser to the cavity, it becomes very difficult to suppress the laser frequency noise at high Fourier frequencies (see section 4.6). As the cavity passively filters out laser noise above the cavity pole, there lies the benefit of working with lower cavity poles. We chose a cavity pole of $\approx 40kHz$ (fig. 2.8) in this initial phase of the experiment, largely due to the uncertainty in cryogenic properties of the $\text{SiO}_2:\text{Ta}_2\text{O}_5$ optical coatings we used; this was low enough to start filtering the frequency noise where we begin to worry about the bandwidth of our servos, but not so low that unexpected shifts in coating parameters would create an un-lockable cavity.

The **radius of mirror curvature** and **spacer length** directly limits the *spot size of the Gaussian beam* on the mirror, which influences how large the coupling of *thermal noise* at the mirror surface into the experiment is. Since what we actually detect is changes in the average phase of the optical field we use to probe the mirror surfaces, the coupling of the thermal noise into this phase decreases as we sample a larger area. An experiment designed to detect this thermal noise

should have a very small spot size, and this goal has led to the design of short optical cavities such as [23]. The initial goal of this work was to make a very low noise cavity, and thus we chose a medium sized spot of $330\ \mu\text{m}$ (see table 2.2 for measured numbers).

The **spacer diameter** and **spacer length** also determine the *seismic sensitivity of the cavity*. In general, we want a stiffer cavity (shorter and thicker) in order for it to be less susceptible to seismic motion. The choice of how we hold the cavity (see section 3.1) ends up being the dominant factor in determining the seismic susceptibility, but we still need to choose a reasonable aspect ratio for the length:diameter, and not have the bore hole where the beam propagates be so large as to compromise the structural integrity of the cavity. We chose an aspect ratio very close to 2:1, and a 4:1 ratio for spacer diameter to bore hole diameter.

Thermo-optic noise in the mirror substrate scales with **mirror thickness**, but this is small compared to coating thermal noise, and we are free to choose this parameter around practical considerations.

The cavity optical properties can be seen in table 2.2.

2.2 Cavity Assembly

The silicon boules used to make the spacers were single crystal boron doped (100) axis with 2" diameter and 4" length, and were grown using the Cz process. Boules were provided by Virginia Semiconductor¹. The boules had a specified resistivity range of 1 to 100 Ohm-cm, which corresponds to² doping levels of 1.4×10^{16} to $1.3 \times 10^{14}\ \text{cm}^{-3}$. The boules were machined down to 1.9" diameter by 3.9" length cylinders, with a 1/2" bore hole through the middle, and a 1/8" vent hole for vacuum considerations. The faces were polished to < 0.005 " parallelism over their length. The outer surface of the spacer was polished; the bore hole was not. Machining was done by Light Machinery³.

The substrates were purchased from Coastline⁴ and had 1/4" thickness, 1 m curvature, and 1" diameter. AR and HR coatings were done by ATFilms⁵.

Portions of the cavity assembly process are shown in fig. 2.6. It is common for these cavities to be assembled via optical contacting [24], and a glue joint is assumed to be lossier than techniques like optical contacting and hydroxide bonding (see [25] for a comparison). While it seems very reasonable to assume that the bulk loss of glue is larger than the bulk loss of the material in the neighborhood of an optical contact or hydroxide joint (as is the case with TRA-DUCT 2902 epoxy [26]), whether or not you are limited by this loss is dependent on the specific geometry of the joint. We decided to fabricate the first generation of cavities using adhesive, and a carefully developed

¹Virginia Semiconductor — <https://www.virginiasemi.com>

²<http://www.cleanroom.byu.edu/ResistivityCal.phtml>

³Light Machinery — <https://lightmachinery.com/>

⁴Coastline Optics — <http://www.coastlineoptics.com>

⁵Advanced Thin Films — <http://advancedthinfilms.com>




Figure 2.6: Silicon Cavity Assembly. From top left to bottom right: Machined and polished silicon boule, silicon blank on wafer demonstrating gluing technique, jig used for cavity + mirror assembly, and assembled silicon spacer with silicon mirror glued to front face.

technique (shown in part in fig. 2.6) in which we restrict the adhesive to a ring around the outside of the mirror. So long as there was no glue between the substrate and spacer, and the total volume of glue used was very low, the thermal noise contribution from the glue should be negligible. Based on what we know, however, future cavities should still be constructed with a technique which is less likely to be lossy (such as optical contacting).

2.3 Cavity Properties


	W(300 K)	E(300 K)	W(123 K)	E(123 K)	
f_1	4.79 ± 0.13	4.58 ± 0.17	4.83 ± 0.11	4.41 ± 0.08	MHz
f_{hom}	218 ± 5.7	212 ± 3.4	218 ± 7.4	210 ± 4.0	MHz
f_{cav}	41.3 ± 0.1	37.3 ± 0.1	-	-	kHz
f_{FSR}	1.521 ± 0.057	1.479 ± 0.060	1.521 ± 0.062	1.46 ± 0.038	GHz
L_{cav}	9.86 ± 0.37	10.14 ± 0.41	9.86 ± 0.40	10.23 ± 0.27	cm
g_{cav}	0.900 ± 0.010	0.900 ± 0.010	0.900 ± 0.012	0.900 ± 0.007	
R_{mirror}	0.990 ± 0.04	1.018 ± 0.04	0.990 ± 0.04	1.027 ± 0.03	m
w_0	326 ± 10	331 ± 10	326 ± 11	332 ± 7	μm
\mathcal{F}	18420 ± 45	19830 ± 53	18420 ± 45	19650 ± 53	
Loss	171 ± 6	158 ± 6	171 ± 7	160 ± 4	ppm
Abs	24 ± 1	-	-	-	ppm

Table 2.2: Measured optical properties of the silicon cavities. No quantity changed by more than the measurement error during the cooldown. The largest expected change in any property is 1×10^{-4} .

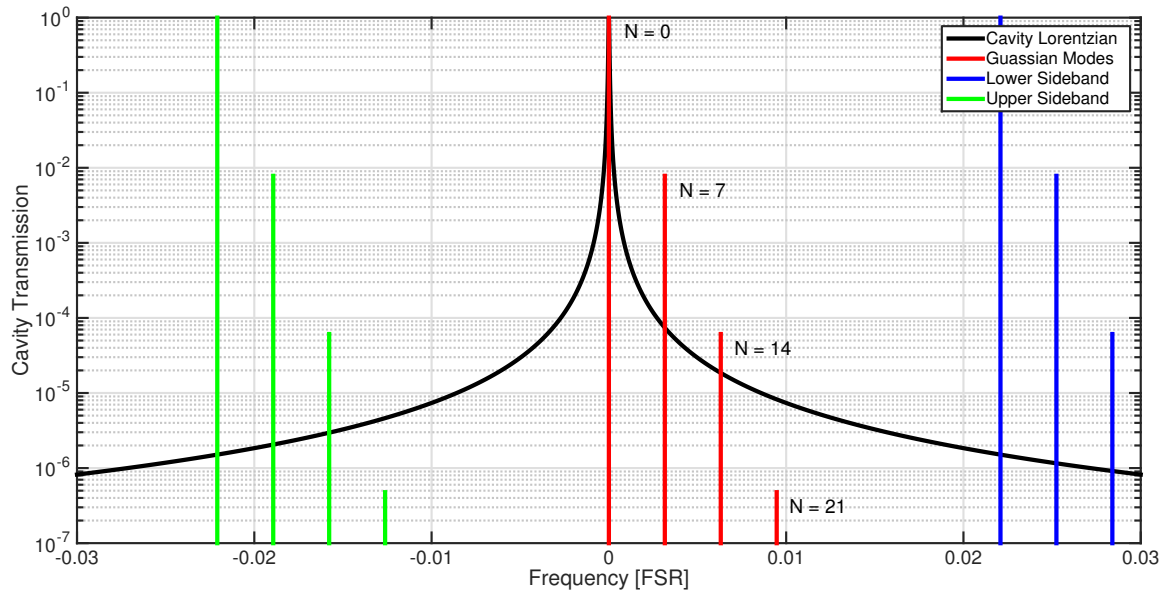
The optical quantities which we measured are: The free spectral range of the cavity (f_{FSR}), the cavity pole (f_{cav} , shown in fig. 2.8), the higher order mode spacing (f_{hom} , shown in fig. 2.7), and the frequency splitting of the higher order modes (ν_{split}). The measured cavity parameters, as well as those given by eq. (A.1) through eq. (A.6), are shown in table 2.2.

We measured the f_{FSR} by sweeping the frequency of the interrogating laser and looking at the transmitted power. The measurement was made on an oscilloscope and calibrated using the frequency spacing of phase modulation sidebands on the light. By measuring distance between the zeroth and seventh order spatial modes, $f_1 = 7 \times f_{\text{hom}} - f_{\text{FSR}}$, and using our measurement of f_{FSR} , we were able to solve for f_{hom} . The measurement of f_1 was also made on a scope, and calibrated against the phase modulation sidebands. The frequency splitting of higher order spatial modes was measured by comparing the frequency splitting of the TEM(01) and TEM(10) modes with the (now known) higher order mode spacing.

We measured the cavity pole (f_{cav}) by taking an optical transfer function of the cavity's response to intensity fluctuations out to high frequency, and fit a simple pole to the results (shown in fig. 2.8). The intensity modulation was done with an acousto-optic modulator (AOM). It is common to measure the pole using a simple cavity ring down technique, but a proper swept sine transfer function



(a)



(b)

Figure 2.7: Gaussian mode spacing based on the measured cavity parameters table 2.2. The black trace is a simple Lorentzian with a HWHM of f_{cav} . The increasing order of the gaussian modes is indicated by decreasing height, and the mode number is indicated on the right hand side of (a).

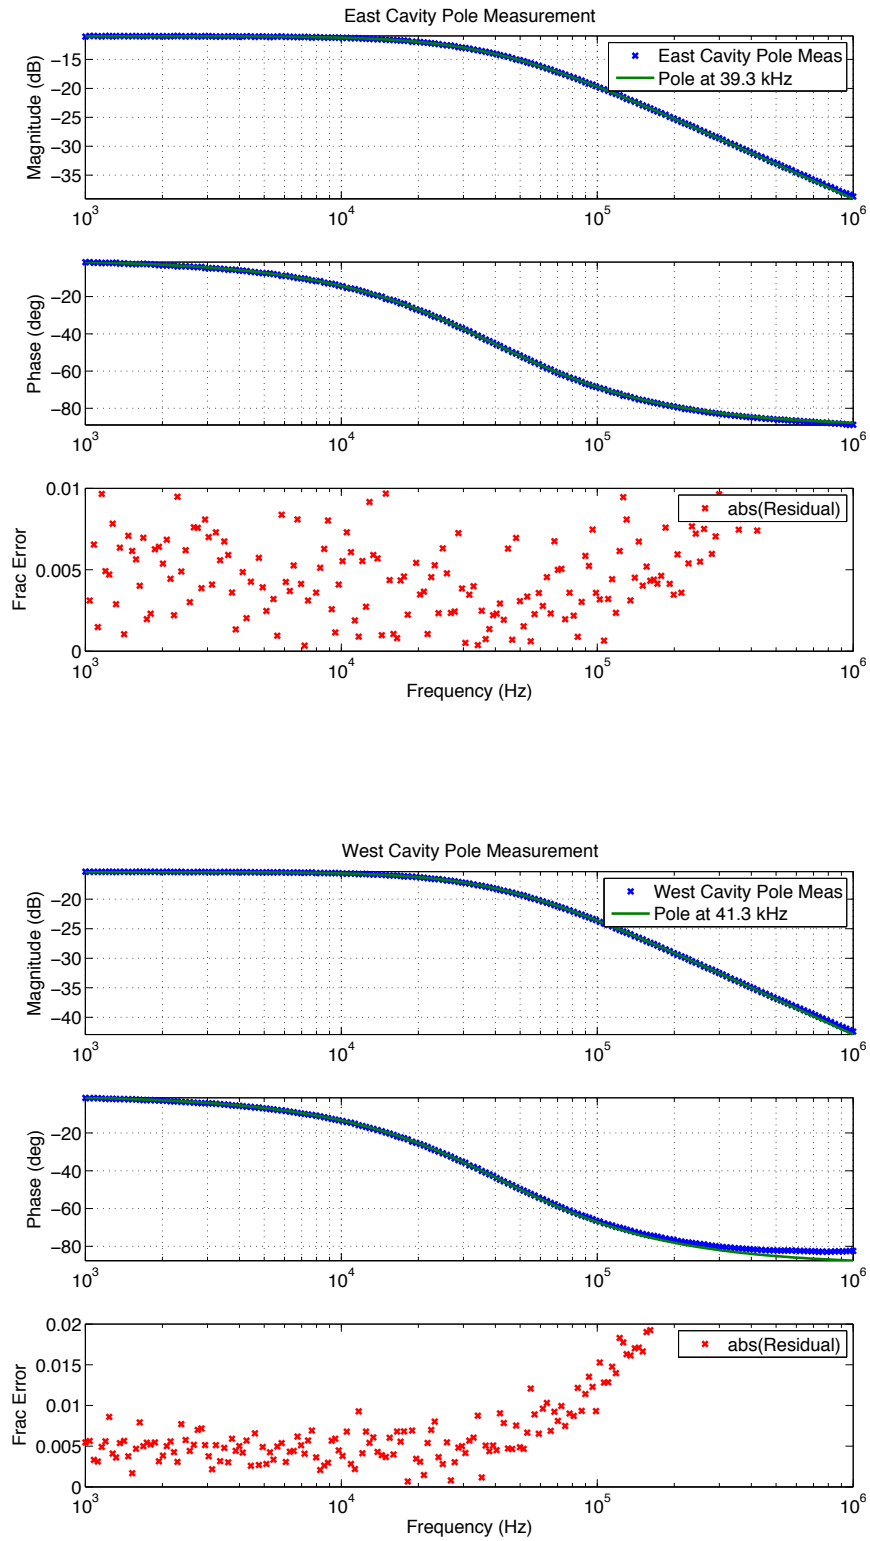


Figure 2.8: Measurements of the cavity pole for both East and West cavities. The high frequency deviation in phase is due to a mismatch in the input and transmitted photodetectors.

with proper SNR over the feature should out perform a time domain measurement of the exponential decay.

The physical dimensions of the cavity shrink by a factor of $\int_{123K}^{300K} \alpha_{Si}(T)dT \approx 1 \times 10^{-4}$ as we cool from 300 K to 123 K. We expect a 1×10^{-4} decrease in L_{cav} , R_{mirror} , f_{FSR} , and f_{cav} ⁶, no change in g_{cav} , and a 1×10^{-4} increase in f_{hom} . The expected differences in these quantities from cooling down are all significantly less than the measurement errors shown in table 2.2.

In summary: we engineered a 10 cm long cryogenic silicon Fabry-Perot cavity with low coating thermal noise with the goal of a low linewidth coating Brownian noise limited cavity. We chose modest optical parameters ($\mathcal{F} \approx 2 \times 10^4$, corresponding to a mirror transmission of $\approx 20ppm$). As we will see in chapter 4, making a low noise coating thermal noise limited cavity was harder than anticipated, and we will recommend shifting to a ≈ 1 cm short cavities with high coating noise in order to directly detect this.

⁶The 1×10^{-4} decrease in f_{cav} is solely based on the length change - changes due to varying loss as a function of temperature are unknown

Chapter 3

The Experiment


The experiment is a precision differential measurement of cavity length. In this chapter we detail the optical and mechanical design choices made:

- We rested a pair of (similar length) single crystal silicon cavities on rigid (PEEK plastic) support points, attached to a rigid platform to make vibrations common mode.
- We suspended the rigid platform from a downward facing LN₂ cooled optical breadboard by BeCu springs, which we damped using 0000 Cu wool.
- We used a custom designed 77 K cryostat with sapphire optical windows to couple 1550 nm light into the cavities.
- We surrounded each cavity with an individually temperature controlled radiation shield, which we used to control the temperature of the cavities at 123 K.
- We locked a 1550 nm laser to each cavity using PDH loops so that the laser frequency followed cavity length.
- We interfered the transmitted beams, which produced an RF beat with a frequency determined by the (changing) differential cavity length.
- We used a low noise phase-locked loop to convert the RF beat we detect into a signal proportional to the differential length.


3.1 The Supports

In order to study the cavities shown in fig. 2.6, we need to actually touch them in some way. How we choose to touch the cavities sets our seismic susceptibility (see section 4.3), the minimum mechanical loss of the cavity body modes (see chapter 5), and consequently the brownian thermal noise of the cavity (see eq. (4.4)).


The competing techniques for suspending rigid cavities (shown in fig. 3.1) are cradle style suspensions with springs or wires, vertically orienting the of the cavity [7], and rigid support points. We ruled out vertically orienting the cavity due to the complication and cost associated with periscoping a beam into the cryostat in a low noise way, though this would likely be a very good way to minimize the mechanical loss of the cavity. A single wire loop based suspension has been shown to support a very low mechanical loss [27], though the experience of the bar detector community led us to believe that the resultant motion of the cavity's optic axis would be too difficult to deal with [21]. We opted not to use a double wire loop cradle suspension due to our group's experience with up-converted differential cavity motion from swinging limiting experimental noise. For this work we chose a relatively simple design of rigid contact points on a horizontally oriented cylindrically symmetrical cavity (shown in fig. 3.2). For details on the performance of the supports and recommendations for future upgrades to them see section 4.3 and chapter 5.



(a) LIGO style spring and wire cradle reference cavity suspension. Picture courtesy of Frank Seifert.



(b) Cylindrical cavities supported by four rigid contact points in [23].



(c) The silicon football cavity produced in a collaborative effort by JILA and PTB [7]. Picture courtesy of Thomas Kessler.

Figure 3.1




Figure 3.2: A rendered image of the the cavity support structure, with and without a model cavity inside. Everything but the small support points themselves, which can be seen in (a), is made of aluminum.

3.2 The Suspension

Having established how we will touch the cavities in section 3.1, we determine how we will attach these supports to the cryostat (and thus to the optical table). The design philosophies we entertained were vibrationally isolating the cavities from the optical table by some sort of pendulum or spring based system (as done with the LIGO test masses and fig. 3.1a) versus rigidly attaching the cavities to the optical table to benefit from common mode rejection below the mechanical resonances of the experiment. A rigid suspension will generally be quieter below the mechanical resonances of the system, where a spring or pendulum suspension will provide better isolation above its primary resonance. Where each one is better depends on how low we can make the resonance of for isolation (limited in this case by space constraints inside the cryostat), and how high we can push the mechanical resonances of a rigid suspension. There were two generations of suspension used over the course of this work, for information on on the reasons behind the upgrade see appendix B.

The upgraded suspension design (shown in section 3.2) was a compromise between the above mentioned design philosophies. The cavities were both rigidly attached to a platform, and the platform was suspended inside the cryostat with BeCu springs. Differential motion between the optic axes is mitigated by having them both attached to the platform, whereas they are still seismically isolated from the cryostat and ground by the springs.

BeCu is used in often chosen for cryogenic applications due to its low mechanical loss at cryogenic temperatures. This property is actually undesirable for this use. A low loss material for the spring is associated with longer ring down times due to mechanical impulses (which occur in any realistic lab environment), and higher rms motion of the cavities. To lower the Q of the platform modes, we stuffed #0000 copper wool inside the top 1/4 of the spring and in between the side of the spring




Figure 3.3: Cryogenic suspension used in this work. Both cavities are rigidly attached to the suspension platform to improve common mode rejection of cavity motion. Suspension modes are damped by #0000 Cu wool as shown on the left.

and its support from the cryostat (see section 3.2). The Cu wool was packed such that the portion of the spring packed with wool appeared to move freely upon excitation of the platform. For details on the performance of the suspension and recommended upgrades see section 4.2 and section 4.3.

3.3 The Cryostat

We commissioned Precision Cryogenics¹ to build the optical cryostat. For ease of use we chose a simple two volume LN2 cryostat, with a downward facing Al cold plate with Heli-Coil inserts, and AR coated optical windows to couple into the cavities. We designed the cryostat almost entirely out of aluminum to minimize the spatial temperature gradients seen with stainless steel vacuum chambers. We could not use the industry standard ConFlat seals for the optical windows, as aluminum is too soft to use with a Cu o-ring, and we did not want to temperature cycle a rigid Al-SS weld joint.

For the windows, AR coated sapphire substrates were brazed to an industry standard stainless steel fixture, which was welded to a custom o-ring flange which used 0.030" groove for an indium seals. Given the additional difficulty of replacing windows with a custom flange design, future experimental design should avoid using an all-aluminum cryostat and rather solve the problem of spatial thermal gradients through radiation shielding or improved active temperature control. Electroplating the stainless steel surfaces is another potential solution, which gives us control over

¹Precision Cryogenics — <http://www.precisioncryo.com>

both the thermal conductivity of the vacuum chamber surfaces, as well as their emissivity, for designing the temperature control.

Additionally, the LN₂ chamber in the initial cryostat design was wrapped in superinsulation, which is an easy way to decrease the radiative load between the outer walls of the cryostat (at 300 K) and the cold chamber (at 77 K). We observed very long pump down times², which we attribute to the many layers of superinsulation gathering water and He. As the thermal loading between 300 K and 77 K does not appreciably change the LN₂ hold time of the experiment, we would rather it not exist at all. We advise that future generations of the cryostat avoid using superinsulation.

3.4 Thermal Shielding and Control


In order to operate the cavities at 123 K (the first zero crossings in $\alpha_{\text{Si}}(T)$), we need a way to reliably sense and control the temperature of the cavities. To thermal isolate the cavities from the cold chamber of the cryostat, we surrounded them with a 1/8" thick cylindrical aluminum radiation shield with a 1/2" diameter aperture at each end (shown in fig. 3.6). We sensed the temperature of the radiation shield using a 1000 ohm platinum RTD, mounted to a thermally conductive Kapton PCB (shown in fig. 3.5), which was attached to the radiation shield with GE 7031 varnish. The temperature sensor was wired in a 4-lead configuration with 36 gauge polyamide insulation, and read out with a SR CTC100 temperature sensing and control unit. We affixed a 2 in² 50 ohm Kapton clad resistive heating element next to the temperature sensor for each cavity. The temperature control was done using PID loops, digitally tuned by the CTC100, to feed the each temperature sensor signal back to its respective heater.

The wiring for all heaters and temperature sensors was routed through a block of LEMO connectors which was attached to the the cold plate in order to facilitate separation of the platform + cavity + radiation shield assembly from the cryostat. The wiring between the LEMO block and the room temperature breakout was wrapped in indium, and clamped to the cold plate in order to provide better thermal isolation for the radiation shields from room temperature. We read out 5 temperature sensors and two heaters in total. For details on motivations behind, and performance of the temperature control, see section 4.5.


3.5 Optical Layout

Now that we have discussed how we construct and cool a silicon cavity to support length stability, we need a way to measure its length. Keeping with the general experimental rule, differential

²In the course of upgrading the experiment, a significant amount of work is done with the vacuum chamber open. Depending on how long the cryostat had been open to air, this long pump down time could be weeks long.



(b) Rendered SolidWorks model of as-built cryostat. The suspension and radiation shield assembly shown inside the cold chamber is the first generation of both as described in appendix B.



(a) Cryostat design.

Figure 3.4




Figure 3.5: Shown here is the temperature sensor and heater, as fixed to the radiation shield. We printed a custom Kapton PCB to make a mountable 4-lead temperature sensor which could be attached with easy to remove varnish. The heater wires are broken in this image.

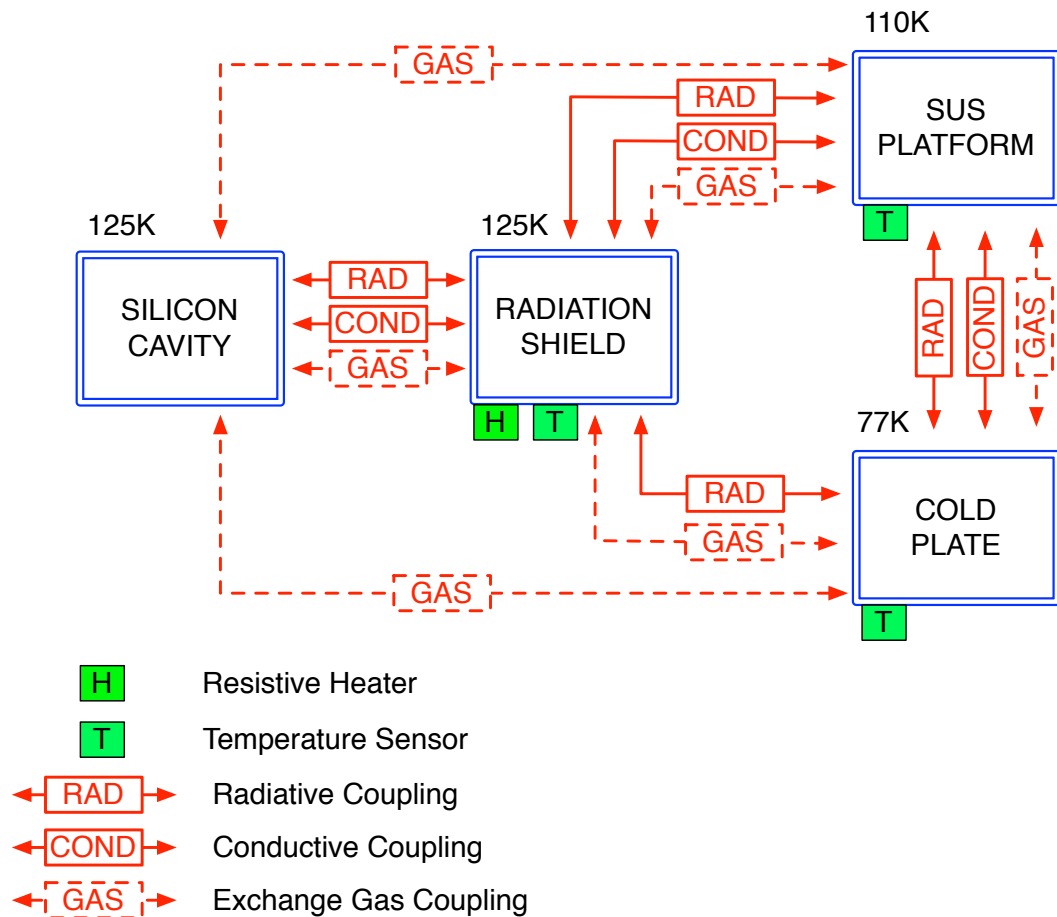


Figure 3.6: A diagram of the thermal coupling. The radiative and conductive thermal couplings are called out as solid red lines, and temperature sensors and heaters are shown in green. The dashed line represents exchange gas (N_2) which we backfill into the volume with the cavities in order to cool them quicker. During low noise operating the dashed line is not large enough to be relevant.

measurements are far easier to make than absolute measurements. A commonly used technique [28, 23, 29, 30, 31] for measuring differential cavity length changes is to lock a laser to each cavity (so that fractional changes in length become fractional changes in laser frequency), and then compare the two fields via interferometric measurement. If the fields are close enough in frequency, then their RF beat can be detected on a photodiode, and the frequency noise of this signal can be analyzed using a frequency discriminator (such as a phase-locked loop as shown in fig. 3.10 and used in this work, or alternately a self delay homodyne measurement [32]). If the fields are separated by many GHz, then advanced measurement techniques like fiber frequency comb are needed to compare the two [33, 34]. We chose the cavity length and laser frequency to facilitate direct detection of the RF beat, as frequency combs add significant complication to the optical setup.

We locked a 1550 nm laser to each cavity using the standard PDH technique [35]. We added phase modulation sidebands at f_m to each input beam using a free space EOM, and read out the cavity retroreflection with a resonant RF photodetector, which we then demodulated at f_m , producing a linear signal in volts as a function of offset from perfect resonance (perfect resonance corresponds to an exact integer number of wavelengths fitting inside the cavity). We controlled the laser frequency by actuating on the diode current (which has a strong response in Hz/mA appendix E). Within the bandwidth of the PDH loops, fractional changes in the cavity length are imprinted upon the interrogating light frequency as:

$$\frac{\delta\nu}{\nu} = \frac{\delta L}{L}, \quad (3.1)$$

where ν is the optical frequency (1.93×10^{14} Hz for 1550 nm light), L is the length of the cavity (≈ 10 cm in this work), and the δ 's are small changes in each quantity.

To help understand the optical layout (shown in fig. 3.7 and fig. 3.9), we trace one of the individual input paths (omitting mode matching lenses and steering mirrors): the light is produced in a fiber coupled laser diode, exits the fiber as a free space beam at the end of a collimator, and is linearly polarized at this point. Then it passes through a 1/2 wave plate and a faraday isolator, which are used in conjunction to set the the power level on the table. The light rejected from the faraday is sent onto a razor blade stack beam dump. The transmitted p-polarized field is sent to an acousto-optic modulator (AOM - used as a fast intensity actuator such that everything but the zeroth order diffraction is dumped). It then proceeds to a 1/2 wave plate, then an EOM, and then another half wave plate (the EOM polarization axis drifts around depending on temperature). The EOM produces phase modulation sidebands at $f_m \approx 33$ MHz for use in the PDH locking scheme. (We refer to the remaining light at the unshifted optical frequency as the carrier.) The p-polarized field is transmitted through polarizing beam splitter, and a portion of the input power to the cryostat is picked off with partially transmissive optics as an input power monitor. After the partial reflectors,

the field proceeds through a 1/2 wave plate, a 1/4 wave plate, and both cryostat windows, and finally is reflected off the cavity mirror. The 1/4 wave plate is used to make the light incident on the cavity circularly polarized, and the 1/2 wave plate corrects for axially asymmetric dispersion in the cryostat windows. The prompt retro-reflection, the sidebands, and the carrier which resonated in the cavity all proceed back along the optic axis through both windows, the 1/4 wave plate, the half wave plate, and are now s-polarized. At the polarizing beam splitter, the field is now all reflected, and detected with a resonant photodetector. The fields all mix, and we demodulate the signal at f_m , generating the error signal (a linear signal dependent on the offset between carrier and cavity length). This signal is used to control laser frequency via laser current. For more details on the PDH loops see section 4.6.

The light resonating inside the cavity is transmitted through the end mirror of the cavity, passes through both cryostat windows, a half wave plate (to match polarizations with the light from the other cavity), and a quarter wave plate (recalling that the light inside the cavity is circularly polarized). The light then encounters a 70:30 beamsplitter, where 30% of the field is transmitted onto a photodetector to monitor the transmitted power. The reflected path is steered towards a 50:50 beamsplitter, where it spatially overlaps with the field transmitted through the other cavity, and both are steered onto a broadband RF photodiode, where they interfere and produce a signal:

$$V(t) \propto \sin(2\pi\Delta\nu t), \quad (3.2)$$

where the beat frequency $\Delta\nu$ is the differential change in the frequency of light transmitted through each cavity related to differential changes in cavity length by eq. (3.1):

$$\Delta\nu = \delta\nu_2 - \delta\nu_1 = \frac{\nu}{L_{\text{cav}}} (\delta L_2 - \delta L_1) = \frac{1.93 \times 10^{14} \text{ Hz}}{10 \text{ cm}} (\delta L_2 - \delta L_1), \quad (3.3)$$

denoting the two paths with the subscripts 1 and 2. Ignoring the difference between optical frequency and cavity length (shown in table 2.2) only causes a 2% systematic error in the calibration of the beat, which is negligible. This beat is read out using a PLL (phase-locked loop) as shown in fig. 3.10.

For the experimental layout shown in fig. 3.9 to be low noise, we must limit the amount of light backscattered into the optical path. Wherever the light is incident on a surface, there is some degree of backscatter. To minimize this, it is important that:

- All optics (both reflective and transmissive) should be as clean and smooth as possible. Rough surfaces and dirty optics generate significant scatter. The cryostat windows are not exempt from this rule, and happen to be a particularly bad offender.
- All terminal beams (especially reflections off the photodetectors, where we are most sensitive

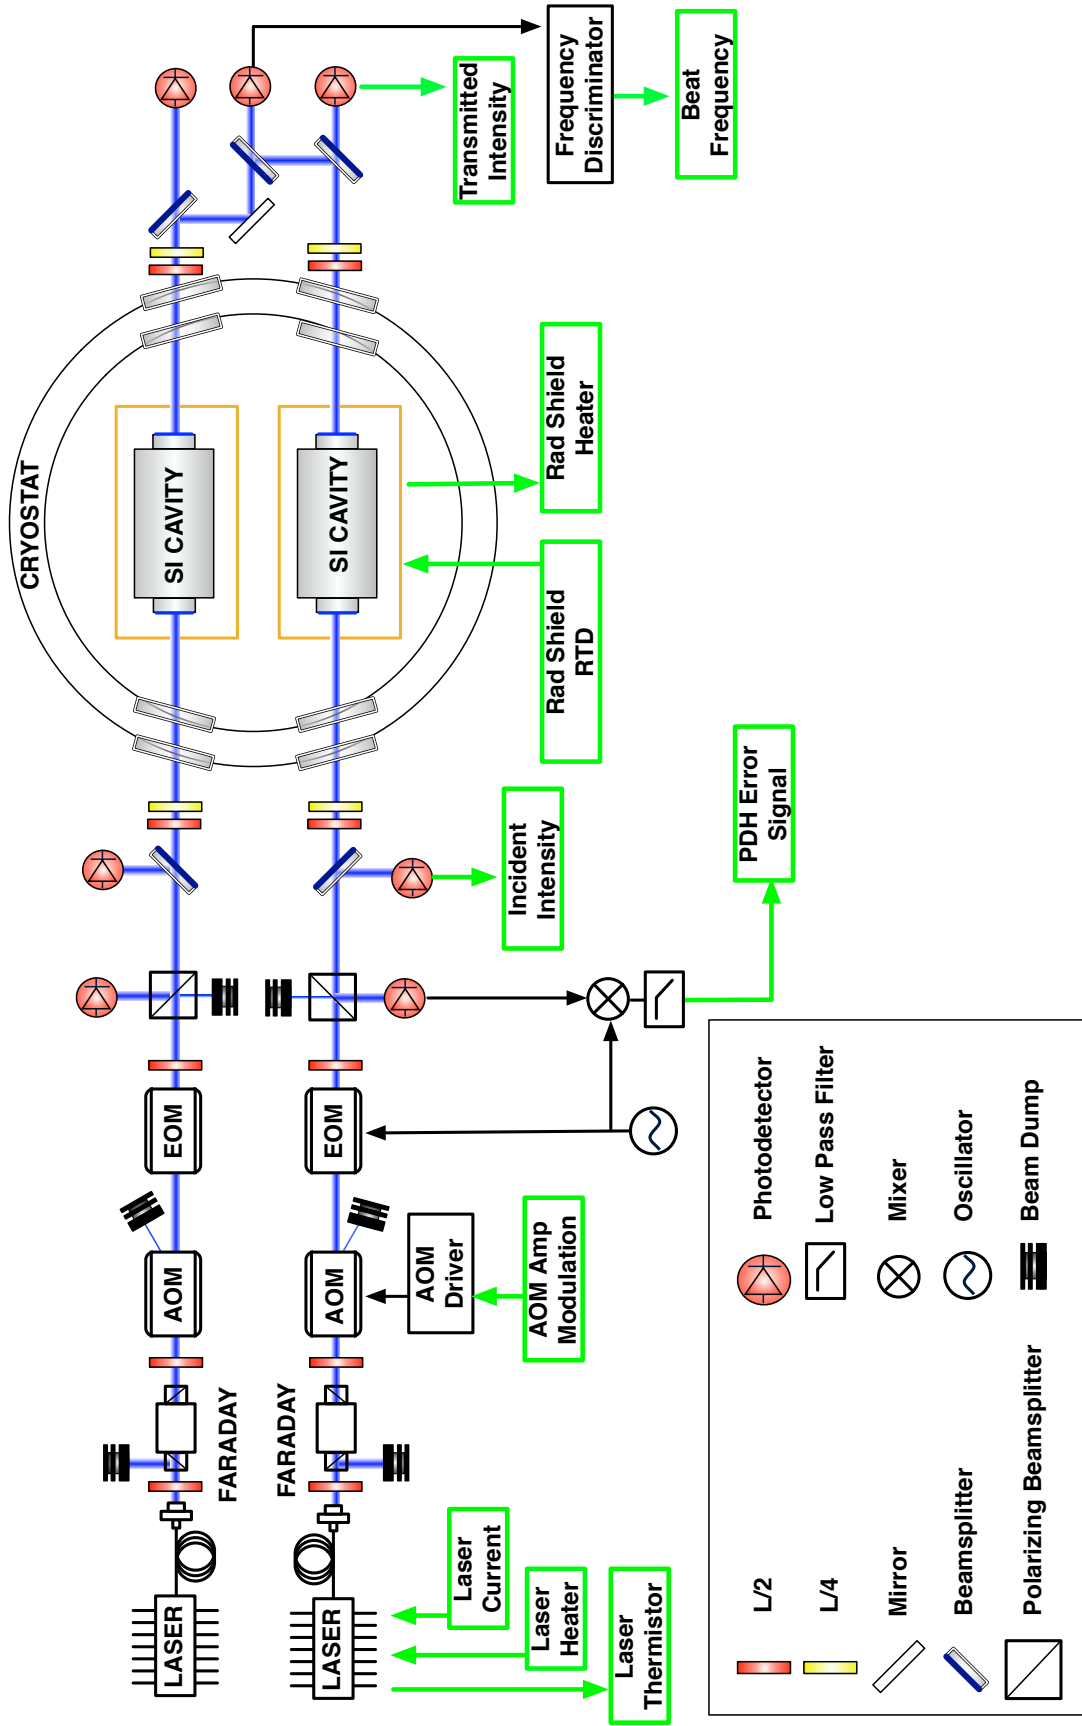


Figure 3.7: Experimental Diagram with sensing and control channels highlighted. Duplicate labeling was omitted, but each path has identical sensing and actuation.




Figure 3.8: Experimental setup. See fig. 3.9 for blow up of parts A,B,C.

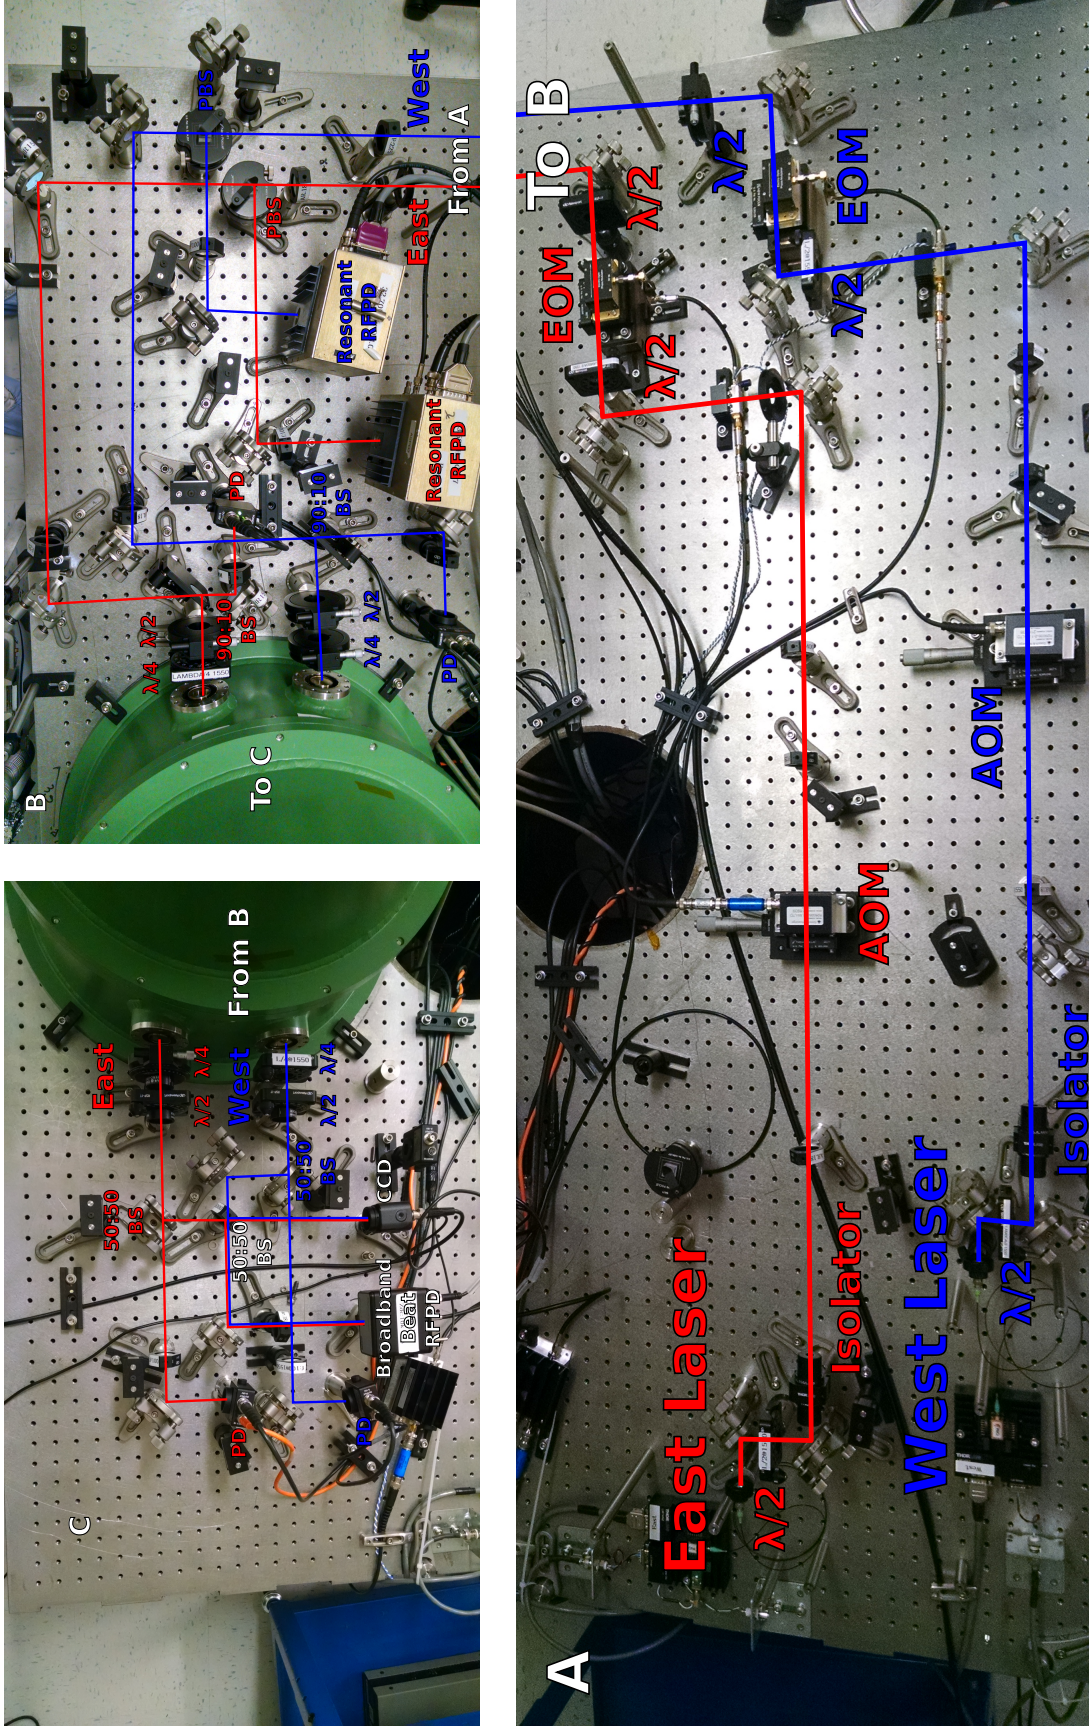


Figure 3.9: Panel A shows the lasers and EOMs. Panel B shows the cavity input and PDH retro-reflection optics. Panel C shows the cavity transmission and beat readout.

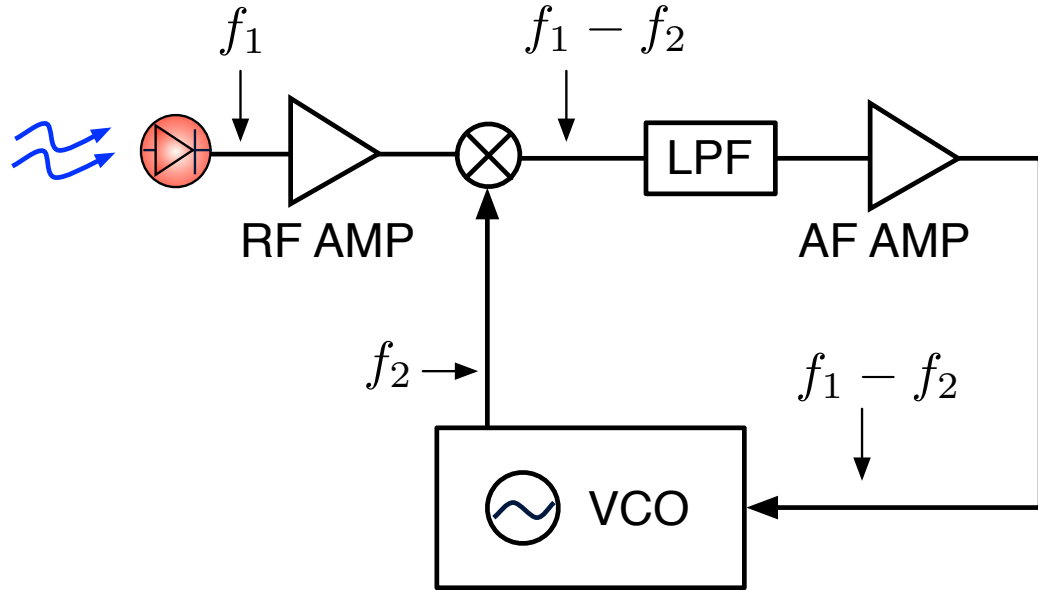


Figure 3.10: PLL used for beat frequency readout. When the frequency noise in the VCO is lower than the frequency noise in the beat, then the control signal in the PLL is a readout of the beat frequency (and thus cavity length fluctuations).

to backscatter) are dumped on a black non-scattering surface such as razor-blade stack beam dumps.

- All optics are angled or wedged to spatially separate stray beams due to multiple reflections which occur inside all transmissive optics.

Hunting for scatter sources and mitigating them is a crucial part of the experimental (re)design process. See section 4.2 for coupling levels and details on how scatter spoils experimental noise.

3.6 Lessons learned

There are a number of fundamental design choices which we would make differently based on what we learned over the course of the experiment. These are:

- Cavity length. Long cavities with large spots for ultra-stable cavity with low coating thermal noise \rightarrow short cavities with small spots to increase coating thermal noise and make it easier to measure.
- Cavity construction. Affixing the mirror substrates to the spacer with varnish \rightarrow optically contacting the mirror directly to the spacer.
- Cavity support. Rigid supports which touch cavity \rightarrow isolating cavity body mode to support high mechanical Q.

- Cryostat construction. Making an all aluminum cryostat to reduce spatial thermal gradients → using an all stainless steel cryostat.
- Suspension damping. Using Cu wool to damp the BeCu suspensions → make springs out of a material with low mechanical Q factor and use eddy current damping.
- Window flanges. Non wedged sapphire windows with e beam coatings brazed into custom flanges → super-polished wedged fused silica optics in industry standard (Conflat) flanges with ion beam sputtered coatings, and a tightly controlled workflow for protecting window surfaces.

We discuss the reasons for these changes throughout this thesis.

Chapter 4

Noise Budget

We use (one sided) amplitude spectral density to analyze the experimental noise. When talking about frequency stability, it is standard to work in units of Hz/ $\sqrt{\text{Hz}}$ rather than canceling out the units, as the numerator represents frequency deviations and the denominator Fourier frequency. To help intuitively understand spectral densities, it helps to express Parseval's theorem in the convention:

$$X_{\text{rms}} = \sqrt{\frac{1}{T} \int_0^T [X(t)]^2 dt} = \sqrt{\int_{1/T}^{\infty} S_X(f) df}, \quad (4.1)$$

where $X(t)$ is the time series of some physical variable, T is the time over which we observe it, and $S_X(f)$ is the power spectral density of X ($\sqrt{S_X(f)}$ is the ASD). For this work, all spectra were created using a Hann windowed FFT. For further discussion on the mathematical reasoning behind spectral densities, see chapter 6 of [36].

The experimental noise budget for this work is shown in fig. 4.1. The limiting noise sources are well understood, except between 200 Hz and 1 kHz, and are discussed in depth in this chapter. The evolution of the experimental noise over the course of this work is shown in fig. 4.2.

4.1 Thermal Noise

4.1.1 The Fluctuation Dissipation Theorem

As shown in 1.1, the way in which Brownian motion couples into a variable we are trying to measure can be expressed in terms of mechanical properties of the system. As the mechanical impedance, $Z(f)$, of a system can be hard to calculate, we use the Levin formalism for the fluctuation dissipation theorem [37]. When you are trying to measure some physical parameter (in our case the average phase front of an optical field as a proxy for changes in mirror position), you sinusoidally push on the conjugate variable (field amplitude), and by examining the amount of energy dissipated per cycle, you know how much the original variable (the phase front) is moving due to thermodynamic

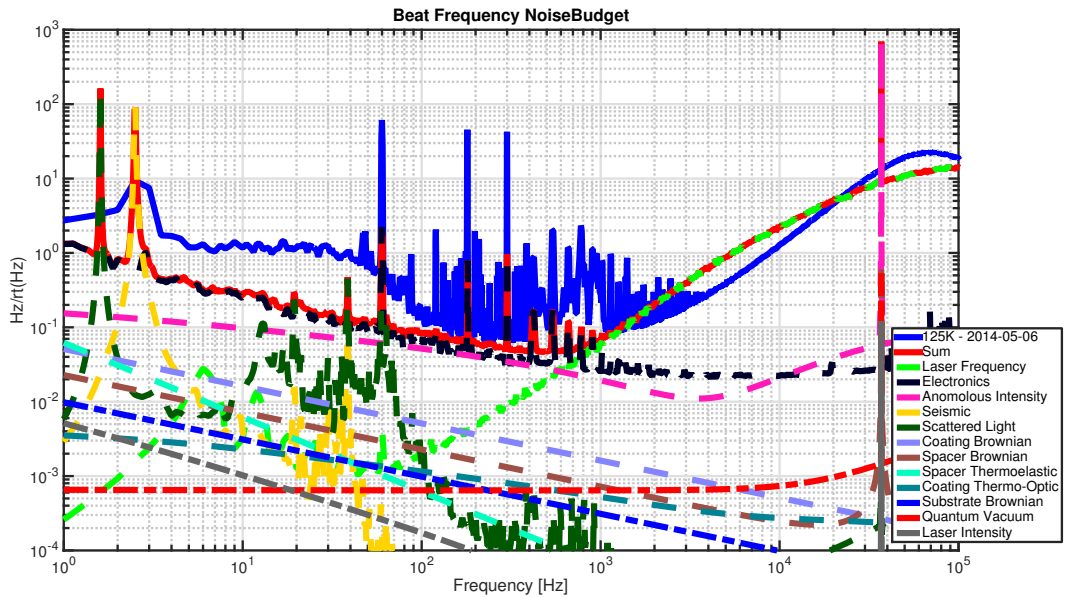


Figure 4.1: Experimental Noise Budget - total experimental noise is the quadrature sum of all terms.

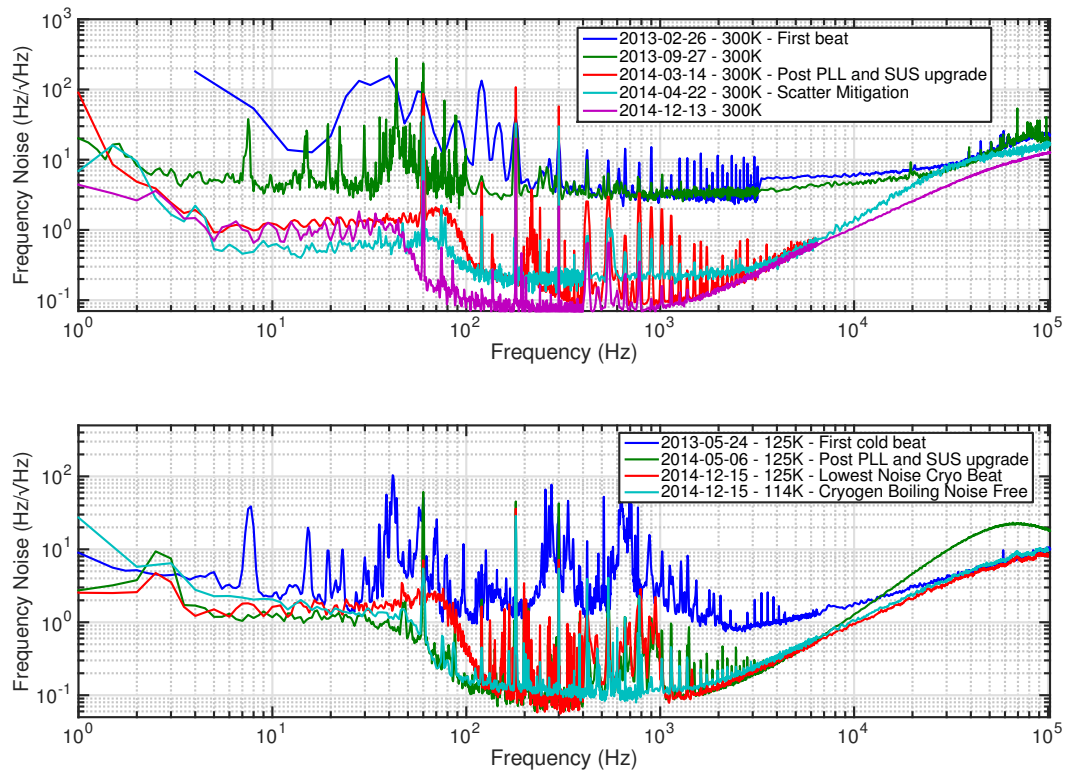


Figure 4.2: Evolution of the experimental noise over the course of this work.

fluctuations. In equation form, this reads:

$$S_x(f) = \frac{2k_B T W_{\text{diss}}}{\pi^2 f^2 F_0^2}, \quad (4.2)$$

where x is the variable you are interested in measuring, k_B and T are Boltzmann's constant and the temperature, respectively, F_0 is the (small) magnitude of the sinusoidal applied force, and W_{diss} is the average power lost to dissipative forces due to the oscillating force. For each dissipative mechanism in the system you are examining, you must treat the coupling mechanisms coherently. We examine both the damping forces associated with the mechanical loss angle, as well as those from thermoelastic effects, where Zener damping comes into play.

4.1.2 Brownian Noise

The simplest, and earliest understood thermal noise source is referred to in the literature as Brownian thermal noise. In the case of a reference cavity, when you push on the mirror surface with a gaussian pressure profile, there is elastic damping in the mirror coatings, in the mirror substrates, and in the cavity spacer. The mechanical loss of the various materials is what determines the level of Brownian thermal noise in each part of the cavity.

For the case of the mirror substrate and coating, the power spectrum of noise can be expressed as:

$$S_x(f) = \frac{4k_B T}{\pi^{3/2} f} \frac{1 - \sigma^2}{Y w_z} \left[\phi_{\text{sub}} + \frac{2}{\pi^{1/2}} \frac{1 - 2\sigma}{1 - \sigma} \frac{d_{\text{coat}} \phi_{\text{coat}}}{w_z} \right], \quad (4.3)$$

where σ is the Poisson ratio, Y is the Young's modulus of the mirror substrate, w_z is the beam radius at the mirror, d_{coat} is the thickness of the optical coating, and ϕ_{sub} and ϕ_{coat} are the bulk mechanical loss angles of the substrate and coating, respectively [38]. The actual mechanical loss in the coating is a complicated function of material properties and coating thickness, which is, in general, not a volumetric average of the bulk loss in the coating materials. eq. (4.3) was modified to Harry et al [39] separate ϕ_{coat} into ϕ_{para} and ϕ_{perp} , and then corrected by Hong et al [40] using ϕ_{bulk} and ϕ_{shear} to avoid negative energy densities in the calculation of energy lost per cycle when applying the Levin approach eq. (4.2). For the purposes of this work, we will use eq. (4.3) for the coating Brownian noise calculation, as it requires the most limited knowledge of cryogenic properties of the individual coating materials, many of which are unknown. It is important to note that the loss used in this calculation was that for SiO₂:Ta₂O₅ at 300 K, though we are currently pursuing estimates of material properties to enable a better estimate of the cryogenic coating Brownian noise. We believe that the mechanical loss in the coatings gets worse at 123 K, which increases the coating Brownian noise an unknown amount. Obtaining a reliable number for ϕ_{coat} at 123 K should be considered a very high priority.

For the spacer itself, the brownian noise is expressed by Kessler et al [41] as:

$$S_x^{\text{spac}}(f) = \frac{4k_B T}{2\pi f} \frac{L}{2\pi(R^2 - r^2)Y} \phi_{\text{spac}}, \quad (4.4)$$

where L , R , and r are the length, outer radius, and inner radius of the cylindrical spacer, and ϕ_{spac} is its mechanical loss angle. COMSOL simulations for room temperature SiO₂ cavities with similar geometry were done in chapter 2 of [12]. The results of their simulation agree with 4.4.

It should be noted that the actual loss angle used in this calculation, $\phi_{\text{space}} = Q^{-1}$, is the inverse mechanical quality factor of the *actual supported spacer's first body mode*. This is potentially much less than the theoretical minimum loss of the material used, and in the case of this preliminary study, was worse by five orders of magnitude because of the support scheme used (see chapter 5). In the cavity geometry we used, the spacer Brownian noise dominates the other thermal noise sources at a spacer Q of 7×10^3 , which is less than the as-built experimental body mode Q. The suspension should be upgraded in order to suppress the spacer Brownian noise, but it does not dominate the other thermal noise sources.

4.1.3 Thermo-elastic Noise

In addition to the damping from mechanical friction described by ϕ , there is also damping due to the thermoelastic effect. Local fluctuations in strain produce fluctuations in temperature based on the coefficient of thermal expansion. When you apply a sinusoidal pressure force to the mirror surface as in [37], and create varying strain inside the cavity, you also create thermal gradients. Energy is non-reversibly dissipated as the thermal gradients flow towards equilibrium, which gives another mechanism for W_{diss} .

For thermoelastic noise in the substrate, we followed the calculations for the LIGO mirrors begun by Liu and Thorne [42], and corrected for finite substrate size by Cerdonio et al [43]. At high frequencies, these two solutions converge, and can be written as:

$$S_x(f) = \frac{8}{(2\pi)^{5/2}} \alpha^2 (1 + \sigma)^2 \frac{k_B T^2}{(\rho C)^2} \frac{\kappa}{f^2 r_0^3}, \quad (4.5)$$

where α is the (temperature dependent) coefficient of thermal expansion, σ is Poisson's ratio, ρ is the density, C is the specific heat capacity, κ is the thermal conductivity, and $r_0 = \omega_0/\sqrt{2}$ is the gaussian beam waist divided by $\sqrt{2}$. The work of Cerdonio et al showed that below a certain corner frequency, $f_c = \kappa/(\pi\rho C\omega_0^2)$, the slope of thermoelastic noise in a substrate changed from f^{-2} to $f^{-2/5}$. In the silicon cavities used for this experiment, $\omega_0 = 340 \mu\text{m}$, and the corner frequency is $f_c = 80 \text{ Hz}$.

When dealing with Silicon, the $\alpha^2\kappa$ scaling is far more significant in determining the thermoelastic noise level than the low frequency corrections of Cerdonio et al. While it is theoretically possible to

eliminate this noise to first order, in reality we can only control the temperature of the cavities so well. The already large thermal conductivity of Silicon increases by a factor of a few when moving from 300 K to 123 K. To provide an intuitive understanding we look at the ratio of $\alpha^2\kappa$ of silicon at 123 K to that of SiO₂ (commonly used for ultra stable cavities at room temperatures) at 300 K:

$$\frac{\alpha_{Si}^2(123K)\kappa_{Si}(123K)}{\alpha_{SiO_2}^2(300K)\kappa_{SiO_2}(300K)} \approx 0.5 \left(\frac{\delta T}{K} \right), \quad (4.6)$$

where $\delta T/K$ is the number of degrees, in Kelvin, which we operate off the zero crossing temperature of silicon. So long as the temperature is held within 1.4 K of the zero crossing, the thermoelastic noise in the mirror substrates in a silicon cavity operating at 123 K will be less than that in an identical setup using fused silica running at 300 K. Fused silica and silicon have similar densities so we omitted that ratio for simplicity. If one wanted to directly measure thermoelastic noise in a substrate, operating at the temperature at which $\alpha^2\kappa$ has a maximum value would be a good approach. Examining the properties of silicon (see 2.1 and 2.3), this maximum will be somewhere around 300 K.

Most of the detailed analytic calculations for these various thermal noise sources are driven by groups trying to understand the noise sources that couple into the LIGO gravitational wave detector. As the LIGO mirrors are suspended rather than separated by a spacer, there has been little interest in deriving an analytic form for this noise source. The same scaling arguments we use to understand the thermoelastic noise in the substrate apply here. So long as we hold the temperature close enough to the zero crossing in α , this noise will not matter. Based on the geometric (cylindrical) similarity of the spacer and substrate, it seems reasonable to assume the loss mechanisms have similar forms, and the ratio of spacer thermoelastic noise in silicon to that in SiO₂ has the same $\alpha^2\kappa$ scaling. Thus we believe that that thermoelastic noise in the spacer will be less than that in a SiO₂ spacer at 300 K.

If one wants to correctly determine the thermoelastic noise in a cavity, a good method is to use a finite element modeling program such as COMSOL to simulate the thermoelastic dissipation when you push on the mirrors with a gaussian pressure profile. This is simply applying the Levin recipe for finding the thermal noise numerically, rather than analytically. This was done for SiO₂ cavities at 300 K by Chalermongsak et al [12] and shown to be small compared to the other thermal noise sources. Based on the logic outlined above we believe the spacer thermoelastic noise will not be limiting.

4.1.4 Thermo-refractive Noise

The thermo-refractive noise was calculated using the equations presented for finite cylindrical mirror substrates by Heinert et al [44], which expands upon the work done by Braginsky et al [45] for an infinite sized cylindrical mirror substrate. For linear cavities using PDH locking, only the path

length change in transmission matters, since fluctuations in the reflected path length are common mode to both the carrier and sidebands (so long as carrier and sideband see the same local index of refraction, which is a good approximation to 1 part in 10×10^7).

Using a value of $\beta = 9.8 \times 10^{-5}$ [46], we found the thermo-refractive noise in the substrate to be very small compared to the other thermal noise sources: above 1 kHz the noise is flat at a level of 2.3×10^{-7} Hz/Hz^{1/2}, and below this it is decreasing in value with a positive slope in f . The experimental topology is chosen because it is much more sensitive to cavity length fluctuations than optical path length fluctuations (such as those occurring from every optical mount on the table fluctuating), and thermo-refractive noise in the mirror substrate couples is just an optical path length change, so it is unsurprising to find that this noise level is so small.

There has been consideration of making linear cavities by machining a single piece of silicon and coating the ends, and more advanced geometries making use of total internal reflection [47]. Both ideas would have much higher thermo-refractive noise than that in a hollow cavity.

4.1.5 Thermo-Optic Noise

Thermoelastic and thermo-refractive noise in the optical coatings must be treated coherently, as described by Evans et al [48]. They use the analytic form of statistical fluctuations in temperature within the coating due to the thermal gradients, which was calculated by Levin in [49], and then coherently treating the resulting changes to coating reflectivity, and both thermoelastic and thermo-refractive changes to the optical path length.

Work has gone into designing coatings which exploit the different sign in α and β of the coatings to provide cancellation of thermal noise for LIGO, as well as in reference cavities [23].

For the SiO₂:Ta₂O₅ quarter wave stack coatings used in this experiment, there is no expected cancellation of thermo-optic noise, and the coating brownian noise dominates. To estimate the thermo-optic noise level, we used code written by Matthew Abernathy, which follows the recipe laid out in [48], with several corrections. For the initial estimates, we used the room temperature properties for SiO₂:Ta₂O₅ (also given in [48]). Better estimates should be made using the real 123 K material properties of the coatings, though we expect that coating brownian noise will still dominate.

4.1.6 Silicon Carrier Density Noise

A noise source which has recently been studied by Heinert et al [50] is that caused by fluctuations in free carrier density of semiconductors used as transmissive optics. At frequencies higher than the effective diffusional flow of carriers, this is expressed as:

$$S_z(\omega) = 4 \sqrt[3]{\frac{3}{\pi^7} \frac{H}{r_0^4} D \frac{\gamma^2 m k_B T}{\hbar^2} \sqrt[3]{n} \frac{1}{\omega^2}}, \quad (4.7)$$

where D is the diffusion constant for free carriers, H is the thickness along the optic axis, m is the electron mass, and n is the density of charge carriers. The silicon substrates used in the cavities are p-type doped Cz grown, with a (100) crystal axis, though we do not explicitly know their resistivity. Assuming the p-type dopant is Boron, a very common dopant used in the semiconductor industry, we make the most pessimistic estimate for charge carrier density we can: resistivity = 0.005 Ohm-cm, which corresponds to¹ a density of holes of $n_h = 2 \times 10^{19} \text{ cm}^{-2}$. The amplitude of the noise scales as the sixth root in n_h , so the calculation is fairly forgiving to uncertainties in carrier density. Boron and phosphorous have very similar ionization energies (44 and 46 meV, respectively), and so the calculation used to derive eq. (4.7), which assumed a phosphorous dopant holds for boron.

The upper bound on amplitude spectral density of noise for our cavities is closely approximated by:

$$\sqrt{S_f(f)} = 3.4 \times 10^{-8} \frac{f}{1 + f/(9 \text{ kHz})} \text{ Hz}/\sqrt{\text{Hz}}. \quad (4.8)$$

Below the corner frequency of 9 kHz, the slope is closer to $f^{9/10}$. The coating Brownian noise dominates this below 20 kHz. If the experimental sensitivity reaches this level at high frequency in future upgrades, we might uncover this noise.

As with thermo-refractive noise in the substrates, the experiment is designed to be insensitive to these effects, and only the transmission is relevant for the noise. Any future cavities which considered having silicon on the inside of the cavity would need to consider doping more carefully, as they would be much more sensitive to this noise process.

4.2 Scatter

Below 100 Hz, the experiment is limited by scatter noise. This occurs when some spurious scattering path recombines with the retro-reflected cavity light, and beats with it. The relative phase between the cavity retroreflection and the spurious scatter path looks like cavity length changes, but have a different morphology. Since the scatter path is not a length stabilized path, the motion is often multiple wavelengths, and what couples into the error signal is the up-converted noise as the spurious path wraps through multiple fringes. We were able to move the corner frequency of the scatter bump up and down by exciting the swinging modes of the cavity platform suspension. We also were able to audibly identify fringe wrapping by listening to the beat signal on headphones.

We systematically excited each optic in the setup with broadband acoustic noise injections using a PZT wand powered by a HV amplifier. Listening to the beat output on headphones while acoustically exciting the optics allowed us to rule out individual optics as the main scatter point. We were

¹Using the online calculator available at <http://www.cleanroom.byu.edu/ResistivityCal.phtml>, which relates silicon resistivity and doping.

unable to significantly change the low frequency noise by exciting any individual optic on the table layout. We believe that the scatter offender is the actual cryostat windows, which we were unable to individually excite. We can determine the platform RMS motion from the corner frequency of the scatter to be $5.4 \mu\text{m}_{\text{rms}}$, which is 3.5λ . See appendix D for more information on why we suspect the cryostat windows, and scatter measurements.

There are several measures available for improving the scattering noise. The corner frequency of the scatter noise can be lowered by further damping the suspension, which dominates the path length change between the cavity retro-reflection and the scatter path. The height of the scatter peak can be lowered by lowering the scatter, which we believe to be the windows. New windows should be made, with extra care made to protecting coating fidelity in the fabrication process. Advanced techniques such as adding a PZT shaker to the offending scatter object in order to frequency shift the scattered light [51] can also be used, but the above mentioned changes should be made first.

4.3 Seismic

Seismic motion dominates the beat frequency noise at the six pendulum resonances of the suspension, the violin modes of the suspension springs, and via scatter (see section 4.2) below 100 Hz.

As discussed above, the higher the rms motion of the platform (and thus cavities), the higher the frequency of the scatter hump will be. The bulk motion can be decreased without changing the support scheme shown in section 3.2 if we further damp the modes of the suspension. We can try to use the current damping design, change the spring material, or add advanced active and/or passive damping techniques.

The swinging modes of the suspension platform are damped where the spring meets the cold plate by friction between the BeCu spring, the 0000 Cu wool, and the Al suspension bracket (shown in section 3.2). The mounting bracket can be easily reengineered to facilitate damping of these swinging modes by having it girdle the spring. We could also attack the problem at its source and attempt to find a lossier material to make a suspension spring from than BeCu. BeCu was initially chosen to minimize the thermal noise in the suspension as this has long been a very significant thermal noise source for the LIGO experiment [52, 53, 54, 55]. This is considered a mistake, as the significant mechanical distance between the cavities and the suspension makes this an irrelevant noise source; lowering the Q of all spring modes would improve the experimental noise at the scatter hump.

The seismic coupling path we expect to dominate in the sensitive band of the experiment (audio frequencies above the scatter hump) is vertical seismic motion coupling to vertical platform motion. The cavity bends under the DC force of gravity, which causes a change in the length of the optic axis. Below the mechanical resonance of the bending mode (tens of kHz), this behavior describes

each cavity’s acceleration sensitivity well (in [m/g]). Shown in fig. 4.5 is a 3D render of the cavity bending under vertical acceleration, as well as the results of FEA simulations using COMSOL which parametrize the sensitivity as a function of support position. The cavity support points were separated longitudinally by 2.5 inches, and were all 30 degrees off vertical.

If we assume that the only difference between the two cavity assemblies is the contact points for the support, then we can propagate uncertainties in the support points into an estimate of the seismic common mode rejection (CMR).

Shown in section 4.3 is the vertical acceleration measured at the top of the cryostat for the various states of the experiment (e.g., turbo pump on/off, and LN2 boiling noise). The measurement was done using a set of Wilcoxon accelerometers bolted to the top of the cryostat with a single 1/4-20 screw. The noise was estimated by measuring the coherence between sensors placed on opposite sides of the cryostat lid. The amplitude spectral density of noise shown in section 4.3 is:

$$\sqrt{S_N(f)} = [1 - \gamma_{z_1, z_2}^2(f)]^{1/2} \sqrt{S_{z_1}(f)}, \quad (4.9)$$

where z_1, z_2 are the two measurements of vertical acceleration, $\gamma_{z_1, z_2}^2(f)$ is the coherence function between the two, N is the measurement noise floor, and $\sqrt{S_N(f)}$ is the amplitude spectral density of N (as discussed in chapter 6 of [56]).

To propagate the vertical seismic noise down the top of the cryostat to the cavity support points, we look at the signal flow through the cryostat assembly. This is highlighted in fig. 4.3. The cold plate (upon which the LN2 chamber lies) is connected to the top of the cryostat via a 10” tall 10” diameter x 0.030” thickness G10 cylinder. The only portion of the vertical transfer function chain which has features in the bandwidth of our seismic noise measurement is the spring suspension. The bounce mode is at 2.6 Hz, with $Q \approx 250$ (see appendix B for more). The suspension platform is a 1/4” thick 8” x 10” anodized aluminum optical breadboard, and rigidly ties the two cavities together. Below 770 Hz (the first resonant frequency of the platform), the two cavities assemblies experience roughly identical vertical acceleration. The bounce modes of the softer (PTFE, PEEK) materials making up the support structure are above ≈ 500 Hz, and we ignore them for this estimate. To translate the seismic noise measured at the top of the cryostat (fig. 4.4) to frequency noise in the cavities (fig. 4.5) we did the following:

$$\sqrt{S_x(f)} = \sqrt{S_z(f)} \times \frac{1}{1 - if/2.6 \times 250 - f^2/2.6^2} \times \sigma_m \frac{\partial}{\partial m} H_{\text{vert}}(m; f) \times \frac{1.93 \times 10^{14} \text{Hz}}{10 \text{cm}}, \quad (4.10)$$

where x is cavity length, z is vertical acceleration noise at the top of the cryostat, m is the distance between the s for one cavity, σ_m is our uncertainty in the support distance between the two cavities, and $H_{\text{vert}}(m; f)$ is the cavity’s modeled length sensitivity to vertical acceleration (shown

in fig. 4.5). $H_{\text{vert}}(m; f)$ is a transfer function, but we only use the DC value reported by our FEA modeling. The simulation shows that this noise source is not be limiting even without the the low frequency common mode rejection of the suspension platform.

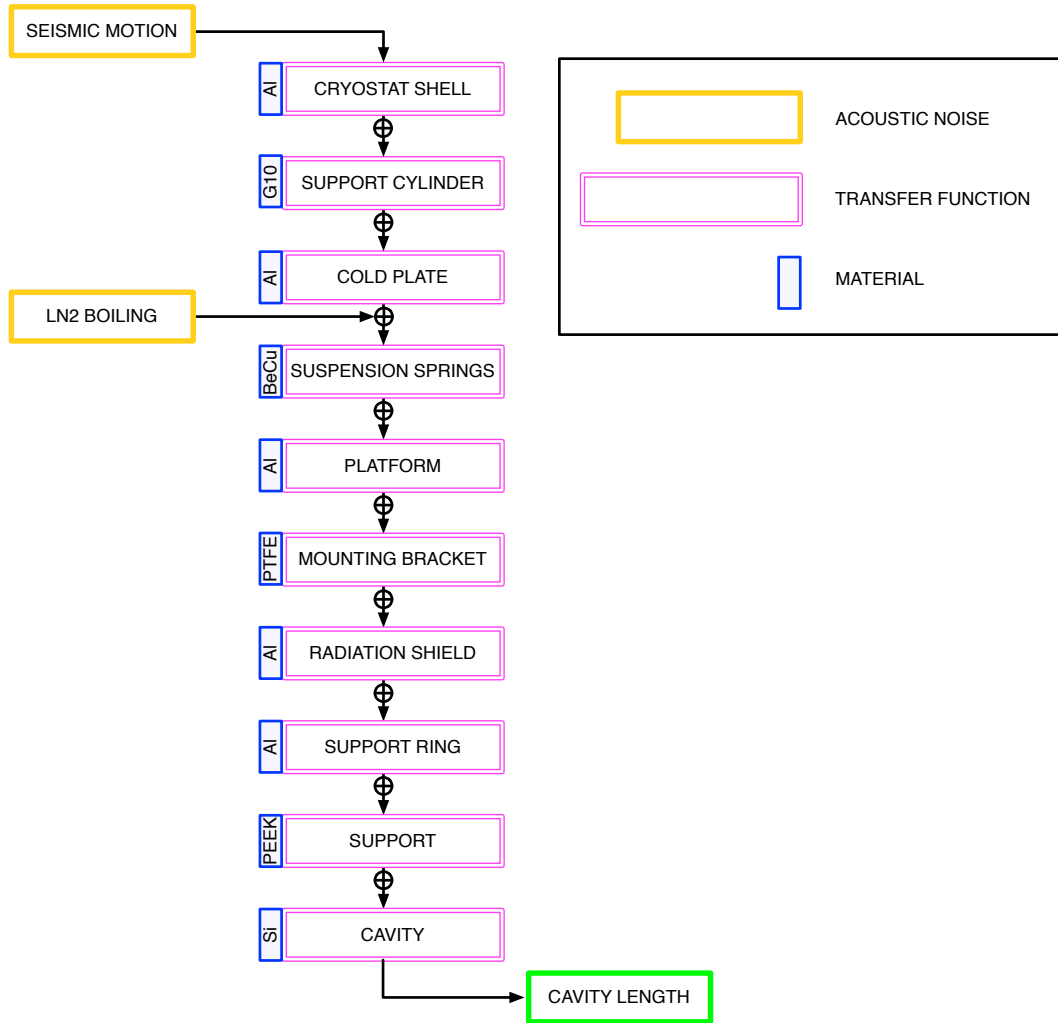


Figure 4.3: Here we show the mechanical propagation of the seismic noise from the top of the cryostat into cavity length. The pink blocks represent mechanical transfer functions, and the nodes connecting them represent the motion in some (multivariable) coordinate system.

By making measurements after the LN2 boiled off, but before the cavities warmed up, we were able to generate beat spectra with and without the boiling noise. We compare these spectra to the seismic noise measured at the top of the cryostat, propagated to the cavities using eq. (4.10). At low frequency, the measured seismic noise propagate through the model is consistent with the measured beat noise. We expect the rms motion to be higher while LN2 boiling noise is present, though we are not certain if this is the cause of disparity between the beat measurements shown here.

One noteworthy thing we did not take into account in eq. (4.10) was the violin modes of the suspension springs. These are difficult to model, and act as mechanical shorts between the cold plate

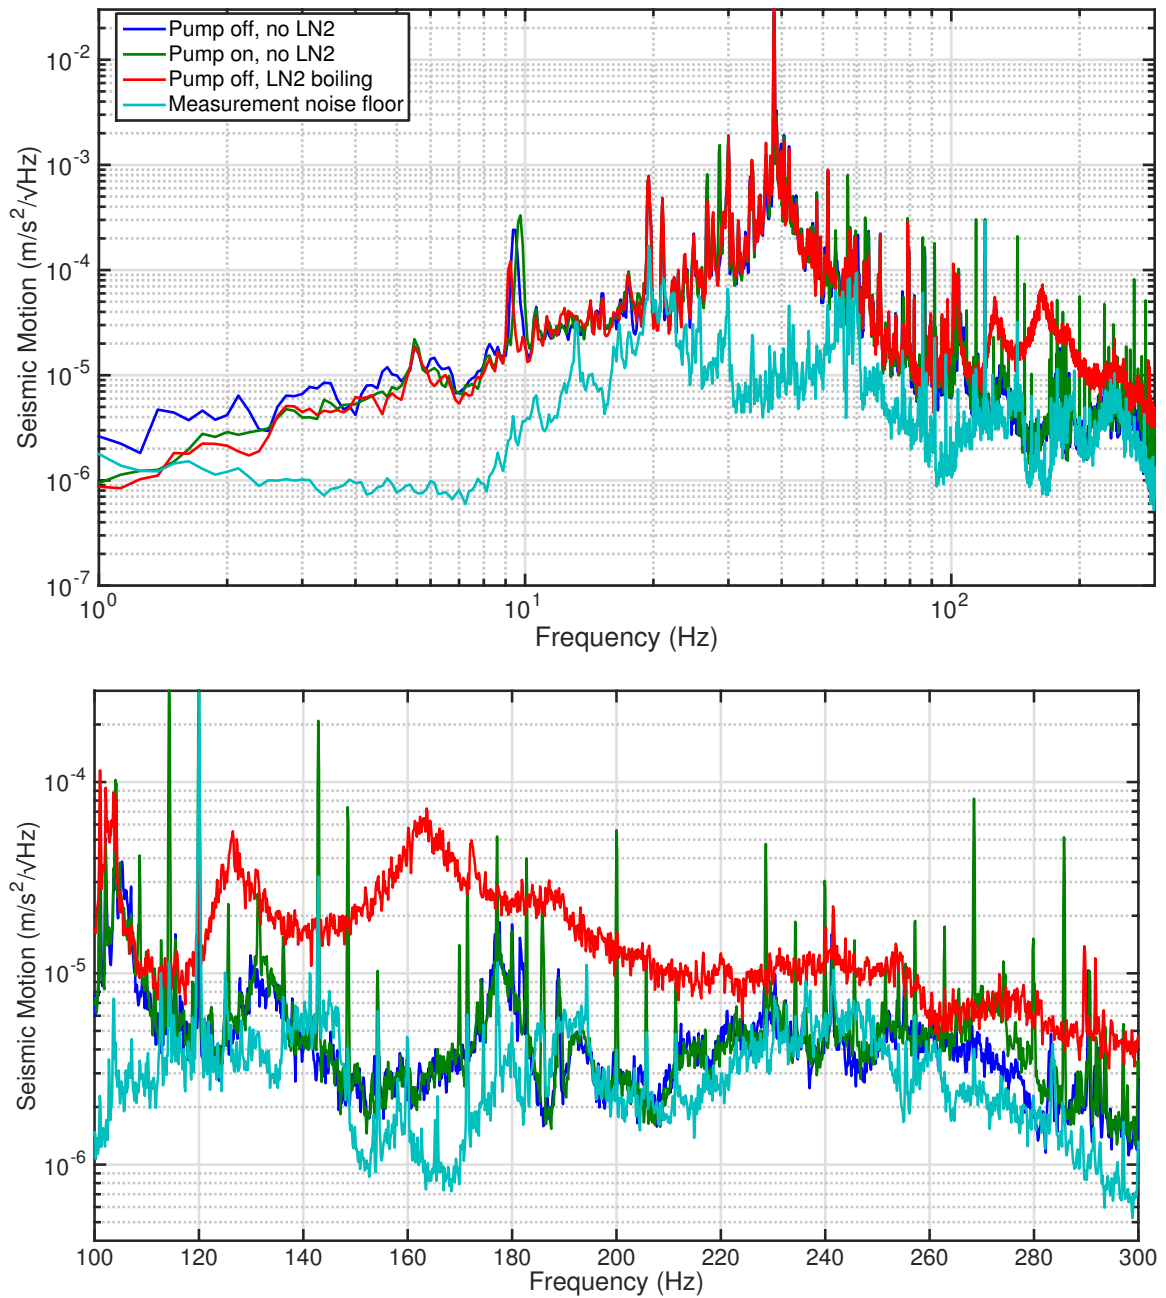


Figure 4.4: Vertical motion measured at the lid of the cryostat for the different states of the experiment. All low noise beat spectra were taken in the Pump off, no LN2 boiling state. The noise floor was determined by measuring the coherence between two Wilcoxon accelerometers on opposite sides of the cryostat lid. At most frequencies from 2 Hz to 300 Hz, we detected the acceleration of the cryostat lid. A zoomed in plot of the frequencies where the LN2 boiling noise increases is also shown here.

and the platform, spoiling the f^{-2} isolation normally seen above the bounce resonance. We believe that the forest of peaks we see from 400 Hz to 1 kHz in fig. 4.6 are the violin modes of the springs being excited by the LN2 boiling noise, and coupling into cavity length. Future work should consider measuring the transfer function through the spring suspension, and should pursue damping.




Figure 4.5: Shown is the results of COMSOL simulations determining sensitivity of the cavity length to vertical acceleration. On the left is a 3d render of the cavity bending about the support points. On the right is the simulation results, in kHz/g. Shortening of the cavity axis due to cavity sag is treated coherently with the elongation of the cavity axis due to mirror tilt. The coupling for this experiment was near the maximum shown above; we placed the cavity supports 2.5 inches from each other, and 60° below the horizontal. The frequency noise due to this mechanism coupling vertical motion to length was not limiting by over a factor of 10. For simulation results with more mounting positions and angles see appendix D. COMSOL simulation performed by Evan Hall. Left image courtesy Brian Dawes.

The above indicates that there is little benefit from reducing the seismic sensitivity of the experiment, and thus should be considered low priority. It should be noted that the choice of support points and materials also influences the Q of the cavity’s mechanical modes. See chapter 5 for more.

4.4 Intensity

We write the field at a given spatial location as $E(t) = A_0 e^{-i\omega t}$, and express the fluctuations in frequency and amplitude as:

$$E(t) = A_0(1 + \delta A(t)/A_0)e^{-i(\omega + \delta\omega(t))t}. \quad (4.11)$$

It is conventional to characterize fluctuations in the field amplitude as relative intensity noise (RIN), where $\text{RIN} = |\delta A(t)/A_0|^2$. This presents as fluctuations in photocurrent produced at the photodetector, normalized by the DC photocurrent. As mentioned in section 3.5, changes in laser




Figure 4.6: The estimated noise in the beat due to vertical seismic motion measured in fig. 4.4. We propagated the vertical motion to cavity length as described in eq. (4.10), using a very pessimistic estimate of 21 dB for the common mode rejection. Shown on the right y-axis is vertical acceleration at the cavity support points.

frequency are how we measure changes in cavity length, and thus anything which couples RIN and frequency noise interests us.

The RIN we measured in our setup is shown in fig. 4.10. The two questions we wish to ask at this point are “*does the RIN limit our experimental noise?*” and “*why is the RIN in transmission so high?*”. We investigated optical offsets at the PDH error point, net PDH offsets², photothermal absorption, radiation pressure, PLL RIN sensitivity, beam jitter, and polarization noise. We detail these investigations below.

In the experiment (e.g., after the light is generated) RIN couples into laser frequency at the PDH error point section 4.6, actual cavity length via intracavity power section 4.9, and the beat frequency due to PLL offsets section 4.6. These mechanisms for RIN coupling into the beat frequency are all shown in fig. 4.7, and are described in more detail in their respective sections. None of the linear mechanisms for RIN were able to explain the excess noise seen in fig. 4.1.

In the laser diode itself, laser current causes correlation between intensity and frequency, as it modulates both. As we are feeding back to laser current to stabilize the lasers, it is reasonable to ask whether or not locking the lasers to our cavities increases the RIN. The laser current has DC couplings of ≈ 30 MHz/mA and ≈ 0.05 RIN/mA (differed for each laser). The free running RIN

²It is important to distinguish between a net offset in the point at which we lock to the cavity and an optical offset. We can use an electronic offset to null an optical offset in the PDH loop to eliminate the conversion of frequency noise to intensity noise on transmission, but the intensity noise has already been converted to frequency noise at the error point — nulling the offset cannot undue this pollution of the error signal




Figure 4.7: Typical coupling levels of RIN transmission to frequency noise at the beat. At low frequencies, coupling due to optical offset at the PDH error point and absorption at the coatings dominates the transfer function. At high frequencies, the residual offset of the PLL loop dominates. Actual values of the transfer functions (and thus the net shape) vary over the course of the experiment. See fig. 4.1 to see the level of noise through these coupling mechanisms. See section 4.9 for more discussion on photothermal absorption.

of the laser was not increased when locked to the cavities, and we saw no coherence between input RIN and laser frequency. The correlation caused by feeding back to laser current is negligible.

Even though the known linear couplings of RIN to beat frequency are not currently limiting the noise budget, there is no reason to believe that whatever generates the excess RIN does not couple into the experimental noise via some independent path. It is thus important to investigate the source of the anomalously high RIN seen in transmission.

Laser frequency can couple into laser RIN via a net offset at the PDH error point. The linear coupling from frequency noise to transmitted RIN as a function of net offset is shown in fig. 4.8. We can zero out the offset electronically at the error point, but it will still drift over time. The dominant source of offset drift is changing RFAM levels due to the temperature dependence of the EOM's polarization axis. Using a very pessimistic estimate for offset (a 10% drop in output power), frequency noise in the experiment coupling into RIN via an offset with a level of 1.5×10^{-5} [RIN/Hz] is nowhere near high enough to account for the excess RIN. Additionally, coherence measurements seen in fig. 4.9 indicate that we start to see this coupling around 100 kHz, though it only accounts for 2.5% of the total transmitted RIN.




Figure 4.8: Shown here is the linear coupling of frequency noise to intensity noise due to an offset in the PDH loop. The top figure is cavity transmission, normalized to 100% coupling. The bottom figure is the ratio of the derivative of the cavity transmission to the cavity transmission.

We encountered two clues while investigating the source of the transmitted RIN. Small misalignments in the $\lambda/4$ waveplate (≈ 2 deg) in front of the cryostat and small changes to the gain in the

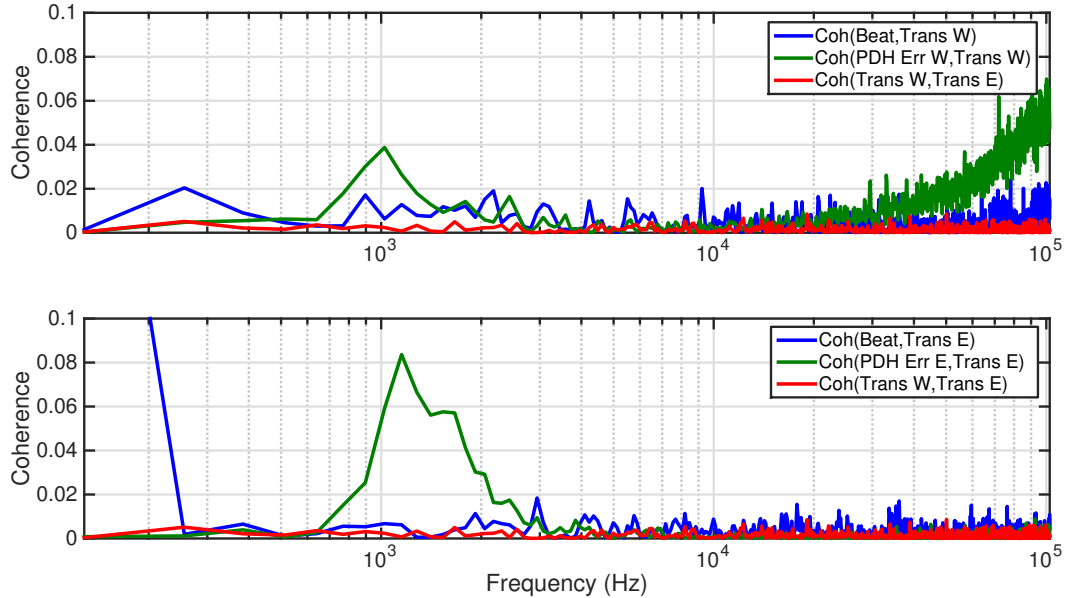


Figure 4.9: Coherence of RIN vs PDH error signals. The PDH error and control signals exhibit high coherence (not shown here) above 1 kHz, thus we can infer that the incident RIN level is not limited by the PDH control signal coupling into RIN via direct power modulation. The low coherence between the beat noise and RIN indicates that the beat is not limited by a linear mechanism which couples RIN to frequency noise. The slight increase seen in $\text{Coh}(\text{PDH Err W, Trans W})$ indicates a small net PDH offset, indicating that 2.5% of the transmitted RIN can be accounted for by a PDH offset in this path.

PDH loops ($\approx 15\%$) both caused the transmitted RIN level to change. The waveplate tuning sets the polarization of the retro-reflected beam at the PDH polarizing beamsplitter, which determines how much power is sent to the PDH PDs and how much is sent back towards the laser. Even though we had a faraday isolator in front of the laser, we observed a tendency to mode hop in the laser which could be replicated by scattering light back at it with an IR viewer card. The combination of these observations makes us suspect some laser feedback induced mechanism such as jitter or polarization. It is unlikely that these are the same mechanisms. The degree level changes of waveplate misalignment correspond to a power drop of 1% on the photodiode, which is ≈ 15 times smaller than the range of gain change needed to induce similar changes to the transmitted RIN spectrum.

Beam jitter couples into transmitted RIN via cavity alignment. The jitter to RIN coupling is a function of misalignment. We introduced a large intentional misalignment (corresponding to a 10% decrease in power), and did not see any increase in the transmitted RIN. Our normal (mis)alignment is within 2% of the maximum power. We did not see any difference in transmitted RIN between these two states, and conclude that linear coupling of beam jitter is not the source of the excess RIN.

Polarization fluctuations couple into RIN wherever we have polarization sensitive optics. We measure the RIN on the light with a pickoff immediately before it enters the cryostat, so the source

of polarization noise would have to be either the cryostat windows or the cavity itself. As the noise looks broadband at least until 100 kHz (see fig. 4.10) and featureless, this source seems very unlikely. For further confirmation, we could remove the cavities and see if the anomalous RIN still exists in transmission.




Figure 4.10: Transmitted RIN vs incident RIN. Neither incident nor transmitted RIN measurements were noise limited in the band shown. Additionally, the transmitted RIN measurements exhibited high coherence with the RIN measured at the beat detection RFPD. We do not know the reason for the anomalously high transmitted RIN.

In summary, we examined the following coupling mechanisms in our investigation of RIN:

- Beam jitter at cavity input \rightarrow RIN due to mode matching into the cavity,
- Laser frequency noise at cavity input \rightarrow RIN due to net PDH offset (fig. 4.8, fig. 4.9),
- Polarization noise at cavity input \rightarrow RIN due to polarizing optics,
- RIN at input \rightarrow cavity length via photothermal absorption and radiation pressure (fig. 4.7),
- RIN at input \rightarrow laser frequency noise via optical PDH offset (fig. 4.7),
- RIN at transmission \rightarrow beat frequency noise via PLL RIN sensitivity (fig. 4.7).

We do not know of any other linear mechanisms that could explain the RIN. Further investigation should include increasing the PDH bandwidth, and improving the laser isolation (as these might mitigate nonlinear couplings). Additionally, it would be good to add photodetectors to monitor the RIN in reflection.

4.5 Temperature

The requirements for precision and accuracy of cavity temperature are set by both the thermoelastic noise and residual linear coupling to cavity length.

Deviations from the zero crossing in $\alpha_{\text{Si}}(T)$ increase the thermal noise contributions of thermoelastic noise in the spacer and substrate as $\alpha\kappa^{1/2}$. The spacer thermoelastic noise curve shown in fig. 4.1 uses an incredibly pessimistic estimate for temperature offset of 10 K. This noise source is still very small compared to the other noise sources at the frequencies where the experiment is sensitive even with a huge temperature offset. Concerns about thermoelastic noise can be ignored when engineering the temperature stabilization.

An offset from the the zero crossing in $\alpha_{\text{Si}}(T)$ also gives rise to linear coupling between temperature and cavity length due to thermal expansion of the spacer. Length fluctuations due to temperature noise coupling through the residual linearly CTE will only be relevant at low frequencies. The spacer responds relatively fast to temperature changes around 123 K with a 5 minute time constant. The coupling between the radiation shield and the spacer, however, has a time constant of ≈ 4 hours. In fig. 4.11 we show the temperature noise of the radiation shield, use it to estimate the temperature noise of the cavity⁴, and then estimate the resultant length coupling due to thermal expansion for several temperature offsets from 123 K. At a temperature offset of 1 K from the 123 K zero crossing in $\alpha_{\text{Si}}(T)$, our model for the residual thermal coupling to cavity length from temperature fluctuations is above the coating Brownian noise.

With minimal effort (e.g., $\sigma T_{\text{cav}} < 1\text{K}$), we will not be limited by residual temperature fluctuations coupling in as thermal expansion based changes to spacer length, and the spacer thermoelastic noise will be negligible. The above is only a partial model: in order to truly say that we will not be limited by the coupling of temperature noise to length, we will need to make an accurate model of the changes to the coating due to thermal expansion at this temperature. This can (and should) be accomplished by direct calculation, which will be very similar to that done by Evans et al [48] when calculating the thermo-optic noise.

4.6 Electronics

Electronics noise couples into the experiment through the PDH servos used to lock the lasers to the cavities, the PLL used for the beat readout, and the temperature control loops used to stabilize the cavities. We discuss the PDH and PLL electronics below, and the temperature sensing noise in section 4.5.

³A simple pole is not quite correct — an accurate model of the transfer function can be obtained by following the work done in [57], with the actual silicon cavities replacing the dummy spacers used in that work

⁴Until we are in the high-Q mechanical mode regime, we may add a temperature sensor directly to the cavity to read it out without spoiling the noise in any way.




Figure 4.11: Temperature noise of the radiation shield, as read out by the temperature sensor shown in fig. 3.5. We can see the effect of the temperature servo below 2×10^{-3} Hz - above this the temperature is unstabilized. This temperature noise is then filtered by the thermal transfer function between the radiation shield and the cavity. We approximate this transfer function as a simple pole with a 2.5 hour time constant³, and apply that filter to the stabilized radiation shield temperature noise, generating a rough estimate of cavity temperature noise. Using the linearized eq. (2.1), we show the resultant frequency noise due to thermal expansion for several temperature offsets from the zero crossing around 123 K. With minimal effort this residual coupling will be negligible compared to any thermal noise source at frequencies where the experiment is sensitive ($f > 10\text{Hz}$).




Figure 4.12: Electronics noise budget. Shown are the main contributors to the electronics noise. The electronics noise is lower than the experimental noise seen in fig. 4.1, and higher than the thermal noises we are trying to measure; we must mitigate the electronics noise to facilitate detecting thermal noise in the experiment. A combination of redistributing the gain and upgrading the PDH servo and current driver will yield at least a 10x lower noise in the PDH loops. The PLL noise can be lowered by upgrading to a lower noise VCO. The dark noise of the beat frequency readout is too low to show up on this plot ($\approx 8 \times 10^{-8} \text{ rad}/\sqrt{\text{Hz}}$).

The PDH loops are shown in fig. 4.13. We used a pair of RIO⁵ PLANEX ([58]) lasers, controlled by a modified low noise driver (described in [59], with schematic and modifications shown in appendix E). The PLANEX is a 5 kHz line width fiber coupled diode laser with an on-chip Bragg grating. The current driver has low noise ($< 1 \text{ nA}/\sqrt{\text{Hz}}$) with high bandwidth (< 5 degrees of phase lag up to 1 MHz) and range (10 % modulation depth).

The PDH loops had a unity gain frequency (UGF) of 200 kHz, and a rough shape of f^{-1} everywhere, except from 60 Hz to 6 kHz, where they had a f^{-2} boost. The actual open loop transfer function is shown in fig. 4.14. The residual RMS of the error signal was up to 10% of the error signal's peak to peak voltage on sweep and was dominated by the noise from 100 kHz to 400 kHz (as shown in fig. 6.3). This is large, and indicates that we should improve the PDH loops in the future to improve the linearity of the system. The direct coupling of the residual error signal dominates the experimental noise above 1 kHz, and the coating thermal noise (the largest thermodynamic noise source) at all frequencies. In addition, there is some unknown nonlinear coupling of the residual error signal to the beat readout. When we increased the loop gain, the high frequency noise shifted as expected (more suppression in the loop bandwidth, more oscillation above the UGF), but the noise floor of the most sensitive part of the experiment (between 100 Hz and 1 kHz) increased. This indicates the system is nonlinear.

The main deviation from the shape described above is due to the frequency response of the laser actuator itself. We attribute the odd frequency response and phase delay seen in fig. 4.16 to the physical transfer function of [Hz/mA] inside the PLANEX laser. This is consistent with characterization done in [58]. The shape can be explained as the coherent sum of two paths for laser current to couple into laser frequency. Laser current changes the charge carrier density in the laser, which changes laser frequency. In addition, the total power deposited via resistive heating changes with current in the obvious way. When the thermal wavelength is large compared to the size of the laser chip, the temperature changes (and thus the frequency shifts) will be attenuated. At low frequencies, The [Hz/mA] transfer function is dominated by the charge carrier density path. At high frequencies, the temperature path dominates. The 180° phase change is consistent with these paths having opposite sign.

We show a measurement of the full open loop transfer function of the PDH loop in fig. 4.14, as well as the sum of measurements of its components. In figure fig. 4.15 we show a phase budget for the loop to reveal how much each component is limiting its unity gain frequency. We show the transfer function of the driver and laser in fig. 4.16. For transfer functions of other components, see appendix E.

The PDH loops should be improved by lowering their input referred noise, and increasing the loop bandwidth (to increase their linearity). The dominant source of noise at the error point is the

⁵Redfern Integrated Optics

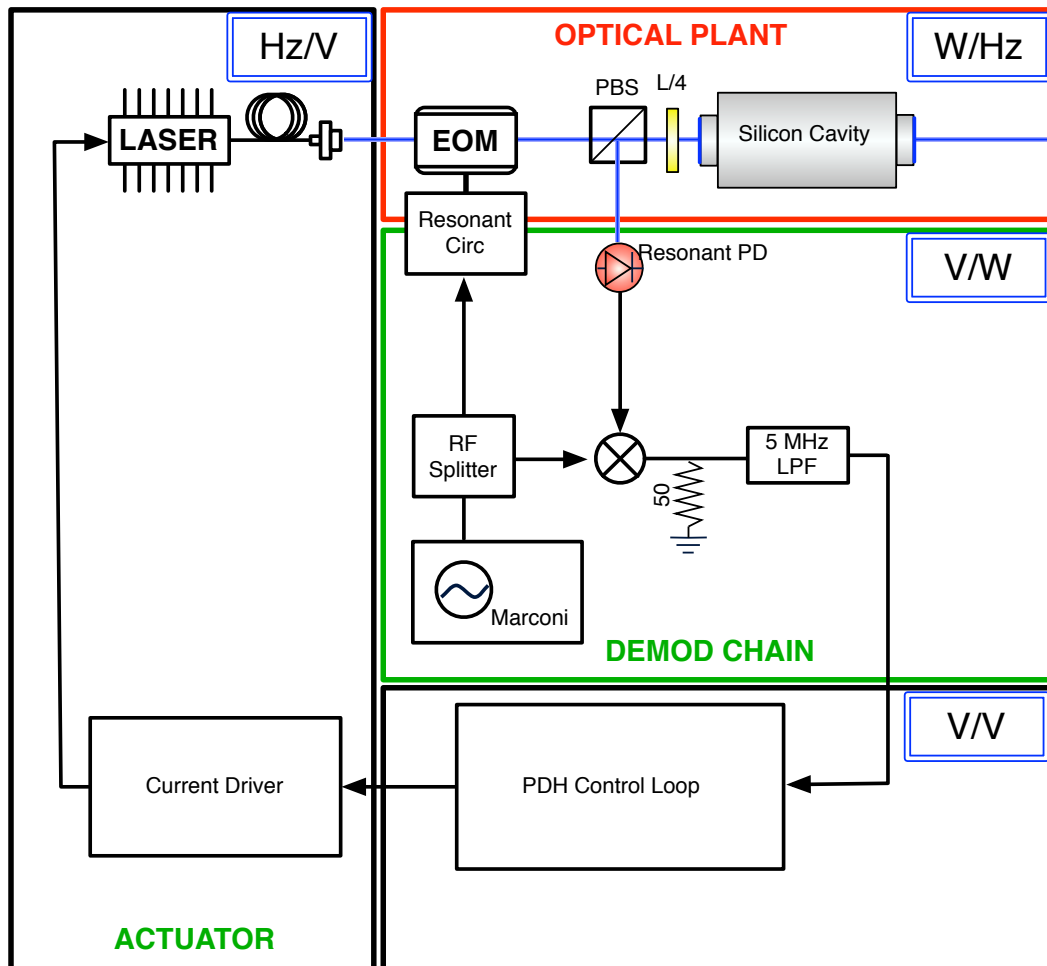


Figure 4.13: Shown here is a simplified schematic of the PDH loops used to lock the lasers to the cavities. The transfer function of the ACTUATOR block is shown in fig. 4.16. The open loop transfer function is shown in fig. 4.14. Future versions of this servo will use a fast phase correction path via an additional Pockels cell to extend the bandwidth.




Figure 4.14: Shown here is the measured open loop transfer function taken during the measurements which produced the noise curve shown in fig. 4.1. The phase budget (which determines how high we can push our loop bandwidth while maintaining loop stability) is shown in fig. 4.15.




Figure 4.15: Here we show the phase budget of the PDH loop, broken into components. The largest offenders of missing phase are the servo box, the PDH low pass filter, and the laser actuator itself. None of these components are fundamentally limiting, and can all be upgraded to decrease their phase delay.




Figure 4.16: The open loop transfer function from the input of the laser current driver to the frequency of light at the output of the laser. We locked one laser to the other with a 10 kHz UGF PLL, and measured at the error point of the PLL, taking care to ensure proper 50 ohm impedance matching and calibration.

input noise of the LB1005 driver ($40 \text{ nV}/\sqrt{\text{Hz}}$). Replacing this with something that has a $10 \text{ nV}/\sqrt{\text{Hz}}$ input noise should be considered high priority. The PDH noise can be further improved by shifting more gain from the current driver to the feedback servo and improving the locking, as we used significantly more range to acquire lock than we did to maintain lock. A sensible set of improvements to make would be to upgrade the laser temperature controller from a Thorlabs⁶ ITC510 to allow for slower scanning through laser resonance, and use the lower range requirements to move gain from the current driver to the PDH control servo.

Further suppressing the PDH residual should also be considered high priority to combat possible nonlinear couplings mentioned above. To increase the bandwidth of the PDH loops, we need a faster actuator. We can use a Pockels cell as a fast phase modulator, and use that in conjunction with laser current for high bandwidth PDH actuation [60].

The PLL (Phase-locked loop) we used to readout the beat between the cavities is shown in fig. 3.10. The limiting noise source is normally the input referred noise of the voltage controlled oscillator we lock to the beat frequency. We used an ifr 2032 as our VCO. The PLL noise floor was set by the VCO. The beat was not limited by the noise of the PLL, though the electronics noise (see fig. 4.12) was limited by this around 3 kHz. This level was set by the low frequency cavity length drift, which limited the ifr function generator input range, which in turn sets the noise floor of the VCO. In order to detect coating thermal noise, we need a PLL with 5 kHz of range above 10 Hz, and the ability to follow the DC drift of the cavities below this. This can be accomplished by embedding a low noise low bandwidth high range PLL inside a low noise, high bandwidth PLL using a custom VCO (such as this one [61] used in advanced LIGO). Alternately, we could change the beat readout method from a PLL an electrical delay based frequency discriminator as described in [32].

It is also in general desirable to increase the bandwidth of the PLL, as that decreases the residual phase fluctuations at the PLL's error point, which in turn makes the readout more linear. The optimal unity gain for this generation of the experiment was 10 kHz; increasing the UGF past this increased the residual phase fluctuations at the error point since the PDH residuals are dominated by frequency noise around 100 kHz. Minimizing the PLL residual can be accomplished by simultaneously upgrading the PLL UGF to very high frequency (1 MHz), and further improving the PDH loops as discussed above, as they were the dominant contribution to phase noise in the beat above 10 kHz. One often considered nonlinearity in a PLL readout is the coupling of AM to phase noise. This is not a significant noise source for us until $\approx 2 \text{ kHz}$, where the coupling in $[\text{Hz}/\text{RIN}]$, rising as f^1 , overtakes photothermal absorption (section 4.9). The potential gains of decreasing the phase residual should be examined in more depth before a serious effort to upgrade the bandwidth is pursued.

⁶Thorlabs




Figure 4.17: Shown here is the calibration for the control point of the PLL (the signal sent to the VCO). In band of the PLL, this is the lowest noise readout for the beat frequency we have. At some point above the UGF, the error point becomes the most sensitive sensor, which is calibrated in radians rather than Hz.

4.7 Quantum Fluctuations

The quantum noise of the experiment is shown in fig. 4.18. The software package Optickle⁷ was used to simulate the optical cavity layout of the experiment. Quantum vacuum fluctuations were injected into the open optical ports of the experiment, and allowed to interact with the cavities and loops.




Figure 4.18: Quantum vacuum fluctuations coupling into experimental frequency noise. Shot noise limiting the PDH photodetection is shown as REFL I. The strong peak in the PDH contributions is due to optomechanical coupling of the cavity body mode. The beat is a phase sensitive measurement, and thus the quantum vacuum noise is not significant until higher frequencies.

4.8 RFAM

Experimental techniques using a phase modulator to add phase sidebands to an optical field can be limited by residual amplitude fluctuations of the optical field being phase modulated. This problem was encountered in FM spectroscopy [62, 63, 64], and a solution was proposed by Wong and Hall [65]. This solution is now used to stabilize the RFAM encountered in ultra-stable cavities [7, 28]. We do not stabilize the RFAM in this experiment, and describe the coupling of RF amplitude modulation (RFAM) for this work in more detail below.

RFAM at the frequency used to generate the PDH sidebands is indistinguishable from actual error point fluctuations (due to changes in laser frequency or cavity length). This can be caused

⁷<https://github.com/Optickle/Optickle>

by scatter-sourced parasitic etalon effects, vibration of anything in the optical path, amplitude fluctuations of the laser field at the phase modulation frequency, and temperature driven drift in the phase modulator's birefringence. We discuss the latter source below, as we believe it to be the dominant effect.

The EOM used to generate the PDH phase sidebands also generates some amount of amplitude modulation:

$$E_0 e^{i(2\pi\nu_{opt}t)} \rightarrow E_0 [1 + \Gamma_{AM} \cos(2\pi f_{mod}t + \delta)] e^{i(2\pi\nu_{opt}t) + \Gamma_{PM} \cos(2\pi f_{mod}t)}, \quad (4.12)$$

where E_0 is the magnitude of the electric field, ν_{opt} is the optical frequency, f_{mod} is the modulation frequency, Γ_{AM} is the depth of amplitude modulation, and Γ_{PM} is the depth of phase modulation. For an ideal phase modulation process, $\Gamma_{PM}/\Gamma_{AM} \rightarrow \infty$. The primary cause of AM in this work was a mismatch between the EOM crystal axis orientation and the polarization axis of the electric field. We can tune the axes via waveplates such that $\Gamma_{AM} = 0$, but the EOM crystal axis orientation is a function of temperature, and it thus drifts out of alignment.

The RFAM gets demodulated along with the PDH signal, and the portion which is in phase with the PDH demodulation signal shows up as an optical offset to the PDH error point. The optical offset determines the coupling from intensity noise to frequency noise at the PDH error point (which ideally is an amplitude insensitive measurement). Correcting this offset electronically (say, by adding in a low noise offset at the PDH servo) can null the cavity offset, and thus the transfer function from frequency noise into intensity noise, but it cannot undo the addition of noise which occurs during photodetection. Additionally, changes to the RFAM level or the asrelative phase between the RFAM and the PDH demodulation frequency both look like frequency noise and are corrected for as such by the loop.

For this work, we used a broadband LiNbO₃ EOM supplied by Thorlabs, to which we attached a resonant circuit for the purpose of achieving greater modulation depth. We achieved a stable modulation depth of $\Gamma_{PM} = 0.45$. Due to the above mentioned temperature drifts, Γ_{AM} would drift up to a value of $\Gamma_{AM} \approx 0.02$ over the course of 10 minutes after minimizing the RFAM (which is proportional to Γ_{AM}^2), which corresponds to an error point offset of up to 10 Hz (depending on the relative phase of the RFAM and the demodulation signal).

Shown in fig. 4.19 is a pessimistic estimate of the frequency noise due to drifting RFAM (measured using a lock-in amplifier) which assumes that the RFAM signal is perfectly in phase with the PDH demodulation frequency. This could be easily improved by stabilizing the temperature of the EOMs (either via active or passive techniques). This is larger than the coating thermal noise starting below 3 Hz, where this experiment is not sensitive.

As shown in fig. 4.7, the transfer function of intensity noise to frequency noise for the typical PDH

offset is comparable to the level of the photothermal transfer function from 10 Hz to 10 kHz. For the purpose of measuring the photothermal transfer function, stabilizing the RFAM would enable more accurate measurements over a larger frequency range, which allows us to measure the absorption loss in the cavity more accurately (see fig. 4.25). If the RFAM was static, this would not be important, and we could just fit for the PDH coupling to get to the optical loss. Drift in the RFAM over the course of a (long) measurement adds hard to characterize systematic errors to the photothermal transfer function measurement. Before trying advanced RFAM stabilization techniques, we should attempt passive thermal isolation of the EOM crystal housing.




Figure 4.19: A pessimistic estimate of the frequency noise at the beat due to drifting RFAM. Here we assume that the RFAM is perfectly in phase with the PDH demodulation.

4.9 Photothermal Response

Optical absorption in the mirror coatings provides another path for intensity fluctuations to coupling into cavity length. The effect of photothermal absorption on the phase profile of a gaussian beam sampling the surface of a mirror has been studied theoretically by Braginsky et al [66] and [67] et al. These equations have been experimentally verified [68, 69, 70]. Farsi et al generated a numerical model of the effect at room and cryogenic temperatures [69], and performed experiments measuring the photothermal effect for $\text{SiO}_2:\text{Ta}_2\text{O}_5$ on a Si substrate using 1064 nm light. Chalermongsak et al [70] used this model (correcting for CTE of the coating) to suppress photothermal noise in $\text{GaAs}/\text{Al}_{0.92}\text{Ga}_{0.08}\text{As}$ Bragg mirrors via coherent cancellation. The details of the calculation in the

Farsi paper are identical to those used to calculate the coating thermo-optic noise by Evans et al [48] if you substitute absorption of laser intensity for statistical fluctuations in temperature

There are a number of sources for intracavity power fluctuations. In principle, the only one which cannot be directly suppressed via engineering is that of shot noise. Increasing the power allows us to gain SNR as \sqrt{N} in the phase detection parts of the experiment as the signal we are trying to detect grows faster than the shot noise. The transfer function in $[\text{Hz}/\text{W}_{\text{intracavity}}]$ is fixed, so increasing the power increases actual length fluctuations in the cavity. The bound on power set by shot noise is well above any power level we considered for the experiment.

We used the numerical modeling done presented by Farsi et al [69], with the modifications for coating expansion in [70] to model our system. We used the temperature dependent material property values for silicon shown in chapter 2, and room temperature coatings properties from [48]. Shown in fig. 4.20 is the modeled photo thermal transfer function, parametrized as a function of temperature. The portion of the model which is most strongly temperature dependent is the thermal expansion of the silicon substrate, shown separately in fig. 4.21. The individual traces and their coherent sums are shown together in fig. 4.22.




Figure 4.20: Modeled temperature dependence of the photo thermal transfer function.

We made a direct measurement of the photo thermal response in one of the cavities. As shown in fig. 4.23, we used an AOM to drive a swept sine in power, and measured the transfer function of transmitted power to beat frequency. The transfer function is the coherent sum of all effects which couple intensity and beat frequency noise as read out by the PLL. These are as follows:




Figure 4.21: Temperature dependence of the portion of the photo thermal transfer function which is due to the substrate expansion. As expected, this goes to zero as we approach the zero crossing in α .




Figure 4.22: The individual effects which are coherently treated in the model are shown here. At 300K, the substrate expansion dominates, but as we approach 123K, the substrate expansion goes to zero, and the coating expansion dominates.

Optical PDH offset, PDH cavity offset, PLL RIN sensitivity, and photo thermal absorption driven length changes. The optical PDH offset occurs due to RFAM at the PDH error point (discussed in section 4.8); driving the intensity of light puts an amplitude modulation envelope on the RFAM, which is indistinguishable from cavity length fluctuation. This couples into the transfer function with a slope of f^0 . The net PDH cavity offset gives rise to Hz \rightarrow RIN coupling based the slope of cavity transmission (zero for perfect resonance). Amplitude noise at the PLL phase detection gives a RIN \rightarrow Hz coupling based on the PLL offset (zero for a perfect PLL). This couples into the transfer function with a slope of f^1 . Amplitude changes inside the cavity coupled with absorption in the mirrors yield the above mentioned photo thermal coupling of RIN \rightarrow Hz. The shape is as shown in fig. 4.20.

Shown in fig. 4.24 are the results of fitting the photo thermal model described above with only absorption in ppm as a free parameter, and the optical PDH offset. The cavity offset frequency to RIN coupling is insignificant compared to the measured level, and the PLL RIN sensitivity, with its slope of f^1 , is not large until higher frequencies. Deviations seen in phase by 1 Hz are likely the effect of the cavity spacer. The measurement could be improved by stabilizing the RFAM, as drifting RFAM over the time scale of the measurement changes the amplitude of the optical offset coupling.

We attempted to measure the photothermal transfer function at 123 K, but were unable to see anything except the PDH offset. To make a lower noise measurement, we let all the LN₂ boil off, slowly warmed the cryostat up, and used the heaters to pause at various temperatures. We were unable to see anything but the PDH offset at 114K as well. Due to a partial vacuum failure, there were no measurements possible between 123 K and 233 K. The absorption loss in ppm vs temperature is plotted in fig. 4.25.

Though Farsi et al [69] have already modeled and made the first measurements of the photothermal transfer function at cryogenic temperatures for SiO₂:Ta₂O₅ coatings on a Silicon substrate, that measurement was performed with 1064 nm light. The gravitational wave community is interested in using 1550 nm light with silicon optics in future generations of gravitational wave interferometers [8]. There is still considerable scientific merit in quantifying the absorption of SiO₂:Ta₂O₅ and crystalline coatings (as done by Chalermsongsak [70]) at 1550 nm using measurements of the photothermal absorption.

A high resolution measurement of optical loss versus temperature based on a temperature dependent photothermal absorption model made using a high finesse cavity would yield insight into where any change in absorption loss happens (for example, due to the mirrors acting as a cryopump and gathering contaminants on the mirror surface, increasing the scatter and absorption). By simultaneously making a high precision measurement of the cavity pole, one can separate the optical loss due to scatter and optical loss due to absorption. Measurements of the cavity pole we made using

the technique shown in fig. 2.8 at various temperatures from 123 K to 300 K showed no significant temperatures dependence.

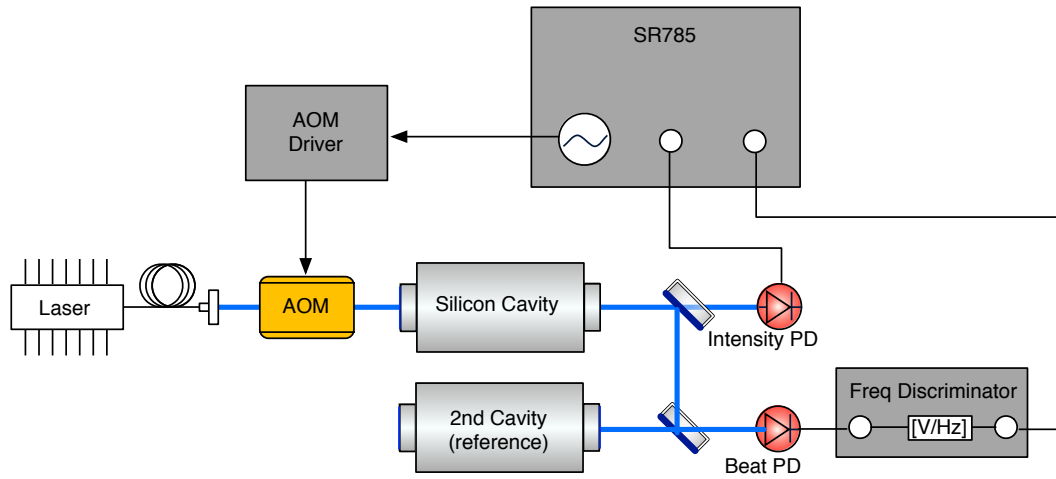


Figure 4.23: Shown is the setup used to measure the cavity's photothermal response. We used an AOM to drive the intensity with 30% modulation depth, and made a swept sine measurement of the beat frequency, as detected by the PLL control signal, and the transmitted intensity.




Figure 4.24: Photothermal transfer function with data, and two free parameter fit. Systematic errors in the fit dominate, as can be clearly seen in the residuals. Error bars were pessimistically estimated based on systematic errors in the fit.




Figure 4.25: Measured absorption loss in ppm vs temperature.

Chapter 5

Towards Macroscopic Quantum Cavity Opto-Mechanics

In addition to setting the level of Brownian thermal noise in the spacer, the quality factor of the cavity's mechanical modes determines whether or not we can explore their quantum behavior.

When viewed as a harmonic oscillator, a mechanical mode with frequency f_m has a thermal occupation number of $n_m = k_B T / \hbar 2\pi f_m$. As $T \rightarrow 0$, $n_m \rightarrow 0$ (as we extract energy from the system, it approaches its ground state). Once an oscillator is sufficiently close to the ground state, we can explore quantum behavior in the system.

There has been significant effort and progress made by the optomechanics community to probe oscillators in their ground state. As it is not feasible to cool an object to absolute zero, the cavity optomechanics community pursues other means [71, 72] of extracting energy from a particular mechanical mode. Some groups have successfully cooled mechanical oscillators into the ground state [73, 74]. There are optomechanics experiments at many length scales now, shown in fig. 5.1. We wish to pursue cavity optomechanics at the kg scale with this experiment (see chapter 8 for a list of experiments we could perform with macroscopic cavities cooled to the ground state).

A high Q silicon Fabry-Perot cavity with an accompanying optical readout would allow us to explore quantum mechanical behavior of a kg scale system, potentially including the ability to measure the evolution of the cavity hamiltonian in realtime. The first step in working towards this goal is to support a high Q body mode in the silicon cavities. In this section we explore:

- Measurement of the body mode ringing down at 300 K and 123 K;
- Most likely culprit for current limitation on Q factor;
- Recommended solution for increasing the Q, with accompanying simulation.

We used an electrostatic driver (ESD), shown in fig. 5.2, to excite the body modes of the silicon cavity. The spacing between the capacitive strips of the ESD and the silicon cavity was roughly

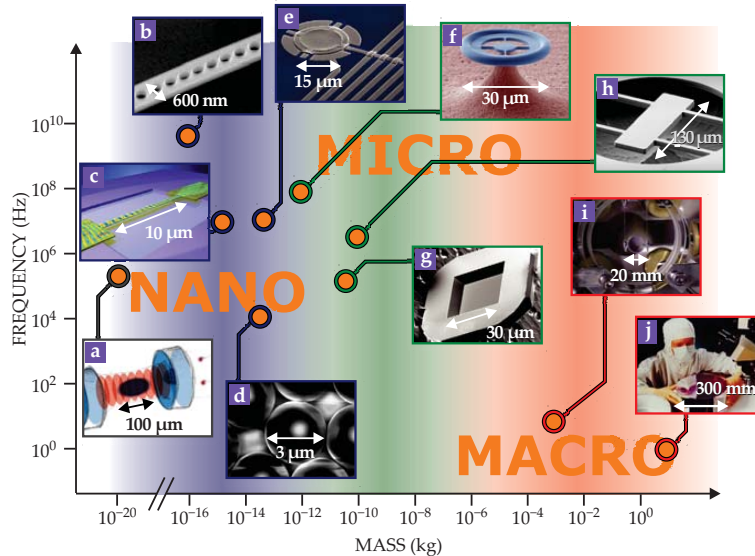


Figure 5.1: The lay of the land in the field of cavity optomechanics, plotted by frequency and mass. Figure reproduced from [75].

equidistant, as discussed in ([76]). We used a 100 V audio amplifier to drive the ESD and ring up one of the cavities. We fed the differential cavity length signal (the output of the beat frequency discriminator) into a lock in amplifier, and demodulated at the ESD drive frequency to measure the body mode in real time. In order locate the body mode, we added large audio sidebands at $\pm 10 Hz$ of the frequency used to drive the ESD, slowly scanning modulation frequency, and watched the beat spectrum on an SR785 spectrum analyzer. A blind search was not overly time consuming; it took ≈ 4 hours to scan over 2 kHz. The measurement set up is shown in fig. 5.3. We measured the body mode Q by fitting the a simple time constant to the ring down after shutting off the drive signal. We fit to the peaks of the absolute value of the ring-down voltage, as shown in fig. 5.4. The results are shown in table 5.1

	f [Hz]	Q
300 K	34756	2×10^4
123 K	34909	7×10^3

Table 5.1: Silicon drum mode quality factor. The shift in frequency is due to changing elastic properties of silicon at cryogenic temperatures. COMSOL simulations indicate this is the drum mode (± 500 Hz). This peak was the largest we found while scanning the drive frequency, which is unsurprising as the TEM₀₀ mode used to sample the mirror surface is very sensitive to these oscillations.

The measured Q (seen in fig. 5.4) of 7×10^3 was four orders of magnitude lower than values obtained in bulk measurements shown in fig. 2.2. We began our investigation with the most likely culprits: the PEEK plastic supports on which the cavity rests.




Figure 5.2: Close up view of ESD used to excite the mechanical modes of the cavity and their attachment to the radiation shield caps. The bore hole in the caps shown on the left is collinear with the optic axis. Separation between the ESD and the cavity was roughly the 1cm (the same distance between the ESD's capacitive plates).

Ringdown measurement

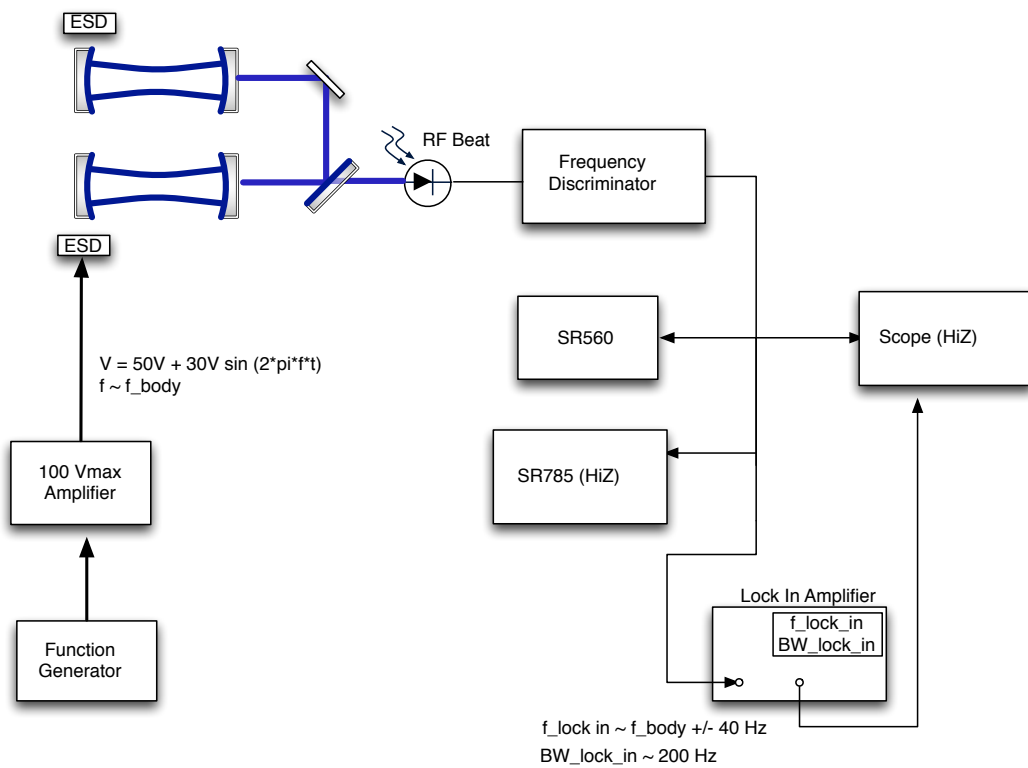


Figure 5.3: Diagram of ring-down measurement setup.




Figure 5.4: Shown here are ring-down measurements done to determine the Q of the cavity's drum (stretching) mode. $Q(300\text{K}) = 2 \times 10^4$. $Q(123\text{K}) = 7 \times 10^3$. In the absence of some additional loss mechanism, we expect this to be as high as 3×10^8 (see fig. 2.2). We believe that loss through the support points (shown in fig. 3.2) is the reason the Q is so low at both temperatures.




Figure 5.5: Shown here are transfer function measurements of the ESD drive and intracavity power into cavity length as measured at the beat, respectively. As a length actuator, the ESD had 600x the range of intracavity power with a maximum drive amplitude of ≈ 12 pm. This scales with the Q of the body mode.

We used a finite element model (COMSOL) to determine the eigenmodes of the cavity supports. First we simulated the eigenmodes of the PEEK supports (shown in fig. 5.6). The proximity of one of the support eigenmodes to the measured body mode frequency table 5.1 is concerning. The mechanical loss of PEEK is shown in fig. 5.7, and is slightly worse at 123 K than at 300 K for dry PEEK (which is what our supports consist of). PEEK's Q of ≈ 100 corresponds to a width of $\approx 360\text{Hz}$ at the mode shown in fig. 5.6. The $\approx 1.4\times$ decrease in PEEK's Q at 123 K is consistent with us measuring a lower body mode Q at 123 K. Additionally, the body mode frequency at 123 K is closer to this eigenmode of the PEEK support, which is also consistent with measuring a lower body mode Q at 123 K.

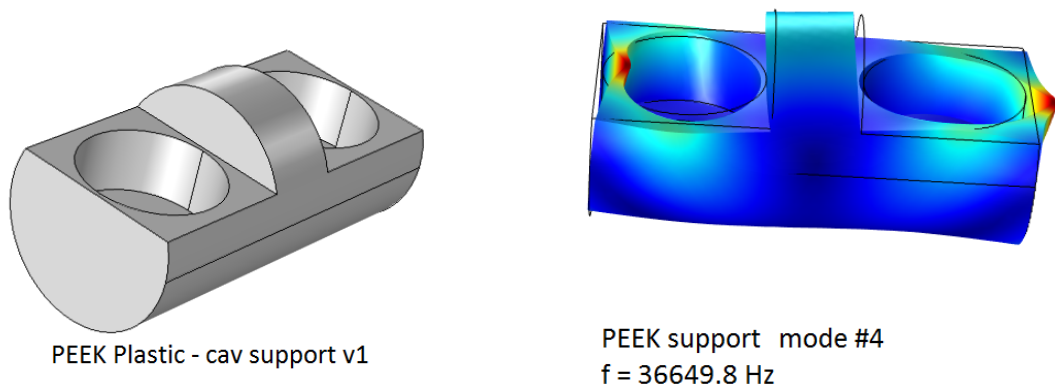


Figure 5.6: COMSOL simulation showing one of the eigenmodes for the PEEK support used to hold the cavities. The proximity between the mechanical response of the support and the measured eigenfrequency of the cavity is unintentional, and can help facilitate extraction of energy from the cavity. We consider this evidence that we need to redesign the supports.

While there is not enough evidence to conclude that the resonant mode of PEEK is enhancing the extraction of energy at the cavity body modes without additional simulation, the proximity of the body mode and support mode should be considered bad design for the purposes of supporting a high Q drum mode. We believe that further simulation of this generation of support would be a waste of time, and simulation of an upgraded suspension (shown in fig. 5.8) should be pursued instead.

It is helpful at this point to remember that we are interested in the vibrations of the cavity at a particular frequency due to thermodynamical fluctuations, which is precisely what the Levin formulation of the FTD [49] calculates. We go over the recipe here, and let this motivate our preliminary heuristic support redesign.

1. Determine the variable we are interested in the fluctuations of (the position of the cavity mirror, with gaussian weighting determined by the beam incident on the mirror).




Figure 5.7: PEEK loss versus temperature. PEEK used in this experiment was dry, which has $1.4 \times$ higher loss at 123 K than it has at 300 K. Figure reproduced from chapter 2 of Polymers at Cryogenic Temperatures [77].




Figure 5.8: Proposed upgrade to cavity supports. The dimensions (width, thickness, height) shown were varied for the eigenfrequency analysis in fig. 5.9 and fig. 5.10. New supports to be made out of Al 6061.

2. Drive the variable conjugate to the one we wish to measure (a gaussian pressure profile).
3. Determine the resultant motion of the mechanical system (defining what the system includes is the tricky part. Including every object mechanically connected to the experiment would include the entire planet and is computationally ridiculous to consider; including only the cavity ignores the fact that energy leaks out through the support points, and gives us the wrong answer).
4. Calculate the total energy lost per cycle (for a COMSOL FEA model, this is done by taking volumetric averages of strain energy density in each part simulated, and multiplying by the loss in each material).
5. Plug this energy lost per cycle into the Fluctuation-Dissipation Theorem.

Examining the third item in the above recipe, we see that we want to minimize the energy per cycle lost in the supports (and in mechanical connections further separated from the cavity). We can do this in a few ways:

- Choose a support point with a smaller coupling from mirror motion to support point at f_m (either by moving the support points closer to the center of the drum mode, or by choosing a higher order mechanical mode which still couples into cavity length and supporting at minima of that transfer function. We pursue the former).
- Alter the support structure to minimize the transfer function through the support, isolating the cavity motion from however we choose to attach the support to the cryostat. We pursue this below. As commonly expressed in the optomechanics community, we wish to maximize the phononic bandgap around f_m in order to localize the cavity's mechanical mode¹.
- Use a material with high Q for the support. It is easy to win big here as PEEK is a terrible material for supporting a high Q body mode. We recommend something like Al 6061 or BeCu.

If we decrease the separation of the support points from 3 inches to 0.5 inches (the closest we can make them using the current support structure), we decrease the magnitude of average mirror motion into support strain by a factor of 6 at f_m . By replacing the lossy PEEK supports with a metal or alloy that has a high Q at 123 K (such as Al or BeCu), we can further decrease the support dependent limitation on body mode Q by the ratio of Q factors. Assuming a cryogenic support material Q of 1×10^7 , we can suppress the loss in the supports by five more orders of magnitude.

¹The optomechanics community tends to use phononic bandgap (as described in [78] as their figure of merit when analyzing mechanical modes and the limitations to their mechanical quality factor due to loss of energy to the support structure. The gravitational wave community tends to approach the problem from the standpoint of transfer functions and mechanical losses as they use the Levin recipe for calculating frequency dependent loss and thermodynamical noise. These are both useful approaches to analyze mechanical systems, though the language of transfer functions makes design of a mode isolating support very intuitive.

This combined improvement of 6×10^5 means that loss in the supports themselves would not limit the body mode Q until it reached $\approx 4 \times 10^9$.

If we just made the changes outlined above, we would not achieve a body mode Q limited by mechanical loss in the silicon itself; the Q would be limited by loss in the structure which houses the support, the screws which hold them together, and friction in the joint they make. We want to minimize the coupling of vibration at the top of the support (where it touches the cavity) to the parts of the support which touch anything else.

The support shown in fig. 5.8 was designed with exactly this in mind. The inverted pendulum response of the new support is a function of the position, width, and depth of the cuts shown. Heuristically, a structure which has an inverted pendulum mode well below the cavity body mode frequency, and does not have any mechanical modes its vicinity should achieve good isolation. That is to say, the support will not transfer energy into the outlying structure, and loss in the outlying support structure will not limit the Q of the cavity body mode.

We used COMSOL to and an eigenfrequency analysis of the new support, under load, for different geometries. The mode shapes are shown in fig. 5.9, with Mode 1 being the pendulum mode of interest. The results of our COMSOL eigenfrequency simulations are shown in fig. 5.10. For all geometries simulated, the first mode remained relatively stationary, and the frequency band around the body mode (35 kHz) remained relatively clear.

Another question we posed is "how skinny is too skinny for the cavity support width parameter?" To address this, we estimated the maximum force the support will experience to be $2g$, and performed an analysis of the von Mises stress [79] which was concentrated in the skinniest part of the support fig. 5.11. For the flimsiest geometry considered (e.g., the tallest and skinniest) the von Mises stress was more than a factor of a few less than the yield strain.

The above should be considered a preliminary study. Before choosing a final geometry for fabrication, future work needs to include a full COMSOL simulation of the body mode Q based on the Levin method as outlined above, where we optimize over the parameters shown in fig. 5.8. Due to the breadth of geometric length scales, this might require significant computational power. For material, we recommend using Al 5056 for its high Q factor of 4×10^7 below 20 K [80], though any easily machinable metal with a high yield strain and high Q would be fine. If possible, we want to use a material with a lower thermal conductivity in order to not lower the time constant between the radiation shield and cavity too much (though this is not crucial, as we can solve the thermal engineering problems other ways).

A figure of merit central to the field is the fQ product. This in part motivates the push towards high Q oscillators [81]. The number of cycles it takes for an oscillator in the ground state to thermally decohere (e.g. the system's occupation number increases by roughly one due to noise entering from




Figure 5.9: Eigenmodes of proposed support geometry under load.

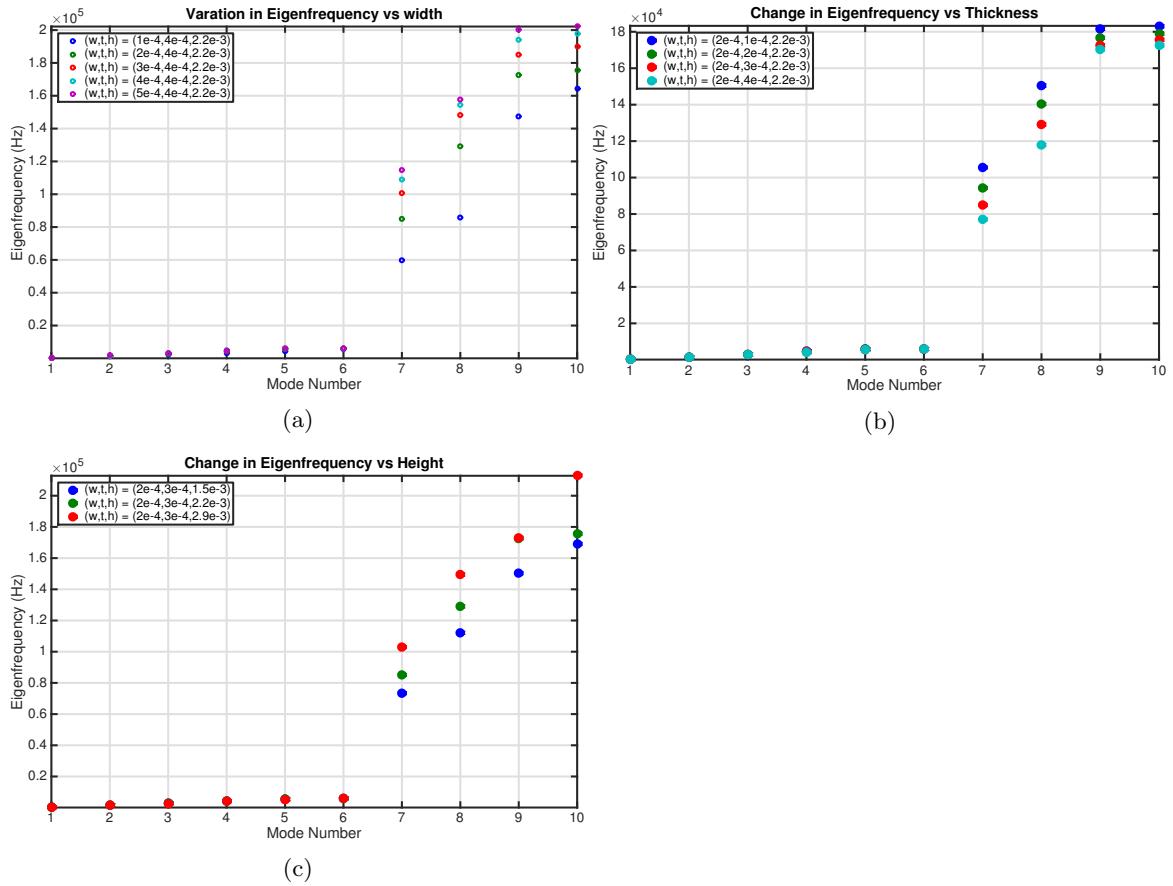


Figure 5.10: Change in eigenfrequency of the first 10 support eigenmodes as a function of geometry. For all geometric parameters simulated, the frequency of the inverted pendulum mode (#1 in fig. 5.9) is low. This is indicative of good mechanical isolation at the cavity body mode frequency. In addition, the eigenmodes of the supports remained distant in frequency from the cavity's drum mode, indicating that the isolation of the inverted pendulum is not spoiled by a mechanical short.




Figure 5.11: COMSOL von Mises yield simulations for proposed cavity support.

the thermal bath) is roughly:

$$N_{cycles} = \frac{Q}{n_m} = \frac{\hbar 2\pi f_m Q}{k_B T} \propto fQ, \quad (5.1)$$

where n_m is the thermal occupation number, k_B is Boltzmann's constant, and f_m is the frequency of the mechanical mode. Equivalently, we can express this as the minimum fQ required to guarantee one cycle of oscillation in the quantum ground state before thermal decoherence, which we plot in fig. 5.12. Every fQ shown in fig. 5.13 can satisfy this criterion easily at $T < 2K$, which is trivially attained using LHe. Additionally, the requirement $Q > 2.4n_m$ enables ground state cooling [82] (this is required in order to reach a ground state occupancy probability $P > 50\%$).




Figure 5.12: Minimum fQ product required such that the oscillator can exist in the ground state for at least one cycle as a function of temperature. This sets useful bounds for the fQ we need to attain at 123 K and 300 K. $fQ_{min}(123K) = 2.5 \times 10^{12}$, and $fQ_{min}(18K) = 3.7 \times 10^{11}$.

In fig. 5.13, we plot the $f \times Q$ that have been achieved by the optomechanics community alongside that achieved in this work. For reference, we add the maximum $f \times Q$ we could hope to achieve with the drum mode with the current cavities. In order to achieve quantum coherence for one cycle with the current cavities, we would need to achieve a body mode Q of 7.2×10^7 . We could also follow in the footsteps of bulk acoustic wave resonators (BAW) [83], and explore which higher order mechanical modes have high Q in order to increase the fQ product. This could provide up to a factor of ≈ 10 improvement, though there is the added complication of coupling strength; there is no guarantee that a particular high Q mechanical mode couples into the (Gaussian beam averaged) distance between the cavity mirrors.




Figure 5.13: $f \times Q$ plot comparing the various optomechanics experiments which I previously generated, with the low $f \times Q$ of the cryogenic silicon cavities shown. The maximum possible $f \times Q$ we might hope to obtain with this geometry is indicated by the arrow. Legend entries references in chapter 5.

The future generation of this experiment will use shorter cavities, which increases the body mode frequency, and in turn makes achieving high fQ easier. A cavity of reasonable size is shown in fig. 5.14, and yields a drum mode frequency of 145 kHz. This could yield an fQ of 4×10^{13} in a 2 K cryostat, which corresponds to 1×10^3 cycles before thermal noise excites it out of the ground state.

Painter[73]
Schwab[84]
Helium Superfluid[81]
UCSB[85]
Kavli Nanotubes[86]
Spain Nanotubes[87]
BAW quartz[88]
Boulder[89]
Silicon (1978)[90]
Nanocrystalline Diamond[91]
Microdisk[92]

Table 5.2: References for legend labels in fig. 5.13.




Figure 5.14: A COMSOL simulation showing a reasonably sized short silicon cavity with an eigenfrequency of 145 kHz. Such a cavity will be used in the future to facilitate direct measurement of the coating Brownian noise. Image courtesy of Rana X. Adhikari.

Chapter 6

Frequency Domain Subtraction

In the optical readout for the experiment, we read out the differential cavity length as:

$$Y = S + \sum n_i, \quad (6.1)$$

where Y is the differential frequency noise as measured by the readout, S is the actual differential cavity length signal, and n_i are the various noise sources which couple into the readout. If we detect the n_i separately from Y using auxiliary channels, we can subtract the noise from the experimental readout and lower the measurement noise floor. When we work in an orthogonal basis $\langle S|n_i\rangle = \langle n_i|n_j\rangle = 0$. In general, the assumption that $\langle n_i|n_j\rangle = 0$ is not true; at frequency bands like 60 Hz¹ and its harmonics this is often untrue, and the correlation between channels is very high. The assumption $\langle S|n_i\rangle = 0$ is generally justified, as S is some difficult to measure signal, and accidentally measuring it with some auxiliary channel is unlikely. We encounter the problem of cross correlation between channels when we attempt to subtract more than one auxiliary channels from our main (cavity length) signal. A mathematically correct (and very computationally expensive) way to solve this is to use the Gram-Schmidt process as done in [93]. To first order, we can just subtract each n_i from Y , and then check the correlations $\langle n_i|n_j\rangle$ explicitly to check for over subtraction and correct for it if necessary.

6.1 Wiener Filtering

The technique to optimally subtract auxiliary channels (often referred to as witness channels) from the signal we are trying to measure is called Wiener filtering, introduced by Norbert Wiener in 1949 [94]. Examples of this technique being implemented include subtracting seismic noise at the Caltech 40m prototype [95] and at the LIGO detectors [96, 97], subtracting auxiliary channel data at the 40m [98] and the LIGO sites [93], and to compare collocated LIGO H1 and H2 detectors to set a

¹60 Hz for the USA and parts of Asia, 50 Hz in other parts of the world.

bound on the stochastic gravitational wave background [99]. See [98] and chapter 5 of [36] for a brief review of the math behind the technique.

The aforementioned implementations of Wiener filtering were focused on the time domain, as the gravitational wave community relies on having access to the interferometer strain time series for their data analysis techniques. In the case of this experiment, we are interested in the spectral density of noise, and do not care about having access to the time series of the signal after filtering out the witness channels. For this reason, we may forego calculating the Wiener filters, which can be computationally expensive, and subtract the signals directly from measurements of their coherence.

6.2 Coherent Subtraction

In general, the channels we measure in lab are:

$$X_i = \sum G_{ij} n_j. \quad (6.2)$$

We make effort to ensure that the sensors we build sense one signal with high signal to noise, and thus it is reasonable to approximate the auxiliary channels as detecting a single noise:

$$X_i = H_i n_i + m_i, \quad (6.3)$$

where H_i is the transfer function of noise into our channel, m_i is its sensing noise, and $S_{H_i n_i}(f)/S_{m_i}(f) \ll 1$. For a single channel with a high SNR, X_1 , the noise output spectrum (call this $S_N(f)$, which is defined in Bendat and Piersol [56] to be the portion of S_Y not due to X_1) is:

$$S_N(f) = [1 - \gamma_{X_1, Y}^2(f)] S_Y(f), \quad (6.4)$$

where $\gamma_{X_1, Y}^2(f)$ is the ordinary coherence function between X_1 and Y . This quantity, how much noise in the measured lab signal Y is not due to X_1 is exactly what Wiener filtering tries to determine. In our case, S is the beat frequency (the output of the PLL frequency discriminator), and the X_i measured were the transmitted intensity noise, PDH error signals, and PDH control signals for both cavities. We measured $\gamma_{X_i, Y}^2(f)$ (fig. 6.1), the coherence of the auxiliary channels with the beat, as well as $\gamma_{X_i, X_j}^2(f)$ (fig. 6.2), the coherence between the auxiliary channels.

Based on the measured coherences, we have two signals (the PDH error signals) that are detected with high SNR and no cross correlation, which couple into our beat. This case is treated in section 7.2.1 of Bendat and Piersol [56], and the noise output spectrum is given by:

$$S_N(f) = [1 - \gamma_{X_1 Y}^2(f) - \gamma_{X_2 Y}^2(f)] S_Y(f), \quad (6.5)$$




Figure 6.1: Coherence measurements taken at 300 K of the beat with all auxiliary channels. The coherence between the PDH error signals and the beat above 1 kHz indicates the beat is limited by unsuppressed laser frequency noise. This is only true if the PDH loops are gain limited here.

where X_1 and X_2 are the two PDH error signals. The noise output spectrum, as well as the calibrated PDH error signals and their quadrature sum plotted vs the measured beat signal can be seen in fig. 6.3. The coherence measurements were done using an SR785 spectrum analyzer, which calculates coherence as:

$$\gamma_{XY}^2 = |\langle S_{XY}(f) \rangle_{rms}|^2 / (S_X(f)S_Y(f)), \quad (6.6)$$

where $|\langle S_{XY}(f) \rangle_{rms}|^2$ is the squared absolute value of the rms averaged cross spectral density between X and Y , and S_X and S_Y are power spectral densities. This estimator of coherence has a fractional downward bias of $1/(1 + SNR)$, where SNR is the signal to noise ratio with which we detect the PDH residual, in our case. As the PDH residual switches from noise limited to gain limited (e.g., the Fourier frequency at which the input referred noise is equal to the residual error signal), we have a downward bias of 50%, causing us to underestimate the coherence. In this regime, eq. (6.5) is the incorrect formula to use even if we had an unbiased estimate of the coherence, as the assumption that we have no input noise is no longer true. The noise output spectrum should not be interpreted as the amount of noise in the beat which is not due to the PDH signals at this frequency, though as we approach higher SNR it is fine.

In the gain limited regime of a servo loop, we can use the error signal as a witness sensor to subtract the residual noise out of some other sensor channel. Once you have pushed a loop as far as it can go due to time delay limitations in the feedback, you can extend the reach of your experiment by using the noise output spectrum.




Figure 6.2: Coherence measurements taken between the auxiliary channels. The low coherence between the East and West PDH error signals (orange and grey traces) indicates we can use eq. (6.5) to subtract them from the beat signal. The high coherence between the PDH error and control signal is what we would expect where the servo is gain limited (a zoomed in view of this is shown on the bottom). The high coherence seen at 60 Hz and its harmonics are due to ground loops. The slight increase in coherence in the West transmitted intensity with the West PDH error signal was due to a slight net PDH offset and can be ignored.




Figure 6.3: Residual PDH error signals and their quadrature sum plotted against the measured beat noise. The $\approx 30\%$ difference seen around 10 kHz is unknown, and not reflected in the coherence measurements. Shown in pink is the noise output spectrum, of the portion of the beat which is not due to the PDH error signals. The subtraction residual is likely limited by the AM sensitivity of the DFD (based on the morphology seen in noise budget fig. 4.1), and would be further suppressed by either mitigating the source of the anomalous RIN or closing an intensity stabilization servo on it. The discontinuities in the pink spectrum are due to noise in the coherence measurements shown in fig. 6.1; when subtracting multiple coherences where the noise in our measurement of coherence is comparable to the noise level we wish to unmask via subtraction, eq. (6.5) can yield negative power, corresponding to an imaginary amplitude. At these missing frequencies, we do not claim to know the beat spectrum, though could easily construct upper bounds from our knowledge of the coherences. It is likely that increasing measurement time would yield a smooth and continuous curve, so long as the noise levels and transfer functions are stationary on the (increased) measurement time scale. This problem of negative power only arises when subtracting multiple variables, which is obvious from eq. (6.4).

The output noise spectrum can be interpreted as the result of Wiener filtering using the residual PDH error signals as witness sensors to subtract the residual laser noise from the beat spectrum, expressed as a power spectral density. This can also be interpreted as how much the beat noise can be improved by increasing the gain of the PDH servos. As we approach higher Fourier frequencies, the phase delay in the PDH loops becomes harder to compensate for (higher order zeros and poles, phase delay due to gain bandwidth limitations in the op-amps, etc), and eventually physically impossible (time delay). Using the residual PDH error signal (or any error signal) as a witness sensor for Wiener filtering (either in the time or frequency domain) allows you to extend the sensitivity of an experiment to higher frequencies than you would otherwise be able to.

For suppressing the PDH residual, we can only eliminate it to the degree which we can sense it; the electronics noise at the error point limits the usefulness of using the PDH error signal residuals for frequency domain subtraction. In this version of the experiment, the sensing noise of the PDH loops is roughly $2 \times 10^{-2} \text{Hz}/\sqrt{\text{Hz}}$ at frequencies where the PDH loops are gain limited (above 10 kHz). In order to actually detect coating thermal noise with this technique, the noise at the error point (and the input noise of all electronics in the PDH loops) must be suppressed below the coating thermal noise limit. This level of noise is 40 to 130 times above the coating thermal noise limit estimated - we would need to lower the noise in the PDH loops by this amount in order for the technique to detect the coating thermal noise estimated for the noise budget presented in fig. 4.1. As mentioned in section 4.1, the coating Brownian noise is likely higher than this estimate by an unknown amount, and we intend on making short cavities with intentionally high coating Brownian noise levels, which will make the above mentioned ratio much smaller. We can likely improve the PDH noise above 10 kHz by up to a factor of 10, which may be enough to use coherent subtraction in order to detect the coating Brownian noise.

Regardless of whether or not this technique can be used to improve our sensitivity to coating thermal noise (and thus improve our error bars on ϕ_{coat} once we eventually detect it with $\text{SNR} > 1$), we should consider using subtraction as part of our effort to detect the mechanical mode(s) of the cavity (the relevance of which is discussed in chapter 5). Once we have a high Q body mode, this could both aid cooling efforts in the case of active cooling (subtracting the PDH residuals from the beat at the body mode frequency could make a very low noise signal to send back to a cavity length actuator), and detection of the mode in the case of passive cooling techniques such as resolved sideband cooling.²

If you can still push the bandwidth of your loop up, it is better to suppress the signal directly to increase its linearity by the means outlined in section 4.6. In addition, it should be noted that as you approach high coherence values, the subtraction technique approaches trying to measure a

²Resolved sideband cooling would require a much lower cavity pole and/or a much higher mechanical frequency — they are close to parity in this experiment.

small number by trying to subtract two large numbers. This technique is not limited to trying to extend the PDH loop bandwidth, though in this generation of the experiment that is the only place we can reap significant benefit in terms of lowering the experimental noise; future generations with different noise limitations may be able to use this technique using different signals at different Fourier frequencies.

Chapter 7

Doubling Noise

Harmonic generation of optical fields in nonlinear optical crystals can be modeled with classical coupled wave equations [100]. We define the fundamental and second harmonic fields as $E_1 = \mathcal{E}_1 \exp(2\pi i\nu_1 t)$ and $E_2 = \mathcal{E}_2 \exp(2\pi i\nu_2 t)$, respectively. In ideal second harmonic generation, $\nu_2 = 2\nu_1$. While the mean frequency ratio has been measured very precisely ($\langle \nu_2/\nu_1 \rangle = 7 \times 10^{-19}$ [101]), we present an upper bound on the frequency noise, or the spectrum of the time dependent quantity $\nu_2 - 2\nu_1$.

Several experiments at the forefront of precision metrology and frequency standards use harmonic generation in their experiments. These include iodine stabilized Nd:YAG lasers [102], optical frequency combs [103, 34], measurement of optical frequency ratios [101, 104], and precision atomic spectroscopy [105]. Many of these experiments provide bounds to any excess frequency noise which might be found in the second harmonic generation process (e.g., from thermodynamic fluctuations in the crystal temperature). As experiments like these push towards lower noise levels, the fundamental noise sources in second harmonic generation may become a relevant noise source.

One low noise area which will use harmonic generation in the future is interferometric detection of gravitational waves. These interferometers require measurements of mirror displacements at the level of $10^{-20} \text{ m}/\sqrt{\text{Hz}}$ [106]. In addition, increasingly complex optical readout schemes are required to reliably control the interferometer. A scheme was developed by Mullavey et al [31], tested at the Caltech 40m interferometer [107], and has been successfully implemented on the LIGO interferometer [108] involves to frequency doubling the 1064 nm carrier. This scheme uses both the fundamental and second harmonic to readout the motion of the mirrors. To determine the feasibility of such a scheme, we need to bound the frequency noise introduced in the second harmonic generation process. In our case, we use a 1064 nm carrier (Nd:YAG), and generate the 532 nm doubled light with PPKTP crystals. In order to meet the stability requirements for this technique, the excess frequency noise must be less than 70 mHz RMS from 10 mHz to 30 Hz [109]. Further in the future, optical configurations using more than one wavelength inside the interferometer to beat standard quantum noise limits of the detection [110] may be used. Fundamental noise limits in harmonic

generation may be relevant there.

7.1 Experimental Setup

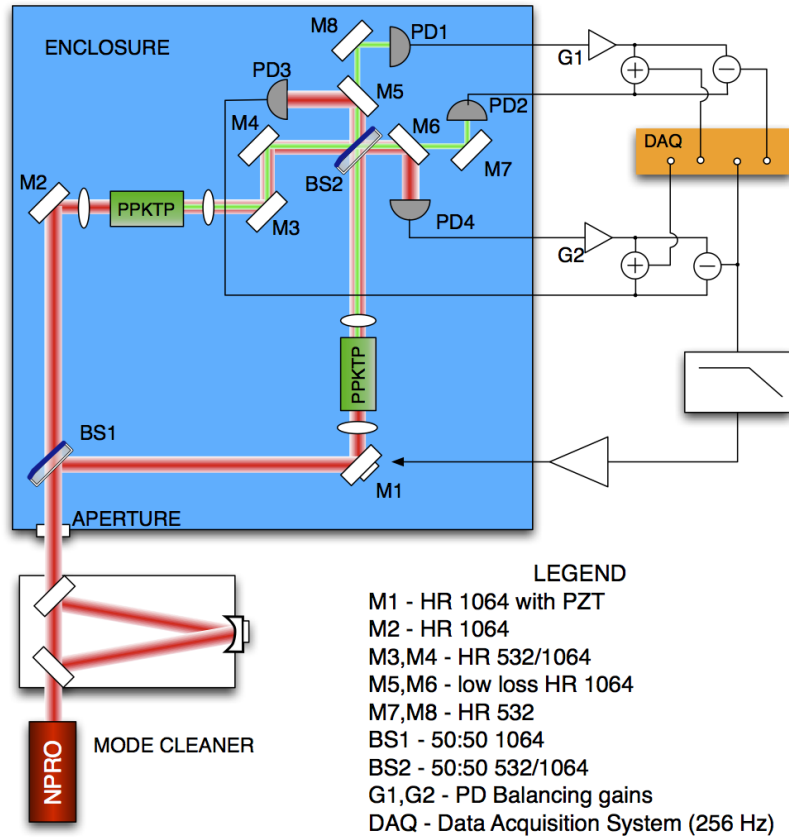


Figure 7.1: Experimental layout - A dual wavelength Mach-Zehnder Interferometer. Some readout mirrors for harmonic separation are not shown. See text for a more complete description.

Fig. 7.1 shows the experimental setup used to measure the uncorrelated frequency noise between the fundamental and the second harmonic. The beam from a 2 W non-planar ring laser at 1064 nm is passed through a spatial filter cavity [111]. The spatially filtered beam then enters the Mach-Zehnder through a hole in an acoustic enclosure. PPKTP crystals (Raicol Crystals) inside temperature stabilized ovens are placed in each arm, with lenses added to mode match to the each crystal. Two additional dichroic mirrors (M3, M4) are placed in one arm for alignment. The beams recombine on a dichroic 50:50 beamsplitter and low loss HR mirrors ($T < 10$ ppm at 1064 nm) are used to separate the fundamental from the second harmonic. Commercial dichroic mirrors (HR532, AR1064) were used to further separate the 532 nm from the 1064 nm light. For the 532 nm and 1064 nm detection, Si (Hamamatsu 1223) and InGaAs (GPD 2000) photodiodes were used, respectively. The Mach-Zehnder arm lengths were adjusted to be mid fringe for both 532 and 1064 nm simultaneously. The

interferometer is locked to this point by applying feedback from the 1064 nm PD difference signal to the PZT on M1 (see Fig. 7.1). We drove the PZT with a small signal at 100 Hz to monitor the calibration (rad/V) at both wavelengths, as any differential frequency noise between the two wavelengths would cause the 532 nm calibration to drift. The calibration of the Mach-Zehnder was confirmed by sweeping through multiple fringes. The 1064 nm and 532 nm Mach-Zehnder length signals were read out with two pair of balanced homodyne detectors. The sum and difference signals were digitized at 256 Hz and processed further offline.

In the non-depleted pump approximation, with imperfect phase matching, the phase relation of the two fields in each arm before the recombining beamsplitter is:

$$\theta_{532} = 2\theta_{1064} - \pi/2 - \Delta k L/2, \quad (7.1)$$

where θ_{1064} and θ_{532} describe the relative phases of the fundamental and the second harmonic, respectively, L is the length of the doubling crystal, and Δk is the phase mismatch, the parameter normally used to describe efficiency of second harmonic generation. In theory, Δk can be arbitrarily small (limited in practice only by the ability to stabilize the temperature of the nonlinear crystal).

Using superscripts to differentiate between the two arms, the difference in the second harmonic phase is thus

$$\theta_{532}^A - \theta_{532}^B = 2(\theta_{1064}^A - \theta_{1064}^B) - (\Delta k^A - \Delta k^B)/2. \quad (7.2)$$

7.2 Results




Figure 7.2: Example time series of the Mach-Zehnder output. The RMS phase noise over the one hour period was 6 mrad RMS at 532 nm and 3 mrad RMS at 1064 nm.

The phase difference ($\delta\theta_{1064}(t) \equiv \theta_{1064}^A(t) - \theta_{1064}^B(t)$) was suppressed by the servo, which had a unity gain frequency of ~ 10 Hz. Typical values of $\delta\theta_{1064}(t)$ and $\delta\theta_{532}(t)$ are shown in Fig. 7.2.




Figure 7.3: On the left we see the full noise budget of the experiment: $\theta_{1064}^A - \theta_{1064}^B$ (in red) (which is the in-loop signal of the Mach-Zehnder), $\theta_{532}^A - \theta_{532}^B$ (in green) (which contains excess phase noise), and the subtraction residual (in black). The total RMS frequency fluctuation in the measurement band is 3 mHz. On the right we have a comparison of this work with previous bounds. These bounds show up as frequency noise [112], timing error [34], and phase noise [103].

The frequency noise amplitude spectral density of the Mach Zehnder is shown in Fig. 7.3 with known noise sources. $\delta\nu_{1064}$ and $\delta\nu_{532}$ are shown in red and green, respectively, where $\delta\nu \equiv \delta\dot{\theta}/2\pi$. Above 10 Hz the phase noise is dominated by a forest of mechanical resonances on the optical table which show up strongly in both $\delta\nu_{1064}$ and $\delta\nu_{532}$. We used the average transfer function over the measurement time, $H(f) = \langle \delta\nu_{532}(f)/\delta\nu_{1064}(f) \rangle$, to subtract the noise which is coherent between the two $\delta\nu$ s. This subtracted level is shown as the black trace in Fig. 7.3. The black trace bounds any noise source which causes frequency noise between the fundamental and the harmonic and is uncorrelated between the two ovens, such as thermodynamic fluctuations in the crystals, or any temperature fluctuation not common to both ovens. Some common mode effects such as intensity dependent phase shifts and temperature fluctuations (and thus phase matching fluctuations) are suppressed by the experimental setup, so these technical noise sources will not be visible in these measurements. However, a pessimistic estimate of the temperature noise coupling assuming no common mode rejection shown in Fig. 7.3 is below the measured excess frequency noise. In addition, the intensity of the fundamental was only stabilized to a level of 2×10^{-6} at 1 Hz, where it can be lowered to the 10^{-8} at 1 Hz with current techniques, so unless the rejection of this effect was more than 46 dB, we can safely ignore it. It is highly improbable that the intensity to frequency noise coupling would be above the level shown in Fig. 7.3. The excess noise below 10 Hz was found to be

correlated with air currents on the table, and would be reduced by moving the setup into a vacuum chamber. The total RMS excess frequency noise of the black trace is 3 mHz RMS in the 10 mHz to 128 Hz band.

7.3 Noise Sources

In addition to the usual technical noise sources, it is worth considering whether there is a more fundamental limit to the relative phase between the fundamental and the harmonic. A rough estimate of thermal noise from thermoelastic (Zener) damping was obtained by directly applying [113] the Fluctuation-Dissipation Theorem. We treat the crystal as an 0.8 mm radius cylinder, and follow the calculation done by Heinert et al. [114]. This yields a spectral density of Δk taking into account both thermorefractive and thermoelastic fluctuations. When expressed as frequency fluctuations, it is well approximated by $2.5/(1 + 500f^{-7/8}) \mu\text{Hz}/\sqrt{\text{Hz}}$ above 10 Hz. Below that, it must be flat or continue to decrease, or else the RMS temperature integral would diverge. Practically speaking, in our band of interest, the temperature fluctuations of the ovens is at least 2 orders of magnitude greater than these fundamental thermodynamic temperature fluctuation. In the future, when researchers seek to make frequency comparisons at better than the 10^{-21} level, these thermal noises will have to be calculated with more accuracy.

7.4 Previous Bounds on Excess Frequency Noise

Previous bounds in the literature have been reported in a number of different ways. In Fig. 7.3 we compare frequency noise [112], timing errors [34], and phase noise [103]. In Fig. 7.4 we compare Allan deviations [115, 104, 116, 101, 34, 102, 117]. While the comparison in Fig. 7.3 is straightforward, some caution should be taken interpreting Fig. 7.4. Our Allan deviation at the 0.1 s time scale is heavily influenced by the high frequency noise in the measurement (10-128 Hz). Since we low pass the signal to acquire data at 256 Hz, we reject noise which would make the Allan deviation increase at all time scales. It should also be emphasized that we only measure relative frequency fluctuations, and that there are some common mode noise sources which the experiment is insensitive to. See [118] for more information on Allan deviations and phase noise.

7.5 Conclusions

In conclusion, we have demonstrated that the RMS frequency fluctuations added from uncorrelated mechanisms between two SHG crystals is less than 3 mHz at time scales over 10 ms. The obvious correlated mechanisms (temperature and intensity noise coupling) are likely insignificant compared




Figure 7.4: The Allan deviations from this experiment (green and black) were obtained from the power spectrum as described in [119]. Bounds from previous work are shown for reference.

to this as discussed in section 7.2. This is low enough to not limit the lock acquisition [31, 120] scheme for Advanced LIGO and other gravitational-wave detectors. Additionally, we have shown that there is no excess noise process at a level which is of interest to those doing precision atomic spectroscopy [105], using frequency combs to transfer optical harmonics [34, 103, 121] and other tests of fundamental physics.

Chapter 8

Conclusions and Future Work

The ultimate goals of this experiment are to make a definitive measurement of the coating thermal noise at cryogenic temperatures, and pursue macroscopic quantum mechanics (MQM) frontiers using these coating thermal noise limited cavities, and the accompanying cavity length readout. In this work we have presented a two cavity system, cooled to 123 K, with a laser locked to each cavity, and a low noise beat readout. In the course of this work, we identified and mitigated numerous noise sources, though more work on mitigating the remaining sources is required to reach the coating thermal noise limit.

The main upgrades which need to be made to the experiment in order to reach the coating Brownian noise limit are:

- Replace cryostat windows with ultra low scatter optics
- Damp pendulum modes of suspension
- Increase bandwidth and lower noise in PDH loops
- Decrease beat frequency readout noise

The details of the above upgrade recommendations are in their respective sections in this thesis. In order to pursue MQM measurements, we need to increase the Q of the cavity's first body mode by four orders of magnitude. To accomplish this, we must alter the method we use to support the cavity as detailed in chapter 5.

Another way we can attain a definitive measurement of coating thermal noise is to change the cavity geometry. If we shift to a shorter cavity, such as the one used in [12, 70], we can use a smaller spot size on the mirror, which increases the coupling level of coating Brownian noise. After a successful measurement the coating noise, we can shift back to large spot sizes if we still wish to make an ultra low noise cavity. In addition, shifting to a smaller cavity increases the eigenfrequency of the first body mode, which makes it easier to attain a high fQ product for the purposes of pursuing MQM measurements.

The longer term MQM goals of the experiment include the following:

- Cooling the silicon cavities to the ground state
- Couple the cavity body modes either with force actuation like an ESD or some optical path, and exploring the free evolution of the coupled hamiltonian as the cavities leave the ground state in a timescale recordable by current ADC technologies
- Isolating the cavities and then tidally coupling them in order to search for anomalous phase delay in G , the gravitational constant (predicted by some alternate theories of quantum gravity)
- Searching for frequency splitting in the body mode (as detailed by by Huan Yang [122]).

Characterizing coating Brownian noise at cryogenic temperatures is an important step in the field of precision metrology, paving the way for both lower noise ultra-stable reference cavities, and lower noise gravitational wave detectors. Measuring quantum mechanical behavior in a macroscopic system is an important milestone in our exploration of the interface between quantum physics and gravity. In this work we made a significant exploratory step towards both of these goals.

Appendix A

Resonator Equations

We reproduce the equations which relate physical dimensions to optical properties described in [123], algebraically rearranged to reflect which quantities we measured in section 2.3. The equations below are used to derive linearly dependent optical properties of the cavity in table 2.2.

The free spectral range (f_{FSR}) of the cavity is given by eq 48 ch11, $f_{\text{FSR}} = c/2L$, where L is the cavity length. Rearranging, we get cavity length as a function of f_{FSR} :

$$L(f_{\text{FSR}}) = \frac{c}{2f_{\text{FSR}}} \quad (\text{A.1})$$

The higher order mode spacing for a symmetrical cavity is given by eq 22 ch19, $2\pi f_{\text{hom}}L/c - (n + m + 1) \cos^{-1}(\pm g) = q\pi$, where g is the cavity g-factor, and $n, m, q \in \mathbb{Z}$. Rearranging, we get:

$$g(L, f_{\text{FSR}}) = \cos\left(\frac{2\pi f_{\text{hom}}L}{c}\right) \quad (\text{A.2})$$

The g-factor of a symmetrical cavity (g) is given by eq 3 ch 19, $g = 1 - L/R$, where R is the mirror radius of curvature. Rearranging, we get:

$$R(L, g) = \frac{L}{1 - g} \quad (\text{A.3})$$

The cavity waist (ω_0) is given by eq 4 ch 19:

$$\omega_0(L, g) = \left(\frac{L\lambda}{2\pi}\right)^{1/2} \left(\frac{1+g}{1-g}\right)^{1/4} \quad (\text{A.4})$$

where λ is the wavelength of light used to interrogate the cavity.

The cavity finesse (\mathcal{F}) for low loss cavities is given by eq 54 ch 11:

$$\mathcal{F}(f_{\text{FSR}}, f_{\text{cav}}) = \frac{f_{\text{FSR}}}{2f_{\text{cav}}} \quad (\text{A.5})$$

where f_{cav} is the cavity pole, or the frequency offset at which the cavity transmission decreases by $1/2$. We can also express the finesse in the more familiar units of parts per million [ppm] of incident power lost per mirror:

$$Loss(\mathcal{F}) = 1 \times 10^6 \frac{\pi}{\mathcal{F}} \quad (\text{A.6})$$

Appendix B

Suspension Redesign

The suspension (see fig. B.2) was initially designed with the goal that the common mode rejection of a semi rigid structure would mitigate the low frequency acoustic coupling. The mechanical resonances of the initial suspension were the limiting noise source from 30 Hz - 2 kHz, as shown in fig. B.5.




Figure B.1: Spring Suspension

We designed a new suspension to mitigate the acoustic coupling seen in the initial suspension: we attached the supports for both cavities to the same rigid platform, and suspended this platform from the cold plate with four BeCu springs, as shown in fig. B.1. This allows for common mode rejection of the noise due to beam overlap we would see if the rms motion of the cavities was differential.

We do not, however, expect any such common mode rejection in the noise from scatter born parasitic interferometric paths. The corner frequencies of such paths scale as the total rms motion of the parasitic path, which will be dominated by the platform motion for us. For this reason it is




Figure B.2: Rigid Suspension

better to damp the springs. We chose to damp the springs by stuffing the top 1/5 of the springs with 0000 Cu wool. The density of packing was adjusted so that the portion of spring contacting the wool would still move freely on excitation.

A problem with this methodology can be seen in fig. B.4: different amplitude oscillations have different damping dynamics (e.g. slip-stick dynamics). Additionally, we do not know how the up-converted noise overlaps with the violin modes of the suspension. It is possible that we are increasing the noise at frequencies where the suspension acts as a mechanical short between the cold plate and the suspension platform. Improved damping could be achieved by adding more Cu wool to the spring, though some estimate of analysis of how the up-converted noise overlaps with the modes of the spring should proceed this. Alternative solutions include using eddy current damping, active damping [124], or different materials and spring geometries. To our knowledge, neither of the aforementioned techniques has been successfully implemented at cryogenic temperatures, so there may be unknown engineering complications. In addition, as BeCu is generally chosen for use at cryogenic temperatures for its relatively high Q-factor, and we want something with a relatively low Q for the suspension, exploring a more sensible material for use as springs at cryogenic temperatures seems promising.




Figure B.3: Spring Suspension Q-reduction - The modes were excited by hand, and a simple HeNe laser photodetector occlusion setup was used to record a time series of the ring down. Above are the spectra for the ring down of each mode.





Figure B.4: Spring Suspension Q-reduction - by averaging in the time domain, we can estimate the mechanical loss (Q^{-1}) as a function of cycle number. The damping achieved clearly decreases with the magnitude of the oscillation. Plot courtesy Nicolas Smith.




Figure B.5: Highlighted here are two traces from the bottom of fig. 4.2: the beat noise spectrum during the first cryogenic measurement, when we were using the suspension shown in fig. B.2, and that once we had switched to the new suspension shown in section 3.2. Examining the experimental noise during its noisiest state (LN2 boiling noise present) allows for easier comparison of the two suspensions. The vibrational noise transferred to cavity length through shaking of the suspension is drastically reduced in the spring based suspensions. The low frequency improvement (below 100 Hz) is due to scatter mitigation. The broadband improvement above 100 Hz is due to upgrading the beat readout.

Appendix C

Window Scatter

We believe that the cryostat window coatings were contaminated based on visible marring upon delivery (see fig. C.1). We cleaned them extensively with the drag wiping technique, inspecting them under white light, and were able to remove all visible evidence of damage.




Figure C.1: Window as received

After cleaning the windows, we measured their transmission, reflection, and scatter. Ideally, we want to measure the Bidirectional reflectance distribution function (or BRDF [125, 126]) for each window. We used a technique similar to Magana et al [127], shown in fig. C.2 to measure the BRDF. The BRDF is the intensity of the light scattered from a surface per solid angle, per input

intensity (dotted with the normal vector), and is a function of input and output angle. We calibrate our scattering sensor (a CCD camera in this case) against a surface with a known BRDF. We can assume that a flat white surface (such as a sheet of paper, or in this case, a clean white lens cloth) is a "Lambertian surface" [126], which scatters isotropically. This means that it has a BRDF of $1/\pi$. We scale the measurement of input power and total power on the camera (CCD counts) such that the Lambertian surface gives us a BRDF of $1/\pi$. See fig. C.3 for CCD images of the calibration and window. The angle of incidence between the laser and window shown in the diagram is the angle between the window's normal vector and the cavity optic axis in normal operation when they are installed in the cryostat.




Figure C.2: BRDF measurement setup. We illuminate the optic under test with a 1550 nm light source, monitor the input, reflected, and transmitted power, and capture the scatter with a 1550nm sensitive CCD camera.

We used a 1550 nm sensitive CCD, and averaged over 10000 images with the laser both on and off. We summed over each pixel, and took the mean of all images. We used the laser off measurements as our dark noise calibration, and subtracted this out. We were only able to create an upper bound on the cryostat window BRDF from this measurement, which is presented in table C.1 along with the transmission and reflection of each window. We present the standard deviation of a 10000 image sample as our error bars. This is likely an overestimate of the error based on how close many of the

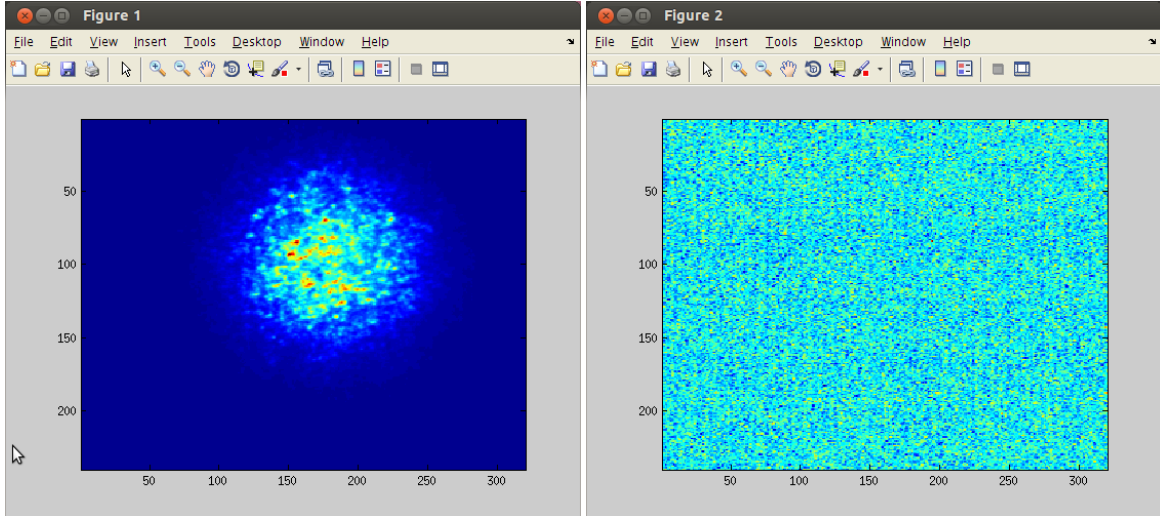


Figure C.3: Images from the camera shown in fig. C.2 used to measure scatter as BRDF. On the left is the backscatter from a white lens cleaning cloth. On the right is the backscatter from one of the cryostat windows after cleaning. The window backscatter measurement seems clearly noise limited. Both images are averaged over many exposures to improve SNR.

measurements are to each other (1×10^{-4} /sr grouping with a 1×10^{-2} /sr error.) Regardless, we were only able to place an upper bound on the window BRDF.

Window	T	R ($\times 10^{-3}$)	BRDF (sr^{-1})
E	0.98	1.38	-0.1833
F	1.00	1.46	-0.0300
G	0.99	1.50	-0.0109
C	0.98	1.38	-0.239
D	1.00	1.31	1.661
H	0.98	1.01	1.400
A	0.97	1.36	1.559
B	0.98	0.83	1.2893
uncertainty	0.01	(0.01×10^{-3})	0.011

Table C.1: Cryostat window properties. The first four BRDF measurements (E,F,G and C) were taken on one day, and the latter four were taken on another. It is clear that an unknown systematic dominated the BRDF measurement.


In order to make a more sensitive measurement and improve the scatter (which is important for reasons highlighted in section 4.2), we need to make future measurements using a sensitive photodetector and lens in place of the CCD. We opt for switching to a PD as sensitive CCDs are not readily available at 1550 nm (most mid range CCDs are just Si or Ge with a phosphor coating to down-convert the radiation to a detectable optical frequency). Alternately, we could just make the measurements at 1064 nm, where high power lasers and sensitive cameras/PDs are readily available. We do not know of any reason the BRDF would drastically change by scaling the frequency up by

33%, so this also would be reasonable.


Appendix D

Cavity Seismic Susceptibility


Results of the COMSOL simulation described in section 4.3 for alternative cavity geometries. The cavity geometry used in this work was roughly 4" length, and 1" radius. COMSOL simulations were performed by Evan Hall.



(a)



(b)



(c)

Figure D.1

Appendix E

Electronics

Here we present the rest of the parts of the PDH loop shown in fig. 4.13.




Figure E.1: Shown here is the optical plant, which is change in rf power at the PDH error point per Hz of frequency offset from cavity resonance.




Figure E.2: Shown here is the transfer function of RF Photodetection and demodulation of the PDH signal.




Figure E.3: LB1005 and low pass filter used in conjunction as the PDH servo.

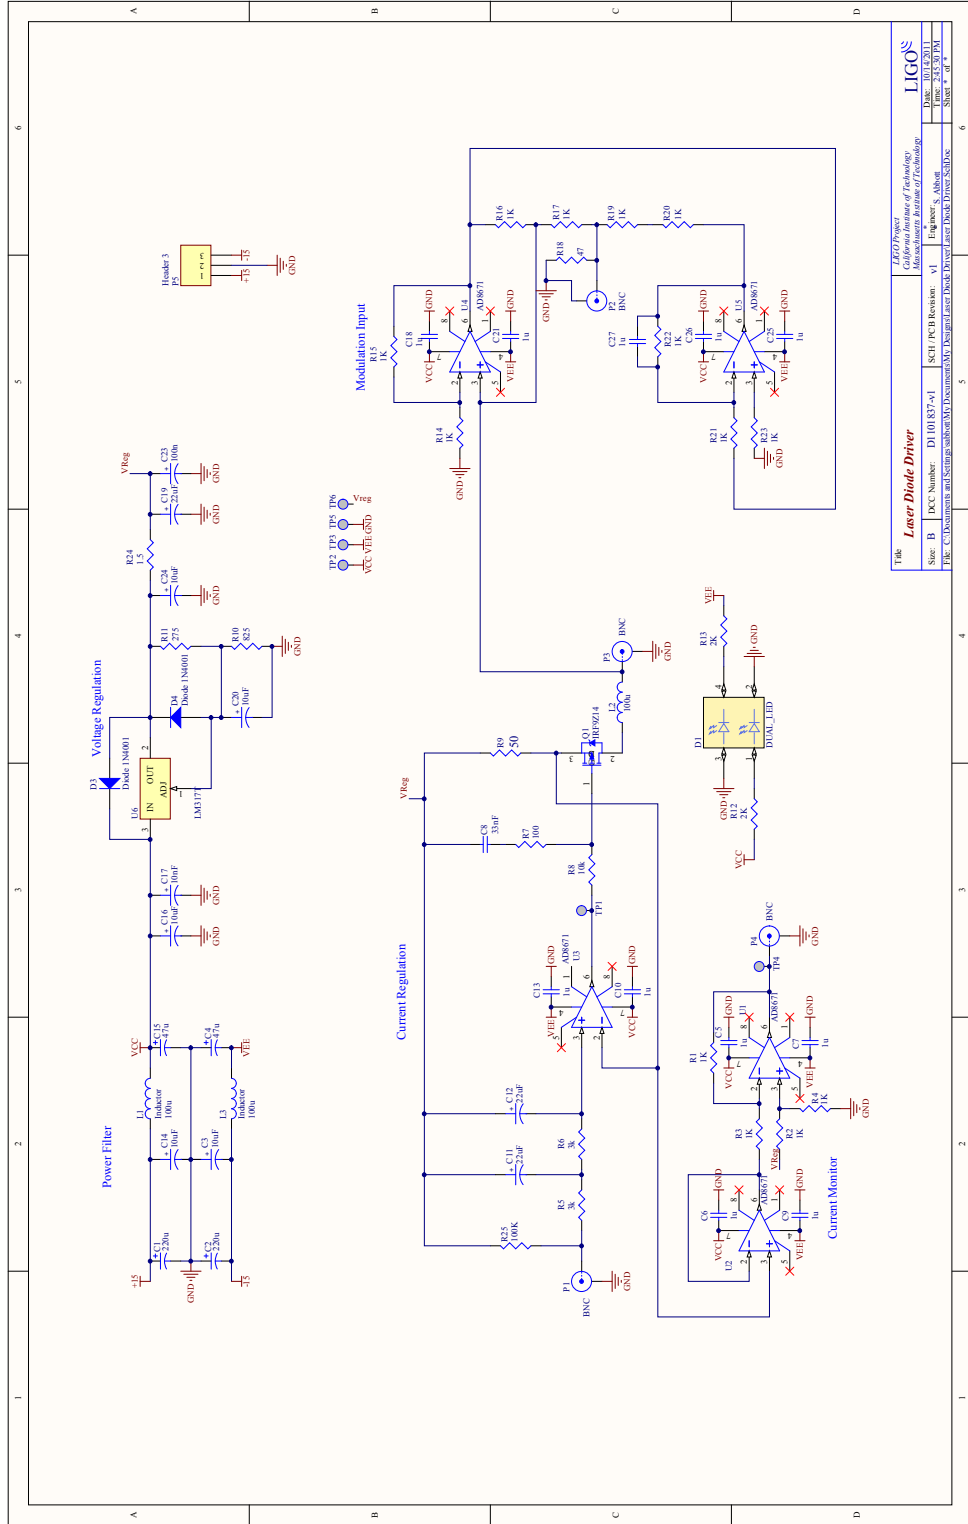


Figure E.4

Appendix F

Photothermal Transfer Functions

Here we show the rest of the fitting which produced the data points in fig. 4.25.




Figure F.1




Figure F.2




Figure F.3




Figure F.4

Bibliography

- [1] The LIGO Scientific Collaboration, J Aasi, B P Abbott, R Abbott, T Abbott, M R Abernathy, K Ackley, C Adams, T Adams, P Addesso, R X Adhikari, V Adya, C Affeldt, N Aggarwal, O D Aguiar, A Ain, P Ajith, A Alemeic, B Allen, D Amariutei, S B Anderson, W G Anderson, K Arai, M C Araya, C Arceneaux, J S Areeda, G Ashton, S Ast, S M Aston, P Aufmuth, C Aulbert, B E Aylott, S Babak, P T Baker, S W Ballmer, J C Barayoga, M Barbet, S Barclay, B C Barish, D Barker, B Barr, L Barsotti, J Bartlett, M A Barton, I Bartos, R Bassiri, J C Batch, C Baune, B Behnke, A S Bell, C Bell, M Benacquista, J Bergman, G Bergmann, C P L Berry, J Betzwieser, S Bhagwat, R Bhandare, I A Bilenko, G Billingsley, J Birch, S Biscans, C Biwer, J K Blackburn, L Blackburn, C D Blair, D Blair, O Bock, T P Bodiya, P Bojtos, C Bond, R Bork, M Born, Sukanta Bose, P R Brady, V B Braginsky, J E Brau, D O Bridges, M Brinkmann, A F Brooks, D A Brown, D D Brown, N M Brown, S Buchman, A Buikema, A Buonanno, L Cadonati, J Caldern Bustillo, J B Camp, K C Cannon, J Cao, C D Capano, S Caride, S Caudill, M Cavagli, C Cepeda, R Chakraborty, T Chalermongsak, S J Chamberlin, S Chao, P Charlton, Y Chen, H S Cho, M Cho, J H Chow, N Christensen, Q Chu, S Chung, G Ciani, F Clara, J A Clark, C Collette, L Cominsky, M Constancio Jr, D Cook, T R Corbitt, N Cornish, A Corsi, C A Costa, M W Coughlin, S Countryman, P Couvares, D M Coward, M J Cowart, D C Coyne, R Coyne, K Craig, J D E Creighton, T D Creighton, J Cripe, S G Crowder, A Cumming, L Cunningham, C Cutler, K Dahl, T Dal Canton, M Damjanic, S L Danilishin, K Danzmann, L Darteze, I Dave, H Daveloza, G S Davies, E J Daw, D DeBra, W Del Pozzo, T Denker, T Dent, V Dergachev, R T DeRosa, R DeSalvo, S Dhurandhar, M Daz, I Di Palma, G Dojcinoski, E Dominguez, F Donovan, K L Dooley, S Doravari, R Douglas, T P Downes, J C Driggers, Z Du, S Dwyer, T Eberle, T Edo, M Edwards, M Edwards, A Effler, H.-B Eggenstein, P Ehrens, J Eichholz, S S Eikenberry, R Essick, T Etzel, M Evans, T Evans, M Factourovich, S Fairhurst, X Fan, Q Fang, B Farr, W M Farr, M Favata, M Fays, H Fehrmann, M M Fejer, D Feldbaum, E C Ferreira, R P Fisher, Z Frei, A Freise, R Frey, T T Fricke, P Fritschel, V V Frolov, S Fuentes-Tapia, P Fulda, M Fyffe, J R Gair, S Gaonkar, N Gehrels, L Gergely, J A Giaime, K D Giardina, J Gleason, E Goetz, R Goetz, L Gondan, G Gonzalez, N Gordon, M L Gorodetsky, S Gossan, S Goler,

C Grf, P B Graff, A Grant, S Gras, C Gray, R J S Greenhalgh, A M Gretarsson, H Grote, S Grunewald, C J Guido, X Guo, K Gushwa, E K Gustafson, R Gustafson, J Hacker, E D Hall, G Hammond, M Hanke, J Hanks, C Hanna, M D Hannam, J Hanson, T Hardwick, G M Harry, I W Harry, M Hart, M T Hartman, C-J Haster, K Haughian, S Hee, M Heintze, G Heinzl, M Hendry, I S Heng, A W Heptonstall, M Heurs, M Hewitson, S Hild, D Hoak, K A Hodge, S E Hollitt, K Holt, P Hopkins, D J Hosken, J Hough, E Houston, E J Howell, Y M Hu, E Huerta, B Hughey, S Husa, S H Huttner, M Huynh, T Huynh-Dinh, A Idrisy, N Indik, D R Ingram, R Inta, G Islas, J C Isler, T Isogai, B R Iyer, K Izumi, M Jacobson, H Jang, S Jawahar, Y Ji, F Jimnez-Forteza, W W Johnson, D I Jones, R Jones, L Ju, K Haris, V Kalogera, S Kandhasamy, G Kang, J B Kanner, E Katsavounidis, W Katzman, H Kaufer, S Kaufer, T Kaur, K Kawabe, F Kawazoe, G M Keiser, D Keitel, D B Kelley, W Kells, D G Keppel, J S Key, A Khalaidovski, F Y Khalili, E A Khazanov, C Kim, K Kim, N G Kim, N Kim, Y.-M Kim, E J King, P J King, D L Kinzel, J S Kissel, S Klimenko, J Kline, S Koehlenbeck, K Kokeyama, V Kondrashov, M Korobko, W Z Korth, D B Kozak, V Kringel, B Krishnan, C Krueger, G Kuehn, A Kumar, P Kumar, L Kuo, M Landry, B Lantz, S Larson, P D Lasky, A Lazzarini, C Lazzaro, J Le, P Leaci, S Leavey, E O Lebigot, C H Lee, H K Lee, H M Lee, J R Leong, Y Levin, B Levine, J Lewis, T G F Li, K Libbrecht, A Libson, A C Lin, T B Littenberg, N A Lockerbie, V Lockett, J Logue, A L Lombardi, M Lormand, J Lough, M J Lubinski, H Lck, A P Lundgren, R Lynch, Y Ma, J Macarthur, T MacDonald, B Machenschalk, M MacInnis, D M Macleod, F Magaa-Sandoval, R Magee, M Mageswaran, C Maglione, K Mailand, I Mandel, V Mandic, V Mangano, G L Mansell, S Mrka, Z Mrka, A Markosyan, E Maros, I W Martin, R M Martin, D Martynov, J N Marx, K Mason, T J Massinger, F Matichard, L Matone, N Mavalvala, N Mazumder, G Mazzolo, R McCarthy, D E McClelland, S McCormick, S C McGuire, G McIntyre, J McIver, K McLin, S McWilliams, G D Meadors, M Meinders, A Melatos, G Mendell, R A Mercer, S Meshkov, C Messenger, P M Meyers, H Miao, H Middleton, E E Mikhailov, A Miller, J Miller, M Millhouse, J Ming, S Mirshekari, C Mishra, S Mitra, V P Mitrofanov, G Mitselmakher, R Mittleman, B Moe, S D Mohanty, S R P Mohapatra, B Moore, D Moraru, G Moreno, S R Morriss, K Mossavi, C M Mow-Lowry, C L Mueller, G Mueller, S Mukherjee, A Mullavey, J Munch, D Murphy, P G Murray, A Mytidis, T Nash, R K Nayak, V Necula, K Nedkova, G Newton, T Nguyen, A B Nielsen, S Nissanke, A H Nitz, D Nolting, M E N Normandin, L K Nuttall, E Ochsner, J O Dell, E Oelker, G H Ogin, J J Oh, S H Oh, F Ohme, P Oppermann, R Oram, B O'Reilly, W Ortega, R OShaughnessy, C Osthelder, C D Ott, D J Ottaway, R S Ottens, H Overmier, B J Owen, C Padilla, A Pai, S Pai, O Palashov, A Pal-Singh, H Pan, C Pankow, F Pannarale, B C Pant, M A Papa, H Paris, Z Patrick, M Pedraza, L Pekowsky, A Pele, S Penn, A Perreca, M Phelps, V Pierro, I M Pinto, M Pitkin, J Poeld, A Post, A Poteomkin, J Powell, J Prasad, V Predoi,

S Premachandra, T Prestegard, L R Price, M Principe, S Privitera, R Prix, L Prokhorov, O Puncken, M Prrer, J Qin, V Quetschke, E Quintero, G Quiroga, R Quitzow-James, F J Raab, D S Rabeling, H Radkins, P Raffai, S Raja, G Rajalakshmi, M Rakhmanov, K Ramirez, V Raymond, C M Reed, S Reid, D H Reitze, O Reula, K Riles, N A Robertson, R Robie, J G Rollins, V Roma, J D Romano, G Romanov, J H Romie, S Rowan, A Rdiger, K Ryan, S Sachdev, T Sadecki, L Sadeghian, M Saleem, F Salemi, L Sammut, V Sandberg, J R Sanders, V Sannibale, I Santiago-Prieto, B S Sathyaprakash, P R Saulson, R Savage, A Sawadsky, J Scheuer, R Schilling, P Schmidt, R Schnabel, R M S Schofield, E Schreiber, D Schuette, B F Schutz, J Scott, S M Scott, D Sellers, A S Sengupta, A Sergeev, G Serna, A Sevigny, D A Shaddock, M S Shahriar, M Shaltev, Z Shao, B Shapiro, P Shawhan, D H Shoemaker, T L Sidery, X Siemens, D Sigg, A D Silva, D Simakov, A Singer, L Singer, R Singh, A M Sintes, B J J Slagmolen, J R Smith, M R Smith, R J E Smith, N D Smith-Lefebvre, E J Son, B Sorazu, T Souradeep, A Staley, J Stebbins, M Steinke, J Steinlechner, S Steinlechner, D Steinmeyer, B C Stephens, S Steplewski, S Stevenson, R Stone, K A Strain, S Strigin, R Sturani, A L Stuver, T Z Summerscales, P J Sutton, M Szczepanczyk, G Szeifert, D Talukder, D B Tanner, M Tpai, S P Tarabrin, A Taracchini, R Taylor, G Tellez, T Theeg, M P Thirugnanasambandam, M Thomas, P Thomas, K A Thorne, K S Thorne, E Thrane, V Tiwari, C Tomlinson, C V Torres, C I Torrie, G Traylor, M Tse, D Tshilumba, D Ugolini, C S Unnikrishnan, A L Urban, S A Usman, H Vahlbruch, G Vajente, G Valdes, M Vallisneri, A A van Veggel, S Vass, R Vaulin, A Vecchio, J Veitch, P J Veitch, K Venkateswara, R Vincent-Finley, S Vitale, T Vo, C Vorvick, W D Vousden, S P Vyatchanin, A R Wade, L Wade, M Wade, M Walker, L Wallace, S Walsh, H Wang, M Wang, X Wang, R L Ward, J Warner, M Was, B Weaver, M Weinert, A J Weinstein, R Weiss, T Welborn, L Wen, P Wessels, T Westphal, K Wette, J T Whelan, S E Whitcomb, D J White, B F Whiting, C Wilkinson, L Williams, R Williams, A R Williamson, J L Willis, B Willke, M Wimmer, W Winkler, C C Wipf, H Wittel, G Woan, J Worden, S Xie, J Yablon, I Yakushin, W Yam, H Yamamoto, C C Yancey, Q Yang, M Zanolin, Fan Zhang, L Zhang, M Zhang, Y Zhang, C Zhao, M Zhou, X J Zhu, M E Zucker, S Zuraw, and J Zweizig. Advanced LIGO. *Classical and Quantum Gravity*, 32(7):074001, 2015.

- [2] T P Heavner, S R Jefferts, E A Donley, J H Shirley, and T E Parker. Nist-f1: recent improvements and accuracy evaluations. *Metrologia*, 42(5):411, 2005.
- [3] C. W. Chou, D. B. Hume, J. C. J. Koelemeij, D. J. Wineland, and T. Rosenband. Frequency comparison of two high-accuracy al^+ optical clocks. *Phys. Rev. Lett.*, 104:070802, Feb 2010.
- [4] T. L. Nicholson, M. J. Martin, J. R. Williams, B. J. Bloom, M. Bishof, M. D. Swallows, S. L. Campbell, and J. Ye. Comparison of Two Independent Sr Optical Clocks with 1×10^{-17} Stability at 10^3 s. *Physical Review Letters*, 109(23):230801, December 2012.

- [5] Kenji Numata, Amy Kemery, and Jordan Camp. Thermal-noise limit in the frequency stabilization of lasers with rigid cavities. *Phys. Rev. Lett.*, 93:250602, Dec 2004.
- [6] Rick L Savage, PJ King, and SU Seel. A highly stabilized 10-watt nd: Yag laser for the laser interferometer gravitational-wave observatory (LIGO) ¹. *Laser Physics*, 8(3):679–685, 1998.
- [7] T. Kessler, C. Hagemann, C. Grebing, T. Legero, U. Sterr, F. Riehle, M. J. Martin, L. Chen, and J. Ye. A sub-40-mHz-linewidth laser based on a silicon single-crystal optical cavity. *Nature Photonics*, 6:687–692, October 2012.
- [8] N. Smith and R. X. Adhikari. A cryogenic silicon LIGO upgrade. Technical Report T1400226, LIGO, 2014.
- [9] Steven D Penn, Peter H Sneddon, Helena Armandula, Joseph C Betzwieser, Gianpietro Cagnoli, Jordan Camp, D R M Crooks, Martin M Fejer, Andri M Gretarsson, Gregory M Harry, Jim Hough, Scott E Kittelberger, Michael J Mortonson, Roger Route, Sheila Rowan, and Christophoros C Vassiliou. Mechanical loss in tantala/silica dielectric mirror coatings. *Classical and Quantum Gravity*, 20(13):2917, 2003.
- [10] Gregory M Harry, Matthew R Abernathy, Andres E Becerra-Toledo, Helena Armandula, Eric Black, Kate Dooley, Matt Eichenfield, Chinyere Nwabugwu, Akira Villar, D R M Crooks, Gianpietro Cagnoli, Jim Hough, Colin R How, Ian MacLaren, Peter Murray, Stuart Reid, Sheila Rowan, Peter H Sneddon, Martin M Fejer, Roger Route, Steven D Penn, Patrick Ganau, Jean-Marie Mackowski, Christophe Michel, Laurent Pinard, and Alban Remillieux. Titania-doped tantala/silica coatings for gravitational-wave detection. *Classical and Quantum Gravity*, 24(2):405, 2007.
- [11] I W Martin, R Nawrodt, K Craig, C Schwarz, R Bassiri, G Harry, J Hough, S Penn, S Reid, R Robie, and S Rowan. Low temperature mechanical dissipation of an ion-beam sputtered silica film. *Classical and Quantum Gravity*, 31(3):035019, 2014.
- [12] Tara Chalermongsak. High fidelity probe and mitigation of mirror thermal fluctuations. 2014.
- [13] C. A. Swenson. Recommended values for the thermal expansivity of silicon from 0 to 1000 k. *Journal of Physical and Chemical Reference Data*, 12(2):179–182, 1983.
- [14] C. J. Glassbrenner and Glen A. Slack. Thermal conductivity of silicon and germanium from 3 k to the melting point. *Phys. Rev.*, 134:A1058–A1069, May 1964.
- [15] R Nawrodt, A Zimmer, T Koettig, C Schwarz, D Heinert, M Hudl, R Neubert, M Thrk, S Nietzsche, W Vodel, P Seidel, and A Tnnermann. High mechanical q-factor measurements on silicon bulk samples. *Journal of Physics: Conference Series*, 122(1):012008, 2008.

- [16] P Flubacher, AJ Leadbetter, and JA Morrison. The heat capacity of pure silicon and germanium and properties of their vibrational frequency spectra. *Philosophical Magazine*, 4(39):273–294, 1959.
- [17] Matthew Hopcroft, William D Nix, Thomas W Kenny, et al. What is the young’s modulus of silicon? *Microelectromechanical Systems, Journal of*, 19(2):229–238, 2010.
- [18] CT Taylor, M Notcutt, EK Wong, AG Mann, and DG Blair. Measurement of the coefficient of thermal expansion of a cryogenic, all-sapphire, fabry-perot optical cavity. *Optics communications*, 131(4):311–314, 1996.
- [19] R Berman, EL Foster, and JM Ziman. Thermal conduction in artificial sapphire crystals at low temperatures. i. nearly perfect crystals. In *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, volume 231, pages 130–144. The Royal Society, 1955.
- [20] Eiichi Hirose, Kieran Craig, Hideki Ishitsuka, Iain W. Martin, Norikatsu Mio, Shigenori Moriwaki, Peter G. Murray, Masatake Ohashi, Sheila Rowan, Yusuke Sakakibara, Toshikazu Suzuki, Kouichi Waseda, Kyohei Watanabe, and Kazuhiro Yamamoto. Mechanical loss of a multilayer tantala/silica coating on a sapphire disk at cryogenic temperatures: Toward the kagra gravitational wave detector. *Phys. Rev. D*, 90:102004, Nov 2014.
- [21] Warren Johnson. Private communication.
- [22] Ashot Markosyan. Private communication.
- [23] T. Chalermongsak, F. Seifert, E. D. Hall, K. Arai, E. K. Gustafson, and R. X. Adhikari. Broadband measurement of coating thermal noise in rigid Fabry-Perot cavities. *Metrologia*, 52:17, February 2015.
- [24] Chris Myatt, Nick Traggis, and Kathryn Li Dessau. Optical fabrication: Optical contacting grows more robust. *Laser Focus World*, 42, 2006.
- [25] Alix Preston, Benjamin Balaban, G Boothe, and Guido Mueller. Stable materials and bonding techniques for space-based optical systems. In *NASA Science and Technology Conference, Session C1 P3*, 2007.
- [26] Greg Harry. Mechanical loss oftra-duct 2902 epoxy. DCC: G1300207, 2013.
- [27] D.F. McGuigan, C.C. Lam, R.Q. Gram, A.W. Hoffman, D.H. Douglass, and H.W. Gutche. Measurements of the mechanical q of single-crystal silicon at low temperatures. *Journal of Low Temperature Physics*, 30(5-6):621–629, 1978.

- [28] Michael J. Martin. Quantum metrology and many-body physics: Pushing the frontier of the optical lattice clock. 2013.
- [29] Kenji Numata, Amy Kemery, and Jordan Camp. Thermal-Noise Limit in the Frequency Stabilization of Lasers with Rigid Cavities. *Phys. Rev. Lett.*, 93:250602, Dec 2004.
- [30] J.L. Hall, Long-Sheng Ma, M. Taubman, B. Tiemann, Feng-Lei Hong, O. Pfister, and Jun Ye. Stabilization and frequency measurement of the I₂-stabilized Nd:YAG laser. *Instrumentation and Measurement, IEEE Transactions on*, 48(2):583–586, apr 1999.
- [31] Adam J. Mullavey, Bram J. J. Slagmolen, John Miller, Matthew Evans, Peter Fritschel, Daniel Sigg, Sam J. Waldman, Daniel A. Shaddock, and David E. McClelland. Arm-length stabilisation for interferometric gravitational-wave detectors using frequency-doubled auxiliary lasers. *Optics Express*, 20(1):81–89, Jan 2012.
- [32] Hewlett Packard. *Phase Noise Characterization of Microwave Oscillators: Frequency Discriminator Method*. Product Note 11729. Sept 1984 1985.
- [33] P Maddaloni, P Cancio, and P De Natale. Optical comb generators for laser frequency measurement. *Measurement Science and Technology*, 20(5):052001, 2009.
- [34] I. Coddington, W. C. Swann, L. Lorini, J. C. Bergquist, Y. Le Coq, C. W. Oates, Q. Quraishi, K. S. Feder, J. W. Nicholson, P. S. Westbrook, S. A. Diddams, and N. R. Newbury. Coherent optical link over hundreds of metres and hundreds of terahertz with subfemtosecond timing jitter. *Nature Photonics*, 1(5):283–287, 05 2007.
- [35] 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(2):97–105, 1983.
- [36] Roger D Blandford and Kip S Thorne. Applications of classical physics.
- [37] Yu. Levin. Internal thermal noise in the LIGO test masses: A direct approach. *Phys. Rev. D*, 57:659–663, Jan 1998.
- [38] N Nakagawa, AM Gretarsson, EK Gustafson, and MM Fejer. Thermal noise in half-infinite mirrors with nonuniform loss: A slab of excess loss in a half-infinite mirror. *Physical Review D*, 65(10):102001, 2002.
- [39] Gregory M Harry, Andri M Gretarsson, Peter R Saulson, Scott E Kittelberger, Steven D Penn, William J Startin, Sheila Rowan, Martin M Fejer, DRM Crooks, Gianpietro Cagnoli, et al. Thermal noise in interferometric gravitational wave detectors due to dielectric optical coatings. *Classical and Quantum Gravity*, 19(5):897, 2002.

- [40] Ting Hong, Huan Yang, Eric K Gustafson, Rana X Adhikari, and Yanbei Chen. Brownian thermal noise in multilayer coated mirrors. *Physical Review D*, 87(8):082001, 2013.
- [41] Thomas Kessler, Thomas Legero, and Uwe Sterr. Thermal noise in optical cavities revisited. *J. Opt. Soc. Am. B*, 29(1):178–184, Jan 2012.
- [42] Y. T. Liu and K. S. Thorne. Thermoelastic noise and homogeneous thermal noise in finite sized gravitational-wave test masses. *Phys. Rev. D*, 62:122002, 2000.
- [43] M. Cerdonio, L. Conti, A. Heidmann, and M. Pinard. Thermoelastic effects at low temperatures and quantum limits in displacement measurements. *Phys. Rev. D*, 63:082003, Mar 2001.
- [44] D. Heinert, A. G. Gurkovsky, R. Nawrodt, S. P. Vyatchanin, and K. Yamamoto. Thermorefractive noise of finite-sized cylindrical test masses. *Phys. Rev. D*, 84:062001, Sep 2011.
- [45] V. B. Braginsky, M. L. Gorodetsky, and S. P. Vyatchanin. Thermo-refractive noise in gravitational wave antennae. *Phys. Lett. A*, 271:303, 2000.
- [46] Bradley J Frey, Douglas B Leviton, and Timothy J Madison. Temperature-dependent refractive index of silicon and germanium. In *SPIE Astronomical Telescopes+ Instrumentation*, pages 62732J–62732J. International Society for Optics and Photonics, 2006.
- [47] K. Arai, M. Abernathy, and D. Chatterjee. Design of a coating-less reference cavity with total internal reflection. Technical Report T1300472, LIGO, 2013.
- [48] M. Evans, S. Ballmer, M. Fejer, P. Fritschel, G. Harry, and G. Ogin. Thermo-optic noise in coated mirrors for high-precision optical measurements. *Phys. Rev. D*, 78:102003, Nov 2008.
- [49] Yuri Levin. Fluctuation-dissipation theorem for thermo-refractive noise. *Physics Letters A*, 372(12):1941 – 1944, 2008.
- [50] D. Heinert, A. Bell, G. Cagnoli, J. Degallaix, G. Gemme, S. Hild, J. Hough, H. Luck, I. J. Martin, S. Rowan, S. P. Vyatchanin, and R. Nawrodt. Free carrier driven noise in transmissive semiconductor optics. Technical Report P1400084, LIGO, 2014.
- [51] Harald Lück, Jerome Degallaix, Hartmut Grote, Martin Hewitson, Stefan Hild, Benno Willke, and Karsten Danzmann. Opto-mechanical frequency shifting of scattered light. *Journal of Optics A: Pure and Applied Optics*, 10(8):085004, 2008.
- [52] Peter R Saulson. Thermal noise in mechanical experiments. *Physical Review D*, 42(8):2437, 1990.

- [53] Phil Willems, Virginio Sannibale, Jaap Weel, and Valery Mitrofanov. Investigations of the dynamics and mechanical dissipation of a fused silica suspension. *Physics Letters A*, 297(1):37–48, 2002.
- [54] S Rowan, J Hough, and DRM Crooks. Thermal noise and material issues for gravitational wave detectors. *Physics Letters A*, 347(1):25–32, 2005.
- [55] SM Aston, MA Barton, AS Bell, N Beveridge, B Bland, AJ Brummitt, G Cagnoli, CA Cantley, L Carbone, AV Cumming, et al. Update on quadruple suspension design for advanced LIGO. *Classical and Quantum Gravity*, 29(23):235004, 2012.
- [56] Julius S Bendat and Allan G Piersol. *Random data: analysis and measurement procedures*, volume 729. John Wiley & Sons, 2011.
- [57] Matthew Arran. Development of a cryogenic silicon reference cavity for laser frequency stabilisation. DCC: T0900526-v1, 2013.
- [58] Kenji Numata, Jordan Camp, Michael A Krainak, and Lew Stolpner. Performance of planar-waveguide external cavity laser for precision measurements. *Optics express*, 18(22):22781–22788, 2010.
- [59] Christopher J. Erickson, Marshall Van Zijll, Greg Doermann, and Dallin S. Durfee. An ultra-high stability, low-noise laser current driver with digital control. *Review of Scientific Instruments*, 79(7):–, 2008.
- [60] M. Frede, P. King, R. Savage, and B. Willke. Advanced LIGO pre-stabilized laser conceptual design document. Technical Report T0050035-02-D, LIGO, 2005.
- [61] C. Angert and D Sigg. Characterization of a voltage controlled oscillator. Technical Report T0900451-v1, LIGO, 2009.
- [62] Gary C. Bjorklund. Frequency-modulation spectroscopy: a new method for measuring weak absorptions and dispersions. *Opt. Lett.*, 5(1):15–17, Jan 1980.
- [63] Marco Romagnoli, M. D. Levenson, and G. C. Bjorklund. Frequency-modulation-polarization spectroscopy. *Opt. Lett.*, 8(12):635–637, Dec 1983.
- [64] M. D. Levenson, W. E. Moerner, and D. E. Horne. Fm spectroscopy detection of stimulated raman gain. *Opt. Lett.*, 8(2):108–110, Feb 1983.
- [65] N. C. Wong and J. L. Hall. Servo control of amplitude modulation in frequency-modulation spectroscopy: demonstration of shot-noise-limited detection. *J. Opt. Soc. Am. B*, 2(9):1527–1533, Sep 1985.

- [66] VB Braginsky, ML Gorodetsky, and SP Vyatchanin. Thermodynamical fluctuations and photo-thermal shot noise in gravitational wave antennae. *Physics Letters A*, 264(1):1–10, 1999.
- [67] M. Cerdonio, L. Conti, A. Heidmann, and M. Pinard. Thermoelastic effects at low temperatures and quantum limits in displacement measurements. *Phys. Rev. D*, 63:082003, Mar 2001.
- [68] M De Rosa, F Marin, F Marino, O Arcizet, A Heidmann, and M Pinard. Experimental investigation of dynamic photo-thermal effect. *Classical and Quantum Gravity*, 23(8):S259, 2006.
- [69] Alessandro Farsi, Mario Siciliani de Cumis, Francesco Marino, and Francesco Marin. Photo-thermal and thermo-refractive effects in high reflectivity mirrors at room and cryogenic temperature. *Journal of Applied Physics*, 111(4):–, 2012.
- [70] T. Chalermongsak, E. D. Hall, G. D. Cole, D. Follman, F. Seifert, K. Arai, E. K. Gustafson, J. R. Smith, M. Aspelmeyer, and R. X. Adhikari. Coherent Cancellation of Photothermal Noise in GaAs/Al_{0.92}Ga_{0.08}As Bragg Mirrors. *ArXiv e-prints*, June 2015.
- [71] P. F. Cohadon, A. Heidmann, and M. Pinard. Cooling of a mirror by radiation pressure. *Phys. Rev. Lett.*, 83:3174–3177, Oct 1999.
- [72] Qiang Lin, Jessie Rosenberg, Xiaoshun Jiang, Kerry J. Vahala, and Oskar Painter. Mechanical oscillation and cooling actuated by the optical gradient force. *Phys. Rev. Lett.*, 103:103601, Aug 2009.
- [73] Jasper Chan, TP Mayer Alegre, Amir H Safavi-Naeini, Jeff T Hill, Alex Krause, Simon Gröblacher, Markus Aspelmeyer, and Oskar Painter. Laser cooling of a nanomechanical oscillator into its quantum ground state. *Nature*, 478(7367):89–92, 2011.
- [74] Aaron D O’Connell, Max Hofheinz, Markus Ansmann, Radoslaw C Bialczak, Mike Lenander, Erik Lucero, Matthew Neeley, Daniel Sank, H Wang, M Weides, et al. Quantum ground state and single-phonon control of a mechanical resonator. *Nature*, 464(7289):697–703, 2010.
- [75] Markus Aspelmeyer, Pierre Meystre, and Keith Schwab. Quantum optomechanics. *Physics Today*, 65(7):29–35, 2012.
- [76] John Miller. On non-gaussian beams and optomechanical parametric instabilities in interferometric gravitational wave detectors. 2010.
- [77] Susheel Kalia and Shao-Yun Fu. *Polymers at Cryogenic Temperatures*. Springer, 2013.

- [78] Roy H Olsson III, James G Fleming, Ihab F El-Kady, Melanie R Tuck, and Frederick B McCormick. Micromachined bulk wave acoustic bandgap devices. In *Solid-State Sensors, Actuators and Microsystems Conference, 2007. TRANSDUCERS 2007. International*, pages 317–321. IEEE, 2007.
- [79] R von Mises. Mechanik der plastischen formänderung von kristallen. *ZAMM-Journal of Applied Mathematics and Mechanics/Zeitschrift für Angewandte Mathematik und Mechanik*, 8(3):161–185, 1928.
- [80] T Suzuki, K Tsubono, and H Hirakawa. Quality factor of vibration of aluminum alloy disks. *Physics Letters A*, 67(1):2–4, 1978.
- [81] LA De Lorenzo and KC Schwab. Superfluid optomechanics: coupling of a superfluid to a superconducting condensate. *New Journal of Physics*, 16(11):113020, 2014.
- [82] Yong-Chun Liu, Rui-Shan Liu, Chun-Hua Dong, Yan Li, Qihuang Gong, and Yun-Feng Xiao. Cooling mechanical resonators to the quantum ground state from room temperature. *Physical Review A*, 91(1):013824, 2015.
- [83] Maxim Goryachev, Daniel L Creedon, Eugene N Ivanov, Serge Galliou, Roger Bourquin, and Michael E Tobar. Extremely low-loss acoustic phonons in a quartz bulk acoustic wave resonator at millikelvin temperature. *Applied Physics Letters*, 100(24):243504, 2012.
- [84] T Rocheleau, T Ndukum, C Macklin, JB Hertzberg, AA Clerk, and KC Schwab. Preparation and detection of a mechanical resonator near the ground state of motion. *Nature*, 463(7277):72–75, 2010.
- [85] Aaron D O’Connell, Max Hofheinz, Markus Ansmann, Radoslaw C Bialczak, Mike Lenander, Erik Lucero, Matthew Neeley, Daniel Sank, H Wang, M Weides, et al. Quantum ground state and single-phonon control of a mechanical resonator. *Nature*, 464(7289):697–703, 2010.
- [86] Edward A Laird, Fei Pei, Wei Tang, Gary A Steele, and Leo P Kouwenhoven. A high quality factor carbon nanotube mechanical resonator at 39 ghz. *Nano letters*, 12(1):193–197, 2011.
- [87] Joel Moser, Alexander Eichler, Johannes Güttinger, Mark I Dykman, and Adrian Bachtold. Nanotube mechanical resonators with quality factors of up to 5 million. *Nature nanotechnology*, 9(12):1007–1011, 2014.
- [88] Serge Galliou, Maxim Goryachev, Roger Bourquin, Philippe Abbé, Jean Pierre Aubry, and Michael E Tobar. Extremely low loss phonon-trapping cryogenic acoustic cavities for future physical experiments. *Scientific reports*, 3, 2013.

- [89] JD Teufel, Tobias Donner, Dale Li, JW Harlow, MS Allman, Katarina Cicak, AJ Sirois, Jed D Whittaker, KW Lehnert, and Raymond W Simmonds. Sideband cooling of micromechanical motion to the quantum ground state. *Nature*, 475(7356):359–363, 2011.
- [90] DF McGuigan, CC Lam, RQ Gram, AW Hoffman, DH Douglass, and HW Gutche. Measurements of the mechanical q of single-crystal silicon at low temperatures. *Journal of Low Temperature Physics*, 30(5):621–629, 1978.
- [91] S Ummethala, P Rath, G Lewes-Malandrakis, D Brink, C Nebel, and WHP Pernice. High- q optomechanical circuits made from polished nanocrystalline diamond thin films. *Diamond and Related Materials*, 44:49–53, 2014.
- [92] DJ Wilson, V Sudhir, N Piro, R Schilling, A Ghadimi, and TJ Kippenberg. Measurement-based control of a mechanical oscillator at its thermal decoherence rate. *Nature*, 524(7565):325–329, 2015.
- [93] Grant David Meadors, Keita Kawabe, and Keith Riles. Increasing LIGO sensitivity by feedforward subtraction of auxiliary length control noise. *Classical and Quantum Gravity*, 31(10):105014, 2014.
- [94] Norbert Wiener. *Extrapolation, interpolation, and smoothing of stationary time series*, volume 2. MIT press Cambridge, MA, 1949.
- [95] Jennifer C Driggers, Matthew Evans, Keenan Pepper, and Rana Adhikari. Active noise cancellation in a suspended interferometer. *Review of Scientific Instruments*, 83(2):024501, 2012.
- [96] Ryan DeRosa, Jennifer C Driggers, Dani Atkinson, Haixing Miao, Valery Frolov, Michael Landry, Joseph A Giaime, and Rana X Adhikari. Global feed-forward vibration isolation in a km scale interferometer. *Classical and Quantum Gravity*, 29(21):215008, 2012.
- [97] Jennifer C Driggers, Jan Harms, and Rana X Adhikari. Subtraction of newtonian noise using optimized sensor arrays. *Physical Review D*, 86(10):102001, 2012.
- [98] Bruce Allen, Wensheng Hua, and Adrian Ottewill. Automatic cross-talk removal from multi-channel data. *arXiv preprint gr-qc/9909083*, 1999.
- [99] J Aasi, J Abadie, BP Abbott, R Abbott, T Abbott, MR Abernathy, T Accadia, F Acernese, C Adams, T Adams, et al. Searching for stochastic gravitational waves using data from the two colocated LIGO hanford detectors. *Physical Review D*, 91(2):022003, 2015.
- [100] J. A. Armstrong, N. Bloembergen, J. Ducuing, and P. S. Pershan. Interactions between Light Waves in a Nonlinear Dielectric. *Phys. Rev.*, 127:1918–1939, Sep 1962.

- [101] Jorn Stenger, Harald Schnatz, Christian Tamm, and Harald R. Telle. Ultraprecise Measurement of Optical Frequency Ratios. *Phys. Rev. Lett.*, 88:073601, Feb 2002.
- [102] Er Jun Zang, Jian Ping Cao, Ye Li, Cheng Yang Li, Yong Kai Deng, and Chun Qing Gao. Realization of Four-Pass I₂ Absorption Cell in 532-nm Optical Frequency Standard. *IEEE transactions on Instrumentation and Measurement*, 56(2):673–676, April 2007.
- [103] G. Grosche, B. Lipphardt, and H. Schnatz. Optical frequency synthesis and measurement using fibre-based femtosecond lasers. *The European Physical Journal D - Atomic, Molecular, Optical and Plasma Physics*, 48(1):27–33, 2008.
- [104] T. Rosenband, D. B. Hume, P. O. Schmidt, C. W. Chou, A. Brusch, L. Lorini, W. H. Oskay, R. E. Drullinger, T. M. Fortier, J. E. Stalnaker, S. A. Diddams, W. C. Swann, N. R. Newbury, W. M. Itano, D. J. Wineland, and J. C. Bergquist. Frequency Ratio of Al⁺ and Hg⁺ Single-Ion Optical Clocks; Metrology at the 17th Decimal Place. *Science*, 319(5871):1808–1812, March 2008.
- [105] T. M. Fortier, N. Ashby, J. C. Bergquist, M. J. Delaney, S. A. Diddams, T. P. Heavner, L. Hollberg, W. M. Itano, S. R. Jefferts, K. Kim, F. Levi, L. Lorini, W. H. Oskay, T. E. Parker, J. Shirley, and J. E. Stalnaker. Precision Atomic Spectroscopy for Improved Limits on Variation of the Fine Structure Constant and Local Position Invariance. *Phys. Rev. Lett.*, 98:070801, Feb 2007.
- [106] B. P. Abbott, R. Abbott, R. Adhikari, P. Ajith, B. Allen, G. Allen, R. S. Amin, S. B. Anderson, W. G. Anderson, M. A. Arain, and et al. LIGO: the Laser Interferometer Gravitational-Wave Observatory. *Reports on Progress in Physics*, 72(7):076901, July 2009.
- [107] Kiwamu Izumi, Koji Arai, Bryan Barr, Joseph Betzwieser, Aidan Brooks, Katrin Dahl, Suresh Doravari, Jennifer C. Driggers, W. Zach Korth, Haixing Miao, Jameson Rollins, Stephen Vass, David Yeaton-Massey, and Rana X. Adhikari. Multicolor cavity metrology. *J. Opt. Soc. Am. A*, 29(10):2092–2103, Oct 2012.
- [108] A Staley, D Martynov, R Abbott, RX Adhikari, K Arai, S Ballmer, L Barsotti, AF Brooks, RT DeRosa, S Dwyer, et al. Achieving resonance in the advanced LIGO gravitational-wave interferometer. *Classical and Quantum Gravity*, 31(24):245010, 2014.
- [109] Adam J Mullavey, Bram JJ Slagmolen, John Miller, Matthew Evans, Peter Fritschel, Daniel Sigg, Sam J Waldman, Daniel A Shaddock, and David E McClelland. Arm-length stabilisation for interferometric gravitational-wave detectors using frequency-doubled auxiliary lasers. *Optics express*, 20(1):81–89, 2012.

- [110] Farid Khalili, Stefan Danilishin, Helge Müller-Ebhardt, Haixing Miao, Yanbei Chen, and Chun-nong Zhao. Negative optical inertia for enhancing the sensitivity of future gravitational-wave detectors. *Phys. Rev. D*, 83:062003, Mar 2011.
- [111] B. Willke, N. Uehara, E. K. Gustafson, R. L. Byer, P. J. King, S. U. Seel, and Jr. R. L. Savage. Spatial and temporal filtering of a 10-w nd:yag laser with a fabry–perot ring-cavity premode cleaner. *Opt. Lett.*, 23(21):1704–1706, Nov 1998.
- [112] Volker Leonhardt and Jordan B. Camp. Space interferometry application of laser frequency stabilization with molecular iodine. *Appl. Opt.*, 45(17):4142–4146, Jun 2006.
- [113] Yuri Levin. Fluctuation–dissipation theorem for thermo-refractive noise. *Physics Letters A*, 372(12):1941–1944, 2007.
- [114] D. Heinert, A. G. Gurkovsky, R. Nawrodt, S. P. Vyatchanin, and K. Yamamoto. Thermorefractive noise of finite-sized cylindrical test masses. *Phys. Rev. D*, 84:062001, Sep 2011.
- [115] Long-Sheng Ma, Zhiyi Bi, Albrecht Bartels, Lennart Robertsson, Massimo Zucco, Robert S. Windeler, Guido Wilpers, Chris Oates, Leo Hollberg, and Scott A. Diddams. Optical Frequency Synthesis and Comparison with Uncertainty at the 10^{-19} Level. *Science*, 303(5665):1843–1845, 2004.
- [116] L. Robertsson. International comparison of $^{127}\text{I}_2$ -stabilized frequency-doubled Nd:YAG lasers between the BIPM, the NRLM and the BNM-INM. *Metrologia*, 38(567), 2001.
- [117] J. E. Stalnaker, S. A. Diddams, T. M. Fortier, K. Kim, L. Hollberg, J. C. Bergquist, W. M. Itano, M. J. Delany, L. Lorini, W. H. Oskay, T. P. Heavner, S. R. Jefferts, F. Levi, T. E. Parker, and J. Shirley. Optical-to-microwave frequency comparison with fractional uncertainty of 10^{-15} . *Applied Physics B: Lasers and Optics*, 89:167–176, November 2007.
- [118] Rubiola. *Phase Noise and Frequency Stability in Oscillators*. Cambridge University Press, 2008.
- [119] International Radio Consultative Committee. Characterization of Frequency and Phase Noise. Report, 1986.
- [120] Kiwamu Izumi, Koji Arai, Bryan Barr, Joseph Betzwieser, Aidan Brooks, Katrin Dahl, Suresh Doravari, Jennifer C. Driggers, W. Zach Korth, Haixing Miao, Jameson Rollins, Stephen Vass, David Yeaton-Massey, and Rana X. Adhikari. Multi-wavelength cavity metrology. arXiv:1205.1257, May 2012.
- [121] O. Terra, G. Grosche, K. Predehl, R. Holzwarth, T. Legero, U. Sterr, B. Lipphardt, and H. Schnatz. Phase-coherent comparison of two optical frequency standards over 146 km using

- a telecommunication fiber link. *Applied Physics B: Lasers and Optics*, 97:541–551, 2009. 10.1007/s00340-009-3653-2.
- [122] Huan Yang, Haixing Miao, Da-Shin Lee, Bassam Helou, and Yanbei Chen. Macroscopic quantum mechanics in a classical spacetime. *Physical review letters*, 110(17):170401, 2013.
- [123] Anthony E Siegman. *Lasers*. Mill Valley, CA, 37, 1986.
- [124] Stuart Aston. BOSEM design document. DCC:T050111, 2009.
- [125] Kenneth E Torrance and Ephraim M Sparrow. Theory for off-specular reflection from roughened surfaces. *JOSA*, 57(9):1105–1112, 1967.
- [126] John C Stover. *Optical scattering: measurement and analysis*, volume 2. SPIE optical engineering press Bellingham, WA, 3rd edition, 2012.
- [127] Fabian Magaña-Sandoval, Rana X Adhikari, Valera Frolov, Jan Harms, Jacqueline Lee, Shannon Sankar, Peter R Saulson, and Joshua R Smith. Large-angle scattered light measurements for quantum-noise filter cavity design studies. *JOSA A*, 29(8):1722–1727, 2012.