

A Fully Crystalline Cryogenic Reference Cavity

by

Dhruv Kedar

B.A., University of Chicago, 2013

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
Department of Physics
2023

Committee Members:

Jun Ye, Chair

Andrew Ludlow

James K. Thompson

Thomas R. Schibli

Svenja Knappe

Kedar, Dhruv (Ph.D., Physics)

A Fully Crystalline Cryogenic Reference Cavity

Thesis directed by Prof. Jun Ye

To date, neutral atom optical lattice clocks have demonstrated the highest precision measurements of fractional frequency shifts. Frequency noise of the optical local oscillator, used to drive the clock transition, is the primary limitation for how quickly these clocks can measure a quantity. Improving the stability of these oscillators is therefore crucial in enabling efficient characterization of physics at smaller energy scales. State-of-the-art optical oscillators employing cryogenic reference cavities have been used to push this frontier, enabling record-level stabilities and precision of neutral atom optical clocks. The exceptional stability of these oscillators has resulted in foundational advances in optical frequency metrology, establishing these systems as worthy scientific pursuits in expanding the role of measurement science.

The fundamental performance of cryogenic cavities utilizing crystalline spacers and substrates has been limited by the Brownian thermal noise associated with mechanical dissipation of the mirror coatings. Recently, crystalline $\text{Al}_{1-x}\text{Ga}_x\text{As}/\text{GaAs}$ coatings have emerged as a promising candidate for reduced coating thermal noise. We present measurements of the frequency noise of a fully crystalline cryogenic reference cavity with $\text{Al}_{0.92}\text{Ga}_{0.08}\text{As}/\text{GaAs}$ optical coatings.

We report on previously unmeasured birefringent noise of crystalline coatings associated with anti-correlated frequency fluctuations between the two polarization modes of the cavity, and identify variables that affect its magnitude. We present phenomenological characterizations of this noise and implement an interrogation scheme that cancels it by simultaneous probing of both polarization modes. These results challenge our current understanding of noise processes generating length fluctuations and introduce our discovery of novel light induced frequency noise.

Dedication

To family and friends for their unconditional love and support. You've taught me what that means. To J. M. W, M. R. C., and J. H., I wouldn't be here without you.

Acknowledgements

Like most sections in this thesis, this will be extremely verbose. But, there is a village to thank.

First, Jun and Debbie. Debbie took on the task of accepting me into her group and JILA when I had few experimental skills. She quickly taught me that there are no parts of the experiment that one can afford to shy away from, i.e. “black boxes” should not really be black boxes. On several occasions Debbie came down to the lab to walk me through various parts of the machine that I could not follow by myself. After the rotating gas machine was decommissioned, Jun accepted me into his group. I am indebted to Jun and Debbie for resolving to find a happy home for her remaining students. Both have been some of the greatest advocates for the well being of their students.

Jun is one of the most incredible scientists I have come across; coupled with an infectious attitude towards stoking our excitement in lab, it is no surprise that the group has developed into a powerhouse capable of tackling the most ambitious of experiments (or Jun’s words, ‘nontrivial’). I will miss many of your poetic physics-related metaphors and I cannot thank you enough for taking the time to nurture me into the scientist I am today. I drifted through undergraduate lab courses thinking that measurement noise was an irritating byproduct of taking data but I will never forget your statement that “without noise, physics would be boring”. Both Debbie and Jun taught me that values and methods of communicating your work effectively and I hope I can do you both justice.

The stable lasers team of Wei Zhang, Eric Oelker, John Robinson, and Will Milner, thank

you all for being steadfast colleagues who are willing to wade through the muddy trenches of cavity work. Eric Oelker in particular brought in his LIGO expertise and taught me nearly everything I know about low noise measurements. Thank you for teaching me that most things can be simulated, and when they can't you build the experiment. Wei is one of the kindest and hardest workers that I will likely ever know, I always appreciate you taking the time to relay historical knowledge of Si4. Si6 has largely been built off the tenacity of grad students past, John and Will. Both are extremely careful scientists who find pleasure in the minute details and I can always count on you to engage on peculiarities that others might find trite.

Thanks to the rest of the strontium group members who I've interacted with over the years: Colin Kennedy, Toby Bothwell, Shimon Kolkowitz, Sarah Bromley, Ed Marti, Sara Campell, Aki Goban, Christian Sanner, Ross Hutson, Lindsay Sonderhouse, Lingfeng Yan, Alex Aeppli, Maya Miklos, Yee Ming Two, Will Warfield. You've all been great mentors and co-workers and I wish I had the space to single out memorable interactions I've had with all of you. The larger Ye group family on the combs, KRb, YO, and OH, I've enjoyed your friendships, knowledge, and meaningful conversations we've shared.

Jan Hall has always been a kind presence in S1B20 and some of my fondest memories from that room are the long impromptu conversations we would have on cold cavity drift rates or the merits of particular voltage regulators. The electronics shop has also been invaluable in all the experiments in the Ye group. It turns out that pushing frontiers of frequency metrology go hand in hand with pushing frontiers of circuit stability. Terry Brown has been an invaluable mentor on this regard. Everything I know, and do not know with regards to electronics can be attributed to him so please seek him out if my circuits present issues. Much of the work in this thesis would have been significantly more difficult to realize without his magic resonant photodetectors. James Fung-a-Fat and Carl Sauer have always been reliable resources when it comes to understanding or repairing the wizardry of JILA PCBs. The machine shop has similarly been invaluable to Ye group research efforts, and a consistent source of pleasure to interact with. James, Hans, Kyle, Calvin, and Adam are always available for assistance with any mechanical problem and can be counted on to add

an appropriate level of humor to any issue. Adam, Calvin, and James in particular constructed some integral components for the Si6 cryogenic environment. Computing gurus JR, Corey, Dan, and Eric provided support necessary to keep the lab work progressing forward. Special thanks go to the JILA purchasing department who have gracefully handled numerous rush orders with complaint and who can always be counted on to reliably advocate on our behalves with respect to non-CU entities. Krista and Amy also deserve a special recognition for patiently dealing with my administrative ineptitude. JILA could not ask for a better visionary than Beth, thanks allowing us to spend most of our efforts focusing on science.

Our cavity counterparts at PTB have been a delight to collaborate with. Anyone will be hard pressed to find a kinder and more generous set of individuals than Uwe, Thomas, and Fritz - each of whom are treasure troves of knowledge relating to the history and development of the silicon cavity effort. Recent addition Jialiang You was a close collaborator in understanding the novel noise mechanisms associated with crystalline coatings. Without his tenacity and propensity for careful measurements, our current understanding of AlGaAs coating effects would be lacking. Garrett Cole has also been a great collaborator, always interested in exploring new phenomena related to his crystalline coatings.

It takes a special environment to be able to call everyone you interact with a friend. By necessity, JILA has been a home but I could not ask for a better environment to spend a good portion of my life. Several of my friendships have been formed in this building, and thank you all for showing me that there is a world outside of lab. To the of dodos of Rocky Mountain Runner, the wise old sages of Boulder Track Club, and fearless minions of OSMP, thank you for keeping me sane and helping me reassess what is physically and mentally possible. To my family, thank you for being there and always supporting my interests. And to Boba, Attila, and Sunny, thanks for showing me that there's another component to happiness.

Contents

Chapter

1	Introduction	1
1.1	Thesis outline	3
2	Oscillators and stabilization	10
2.1	Fabry-Perots	10
2.2	Thermal noise in cavities	14
2.2.1	Normal mode method	16
2.2.2	Direct approach	19
2.2.3	Temperature dependent damping	23
2.3	Cavity locking	24
2.4	Frequency stability of oscillators	30
2.5	Frequency stability of atom-steered oscillators	36
3	Experimental setup	43
3.0.1	Silicon in cavities	43
3.0.2	AlGaAs coatings	46
3.0.3	Cryostation	52
3.0.4	Noise measurements	56
4	Technical noise sources	59
4.1	Laser stabilization	59

4.2	Residual amplitude modulation	67
4.2.1	RAM from the modulator	69
4.2.2	AM and PM on the phase reference	72
4.2.3	Etalon effects	75
4.3	Vibration sensitivity	78
4.3.1	Measuring cavity sensitivity	81
4.3.2	Improving the sensitivity	84
4.3.3	Feedback to AVI	88
4.4	Thermo-expansive length stabilization	91
4.4.1	Heat flow	93
4.4.2	Precision thermometry	102
4.4.3	Cooldown tests	104
4.4.4	Thermal time constants	109
4.4.5	Temperature stabilization	111
4.5	Vacuum pressure	113
4.6	Dissemination noise	117
4.7	Technical noise budget	120
5	Frequency noise of the cavity	123
5.1	Photo-birefringent noise	124
5.1.1	Photothermal effect	124
5.1.2	Photothermal effects in crystalline coatings: transient response	128
5.1.3	Photothermal effects in crystalline coatings: coherent response	133
5.2	Frequency instability of the cavity	138
5.2.1	Mode area dependence	144
5.2.2	Correlation of light induced noise	149
5.2.3	Dual frequency locking	150

5.2.4	Residual noise	154
5.3	Direct measurement of birefringent noise	158
5.3.1	Two optical paths in one setup	159
5.3.2	Asynchronous detection of cavity error signals	160
5.3.3	Simultaneous detection	169
5.4	Long term instability	174
5.4.1	Etalon identification	175
5.4.2	Pointing stability	180
5.4.3	Top setup	182
5.5	Drift rate	184
6	Novel applications of crystalline cavities	188
6.1	Ratio Comparison	188
6.2	Dark matter searches with oscillators	194
6.3	Timescale	201
6.4	Gravitational wave detection	204
7	Conclusion	209
	Bibliography	213
	Appendix	
A	Dick effect limited instability and laser noise	227
B	Analog locking	233

Tables

Table

2.1	Relations between the fractional frequency noise spectrum $S_y(f)$, variance $\sigma_y^2(\tau)$, and modified variance mod $\sigma_y^2(\tau)$. Calculation of these relations is taken from [1]. . .	32
3.1	Comparison of bulk substrate and spacer materials.	46
3.2	Parameters of the AlGaAs/GaAs coatings	51
3.3	Resonances of Si3 identified in the noise model, with center frequency f_i , amplitude a_i , and width Γ_i	58
4.1	Comparison of low frequency AC sensitivities and DC sensitivities.	84
5.1	Fast and slow time constants (in seconds) of a double exponential fit the data in Fig.5.7, and analogous data taken when locked to the other cavity mode.	131

Figures

Figure

- 1.1 **Left:** Plot of the time record of all the SRI systematic shifts. Changes in atom number, ambient temperature, or magnetic field all result in corrections to the clock frequency, and their total magnitude is shown over a six hour data campaign. The clock achieves over 98% uptime over the course of this single comparison day and slight gaps in the data indicate brief periods where the laser is not locked to the atoms. **Right:** The same data is plotted as a fractional instability normalized to the Sr clock frequency. The individual contributions of density shift (blue), BBR (red) and second order Zeeman shift (yellow) are shown as the dashed curves. For operation times up to 10^4 s, fluctuations in systematic offsets are bounded below 4×10^{-19} 5
- 1.2 Density shift evaluations. Density shift measurements were performed over four weeks at the same trap conditions. The weighted mean of all measurements is shown with a dashed, black line, with a reduced chi-squared of 1.07 5
- 1.3 Non-synchronous comparison (i.e. interrogation times of the local oscillator do not overlap such that there is sensitivity to laser noise) with the JILA 3D optical lattice clock demonstrates that the beat between the two clocks averages below the quoted total systematic uncertainty. The black line is a white noise $\tau^{-1/2}$ fit to the single clock instability. 6

- 2.1 **Left:** Calculation of transverse modes for Si6. The x-axis refers to a total mode number ($m + n$). The closest degenerate mode to HG_{00} is one where $m + n = 10$, and this is still 6.4 MHz away. **Right:** We measure the cavity couplings to some of the low order modes with $m + n \leq 3$. The highest measured coupling is to HG_{00} with $\eta = 0.5$. The sum of couplings for displayed modes is still less than 1, indicating that there is a finite coupling to other higher modes. We were unable to measure these couplings, indicating that the residual $\eta \sim 0.4$ is likely distributed over very many modes. 14
- 2.2 **Left:** We calculate the coupling efficiency for an incident Gaussian beam to the HG_{00} mode of a cavity. Both beams have the same waist w_0 but the incident beam is displaced from the cavity mode in the transverse direction. When the radial displacement δr is equivalent to the cavity waist, the coupling has already dropped to 0.3. **Right:** We now assume that the incident field is perfectly overlapped with the cavity mode in the transverse direction, but the waist of the incident mode w_{in} differs from waist of the HG_{00} mode. We see that the coupling is much less sensitive to variations in w_{in} 15
- 2.3 We consider an oscillator with resonance at 1 kHz and varying mechanical loss angle ϕ . **Top:** Transfer function $\text{Re}[\chi(\omega)]$. **Bottom:** Displacement noise spectrum $S_x(\omega)$. In both cases a lower value of ϕ results in a sharper resonance but only in the displacement noise can we see an effect of the quality factor on the low frequency region. Mechanical resonances of the cavity components are always designed to be above 100 s of Hz so $\omega \ll \omega_0$ is the regime of interest. 18
- 2.4 Reproduced from [2], the x 's denote the noise amplitude contribution from the first 100 normal modes of a 10 cm mirror with a 2.2 mm beam spot. The solid line illustrates the cumulative noise of all these modes and the summation has still not converged over the 100 modes. Notably the mirror diameter is crucial for this calculation whereas it is unnecessary in the direct approach. 20

2.5	Evidence of a temperature dependent loss peak in crystalline silicon from [3]. A feature at 160 K is discovered to be associated with impurities in the silicon. An acid etch of the ingot appears to remove the loss peak.	23
2.6	Comparison of lineshapes of the various locking schemes discussed. The cavity transmission function (red) is used to enact side of line locking. Its derivative is used in White's method [4] for modulation the laser frequency. The true dispersive lineshape is obtained via the PDH [5] and Hänsch-Couillard schemes [6].	26
2.7	Top left: Relative power in the carrier, first, second, and third order sidebands for varying modulation index β . The total optical power is normalized to unity. At $\beta = 1$, roughly half the power is contained in the carrier while each first order sideband contains 25%. Top right: Discriminator slope $\mathcal{D}^{-1}$ for different modulation depth. Again, maximum sensitivity is achieved close to $\beta = 1$. Note that there are values of modulation index where there is nearly no sensitivity to the PDH signal. At $\beta = 2.3$, power in the carrier drops to zero meaning no light enters the cavity. At $\beta = 3.8$, there is still power in the carrier but there is no electric field oscillating at 1Ω . A PDH signal can be recovered by using the 2Ω sidebands as a phase reference. Bottom: Error signals measured for varying modulation index. Curves are offset vertical and horizontally for clarity.	29
2.8	Calculation of the frequency domain filter functions $ W_a(f) $ and $ W_m(f) $ for the Allan and modified Allan variances. The modified Allan variance has less susceptibility to high frequency noise	33
2.9	Calculation of the time domain filter functions $w_a(t)$ and $w_m(t)$ for the Allan and modified Allan variances.	35
2.10	Simulation that compares the variance and modified variance for different values of α where $S_y(f) \propto f^\alpha$	35

- 2.11 **Top:** We simulate datasets with frequency noise spectrum $S_y(f) \sim f^\alpha$ for $\alpha = -2, -1, 0$. Discrete frequency jumps of the same magnitude are added to each time series. **Bottom:** Solid circles indicate the modified variance calculated for the corresponding datasets from the plot above. Dashed lines are variances of the unperturbed white, pink, and brown noise without the discrete frequency jump. The frequency jump has the effect of deteriorating the stability of all three time series but also changes the long term instability to scale as $\sigma_y(\tau) \sim \tau^{0.5}$. Data that is dominated by white frequency noise and flicker frequency noise can still have variances that corresponding to random walk frequency noise if there is a discrete perturbation. 37
- 2.12 Calculation of the time domain filter functions $r_{Rabi}(t)$ and $R_{Ramsey}(t)$ for the two spectroscopic methods. In both cases we assume a cycle time of 5 s, a spectroscopy time of 3 s, and in the case of Ramsey a π pulse time of 100 s. 40
- 2.13 We calculate the Dick effect via Eq.2.66 for different laser spectra $S_y(f) = 10^{-33} f^\alpha$ with T_c . For reference, $S_y(f) = 10^{-33} f^{-1}$ is the thermal noise limit of Si3. The corresponding Allan deviation is insensitive to the type of noise that dominates the laser spectrum and always follows $\sigma_y(\tau) \sim \tau^{-1/2}$. However the amplitude of the Dick effect is determined by the slope of the laser spectrum, with a lower variance associated with larger values of α . Solid lines indicate variances calculated with the Rabi sensitivity function and dotted lines are calculated with the Ramsey sensitivity. As α decreases, Ramsey spectroscopy provides increasing improvements in stability over Rabi spectroscopy. 41
- 2.14 The various amplitude coefficients $|g_k/g_0|^2$ are calculated for different spectroscopic method. All traces use $T_c = 5$ s, $\tau_d = 4$ s. The gold curve uses a conservative $t_\pi = 0.2$ s whereas the green trace assumes infinitely short pulses of $t_\pi = 0$. The majority of the Rabi sensitivity is concentrated around frequencies of $1/T_c$ whereas Ramsey spectroscopy has a higher contribution from the high frequency harmonics of the cycle time. 42

3.1	Loss angle ϕ temperature dependence for commonly used cavity spacer materials.	
	ULE is not shown on this plot but has a measured room temperature loss angle of 1.6×10^{-5} [7]. Crystalline spacer materials have a ϕ so low that the total Brownian thermal noise limit of crystalline cavities is limited by the loss angle of the optical coating. Figure taken from [8].	45
3.2	Coefficients of thermal expansion for commonly used cavity spacer materials. Between the two crystalline materials of sapphire and silicon, only the latter has a ZCTE. Fused silica can be an attractive option due to a ZCTE around 170 K but the significantly higher loss angle means that a silicon cavity will still have a superior thermal noise limited performance. Figure taken from [8].	45
3.3	Coefficient of thermal expansion of silicon across the three ranges of operation, 124 K, 16 K, and 4 K. At 124 K, $\alpha' = 2 \times 10^{-8} \text{ K}^{-2}$, similar to the CTE slope around the ULE room temperature ZCTE. The slope is considerably shallower at 17 K with $\alpha' = 1 \times 10^{-9} \text{ K}^{-2}$ but requires different cooling technology. There is no zero crossing of α at 4 K but values of 10^{-11} can be achieved, requiring only μK temperature stability for 1×10^{-17} cavity instability.	47
3.4	Two 55 cm silicon ingots grown at IKZ Berlin. Image courtesy of Thomas Legero. .	47
3.5	Ringdown measurement of Si6 transmitted light with the cavity at vacuum and cooled to 4 K. No difference in finesse is measured between 4 K and 16 K. A finesse of 290,000 is measured.	50
3.6	Brownian noise contributions for the four silicon cavities referenced through this thesis. Coating thermal noise is consistently the dominant contribution due to the exceptionally low loss angle of silicon. Coating materials with high mechanical quality factor are therefore instrumental in developing optical oscillators with longer coherence times.	51

3.7	Reflectance of the two crystalline mirrors used on Si6 at a temperature of 4 K. Data is from Garrett Cole. We use a 1542 nm laser to lock to the cavity, centered within the ~ 150 nm passband of the coatings.	52
3.8	The mode splitting of Si6 is observed to vary with temperature and intracavity power. We plot the change in mode splitting normalized to the 8.6 kHz cavity linewidth. . .	52
3.9	A capacity map for our cryostation's cooler displays the cooling power available on each stage. We operate in the lower right corner of the plot with less than 100 mW heatload on stage 2 and roughly 1 W heatload on stage 1.	54
3.10	Vibration noise measured on the AVI with vibration isolation enabled for all three axes.	55
3.11	A solidworks mockup of the Si6 cryostation, cavity, and part of the optical setup. . .	57
4.1	(Left) A comparison of frequency noise between the RIO and NKT obtained via heterodyne beat with cavity stabilized light. (Right) The closed loop gain necessary to push the free running laser noise to the Si6 thermal noise at a given offset frequency.	61
4.2	AOM1 is used for both intensity and frequency feedback but it does not exactly actuate along the two orthogonalized quadratures. A transfer function of the cross coupling between frequency to optical power incident on the cavity is shown in (a). Intensity and frequency servo loops are also shaped to suppress any cross couplings. Feedback to the AOM frequency is highpassed with a 50 Hz corner while the intensity servo has an integrator enabled at a 100 Hz corner. The measured cross coupling of the AOM between frequency and intensity is small enough that the control signal from each cannot introduce noise in the opposite quadrature above the Si6 thermal noise.	61
4.3	Voltage-to-frequency transfer function of the VCO to AOM on the Si3 and Si6 systems. The delay of the Si6 modulator is slightly lower in frequency and consequently the maximum attainable loop bandwidth is lower.	62

- 4.4 Topology of the Si6 PDH loop. G_C , G_{PZT} , and G_{AOM} denote the loop filters for the common, piezo, and AOM stages. Responses of the actuators are denoted by A_{PZT} and A_{AOM} . Plant responses of the cavity filter function and the limit imposed by the photodetector bandwidth are given by P_{cav} and P_{PD} 62
- 4.5 (a) The input amplification stage of the PDH servo board uses a low-noise FET with only 40dB of gain at low frequency. This pushes the input referred noise of the board to the $10 \text{ nV}/\sqrt{\text{Hz}}$ regime in the low frequency regime. Combined with a large discrimination slope achieved with the typical operation of a few microwatts incident on the cavity (b), the servo noise can be pushed to a frequency noise limit that is insignificant compared to the expected Si6 thermal noise. A large input referred noise can always be alleviated by using a corresponding steeper discrimination, but comes at the cost of increased sensitivity to photothermal effects. 63
- 4.6 (a) The closed loop actuator crossover illustrates the relative loop shaping between the two actuators. The phase measurement is informative in identifying the first, and higher order mechanical resonances of the piezo. 63
- 4.7 Summing an excitation into the closed loop yields the complete transfer function with all plant responses included. In the phase domain, this is particularly useful for identifying plant response poles in the high frequency. Modeling the closed loop phase reveals an additional measured pole at 50 kHz which is identified as the PDH photodetector response. This can then be appropriately compensated by adding a zero in the loop. The 2π phase wrap from 10 - 2000 Hz identifies the four integrator stages implemented with corners throughout that band. 64
- 4.8 Calculation of the fundamental shot noise limited length resolution on Si6. We account for the imperfect cavity coupling. Dashed lines indicated the thermal noise limit of different cavities. 68

- 4.9 Diagram of the differential optical phase that can be acquired between the carrier and modulation sidebands after their generation. Red phases are detected in the RAM and green phases are differential to the PDH path. RAM and PDH feedback to the phase errors measured on each photodetector are applied to the EOM and AOM locations respectively. 72
- 4.10 **Top:** PDH error signals with $\Omega = 5$ MHz, PM of $\phi_+ = \phi_i = \pi/6$, and AM of $\alpha_+ = 1/\alpha_- = 1.5$. The **Left** column displays demodulation phase of to extract the quadrature component, and $\varphi = \pi/3$ is on the **Right**. Curves are offset horizontally for clarity but not vertically. We can see that PM results in a large offset of the baseline signal that is not removed at $\varphi = 0$. For AM this is not the case though nonzero demodulation phase will reveal a large offset. **Bottom:** We show the magnitude of the PDH slope in all three cases. Dashed black lines indicate an offset frequency of half the cavity linewidth. At $\varphi = 0$ PM has the surprising feature that maximum sensitivity to the error signal does not occur at an offset of 0 Hz. Additionally the minimum sensitivity is shifted from the expected location at half the cavity linewidth. 74
- 4.11 Simulating the effect of a 1 MHz linewidth etalon with 20% coupling stationed 3 MHz away from the cavity resonance. A significant distortion of the error signal occurs, and the etalon will always decrease the effective discrimination slope. Coupling to the etalon will typically be lower than the parts per thousand level if all optics have high quality AR coatings. The parameters are purposely exaggerated here to demonstrate the etalon effect. 76

- 4.12 **Top:** Fractional frequency noise of the out of loop RAM characterization. Though far from the projected Si6 thermal noise, the cavity instability should be limited to a level similar to the Si3 thermal noise. This would mark a significant improvement over the instability of the Si4 system with the added benefit that the 4 K closed cycle cryostat makes Si6 a hands-free oscillator. Dark noise corresponds to noise of the resonant photodetector. A high white noise level is visible at frequencies larger than 1 Hz. This specific dataset coupled to the TEM01 mode of the cavity without a phase plate. The poor mode coupling leads to a correspondingly higher contribution of shot noise. **Right:** Modified Allan deviation of traces on the left. 77
- 4.13 Vertical acceleration sensitivity of the cavity. Si6 is supported at the midplane and we measure a sensitivity that is roughly centered at zero. We rotate the cavity on its mount to be near one of these zero crossings. 80
- 4.14 Transfer functions (**left**) and phase responses (**right**) of the three acceleration responses to frequency when driving the AVI along a particular axis. Different rows display the data of shaking the AVI along each direction. The acceleration response to driving the platform along gravity (bottom row) is largely contained along $\hat{V}$. This is less the case for driving along the horizontal axes where we can see a few cross couplings at low frequency. This likely suggests that the platform is torqued in the H1-H2 plane, shaking along the AVI's horizontal axes generates an admixture of excitation along H1 and H2. In the phase domain, we can identify some sort of mechanical resonance of the cryostation near 100 Hz, where phase wrappings are evident in all axes. 83

4.15	The frequency dependence of measured acceleration sensitivities on Si6. The vertical sensitivity is notably higher than the mid- 10^{-12} s value that can be attained by adjusting the spacer rotation. We observed that the values of the horizontal sensitivity had a slight dependence on rotation angle as well. The current orientation of the spacer was chosen to minimize the largest sensitivity (in this case k_{H1}) without significantly increasing the others.	84
4.16	Vibration induced frequency noise measured along the different axes and their respective couplings to the cavity.	85
4.17	Various cavity mounts that we tested for improved sensitivity to acceleration noise.	
	Top left: The original folded cup design from Si4, made from G10 composite.	
	Top right: A simplified structure made from G10. Bottom left: A monolithic stainless steel support machined a block of 316 stainless. Hemispherical divots have been milled out from the top of each post. Spheres of different material can be placed within each divot, the image shows a set of teflon balls resting within the indentations. Bottom right: Our final cavity support made entirely from stainless steel	86

- 4.18 Measured sensitivities for the variety of cavity mounts we considered. Label "New G10 mount" refers to the green G10 cup from Fig.4.17, "Fully S316" refers to the stainless steel mount with rounded tips, and "Divot" is the stainless mount with a hemisphere milled out from the top of the supporting posts. We place stainless steel, nylon, and teflon spheres in the divots and rest the cavity on top. All traces have been smoothed with a 5th-order spline to improve visual clarity. Horizontal sensitivities are the more meaningful quantity since we know that k_V can be zeroed by adjusting the cavity-mount angle. Variation of k_V between the different mounts is due to imperfect placement of the cavity. All k_i steadily increase above 100 Hz, indicating that some other component of the cryostation is contributing to an increased acceleration sensitivity at high frequency. A large rollup in the k_{H1} is visible for all the different cavity supports. At the time of these measurements, we were suspected that this was not representative of the true H1 sensitivity but were unsure of its source. We later understood that driving the AVI along the H1 axis (collinear to the coldfinger axis) would excite a compressive mode along the coldhead's bellows. This generates air currents that modulate the laser phase coherently with the cavity's acceleration response and appear as an artificial k_{H1} signal. We later remove this effect by pulling a vacuum on the cryostation, and its absence is noted in Fig.4.15. 89
- 4.19 Vibrations measured along H1 with the AVI isolation enabled are shown in blue. When we use a lower noise external accelerometer to measure the acceleration spectrum, we can feedback to the AVI by modulating its displacement along H1 to further reduce the measured noise (black). At frequencies below 40 Hz, the black curve is limited by the noise floor of our new in-loop accelerometer. 91

4.20	Plots of the coherent response measured along each axis when driving the AVI along one direction. In the top plot for example, the AVI is driven along H1 and we show the corresponding transfer functions $a_i/\sqrt{\sum_i a_i^2}$. The vast majority of the coherent displacement occurs along H1 at low frequency, though in regions around 50 Hz and 80 Hz, an equal response is measured in the vertical axis. Single axis feedback along H1 would exhibit strong cross coupling to another axis at these frequencies.	92
4.21	Schematic of the thermal circuit that models heat transfer between the 30 K stage and the cavity. Radiative and conductive pathways are placed in parallel, and the low pass filter is completed with the inclusion of materials with high thermal capacitance. An additional lowpass should be included to capture the heat from the 300 K vacuum chamber to the cryogenic 30 K stage, but details this support structure within the Montana Instruments cryostation are proprietary, and are not shared with us. We calculate an active-passive time constant of ~ 10000 s and a passive-cavity time constant of 600 s. Note that the passive-cavity time constant is reduced from the 1200 s modeled and measured in Si4. The G10 cavity support in Si4 is replaced with the aforementioned stainless steel mount which has a two times lower thermal resistivity. Unlike the 124 K system, the thermal capacitance of silicon is very small at 4 K and results in a low thermal time constant between the baseplate of the passive shield and the cavity mirrors.	95
4.22	Volumetric heat capacities of materials typically used in a cryogenic environment. Common materials (dark blue curves) have a vanishingly small heat capacity below 10 K, requiring use of more exotic alloys.	97
4.23	Thermal conductivity of high purity copper with different residual resistivity ratios.	101
4.24	Various gold-plated components for the thermal shields. The radiation shield is on the left, 4 K active shield in the middle, and 4 K passive shield on the right. A baffle with a 5 mm aperture on the radiation limits the room temperature radiation load passing through the windows to the cavity.	101

- 4.25 **Left:** Temperature calibration of the three Cernox thermistors installed on the cryogenic thermal shields. **Right:** Sensitivity of the thermistors. 103
- 4.26 **Left:** A copper disk is used as a dummy cavity, a thermistor attached to it allows us to measure the expected cavity temperature during cooldown. **Right:** The active shield with thermistors installed. Temperature sensors from the passive shield pass through a small slit beneath the lid. A thin layer of vacuum grease between the lid and cylindrical section improves the thermal contact. Cables for the sensors are wrapped around each shield several times for thermal anchoring. This is especially important on the radiation shield so that heat flowing from the 300 K vacuum feedthroughs is dumped on the 30 K stage where the coldhead has more cooling power. Sensors are attached to the shields with VGE-7041 varnish and baked at 150 C for 6-10 hours for curing. 105
- 4.27 During one of the initial cooldown tests, two thermistors were placed on each of the active and radiation shields to measure thermal gradients. By design, the high purity OFHC copper should have suppressed gradients to the mK level. Surprisingly, a residual gradient of 1 K was measured on the radiation shield (left column) and 0.65 K on the 4 K active shield (right column). We were unable to find a proper explanation for this, it is possible that curing the epoxy that attached each thermistor to the shield introduced resistance offsets via mechanical stresses. This is another argument for securing the thermistors with a mechanical clamp in future systems. The bottom row displays an Allan deviation of the gradient fluctuations. The active shield gradient is remarkably static with fluctuations slightly over the 10 μ K level over 1000 s. 106
- 4.28 The cryostation is heated from 4 K to room temperature while we read out thermistors on the top and bottom of the active shield. The difference in measured temperature (right) varies significantly over the warmup process. 107

4.29	Temperatures on the thermal shields and cryocooler stages are continuously measured during the final cooldown.	108
4.30	Thermal time constants In a test cooldown of the system, the cavity is replaced with a a copper disk of the same thermal mass. Thermal sensors are attached to the testmass, inner 4K shield, and outer 4K shield. The temperature of the outer 4K shield (active) is stepped via a ring heater wound around its base, and we monitor the subsequent temperature increases of the inner 4K shield (passive) and testmass. We extract the time constants between the different stages as well as the thermal gradients between them. The ring heater on the active shield base also allows for active temperature stabilization to either the active or passive thermistor.	110
4.31	Top left: Time series of the inloop and witness thermistors placed on the 4 K active shield (loop is closed on the inloop thermistor). Top right: Instability of the two thermistor readings. Bottom left: The witness' temperature fluctuations are converted to frequency fluctuations via our measured CTE. The brown curve is the witness variations filtered by the passive platform and cavity mount. Bottom right: Instability of the traces on the left.	112
4.32	Vapor pressure of most gases as a function of temperature, taken from [9]. Below 20 K, the only residual gas aside from Helium with a pressure above 10^{-13} Torr is H_2 , which has a vapor pressure of $\sim 10^{-6}$ Torr at 4 K. At the UHV range, most curves are extrapolated with decreasing temperature.	114

- 4.33 Schematic of the modifications for additional pumping on the cryostation. Montana Instruments provides a control unit connected to the cryostation with a long 2 m, 0.375" ID vacuum bellows. Within the control unit is a small roughing pump that is valved off from the bellows with a case valve contained within the control unit. A vent valve can be left open to the environment or connected to a nitrogen canister to vent the system to ambient pressure. We removed this connection by separating it from the pumping bellows with a gate valve. Instead, we add a tee that is connected to a turbo and scroll pump. After operating these two for a few days, we valve them off and activate a small 2 L/s diode pump. 116
- 4.34 **Left:** Frequency noise of Si3-Si6 beatnote when the vacuum pumps are running and when shut off. There is no additional noise measured due to vibrations of the pumps. **Right:** Fractional instability of the beatnote data on the left. In the low frequency regime, we can slowly displace the AVI platform slightly above the ambient vibration noise and look for corresponding changes in the cavity noise. Below 1 Hz, accelerations coupling directly to the cavity should generate frequency perturbations significantly below the thermal noise limit. Any observed increase in cavity noise could therefore be an indication of poor vacuum within the cryostation. We do not observe any such effects. 117
- 4.35 The fiber noise cancellation servo is locked with high bandwidth of ~ 150 kHz. There is little noise present at this high offset frequency, but significant frequency noise is present at 1 kHz and below. The high servo bandwidth is necessary for providing the large servo gain needed to suppress the in-loop error to the cavity thermal noise limit. 119

4.36	Schematic of the setup used to measure the out-of-loop fiber noise. Light from the NKT laser is split into two paths, one going to a fiber noise cancelled interferometer, and the other serving as a frequency reference. The fiber noise cancelled link, which typically distributes light to the Si3-Si6 photodetector, is instead sent to a separate out of loop photodetector to generate a heterodyne beatnote with the reference light. Frequency noise of this beatnote measured on the OOL PD quantifies the frequency noise associated with the light dissemination.	120
4.37	Measurement of the dissemination noise using the schematic shown in Fig.4.36. Left: Fractional frequency noise of the measurement. Right: Fractional Allan deviation of the measurement. Dissemination noise is well below the projected thermal noise limit of Si6.	121
4.38	Compilation of technical noise sources measured on Si6. The frequency noise due to intensity fluctuations is discussed in the following chapter.	122
5.1	Coating and substrate diffusion lengths are plotted as solid red and green lines respectively. Dotted and dash dotted red lines indicate two length scales relevant to the coating diffusion, i.e. the beam size and the coating thickness.	125
5.2	Left: Refractive indices of the high and low index GaAs and AlGaAs layers. Right: Temperature dependence of the refractive indices.	127
5.3	Calculated photothermal transfer functions for the elastic spacer, substrates, and coating responses, and the thermo-refractive coating effect. In the low frequency regime where we aim to be thermal noise limited, photothermal effects from the spacer are the dominant contribution	128
5.4	Transient response of dropping the transmitted intensity from 300 nW to 100 nW for the E_1 mode (solid line) and the E_2 mode (dotted line). Intensity is dropped at ~ 10 seconds, and there is a fast transient, followed by a slow settling of the frequency that exhibits opposite sign for the two polarizations.	129

- 5.5 Transient response of the E_1 when changing the transmitted power between different values. A clear asymmetry in time constants is visible between the left and right panels. Increasing laser power results in a much faster frequency response than decreasing the power, indicating that there are different photothermal transfer functions depending on the sign of the intensity change. 129
- 5.6 Frequency responses renormalized by the magnitude of the transmitted power change. Solid lines indicate curves taken when decreasing the optical power, dashed lines indicate an increase. Curves of same color are taken with the same set of initial and final power values. It is clear that there is no singular transfer function $\delta\nu/\delta P$ that captures the frequency response. 130
- 5.7 **Left:** Single and double exponential fits of the data in Fig.5.5. **Right** Frequency residuals of a single exponential fit (pink) and the double exponential fit (grey). . . . 132
- 5.8 **Left:** Measured photothermal transfer function of the cavity response when modulating the transmitted power while locked to a single polarization mode of Si6. Data is taken at an operating temperature of 4 K. **Right:** Renormalizing the left set of curves by $\sqrt{P_{transmitted}}$ 134
- 5.9 Converting the photothermal transfer function from units of frequency deviation to length deviation allows us to directly compare the coating response of Si5 (stars and solid lines) and Si6 (circles and dotted lines) without complications from the different cavity lengths. Renormalizing by the transmitted power produces a similar set of transfer functions for the two systems, despite over an order of magnitude span in power. 136

- 5.10 Measurements of the cavity frequency noise for different transmitted powers, ranging from a minimum of 90 nW to 1000 nW transmitted. When this data was taken, optical scatter off one of the input vacuum optics contributed a large noise hump at 100 mHz offset. Cavity noise at frequencies below this offset is limited by this technical source. Below a cavity transmission of 90 nW, electronic noise sources start to become significant and no further reduction in frequency noise is observed. A dotted red line provides a $1/f^{1.5}$ guide to eye, well above the projected Si6 thermal noise (black line). 139
- 5.11 **Left:** Another set of frequency noise measurements taken with different intracavity powers. **Right:** Curves on the left are normalized by $\sqrt{P_c}$ and collapse to a single noise limit despite the factor of 10 difference in optical power. A dotted red line is displayed as a $1/f^{1.5}$ guide to the eye 140
- 5.12 Light-induced noise at a cavity temperature of 4 K (solid lines) and 16 K (dotted lines). Curves of same color are measured with same intracavity power. We do not measure any changes in coating loss and transmission between the two temperatures. 141
- 5.13 **Top:** Light-induced noise measured on Si5 at 124 K for different intracavity powers. **Bottom left:** Comparison fractional frequency noise of Si5 and Si6 with the $\sqrt{P_c}$ rescaling implemented. The light induced noise in Si5 is over 10 dB lower than that in Si6 due to the longer cavity length. **Bottom right:** Renormalizing by the cavity lengths allows us to compare the absolute coating noise between the two system. . . 143
- 5.14 Merit factor $g_{m,n}$ for different spatial modes in the rectangular basis is plotted on the top, and $g_{\rho,l}$ for the cylindrical basis is plotted for the bottom. Rectangular modes $HG_{0,i}$ and $HG_{i,0}$ have the same scaling factor due to separability along the two axes. This is not the case for cylindrical modes where increasing the azimuthal mode index leads to a larger effective mode area. 145

- 5.15 Generation of higher order modes with phase plates. The top row shows the phase plate (middle) needed for converting the fundamental mode (left) into the HG_{11} mode (intensity profile on the right). Yellow regions of the mode converter imprint an additional π phase shift when compared to the purple regions. The bottom row shows a similar procedure for generating the LG_{03} mode. 146
- 5.16 When using higher order modes, sensitivity to clipping is increased. Not only is the mode area greater, beam size will be larger at all locations along the propagation axis. Here, we compare the beam size for the LG_{00} , LG_{20} , and LG_{02} modes at various locations before the cavity. Locations of various optical apertures are shown as dashed lines at their corresponding distances from the cavity input mirror. A pink dot indicates the size of the 4K active window at its location 9 cm from the cavity. The significantly larger beam radius of the LG_{02} mode makes it much more susceptible to clipping on this optic. 147
- 5.17 Mode area effect on birefringent noise. We separately measure the birefringent noise of the HG_{00} and HG_{01} modes by addressing a single polarization of the cavity. The ratio of birefringent noise S_{biref}^{01} and S_{biref}^{00} for the two modes reveals the dependence on mode area. If the birefringent noise is independent of the mode area, we expect a ratio of 1. We find a ratio of 0.71, consistent with the expectation of 0.75 if the noise scales with mode area $g_{mn}w_0^2$. A large peak visible at 0.057 mHz is the result of a ground loop that appeared in the S_{biref}^{01} dataset and was omitted in the fit. . . . 148
- 5.18 Data from the bottom right of Fig.5.13 is now divided by the mode area of the Si5 and Si6 cavity. 148

- 5.19 **(a)** Optical layout for probing the Si6 system with active RAM suppression at a modulation frequency (5 MHz) where the cavity error signal is encoded in a traditional PDH locking setup. Dual-tone probing is achieved by driving AOM2 with two RF tones separated by the birefringent mode splitting of Si6 (green and purple sets of tones). A half waveplate before the cavity is aligned to couple an equal mixture of s- and p-polarized light to the two polarization modes of the cavity, labeled $\hat{E}_1$ and $\hat{E}_2$. The collinear beams are then launched upwards with a 45° mirror to couple into the cavity. The reflected cavity error signal has equal contributions from noise of the two polarization modes, and any anti-correlated noise is automatically rejected. **(b)** Frequency landscape of the tones applied to the Si6 cavity. The two tones applied to AOM2 separate the carrier and the two sidebands each into two tones separated by the birefringent mode splitting of the cavity. The purple tones shown address one polarization mode ($\hat{E}_1$) while the green tones simultaneously address the orthogonal mode ($\hat{E}_2$). 151
- 5.20 **Left:** At $t = 0$, the laser is initially dual frequency locked to E1 and E2 modes. Around $t = 10$, the tone addressing E2 is removed so that the laser is only locked to E1. A frequency jump corresponding to Δ is observed. The E2 tone is then reintroduced, and the E1 tone is removed. The laser frequency, distributed after AOM1, is continuously monitored in a beat with Si3. **Right:** A cartoon indicating the location of the AOM2 tones relative to the two polarization modes. Colors correspond to the data on the left. With $\Delta = 0$, there is no frequency jump when switching from the dual mode lock to the single mode lock, indicating that the spacing of the AOM2 tones is equivalent to f_{Biref} 155
- 5.21 Residual noise $S_{\text{residual}} + S_{\text{Brown}}$ after cancellation of the birefringent noise at both 4 K and 16 K shows no temperature dependence. 157

- 5.22 Fractional frequency instability of Si6 of birefringent noise (red squares) and residual noise at 16 K (blue circles). The residual noise is measured over a continuous 24000 s dataset. 158
- 5.23 The optical scheme of [10] splits the beam across two different paths that are prepared in frequency and polarization to address the birefringent modes. Beams are recombined on orthogonal ports of a polarizing beam splitter and coupled into the cavity. A Faraday rotator before the cavity swaps the polarizations of each incoming beams such that reflected light from the cavity fast axis is measured on the same path as incident light addressing the cavity slow axis. A heterodyne beatnote between the birefringent modes can be detected in cavity transmission after interfering the two polarizations. 160
- 5.24 **Left:** schematic of a proposed optical setup. Ideally, a switchable polarizing element could rotate the incident light to the fast or slow axis of the cavity in sync with the data acquisition rate. **Right:** A cartoon illustrating the concept of the asynchronous birefringence measurement. An AOM switches the laser frequency to alternate between addressing the two modes. The demodulated signal is recorded on digital multimeter and the corresponding phase error of each cavity mode can be reconstructed to yield the information we want. 162
- 5.25 The Si6-Si3 heterodyne beatnote (blue) where two lases are locked to the two separate cavities. Probing the Si6 cavity with a laser locked to Si3 recovers the Si6 noise spectrum with the Si3 noise included. This signal is measured in cavity reflection and then converted to frequency noise by calibrating the discrimination slope (pink). Both traces encode the PDH error signal at the same modulation frequency of 5 MHz. 164
- 5.26 Attenuation factor in transmission of power coupled into the cavity. 164

5.27 **Left:** Shot noise limited length resolution in a 1 Hz bandwidth for the Si6 cavity.

Reducing the optical power in the carrier will decrease the magnitude of the birefringent noise, but also hinder the ability to measure that noise at high frequency.

Lowering the measurement bandwidth can mitigate this. **Right:** Johnson noise equivalent for a particular level of shot noise. If the detector noise floor is limited by Johnson noise and other op amp noise sources are below it, this yields the noise

equivalent power. 165

5.28 Fractional frequency noise of the Si3-Si6 beatnote is computed with an evenly sam-

pled time series of beat (black). The effect of aliasing the data, or introducing deadtime between measurements can be simulated by extracting a time series of every alternate point (e.g. all even timestamps), doubling the effective acquisition period to account for deadtime, and computing the power spectrum (blue). The low frequency noise of the beatnote is fully recovered, but an artificial white noise floor is evident at higher frequencies. Increasing the deadtime or sampling even fewer points has the same effect, but accentuates the magnitude of the white noise. In the asynchronous measurement, datapoints will alternate between addressing the fast and slow axes of the cavity. The degree of common mode rejection for noise present in both modes can be estimated by taking the difference of two signals aliased from the same dataset. In orange, we show the difference channel between an aliased signal composed of all even points of the Si3-Si6 timeseries and an aliased signal composed of all odd points. Low frequency noise in the difference is suppressed to the level of the white noise which is incoherent to the two datasets. Over 20 dB of common noise rejection is visible at the lower range of 100 mHz. 166

- 5.29 Sensitivity to the AC components (blue) from the Pound-Drever-Hall modulation scheme is contrasted to sensitivity obtained from the DC components (green) via the side of fringe interrogation. At a frequency indicated by the dashed magenta line, AC and DC components should yield equivalent sensitivity to cavity noise. The cavity transmission function is overlaid in red. 167
- 5.30 **Left:** Cavity noise measurements from the AC (modulation at 5 MHz) and DC signals. The DC signal is slightly smaller in amplitude due a smaller sensitivity to cavity noise at this particular detuning (at the cavity pole). **Right:** Detunings for the AC and DC measurements are now swapped, so that each measurement is insensitive to the amplitude noise converted by the cavity. 168
- 5.31 **Top:** Transimpedance gain of detector for measuring birefringent noise measured by Terry Brown. X-axis should be in kHz. **Bottom:** Current noise of the detector. We use the typical Hamamatsu G12180-003 InGaAs PIN with 1.1 A/W responsivity at 1550 nm. At 770 kHz, noise is equivalent to shot noise of 220 nW. 170
- 5.32 Laser-cavity error signals of demodulating the two tone electric field at the birefringent splitting 2Δ . Different curves indicated a different extraction phase ϕ . All curves are shifted horizontally for clarity though there is no vertical offset added. **Left:** Error signals when $\tilde{E}_1$ is transmitted. **Right:** Error signals when $\tilde{E}_2$ is transmitted. Since components from the cavity modes oscillate out of phase with each other, the large residual amplitude modulation is removed. However, the signal retains its symmetry about zero and is three times smaller. Detuning the laser by 2Δ to measure laser-cavity noise of a single mode no longer has any sensitivity when the polarizer is orientated to transmit $\tilde{E}_2$. In both cases, sensitivity to laser noise can be obtained by detuning the probe to the side of both lines, similar to the side-of-line probing demonstrated in the previous section. Alternatively, locking to both modes simultaneously and reducing the splitting of polarization tones to $2\Delta - \Gamma_c$ will naturally shift the absorptive signal by $\Gamma_c/2$ where there is sensitivity to detuning. 173

- 5.33 **Left column:** Sensitivity to noise on the mode splitting when $\tilde{E}_1$ is transmitted.
Right column: Sensitivity to noise on the mode splitting when $\tilde{E}_2$ is transmitted.
Top row: Laser is tuned such that both modes are addressed simultaneously. **Bottom row:** Laser is detuned by 2Δ so that only one mode is addressed. We regain sensitivity to the birefringent noise with the laser resonant with a single mode, but this sensitivity depends strongly on whether $\tilde{E}_1$ or $\tilde{E}_2$ is transmitted. 175
- 5.34 Si3-Si6 beatnote with krypton gas flowed into optics enclosure at different times.
The blue segment highlights the steady state without any deliberate perturbation to the optical path length. A short burst of gas is released at the beginning of the orange segment. The flow is continuous for the duration of the green segment and then shut off at the end. A clear chirp in the cavity frequency is visible as the gas dissipates. 176
- 5.35 **Top:** Frequency response of the cavity to a step of the vacuum chamber temperature.
Bottom: Residuals of an exponential fit to the frequency response above display a modulation indicative of an etalon in the system. The lack of a significant time constant between the temperature step and etalon response suggests that one of the etalon surfaces is not thermally isolated from the 300 K stage. Modeling the vacuum chamber temperature as a phase ramp, we can largely reproduce the cavity frequency response by fitting the coupling of temperature to phase. In the orange trace, we correspond a 100 mK vacuum chamber fluctuation to a 2π phase shift. Active suppression of the 2 Hz peak-to-peak etalon amplitude below the 10^{-16} level would require $\sim$ mK stabilization of the window temperature. This is difficult, an easier approach is to attempt to reduce the etalon finesse. 178
- 5.36 Sweeping the laser frequency, several transverse modes visible in the scan as well as a modulation on the baseline indicating the presence of an etalon. A fit to this reveals the etalon's free spectral range to be 160 MHz which corresponds to a cavity length of 0.9 m. 179

- 5.37 **Top:** Dedrifted Si3-Si6 beatnote overlaid with the PDHDC signal multiplied by the horizontal sensitivity coupling. **Bottom:** Instability of the beat, and PDHDC measurements when propagated through the horizontal and vertical transfer functions. Some sections of the beatnote instability between 100-1000 s appear to be well described by the horizontal pointing deviations. For longer averaging times, the cavity instability might result from a linear combination of vertical and horizontal pointing fluctuations. 181
- 5.38 Modulating the voltage applied to the piezo allows us to measure a transfer function between the cavity frequency response and changes in the PDH DC signal. We do this for both axes to measure the system's sensitivity to beam pointing. **Left:** The piezo is place in the mirror's horizontal pin. We drive it with a sawtooth and simultaenously record the frequency deviation. Red trace is the frequency response, and the green trace is the PDH DC voltage converted to frequency via our measured coupling. **Right:** Same data, except the piezo is now driving the mirror's vertical pin. 182
- 5.39 **Left:** The light is coupled to the cavity from the bottom setup and the transmitted beam is imaged from the top. **Right:** Light is coupled from the top and imaged on the bottom. 183
- 5.40 Fig.4.12 is reproduced but now we include a measurement of RAM from the top setup. The modified allan deviation of RAM instability is now consistently better than the Si3 thermal noise from 0.5-10 s, and the fractional frequency noise shows a considerable 6 dB improvement around 100 mHz. No effort was put into improving the long term instability of RAM in the top setup, so we expect that this can be improved. The top setup couples to the HG_{00} mode while the bottom setup couples to HG_{01} 185

- 5.41 A five day measurement of the Si3-Si6 beatnote with the Si3 drift removed. During this period the Si3 drift was measured by the SrI lattice clock so its sign and magnitude is known. The dotted line is a fit to the entire dataset and finds a fractional drift rate of $-2.3 \times 10^{-20}/\text{s}$, whereas the dashed line fits to the latter two thirds where a more linear behaviour is observed. This segment finds a fractional drift rate of $-5.6 \times 10^{-20}/\text{s}$ 187
- 6.1 We simulate a continuous dataset of different types of noise, $S_y \sim f^0, f^{-1}, f^{-2}$. The performance of the periodogram and Lomb-Scargle periodogram estimators is compared. The Lomb-Scargle periodogram consistently underestimates the noise spectrum by a factor of 2 in amplitude. A correction therefore needs to be included in order to obtain a comparable result. 193
- 6.2 Exclusion limits for d_e measured between the three optical clocks during the BACON frequency intercomparison data campaign. The dataset is the same as that from [11]. The most stringent bounds on d_e are set by the Yb/Sr frequency ratio despite the relatively higher sensitivity of the Al^+/Yb comparison. 195
- 6.3 **Left:** Comparison of the data record of the H-maser measured against the silicon cavity frequency, and the 1D Sr lattice clock measured against the silicon cavity. The time series of the maser-cavity frequency ratio is decimated by a factor of 720 for clarity. The short term instability of the atom-cavity measurement is far superior to that of the maser-cavity ratio, but the enhanced sensitivity of the maser to the scalar DM oscillation will enable it to set a competitive bound. **Right:** Frequency noise of the f_H/f_{cav} (blue) and the f_{Sr}/f_{cav} (red) ratios calculated with the Lomb-Scargle method. A noise model for each ratio derived from the frequency noise contributions of the cavity, maser, and 1D Sr lattice is shown in black. 197
- 6.4 A cartoon illustrating the dependence of the H maser, silicon cavity, and strontium clock transition on α and m_e 197

6.5	The ratio of the observed DM signal compared to the actual signal amplitude when considering a nonzero phase δ to the DM field.	198
6.6	The constraints placed in this work are the blue curve for the maser-cavity frequency measurement, and the red curve denoting the atom-cavity measurement. Purple lines indicated limits set from equivalence principle tests [12, 13]. Black curves are exclusion bounds determined from atom-atom frequency comparisons [14, 15, 16]. .	199
6.7	Exclusion limit for the electron mass coupling d_{me} obtained from the maser-cavity frequency comparison.	200
6.8	Long term instability of the silicon cavity (orange and red points) with our model for the oscillator noise (dashed black line). In a single day's measurement, particular Si3 datasets have displayed thermal noise limited performance out to 10^4 seconds. This low instability is naturally degraded with the inclusion of significantly more data. A similar effect is apparent in the Si6 long term behaviour, short datasets of 6-8 hours demonstrate instabilities of $\text{mod } \sigma_y < 1 \times 10^{-16}$ for 10^3 s, though a week long dataset shows $\text{mod } \sigma_y = 1 \times 10^{-16}$ at 200 s. Our intent is to point out that length of the dataset is critical when assessing the long term cavity behaviour. . . .	202
6.9	Figure adapted from [17]. Stability of the optical and microwave timescales is evaluated when measuring against an optical Sr standard for two duty cycles. Measurement of the OLO for 1 hour per day is still superior to the microwave timescale when measured against the strontium clock 12 hours/day. This is directly related to the superior short term thermal noise limited performance of Si3 and the excellent long term control. We additionally display the significant advantage enabled by measuring the microwave timescale against an optical reference compared to a microwave reference.	204
6.10	Top: Bottom: Curves from the top plot are converted to a strain sensitivity by normalization with the cavity's mechanical response to a gravitational wave. The aLIGO strain sensitivity is displayed in red.	207

7.1	We plot the total Brownian thermal noise limit for MJM (purple), Si2/Si3 (green), the limit for a 21 cm silicon cavity at 16 K with dielectric coatings using the measured loss angle from [18], and the limit for a 21 cm silicon cavity at 16 K with AlGaAs/GaAs coatings using the measured loss angles from [19, 20]. A dashed black line indicates the instability where an oscillator has experienced a phase excursion of $\phi_{\text{rms}} = 1$ rad.	211
A.1	The laser has a white phase noise spectrum where $S_y(f) \sim f^2$. Top row shows the laser $S_y(f)$ and the corresponding modified Allan variance. The bottom row shows the Dick effect limited 1 s clock stability that can expected for different dead times and pulse times. Different behaviour will be visible for Rabi spectroscopy (lower left) and Ramsey spectroscopy (lower right).	229
A.2	The laser has a white frequency noise spectrum where $S_y(f) \sim f^0$	230
A.3	The laser has a flicker frequency noise spectrum where $S_y(f) \sim f^{-1}$	231
A.4	The laser has a random walk frequency noise spectrum where $S_y(f) \sim f^{-2}$	232
B.1	Amplitude of different sidebands generated via phase modulation on an EOM. . . .	234

Chapter 1

Introduction

For the majority of us, our introduction to physics involves identifying the forces exerted on an object and collecting these terms into an equation that captures the evolution of that object's state. In a closed system where the energy is conserved and phase space is a natural description, this amounts to identifying the Hamiltonian and allows us to calculate useful observables. For many classes of problems in classical mechanics, we can exactly calculate all the desired quantities; a canonical example is an oscillating pendulum where the Hamiltonian consists solely of the mass' kinetic energy and its potential energy in a gravitational field. The efficacy of this description depends on how accurately the model captures the physical dynamics. We can often benchmark our calculations against a priori knowledge of the exact ground state, or the exact equations of motion.

In the quantum realm, this can be more difficult. We might want understand a particular process and generate a Hamiltonian as an ansatz for the system under study. If the system is especially complicated, we might not have exact knowledge of the ground state. Indeed in quantum systems with large interactions (quantum Hall physics, superconductivity, magnetism), the complexity is such that these problems may be intractable to solve numerically. An alternative approach has been to use quantum systems as 'simulators', where a precisely controllable system can be used to engineer a relevant Hamiltonian and compute the associated state variables [21, 22, 23]. The first production of degenerate gases of neutral atoms heralded the ability to control a system's interactions and energy scales to the point where its quantum state can be subjected to carefully

chosen Hamiltonians [24, 25]. We've now developed the capability to reconstruct intractable ground states by preparing quantum resources and subjecting them this Hamiltonian [26, 27]. Quantum simulation with degenerate gases has a well established history of lending access to global observables, and the recent advent of quantum microscopes [28, 29] has now also enabled us to extract local observables.

Despite the power of this method, it can be argued that this is not the most natural way to extract a system's ground state. In introductions to quantum mechanics, we usually learn to solve quintessential problems (e.g. the quantum harmonic oscillator) by diagonalizing the Hamiltonian, providing us with direct access to the energy eigenstates. This may seem to be a preposterous suggestion for studying real physical systems - how does one actually diagonalize the shifted energy spectrum of a many body system when an atom interacts with light?

This concept is in fact a central tenet of metrology, where spectroscopy is used as a tool to measure these energy shifts in the frequency domain [30]. Clever spectroscopic methods have been developed to enable us to resolve finer and finer interactions, i.e. probe smaller and smaller energy scales. Of course, as with all measurements, a quantity can only be resolved if the measurement device does not limit the measurement resolution. In principle, we would like an infinitely smaller ruler - a delta function of sorts - to be able to study arbitrary systems, investigate interesting Hamiltonians, and probe new phenomena at the smallest length scales possible.

This is perhaps a peculiar, inverted approach to atomic physics. but frequency metrology with atoms has enabled us to make extremely high resolution measurements of frequency shifts and Hamiltonian ground states [31, 32, 33]. Particular atoms support so called 'clock transitions', where an excited state lifetime is nearly metastable and optical transitions to the excited state can therefore support some of the highest natural quality factors available to us [34, 35, 36]. The same transition makes these atoms suitable as clocks where an atom can be placed in a superposition of ground and excited states of the clock transition, precess at an optical frequency, and precisely measure an optical phase shift corresponding to some external potential. This is an interesting feature of these systems - the longer we allow an atom to precess in that superposition and accumulate

phase, the finer our ability to resolve a smaller phase shift. In effect, our measurement device (in this case a shift of the clock frequency) becomes more precise the longer we measure for.

It appears that we should simply measure forever, to achieve our goal of letting our measurement precision approach that of a delta function in the frequency domain. In reality, this is technically demanding and requires exquisite control of our atomic system. An alternative approach is to use many identical copies of the system so that the measurement precision is naturally improved by incoherent averaging. This reveals the power of a many-body approach where a coherent preparation of these identical copies allows us to leverage improved statistics associated with coherent averaging of a measurement. Optical clocks are well suited to these problems because the coherence between atoms exists for a long period of time, tens of seconds in particular demonstrations [37, 38]. Our ability to coherently interrogate these atoms for the same duration however is a formidable problem, and remains one of the primary hurdles in optical frequency metrology. Optical transitions must be driven with an optical local oscillator (OLO), in this case a laser. Any OLO of the correct color can interrogate the atoms and using them as a frequency discriminant forms the basis of the optical clock. Stability, the capability of the clock to quickly measure small frequency shifts, is directly tied to the coherence time of the oscillator. The impressive coherence times afforded in atoms with optical clock transitions can only be utilized if the OLO introduces a minimal amount of phase noise during clock spectroscopy. This thesis will detail our efforts to push the frontiers of atom-light coherence by developing a next generation oscillator stabilized to a cryogenic, crystalline Fabry-Perot resonator.

1.1 Thesis outline

The first half of my graduate career focused on the atom side of the clock. With the development of state-of-the-art silicon cavities [39, 40, 41, 42, 43], it suddenly became feasible to average to the systematic uncertainty our Strontium lattice clocks in a matter of minutes [44]. Instead of the cavity being the bottleneck for achieving groundbreaking atomic precision, the quantum resource was now the primary limitation. There was thus a strong motivation to dedicate effort to improving

the atom-atom coherence. Identifying limitations to this parameter typically require an evaluation of all systematics that couple to the atomic frequency. For metrology, it is useful to evaluate the entire systematic error budget, quoting an accuracy and uncertainty for each mechanism.

This resulted in a new SrI accuracy evaluation with a fractional uncertainty of 2×10^{-18} [45]. Blackbody radiation and AC stark systematics were addressed in a novel way that had not been approached in previous Strontium lattice clocks. Active temperature stabilization of the entire vacuum chamber resulted in a residual gradient of less than 100 mK, that was stable to better than 1 mK over a 10,000 second measurement. The effects of this extended far beyond precise control of the thermal environment - loading into the optical lattice and consequent mode occupations and sample temperature were well controlled and highly reproducible. The effective optical potential sampled by the atom therefore allows for an AC stark shift that remains consistent over the course of weeks. Density shifts coefficients for example no longer varied over the course of a day, and exhibit more consistency on a day to day basis. On a day timescale, we calculate point-by-point corrections to the atomic frequency to the major systematics that may fluctuate between each interrogation, or longer. Over a 6 hour dataset, the total corrections that need to be applied are displayed in Fig.1.1. The clock is precise to the 10^{-19} level for the entire fractional Allan deviation record of 10^4 seconds. The remarkable point to note is the atomic ensemble and environment are passively controlled to this level over the course of the measurement, meaning that a comparison with a more stable frequency reference would presumably allow SrI to average to the 10^{-19} *without* providing any corrections to the strontium atomic frequency.

These are all in-loop measurements of a sense. We measure a coupling to a systematic effect, and use a detection channel for the corresponding environmental or ensemble property to calculate the corresponding frequency shift. This is an important component of assembling the clock's uncertainty budget, but a true evaluation of the level of frequency control can only occur by measuring against another frequency reference. In this case, we measured against a 3D optical lattice clock, SrII.

The SrI-SrII frequency comparison was performed asynchronously, meaning that the mea-




Figure 1.1: **Left:** Plot of the time record of all the SrI systematic shifts. Changes in atom number, ambient temperature, or magnetic field all result in corrections to the clock frequency, and their total magnitude is shown over a six hour data campaign. The clock achieves over 98% uptime over the course of this single comparison day and slight gaps in the data indicate brief periods where the laser is not locked to the atoms. **Right:** The same data is plotted as a fractional instability normalized to the Sr clock frequency. The individual contributions of density shift (blue), BBR (red) and second order Zeeman shift (yellow) are shown as the dashed curves. For operation times up to 10^4 s, fluctuations in systematic offsets are bounded below 4×10^{-19} .




Figure 1.2: Density shift evaluations. Density shift measurements were performed over four weeks at the same trap conditions. The weighted mean of all measurements is shown with a dashed, black line, with a reduced chi-squared of 1.07

surement stability was sensitive to the noise of the local oscillator. The level of precision achieved is characteristic of the excellent systematic control of both clocks but the time taken to average to this value is a feature of the local oscillator's low instability. Similarly, an accuracy evaluation can




Figure 1.3: Non-synchronous comparison (i.e. interrogation times of the local oscillator do not overlap such that there is sensitivity to laser noise) with the JILA 3D optical lattice clock demonstrates that the beat between the two clocks averages below the quoted total systematic uncertainty. The black line is a white noise $\tau^{-1/2}$ fit to the single clock instability.

only be performed by measuring against another accurate reference. After converting the old SrI machine into an accurate and precise lattice clock, we participated in a three way frequency ratio intercomparison with the Al^+ ion clock and Yb lattice clock at NIST.

These results are more comprehensively detailed in [46] so we won't delve further into them here. Instead, we'll turn our eye to the technology that has enabled all of these measurements of record breaking precision, the local oscillator. With the impressive progress in improving atom-atom coherence [38, 37], the development of better flywheels have naturally received less attention. By themselves lasers stabilized to cryogenic monocrystalline silicon cavities have realized linewidths of a few mHz at $1.5\mu\text{m}$ [39, 40, 42, 43]. However, particular research directions are only possible with the development of better continuous optical phase references [17, 47]. These studies have enshrined high stability optical oscillators as key components in pushing frontiers of optical metrology.

These cavities also occupy a unique space in frequency metrology. Nearly all areas of research leverage the quantum properties of atoms for sensing and measurement applications, but here we have a classical resource that so far has no competitive quantum counterpart. The development

of quantum based optical oscillators has been an intensive area of study [48, 49, 50, 51, 52] but optical interferometry has still remained a superior method for low noise length and frequency measurement. It is in fact remarkable that the leading results in low noise length resolution are achieved by stabilization to a macroscopic, classical crystal. A significant advantage of working with the classical resource is that error correction between the the oscillator (free running laser) and the spectroscopic feature (the cavity) is performed at high bandwidth. More importantly, the spectroscopic feature is always present, and can be interrogated continuously without any deadtime associated with preparing the resource. In concept, the cavity can provide a limitless signal to noise ratio if available optical power is not a concern.

The field has progressed to the point where the achievable frequency stability is limited by fundamental fluctuations of the cavity materials due to Brownian thermal noise. Crystalline cavities have refined this to a limitation that originates nearly entirely from the coating Brownian thermal noise. In our bandwidth of interest, frequency noise of our 10 cm-scale resonator is fundamentally limited by length fluctuations of the thin 10 μm Bragg stack. A similar issue plagues the LIGO gravitation wave detector, where detection at 100 Hz is largely limited by coating thermal noise [53]. In this case the contrast is even more stark where the 10 μm coating places a constraint on the strain noise budget of an interferometer with 4 km long arms. Next generation detectors have proposed to circumvent this limitation by operating at cryogenic temperatures [54, 55, 56], but this introduces more complications than it solves.

On the opposite length scale, experiments with microresonators [57, 58, 59] interrogate displacement noise at the higher bandwidth of $10^3 - 10^5$ Hz, yet are still hampered by thermal noise of the mechanical oscillator. Coating thermal noise is therefore a ubiquitous mechanism of concern within many fields employing optical interferometry techniques. Inventive solutions to reduce its contributions have been investigated, including the interrogation of higher order spatial modes [60], or the design of resonators that support flat topped mesa modes [61].

These proposals are still difficult to implement and so far the primary research focus has been the development of optical coatings with low mechanical loss. Crystalline materials in the bulk have

demonstrated quality factors significantly higher than commonly used amorphous substitutes. The switch from ultra low expansion glass and fused silica spacers and substrates to ones made of silicon heralded significant advances in low noise interferometry [62]. A similar investigation with optical coatings has lead to the development of $\text{Al}_{1-x}\text{Ga}_x\text{As}/\text{GaAs}$ crystalline coatings [19]. Measured room temperature loss angles in these thin films are ten times lower than those measured with amorphous $\text{SiO}_2/\text{Ta}_2\text{O}_5$ dielectric stacks [19, 43, 63]. Fully crystalline cavities (spacer, substrates, mirrors) have the potential to realize even lower laser linewidths and enable novel demonstrations and applications of enhanced atom-light coherence.

The scope of this thesis investigates novel light-induced noise mechanisms associated with one of the first realizations of a fully crystalline Fabry-Perot resonator. We operate a 6 cm silicon cavity with crystalline $\text{Al}_{0.92}\text{Ga}_{0.08}\text{As}/\text{GaAs}$ coatings (henceforth referred to as Si6) in a closed cycle cryostat that enables operation down to base temperatures of 4 K. Reducing our technical noise contributions and using the measured room temperature mechanical loss angle of these coatings, we expected a fractional Brownian thermal noise limited performance of 1.3×10^{-17} for Si6. We instead observed unexpected dynamic and static light-induced frequency shifts that limited the cavity stability to ten times this value. We determined these effects to be unique features of the crystalline coatings and in comparison with a similar AlGaAs coating silicon cavity operating at 124 K, we identified physical parameters that set the magnitude of this noise, and corresponding ways to mitigate it. A unique interrogation scheme is implemented to suppress these effects, and we study the residual noise present after cancellation.

The structure of this work is as follows

- In Chapter 2, we'll present the concepts necessary to understand thermal noise and to consider when designing a Fabry-Perot cavity. We will lay out the formalism used when transferring the oscillator stability to an atomic reference.
- In Chapter 3, we revisit the decision to use silicon as a cavity material and the advantages of crystalline coatings over dielectric coatings. We will provide a brief overview of the

cryogenic environment and the methods we use to measure the Si6 frequency noise.

- Chapter 4 will cover the various noise sources that can reduce the stability of the laser locked to the resonator. We will detail our approach to mitigating each of these contributions.
- Chapter 5 concerns the actual frequency noise measurements of Si6. We will present our observations of static and dynamic light-induced effects identified in the crystalline semiconductor AlGaAs/GaAs coatings
- In Chapter 6, we will present novel uses of ultrastable resonators and detail the experiments that were only possible with the excellent long term stability afforded by cryogenic crystalline resonators.

Chapter 2

Oscillators and stabilization

The laser is an integral part of the optical lattice clock and the silicon cavity in particular has enabled nearly all of the results previously discussed. Here we will present an introduction into cavity design, the fundamental noise limit of a resonator, and how the laser noise spectrum determines clock stability and rotation fidelity.

2.1 Fabry-Perots

The cavity is a linear system, much simpler conceptually than many of the experiments undertaken in JILA. The formalisms of incident, reflected, transmitted fields, longitudinal and transverse modes, mode couplings, etc. are all well understood, and the true challenges lie with the task of precision interferometry. There are many detailed resources on cavities [64, 65] that derive the wave equation solutions for the electric fields supported by the resonator. Here we will only present some of the main results and key features that need to be considered when designing and characterizing the cavity.

The response of the Fabry Perot cavity is entirely determined by mirror parameters and the physical distance between mirrors. Any resonator of length L in vacuum will support a mode of frequency $\nu = c/2L$. However, a continuum of longitudinal modes is in fact supported with frequencies $\nu_i = nc/2L$. These will be evenly spaced by the free spectral range (FSR) of the resonator

$$\text{FSR} = \frac{c}{2L} \tag{2.1}$$

The width of these resonances however is related to both the mirror parameters and the cavity length. The linewidth of the cavity is given as

$$\Delta\nu = \frac{\text{FSR}}{\mathcal{F}} \quad (2.2)$$

The cavity finesse $\mathcal{F}$ is a measure of the mirror losses. For a single mirror, the electric field reflectivity, transmissivity, and loss are r_i, t_i, l_i respectively. These parameters share the relation

$$r^2 + t^2 + l^2 = 1 \quad (2.3)$$

The finesse depends entirely on the reflectivity and is

$$\mathcal{F} = \frac{\pi\sqrt{r_1 r_2}}{1 - r_1 r_2} \quad (2.4)$$

The linewidth of the cavity determines the storage time of a photon as $\tau_s = (2\pi\Delta\nu)^{-1}$. If the cavity length is known, the finesse can be measured via cavity ringdown where the exponential decay of optical power leaving the cavity yields the characteristic storage time τ_s .

The mirror transmissivities are needed for calculating the circulating and transmitted optical power through the cavity. The corresponding electric fields are found to be [64]

$$E_{trans} = \frac{t_1 t_2 e^{-i\phi/2}}{1 - r_1 r_2 e^{-i\phi}} E_{inc} \quad (2.5)$$

$$E_{circ} = \frac{t_1 e^{-i\phi/2}}{1 - r_1 r_2 e^{-i\phi}} E_{inc} \quad (2.6)$$

$$E_{refl} = \frac{r_1 - r_2(r_1^2 + t_1^2)e^{-i\phi}}{1 - r_1 r_2 e^{-i\phi}} E_{inc} \quad (2.7)$$

Where the round trip phase shift $\phi = kL$. A special case of interest is the symmetric cavity where $r_1 = r_2$ and $t_1 = t_2$. When the phase shift is a multiple of 2π , the reflected power drops to zero and nearly all the incident power is transmitted (limited by the mirror losses l_1, l_2). In this case, the transmitted field is

$$E_{trans} = \frac{t_1 t_2}{t_1 t_2 + l_1 l_2} E_{inc} \quad (2.8)$$

When the transmissivity is on the same order as the losses, less than half the incident field is outcoupled from the cavity. As we will see, this is the case for the crystalline coatings in Si6 where

$t_i \sim l_i$. The ratio of circulating power to incident power $|E_{circ}|^2/|E_{inc}|^2$ is also referred to as the cavity gain, and describes the build-up of power coupled into the resonator.

This representation, while useful for describing the ratios of fields within and outside the resonator, does not provide a realistic description of the longitudinally varying electric field. The ansatz for these equations assumed outcoupled and intracavity fields of a traveling plane wave. From experience we know that there is a finite radial extent to these waves so a better approximation is to retain the fast $e^{i\omega t}$ oscillations longitudinally while treating the transverse extent as a slowly varying field.

It is well known that that beams with Gaussian intensity profile are solutions of the paraxial Helmholtz equation for a wave propagating within a resonator. The general solution for the intensity profile with azimuthal symmetry in the transverse direction r is [65]

$$I(r, z) = \frac{2P}{\pi w^2(z)} \exp \left[-\frac{2r^2}{w^2(z)} \right] \quad (2.9)$$

We can see that the beam size $w(z)$ is dependent on the longitudinal position, given by

$$w(z) = w_0 \sqrt{1 + \left(\frac{z}{z_r} \right)^2} \quad (2.10)$$

For Rayleigh range

$$z_r = \frac{\pi w_0^2}{\lambda} \quad (2.11)$$

The waist w_0 is the primary feature characterizing the beam size and is determined entirely by the cavity length and mirror radius of curvature.

$$w_0 = \sqrt{\frac{\lambda L}{\pi}} \left(\frac{g_1 g_2 (1 - g_1 g_2)}{(g_1 + g_2 - 2g_1 g_2)^2} \right)^{1/4} \quad (2.12)$$

Factor's g_i are calculated for a single mirror as

$$g_i = 1 - \frac{d}{R_i} \quad (2.13)$$

Stability of a resonator is determined by the criterion $0 \leq g_1 g_2 \leq 1$. For a cavity outside this parameter range, most light coupled into the resonator will eventually be reflected away from the cavity axis.

The wavefront of the Gaussian beam experiences a longitudinally varying phase shift when compared to the wavefront of a plane wave. This is the Gouy phase shift and is calculated as

$$\phi(z) = -\arctan\left[\frac{z}{z_r}\right] \quad (2.14)$$

The Gaussian beam is not the only solution to the paraxial wave equation. This relies on the assumption of cylindrical symmetry of the resonator mode and when this is broken, a basis of solutions is provided by the orthonormal set of Hermite-Gauss (HG) modes. The spatial profile of any beam can be decomposed as a linear combination of these modes. The spread in transverse momentum is larger for these higher order and correspondingly they acquire a larger Gouy phase shift of

$$\phi_{mn}(z) = -(m+n+1)\arctan\left[\frac{z}{z_r}\right] \quad (2.15)$$

Where m, n characterize the order of the mode along the two transverse directions. We can see that higher order modes acquire a larger phase shift. This manifests as a frequency shift for these higher order modes, with mode spacing of

$$\Delta\nu_{mn} = \text{FSR}(m+n+1)\left[\frac{\phi(z_1) - \phi(z_2)}{2\pi}\right] \quad (2.16)$$

Where z_1, z_2 are the locations of the two cavity mirrors.

The Si6 cavity has a spacer length of 6 cm and mirrors with 1 m radius of curvature, corresponding to a FSR of 2.5 GHz. The transverse mode spacing is roughly 250 MHz, well outside the cavity linewidth and offset frequency of the modulation sidebands. This is an important consideration since low-order modes near the fundamental HG_{00} may have finite coupling. If they are close to the HG_{00} mode in the frequency domain, the resonant feature being probed might experience line pulling effects as coupling to these other modes fluctuates with slow drifts in beam pointing.

High coupling to a particular mode is obtained by shaping the beam size, pointing, and location to overlap with the corresponding mode supported by the cavity. For the HG_{00} mode, it is desirable for the incident beam to have a size corresponding to the size of the cavity mode and

the input mirror, $w(z_1)$. This maximizes the overlap in the longitudinal direction. The general expression for the coupling to a cavity mode with electric field E_{mn} is

$$\eta = \frac{\int_{-\infty}^{\infty} |E_i E_{mn}|^2 dx dy}{\int_{-\infty}^{\infty} |E_i|^2 dx dy \int_{-\infty}^{\infty} |E_{mn}|^2 dx dy} \quad (2.17)$$

With incident field E_i . In general, it is difficult to calculate this quantity ab initio. Instead, it can be easily measured by locking to each higher order mode and extracting the coupling. A demonstration of this is illustrated in Fig.2.1.




Figure 2.1: **Left:** Calculation of transverse modes for Si6. The x-axis refers to a total mode number ($m + n$). The closest degenerate mode to HG_{00} is one where $m + n = 10$, and this is still 6.4 MHz away. **Right:** We measure the cavity couplings to some of the low order modes with $m + n \leq 3$. The highest measured coupling is to HG_{00} with $\eta = 0.5$. The sum of couplings for displayed modes is still less than 1, indicating that there is a finite coupling to other higher modes. We were unable to measure these couplings, indicating that the residual $\eta \sim 0.4$ is likely distributed over very many modes.

2.2 Thermal noise in cavities

We're well aware that systems in equilibrium can be well described by their statistical averages, and in the strontium experiments we make use of the fact that there are deviations around these averages (e.g. a distribution in the states of an optically trapped atom). Fluctuations of equilibrium systems are important considerations in the quantum world, but perhaps less noted in the macroscopic classical world. Nevertheless, we are deeply interested in this.




Figure 2.2: **Left:** We calculate the coupling efficiency for an incident Gaussian beam to the HG_{00} mode of a cavity. Both beams have the same waist w_0 but the incident beam is displaced from the cavity mode in the transverse direction. When the radial displacement δr is equivalent to the cavity waist, the coupling has already dropped to 0.3. **Right:** We now assume that the incident field is perfectly overlapped with the cavity mode in the transverse direction, but the waist of the incident mode w_{in} differs from waist of the HG_{00} mode. We see that the coupling is much less sensitive to variations in w_{in} .

A central tenet of statistical mechanics discovered in 1905 by Einstein was the existence of finite RMS motion of particles in a thermal bath. This was refined in the 1950s as the notion that any coupling between a system and a bath gives rise to a back-action on an observable of the system [66, 67]. Stated another way, there is a fundamental relationship between the behaviour of molecules in thermal equilibrium with the environment and fluctuations of a macroscopic observable. A familiar example is the Johnson noise we measure of a resistor in thermal equilibrium with its environment - dissipation through the resistive material gives rise to a proportional level of current fluctuations measured across the resistor. Capacitors however don't contribute to the current noise, so a general formulation for Johnson noise equates the amplitude of the fluctuations to the real part of the circuit impedance Z .

The general concept associating these two quantities is the fluctuation dissipation theorem which relates the fluctuations of an energy spectrum to the imaginary component of the system's response function.

$$S_x = \frac{4k_B T}{(2\pi f)^2} \text{Re}[1/Z] \quad (2.18)$$

The left hand side represents the power spectrum of length noise fluctuations which are familiar to us, the right side contains the mechanical impedance Z of the system under study. This will end up setting a fundamental limit for the length resolution that we can resolve in any experiment using optical interferometry. A demonstration of this relation is provided next.

2.2.1 Normal mode method

Different types of damping lead to different calculations of the impedance, so intuition for the physical mechanism driving the friction between particles was important. The historical distinction was between external velocity damping (viscous drag exerted on a particle in fluid-like environment) vs internal damping (found to obey Hooke's law). The latter has been accepted as the appropriate formalism describing thermal noise in optical resonators. One of the first applications to optical interferometry was presented in [68]

In these calculations, a particular form of Hooke's law is used

$$F_{th} = -k[1 + i\phi(\omega)]x \quad (2.19)$$

The imaginary component $i\phi(\omega)$ indicates a phase lag between the driving force and material response. The existence of a nonzero ϕ implies that the material is not perfectly elastic given the delay between an applied stress and the resulting strain. Structural damping refers to a frequency independent ϕ which is what we'll consider for now. The loss angle ϕ is used interchangeably with the mechanical quality factor $Q = 1/\phi$. Q denotes the ratio of energy stored in the oscillator to energy lost in a single cycle and the origin of this quantity is nebulous. Not only is it material dependent, but we will see later that it can vary significantly on the the method of synthesizing a material.

Analogous to the γ of the damped driven pendulum, we can see that ϕ quantifies the energy loss of the system when driven by an external force. As we derive the transfer function of the

response, we will find that this is also equivalent to the damping ratio γ that quantifies the width of the system resonance.

Solving the system's equation of motion

$$m\ddot{x} + k(1 + i\phi) = F \quad (2.20)$$

With a generic normalized solution to the systems motion $x(\omega) = e^{i\omega t}$, the transfer function $\chi(\omega)$ of the oscillator is solved as

$$\chi(\omega) = \frac{\omega_0^2(1 + i\phi)}{(\omega_0^2 - \omega^2) + i\phi\omega_0^2} \approx 1 \quad (2.21)$$

In the limit where ω is much less than the resonance frequency $\sqrt{k/m}$. Note that the response is flat in this region. The impedance is found to be

$$Z = \frac{F}{\dot{x}} = \frac{(k - m\omega^2) + ik\phi}{i\omega} \quad (2.22)$$

$$\frac{1}{Z} = \frac{k\phi\omega + i(k\omega - m\omega^3)}{(k - m\omega^2)^2 + k^2\phi^2} \quad (2.23)$$

Substituting back into Eq.2.18, the fluctuations in displacement are determined as

$$S_x(\omega) = \frac{4k_B T \omega_0^2 \phi}{m\omega[(\omega_0^2 - \omega^2)^2 + \omega_0^4 \phi^2]} \quad (2.24)$$

in the low frequency limit where $\omega < \omega_0$ we find

$$S_x(\omega) = \frac{4k_B T \phi}{\omega k} \quad (2.25)$$

since the loss angle $\phi \ll 1$. The intuition here is the same, a stiffer spring constant reduces the overall amplitude of the thermal noise. There is an important point not to be missed - the low frequency limit of the oscillator's transfer function is flat and does not depend at all on the mechanical damping. The quantity ϕ only affects the resonance width. However, ϕ is enormously important in determining the amplitude of displacement *fluctuations*. Unlike the oscillator response, the inherent noise in the oscillator position is dependent on the quality factor of its spring.

The spring constant is a convenient approximation - for a solid material in the elastic regime, we can substitute $k \rightarrow EA/L$, for a Young's modulus E , cross sectional area A and length L so




Figure 2.3: We consider an oscillator with resonance at 1 kHz and varying mechanical loss angle ϕ . **Top:** Transfer function $\text{Re}[\chi(\omega)]$. **Bottom:** Displacement noise spectrum $S_x(\omega)$. In both cases a lower value of ϕ results in a sharper resonance but only in the displacement noise can we see an effect of the quality factor on the low frequency region. Mechanical resonances of the cavity components are always designed to be above 100 s of Hz so $\omega \ll \omega_0$ is the regime of interest.

that

$$S_x(f) = \frac{2k_B T}{\pi f} \frac{\phi L}{EA} \quad (2.26)$$

Remarkably this simple picture provides an excellent approximation of the thermal noise calculated for the cavity spacer [7]:

$$S_{x,sp}(f) = \frac{2k_B T}{f} \frac{L}{\pi^2(R^2 - r^2)} \frac{\phi}{E} \quad (2.27)$$

for a spacer with central bore of radius r and outer radius of R .

The complication arises when we need to consider a system comprised of several coupled springs. The system's response to an applied force is determined by the superposition of the

responses of each of its normal modes. Including contributions from the higher order modes is often necessary for convergence of the displacement noise.

This is not critical in our particular example of the cavity spacer, but the normal mode approach is found to underestimate the total thermal noise contribution particularly when the forces are not homogeneously distributed over the material [69]. In our intuitive picture of the spacer behaving as a solid spring, we can convince ourselves that the internal frictional forces are uniformly distributed across its volume to generate displacement fluctuations. In the case of the cavity mirrors, sections of the coating far from the beam location should not contribute significant noise.

It is in these nuanced scenarios where the normal mode summation approach is exceedingly difficult, and a solution to this problem was proposed by Yuri Levin in the oft-called "direct approach" [69].

2.2.2 Direct approach

The significance of Levin's approach to the FDT calculation of thermal noise is well-stated by [70]:

Until recently, the only method for evaluating Z was based on an expansion of the vibrating state of the mirror on a basis of normal modes. However, computation of the eigenfrequencies, determination of associated eigenmodes and of their coupling with the main beam, is a very heavy task in terms of computational power.

Gillespie and Raab used a multi-order mode summation to calculate coating thermal noise [2]. In the low frequency limit, the total displacement noise of a surface is calculated via the individual contributions of the n modes

$$S_x = \sum_i^n S_{x,i} = \sum_i^n \frac{4k_B T}{\omega} \frac{\phi_i}{\alpha_i m \omega_{0,i}^2} \quad (2.28)$$

Where the $\omega_{0,i}$ are the resonances of each individual mode. The coupling to each mode i is parameterized by an effective mass a_i (computed independently for each mode i) that weights the relative amplitude contributions for the total displacement noise.




Figure 2.4: Reproduced from [2], the x 's denote the noise amplitude contribution from the first 100 normal modes of a 10 cm mirror with a 2.2 mm beam spot. The solid line illustrates the cumulative noise of all these modes and the summation has still not converged over the 100 modes. Notably the mirror diameter is crucial for this calculation whereas it is unnecessary in the direct approach.

The herculean approach necessary to provide an accurate approximation of thermal noise cannot be overstated. Each mode may have a different spring constant due to varying resonances ω_0 , correspondingly affecting the amplitude of each noise spectrum $S_{x,i}$. Fig.2.4 illustrates the summation over the first 100 mechanical modes of the coating and convergence is still not evident. Each mode is statistically independent, and requires an additional calculation to determine the mode amplitude. Additional complications such as displacing the beam away from the mirror require recomputation of all modes since the relative contributions of the various geometrical eigenmodes will change. Ultimately one finds that S_x through the summation method is insensitive to these effects but that comes at the cost of a highly intensive calculation.

The direct approach is far simpler and uses Levin's realization that mechanical impedance Z can be derived by calculating the mirror's deformation due to an incident pressure with profile of the intracavity optical mode. This pressure is generally considered as $P(r) = F_0 I(r)$ with amplitude

F_0 and profile

$$I(r) = \frac{1}{\pi w^2} \exp\left(-\frac{r^2}{w_0^2}\right) \quad (2.29)$$

This pressure generates a strain U within the material and the displacement fluctuations can be found as

$$S_x = \frac{4k_B T}{\pi f} U \phi \quad (2.30)$$

Notable distinctions exist between this form and Eq.2.18. The denominator in Eq.2.30 is now proportional to f^{-1} rather than f^{-2} and an explicit dependence on the loss angle ϕ is now apparent. A direct comparison can be made to Eq.2.25, where the spectrum only scales with f^{-1} due to a linear frequency dependence of $\text{Re}[1/Z]$. Using our prior knowledge of the FDT applied to the cavity spacer, we know that we should expect $S_x \propto \phi/f$ so we can infer that the elastic energy stored in the material (U) will not depend on either of these quantities. Unlike the summation method, we don't need to know the mirror size or beam location relative to the coating center.

Levin's approach provides a much more intuitive picture where we *only* need to calculate the amplitude response to the applied pressure, not the damping ratio or frequency dependence of the strain. The calculation is complicated and can be found in [70] but the general idea is to integrate the strain density u as

$$U = \int_V u dV = -\frac{1}{2} \int_S P(r) u_z(r) dS \quad (2.31)$$

Depending on the approximations made, one can find two different solutions and this will be important as we make the distinction between thermal noise due to the optical coating and thermal noise due to the cavity substrate. The optical pressure is incident on both surfaces with the same profile - all that differs is the consideration of whether or not the beam size is comparable to the depth of the surface. A nuance to this calculation is that the correlation length of the coating response must be considered; a force exerted towards the edge of the coating will only generate phase fluctuations on the optical mode of the coating response is coherent over that distance.

Jumping ahead, what we will find [70, 71] is that the notable difference between the substrate and coating thermal noises is a relative factor d/w_0 where d is the coating thickness.

$$S_{x,sub}(f) = \frac{2k_B T}{f} \frac{1 - \sigma^2}{\pi^{3/2}} \frac{1}{w_0 E} \phi \quad (2.32)$$

$$S_{x,coat}(f) = \frac{4k_B T}{f} \frac{1 - \sigma_{sub}^2}{\pi^2} \frac{1 - 2\sigma_{coat}}{1 - \sigma_{coat}} \frac{d}{w_0^2 E} \phi \quad (2.33)$$

For a rigid cryogenic cavity, the total thermal noise is therefore the sum of all these terms

$$S_x^{tot}(f) = S_{x,sp}(f) + 2S_{x,sub}(f) + 2S_{x,coat}(f) \quad (2.34)$$

Here the Poisson ratio is denoted by σ . When w_0 is *small* compared to the thickness of the substrate, the surface is assumed to be an infinite half space and this assumption is used to calculate the substrate's displacement fluctuations. This approximation is correct up to a constant but a correction accounting for its finite thickness can be applied as $[1 + \mathcal{O}(\frac{w}{D})]$ for substrate thickness D . A potential point of confusion is the scaling of coating thermal noise with $1/w_0^2$. This is not based on the intuition that the applied pressure encompasses a mode area w_0^2 , but rather that coating noise is suppressed by d/w_0 relative to the substrate noise via approximations in the calculation [71].

In real coatings, the mirror is composed of alternating stacks of high and low refractive index materials. Often the loss angles of these two materials are different so the coating thermal noise must be decomposed into separate contributions for each set of layers. Eq.2.33 would imply that the coating thermal noise would be d/w_0 times smaller than $S_{x,sub}$ but we will see in a later section on the AlGaAs coatings that this is rarely the case. Contrary to what has been insinuated so far, loss angle is not innate to a particular material and often varies by orders of magnitude depending on its size and the way it is made. The loss angle for coatings is often orders of magnitude larger than the loss angle of a bulk material, so $S_{x,coat}$ is generally a significant source of length fluctuations.

2.2.3 Temperature dependent damping

One point glossed over in Eqs.2.27,2.32,2.33 is the existence of a temperature dependent loss angle. Reduction of thermal noise is not always as straightforward as reducing the operating temperature. Particular materials have pronounced peaks in ϕ at low temperature, meaning that reducing T can lead to a surprising increase in the thermal noise limit. For this reason, temperature dependent loss angles of a material are often measured in ringdown experiments of mechanical cantilevers. The infrastructure and effort is significantly lower than operating a cavity and measuring the thermal noise over a continuous spectrum from 4 K to 300 K.




Fig. 1. Internal friction measurements as a function of temperature.

Figure 2.5: Evidence of a temperature dependent loss peak in crystalline silicon from [3]. A feature at 160 K is discovered to be associated with impurities in the silicon. An acid etch of the ingot appears to remove the loss peak.

The sources of these temperature dependent losses are not fully understood. Several models have been developed to explain particular features evident in some materials and the consensus is that there is not a singular process that gives rise to these effects [72].

Thermally activated transitions have an associated activation energy, and as the ambient temperature approach that energy, a loss peak can often be observed in the material's loss angle that

increases with lower ambient temperature [73]. Phonon-phonon scattering is another process that has been identified to be a source of additional dissipation in low temperature crystalline sapphire [74]. Chemical impurities have also been discovered to be a source of temperature dependent peaks in ϕ . The particular impurities can be very difficult to identify but different fabrication methods of the material have been shown to remove some of these features [3].

The point of these details is to demonstrate the complexities of understanding material loss angles. A significant amount of research has been devoted to this topic but there is no singular model to capture temperature dependence of ϕ . Knowledge of a material's ϕ at a particular temperature can only be identified through measurement, and the value of ϕ itself depends on the material choice. It is therefore not a generality that the room temperature value can be extended down to the cryogenic regime.

Later, we will discuss the application of crystalline coatings to a cryogenic silicon cavity at 4 K (Si6). Room temperature measurements of the coating loss angle have been performed, but few available results exist for the cryogenic regime. The true thermal noise limit for Si6 is therefore a mystery, but we will improve technical noise sources under the assumption that it follows its room temperature value.

Thermo-elastic noise is another fundamental source we have not discussed, where temperature fluctuations drive gradients across the mirror surface leading to local deformations. This effect scales as $1/T^2$ and is negligible in cryogenic systems, so we will not cover it here. It is worth noting that the thermo-elastic and thermo-refractive effects that comprise thermo-optic are coherent, and can therefore be coherently cancelled with clever design of the coating Bragg stack [75]. This is fundamentally different from Brownian thermal noise which is incoherent.

2.3 Cavity locking

Transferring the cavity stability to optical signal that can be distributed and used is a nontrivial task that was a significant barrier to optical interferometry in the late 20th century¹. Locking

¹ Technical details relating to cavity locking implemented on Si6 will be discussed in the next chapter

a laser to a reference requires a dispersive signal that allows a servo to identify the correct sign of frequency corrections. Transmission through the cavity linewidth provides an absorptive signal so locking at the transmission peak does not reveal information about the parity of the frequency error. The side of the line however can be approximated as linear so the transmitted power when detuned at the cavity half linewidth can be a suitable locking point. For a while this "side of line" locking was a ubiquitous method of frequency locking but it is highly sensitive to fluctuations in the DC value of the transmitted power. White [4] noted that the derivative of the cavity reflection function provides a suitable frequency discriminant for locking to the transmission peak, and enacted a scheme where the laser is dithered within the cavity resonance to scan out the derivative of the cavity transmission. The cavity converts frequency modulation into amplitude modulation which can be easily recovered through photodetection. Similar to side of line locking, the error signal quickly drops to zero, making it difficult to correct for frequency excursions larger than a half linewidth. The bandwidth of the lock is also limited by the cavity due to the delay associated with the intracavity storage time.

Hänsch and Couillard developed a unique method of generating a dispersive error signal by placing a linear polarizer within a cavity and encoding frequency fluctuations in polarization ellipticity [6]. This method however is not applicable to rigid cavities where a solid spacer prevents access between the mirrors. However, this was an important demonstration that an error signal could be obtained interferometrically between a beam reflected from the cavity input mirror and the leakage beam that entered the cavity. This method makes use of the imaginary component of the cavity reflection function - the beam that enters the cavity experiences a frequency dependent phase shift whereas the beam that does not can serve as a phase reference.

Both of these methods suffer from the issue of detecting a signal at baseband, where technical noise of the laser and RF environment can corrupt detection of the cavity signal. A clever technique was pioneered by Pound [76] for locking an oscillator to microwave cavities, where a phase modulator imprints sidebands nonresonant with the cavity. The error signal is then encoded high in the RF regime where technical effects are negligible. This was adapted to optical oscillators by Drever and

Hall and the resulting Pound-Drever-Hall (PDH) technique [5] has become the standard approach for transferring cavity stability to an oscillator.




Figure 2.6: Comparison of lineshapes of the various locking schemes discussed. The cavity transmission function (red) is used to enact side of line locking. Its derivative is used in White’s method [4] for modulation the laser frequency. The true dispersive lineshape is obtained via the PDH [5] and Hänsch-Couillard schemes [6].

A phase modulator places sidebands at an offset frequency of Ω around a carrier of frequency ω with modulation index β . The incident electric field becomes

$$E_i = E_0 e^{i(\omega t + \beta \sin(\Omega t))} \quad (2.35)$$

Expansion of the phase modulation term to first order with Bessel functions gives

$$E_i = E_0 (J_0(\beta) e^{i\omega t} + J_1(\beta) e^{i(\omega + \Omega)t} - J_1(\beta) e^{i(\omega - \Omega)t}) \quad (2.36)$$

Note that the sidebands have an opposite amplitude of electric field. As we will see later, this ensures that the baseline value of the error signal far from the cavity resonance is close to zero. Additionally, there will be no photocurrent from the cavity signal when locked to the center of the resonant feature.

The field reflected from the cavity E_r will be modified by complex components relating to

the cavity transmission and phase shift

$$E_r = E_0(J_0(\beta)e^{i\omega t}R(\omega) + J_1(\beta)e^{i(\omega+\Omega)t}R(\omega+\Omega) - J_1(\beta)e^{i(\omega-\Omega)t}R(\omega-\Omega)) \quad (2.37)$$

Our photodetector will measure the optical power $P_r = E_r^*E_r$. Components of the optical power will be present at DC, and also oscillate at frequencies $\Omega, 2\Omega$, etc. This is the key advantage of the PDH method: the cavity error signal is present at the sideband frequency Ω which is generally chosen to be much larger than the cavity linewidth so that it serves as the phase reference for the phase shifted intracavity light. Components at DC and 2Ω are filtered out so that the only signal we receive is the power at Ω

$$P_r = E_0^2 J_0(\beta) J_1(\beta) \left[(R(\omega)R^*(\omega+\Omega) - R^*(\omega)R(\omega-\Omega))e^{-i\Omega t} \right. \quad (2.38)$$

$$\left. + (R^*(\omega)R(\omega+\Omega) - R(\omega)R^*(\omega-\Omega))e^{i\Omega t} \right] \quad (2.39)$$

Demodulating the signal at Ω recovers the PDH error signal

$$V_r \propto 2E_0^2 J_0(\beta) J_1(\beta) \left[\text{Re } \chi(\omega) \cos \varphi + \text{Im } \chi(\omega) \sin \varphi \right] \quad (2.40)$$

Where φ is the demodulation phase between the error signal at Ω and the local oscillator driving the EOM with modulation Ω . The quantity

$$\chi(\omega) = R(\omega)R^*(\omega+\Omega) - R^*(\omega)R(\omega-\Omega) \quad (2.41)$$

Provides the dispersive signal shown in Fig.2.6.

There are a few observations we can make that will provide some intuition into the nature of the error signal. First, we can see that the signal is maximized when the Bessel product $J_0(\beta)J_1(\beta)$ is maximized. This occurs for a modulation depth of $\beta \sim 1$, where the power in the carrier is 3 dB higher than the sum of power contained in the two sidebands. Second, the imaginary component of $\chi(\omega)$ contains the dispersive signal so the error signal is maximized when $\phi = \pi/2$. This condition is typically enforced by adding delay lines of appropriate length if the demodulation oscillator is analogue, or by directly adjusting the signal phase if it is digitally synthesized.

The approximate slope of the frequency discriminator via the PDH method can be obtained by making a few approximations. When the sidebands are placed far outside the cavity transmission feature, it is reasonable to set $R(\omega \pm \Omega) \approx 1$. When the laser is within the cavity linewidth, we can approximate

$$R(\omega) = -\omega \frac{\omega + i\Gamma/2}{(\Gamma/2)^2 + \omega^2} \approx -i \frac{2\delta\omega}{\Gamma} \quad (2.42)$$

By additionally setting $E_0^2 J_0(\beta) J_1(\beta) = \sqrt{P_c P_s}$ for carrier power P_c and sideband power P_s , we find that the slope

$$P_r = 4\sqrt{P_c P_s} \text{Im } R(\omega) \approx 8 \frac{\sqrt{P_c P_s}}{\Gamma} \delta\omega \equiv \mathcal{D} \delta\omega \quad (2.43)$$

This quantity $\mathcal{D}$ is often referred to as the discrimination slope and dictates the magnitude of the PDH signal for cavity locking. More optical power increases the sensitivity but often at the cost of increased sensitivity to local heating effects in the mirror coating. A narrower linewidth also provides a larger slope, which is achieved by making a longer cavity with higher finesse mirrors.

One interesting point of note is the difference in discrimination slopes between side of fringe locking and PDH. The Lorentzian formula for the power in cavity transmission, and the slope is

$$P(\omega) = \frac{P_t}{1 + \left(\frac{\omega}{\Gamma/2}\right)^2} \quad (2.44)$$

$$P'(\omega) = \frac{8\omega}{\Gamma^2} \left[1 + \left(\frac{\omega}{\Gamma/2}\right)^2\right]^{-2} P_t \quad (2.45)$$

At the half linewidth where one operates for side of fringe locking

$$P'(\omega = \Gamma/2) = \frac{P_t}{\Gamma} \quad (2.46)$$

Compare this to the PDH discriminator $\mathcal{D}$, where we operate at the optimal point $P_c = 2P_s$. Then we have

$$\mathcal{D} = \frac{\sqrt{8}P_0}{\Gamma} \quad (2.47)$$

There is a slight advantage in signal with the PDH method due to the fact that only half the incident power is used in side of fringe locking. However, the transmitted power P_t is distinct from




Figure 2.7: **Top left:** Relative power in the carrier, first, second, and third order sidebands for varying modulation index β . The total optical power is normalized to unity. At $\beta = 1$, roughly half the power is contained in the carrier while each first order sideband contains 25%. **Top right:** Discriminator slope $\mathcal{D}^{-1}$ for different modulation depth. Again, maximum sensitivity is achieved close to $\beta = 1$. Note that there are values of modulation index where there is nearly no sensitivity to the PDH signal. At $\beta = 2.3$, power in the carrier drops to zero meaning no light enters the cavity. At $\beta = 3.8$, there is still power in the carrier but there is no electric field oscillating at 1Ω . A PDH signal can be recovered by using the 2Ω sidebands as a phase reference. **Bottom:** Error signals measured for varying modulation index. Curves are offset vertical and horizontally for clarity.

the incident power P_0 as seen in Eq.2.8. Even if the cavity coupling is higher, high mirror loss will lead to a P_t smaller than P_0 .

Recall that our calculation of $\mathcal{D}$ made the approximation that the laser was well within the laser linewidth. For frequency excursions outside the cavity linewidth, the cavity pole becomes a

significant feature and has the effect of lowpassing the signal. The frequency dependent optical gain then becomes

$$\mathcal{D}_f(f) = \frac{\mathcal{D}}{\sqrt{1 + 4f^2/\Gamma^2}} \quad (2.48)$$

We've identified the magnitude of the signal that can be recovered from the cavity, it is also instructive to see how this relates to the noise that is measured. The signal that arrives on our photodetector is fundamentally limited by the optical shot noise of the incident light. If the carrier is fully coupled through the cavity, the shot noise limited spectrum of power fluctuations at the detector will be

$$S_P = 2h\nu(2P_s) \quad (2.49)$$

Shot noise will be indistinguishable from the amplitude fluctuations that comprise the error signal. This will therefore set a limit on the fundamental level of frequency noise that we can resolve and we calculate this limit by simply using the magnitude of the cavity's frequency discriminant

$$S_f = \frac{S_P}{\mathcal{D}^2} = \frac{4h\nu P_s}{64} \frac{\Gamma^2}{P_c P_s} = \frac{h\nu \Gamma^2}{16P_c} \quad (2.50)$$

The spectrum of frequency fluctuations S_f is set entirely by the cavity linewidth and the power contained in the carrier. This sets a fundamental limit on the level of frequency noise that the light allows us to resolve. For a cavity with 10 kHz linewidth at 1542 nm, a carrier of 100 nW only allows a fractional frequency resolution of 1×10^{-17} at a 1 Hz bandwidth. This is surprisingly not far from the optical powers we tend to use in the Si3 and Si4 systems, next generation cavities will need to be mindful of this.

2.4 Frequency stability of oscillators

We've discussed the fundamental noise sources that limit the cavity instability, as well as fundamental sources introduced when transferring the cavity stability to a laser. Lower noise is always better in terms of making a more stable oscillator, but we should understand why this might

be important. In particular, when the laser is distributed to the strontium experiments, what parts of the frequency spectrum are important and which parts should we focus our efforts on improving?

First, let's consider how we convert the frequency noise spectrum into a more intuitive quantity. Measurement precision is more intuitively thought of in terms of fractional resolution at a particular timescale, e.g. a standard HP33401 multimeter used throughout the lab can measure with ppm resolution over 1 second of measuring. How the resolution increases with integration or averaging time depends on the noise mechanism limiting the measurement stability. More generally, the spectrum of noise will determine whether the measurement stability worsens, improves, or stays constant. The sample variance $\sigma_y^2(\tau)$ was developed as a metric to characterize this stability for an averaging time τ [77]. The Allan variance is a particularly useful form of variance that we discuss here which does not diverge for particular noise processes in the $f \rightarrow 0$ limit. This can be directly related to the fractional frequency noise spectrum of a quantity $S_y(f)$ through

$$\sigma_y^2(\tau) = 2 \int_0^\infty S_y(f) \frac{\sin^4(\pi f \tau)}{(\pi f \tau)^2} df \quad (2.51)$$

It is difficult to extract an intuitive relation between $\sigma_y^2(\tau)$ and $S_y(f)$ so let's instead consider various slopes on the spectrum of $S_y(f)$ and calculate the corresponding behaviour of the variance.

We can express

$$S_y(f) = \sum_{\alpha=-2}^2 h_\alpha f^\alpha \quad (2.52)$$

This summation contains the five typical characterizations of frequency noise: random walk ($\alpha = -2$), flicker frequency noise ($\alpha = -1$), white frequency noise ($\alpha = 0$), flicker phase noise ($\alpha = 1$), and white phase noise ($\alpha = 0$). The relation between the variance and frequency noise spectrum for equivalent values of α is displayed in Table 2.1. Note that we have used the relation between phase noise and absolute frequency noise $S_\nu(f)$

$$S_\nu(f) = f^2 S_\phi(f) \quad (2.53)$$

There are a few interesting points to note from the variance calculation. White frequency noise in variance scales as $\sigma_y^2(\tau) \sim 1/\tau$ and can thus be averaged down by measuring longer. While

fundamental sources like photon shot noise appear to be a stringent limitation in the noise spectrum, they do not limit the measurement precision - they only limit the amount of time it takes to average down to a particular precision. A lower white noise floor is nevertheless beneficial for the purposes of reducing the measurement time.

Thermal noise in the other hand which scales as $S_y(f) \sim 1/f$ cannot be averted by integrating the measurement over longer times. We can see that $\sigma_y^2(\tau) \sim h_{-1}$ and is independent of τ . An oscillator that is thermal noise limited will therefore never be able to measure quantities below a particular precision. This highlights the particular need to mitigate thermal noise with clever strategies when using oscillators in metrology.

Type of noise	α	$S_y(f)$	$\sigma_y(\tau)$	mod $\sigma_y(\tau)$
Random walk frequency noise	-2	$h_{-2}f^{-2}$	$\frac{2}{3}\pi^2 h_{-2}\tau$	$0.82\sigma_y^2(\tau)$
Flicker frequency noise	-1	$h_{-1}f^{-1}$	$2\ln(2)h_{-1}$	$0.67\sigma_y^2(\tau)$
White frequency noise	0	h_0f^0	$\frac{1}{2}h_0\tau^{-1}$	$0.5\sigma_y^2(\tau)$
Flicker phase noise	1	h_1f^1	$\frac{1+3\ln(2\pi f_H\tau)}{4\pi^2}h_{-1}\tau^{-2}$	$\frac{6.25}{8\pi}h_{-1}\tau^{-2}$
White phase noise	2	h_2f^2	$\frac{3f_H}{4\pi^2}h_{-2}\tau^{-2}$	$\frac{3}{8\pi}h_{-2}\tau^{-3}$

Table 2.1: Relations between the fractional frequency noise spectrum $S_y(f)$, variance $\sigma_y^2(\tau)$, and modified variance mod $\sigma_y^2(\tau)$. Calculation of these relations is taken from [1].

In the high frequency regime the variance does not discriminate between flicker and white phase noise, both scale as $1/\tau^{-2}$. These can originate from entirely different processes and so it is useful to be able to identify which type of noise may be a limitation. Furthermore, the magnitude of both noise processes is dependent on the bandwidth of the frequency noise measurement - a high pass filter with a corner f_H is needed to avoid a divergence in the calculation of the variance (particularly a divergence in the integral of Eq.2.51). The "modified variance" mod $\sigma_y^2(\tau)$ was developed for this purpose as it does not have these shortfalls. Equivalent conversion from the frequency noise spectrum is displayed in Table 2.1. The modified Allan variance is calculated as

[78]

$$\text{mod } \sigma_y^2(\tau) = 2 \int_0^\infty S_y(f) \frac{\sin^6(\pi f \tau)}{(\pi f n \tau)^2 \sin^2(\pi f \tau / n)} df \quad (2.54)$$

Where $\tau = n\tau_0$ and τ_0 is the minimum data interval.




Figure 2.8: Calculation of the frequency domain filter functions $|W_a(f)|$ and $|W_m(f)|$ for the Allan and modified Allan variances. The modified Allan variance has less susceptibility to high frequency noise

Oftentimes a time series of the signal is measured directly so it is desirable to be able to calculate the variances directly from this data. To see how this follows from Eqs. 2.51, 2.54, note that the quantity within the integral can be interpreted as a filter to the frequency spectrum $S_y(f)$. We can define separate filters for the variance and modified variances as $W_a(f)$, $W_m(f)$ respectively

$$W_a(f) = \sqrt{2} \frac{\sin^2(\pi f \tau)}{(\pi f n \tau)} \quad (2.55)$$

$$W_m(f) = \sqrt{2} \frac{\sin^2(\pi f \tau) |\sin(\pi f \tau)|}{(\pi f n \tau)^2} \quad (2.56)$$

Such that the variance is calculated as

$$\text{mod } \sigma_y^2(\tau) = \int_0^\infty S_y(f) |W(f)|^2 df \quad (2.57)$$

We can relate both frequency domain filters to time domain filters via Parseval's theorem.

The variances can then be obtained directly from the time series of data $y(t)$ by

$$\sigma_y^2(t) = \left\langle \left[\int_{-\infty}^{\infty} y(t) w_a(t) dt \right]^2 \right\rangle \quad (2.58)$$

$$\text{mod } \sigma_y^2(t) = \left\langle \left[\int_{-\infty}^{\infty} y(t) w_m(t) dt \right]^2 \right\rangle \quad (2.59)$$

In the time domain, we can calculate these filters to look like

$$w_a(t) = \begin{cases} \frac{1}{\sqrt{2}\tau} & 0 < t \leq \tau \\ \frac{1}{\sqrt{2}\tau} & \tau < t \leq 2\tau \\ 0 & t \notin (0, 2\tau] \end{cases} \quad (2.60)$$

$$w_m(t) = \begin{cases} -\frac{1}{\sqrt{2}\tau^2}t & 0 < t \leq \tau \\ \frac{1}{\sqrt{2}\tau^2}(2t - 3\tau) & \tau < t \leq 2\tau \\ -\frac{1}{\sqrt{2}\tau^2}(t - 3\tau) & 2\tau < t \leq 3\tau \\ 0 & t \notin (0, 3\tau] \end{cases} \quad (2.61)$$

A illustration of these filters is given in Fig.2.9. The $w_a(t)$ has a flat weighting within the intervals of width τ around $t = \tau$, whereas $w_m(t)$ has a sort of triangular sensitivity. These are correspondingly called Π and Λ filters after the apparent shape of their sensitivities. The shape of these filters also provides some intuition as to why the the modified variance may provide less sensitivity to low frequency noise sources as seen in Table 2.1. The resulting variances for the different types of noise presented in Table 2.1 are compared in Fig.2.10.

These techniques are all useful for characterizing continuous noise processes, i.e. noise that is consistently present in the frequency noise spectrum. In reality, measurement of the cavity frequency will occasionally encounter aperiodic perturbations such as frequency jumps. This can have the effect of deteriorating the stability, but also making it look like there are additional noise




Figure 2.9: Calculation of the time domain filter functions $w_a(t)$ and $w_m(t)$ for the Allan and modified Allan variances.




Figure 2.10: Simulation that compares the variance and modified variance for different values of α where $S_y(f) \propto f^\alpha$.

processes at low frequency. An example is shown in Fig.2.11. A simulation of flicker frequency noise

has a series of discrete frequency jumps included. The modified variance already departs from the expected mod $\sigma_y(\tau) \sim \tau^0$ at short averaging times and scales as $\sigma_y(\tau) \sim \tau^{0.5}$. Without access to the time series, one could incorrectly infer that the dataset has a significant amount of random walk frequency noise. The variance and modified variance do not uniquely determine the nature of frequency excursions in a dataset. One should be aware that these may not be the correct tools to analyze non-continuous noise processes.

2.5 Frequency stability of atom-steered oscillators

We spent a considerable length going over the application of filter functions to a dataset, allowing us to calculate variances of the oscillator's stability. There are many software options that will do this for us automatically so it may appear to have been an unnecessary effort. However, the process of filtering a time series to calculate the free running oscillator stability is the exact process that we will use to determine the stability of an oscillator steered to an atomic reference.

Coherent spectroscopy techniques in clocks place an atom in a superposition between two states where it can evolve under the Hamiltonian of the laser's electric field. The spectroscopic lineshape relates fluctuations in the relative excited state - ground state population to a phase error in the oscillator. This mapping is dependent on the duty cycle T_c and spectroscopic time τ_d for which the atom is interrogated. A "dead time" during which the atom is not interrogated leads to the aliasing Dick effect [79, 80], where the laser phase evolves without being tracked by the atom. The stability of the atom-steered oscillator is dependent on these parameters since the atom is not continuously used to correct phase errors of the laser and is therefore blind to phase evolution during the dead time.

This mapping from the free running oscillator stability to the atom-steered oscillator stability is accomplished through filter functions - in particular, the sensitivity function of the atom. Rabi and Ramsey spectroscopy have different sensitivity functions and weight the f^α types of noise differently (similar to our investigation of variances from above). The Dick effect limited variance will also be different when the atom is interrogated with the two spectroscopic methods.




Figure 2.11: **Top:** We simulate datasets with frequency noise spectrum $S_y(f) \sim f^\alpha$ for $\alpha = -2, -1, 0$. Discrete frequency jumps of the same magnitude are added to each time series. **Bottom:** Solid circles indicate the modified variance calculated for the corresponding datasets from the plot above. Dashed lines are variances of the unperturbed white, pink, and brown noise without the discrete frequency jump. The frequency jump has the effect of deteriorating the stability of all three time series but also changes the long term instability to scale as $\text{mod } \sigma_y(\tau) \sim \tau^{0.5}$. Data that is dominated by white frequency noise and flicker frequency noise can still have variances that corresponding to random walk frequency noise if there is a discrete perturbation.

The Dick effect limited variance has analogues to the Allan variance from Eq.2.58. The filter function $r(t)$ for a spectroscopic sequence is multiplied by the time series of the laser frequency ($y(t)$ as discussed earlier). The end result however is not a variance of the laser frequency, rather a variance of the excitation fraction δf_{exc} during the atomic state readout. Quantitatively, this is expressed as [81]

$$\delta f_{exc} = \frac{1}{2} \int_0^{T_c} y(t)r(t)dt \quad (2.62)$$

Note that the integral runs over the cycle time T_c , not the spectroscopic duration τ_d . Preparation and readout of the atomic sample takes a finite amount of time and consequently most clocks operate with $\tau_d < T_c$. Similar to our depiction of the variance filter functions in Fig.2.9, this implies a lack of sensitivity to laser frequency noise over a duration $T_c - \tau_d$. The sensitivity function will acquire a value of zero (see Fig.2.12 for an illustration) and this dead time leads to the aliasing effect of Dick noise.

During a cycle of time T_c where interrogation of the oscillator begins at time $T_c - \tau_d$ for duration τ_d , the sensitivity functions for Rabi and Ramsey spectroscopy take the form

$$r_{Rabi}(t) = \begin{cases} 0 & 0 < t \leq T_c - \tau_d \\ \frac{2.509\pi^2}{\beta^3} \left(\sin(\beta) - \sin\left(\frac{t-T_c+\tau_d}{\tau_d}\beta\right) + \sin\left(\frac{t-T_c}{\tau_d}\beta\right) \right) & T_c - \tau_d < t \leq T_c \end{cases} \quad (2.63)$$

$$r_{Ramsey}(t) = \begin{cases} 0 & 0 < t \leq T_c - \tau_d \\ \sin\left(\left[t - T_c + \tau_d\right]\frac{\pi}{t_\pi}\right) & T_c - \tau_d < t \leq T_c - \tau_d + t_\pi/2 \\ 1 & T_c - \tau_d + t_\pi/2 < t \leq T_c - t_\pi/2 \\ \sin\left(\left[t - T_c + t_\pi/2\right]\frac{\pi}{t_\pi} + \frac{\pi}{2}\right) & T_c - t_\pi/2 < t \leq T_c \end{cases} \quad (2.64)$$

(2.65)

In the Rabi sensitivity, $\beta = \sqrt{\pi^2 + 2.509^2}$ and t_π refers to the time for a π pulse in the Ramsey sequence.

First, consider the case where $\tau_d = T_c$. We can see that the Rabi sensitivity function does not uniformly weight all portions of the laser frequency. The consequence of this is a decreased ability to track the laser phase. Phase excursions in the frequency range where the laser is dominated by $1/f$ noise will not be perfectly measured by the atom during Rabi spectroscopy. If it was possible to constantly interrogate the atom with sequential Rabi pulses, we would still find errors in our measured laser phase. Ramsey spectroscopy does not have this issue, as we see that the sensitivity is constant over the duration of atomic interrogation. However, the finite time required to drive an atom to $f_{exc} = 0.5$ leads to a softening of the curve at the beginning and end of the coherent spectroscopy. This will also have the effect of reducing the sensitivity to laser noise in some frequency band, but less so than in the case of Rabi spectroscopy.

With the introduction of deadtime ($\tau_d < T_c$), the picture becomes even less intuitive. Previously we found that the variance of a signal is expressed as the integration of the sensitivity function and the laser spectrum $S_y(f)$. In reality, the signal is sampled by a measurement device with a particular rate and the variance becomes a summation of the spectral density measured over integer multiples of each frequency bin.

The Dick effect follows the same formulation. The variance is calculated as a summation over discretized frequency bins for integer multiples of $1/T_c$. The sampled spectrum is then [82, 83],

$$S_d(f) = \sum_{k=1}^{\infty} \left| \frac{g_k}{g_0} \right|^2 S_y(k/T_c) \quad (2.66)$$

The laser spectrum is periodically sampled every k/T_c and the bin amplitude g_k is given by

$$g_k = \frac{1}{T_c} \int_0^{T_c} g(t) e^{-2\pi i k t / T_c} dt \quad (2.67)$$

The Dick effect limited variance is related to this sampled spectrum by

$$\sigma_y^2(\tau) = \frac{S_d}{\tau} \quad (2.68)$$

This model is insensitive to the type of spectroscopy being performed, only the sensitivity $g(t)$ will change. Surprisingly one finds that the sampled spectrum S_d is white frequency noise and

the Dick effect limited variance will then simply scale as $\sigma_y^2(\tau) \sim \tau^{-1}$. As we show in Fig.2.13, the shape of the laser spectrum has no determination on the relation between averaging time and variance. The amplitude of the variance is the only component that varies with the spectroscopy sensitivity function and the spectral components of the laser.




Figure 2.12: Calculation of the time domain filter functions $r_{Rabi}(t)$ and $r_{Ramsey}(t)$ for the two spectroscopic methods. In both cases we assume a cycle time of 5 s, a spectroscopy time of 3 s, and in the case of Ramsey a π pulse time of 100 s.

What does this mean in terms of the impact on the experimenter? If the Dick effect limited variance always scales as $1/\tau$, how should the laser spectrum be shaped? Intuition is gained by looking at the various Fourier components g_k . Fig.2.14 calculates a series of these for different spectroscopic sequences with the same cycle and interrogation time. The spectral window around $1/T_c$ has the largest amplitude in the summation of Eq.2.66 but harmonics of higher k are weighted differently depending on the spectroscopic sequence. For a given evolution time τ_d , Rabi spectroscopy can be sensitive to certain spectral windows where Ramsey spectroscopy has no sensitivity. Within Ramsey sequence, the π -pulse time can be intentionally chosen to shift the spectroscopic sensitivity to windows where the laser spectrum may have less noise, e.g. avoid mechanical resonances of the cavity or large ground loops introduced by electronics.




Figure 2.13: We calculate the Dick effect via Eq.2.66 for different laser spectra $S_y(f) = 10^{-33}f^\alpha$ with T_c . For reference, $S_y(f) = 10^{-33}f^{-1}$ is the thermal noise limit of Si3. The corresponding Allan deviation is insensitive to the type of noise that dominates the laser spectrum and always follows $\sigma_y(\tau) \sim \tau^{-1/2}$. However the amplitude of the Dick effect is determined by the slope of the laser spectrum, with a lower variance associated with larger values of α . Solid lines indicate variances calculated with the Rabi sensitivity function and dotted lines are calculated with the Ramsey sensitivity. As α decreases, Ramsey spectroscopy provides increasing improvements in stability over Rabi spectroscopy.

The point of this is to illustrate the importance of tailoring the spectroscopic sequence to interrogate the 'quieter' windows of the laser spectrum. Global benchmarks for laser noise such as laser linewidth do not uniquely determine the noise spectrum and often have little correspondence to the Dick effect limited variance that is realized in clock spectroscopy. Recent results in qubit control have come to identify the role that the local oscillator noise plays in determine qubit fidelity when driving rotations between the qubit states [84]. The physical process is exactly the same as what we have described in the past two sections - a filter function samples particular spectral

components of the local oscillator to determine a variance in excitation probability (2.62). Driving fast qubit rotations does not make the atom immune to this process, the spectral window for sensitivity is simply shifted to a higher bandwidth.




Figure 2.14: The various amplitude coefficients $|g_k/g_0|^2$ are calculated for different spectroscopic method. All traces use $T_c = 5$ s, $\tau_d = 4$ s. The gold curve uses a conservative $t_\pi = 0.2$ s whereas the green trace assumes infinitely short pulses of $t_\pi = 0$. The majority of the Rabi sensitivity is concentrated around frequencies of $1/T_c$ whereas Ramsey spectroscopy has a higher contribution from the high frequency harmonics of the cycle time.

Chapter 3

Experimental setup

Having presented the primary concepts necessary to design, operate, and characterize a cavity, we'll discuss now these realizations in the Si6 system. We'll detail out choice of silicon as a spacer material, operation of the 4 K cryostat, unique features of crystalline AlGaAs/GaAs mirrors, and characterization of the Si3 cavity which we use as a reference oscillator.

3.0.1 Silicon in cavities

It's evident that three key parameters that dictating the Brownian thermal noise limit are E, ϕ , and T . For improved oscillator stability, a tempting option is to reduce the cavity temperature. The $S_y \sim T$ scaling indicates that modest reductions below room temperature will have little effect, rather cooling into the cryogenic regime can lead to substantial improvements. Technical issues of cryogenics aside, the advantage of ULE cavities are their relatively low insensitivity to room temperature fluctuations. A material's length sensitivity δL to temperature changes δT is given as

$$\Delta L/L = \alpha(T)\delta T + \frac{1}{2}\alpha'(T)^2\delta T^2 + \mathcal{O}(\delta T^3) \quad (3.1)$$

For a coefficient of thermal expansion (CTE) α . Particular materials have $\alpha = 0$ at certain temperatures, and ULE has a zero crossing of its CTE at 20°C. Residual temperature sensitivity is dominated by the much smaller second order sensitivity of $\alpha'(T) = 17 \times 10^{-9} \text{ K}^{-2}$ [85]. Temperature fluctuations of 0.1 mK, routinely achievable with active temperature stabilization, therefore only

leads to length fluctuations of 1×10^{-16} , close to the thermal noise limit of most room temperature ULE cavities. Without a zero crossing of the CTE, equivalent magnitude of thermal fluctuations would require $\alpha(T) = 1 \times 10^{-12}$ which is an absent property for room temperature materials.

The existence of the zero CTE (ZCTE) is therefore seen as a soft requirement for thermal noise limited cavity operation. This is also relevant in the cryogenic regime, where temperature stabilization is often more difficult due to the added thermal fluctuations of the cryocooler and the reduced thermal mass of most materials. The ZCTE property is not ubiquitous to all materials - rather, few are found to have it.

Fortuitously, silicon is a material that has been identified to have not one but two ZCTEs located at 124 K and 16 K [86] (Fig.3.3). Operation at these temperatures is achieved with different cooling technologies: the Si3 cavity in our lab operates at 124 K where temperature stabilization is accomplished by boiling off 77 K liquid nitrogen gas, and 16 K operation can be achieved with a cryocooler using cold helium circulation. The Si6 cavity operates at both 16 K and 4 K, where the CTE is not zero but approaches 1×10^{-11} and has a inconsequential α' sensitivity. The silicon loss angle at these temperatures is found to be at the 10^{-8} level [87], three orders of magnitude lower than the ULE loss angle of 1.6×10^{-5} [7]. Note that many materials are found to have temperature dependent ϕ so operation temperature of the cavity is an important point to consider.

A large Young's modulus is also beneficial for reduced thermal noise. ULE is already fairly stiff with $E = 68$ GPa but our typical crystalline materials are again found to have advantageous properties. Silicon is found to have a Young's modulus between 130-194 GPa depending on the crystal orientation [88], while sapphire has a considerably higher value of $E = 400$ GPa. The anisotropic structure of the crystal is leveraged in unique ways which will be detailed in the following chapter.

Ultimately we choose to use silicon as our spacer material and this has proved to be a fruitful collaboration with our colleagues at PTB, Germany. The Si6 spacer is grown as a high purity monocrystalline ingot using the float-zone method. X-ray crystallography identifies the stiffest $\langle 111 \rangle$ axis along which the ingot is to be grown. A check of the crystal purity is done by measuring




Figure 3.1: Loss angle ϕ temperature dependence for commonly used cavity spacer materials. ULE is not shown on this plot but has a measured room temperature loss angle of 1.6×10^{-5} [7]. Crystalline spacer materials have a ϕ so low that the total Brownian thermal noise limit of crystalline cavities is limited by the loss angle of the optical coating. Figure taken from [8].




Figure 3.2: Coefficients of thermal expansion for commonly used cavity spacer materials. Between the two crystalline materials of sapphire and silicon, only the latter has a ZCTE. Fused silica can be an attractive option due to a ZCTE around 170 K but the significantly higher loss angle means that a silicon cavity will still have a superior thermal noise limited performance. Figure taken from [8].

the resistivity along the length of the spacer, and a value of 30,000 Ω cm is found. Subsequently it is sent to several US and European vendors for machining, polishing, lapping the ends, and etching the surface.

The spacer is cut with the same geometry as Si4. With the operating temperature and material properties being the same, we can calculate the spacer Brownian noise as $S_{sp} = 8 \times 10^{-38}/f$.

Substrates are cut from the same ingot with same orientation, $\langle 111 \rangle$ being normal to the coating surface. In 2014, Coastline optics was only able to polish substrates to a maximum of 1 m radius of curvature. With a cavity length of 6 cm, we can calculate the waist of the fundamental TEM_{00} mode supported by the cavity and find $w_0 = 286 \mu\text{m}$. At the mirror locations, the spot size is found to be $w(L/2) = 294 \mu\text{m}$. We now have all the parameters needed to calculate the substrate thermal noise and find $S_{sb} = 2 \times 10^{-35}/f$. So far, the total fractional thermal noise limit of the cavity is found to be

$$S_{sp} + 2S_{sb} = 4 \times 10^{-35}/f \quad (3.2)$$

This is remarkably low and would imply a fractional frequency instability of $\sigma_y = 6 \times 10^{-18}$, far lower than realized in state of the art optical resonators. As we will now see, the "dominant" limitation to the realized frequency noise is due to the coating.

ULE		Silicon		Sapphire	
ϕ	1.6×10^{-5}	ϕ	1×10^{-8}	ϕ	2×10^{-8}
E	68 GPa	E	190 GPa	E	400 GPa
σ	0.18	σ	0.29	σ	0.28

Table 3.1: Comparison of bulk substrate and spacer materials.

3.0.2 AlGaAs coatings

A key demonstration of Si2, Si3, and Si4 was thermal noise limited performance with instabilities in the mid 10^{-17} s. We've seen that silicon substrate and spacer noise is incredibly small, so it stands to reason that the coating thermal noise overshadows these other two. The larger thermal




Figure 3.3: Coefficient of thermal expansion of silicon across the three ranges of operation, 124 K, 16 K, and 4 K. At 124 K, $\alpha' = 2 \times 10^{-8} \text{ K}^{-2}$, similar to the CTE slope around the ULE room temperature ZCTE. The slope is considerably shallower at 17 K with $\alpha' = 1 \times 10^{-9} \text{ K}^{-2}$ but requires different cooling technology. There is no zero crossing of α at 4 K but values of 10^{-11} can be achieved, requiring only μK temperature stability for 1×10^{-17} cavity instability.



Figure 3.4: Two 55 cm silicon ingots grown at IKZ Berlin. Image courtesy of Thomas Legero.

noise of the coating can be attributed to one quantity, the loss angle ϕ .

Loss angles of optical coatings are consistently found to be larger than that of the spacer and 3-4 orders of magnitude higher than silicon. AlGaAs/GaAs thin film coatings have a measured room temperature loss angle of 2.5×10^{-5} [19, 89]. Compared to silicon, the larger coating loss angle would imply a 1000 times larger thermal noise when compared to S_{sb} but recall that coating thermal noise is modified by a factor of d/w , roughly 1/30 in Si6. It is peculiar that the ϕ of a crystalline coating is slightly larger than that of the ULE spacer when we've stated that crystalline materials should have higher mechanical quality factors than amorphous ones.

This is an issue particular to the thin film nature of the coating. The bulk-melt SiO_2 substrates of the LIGO test masses have a measured Q of 10^7 at room temperature [90] but IBS sputtered SiO_2 films have a $\phi = 4 \times 10^{-5}$ [91]. The origin of significantly higher ϕ in thin films has been the source of intensive study. Theories that the majority of dissipation could be occurring between the interfaces of the many layers were debunked in a study where the total coating thickness was kept constant and the number of layers was varied [92]. Similarly, varying the coating thickness and observing a proportional increase in coating thermal noise was crucial in removing the substrate-coating interface as an additional source of dissipation.

In conventional dielectric coatings alternate quarter wave stacks of high index Ta_2O_5 and low index SiO_2 materials use interference of the Fresnel reflections to generate a high reflectivity surface. The Ta_2O_5 loss angle is found to be $(4.2 \pm 0.4) \times 10^{-4}$ [93], significantly higher than that of the silica layers. Reducing the coating's effective loss angle is therefore best accomplished by reducing the loss of the high index layers, or finding a substitute material with lower loss. Consequently there is still a dedicated effort to fabricate high- Q amorphous coatings by varying the dopants and annealing process

Titania has been found to be a promising dopant within the tantala-oxygen structure for reducing the overall loss [94]. Introduction of up to 40% TiO_2 has been found to reduce the tantala loss angle by over a factor of two to $\phi = (1.8 \pm 0.1) \times 10^{-4}$. Loss angles between samples of tantala from other vendors have differed, suggesting that the fabrication process may be extremely

important in engineering low loss [95]. This has led to an interest in investigating loss angles of coatings with different annealing and sputtering parameters. Extended annealing time of SiO_2 coatings has demonstrated a reduction in ϕ [96] but the improvement varies depending on the vendor.

Recent studies of Ta_2O_5 mixed with GeO_2 and annealed at high temperature have emerged as promising candidates for the high index layers [97]. Loss angles as low as $\phi = 1 \times 10^{-4}$ have been measured in these samples. Fine tuning annealing parameters has also demonstrated a slight decrease in Ti_2O_5 loss angle [98]. Distinction between high temperature deposition and high temperature annealing and varying the two separately has also been beneficial.

Next generation gravitational wave detectors such as the aLIGO+ upgrade requires a coating thermal noise corresponding to a loss angle of 0.9×10^{-4} . Future cryogenic detectors such as Voyager [56] and the Cosmic Explorer [99] operating at 124 K require even lower loss angles of $\phi < 6 \times 10^{-5}$. Exploring the parameter space of doping and material deposition has spawned promising advances in amorphous coating technology, but still far from the requirements of these future experiments. A deeper understanding of coating loss angles is needed.

Several studies have now revealed that the emergence or suppression of long range order in amorphous materials is correlated with the magnitude of a material's loss angle [100, 101]. Bulk glasses that are close to being perfectly amorphous can still have a remarkably high Q , but this is reduced in thin film amorphous layers where the emergence of some order makes it an imperfect glass [102].

A similar phenomenon exists in crystalline materials, where the long range order of a crystal supports a high mechanical quality factor. Deviation from perfectly amorphous or perfectly crystalline material appears to result in a degradation of quality factor. Measurements of bulk GaAs at cryogenic temperatures on a large 3" diameter by 1" thick disk found a $Q \sim 10^7$ [103]. Note that this is reduced by a factor of over 200 when epitaxially grown as a thin film. Still, the AlGaAs coating loss angle is sufficiently low to meet the requirements of next generation interferometers and cavities. Further reductions in the thin film ϕ have not been thoroughly explored for this reason,

and it is possible that the high index GaAs and low index AlGaAs layers may also have differing Q 's.

AlGaAs/GaAs coatings are not the only demonstration of high- Q single crystal structures that can be used in mirrors. Measurements of InGaP membranes in cavity optomechanics experiments [104] reveal a $Q = 2 \times 10^6$ at room temperature and 17 K. Measurements are conducted at offset frequencies $f > 100$ kHz far from the interrogation band for ultrastable resonators, but this demonstrates that there are other phosphides and nitrides that could be studied as next generation coatings.




Figure 3.5: Ringdown measurement of Si6 transmitted light with the cavity at vacuum and cooled to 4 K. No difference in finesse is measured between 4 K and 16 K. A finesse of 290,000 is measured.

The mirror coatings in Si6 are composed of an 11.6 μm -thick multilayer stack of alternating high-index ($n = 3.3$) GaAs and low-index ($n = 2.9$) $\text{Al}_{0.92}\text{Ga}_{0.08}\text{As}$, and are oriented with parallel coating polarizations. The GaAs $\langle 100 \rangle$ axis is normal to the coating, and parallel to the crystalline $\langle 111 \rangle$ axis of the crystalline silicon. Two near-planar mirrors of 1 m radius of curvature realize cryogenic transmission (T) of 11.8 ppm and loss (L) of 10 ppm combined for both mirrors, supporting a finesse of 290,000. We note that at ambient temperature and pressure, we measure a finesse of 220,000 which increases to 240,000 upon pulling vacuum on the cavity.

Coating parameter	Value
Loss angle (ϕ)	2.5×10^{-5}
Young's modulus (E)	85 GPa
Poisson ratio (σ)	0.32
Coating thickness (d)	$11.6 \mu\text{m}$
Mirror radius curvature	1 m
Beam waist (w_0)	$286 \mu\text{m}$

Table 3.2: Parameters of the AlGaAs/GaAs coatings

We measure a 770 kHz birefringent splitting of the polarization modes, well resolved outside of the 8.6 kHz cavity linewidth. A slight temperature and optical power dependence of the mode splitting has been observed (Fig.3.8).




Figure 3.6: Brownian noise contributions for the four silicon cavities referenced through this thesis. Coating thermal noise is consistently the dominant contribution due to the exceptionally low loss angle of silicon. Coating materials with high mechanical quality factor are therefore instrumental in developing optical oscillators with longer coherence times.




Figure 3.7: Reflectance of the two crystalline mirrors used on Si6 at a temperature of 4 K. Data is from Garrett Cole. We use a 1542 nm laser to lock to the cavity, centered within the ~ 150 nm passband of the coatings.




Figure 3.8: The mode splitting of Si6 is observed to vary with temperature and intracavity power. We plot the change in mode splitting normalized to the 8.6 kHz cavity linewidth.

3.0.3 Cryostation

The cavity is housed in a custom-designed cryostation from Montana Instruments capable of reaching base temperatures of 3.5 K. An RDK-101D GM cryocooler from Sumitomo Heavy Industries provides cooling power to the sample space within the cryostation. Temperatures of

the stage 1 and 2 are dependent on the applied heat load, a capacity map of the cooling power is provided in Fig.3.9. Montana provides thermistors installed on the stage 1 and 2 of the cryo, and measurements of their temperatures provide our best estimate of the heat load on each shield. With the cavity in place we typically find a stage 1 temperature of 28 K and stage 2 of 2.7 K, indicating heat loads of roughly 1 W and <30 mW respectively. A total of three thermal shields are used to successively step down the base temperature and reduce the radiation heat load. The outermost shield is attached to the 30 K platform, and the two inner shields are both at 4 K. The outer of the two is used for active stabilization with a 50 Ohm manganin heater wrapped around its base, and the inner passively damps thermal fluctuations at 4 K. Thermal shields supplied by Montana are made of 6061 aluminum, though these have been substituted with high purity copper shields constructed by the JILA machine shop. Details on this will be provided in the next chapter.

Thermal sensors on each shield allow for measurement of time constants for heat flow between stages and as witnesses for temperature fluctuations on each surface. A 20 pin vacuum feedthrough at the cryostation base allows us to read these out with our own measurement devices.


Four sets of UVFS superpolished windows from Coastline optics are placed above and below the cavity, one for the 300 K vacuum chamber, one for the 30 K radiation stage, one for the 4 K active stage, and one for the 4 K passive stage. Each is mounted with a 6° wedge relative to the cavity substrate, and offset in rotation to prevent low angle scatter and etalons between windows. A serrated aluminum clamp holds each window in place with a thin layer of Apiezon-N vacuum grease between the optic and mount to ensure good thermal contact and prevent stresses from developing across the window during cooldown.

A large brass box with acoustic foam on its interior encloses the optics to lowpass any room temperature and pressure fluctuations. A secondary box was installed on the AVI around the optics further dampen pressure fluctuations. Quarter inch thick rubber curtains around the optical table serve as a final layer for decoupling the optics environment from the lab.

The coldhead generates significant vibrations while operating, so care has to be taken to decouple this from the cavity as much as possible. The cryocooler sits atop a pillar directly mounted

RDK-101D Cold Head Capacity Map (60 Hz)

With F-20L Compressor and 3 m (10 ft.) Helium Gas Lines



Specifications subject to change without notice.
Note: Capacity maps for reference only.

©SHI Cryogenics Group 7/21

Figure 3.9: A capacity map for our cryostation's cooler displays the cooling power available on each stage. We operate in the lower right corner of the plot with less than 100 mW heatload on stage 2 and roughly 1 W heatload on stage 1.

onto the concrete foundation of the JILA basement. The cryostation meanwhile sits atop a commercial TS-140 active vibration isolation (AVI) platform that provides acceleration noise feedback from 5-200 Hz. This is placed atop a long three foot high optical table that is shared by Si3. Note that there is not strong coherence between acceleration noise measured on the two systems. Vibrations inevitably will couple through the ground from the coldhead to cavity but this result in an amplitude reduction is the ideal mechanism if couplings cannot be avoided. The only direct mechanical link is through the horizontal copper braid that pumps heat from the cryostation to the cooler. The link design is not shared with us, but we suspect it is a braid that can be compressed and tensioned along its axis. Torquing the AVI has reduced vibration noise measured near the cavity, indicating that the coldfinger is more rigid along its transverse axis.

The helium compressor attached to the coldhead also generates significant vibrations and is a considerable heat source. We place this outside the lab and run long 10 m vacuum insulated hoses




Figure 3.10: Vibration noise measured on the AVI with vibration isolation enabled for all three axes.

to the cryo to decouple vibrations between the two sides. The system is entirely closed cycle and

has a hands-free operation once cooled to cryogenic temperatures. The coldhead needs a routine maintenance every 20,000 hours of runtime, though this is extended if the system is rarely thermally cycled between 4 K and 300 K. The maintenance procedure involves warming the system back to room temperature to exchange the coldhead for a new identical model, but the system is not vented to ambient pressure. Nevertheless, care has to be taken during the warmup to avoid the formation of ice on the colder surfaces of the cavity and 4 K windows.

A beryllium-copper bellows surrounds the coldfinger, connecting the chamber and coldhead at room temperature. Pulling vacuum on the system compresses the bellows so the AVI often needs to be tensioned away from the coldhead with a mechanical spring.

These all contribute to reducing the vibration noise measured at the cryostation, and further steps are taken to reduce the cavity length sensitivity to accelerations. Many of these design choices have been enacted with deliberation and the following chapter will expand on this to explain their merits. A schematic illustrated in Fig.3.11 displays several of these key features. The vacuum chamber and thermal shields are hidden from view to reveal the cavity, its support, and charcoal adsorbers mounted to the lid of the passive shield. The wedged window mounts, visible above the cavity, are oriented to suppress the formation of etalons.

3.0.4 Noise measurements

We're fortunate to have Si3 in continuous operation at a nearly identical wavelength. A heterodyne beat with Si6 appears in the RF domain where we can easily measure the frequency noise with commercial electronics. In [41], Si3 was demonstrated to be thermal noise limited to a fractional instability of 3.7×10^{-17} from 0.2-1000 s. Later studies [17] extended the thermal noise performance out to 10^4 seconds where the phenomenal precision and uptime of the Gen-2 SrI allows us to continuously measure its long term instability. Short term instability can be measured via three-cornered hat analysis with the 40-cm 698 nm ULE cavity. Eric Oelker's fantastic work on




Figure 3.11: A solidworks mockup of the Si6 cryostation, cavity, and part of the optical setup.

the Si3 system resulted in a cavity with noise parameters of

$$S_y = \sum_{i=-1}^0 h_i f^i + \sum_{i=1}^N \frac{a_i}{1 + \left(\frac{f-f_i}{\Gamma_i/2}\right)^2} \quad (3.3)$$

In a band of 1mHz - 100 Hz. The first term consists of a white noise floor and the 1/f thermal noise limit with fractional values

$$h_{-1} = 1.7 \times 10^{-33} \text{ 1/Hz} \quad (3.4)$$

$$h_0 = 4 \times 10^{-34} \text{ 1/Hz} \quad (3.5)$$

The second term in Eq.3.3 details various mechanical resonances with parameters given in Table 3.3.

At higher frequency where the 124 K cavity is no longer thermal noise limited, the nearly 10x lower linewidth of Si3 suggests that its white noise floor should easily be lower than that of Si6.

Index	f_i (Hz)	a_i (1/Hz)	Γ_i (Hz)
1	5.7	7×10^{-34}	1.0
2	12.7	1.5×10^{-34}	1.5
3	20.0	4×10^{-34}	0.1
4	30.0	5×10^{-34}	0.1
5	40.0	5×10^{-34}	0.1
6	45.0	1×10^{-34}	4
7	55.0	4×10^{-34}	1.2

Table 3.3: Resonances of Si3 identified in the noise model, with center frequency f_i , amplitude a_i , and width Γ_i .

This is all to say that we have a robust characterization of the Si3 noise, and direct beat with Si6 will result in a composite spectrum that is likely limited by Si6 at higher frequency, and hopefully limited by Si3 thermal noise at lower frequency. All noise measurements are made by beating light from the two cavity distributed onto a separate breadboard with fiber phase noise cancelled links. The Si6 laser is locked to a mode of the 6 cm cavity that minimizes the RF beat with Si3. For the TEM00 mode, this occurs at 600 MHz and is downconverted to 10 kHz where we use a 4th order active lowpass filter with 1 kHz passband to attenuate any high frequency tones. A time series of this 10 kHz beat is then logged on a $K + K$ FXE80 deadtime-free frequency counter.

Chapter 4

Technical noise sources

The cavity is designed with the Brownian thermal noise in mind as the ultimate limit to the laser stability. Numerous technical noise sources can prevent this realization, and each presents its particular challenges that require particular methods for mitigation. We'll now discuss the noise sources present in Si6 and detail our approaches to managing each of them.

4.1 Laser stabilization

The first question is whether, for high-stability experiments, one must use a laser with special properties, for example a massive invar construction. The answer both from theory and experiment is no, absolutely no. - Jan L. Hall

The main reason for choosing silicon as a spacer material is attributed to its unique physical properties, but another part of the choice was predicated on the wide availability of high quality optics at a telecom band. Narrow linewidth fiber lasers are readily available allowing for direct locking to the cavity. Not only does this preclude the need for prestabilization of a noisier free running laser, but the output mode is cleanly TEM00 so no complex beam shaping is required.

A number of compact fiber components are available at 1550 - the entire setup could in principle be made in fiber were it not for the need to couple into the cavity with free space mirrors. Nevertheless, we incorporate as much of the optical setup in polarization maintaining (PM) fiber as possible. Light from the NKT passes through a fiber isolator to prevent backscattered light from seeding the laser, and then continues to a high bandwidth 40 MHz fiber acousto-optic modulator

(AOM1). A 90/10 fiber splitter distributes the majority of the light through a fiber phase noise cancelled link to form a heterodyne beatnote with another ultrastable laser in the lab. The remaining 10 percent of light bypasses this setup and passes through an EOspace fiber electro-optic modulator (EOM) with 1 GHz bandwidth and half wave voltage of $V_\pi = 5$ V. A non-polarizing isolator was installed between this EOM and the free-space collimator after identifying a large etalon between a free-space optic and an in-fiber surface.

The light is then outcoupled into the free space path where two mirrors are used to align the beam to the cavity mode. An 8.1mm Schafer-Kirchhoff collimator and 750 mm lens are used to mode match into the cavity given the length of the free space path. A 45 degree mirror launches the beam vertically into the cavity and a Faraday rotator and PBS distribute the reflected light to a detector for PDH feedback. A half waveplate after the coupling mirrors allows for alignment to either of the two polarization modes, given the large birefringent splitting of the AlGaAs coatings. Between the two steering mirrors a plate splitter picks off some light for RAM detection. All mirrors and lenses are superpolished to $\lambda/10$ roughness and cube beamsplitters are two optically contacted prisms.

While the Si4 system used a RIO external cavity laser with ~ 1 kHz linewidth, the Si6 system replaced this with a significantly lower noise NKT distributed feedback fiber laser. The frequency noise of the two lasers is compared in Fig.4.1. The significantly lower linewidth relaxes the bandwidth constraints on the servo loop locking the laser to the cavity. Rather than relying on the inherently large bandwidth offered with current modulation, the Si6 servo loop actuates only on an AOM at higher frequency (see Fig.4.2 for AOM transfer function and Fig.4.3 for the modulator's AM-FM cross coupling), and a piezo within the DFB package at low frequency. Conveniently, the NKT piezo has a dynamic range of a full free spectral range of Si6 so a TEM00 mode can always be found within its tuning range.

The JILA school of laser locking involves throwing enough loop gain at the laser to suppress its frequency noise to a sufficient level. Loop bandwidth often pushes the edges of what is allowable with actuator or electronics delay, with integration stages tuned to aggressively high corners. This




Figure 4.1: (Left) A comparison of frequency noise between the RIO and NKT obtained via heterodyne beat with cavity stabilized light. (Right) The closed loop gain necessary to push the free running laser noise to the Si6 thermal noise at a given offset frequency.




Figure 4.2: AOM1 is used for both intensity and frequency feedback but it does not exactly actuate along the two orthogonalized quadratures. A transfer function of the cross coupling between frequency to optical power incident on the cavity is shown in (a). Intensity and frequency servo loops are also shaped to suppress any cross couplings. Feedback to the AOM frequency is highpassed with a 50 Hz corner while the intensity servo has an integrator enabled at a 100 Hz corner. The measured cross coupling of the AOM between frequency and intensity is small enough that the control signal from each cannot introduce noise in the opposite quadrature above the Si6 thermal noise.

is a nice approach that is generalizable to most problems, but Eric Oelker introduced some finesse to this problem by introducing us to some of the methods used in the LIGO community. The narrow linewidth of the NKT laser allows us close the loop with a respectable 150 kHz bandwidth, but the second integration stage is only enabled at the kHz decade. The loop topology of our PDH servo is illustrated in Fig.4.4, it consists of a common gain stage with integrators on the two actuator paths, and then splits into separate loop filters for the piezo and AOM. Canonically, the




Figure 4.3: Voltage-to-frequency transfer function of the VCO to AOM on the Si3 and Si6 systems. The delay of the Si6 modulator is slightly lower in frequency and consequently the maximum attainable loop bandwidth is lower.




Figure 4.4: Topology of the Si6 PDH loop. G_C , G_{PZT} , and G_{AOM} denote the loop filters for the common, piezo, and AOM stages. Responses of the actuators are denoted by A_{PZT} and A_{AOM} . Plant responses of the cavity filter function and the limit imposed by the photodetector bandwidth are given by P_{cav} and P_{PD} .

piezo provides low frequency feedback with the gain rolled off well below its mechanical resonance of 30 kHz while the VCO addressing the AOM is highpassed with a 50 Hz corner. Common stage integrators both have corners at 10 kHz below; laser noise is suppressed well below thermal noise (Fig.4.5), but this is largely due to the fact that the native laser noise is already quite low. Such a scheme would not work for locking a free running ECDL directly to a narrow linewidth cavity, even with access to high bandwidth current modulation on the diode. Additionally, this works nicely on the Si6 system but the limited feedback at high frequency can occasionally make it difficult to work on the optical table with the laser locked to cavity (high frequency mechanical perturbations occur in a region of low loop gain).




Figure 4.5: (a) The input amplification stage of the PDH servo board uses a low-noise FET with only 40dB of gain at low frequency. This pushes the input referred noise of the board to the $10 \text{ nV}/\sqrt{\text{Hz}}$ regime in the low frequency regime. Combined with a large discrimination slope achieved with the typical operation of a few microwatts incident on the cavity (b), the servo noise can be pushed to a frequency noise limit that is insignificant compared to the expected Si6 thermal noise. A large input referred noise can always be alleviated by using a corresponding steeper discrimination, but comes at the cost of increased sensitivity to photothermal effects.




Figure 4.6: (a) The closed loop actuator crossover illustrates the relative loop shaping between the two actuators. The phase measurement is informative in identifying the first, and higher order mechanical resonances of the piezo.

Nonetheless, it is easier to focus on low noise performance if the bandwidth of the servo does not need to be high. At the input of the servo, noise from the initial amplification stage is indistinguishable from laser noise, so the board noise needs to be sufficiently low as to not write this back onto the laser at a significant level. This input referred noise (Fig.4.5 (left)) is easily designed to be low since we don't require much amplification in the initial stages of the servo. Generally the loop should be able to suppress noise arising from any later stages or actuators. One useful diagnostic tool was to terminate each amplification stage with a $10 \text{ k}\Omega$ resistor and measure the




Figure 4.7: Summing an excitation into the closed loop yields the complete transfer function with all plant responses included. In the phase domain, this is particularly useful for identifying plant response poles in the high frequency. Modeling the closed loop phase reveals an additional measured pole at 50 kHz which is identified as the PDH photodetector response. This can then be appropriately compensated by adding a zero in the loop. The 2π phase wrap from 10 - 2000 Hz identifies the four integrator stages implemented with corners throughout that band.

coherence between the input and output referred noise. No coherence indicates that the amplifier's noise surpasses that of the resistors Johnson noise.

Deliberately driving excitations in particular locations of the board can give access to particular properties of the loop. For example, consider an excitation e_{PZT} summed in at the beginning of the subloop actuating solely on the piezo. Measuring the transfer function Z_4/Z_3 across the input and output of the excitation sum point yields the closed loop transfer function between the

two actuators.

$$Z_3 = G_c X - P_{PD} P_{Cav} G_C Y$$

$$Z_4 = Z_3 + e_{PZT}$$

$$\begin{aligned} Y &= Z_4 G_{PZT} P_{PZT} + Z_3 G_{AOM} A_{AOM} \\ &= (Z_3 + e_{PZT}) G_{PZT} A_{PZT} + Z_3 G_{AOM} A_{AOM} \\ &= Z_3 (G_{PZT} A_{PZT} + G_{AOM} A_{AOM}) + e_{PZT} G_{PZT} A_{PZT} \end{aligned}$$

$$\begin{aligned} Z_3 &= G_c X - Z_3 (P_{PD} P_{Cav} G_C) (G_{PZT} A_{PZT} + G_{AOM} A_{AOM}) \\ &\quad + e_{PZT} (P_{PD} P_{Cav} G_C) G_{PZT} A_{PZT} \\ &\equiv G_c X - Z_3 \beta (G_{PZT} A_{PZT} + G_{AOM} A_{AOM}) + e_{PZT} \beta G_{PZT} A_{PZT} \end{aligned}$$

$$\frac{Z_4}{Z_3} = 1 + \frac{e_{PZT}}{Z_3} = \frac{G_c X + e_{PZT} G_{AOM} A_{AOM} \beta}{G_c X - e_{PZT} G_{PZT} A_{PZT} \beta}$$

For large gain $e_{PZT} G_{PZT} A_{PZT} \beta \gg G_c X$ (where we let $\beta = P_{PD} P_{Cav} G_C$), so to first order we find

$$\frac{Z_4}{Z_3} = -\frac{G_{AOM} A_{AOM}}{G_{PZT} A_{PZT}} \quad (4.1)$$

Note the inclusion of A_{AOM} and A_{PZT} , indicating that this measurement includes the actuator response. Without any knowledge of the actuator gains, this allows an empirical method of loop tuning. The measured loop crossover on Si6 is shown in Fig.4.6.

Similarly, measuring the transfer function across an excitation summed into the servo input returns the entire closed loop transfer function with all plant responses included. Suppose we sum an excitation e_C and measure the transfer function Z_1/Z_2 . All servo stages and actuator responses can be lumped together as G_T for brevity. Then:

$$Z_2 = Z_1 + e_C$$

$$Z_1 = X - P_{PD}P_{cav}Y = X - e_C P_{PD}P_{cav}G_T - Z_1 P_{PD}P_{cav}G_T$$

$$\begin{aligned} \frac{Z_1}{Z_2} &= \frac{X - P_{PD}P_{cav}G_T}{X + 1} \\ &\approx -P_{PD}P_{cav}G_T \end{aligned}$$

Measurement of this transfer function yields information about all plant responses, which can be particularly useful if some are unknown. A measurement of the closed loop transfer function of the Si6 PDH servo loop is shown in Fig.4.7.

4.1.0.1 Shot noise limit

In the absence of technical noise, infinite gain at DC does not imply that the laser linewidth will be arbitrarily small. There are a number of mechanisms where the servo writes noise onto the laser, e.g. a noisy photodetector will set a limit on frequency noise measurement at some level. Generally the solution is simply to increase the optical power collected from the cavity since each additional photon carries more length information. As such, it is instructive to consider how much information a single photon will carry, especially since a number of measurements made later on had to take this into consideration.

For light incident on the cavity with modulation sidebands of power P_s each and carrier P_c , the photon shotnoise on the detector will be

$$S = \sqrt{h\nu(2P_s + \epsilon P_c)B}$$

In a bandwidth B with cavity coupling $1 - \epsilon$ and optical frequency ν . From earlier, we see that an approximation of the discriminator is given by

$$D = \frac{8\sqrt{P_s P_c}}{\nu_c}(1 - \epsilon)$$

In the limit $2P_s \gg \epsilon P_c$, fractional length resolution per bandwidth is

$$\frac{S}{DB} = \frac{\sqrt{h\nu_c}}{\sqrt{32\nu P_c}} \quad (4.2)$$

In a 1 Hz bandwidth with a maximized discriminator (no loss due to coupling), collecting 1 nW of power in the carrier over a 1 Hz bandwidth resolves a fractional frequency of 9×10^{-17} . A single photon over the same bandwidth is enough to resolve the cavity linewidth to 1 part in 5. The noise equivalent power of most photodetectors often necessitates significantly more optical power for the shotnoise limit to surpass the technical floor of the detector, but this is all to say that there is an incredible amount of information contained in the collection of each photon. Given the few μW 's of power typically used in these systems, shot noise limited resolution is very difficult to attain and having this as a limitation is typically an indicator of a separate problem. White frequency noise that is identified in the cavity noise spectrum is likely to be technical in origin.

An alternative method for determining requisite optical power is to push electronic noise to the thermal noise level. If some voltage noise V_0 is measured in the detection chain, the discriminator slope is chosen to be large enough such that this becomes insignificant.

$$\frac{V_0}{D_V} = \frac{V_0}{\mathcal{R}GD} < S_{CTN} \quad (4.3)$$

For detector responsivity $\mathcal{R}$ and transimpedance gain G .

4.2 Residual amplitude modulation

The cavity can be designed to support a pristine optical phase, but transferring its stability to the laser is often a significant hurdle for state of the art systems. Residual amplitude modulation (RAM) encompasses a variety of parasitic effects that prevent the experimenter from correctly tracking the phase of the cavity's optical field. In short, all of these effects introduce amplitude modulation at the frequency band where the cavity information is contained. We've previously seen that a strength of the PDH technique is to encode the cavity information high in the RF where technical noise sources are less dominant, thus the residual AM at the modulation frequency




Figure 4.8: Calculation of the fundamental shot noise limited length resolution on Si6. We account for the imperfect cavity coupling. Dashed lines indicated the thermal noise limit of different cavities.

is generally independent of the laser's spectral characteristics. Unfortunately this same method of shifting the PDH signal to a heterodyne beat opens one up to issues where the light used as a local oscillator does not provide a clean reference signal. Dispersive effects between the carrier and frequency offset sidebands generate unwanted amplitude modulation (the RAM signal) which on cavity resonance, are indistinguishable from the cavity error signal. RAM is the all encompassing term that captures the issues of having an imperfect phase reference. We'll now discuss mitigation of these effects.

It's important to note that there are approaches other than PDH that have been developed to eliminate some sources of RAM that are common in PDH. A crucial component of the PDH method is that the modulation sidebands are antiphase, allowing the baseline off-resonant signal to in principle be centered at zero. Alternative methods of tilt-locking [105] and squash locking [106] have explored using a phase reference that is spatially distinct from the resonant cavity mode, rather than spectrally distinct as found in PDH. In these schemes, the phase reference is a higher order nonresonant HG mode, where lobes of the spatial mode contain relative phase differences of π . Interference between the HG_{00} mode and the HG_{01} mode is detected on a quadrant photodetector and the anti-phase lobes of the phase reference reproduces the zero baseline signal when off resonant from the cavity. The tradeoff here is that the RAM and cavity signals are encoded at baseband

and both are now highly sensitive to optical scatter and alignment drifts. Shot noise limited RAM detection in a homodyne configuration is very difficult, so we opt for the traditional PDH detection to lock to the cavity.

4.2.1 RAM from the modulator

The phase modulation device which generates the sidebands is often a major source of RAM. There is an irony that the modulator allowing us to enjoy heterodyne detection is itself a source of parasitic effects, but there are now several developed techniques for dealing with this. A brief overview of various electro-optic modulation devices can be found in [107] and they can typically be grouped into free space and fiber based components. A LiNbO₃ optical waveguide modulator is used for Si6, where we make use of the low V_π drive voltage but do not benefit from the spatially separated polarization components particular to tetragonal $\bar{4}2m$ crystals like dihydrogen phosphates.

In trigonal crystals of the $3m$ class (of which LiNbO₃ is included), the refractive indices n_o, n_e for the extraordinary and ordinary waves differ due to different electro-optic coefficients r_{13}, r_{33} for the two polarizations [65]. Taylor expansion for the field dependence of n_o, n_e returns

$$n_o(E) = n_o - \frac{1}{2}n_o^3 r_{13} E \quad (4.4)$$

$$n_e(E) = n_e - \frac{1}{2}n_e^3 r_{33} E \quad (4.5)$$

We can see that the refractive index will have a DC value, and an AC term oscillating at the applied electric field. For LiNbO₃, one finds $r_{13} = 8.6$ ppm/V and $r_{33} = 30.8$ ppm/V [108] so there will be significant differences between $n_o(E)$ and $n_e(E)$ in the AC term. The corresponding phase shift for the two axes differs by a DC phase shift $\phi_{o,e}$ arising from the natural birefringence of the crystal, and an AC phase shift with the modulation index $\delta_{o,e}$ for the two polarizations. This is

expressed as

$$\Phi_o = \phi_o + \delta_o \sin \Omega t \quad (4.6)$$

$$\Phi_e = \phi_e + \delta_e \sin \Omega t \quad (4.7)$$

In the rotating frame of the laser, the electric field at the crystal output will be

$$E_{out} \propto c_o \exp[i(\phi_o + \delta_o \sin \Omega t)] + c_e \exp[i(\phi_e + \delta_e \sin \Omega t)] \quad (4.8)$$

Keeping terms first order in modulation frequency Ω and using $\Delta\phi = \phi_e - \phi_o$, the detected optical power will be

$$P_{out} \propto c_e^2 + c_o^2 + 2c_o c_e \sum_{k=-\infty}^{\infty} J_k(\delta_e - \delta_o) \cos(\Delta\phi + k\Omega t) \quad (4.9)$$

This can be separated into three signals. For $k = 0$:

$$P_0 = c_e^2 + c_o^2 + 2c_o c_e J_0(\delta_e - \delta_o) \cos(\Delta\phi) \quad (4.10)$$

For k of even order, optical power at $k\Omega$ is :

$$P_{k-even} = 2c_o c_e J_k(\delta_e - \delta_o) [\cos(\Delta\phi + k\Omega t) + \cos(\Delta\phi - k\Omega t)] \quad (4.11)$$

$$= 4c_o c_e J_k(\delta_e - \delta_o) \cos \Delta\phi \cos k\Omega t \quad (4.12)$$

Similarly for k of odd order:

$$P_{k-odd} = 4c_o c_e J_k(\delta_e - \delta_o) \sin \Delta\phi \sin k\Omega t \quad (4.13)$$

Phase shifts $\phi_{o,e}$ can be time varying, as well as prefactors $c_{o,e}$ which are identified in [109] to be alignment factors between the crystal and an output polarizer. Changes in the ambient environment can easily drive rotations of these axes, thus the power oscillating at DC generally does not share the same noise spectrum as power at the AC components (despite the shared terms). Though active and passive stabilization of the crystal can help decouple it from the environment, it is typically not enough to enable sub-ppm control of P_k that is required for the current generation of ultrastable lasers.

However, power oscillating at the modulation frequency Ω contains a term sinusoidal in the differential phase shift $\delta\phi$. For $\delta\phi = 0$, the RAM signal is effectively suppressed, and this condition can be actively forced by feeding a DC bias voltage to the EOM. The net effect is to introduce another DC phase shift ϕ_{DC} that modifies $\Delta\phi \rightarrow \Delta\phi + \phi_{DC}$. Power at P_{k-odd} is monitored and the DC bias is chosen such that $\phi_{DC} = -\Delta\phi$ and

$$P_{k-odd} \propto \sin(\Delta\phi + \phi_{DC}) = 0 \quad (4.14)$$

In this manner, amplitude modulation due to the crystal birefringence can be actively suppressed below the ppm level. Typically this is enacted by sampling part of the optical field on a photodetector before the light sees the cavity resonance. As long as there are no polarization sensitive elements between the detector pickoff and cavity, differential RAM due to crystal birefringence should be insignificant.

This is a powerful technique for actively suppressing amplitude modulation present around an offset frequency of Ω [110], but only occurs up to the point of detection before the beam interacts with the cavity. Any signal after cavity reflection is convoluted by its transfer function $R(\omega)$. The weakness of this technique is that feedback cannot easily occur at the same point where the PDH signal is detected¹.

The location of RAM detection is critical since there is invariably a differential signal that will be measured at the PDH detector. Fig.4.9 illustrates the RAM contributions that can be accumulated after each surface. The signals measured on each detector are

$$\phi_{\text{RAM}} = \phi_1 + \phi_2 + \phi_3 + \phi_4 + \phi_5 + \phi_6 + \phi_7 + \phi_8 + \phi_9 \quad (4.15)$$

$$\phi_{\text{PDH}} = \phi_{\text{RAM}} - (\phi_6 + \phi_7 + \phi_8 + \phi_9) + \phi_{10} + \phi_{11} + 2\phi_{12} + 2\phi_{13} + 2\phi_{14} \quad (4.16)$$

$$+ \phi_{15} + \phi_{16} + \phi_{17} + \phi_{18} \quad (4.17)$$

When the RAM servo is engaged, noise canceled up to its photodetector is removed, so the remaining

¹ For highly birefringent resonators where the mode splitting is resolved outside the linewidth, the RAM signal can be measured after cavity reflection the PDH detector by encoding it in a nonresonant orthogonal polarization. This possibility is an attractive feature for cavities with crystalline coatings, though we will later discuss why this is not implemented in Si6




Figure 4.9: Diagram of the differential optical phase that can be acquired between the carrier and modulation sidebands after their generation. Red phases are detected in the RAM and green phases are differential to the PDH path. RAM and PDH feedback to the phase errors measured on each photodetector are applied to the EOM and AOM locations respectively.

phase error accumulated at the PDH detector location is

$$\phi_{\text{PDH}} = -(\phi_6 + \phi_7 + \phi_8 + \phi_9) + \phi_{10} + \phi_{11} + 2\phi_{12} + 2\phi_{13} + 2\phi_{14} \quad (4.18)$$

$$+ \phi_{15} + \phi_{16} + \phi_{17} + \phi_{18} \quad (4.19)$$

RAM feedback can even have deleterious effects if noise present in the RAM path is not shared in the PDH path.

4.2.2 AM and PM on the phase reference

We haven't specified where these additional phases in ϕ_{PDH} can manifest as. We can generally approximate any small amplitude modulation α or phase modulation ϕ present on the sidebands on the incident electric field as

$$E_i = E_0 \left(J_0(\beta) e^{i\omega t} + \alpha_+ J_1(\beta) e^{i(\omega+\Omega)t} e^{i\phi_+} - \alpha_- J_1(\beta) e^{i(\omega-\Omega)t} e^{i\phi_-} \right) \quad (4.20)$$

To understand how this affects the detected signal, we don't need to involve the cavity. Instead we can treat the laser as being detuned from the cavity resonance so that we only measure a reflection off the input mirror. The detected signal oscillating at Ω is extracted as

$$P_\Omega = E_0^2 J_0(\beta) J_1(\beta) \left(\alpha_+ e^{-i\Omega t} e^{-i\phi_+} - \alpha_- e^{i\Omega t} e^{-i\phi_-} + \alpha_+ e^{i\Omega t} e^{i\phi_+} - \alpha_- e^{-i\Omega t} e^{i\phi_-} \right) \quad (4.21)$$

$$\propto \text{Re} \left[\alpha_+ e^{-i\phi_+} - \alpha_- e^{i\phi_-} \right] \cos(\Omega t) + \text{Im} \left[\alpha_+ e^{-i\phi_+} - \alpha_- e^{i\phi_-} \right] \sin(\Omega t) \quad (4.22)$$

Demodulating to baseband, we find

$$V \propto \text{Re} \left[\alpha_+ e^{-i\phi_+} - \alpha_- e^{i\phi_-} \right] \cos(\varphi) + \text{Im} \left[\alpha_+ e^{-i\phi_+} - \alpha_- e^{i\phi_-} \right] \sin(\varphi) \quad (4.23)$$

The phase φ is between the local oscillator and PDH signal. We can consider the limiting cases to gain some intuition into how AM and PM affect the RAM signal. When $\alpha_+ = \alpha_-$, the AM is common to both sidebands and only has a second order contribution to the RAM signal by varying the discrimination slope. An imbalance of the form $\alpha_+ \neq \alpha_-$ results in the traditional feature where the baseline changes to a nonzero value. In the case of PM, $\phi_+ = -\phi_-$ pushes the electric field phasors away from each other, but the net vector is only perturbed in magnitude to second order. Conversely $\phi_+ = \phi_-$ moves both field vectors together, resulting in a rotation of the total electric field E_i . In-phase PM can therefore have a tremendous effect on the RAM signal. Notably, active compensation of PM cannot be achieved by deliberately applying AM on the sidebands. The two quadratures are distinct, and the power of the EOM DC control is that a single actuator can be used to squash RAM originating from both quadratures.

Once the cavity is introduced, AM and PM have distinctly different ways of modifying the PDH signal. Fig.4.10 illustrates some examples of how unwanted modulation on the sidebands can corrupt the sensitivity to the cavity error.

In reality, it is rare that parasitic modulation on the PDH sidebands can have such a correlated effect. It is more typical that we see $|\phi_+| \neq |\phi_-|$ and $|\alpha_+| \neq |\alpha_-|$. Optical scatter and numerous




Figure 4.10: **Top:** PDH error signals with $\Omega = 5$ MHz, PM of $\phi_+ = \phi_i = \pi/6$, and AM of $\alpha_+ = 1/\alpha_- = 1.5$. The **Left** column displays demodulation phase of to extract the quadrature component, and $\varphi = \pi/3$ is on the **Right**. Curves are offset horizontally for clarity but not vertically. We can see that PM results in a large offset of the baseline signal that is not removed at $\varphi = 0$. For AM this is not the case though nonzero demodulation phase will reveal a large offset. **Bottom:** We show the magnitude of the PDH slope in all three cases. Dashed black lines indicate an offset frequency of half the cavity linewidth. At $\varphi = 0$ PM has the surprising feature that maximum sensitivity to the error signal does not occur at an offset of 0 Hz. Additionally the minimum sensitivity is shifted from the expected location at half the cavity linewidth.

other effects can generate these imbalances, but a highly common source are parasitic etalons within the optical path. Issues relating to this have been thoroughly studied in [111], so we'll just provide a brief overview here.

4.2.3 Etalon effects

An etalon is an unintentional cavity supported by two reflecting surfaces, with finesse typically ranging from 0.1 to 10. The length between these surfaces is never controlled as carefully as that of the silicon cavity, and as it drifts in frequency, each of the tones in the PDH triplet will experience a time varying phase offset. For an etalon centered at frequency δ , the field transmitted will have the form

$$\begin{aligned} E_t = E_0 & \left(J_0(\beta) e^{i\omega t} (1 - R_e(\omega - \delta)) \right. \\ & + J_1(\beta) e^{i(\omega+\Omega)t} (1 - R_e(\omega + \Omega - \delta)) \\ & \left. - J_1(\beta) e^{i(\omega-\Omega)t} (1 - R_e(\omega - \Omega - \delta)) \right) \end{aligned}$$

This is identical to the field transmitted through Si6, except we now use the etalon's reflectance function $R_e(\omega - \delta)$ rather than the cavity's $R(\omega)$. Instead of locking a laser to one cavity, we will effectively be stabilizing to the composite system of Si6 and the etalon. As this is a parasitic effect, care is taken to avoid having parallel high reflectivity surfaces. Consequently, its finesse is typically poor, meaning the linewidth will be significantly larger than that of the cavity. This ends up being a significant issue - since its length is not systematically controlled, drifts in the etalon length (even far from its resonance peak) will cause a continuous differential offset between the PDH sidebands. As discussed earlier, this will directly manifest as a RAM effect, degrading the cavity stability. Fortunately this effect can be easily identified by observing a small modulation of the error signal, evident in Fig.4.11. Observing a full modulation of the etalon fringe allows us to extract its length, and correspondingly identify the optics generating the reflections.

Given the low thermal noise limit in Si6, we are wary of all these effects. Fiber components are passively isolated from temperature fluctuations and we actively feedback to our EOM voltage to stabilize the detected RAM. Despite active stabilization of the EOM temperature, we observe large levels of free running RAM below 1 Hz when the DC voltage servo is not engaged. The fractional cavity instability climbs to the 10^{-15} level without any RAM compensation, the servo is therefore activated with a ~ 300 Hz bandwidth to ensure that there is sufficient gain near DC.




Figure 4.11: Simulating the effect of a 1 MHz linewidth etalon with 20% coupling stationed 3 MHz away from the cavity resonance. A significant distortion of the error signal occurs, and the etalon will always decrease the effective discrimination slope. Coupling to the etalon will typically be lower than the parts per thousand level if all optics have high quality AR coatings. The parameters are purposely exaggerated here to demonstrate the etalon effect.

The vast majority of this appears to originate after the EOM - detecting the RAM right after the light is outcoupled from a fiber indicated similar levels of noise. In the free space path, all mirrors and lenses are superpolished with a 20-10 scratch-dig to inhibit unwanted optical scatter. We take care to tilt all optical surfaces so that no etalons can be supported. Despite this, we still identified a large etalon between one of the optical surfaces within or on the cryostation and an in-fiber component. Details of this characterization can be found in the next chapter. With the servo engaged, we measure the residual RAM effect on the PDH detector in Fig.4.12.

The RAM measured at the PDH detector is consistent with the Si3 level, far from the projected thermal noise of Si6. Though the noise is controlled at the 1 ppm level, we're impaired by the relatively large linewidth of Si6. The measured finesse is nearly 60% larger in Si4, meaning that an equivalent level of RAM relative to the cavity linewidth would reach the $2\text{-}3 \times 10^{-17}$ range.

For a 21 cm cavity with 500,000 finesse, our measured out of loop RAM would be well below the Si6 thermal noise.

We don't have a solid idea of where the residual signal comes from. RAM control at the low 1×10^{-7} has been demonstrated [112], but this is successful when the majority of RAM originates in components that can be easily altered or moved. At some point we placed a mirror before the input vacuum optic and reflect the beam off its surface and the measured RAM was lower, suggesting that the dominant contribution arises from an in-vacuum surface. In the next chapter, we note that changing the input beam pointing to couple to the cavity can lead to a substantial improvement in frequency stability. We suspect that one of the vacuum windows or the input cavity substrate is dirty or scatters light. The prospect of locking to the cavity from the top is very appealing, as we may be able to avoid these issues. There are however other technical challenges associated with that scheme.




Figure 4.12: **Top:** Fractional frequency noise of the out of loop RAM characterization. Though far from the projected Si6 thermal noise, the cavity instability should be limited to a level similar to the Si3 thermal noise. This would mark a significant improvement over the instability of the Si4 system with the added benefit that the 4 K closed cycle cryostat makes Si6 a hands-free oscillator. Dark noise corresponds to noise of the resonant photodetector. A high white noise level is visible at frequencies larger than 1 Hz. This specific dataset coupled to the TEM01 mode of the cavity without a phase plate. The poor mode coupling leads to a correspondingly higher contribution of shot noise. **Right:** Modified Allan deviation of traces on the left.

4.3 Vibration sensitivity

The large stiffness of silicon as a cavity material has already become apparent due to its role in reducing the thermal noise contribution of the spacer. Our intuition supports the concept that this should be beneficial when the cavity is in the presence of any compressive force. While the cavity is a rigid body, its length will still be modified as

$$\frac{\Delta L}{L} \propto \frac{F}{EA} \quad (4.24)$$

for a force F and cross sectional area A . This simple model treats the cavity as a bulk elastic spring, with spring constant determined by the material stiffness E . With a stiffer spacer, an applied force induces a smaller fractional length change. This applies to forces along all three axes of the cavity, though the length changes due to horizontal forces are modified by the Poisson's ratio σ which relates the applied stress to an orthogonal strain. Generally the sensitivity to vibrations can vary between the vertical and two horizontal axes. Unlike thermal noise, this does not impose any fundamental limit but rather eases our task of seismic isolation and reducing those compressive forces.

We opted to cut the silicon spacer along the stiffer $\langle 111 \rangle$ axis, though not for the purpose of simply reducing the bulk spring constant or reducing spacer thermal noise. Silicon has a diamond-like cubic lattice structure with atoms arranged tetrahedrally. There is a three-fold rotational symmetry of the stiffness matrix C when oriented along $\langle 111 \rangle$ [113]

$$C_{\langle 111 \rangle} = \begin{bmatrix} 194.4 & 54.3 & 44.8 & 0 & -13.5 & 0 \\ 54.3 & 194.4 & 44.8 & 0 & 13.5 & 0 \\ 44.8 & 44.8 & 204.0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 60.5 & 0 & 13.5 \\ -13.5 & 13.5 & 0 & 0 & 60.5 & 0 \\ 0 & 0 & 0 & 13.5 & 0 & 70.0 \end{bmatrix} \times 10^9 \text{Pa} \quad (4.25)$$

When the cavity is suspended or supported with similar rotational symmetry, the anisotropic structure of the crystal means that the vertical sensitivity will vary sinusoidally in magnitude with

a 120° period. By varying the angle between the cavity's horizontal axes and the three points of support, we can tune the magnitude of the vertical sensitivity making this an effective knob for reducing the cavity's frequency coupling to accelerations. This phenomena was previously measured in the first realizations of cryogenic silicon cavities [39] and it is now well established that rotating the it allows us to reduce the vertical sensitivity. The cavity vent holes are drilled along the $\langle 2 - 1 - 1 \rangle$, $\langle -1 - 12 \rangle$, and $\langle -12 - 1 \rangle$ crystal axes, so the approximate angle between the cavity mount and $\langle 2 - 1 - 1 \rangle$ can be easily determined when viewing the system from the top. Less consideration has been given as to whether this circumstance of reduced sensitivity can occur for the two horizontal directions. Our FEM simulations suggest that we should expect the two horizontal sensitivities to vary sinusoidally with 360° period. Our measurements have clearly indicated a slight rotational dependence of these couplings, but the magnitude appears to depend strongly on the geometry and material of the support structure.

This can leave the cavity to be more susceptible to horizontal perturbations. Simulating the mechanical eigenmodes of Si6 confirms this as the lowest resonances of the cavity are bending modes. The resonant frequency is still very high however, with the aforementioned resonances occurring near 27 kHz. It is more likely that weak low frequency resonances in the support structure could be a limitation, so a proper design of the mount is crucial. Ideally, mechanical resonances of the combined mount and cavity should still be high, > 100 Hz. The low form factor of Si6 and the 4 K cryostation is fortunate in this regard, since the total system mass cannot reach the levels of the 21 cm Si2 and Si3 systems².

With the constraint of needing to mount the cavity at three points, we orient it vertically underneath its supporting ring. Our intuition suggests that we should support it at the midplane, where vertical forces act oppositely on the two spacer halves and lead to a $\Delta L = 0$ change in the cavity length. This has traditionally been used in ULE cavities [114, 115], where vertically oriented football-shaped spacers are supported at the midplane at three points to reduce overconstraining

² Si2 and Si3 measured a low lying mechanical resonance at 40 Hz originating from a tilting eigenmode of their PEEK cavity mount




Figure 4.13: Vertical acceleration sensitivity of the cavity. Si6 is supported at the midplane and we measure a sensitivity that is roughly centered at zero. We rotate the cavity on its mount to be near one of these zero crossings.

and the vertical acceleration sensitivity. In the crystalline cavity, this has the additional advantage of manifesting as a zero crossing of the vertical sensitivity, meaning that the cavity length is not perturbed when subjected to acceleration along the spacer. As the cavity is supported farther away from the midplane, the sensitivity curve is displaced upwards or downwards in magnitude to the point where the desired zero crossing will not appear.

This occurred in the Si2 and Si3 systems where the supporting ring on the spacer was intentionally offset from the midplane for the purpose of placing the zero crossing at the shallowest sections of the sinusoidal sensitivity. Measurements of the sensitivity do not exactly reproduce simulated results and it was found that no orientation of the cavity resulted in a nulled vertical sensitivity. The mass imbalance between the top and bottom sections of the cavity relative to the mounting plane drive the finite acceleration coupling. A silicon ring placed around the top of the cavity spacer revived the zero crossing [39].

The crystalline structure of the spacer has thus been beneficial for factors unrelated to the low mechanical loss angle. Careful consideration given to the crystal orientation has allowed us to precisely control some components of the cavity's response to accelerations. A significant component however is also due to the geometry of the cavity mount. This has been extensively examined in ULE systems [107, 116, 117, 118, 119], but it has been less studied in crystalline systems where we

are constrained to the tripod. We will next discuss some of our related findings.

4.3.1 Measuring cavity sensitivity

To measure the cavity's response to vibrations, we want to provide an acceleration to the cavity and measure the frequency response. Along the two horizontal axes, we can easily perform a DC measurement of the sensitivity k by tilting the optical table along the two directions H1 and H2. This introduces an acceleration due to gravity as $g \sin \theta$ for tilt angle θ . The resulting frequency shift Δf returns the fractional sensitivity as

$$k = \frac{\Delta f}{f} \frac{1}{g \sin \theta} \quad (4.26)$$

There is no trivial analogue to performing a DC measurement of the sensitivity along the Z axis since tilting the table barely modifies the vertical projection of gravity $g \rightarrow g \cos \theta \approx g$.

An easy solution is to modulate a force along all three directions, and measure the coherent cavity frequency response to obtain k_{H1}, k_{H2}, k_V . Sweeping this modulation frequency, we can map out the spectrum of the cavity's AC acceleration sensitivity. The active vibration isolation platform supporting the cryostation has the option for modulating the table displacement along all three cartesian axes. Driving the AVI along a particular direction can drive accelerations along all three axes due to cross couplings of the platform.

Previously we've suppressed these complications by identifying a drive frequency where these cross couplings are minimized and measuring the cavity response. This has the benefit of being procedurally simple but does not provide access to the full frequency dependence of k_{H1}, k_{H2}, k_V . If we instead modulate at an arbitrary frequency regardless of cross couplings and measure the acceleration response a_{H1}, a_{H2}, a_V along all three directions (3×1 matrix $\mathbf{A}$, the frequency response will be

$$\frac{\Delta f_i}{f} = \sum_{j=H1,H2,V} a_i e^{i\phi_j} k_j \quad (4.27)$$

Each term has an associated phase that will be frequency dependent. We cannot identify any individual components with this single measurement, but if we drive the AVI along the three

orthogonal directions we can construct the 3×3 matrix of measured sensitivity and phase responses $\mathbf{k}$. The 3×1 diagonalized vector of acceleration sensitivity $\mathbf{k}_D$ is calculated as

$$\mathbf{k}_D = \mathbf{P}^{-1} \mathbf{A}^{-1} \begin{bmatrix} \Delta f_{H1}/f \\ \Delta f_{H2}/f \\ \Delta f_V/f \end{bmatrix} \mathbf{P} \quad (4.28)$$

By finding the invertible matrix $\mathbf{P}$.

Performing lock-in detection of the cavity with our modulation frequency allows us to easily make this measurement over a broad band of frequencies by simultaneously shifting the drive and demodulation. This requires an analog signal corresponding to the cavity frequency. We phase lock a VCO to the Si3-Si6 beatnote with high bandwidth such that the servo in-loop is well below the frequency response we aim to resolve. The control signal sent to the VCO is then used as our measure of the cavity noise. A standard set of data is illustrated in Fig.4.14 where we measure the phase and transfer function of our accelerometer signal to the VCO.

It is difficult to extract useful information on the magnitude of the acceleration sensitivity from this plot, but performing the diagonalization of Eq.4.28 returns the significantly more meaningful k_{H1}, k_{H2}, k_V (Fig.4.15). Notably, there is a strong frequency dependence of all three sensitivities. Their magnitudes are nearly ten times larger above 100 Hz, and even in the range of 10-60 Hz they can fluctuate by a factor of 3 or 4. We don't have a good intuition why this occurs. Near a mechanical resonance, we expect a peak in the sensitivity but this should monotonically decrease away from the resonant frequency with a width related to its quality factor. Montana instruments has shared that the 300 K to 30 K support structure is largely made from G10, which could result in a broader resonance. Nevertheless, we have confidence in the AC acceleration sensitivity data since in the low frequency limit our measurements of k_{H1} and k_{H2} approach the values measured with our DC tilt results (Tab.4.1). We naturally expect this to occur, so for the purposes of calculating the contributions of vibrations to frequency noise, we assume that the AC sensitivities measured at 10 Hz retain the same value down to DC.

These sensitivities can be multiplied by our measured vibration noise along each axis to




Figure 4.14: Transfer functions (**left**) and phase responses (**right**) of the three acceleration responses to frequency when driving the AVI along a particular axis. Different rows display the data of shaking the AVI along each direction. The acceleration response to driving the platform along gravity (bottom row) is largely contained along $\hat{V}$. This is less the case for driving along the horizontal axes where we can see a few cross couplings at low frequency. This likely suggests that the platform is torqued in the H1-H2 plane, shaking along the AVI's horizontal axes generates an admixture of excitation along H1 and H2. In the phase domain, we can identify some sort of mechanical resonance of the cryostation near 100 Hz, where phase wrappings are evident in all axes.

obtain the vibration induced frequency noise on the cavity. The cavity frequency will be unable to distinguish between noise from orthogonal components, so we calculate the total contribution to

	DC measurement	AC measurement
$k_V [g^{-1}]$	-	1.0×10^{-11}
$k_{H1} [g^{-1}]$	$0.8(3) \times 10^{-11}$	0.7×10^{-11}
$k_{H2} [g^{-1}]$	$1.2(6) \times 10^{-11}$	2.2×10^{-11}

Table 4.1: Comparison of low frequency AC sensitivities and DC sensitivities.




Figure 4.15: The frequency dependence of measured acceleration sensitivities on Si6. The vertical sensitivity is notably higher than the mid- 10^{-12} s value that can be attained by adjusting the spacer rotation. We observed that the values of the horizontal sensitivity had a slight dependence on rotation angle as well. The current orientation of the spacer was chosen to minimize the largest sensitivity (in this case k_{H1}) without significantly increasing the others.

its fractional noise as

$$S_y = (a_V k_V)^2 + (a_{H1} k_{H1})^2 + (a_{H2} k_{H2})^2 \quad (4.29)$$

Note that we are not considering the relative phases between these terms, so Eq.4.29 is only an upper limit on the true vibration-induced frequency noise. Our projection depicted in Fig.4.16 suggests that the noise contribution due to vibrations should not limit measurement of Brownian thermal noise at frequencies below 1 Hz.

4.3.2 Improving the sensitivity

The acceleration sensitivities shown in Tab.4.1 and Fig.4.14 were the result of many modifications to the Si6 cavity support, most through trial and error. Initial measurements of the vibration




Figure 4.16: Vibration induced frequency noise measured along the different axes and their respective couplings to the cavity.

coupling revealed sensitivities ten times worse, identical to the values measured in Si4. We'll review some of the more impactful upgrades here.

It's useful to introduce some historical context - in Si4 the cavity-support rotation angle was optimized to minimize the vertical sensitivity ($k_V = 5 \times 10^{-12} \text{ g}^{-1}$), but large horizontal sensitivities of $k_{H1} \sim k_{H2} = 2 \times 10^{-10} \text{ g}^{-1}$ were significant limitations. There was no clear indication why these values might be so poor. One of the cavity substrates was contacted off-center to the spacer borehole, and a horizontal displacement of the beam from the cavity axis could beget a $4 \times 10^{-11} \text{ /g/mm}$ sensitivity. Standard machining tolerances of $100 \text{ } \mu\text{m}$ could easily lead to a displaced substrate curvature, but we would expect a residual sensitivity of $4 \times 10^{-12} \text{ /g}$. The 50 times larger value measured on k_{H1} and k_{H2} would be unphysical. This claim also lost its attraction when we swapped the Si4 cavity for the Si6 cavity and remeasured the exact same sensitivity curves. It became clear that the cavity itself was not the primary encumbrance.

Retaining the same level of vibration-induced noise on Si6 would have prevented a thermal noise limited laser performance at 1 Hz, so it was imperative that we understand the source of this high value. A number of studies have investigated the design of the cavity spacer with the intent of reducing the sensitivity. With the Si6 spacer already machined and contacted to substrates, this is




Figure 4.17: Various cavity mounts that we tested for improved sensitivity to acceleration noise. **Top left:** The original folded cup design from Si4, made from G10 composite. **Top right:** A simplified structure made from G10. **Bottom left:** A monolithic stainless steel support machined a block of 316 stainless. Hemispherical divots have been milled out from the top of each post. Spheres of different material can be placed within each divot, the image shows a set of teflon balls resting within the indentations. **Bottom right:** Our final cavity support made entirely from stainless steel

not an option. Instead, we suspect there may be gains from modifying the support structure. The original Si4 cavity support was composed of three thin, nested concentric cups of G10 fiberglass glued together. Three tiny 1 mm tall flat rods equally spaced on the top of the outer cup support the cavity. The intent was to maximize the effective path for heat flow, lengthening the thermal time constant. Unfortunately, G10 has a relatively low bulk modulus of 18 GPa that is characteristic of most laminates. The low mechanical quality factor of the material combined with its three times

larger effective length transforms the cup into a lossy mount with a low spring constant.

Our first avenue for improvement was to increase the structural spring constant by substituting the nested layer design for a hollowed out half-cylinder. The cavity points of contact need to maintain a 120° symmetry for zeroing the vertical sensitivity, so three longer posts tapered to a point are mounted to the top of the cup. This is hard to visualize so a photo of the structure installed in the cryostation is depicted in the top right of Fig.4.17. Despite the green color, the support is still made of G10, but the improved design has pushed the mechanical resonant frequencies much higher. Remeasuring the sensitivity spectrum showed a dramatic improvement, with both $k_{H1}, k_{H2} < 1 \times 10^{-10}$ below 100 Hz.

Clearly the design of the support structure was of significant importance for the acceleration sensitivities. We were largely concerned with reducing the horizontal sensitivities since we know that we can tune k_V close to a minimum. Since increasing the mount's spring constant had a considerable impact on the values of k_{H1} and k_{H2} , it stands to reason that it could be helpful to stiffen the support even further. We decided to test out an analogous design made from stainless steel instead of G10. The larger Young's modulus of 200 GPa allows us some leeway with the geometrical details, we opted to increase the post length and minimize the size of the surrounding cup to increase the thermal time constant.

Si2 and Si3 have had success with resting the cavity on PEEK supports which is less structurally stiff than G10, so we were curious whether the material in contact with the silicon spacer had a significant effect on the measured sensitivity. In our first design of the stainless mount, we cut a hemispherical divot into the top of each post so that we could place small spheres of different material in here, rest the cavity on top, and measure its coupling to accelerations. In principle this should provide us with a methodical method for studying the importance of this contacting material. In reality, this process of exchanging the sphere material, replacing the cavity, and re-measuring the sensitivity is extremely tedious. We only ended up performing this test for teflon, nylon, and stainless steel spheres. We surprisingly observed a small dependence of k_{H1}, k_{H2} on the tip material with teflon balls providing the smallest sensitivity and nylon balls returning the

largest. This ran counter to our recently developed intuition that stiffer materials lead to a lower coupling.

The stainless steel component of this mount was monolithic and machined from a single block of SS316 with a computer controlled CNC. The cutting process lead to some warping of the stainless posts and we worried that this may have some adverse effect on the sensitivity. We opted to build a second stainless mount (posts, tips, and base all made from SS316) by rounding the tips of rolled stainless rods and pressing them into a machined baseplate (lower right of Fig.4.17). Stresses introduced in the cutting and assembly process are much lower, and undesired tilts of the stainless posts can be controlled with much better precision. Remeasuring the sensitivity with this new mount revealed that we could attain much lower values with $k_{H1}, k_{H2} \sim 1 \times 10^{-11}$. This is an over 20 times improvement in sensitivity over the $k_{H1}, k_{H2} \sim 1 \times 10^{-10}$ measured on Si4. The corresponding frequency noise contribution from vibrations will be over 20 dB lower than measured on the previous system. This is more than sufficient for the stricter requirements from the Si6 projected thermal noise so we opt to use this support for all measurements going forward (Fig.4.15 is the sensitivity spectrum measured after cooling the cavity to 4 K). All our intuition here is empirical but the takeaway seems to be that stiffer materials and a properly constructed mount are a path towards a lower vibrational sensitivity.

4.3.3 Feedback to AVI

Reducing the cavity sensitivities proved to be a fruitful method, but further reduction has proved to be difficult. There is no quantitative barrier preventing this, historically it is rare that cavities attain $k_i < 1 \times 10^{-11} \text{ g}^{-1}$. If the total vibrations are prohibitive for measuring thermal noise, trying to reduce the acceleration noise can be a better avenue.

The necessary approach depends on the exact spectrum of vibration noise. In our system, a_V, a_{H1}, a_{H2} are roughly the same magnitude, but k_{H2} is the largest sensitivity. The frequency noise contribution $a_{H2}k_{H2}$ is therefore larger than the other two terms, so reducing a_{H2} should lead to a corresponding decrease in the vibration-induced frequency noise. Reducing a_{H2} means




Figure 4.18: Measured sensitivities for the variety of cavity mounts we considered. Label "New G10 mount" refers to the green G10 cup from Fig.4.17, "Fully S316" refers to the stainless steel mount with rounded tips, and "Divot" is the stainless mount with a hemisphere milled out from the top of the supporting posts. We place stainless steel, nylon, and teflon spheres in the divots and rest the cavity on top. All traces have been smoothed with a 5th-order spline to improve visual clarity. Horizontal sensitivities are the more meaningful quantity since we know that k_V can be zeroed by adjusting the cavity-mount angle. Variation of k_V between the different mounts is due to imperfect placement of the cavity. All k_i steadily increase above 100 Hz, indicating that some other component of the cryostation is contributing to an increased acceleration sensitivity at high frequency. A large rollup in the k_{H1} is visible for all the different cavity supports. At the time of these measurements, we were suspected that this was not representative of the true H1 sensitivity but were unsure of its source. We later understood that driving the AVI along the H1 axis (collinear to the coldfinger axis) would excite a compressive mode along the coldhead's bellows. This generates air currents that modulate the laser phase coherently with the cavity's acceleration response and appear as an artificial k_{H1} signal. We later remove this effect by pulling a vacuum on the cryostation, and its absence is noted in Fig.4.15.

that we need to do a better job of noise feedback than the AVI. The current limitation is the noise limit of the in-board sensors, meaning that a lower noise accelerometer would enable improved feedback. Similar to how we would excite the AVI by driving it with a synthesizer, we can perform broadband noise cancellation by instead feeding it with a servo output's control signal.

In Fig.4.19, we place a high-sensitivity accelerometer on the cryostation split plate and use the measurement as an error signal to a servo. The servo output modulates displacement along the H1 axis, and a second accelerometer near the same location serves as an out of loop measure for the reduction in vibration noise. We measure a clear reduction in measured acceleration noise, though we again are limited by the noise floor of our improved detectors and receive a 3dB penalty in terms of unwanted noise added by the servo.

The principle of this method is simple, but in this test we neglected to account for the feedback range of the AVI. The platform's vibration isolation is highpassed with a 2 Hz corner as it is difficult to distinguish between tilt and displacement below this frequency. Extending the feedback to lower frequency would require a tiltmeter on the optics table to measure the slowly varying inclination of the cryostation.

This demonstration also assume that feedback to the H1 axis only couples to the vibration noise measured along H1. In reality, we previously identified cross couplings between drive and response along the cryostation axes. To truly minimize acceleration noise along H1, it is necessary to apply a control signal to H2 and V as well. The matrix delineating all couplings is our previously measured as $\mathbf{A}$. Feedforward on the vibration noise requires knowledge of both real and imaginary components of $\mathbf{A}$, as well as the frequency dependence of all k_i . A feedback scheme will be significantly simpler since correction of accelerations measured on all axes can be propagated through $\|\mathbf{A}\|$.

The frequency dependence of all nine terms is illustrated in Fig.4.20. Driving the AVI along a particular axis clearly couples to the other two, but generally with a small magnitude. If we're not trying to reduce the acceleration noise by a significant amount, it is actually appropriate to only perform single axis feedback, e.g. drive in H1 to reduce noise along H1, drive in H2 to reduce

noise along H2, etc. In H2 for example, cross couplings between H1 and V are 10 dB smaller than the H2-H2 coupling. If we only need a factor of two reduction in the amplitude spectrum of a_{H2} , this can be done without any propagation through **A**. Vibration noise on Si6 is low enough that we don't need to try any of these schemes, but future cavities with a projected thermal noise below the 10^{-17} level will likely have to consider this.




Figure 4.19: Vibrations measured along H1 with the AVI isolation enabled are shown in blue. When we use a lower noise external accelerometer to measure the acceleration spectrum, we can feedback to the AVI by modulating its displacement along H1 to further reduce the measured noise (black). At frequencies below 40 Hz, the black curve is limited by the noise floor of our new in-loop accelerometer.

4.4 Thermo-expansive length stabilization

With the finite $\sim 10^{-11} \text{ K}^{-1}$ CTE of silicon at 4 K, temperature fluctuations need to be limited to the μK level to enable the thermal noise limited performance of Si6 at $\sim 1 \times 10^{-17}$. Measuring temperature at this level is well below the precision of commercial thermometers. We instead need to be conscious of temperature couplings in system and try to strategically passively dampen them by careful design of our cryostation environment. At 4 K, our choice of materials for the cavity support will become important. Thermal isolation stages can lowpass temperature variations to the nK level, and this corner frequency will be determined by the thermal resistance and capacitance of the system. We will now discuss our approach to meeting these requirements,




Figure 4.20: Plots of the coherent response measured along each axis when driving the AVI along one direction. In the top plot for example, the AVI is driven along H1 and we show the corresponding transfer functions $a_i / \sqrt{\sum_i a_i^2}$. The vast majority of the coherent displacement occurs along H1 at low frequency, though in regions around 50 Hz and 80 Hz, an equal response is measured in the vertical axis. Single axis feedback along H1 would exhibit strong cross coupling to another axis at these frequencies.

and begin with a model of the thermal couplings within the Si6 system.

4.4.1 Heat flow

In our VHV vacuum, convective heat transfer through gas is minimized. Heat flow can occur conductively through the cavity support and through radiative exchange between surfaces at different temperature. We'll consider radiative exchange first.

4.4.1.1 Radiative exchange

Heat exchange between two blackbodies can be calculated via the Stefan-Boltzmann law

$$\dot{q} = \sigma_B A (T_o^4 - T_i^4) \quad (4.30)$$

For surface area A , Stefan-Boltzmann constant σ_B , and temperatures T_o, T_i of the two surfaces. The description is complicated by the fact that the two surfaces are often not perfect emitters, and are instead "grey bodies" with emissivity $0 < \epsilon < 1$. Furthermore, the two surfaces may differ in their total surface area, e.g. a large vacuum chamber around a smaller thermal shield. The "interaction" between these two bodies cannot be neglected, and it is necessary to calculate the form factor $F_{o \rightarrow i}$ that determines the fraction of energy radiation from surface o and received by surface i [120]. This will be unique to the geometry of the two surfaces. For concentric cylinders with an inner surface i and outer surface o , one finds

$$F_{o \rightarrow i} = A_o \left(\frac{1}{A_o \epsilon_i} + \frac{1}{A_i \epsilon_o} - \frac{1}{A_o} \right) \quad (4.31)$$

For this configuration it's clear that $F_{o \rightarrow i} \neq F_{i \rightarrow o}$ so that the corresponding radiation energies $Q_{o \rightarrow i} \neq Q_{i \rightarrow o}$. This is due to the ability of the larger concave surface o to radiate back onto itself. Thus the total shape factor $F_{o \rightarrow i} + F_{o \rightarrow o} + F_{i \rightarrow o} = 1$. The radiative exchange between grey bodies is then found as

$$\dot{q} = \sigma_B \left(\frac{1}{A_o \epsilon_o} + \frac{1}{A_i \epsilon_i} - \frac{1}{A_o} \right)^{-1} (T_o^4 - T_i^4) \quad (4.32)$$

Low emissivities and similar temperatures between the two surfaces will minimize the heat flow. The majority of radiative exchange in our systems occurs between thermal shields due to

their large surface areas. We use three thermal shields to step down the radiative coupling and maintaining the two active and passive shields at 4 K suppresses any heat radiative heat flow between them and the cavity. As we'll elaborate in a later section, we gold plate all shields to reduce the emissivity and lower the heat flow. The majority of heat exchange occurs between the room temperature vacuum chamber and the radiation shield due to the large (300 K)⁴ heat load. The stage 1 temperature of the cryocooler consequently needs to have a sufficiently large cooling capacity. Room temperature fluctuations can easily pull the temperature of the radiation shield, so we heat it to 42-45 K with a 50 Ω manganin heater wrapped around the base³. The diurnal lab temperature fluctuations then have a suppressed coupling to the 30 K stage.

The heater has enough capacity to maintain the shield temperature above its steady state value, but this doesn't relax our need to keep this stage cold. A hotter shield temperature of 60 K will have a five times larger radiative heat exchange with the 4 K active shield which can be large enough to overwhelm the cryo's stage 2 cooling capacity. This results in a higher base temperature of the the system, and the hotter cavity will have a larger CTE, making it more susceptible to temperature fluctuations. In the 21 cm system, multi-layer insulation foil is used between thermal shields to significantly lower heat flow by introducing intermediate surfaces $T_1, T_2, T_3, \dots$, each of which have their own inside and outside surface emissivities $\epsilon_i^o, \epsilon_i^i$. Radiative exchange is then reduced from $\sim (\epsilon_o^i T_o^4 - \epsilon_i^o T_i^4)$ to $(\epsilon_o^i T_o^4 - \epsilon_1^o T_1^4) + (\epsilon_1^i T_1^4 - \epsilon_2^o T_2^4) + (\epsilon_2^i T_2^4 - \epsilon_3^o T_3^4) + \dots$. Insufficient space between the thermal shields on Si6 prevents us from including this in our system, but this is a powerful method to consider. Proper choice of the $\epsilon_i^o, \epsilon_i^i$ can lead to a significantly lower radiative heat load.

Lowering the radiative exchange between stages is always beneficial, but it's useful to think about this in the context of the overall heat flow, both conductive and radiative. A frequent analogy is to adapt the language of electrical circuits. The thermal counterpart to Ohm's law is

$$\Delta T = QR \tag{4.33}$$

³ Despite the cryo stage 1 being at 30 K, the radiation shield attains a lowest temperature of 40 K

For heat density q . A thermal resistance R can be used to quantify an impedance to heat flow through the system. If there are two paths for heat to flow (conductive and radiative exchange), preference is give to whichever has the lower resistance. The effective resistance across both paths will be

$$R_{eff} = \left(\frac{1}{R_{radiative}} + \frac{1}{R_{conductive}} \right)^{-1} \quad (4.34)$$

So if $R_{radiative} \gg R_{conductive}$, there is little to be gained by further suppressing radiative exchange.

Radiative resistance between grey bodies is

$$R_{radiative} = \left[\left(\frac{1}{\epsilon_i A_i} + \frac{1}{\epsilon_o A_o} + \frac{1}{A_o} \right) \left(\sigma_B * (T_i^2 + T_o^2)(T_i + T_o) \right) \right]^{-1} \quad (4.35)$$

In Si6, radiative resistance between all surfaces is considerably higher than the resistance through the conductive path. To good approximation, heat flow can be understood by only considering the conductive exchange.




Figure 4.21: Schematic of the thermal circuit that models heat transfer between the 30 K stage and the cavity. Radiative and conductive pathways are placed in parallel, and the low pass filter is completed with the inclusion of materials with high thermal capacitance. An additional lowpass should be included to capture the heat from the 300 K vacuum chamber to the cryogenic 30 K stage, but details this support structure within the Montana Instruments cryostation are proprietary, and are not shared with us. We calculate an active-passive time constant of ~ 10000 s and a passive-cavity time constant of 600 s. Note that the passive-cavity time constant is reduced from the 1200 s modeled and measured in Si4. The G10 cavity support in Si4 is replaced with the aforementioned stainless steel mount which has a two times lower thermal resistivity. Unlike the 124 K system, the thermal capacitance of silicon is very small at 4 K and results in a low thermal time constant between the baseplate of the passive shield and the cavity mirrors.

4.4.1.2 Conductive exchange

Thermal isolation between the various stages with conductive heat flow is modeled as a series of low pass filters with a corner frequency determined by the support geometry and material thermal properties. A rod for instance will have a thermal resistance calculated analogously to the electrical resistance. Rather than use the resistivity ρ , thermal conductivity k yields a thermal resistance of $R = L/(kA)$ for length L and cross sectional area A . Thermal capacitance is simply the product of a material's heat capacity C_p and it's total volume V (often expressed as $\rho \cdot m$). A long thermal constant will maximize the product RC . With our commercial cryostation, we're clearly limited with the total volume available for packing material into the enclosure. It may appear that a more practical approach would simply be to prioritize the thermal resistance. In the presence of heat flow q , high resistivity will lead to larger thermal gradients as

$$q = k\Delta T \tag{4.36}$$

Within the cryostation, stage 1 and stage 2 of the cryocooler fix a minimum base temperature, so structures with large thermal resistance will find themselves hotter. This can be troublesome at 4 K operation where we want the cavity temperature to be as low as possible for reduced thermal noise and thermal expansion sensitivity. We can gather that mindful selection of our materials will be important for ensuring good thermal isolation.

Unfortunately, the thermal properties of most materials asymptotically approach zero at low temperatures. The Debye model, which provides an accurate description of phonon excitations in solids, estimates the low temperature thermal dependence of heat capacity to scale as $C_p \propto T^3$. This makes it extremely difficult for cryogenic systems to match the thermal time constants found in room temperature systems. The benefits provided by a reduced Brownian thermal noise can be offset by poor thermal instability, especially if the the temperature control is not precise enough to enable realization of thermal noise limited system.




Figure 4.22: Volumetric heat capacities of materials typically used in a cryogenic environment. Common materials (dark blue curves) have a vanishingly small heat capacity below 10 K, requiring use of more exotic alloys.

4.4.1.3 Material properties

So how do we engineer a good thermal isolation? Fortunately the NIST cryogenics division maintains a database of materials suitable for cryocooler applications, as well as several useful associated properties. Fig.4.22 displays the heat capacities for most materials that might be placed in a cryogenic environment. At 4 K, a few materials stand out as promising choices. First, we can see that a group of Gadolinium-based materials reveal a drastic change in C_p close to our desired operation temperature. Particular rare-earth magnetic materials exhibit an antiferromagnetic phase transition at low temperatures, and an anomalously large heat capacity will be evident close to the critical point.

Gadolinium falls in this class of materials, and maintains this property when formed as the compounds $\text{Gd}_2\text{O}_2\text{S}$ (GOS) and GdAlO_3 (GAP). Addition of gadolinium into different mixed

oxides shifts the temperature of the phase transition, providing a unique method of tuning the heat capacity. Unfortunately neither GOS nor GAP are metallic alloys, instead they are formed as compounds that are sintered into ceramics [121]. Most of these materials have their thermal properties measured while being used as "regenerator materials", meaning that the compound will be formed into particulate spheres with $\sim 100\mu\text{m}$ diameter and used as a heat exchanger for the helium gas that flows through it. Synthesis via sintering is nontrivial and requires powders of the oxide and material to be heated and pressurized with particular granule sizes to form large structures of the compound. At the time of our interest, no commercial vendors were found that could produce large samples of GOS. It was additionally unclear whether a large ceramic block of GOS could be easily machined, we found no literature on this. Most of these exotic materials may be suitable for use as fine powders in cryogenics, but it seems that there is still significant engineering required to make them useful for our applications.

HoCu_2 is a metallic alloy that is easier to synthesize, and has a consistently large C_p over a large temperature range. It is generally useful as a high heat capacity material, but more importantly is produced by our commercial cryostation vendor Montana Instruments. This has become our high- C_p material of choice for 4 K operation since they produce and machine the alloy to our specifications.

Commercial vendors are capable of providing arbitrary alloys of most metals via the vacuum arcing process. This is capable of providing large bulk samples of the alloy, and appeared to be an attractive means for getting HoCu_2 samples that we could then machine into cryostation components. We ordered a large sample of HoCu_2 from ACI alloys that was arc-cast into a few $40\text{ mm} \times 12\text{ mm}$ cylindrical disks. Our intent was to test capabilities of machining this sample, though we quickly found this to be much more difficult than expected.

Initial handling and inspection of the disk suggested that the arc casting process formed a sample with several non-uniform grain boundaries. Working the sample with an EDM cutter ripped out large chunks of it, and ultimately we did not find a method for consistently shaping the material. Likely the process used for forming the alloy is not ideal. HoCu_2 can also be formed through an

atomization process which works to combine the individual components at high temperature and then quickly solidify them. This ensures a more consistent material with the desired properties from the alloy, the vacuum arc process often leads to an inconsistency in phase of the bulk material leading to gradient properties. Alloyed samples of HoCu₂ were not used in the final redesign of the cryostation.

At 4.5 K HoCu₂ has $C_p = 0.3 \text{ J/cm}^3\text{K}$, so a time constant of 1000 s for a disk with a 10 cm^3 volume requires a thermally resistive component with $R > 300 \text{ W/K}$. For cavity support structures, we might often opt for a design where cross sectional area $A \ll L$, or $L \sim 10A$. For our desired time constant, the required thermal conductivity would be $k = 0.03 \text{ W/mK}$. The point of this is to illustrate that our lack of flexibility in choosing high- C_p materials means we will have to be fairly stringent with material choice for the thermally resistive components. Copper and steel, which may typically be used for thermal isolation in room temperature systems, have $k = 10 \text{ W/mK}$ and $k = 0.3 \text{ W/mK}$ respectively. Neither of these would be appropriate for our thermal filter, and generally most metals will not be. The Wiedemann-Franz law provides an empirical relation between a material's thermal conductivity and electrical conductivity σ ,

$$\frac{\kappa}{\sigma} = LT \quad (4.37)$$

For temperature T and proportionality constant $L = 2.4 \times 10^{-8} \text{ V}^2/\text{K}^2$. Metallic compounds which will naturally be electrically conductive, will not be candidates for providing us with high thermal resistance. Rather, we are better off using materials that are electrically insulating like plastics or ceramics. PEEK is a polymer with excellent thermal resistivity and is commonly used in cavity systems. However, its Young's Modulus of 3.6 GPa is very low, and would be prohibitive in applications requiring sharp, high-Q mechanical resonances far from the frequency band of interest. Composite materials with glass or polymers bonded with an epoxy (e.g. G-10 fiberglass) suffer from this same issue, where the epoxy leads to a low thermal conductivity, but is also structurally weak when compared to metals. Thus a tradeoff exists where we often have to balance good mechanical properties with our desired thermal properties.

Carbon fiber reinforced plastic (CFRP) was introduced earlier as a suitable material for optimizing the vibration sensitivity of the cavity support. Its stiffness significantly surpasses that of G-10 and PEEK, but the thermal conductivity is comparable. Cryogenic measurement of this value is inconsistent, since different carbon fiber-epoxy ratios can lead to different values for k (adding more epoxy will reduce the stiffness and the thermal conductivity). Values of $k = 0.03$ W/mK at 4 K seem reasonable [122]. The passive platform is supported by a series of G10 rings epoxied to the active base. We replace these with thinner CFRP rings, gaining a significant increase in active-passive thermal time constant due to the lower thermal conductivity and smaller cross sectional area. The acceleration sensitivity is not compromised due the larger > 100 Gpa Young's Modulus.

4.4.1.4 Thermal shields

From Si4, there were existing concerns that large thermal gradients could be present along the length of the shield. The vendor provided shields were originally 0.5 mm thick aluminum, and John Robinson had found significant coupling of the cavity frequency to thermal fluctuations of the radiation shield. Estimations of the radiation shield temperature suggested that it may be at 80 K, significantly warmer than the stage 1 temperature of 30 K. His solution was to replace this with a custom designed shield machined from copper with 2 mm thickness and polished on the outside to reduce the emissivity. A thermistor placed on its outside surface after the swap recorded a temperature closer to 50 K. It is unclear whether there were residual gradients, but clearly the reduction in emissivity had a significant impact on the thermal environment.

Seeing the outcome of this upgrade, we decided to reconstruct all thermal shields for Si6. We first wanted to increase the thickness to 2 mm for the two 4 K shields, and 3 mm for the radiation shield. We chose 4N oxygen-free high thermal conductivity copper as our material, so the additional thermal mass from the increased shield volume would be negligible. High purity copper is characterized by its residual resistance ratio RRR (ratio of resistivity at finite temperature vs 0 K), and OFHC with RRR of 150 has $k = 2500$ W/mK at 30 K. Compared to the original aluminum




Figure 4.23: Thermal conductivity of high purity copper with different residual resistivity ratios.

shields, our proposed upgrades should reduce any thermal gradient by over 500 times.




Figure 4.24: Various gold-plated components for the thermal shields. The radiation shield is on the left, 4 K active shield in the middle, and 4 K passive shield on the right. A baffle with a 5 mm aperture on the radiation limits the room temperature radiation load passing through the windows to the cavity.

The Montana cryostation does not support thermal shields of our desired thickness - shields must be threaded onto a tapped baseplate and the difference in diameter for the active and passive

baseplates is slightly larger than 3 mm. We consequently adopt a unique approach in the shield design, forming them out of concentric threaded cylinders that can be twisted together. The complex machining required likely introduces many stresses into the components, so we anneal them at 500 C for several hours in an argon backfilled chamber.

As the purity of the copper increases, oxidation occurs more quickly and this would reduce the surface emissivity, leading to a higher radiative coupling. The benefits we gain from improved thermal conduction through the shield is less helpful if there is an increased heat load, essentially negated, so we gold plate all shield components after annealing. The JILA shop process doesn't apply a nickel strike as diffusion barrier since that plating process is line of sight only. Direct plating of gold on copper can reduce the plating emissivity as copper diffuses into the gold but some studies [123] suggest that a thick enough plating can inhibit this process. Surface emissivity is expected to be somewhere between 0.03-0.1 depending on the plating. Note that polished copper can easily reach emissivities in this range, but the process is labor intensive and may not be possible for complicated geometries.

4.4.2 Precision thermometry

A few types of parasitic effects can corrupt accurate and precise readout of the thermistor. With Joule heating, measurement of a thermistor with value R with a current I dissipates a power I^2R . A lower readout current is beneficial for mitigating this effect, but comes at the cost of reduced signal-to-noise in the readout. A thermistor can be chosen to have a lower resistance at the relevant temperature, but this is complicated by the fact that we would like a sufficiently large resistance from room temperature to 4 K.

Platinum based thermistors (PT-100) are a staple for accurately characterizing the room temperature environment around the 1D Sr optical lattice clock, and for measuring the thermal environment around the 124 K cavity. However, they quickly lose sensitivity below 100 K as the resistance drops below 10 Ω , necessitating a different option. Zirconium oxide thin film thermistors are a robust option that are ideal for thermometry between 1 and 100 K. However, they still have

a sizeable $> 30\Omega$ resistance at room temperature, unlike germanium or carbon based thermistors which are often used in cryogenic environments. A typical resistance and sensitivity curve is shown in Fig.4.25.




Figure 4.25: **Left:** Temperature calibration of the three Cernox thermistors installed on the cryogenic thermal shields. **Right:** Sensitivity of the thermistors.

Self heating effects are more significant below 4K where the thermal mass of the support structure is greatly reduced compared to all higher temperatures.

Reading out an accurate value of the thermistor is also nontrivial as the readout cables need to be electrically conductive but thermally resistive. A normal copper wire cannot be used as it will provide a thermal short from the cryogenic environment to the room temperature feedthrough on the vacuum chamber. We instead use phosphor-bronze wire as it provides strong thermal isolation with a bulk thermal resistivity of $k = 1.6 \frac{W}{m \cdot K}$ at 4 K temperature. A conservative estimate for 36 AWG gauge with cross sectional area 0.0127 mm^2 and a length of 10 cm will flow $5 \mu\text{W}$ of heat between the 4 K and 30 K stages. The stage 1 of the cryo can provide significantly larger cooling power, so the estimated heat flow can be safely managed. Electrical resistivity appears sufficiently small with $\rho = 11\mu\Omega \cdot \text{cm}$, but the same gauge of wire yields a resistance per unit length is $R/L = 8.7 \Omega/\text{m}$. This is sizeable compared to the thermistor resistance as a 10 cm strand of cable will appear as a nearly 1 K temperature offset at 10 K. These resistive contributions of the cables are eliminated by performing a 4 wire measurement of the thermistor resistance.

Cables were soldered to the thermistors with 60/40 lead/tin and residual flux was cleaned off

with an acetone-dipped swab. The dissimilar junction between the readout wires and the thermistor can generate an electromotive force via the Seebeck effect which provides a slowly varying voltage offset. An upgraded measurement scheme of using an AC bridge to perform the 4-wire readout was planned for a next generation upgrade of the thermometry, but was not implemented in this work. A MKT50 from Anton Paar was used for the majority of low noise resistance measurements.

The Si3 system clamps thermistors to the thermal shields with a small retaining bar screwed into the copper shield. The small volume of the Si6 cryostat prevented us from designing thermal shields with > 1 mm of clearance between the stages and we could not use the same technique given the 0.8 mm vertical profile of the Cernox thermistors. Lakeshore technicians recommend indium solder as the ideal means of affixing the thermistors to a metal surface for maintaining UHV conditions but initial attempts to solder the thermistors to the thermal shields were unsuccessful. Silver based epoxy was suggested as another alternative. The Montana Instruments cryostat makes liberal use of cryogenic varnish to hold various components together so we decided to use this, given that it worked well for Si4 operation. A thin film of Apiezon N-type grease was placed underneath the thermistor, with IMI-7031 varnish applied to the edges. The readout cables were additionally clamped to the shields in a few places for stress relief and the structure was cured at 150C for 12 hours.

4.4.3 Cooldown tests

With the multitude of design changes implemented on the cryostation, we worried that the additional material might prevent us from reaching a base temperature of 4 K. We opted to cool the system down to its lowest achievable base temperature before inserting the Si6 cavity. With the full array of temperature sensors in place, this would allow us to measure thermal time constants between various stages as well as identify any thermal shorts or strong couplings to the environment.

Multiple sensors were placed on the active, passive, and radiation shields to characterize thermal gradients across them. The silicon cavity was replaced with a copper disk designed to give the same thermal mass at 4 K and a thermistor was also placed on this to characterize the expected




Figure 4.26: **Left:** A copper disk is used as a dummy cavity, a thermistor attached to it allows us to measure the expected cavity temperature during cooldown. **Right:** The active shield with thermistors installed. Temperature sensors from the passive shield pass through a small slit beneath the lid. A thin layer of vacuum grease between the lid and cylindrical section improves the thermal contact. Cables for the sensors are wrapped around each shield several times for thermal anchoring. This is especially important on the radiation shield so that heat flowing from the 300 K vacuum feedthroughs is dumped on the 30 K stage where the coldhead has more cooling power. Sensors are attached to the shields with VGE-7041 varnish and baked at 150 C for 6-10 hours for curing.

cavity temperature and time constants.

The initial cooldown thankfully brought the test mass down to cryogenic temperatures, though it settled to 5.8 K, slightly hotter than our expected 4 K. The CTE of silicon starts to rapidly increase above a temperature of 5 K so some minor improvements were needed in the architecture of the thermal environment. Surprisingly, we identified large, nearly 1 K temperature gradients across the active shield despite the increased shield wall thickness and the significantly higher thermal conductivity of high purity OFHC. Fluctuations between the bottom and top of the shields were highly correlated however, with a negligible thermal time constant. No amount of numerical simulations could reproduce this gradient and given the strong coherence of thermistors

across the shields, it's likely that the thermistors themselves had a sizeable thermal offset. When warming the chamber to room temperature, the gradient was clearly non-monotonic (Fig.4.28) and seemed to take unphysical values, e.g. changing sign at various temperatures. It is possible that the first order discontinuities in the warmup curve were the result of stress release in the sensor - cable junction. In hindsight it would have been useful to perform a huddle test of all thermistors during a practice cooldown to identify any offsets and jumps in their temperature readings, but the mounting stresses would no doubt be different after installing them on their respective surfaces.




Figure 4.27: During one of the initial cooldown tests, two thermistors were placed on each of the active and radiation shields to measure thermal gradients. By design, the high purity OFHC copper should have suppressed gradients to the mK level. Surprisingly, a residual gradient of 1 K was measured on the radiation shield (left column) and 0.65 K on the 4 K active shield (right column). We were unable to find a proper explanation for this, it is possible that curing the epoxy that attached each thermistor to the shield introduced resistance offsets via mechanical stresses. This is another argument for securing the thermistors with a mechanical clamp in future systems. The bottom row displays an Allan deviation of the gradient fluctuations. The active shield gradient is remarkably static with fluctuations slightly over the $10 \mu\text{K}$ level over 1000 s.

Ultimately, the first cooldown taught us that our careful attempts to thermally anchor ther-




Figure 4.28: The cryostat is heated from 4 K to room temperature while we read out thermistors on the top and bottom of the active shield. The difference in measured temperature (right) varies significantly over the warmup process.

mistor cables to various stages were insufficient. Stepping temperatures across the radiation shield revealed that the testmass temperature would increase with a slightly shorter thermal time constant than the passive shield, indicating heat flow between the testmass and radiation shield through the sensor wires. We went through a few iterations of cooling down, identifying an issue, warming up, and making minor improvements. After a few thermal cycles, the cryogenic varnish would peel off the shield resulting in a sensor dangling in free space, or worse, shorting two shields together. This technique worked well given the minimal clearance between shields, but in the future it would be preferable to use a mechanical clamp with the sensor sandwiched between layers of indium foil. Ultimately, it was crucial to wrap the readout wires around each thermal shield several times over to provide a proper thermal anchoring of the sensor.

As we learnt more through this process, the active-passive time constant increased by a factor of two, from 5000 s during the initial test to 11000 s on the penultimate cooldown. Our model of the thermal resistance circuit predicts a ~ 10000 s time constant between the active and passive stages. This seemed to indicate success in removing unwanted thermal links and the diminishing returns dissuaded us from further tests. At this point, we decided to cooldown the cavity rather than the testmass. In our final configuration, we ended up installing a single sensor on each of the radiation, active, and passive shields.




Figure 4.29: Temperatures on the thermal shields and cryocooler stages are continuously measured during the final cooldown.

4.4.3.1 Cooldown procedure

Cooling down the cryostation with the cavity in place is a religious experience. The relatively large linewidth of Si6 already makes RAM control difficult and any further reduction would be disastrous. We try to apply standard cleanroom procedures, with 1/4-inch rubber curtains installed around optics enclosure and a HEPA operated inside. The optical table and cryostation are thoroughly cleaned of any dust and everything but the cavity mount is covered with UHV foil. Kapton tape covering the vent holes of the spacer is quickly removed and the cavity needs to be set on its support at the measured angle that minimized total acceleration couplings. The passive shield and its lid are then immediately placed over the cavity, hopefully giving some protection from dust. The active shield is threaded into its mount, cables for the passive sensor are fed through notches in the lid, and then thermally anchored to exterior of the active shield. The same procedure is done for the radiation shield, with thermal anchoring occurring in several locations to impede direct heat flow from room temperature to 4 K. With the cryostation closed, we quickly measure all thermistors and heater resistances to confirm that there are no electrical shorts. After realigning to the cavity, the finesse is remeasured to ensure that there is no contamination of the optical coatings. Vibration sensitivity is also measured to ensure that the cavity is placed at the correct angle relative to the

support structure. Vacuum can then be pulled on the cavity, and the cooldown to 4 K is initiated⁴

4.4.4 Thermal time constants

With the cavity at 4 K, we can verify thermal time constants between various stages by performing a discrete temperature step and recording the cavity frequency response. With the active shield held at a fixed value after a temperature step, the thermal behaviour of the passive shield and cavity follows

$$C_{act}\dot{T}_{pass} = \frac{1}{R_{act-pass}}(T_{act} - T_{pass}) - \frac{1}{R_{pass-cav}}(T_{pass} - T_{cav}) \quad (4.38)$$

$$C_{pass}\dot{T}_{cav} = \frac{1}{R_{pass-cav}}(T_{pass} - T_{cav}) \quad (4.39)$$

Transforming the above ODEs into Laplace space, we find that the cavity response is well described by a second order low-pass filter with poles corresponding to the active-passive and passive-cavity time constants (A step response with the dummy cavity shown in Fig.4.30 highlights this effect). Our final measurements of the time constants between the various 4 K stages were $\tau_{active-passive} = 12000$ s, $\tau_{passive-cavity} = 300$ s. The amplitude of the cavity's frequency response tells us its sensitivity to thermal fluctuations, and we find a CTE of 4.5×10^{-11} K⁻¹. Notably this is larger than the 2×10^{-11} found in Si4, indicating that Si6 spacer is sitting at a higher temperature. Our cooldown tests indicated that the cavity equilibrated close to 4.9 K as opposed to the lower 4 K in Si4. The cavity's length response to thermal fluctuations will be 2.5 times larger, but the thermal time constant between the active platform and cavity has been increased nine-fold. There is still a net gain in suppressing temperature variations at 4 K.

A similar step response can be performed with the radiation shield, though in this case the active shield is actively servoed to prevent any confounding coupling between the radiation and active platforms. At 4 K, a radiation shield - cavity coupling of 10 Hz/K was measured with a time

⁴ The cavity frequency will move by several MHz as it cools to 4 K. As the cavity warms from 4 K and 16 K, we measure a frequency shift of 800 kHz. This contrasts with the 1.6 MHz that would be expected from integrating the silicon CTE with data from [124]




Figure 4.30: **Thermal time constants** In a test cooldown of the system, the cavity is replaced with a copper disk of the same thermal mass. Thermal sensors are attached to the testmass, inner 4K shield, and outer 4K shield. The temperature of the outer 4K shield (active) is stepped via a ring heater wound around its base, and we monitor the subsequent temperature increases of the inner 4K shield (passive) and testmass. We extract the time constants between the different stages as well as the thermal gradients between them. The ring heater on the active shield base also allows for active temperature stabilization to either the active or passive thermistor.

constant of 3000 s. Upon warming up to 16 K, we were unable to identify any frequency response to the temperature change, reflecting the lower CTE of the cavity near it's zero crossing⁵.

A step response on the vacuum chamber revealed a frequency coupling of 200 Hz/K, significantly higher than the 4 Hz/K measured in Si4. Our radiation shield was measured to be below 50 K, so radiative coupling to the vacuum chamber could not explain our significantly higher coupling measured. In [43], it was speculated that the reduced room temperature coupling of 4 Hz/K was largely due to the thermal upgrades that resulted in a significantly lower radiation shield temperature. One key improvement is that the manufacturer supplied optical windows were exchanged with superpolished BK-7 windows in the hopes of reducing optical scatter. During the Si6 upgrades, these were replaced with thinner superpolished UVFS windows due to concerns over the BK-7 performance deep in the cryogenic regime. UVFS transmits more of the 300 K blackbody spectrum than BK-7 but not enough to explain the 50 times larger room temperature coupling. It

⁵ At the time of this measurement we were limited by an etalon in the system that modulated the beatnote by an amplitude of 2 Hz, it was difficult to identify any exponential setting with this present

is more likely that heat flows through the temperature sensor cables at the 300 K feedthroughs to the cryogenic stages. The exact path is unclear, but suggests a possible mechanism for the higher coupling to room temperature fluctuations.

4.4.5 Temperature stabilization

With our couplings identified, we now know the level of stabilization needed on each stage to attain thermal noise limited performance. Coupling from the 30 K stage only needs temperature stabilization to 0.5 mK. We accomplish this by feeding back to the power applied to a resistive heater wrapped around the base of the radiation shield. This loop is closed with ~ 1 Hz bandwidth, and along with the 3000 s time constant, we are confident that this does not contribute any instability above thermal noise out to 1000 s. In Fig.4.27, we identified that the gradient along the shield is stable below the 0.5 mK for all measurement times.

Stabilization at 4 K is a more daunting task due to the larger coupling via the cavity CTE. The 4 K active shield also has a resistive heater wrapped around the platform base, and we can use this to close a loop on a thermistor placed on this shield. In one of our test cooldowns, we placed a second witness thermistor on this shield to view out of loop temperature fluctuations with 4 K temperature stabilization enabled. The witness performance, visible in Fig.4.31, closely follows the inloop stability though it seems to saturate at $10 \mu\text{K}$ temperature variations. We're unsure whether this is due to the noise floor of our measurement device or if this was an actual differential fluctuation. The witness stability is insufficient for the cavity performance we require, limiting it at the 4×10^{-16} level. We can however make use of the long thermal time constant between active and passive stages, and calculate the lowpassed temperature fluctuations that we would expect at the cavity. Using our measured CTE, we can calculate the technical contribution of 4 K fluctuations and find that we would expect the thermal noise limited performance to 200 s.

The cavity can still experience temperature fluctuations that are not entirely captured by the variations on the active shield. On the Si4 system, the 4 K platform was stabilized to the passive shield temperature and saw improved frequency stability. On Si6, a thermistor was placed on the




Figure 4.31: **Top left:** Time series of the inloop and witness thermistors placed on the 4 K active shield (loop is closed on the inloop thermistor). **Top right:** Instability of the two thermistor readings. **Bottom left:** The witness' temperature fluctuations are converted to frequency fluctuations via our measured CTE. The brown curve is the witness variations filtered by the passive platform and cavity mount. **Bottom right:** Instability of the traces on the left.

passive shield and we could feedback to variations on its reading with a slow $< \text{mHz}$ loop bandwidth. For the experiments described in the next chapter, we never observed any coherence between the passive shield temperature and the cavity instability (perhaps due to the larger active-passive time constant). We therefore stabilize the 4 K platform to the active shield temperature, and do not read out the thermistor placed on the passive shield. Excitation with $100 \mu\text{A}$ of current was enough to slightly heat the cavity, so we disable the readout to remove an additional heat source to the cavity. The potential for stabilizing the passive shield temperature is attractive when the cavity is operated at 16 K, but in studies of the long term instability we again never saw evidence that thermal variations were a dominant noise source.

4.5 Vacuum pressure

So far we've discussed physical effects that can vary the cavity length by coupling to the spacer. Our true intent is to keep the optical path between the spacer to be constant. Gases within the spacer borehole vary the vacuum index of refraction n_0 , and will easily modify it by coupling to pressure and temperature.

Variations of air's refractive index have been well characterized with the Edlen equation [125], so we can get a sense of what vacuum stability we may need within the cavity to achieve thermal noise limited performance. The absolute refractivity n of air at 15°C and 1 atm varies dispersively with wavenumber $\sigma[\mu\text{m}]^{-1}$ as [126]

$$(n - 1)_s \times 10^8 = 8342.13 + 2406147(130 - \sigma^2)^{-1} + 15997(38.9 - \sigma^2)^{-1} \quad (4.40)$$

The 3σ uncertainty of this equation applies to visible radiation, but for just attaining a rough quantity we can extend it to the IR. At $1.54 \mu\text{m}$, $(n - 1)_s = 27632 \times 10^{-8}$. The pressure and temperature dependence of the refractive index is modified as

$$(n - 1)_{TP} = \frac{(P/\text{Pa})(n - 1)_s}{96095} \times \frac{1 + 10^{-8}(0.601 - 0.00972 T/^\circ\text{C})(P/\text{Pa})}{1 + 0.003661 T/^\circ\text{C}} \quad (4.41)$$

The rough scaling of n with temperatures T and pressures P at ambient conditions goes as

$$(n - 1)_P \approx 2.8 \times 10^{-7}(P/\text{mbar}) = 2.8 \times 10^{-4}(P/\text{atm}) \quad (4.42)$$

$$(n - 1)_T \approx 1.1 \times 10^{-6}(T/^\circ\text{C}) \quad (4.43)$$

Limiting pressure induced frequency variations to the 10^{-17} level requires the optical path within the cavity to be this stable

$$\frac{\delta\nu}{\nu} = (2.8 \times 10^{-7}/\text{mbar}) \cdot \delta P \quad (4.44)$$

The constraints on the intracavity pressure are stringent, requiring a UHV environment to be maintained with the cavity spacer. Vacuum conductance to the spacer interior is purposely kept to be poor so that it is difficult for dust to enter this region before pulling a vacuum.

However, the cryogenic environment of around the cavity greatly assists us in lowering the background pressure at 4 K by adsorbing any molecules that strike a cold surface. The saturated vapor pressure of most gases decreases dramatically with temperature. The vapor pressure for most common gases is illustrated in Fig.4.32. Aside from hydrogen and helium, there is a negligible amount of other molecules at 4 K. The large surface area of the thermal shields effectively acts as a cryo pump by condensing background gases that collide with them. Small charcoal adsorbers mounted above the cavity serve to capture any remaining molecules that could enter the cavity borehole.




Figure 4.32: Vapor pressure of most gases as a function of temperature, taken from [9]. Below 20 K, the only residual gas aside from Helium with a pressure above 10^{-13} Torr is H_2 , which has a vapor pressure of $\sim 10^{-6}$ Torr at 4 K. At the UHV range, most curves are extrapolated with decreasing temperature.

The variation of the index of refraction at cryogenic temperature is negligible for the cavity frequency, but it's interesting to consider how the dispersion of n from Eq.4.40 can generate a frequency shift between the carrier and modulation sidebands, i.e. do we need to worry about the air's dispersion generating RAM? For modulation tones 5 MHz offset from the carrier, it turns out that this effect contributes a frequency shift $\delta\nu/\nu < 1 \times 10^{-21}$.

Despite the cryo-pumping being sufficient at 4 K, it is naturally less effective between the 300 K vacuum chamber and the 30 K shield. Montana instruments advertises the vacuum environment of their cryostation as being suitable for most applications that may be insensitive to pressure fluctuations outside the sample space. For our system, these effects vary the optical path length and can appear as a frequency shift of the cavity. The residual pressure fluctuations in this environment were sufficient for Si4 to realize its thermal noise limit, but we discovered an important coupling to acceleration noise that could be critical for measuring improved stabilities. In Si4, measurements of the vibration sensitivity always produced unphysical results at low frequency. Measurements of the DC sensitivity via tilting the optical table and AC measurements from the coherent displacement of the AVI would return horizontal sensitivities of $\sim 1 \times 10^{-9}$. Multiplying this by the ambient acceleration noise would result in a frequency noise significantly higher than the thermal noise limit. At the time, we attributed this to some flaw in the measurement that we could not pinpoint.

The same phenomena was observed in Si6, but an improved acceleration sensitivity was critical for achieving thermal noise limited stability. Converse to Si4, our measurements were performed with the cryostation at room temperature and pressure and the large rollup of the k_{H1} sensitivity that was present in Si4 was still visible (Fig.4.18). By accident, we discovered that the abnormally large horizontal sensitivities could be reduced by removing the lid of the cryostation's vacuum chamber, opening the radiation shield directly to the environment. This does not change the steady state atmospheric environment around of the thermal shields or cavity but we inferred that shaking the AVI generated some modulation of the refractive index of air within the chamber. These pressure fluctuations were coherent with the dynamic displacement of the AVI, and could therefore appear as a coherent frequency response of the cavity. Opening the vacuum chamber to the environment eliminated the drum-like modes of the cryostation lid that were being excited and we instead found the cavity's acceleration sensitivity to be largely flat at low frequency (Fig.4.15). We hypothesized that driving the AVI platform compresses the bellows connecting the vacuum chamber to the coldhead and drives air around the thermal shields.

Removing any air from the cryostation should eliminate this effect, yet when we remeasured




Figure 4.33: Schematic of the modifications for additional pumping on the cryostatation. Montana Instruments provides a control unit connected to the cryostatation with a long 2 m, 0.375" ID vacuum bellows. Within the control unit is a small roughing pump that is valved off from the bellows with a case valve contained within the control unit. A vent valve can be left open to the environment or connected to a nitrogen canister to vent the system to ambient pressure. We removed this connection by separating it from the pumping bellows with a gate valve. Instead, we add a tee that is connected to a turbo and scroll pump. After operating these two for a few days, we valve them off and activate a small 2 L/s diode pump.

the acceleration sensitivity after the final installation of the cavity (involved pulling vacuum on the system), the same rollup in acceleration sensitivity was observed. The process of pumping the vacuum chamber was directed entirely by a control unit provided by Montana Instruments. A small roughing pump within the unit would run for a few minutes before valving off the system, upon which it could be cooled to cryogenic temperatures. Exposing some of the vacuum connections within the control unit revealed that a 2.3 m bellows with a 3/8 inch inner diameter connects Montana's roughing pump to the coldhead. Pumping to vacuum chamber was therefore severely conductance limited, with an estimated volume of 5 mL/s removed. Achieving a pressure of 1 mTorr at the cryostatation would require more than 3 hours of continuous pumping, far longer than the control unit determined.

We opted to disconnect the control unit from the long vacuum bellows and instead attach a

turbo backed by a roughing pump. Running this continuously for a three days brings the cryostation to a pressure close to the minimum of the turbo's vacuum range. We then close a valve placed after the turbo (see Fig.4.33 for a schematic) and activate a small 2 L/s Vaclon diode pump. After this process, intentionally shaking the AVI platform results in no changes in the acceleration sensitivity. Furthermore, the current on the diode pump does not show any fast or slow variations. In the limit where vibrations dominate the noise budget of the cavity, the measured frequency noise is now consistent with our measurements of ambient accelerations propagated through the cavity's acceleration sensitivity function. We conclude that refractive index changes due to vibration-driven pressure fluctuations do not contribute to the measured cavity instability.




Figure 4.34: **Left:** Frequency noise of Si3-Si6 beatnote when the vacuum pumps are running and when shut off. There is no additional noise measured due to vibrations of the pumps. **Right:** Fractional instability of the beatnote data on the left. In the low frequency regime, we can slowly displace the AVI platform slightly above the ambient vibration noise and look for corresponding changes in the cavity noise. Below 1 Hz, accelerations coupling directly to the cavity should generate frequency perturbations significantly below the thermal noise limit. Any observed increase in cavity noise could therefore be an indication of poor vacuum within the cryostation. We do not observe any such effects.

4.6 Dissemination noise

When all technical noise sources are below the thermal noise limit, a laser locked to the cavity fringe will inherit the thermal noise limited instability. However, the laser has to be distributed to somewhere for characterization, either via beatnote with another laser or by measurement with

an atomic reference. In our case, we analyze a heterodyne beatnote of Si3-stabilized light and Si6 stabilized light. As we clarified in Fig.4.9 within our discussion of residual amplitude modulation, the location of our detector is crucial for determining what noise gets removed. Similarly, the location we use to pick off light for distribution can have an impact on the frequency noise we measure. Phase shifts picked up in the optical path before the resonator are indistinguishable from cavity noise when measured at the PDH detector. The PDH servo will apply a frequency correction to the laser to nullify this effect, but measuring the light before it encounters these perturbations will effectively write this noise onto it. These undesired effects typically appear in the RAM signal as differential shifts between the carrier and modulation sidebands due to dispersion or etalons, and we've established that active RAM stabilization cannot be implemented on the cavity signal.

A typical solution is to pick off the distribution light as close to the PDH detector as possible so that differential shifts are minimized. The distribution light will typically be 100s of μW in power, as compared to the few μW incident on the cavity. Splitting the light at the PDH detector location is therefore problematic since it will generate huge photothermal effects. In the Si6 system, we associate this with significant frequency perturbations, so this is not feasible. Modulators and fiber components at 1542 nm are plentiful, so we instead opt to pick off light in the fiber path of our optical setup, before it is outcoupled into free space.

A 90/10 fiber splitter after AOM1 (Fig.4.36) sends the majority of the laser's optical power through a 5 m fiber to beat with a laser stabilized to Si3. Significant frequency noise is picked up through environmental perturbations to the fiber, so we implement noise cancellation on this link via feedback to an AOM placed before the 5 m fiber. An in-fiber Faraday mirror forms the phase reference for one of the arms in the noise stabilization interferometer, and is sealed within an acoustic enclosure. The inloop performance of the servo is well below the cavity thermal noise floor (Fig.4.35), but this doesn't equate to the actual frequency noise that will be added to the distributed light.

Extraneous noise encoded at the fiber noise beatnote, or length fluctuations of the reference arm will result in an additional frequency noise contribution. It's therefore important to perform




Figure 4.35: The fiber noise cancellation servo is locked with high bandwidth of ~ 150 kHz. There is little noise present at this high offset frequency, but significant frequency noise is present at 1 kHz and below. The high servo bandwidth is necessary for providing the large servo gain needed to suppress the in-loop error to the cavity thermal noise limit.

an out of loop measurement to identify that the servo is actually removing the noise we would like. One way to quantify this is to perform a loopback measurement, where an additional fiber link at the remote location is noise cancelled and distributes light back to the initial location. This light beats with a reference arm at the initial location, so the heterodyne beat is encoded at the sum frequency of the two AOMs used for each fiber link. This requires twice the infrastructure of a single link, and also involves perturbing the setup at the remote end. In the specific case of the silicon cavities, this beatnote is formed on the frequency comb table and could potentially perturb the comb stabilization to the Si3 1542 nm light.

An easier and equally valid method is to simply take the fiber distributing light to the remote location, and instead bringing it back to the initial location to beat with the reference arm. A diagram detailing this setup is shown in Fig.4.36. The fiber noise beatnote between the reference arm and the distributed light (offset by the FNC AOM frequency of 77 MHz) is measured on a separate out-of-loop photodetector. Note that we distribute light to Si3 before it passes through the fiber EOM. This is to avoid extra modulation tones on the Si3-Si6 beatnote, as well as to avoid the photon shot noise written onto the laser by the RAM servo. There is a significant path length difference between the PDH detector and the distribution fiber, so we may worry about additional

frequency fluctuations as explained earlier. To properly characterize this, we would ideally beat the distributed light with light from the cavity transmission to identify any differential noise. We will discuss this in a later section.

The instability of this beatnote is displayed in Fig.4.37. No deviations above the projected Si6 thermal noise level were observed from averaging times of 100 ms to 1000 s. Metrology applications requiring an accuracy evaluation would also need to know whether there a frequency offset identified, our measurement only indicates that any offset is stable. We find that the light dissemination introduces a fractional frequency offset of $< 1 \times 10^{-21}$.




Figure 4.36: Schematic of the setup used to measure the out-of-loop fiber noise. Light from the NKT laser is split into two paths, one going to a fiber noise cancelled interferometer, and the other serving as a frequency reference. The fiber noise cancelled link, which typically distributes light to the Si3-Si6 photodetector, is instead sent to a separate out of loop photodetector to generate a heterodyne beatnote with the reference light. Frequency noise of this beatnote measured on the OOL PD quantifies the frequency noise associated with the light dissemination.

4.7 Technical noise budget

All the mechanisms described so far will contribute frequency noise that can degrade the cavity stability. Fig.4.38 compiles these to calculate the frequency noise we can expect from technical sources. Ideally these would all be below the thermal noise limit, and the majority of them are with the exception of RAM. We've discussed that the RAM contribution may be generated from some in vacuum scatter, and the only way to rectify this would be to warm the cavity to room temperature, identify and then clean the offending optic. This is an invasive procedure and since




Figure 4.37: Measurement of the dissemination noise using the schematic shown in Fig.4.36. **Left:** Fractional frequency noise of the measurement. **Right:** Fractional Allan deviation of the measurement. Dissemination noise is well below the projected thermal noise limit of Si6.

the current level RAM would already allow for a stability improvement over Si4, we thought it would be best to proceed and measure the frequency noise of Si6.




Figure 4.38: Compilation of technical noise sources measured on Si6. The frequency noise due to intensity fluctuations is discussed in the following chapter.

Chapter 5

Frequency noise of the cavity

Of course, the real test of the achieved performance can only be accomplished with a second, independent detector system. Disappointment is the experimenters' first reward for this measurement. - Jan L. Hall

The objective of the Si6 system was to distribute an improved local oscillator to the strontium experiments where the lower instability could be used to realize record levels of atom-light coherence. Silicon cavities I-IV followed a standard protocol of reducing technical noise sources below the thermal noise limit to achieve a thermal noise limited instability. Recent progress on ultrastable resonators has not yielded any significant surprises in this regard. The work is difficult, but a recipe exists for reducing the oscillator linewidth. Hints that we did not fully understand dielectric effects were noted in Si4 [43] with the surprising observation of a power dependent drift rate. Nevertheless, coating effects were thought to be well understood regarding the laser linewidth.

Cavities with crystalline coatings demonstrated that this was not the case. Si5 and Si6 observed a multitude of bizarre effects, power dependent frequency noise and correlated mode fluctuations that had never been measured in dielectric, amorphous coatings. It took months to even identify these phenomena as coating effects, and only recently has there been an understanding that a dedicated effort is necessary to understand these processes [127]. In this chapter we'll detail our measurements of these novel coating effects and present new methods for studying and suppressing them.

5.1 Photo-birefringent noise

5.1.1 Photothermal effect

In the technical noise budget, we mentioned that intensity noise can drive frequency fluctuations, necessitating stabilization of the laser intensity. This mechanism is not due variations in the discrimination slope, but rather to the heat-induced length change of the cavity.

This was first noticed in the LIGO collaboration where it was discovered that the optical shot noise of an incident beam was driving surface fluctuations on the test mass coatings [128]. Coined the photothermal effect, photons absorbed in the coating would locally heat and deform the surface, generating a displacement that could be read out in the interferometer. The same mechanism is apparent in our Fabry-Perot resonator, photons absorbed within the coating generate an elastic contribution where the surface expands due to a nonzero coefficient of thermal expansion, and a refractive contribution where the nonzero dn/dT of the dielectric stack also changes the cavity length. Eventually the absorbed heat propagates to the substrate and cavity spacer, each of which has their own photothermal contribution.

A formalism from [129] details the photothermal contributions of the cavity, allowing us to calculate length changes originating from each component, and different mechanisms. There are two explicit effects we will consider, the first involves the coating, substrate, and spacer's material thermal expansion. Intuitively, we can expect this elastic response to depend on a few material parameters for each component i : the coefficient of thermal expansion α_i , the thermal conductivity λ_i , density ρ_i , and the heat capacity C_i . For a given beam radius r_0 , we can define a wavenumber

$$k_i = \sqrt{\frac{i\omega\rho_i C_i}{\lambda_i}} \quad (5.1)$$

Where k_i^{-1} denotes a penetration depth for the substrate and coatings. This is a roundabout way to indicate that there is a physically relevant timescale when the penetration depth is equivalent to the beam radius. We define this as the critical relaxation frequency

$$\omega_i = \frac{\lambda_i}{\rho_i C_i r_0^2} \quad (5.2)$$

For the substrate and coatings on Si6, we find $\omega_c = 1700$ Hz, and ω_s to be orders of magnitude higher due to the high thermal conductivity of silicon¹. A plot of the two coating and substrate diffusion lengths in Fig.5.1 compares them with relevant length scales of the cavity.




Figure 5.1: Coating and substrate diffusion lengths are plotted as solid red and green lines respectively. Dotted and dash dotted red lines indicate two length scales relevant to the coating diffusion, i.e. the beam size and the coating thickness.

I'll regurgitate the relevant expressions for the substrate and coating responses from [129]. The photothermal response can also be calculated in the linear regime as a transfer function, with a length change originating from a modulation of the optical power. First, let's define two parameters that we'll re-use in the coating and substrate responses. Making the substitution $\xi_c = k_c r_0$ with coating thickness d , the filter function F attenuates any response at a diffusion length larger than the coating thickness

$$F = \frac{1}{\cosh(\xi_c d / r_0) + \mathcal{R} \sinh(\xi_c d / r_0)} \quad (5.3)$$

¹ I don't quote a number since it's hard to find data on the cryogenic thermal parameters for silicon

The coating elastic response is then calculated as

$$\delta \tilde{X}_c(\omega) = \frac{\delta X_c}{\delta P} = \frac{\alpha_c}{\pi \lambda_c} \frac{\omega_c}{i\omega} \int_0^\infty d\xi \xi e^{-\frac{\xi^2}{2}} F \left[\right. \quad (5.4)$$

$$\gamma_1 \left([\cosh(\xi d/r_0) - \cosh(\xi_c d/r_0)] + \mathcal{R} \left[\frac{\xi}{\xi_c} \sinh(\xi d/r_0) - \sinh(\xi_c d/r_0) \right] \right) \quad (5.5)$$

$$\left. - \gamma_2 \left(\mathcal{R} [\cosh(\xi d/r_0) - \cosh(\xi_c d/r_0)] + \left[\frac{\xi}{\xi_c} \sinh(\xi d/r_0) - \sinh(\xi_c d/r_0) \right] \right) \right] \quad (5.6)$$

With coefficients γ_1, γ_2 defined as

$$\gamma_1 = \frac{1}{2} \frac{1 + \sigma_c}{a - \sigma_c} \left(1 + (1 - 2\sigma_s) \frac{1 + \sigma_s}{1 + \sigma_c} \frac{E_c}{E_s} \right) \quad (5.7)$$

$$\gamma_2 = \frac{1 - \sigma_s^2}{1 - \sigma_c} \frac{E_c}{E_s} \quad (5.8)$$

The substrate elastic response is

$$\delta \tilde{X}_s(\omega) = -\frac{\delta X_s}{\delta P} = \frac{\alpha_s(1 + \sigma_s)}{\pi \lambda_s} \frac{\omega_s}{i\omega} \int_0^\infty d\xi \xi e^{-\frac{\xi^2}{2}} F \left[1 - \frac{\xi}{\xi_s} \right] \quad (5.9)$$

The coating also has a thermorefractive response due to the temperature dependence of the layers' refractive indices

$$\delta \tilde{X}_{tr}(\omega) = \frac{\delta X_{tr}}{\delta P} = \frac{\lambda \beta}{2\pi \lambda_c r_0} \int_0^\infty d\xi e^{-\frac{\xi^2}{2}} \frac{\xi}{\xi_c} \frac{\sinh(\xi_c d/r_0) + \mathcal{R} \cosh(\xi_c d/r_0)}{\cosh(\xi_c d/r_0) + \mathcal{R} \sinh(\xi_c d/r_0)} \quad (5.10)$$

Where the parameter β specifically captures the optical length change with temperature

$$\beta = \frac{1}{4} \frac{(\frac{\partial n_H}{\partial T} + \alpha_H n_H) + (\frac{\partial n_L}{\partial T} + \alpha_L n_L)}{n_H^2 - n_L^2} \quad (5.11)$$

The total photothermal response of the cavity, elastic and thermorefractive effects included, is written as

$$\delta \tilde{X}_T = \delta \tilde{X}_{spacer} + \delta \tilde{X}_s + \delta \tilde{X}_c + \delta \tilde{X}_{tr} \quad (5.12)$$

These expressions are intimidating, but there is some useful intuition to be gained from staring at them. Parameters γ_1 and γ_2 depend solely on properties relating to the deformation and

modulus of the materials. If the coating and substrate have matched $\sigma_s = \sigma_c$ and $E_s = E_c$, we will find that $\gamma_1 = \gamma_2 = 1$. Parameter $\mathcal{R}$ depends on the thermal properties of the coating and substrate and similarly reduces to unity if they share the same thermal parameters. In this case, the elastic response of the coating will be miniscule. It seems surprising that the thin film response will depend so strongly on properties of the substrate but recall that we'll be operating in the low frequency regime where the heat has diffused to a length scale larger than the coating thickness.

It's clearly a tall order to have a coating and substrate share the same material properties if we wish to minimize the coating elastic response. There exists another method - note that the substrate and coating elastic response have opposite sign. If we have the freedom to choose our operating temperature, we can find a regime where the different thermal expansion coefficients α_c and α_s yield $\delta\tilde{X}_s \approx -\delta\tilde{X}_c$. This will also suppress the joint photothermal response of the coating and substrate. However, there will still remain the coating's thermorefractive response which additionally depends on the optical properties of the bragg stack. The most straightforward method of cancelling this effect requires the temperature dependence of the refractive index to have opposite sign for the high and low index layers, i.e. $dn_H/dT + dn_L/dT \sim 0$. The physical picture is complicated with the consideration that the intensity is reduced in each following layer of the stack by the attenuation factor n_L/n_H .




Figure 5.2: **Left:** Refractive indices of the high and low index GaAs and AlGaAs layers. **Right:** Temperature dependence of the refractive indices.

I only point this out to illustrate that it is highly complicated to assemble a traditional

cavity with no response to optical power fluctuations. Photothermal effects are hard to cancel, and an easier route is always to stabilize the laser power. Of course, this task becomes easier if the amplitude of the total photothermal response $\delta\tilde{X}_T$ is small.




Figure 5.3: Calculated photothermal transfer functions for the elastic spacer, substrates, and coating responses, and the thermo-refractive coating effect. In the low frequency regime where we aim to be thermal noise limited, photothermal effects from the spacer are the dominant contribution

5.1.2 Photothermal effects in crystalline coatings: transient response

The formalism presented above represents the best understanding we currently have for modeling photothermal effects in coatings. These models have been benchmarked and verified against measurements on interferometers and cavities employing traditional dielectric coatings. In the Si6 cavity, we made several anomalous measurements that defied some of the key conclusions drawn from the theory of photothermal effects.

The large polarization mode splitting of the cavity enables us to address and lock to either of the two polarizations (I'll arbitrarily call these E_1 and E_2). The sign of the photothermal response is expected to be determined entirely by the sign of $\delta\tilde{X}_T$. Our first measurements on the cavity however demonstrated that there may be additional effects arising with these crystalline coatings.

In Fig.5.4 we show a typical set of measurements where we drop the transmitted intensity of




Figure 5.4: Transient response of dropping the transmitted intensity from 300 nW to 100 nW for the E_1 mode (solid line) and the E_2 mode (dotted line). Intensity is dropped at ~ 10 seconds, and there is a fast transient, followed by a slow settling of the frequency that exhibits opposite sign for the two polarizations.




Figure 5.5: Transient response of the E_1 when changing the transmitted power between different values. A clear asymmetry in time constants is visible between the left and right panels. Increasing laser power results in a much faster frequency response than decreasing the power, indicating that there are different photothermal transfer functions depending on the sign of the intensity change.

the cavity from 300 nW to 100 nW by stepping the RF power applied to an AOM. The two curves show the frequency response when locked to either of the two polarization modes. Notably, there is a fast transient response where the laser frequency instantly increases, and then slowly increases or decreases with an opposite amplitude for the two polarizations. Nothing in our photothermal

model captures the behaviour where orthogonal polarization modes exhibit opposite frequency responses. Given that this was not observed in Si4 and the only change between the cavities was the choice of optical coating, it seems safe to assume that this peculiarity is specific to the switch to a semiconductor, crystalline material.




Figure 5.6: Frequency responses renormalized by the magnitude of the transmitted power change. Solid lines indicate curves taken when decreasing the optical power, dashed lines indicate an increase. Curves of same color are taken with the same set of initial and final power values. It is clear that there is no singular transfer function $\delta\nu/\delta P$ that captures the frequency response.

We can try to simplify the picture by simply stepping the laser power by different amounts when locked to a single polarization mode and observing the frequency response. Transient responses of the E_1 for different increases in transmitted power (Fig.5.5) display a nonlinear behaviour where the time constant of the frequency response appears to vary with the initial and final values of the intensity step. Furthermore, changing the sign of the intensity step (decreasing the transmitted power rather than increasing) shows a settling behaviour with significantly faster time constants. It seems that there is not a single photothermal transfer function that can capture the frequency response for any arbitrary modulation of the laser power. We can at least conclude that the transfer function depends on the sign of the intensity modulation.

To accentuate this point, we renormalize all the frequency responses of Fig.5.5 by the magni-

tude of the optical power step δP . The resulting set of curves in Fig.5.6 show drastically different amplitudes and time constants depending on the sign of the power change, initial value, and final value. We cannot identify a single transfer function $\delta\tilde{X}_T$ that can produce a universal set of curves for the parameters of Fig.5.5 data.

A more complicated model would also be required to explain the varying time constants observed in all transient responses. In Fig.5.7, attempts to fit an exponential decay to each transient results in large frequency residuals and does not appear to capture the trend of the data. We instead find that a double exponential fit of the form

$$f(t) = a \left(\exp\left(-\frac{t+b}{\tau_1}\right) + \exp\left(-\frac{t+b}{\tau_2}\right) \right) + c \quad (5.13)$$

provides a better agreement to the data. This isn't to say that we've proposed this model with some scientific insight, rather it's evident that there is more than one time constant at play. While the fit residuals are smaller, there still appears to be some remaining oscillation around the data so this heuristic model is incomplete.

Time constants returned from the double exponential fit are presented in Table 5.1. The fast transient occurs on timescales of seconds while the longer settling takes roughly an order of magnitude longer. Data appears to be roughly consistent between measurements on the two polarizations, though τ_{slow} varies significantly depending on the conditions of the intensity step.

	E_1		E_2	
	τ_{fast}	τ_{slow}	τ_{fast}	τ_{slow}
100 nW $\rightarrow$ 300 nW	1.9	6.4	2.7	9.3
300 nW $\rightarrow$ 850 nW	2.6	25.6	2.3	21.4
100 nW $\rightarrow$ 850 nW	1.5	12.1	2.0	16.6
100 nW $\leftarrow$ 300 nW	8.9	28.5	10.5	37.6
300 nW $\leftarrow$ 850 nW	0.27	17.2	3.6	47.6
100 nW $\leftarrow$ 850 nW	12.3	59.5	15.3	83.7

Table 5.1: Fast and slow time constants (in seconds) of a double exponential fit the data in Fig.5.7, and analogous data taken when locked to the other cavity mode.




Figure 5.7: **Left:** Single and double exponential fits of the data in Fig.5.5. **Right** Frequency residuals of a single exponential fit (pink) and the double exponential fit (grey).

5.1.3 Photothermal effects in crystalline coatings: coherent response

It's clear that there is no universal transfer function that easily describes the photothermal response of the coating to an intensity step. Despite the several complications that arise when we try to extract a transfer function, we're going to go ahead and measure one to see what else we learn about the coating response. In the previous section, we instantaneously step the intracavity intensity and look for the steady-state frequency response. Now, we'll modulate the intracavity intensity at a particular frequency, and look at the coherent frequency oscillation that appears at that same value.

Ideally we expect to measure something similar to the curves calculated in Fig.5.3. However, we know that the transfer function has an asymmetric response depending on the sign of the intensity change indicating some form of hysteresis, so it is difficult to extract a pure sinusoidal frequency response when given a sinusoidal drive. We attempt to minimize this effect by modulating the optical power with a depth large enough to resolve the signal above the residual frequency noise of the Si3-Si6 beatnote. Depending on the measurement bandwidth, this equates to an amplitude of a few Hz in the frequency response. The optical power through the cavity appears as

$$P_{tot}(t) = P_0 + \delta P(t) \quad (5.14)$$

Where the transmitted power P_0 is significantly larger than the modulated power $\delta P(t)$. The transfer function $\delta\nu(\omega_0)/\delta P(\omega_0)$ is then a measure of the cavity resonance frequency coherently modulated at frequency ω_0 .

A summary of our measurements is displayed in Fig.5.8. We address a single polarization mode of the cavity, and measure the cavity response as we modulate the optical power at different frequencies. In all these measurements we ensure that we are in the linear response regime where a modulation of the optical power always corresponds to the same proportional frequency response. Curiously, we see that the transfer function increases in magnitude with lower frequency for a single curve. Calculations in Fig.5.3 indicate that we should expect a flat frequency response until the spacer contribution becomes significant. It appears possible that the calculation of the




Figure 5.8: **Left:** Measured photothermal transfer function of the cavity response when modulating the transmitted power while locked to a single polarization mode of Si6. Data is taken at an operating temperature of 4 K. **Right:** Renormalizing the left set of curves by $\sqrt{P_{transmitted}}$.

spacer photothermal transfer function could be underestimated, but we see this is not the case as we remeasure the transfer function for different transmitted powers (additionally the frequency dependence of the data does not match the $1/f$ scaling expected from theory). As we increase the optical power the overall amplitude of the TF decreases, though it maintains the same qualitative shape. The observed dependence on the optical power P_0 is again unexpected.

One possibility is that local heating of the coating is significantly changing some material parameters such that the coating response increases. Without a theory that captures the frequency dependence of our observed response, it is difficult to identify coupling of the optical power to the correct material parameters. The response has an inverse dependence with thermal conductivity, increases in optical power could heat the coating and thereby increase k , which would result in a lower overall coating response. Increases in coating heat capacity C_c and coefficient of thermal expansion α_c would have the opposite effect and increase the magnitude of the photothermal transfer function. For this to be the dominant mechanism, thermal conductivity must have a stronger temperature dependence than the other two material parameters. Low temperature phonon models approximate a similar T^3 dependence for both k_c and C_c , so it is unclear how likely this may be.

We can summarize our findings into three key points:

- We measure a photothermal transfer function that changes magnitude depending on the

power P_0 transmitted through the cavity.

- All measured transfer functions share the same approximate $f^{-1/3}$ frequency dependence
- The magnitude of the transfer function is the same for both polarizations, though has opposite phase for the two.

Given that none of these behaviours were observed in Si4, this strongly suggests that a new mechanism is responsible for the observed photothermal response of the AlGaAs coating. All these observations complement our previous measurements of the cavity's photothermal transient response. A crucial point is that we now have a more quantitative measurement of the optical power - frequency coupling, and we can try to heuristically extract some relation between these parameters. In the right panel of Fig.5.8, we renormalize the measured transfer functions by the root of the optical power transmitted through the cavity for those measurements, P_0 . There is no theoretical justification for this, instead we simply notice that this appears to reveal a universal behaviour in the frequency response. To rephrase this, it appears that our observed photothermal transfer function should have an inverse scaling with the optical power of the cavity as

$$\delta\tilde{X}_{measured} \propto \frac{1}{\sqrt{P_{transmitted}}} \quad (5.15)$$

With these being the first measurements made of photothermal effects in crystalline coatings, it is difficult to ensure that this is a reproducible effect and not marred by an anomaly such as a large defect sampled by the optical mode. Fortunately, a similar set of data has been measured on Si5 at PTB. Comparing results allows us to form a picture of the mechanism affecting the transfer function by identifying the variables that might couple to it.

First, we need to remove the confounding effects due to the different spacer lengths of the two cavities. We convert the transfer function of frequency response into length response so that we only consider photothermally induced length changes of the coating. We next treat the Si5 data in a similar manner as the Si6 PT by rescaling all curves with $\sqrt{P_{transmitted}}$. A comparison to the Si6 measurements is seen in Fig.5.9. One of the measured Si5 curves remains displaced from the others,




Figure 5.9: Converting the photothermal transfer function from units of frequency deviation to length deviation allows us to directly compare the coating response of Si5 (stars and solid lines) and Si6 (circles and dotted lines) without complications from the different cavity lengths. Renormalizing by the transmitted power produces a similar set of transfer functions for the two systems, despite over an order of magnitude span in power.

but the rest all share a similar magnitude and frequency dependence. There may easily be other variables that weakly couple to these transfer functions, but it is remarkable that we only need one experimental parameter to find a universal trend between different datasets on two different systems that share a similar coating (albeit at different temperatures).

The presence of this coating response means that we need to be mindful of laser power fluctuations. For cavities with dielectric coatings, a common method for mitigating photothermal effects is to simply reduce the optical power so that the product of the photothermal transfer function and intensity noise is below the fundamental cavity instability. With the crystalline coatings, we see that the magnitude of the transfer function increases as we reduce the optical power. Intensity noise will scale with P_t while the transfer function scales as $P_t^{-1/2}$, so this method still works in principle but its efficacy is reduced. A better practice is to stabilize the optical power in transmission so that the resulting intensity fluctuations produce a negligible frequency shift. This is the approach

that we enact on Si6.

One point which we have not discussed is the polarization dependence of the photothermal response. We find that

$$\delta\tilde{X}_{meas}^{E1} = -\delta\tilde{X}_{meas}^{E2} \quad (5.16)$$

One could imagine that schemes which simultaneously excite both polarization modes of the cavity would have a significantly suppressed photothermal response. The residual effect should be limited by the coating elastic and thermo-refractive perturbations which have an amplitude independent of the light polarization. In the next section, we will discuss methods that enable us to probe the cavity in this manner.

5.2 Frequency instability of the cavity

Our photothermal measurements from the previous section were the first indication that there may be novel noise mechanisms associated with the use of crystalline coatings in the cavity. So far we have only considered responses of the cavity frequency to changes in laser power. Now, we turn our eye to the steady state and study the frequency noise measured on Si6. The majority of measurements detailed in this section are the subject of [63].

Our first measurements of the cavity noise returned instabilities over 20 dB higher than the expected thermal noise limit, with $\text{mod } \sigma_y \sim 2 \times 10^{-16}$. Technical noise sources were suppressed well below this level, so there was no clear indication why the cavity stability was so poor. Consistently remeasuring the cavity noise budget did not yield any new insights. While it was tempting to speculate whether this could in fact be the true thermal limit, it would require the mechanical loss angle of the coating to be one hundred times worse at cryogenic temperatures than the measured room temperature value. Another point of speculation was whether additional loss could arise from the substrate-coating interface due to the CTE mismatch and the transfer from the original GaAs growth wafer.

At PTB, Si5 had already been operating at cryogenic temperatures for nearly a year and the cavity instability was also above the expected thermal noise floor, though it realized a much lower limit of $\text{mod } \sigma_y = 4.5 \times 10^{-17}$. We coincidentally noticed that this instability was lower than that of Si6 by nearly the ratio of the two cavity lengths, suggesting that some coating effect could be in play.

A critical moment was the observation that cavity's frequency noise would vary with the level of power transmitted through the cavity. Increasing the transmitted power would result in a highly repeatable increase in frequency noise, and reducing it would return it to its previous instability. We ruled out any possible saturation effects and contributions due to limited signal to noise typically scale in the opposite direction, where increasing the cavity power and discrimination slope result in a lower electronic noise floor. Varying the modulation depth of the EOM can directly

change the discrimination slope for a constant carrier power. Furthermore, if the carrier power was kept constant, varying the optical power stored in the sidebands had no effect on the observed cavity noise. There is additionally no settling time that we observe with this frequency noise, instantaneous reduction of power appears to instantly reduce the cavity instability. From these observations, we determine that the frequency noise measured on Si6 is not a transient effect, it only cares about the steady state rate of photons entering the cavity.




Figure 5.10: Measurements of the cavity frequency noise for different transmitted powers, ranging from a minimum of 90 nW to 1000 nW transmitted. When this data was taken, optical scatter off one of the input vacuum optics contributed a large noise hump at 100 mHz offset. Cavity noise at frequencies below this offset is limited by this technical source. Below a cavity transmission of 90 nW, electronic noise sources start to become significant and no further reduction in frequency noise is observed. A dotted red line provides a $1/f^{1.5}$ guide to eye, well above the projected Si6 thermal noise (black line).

This was reminiscent of our measurements of the photothermal transfer function, but now the steady state frequency noise we observe has an optical power dependence opposite of what we measured when modulating the intensity. In the latter measurement, increasing the optical power resulted in a reduced photothermal coupling. With the steady state noise, increasing the coupled cavity power results in a higher instability. We can additionally confirm that we are not

limited by photothermal-induced noise by reducing intensity fluctuations via feedback, and seeing no corresponding reduction in frequency noise.

Thus it appears that we are observing a different coating mechanism that now produces extra noise in the steady state cavity response. This clear trend that we observe suggests a strict monotonic relationship between frequency noise and cavity coupled power. Again, we try to rescale the sets of curves in Fig.5.10 to extract the dependency of S_y on transmitted power P_t . Surprisingly, we again find a square root P_t dependence such that

$$S_y \propto \sqrt{P_t} \quad (5.17)$$

This empirical relation suggests a potential mitigation method: significantly lowering the optical power might allow us to reduce these power dependent frequency fluctuations to a level below the thermal noise limit. Unfortunately, an estimation of the necessary optical power proves that this is an unfeasible strategy. At a transmitted power of 100 nW, we measure a lower bound of $S_y = 2 \times 10^{-32}/f$. Supposing we want to reduce this to a level comparable to the projected thermal noise floor of $S_{\text{Brown}} = 1.7 \times 10^{-34}/f$, we would need a 10^4 reduction in optical power, or a transmitted power of 10 pW.




Figure 5.11: **Left:** Another set of frequency noise measurements taken with different intracavity powers. **Right:** Curves on the left are normalized by $\sqrt{P_c}$ and collapse to a single noise limit despite the factor of 10 difference in optical power. A dotted red line is displayed as a $1/f^{1.5}$ guide to the eye

Detection and signal to noise restrictions aside, the shot noise limited length resolution of 10

pW corresponds to a white noise floor of $S_y \approx 10^{-30}$. Only at bandwidths below $100 \mu\text{Hz}$ will the thermal noise limit exceed the shot noise limited resolution and this is clearly an impractical length of measurement time. This strategy is also unfeasible as it is clear from Fig.5.10 and Fig.5.11 that the cavity noise is not $1/f$. We can identify a lower $1/f$ bound which the noise never drops below, but it appears to instead have a slope $S_y \sim 1/f^{1.5}$. This steeper slope indicates that the cavity instability degrades with longer averaging times, so reducing the measurement bandwidth reduces the relative thermal noise contribution.




Figure 5.12: Light-induced noise at a cavity temperature of 4 K (solid lines) and 16 K (dotted lines). Curves of same color are measured with same intracavity power. We do not measure any changes in coating loss and transmission between the two temperatures.

If this noise originates from the fluctuation-dissipation relation, we would expect to see a strong temperature dependence. For the Brownian thermal noise limit scales with temperature, a similar process driving the Si6 noise could also serve as a means for managing it. Unfortunately, it is not an option to lower the temperature of the Si6 cavity below 4K. The 2.5 K stage 2 temperature of the coldhead places a limit on achievable base temperatures, but it is trivial to heat the cavity up. Less than a watt of power on the 4 K active shield is needed to raise the cavity temperature to 16 K, where silicon exhibits zero crossing of the CTE. Comparing this light-induced noise of

the cavity between 4 K and 16 K should provide a large signal to discriminate any temperature dependence.

In Fig.5.12, we compare noise measurements at the two different operating temperatures at identical intracavity powers. To our surprise, we cannot discern a difference, suggesting that in this temperature range light induced noise has no strong temperature dependence. It is possible that there is no explicit dependence on temperature, but rather that the magnitude of the noise is set by some material parameter which has a weak relation with T (the thermo-refractive coating transfer function for example depends on dn/dT). Regardless, our current understanding of fundamental cavity noise terms originates with fluctuation-dissipation, where noise sources are thermodynamically driven. The peculiarity of this light induced noise cannot be overstated as we do not observe the typical $\propto T$ or $\propto T^2$ scaling that is characteristic to known dissipation-induced processes.

Again, as a consistency check, we can compare our results to another system to ensure that this is a reproducible feature. The same effect was measured on Si5, though at a much lower magnitude in fractional units. The fractional frequency light induced noise was measured to be over 10 dB lower (Fig.5.13) than what we measured in Si6. We note that the absolute length noise S_x is the proper quantity to compare, since the noise originates in the coating and the longer spacer length of the 21 cm cavity converts this into a smaller fractional frequency noise. As with fundamental noise mechanisms, this indicates that even longer cavities could advantageously reduce the contribution of this noise mechanism, though technical issues become significant over spacer lengths of 10s of cm. Rescaling the length noise in the bottom right panel of Fig.5.13 brings the two systems into much closer agreement, though there is still a slight discrepancy between the Si5 and Si6 sets of curves. It is reasonable to surmise whether this could be a residual direct temperature dependence, but this would be an odd behaviour. Interpreted in this manner the data would suggest that lower temperature results in an increased noise, contrary to our intuition. Furthermore, the lack of any difference between data at 4 K and 16 K suggests that this is not a typical thermodynamically driven mechanism.




Figure 5.13: **Top:** Light-induced noise measured on Si5 at 124 K for different intracavity powers. **Bottom left:** Comparison fractional frequency noise of Si5 and Si6 with the $\sqrt{P_c}$ rescaling implemented. The light induced noise in Si5 is over 10 dB lower than that in Si6 due to the longer cavity length. **Bottom right:** Renormalizing by the cavity lengths allows us to compare the absolute coating noise between the two system.

It is more likely that there are other parameter differences between Si5 and Si6 that we have not yet considered. The crystalline coatings were grown in the same run, so it seems improbable that stresses within the AlGaAs and GaAs could be any different. The similar static birefringence of the multilayers corroborates this, so the only remaining disparity is the beam waist supported by the cavity. Whereas Si6 uses two mirrors with 1 m radius of curvature, Si5 has a planar mirror on the input end to the cavity. In materials with large electro-optic effects where the optical index varies with applied electric field, intensity - not power - is the proper variable to describe the material response. It stands that a similar relation may apply here, and that we should instead be normalizing the noise curves by intensity. Therefore we need to verify that there is in fact a scaling with mode area.

5.2.1 Mode area dependence

Once the cavity is constructed and the thermal environment is kept constant, the only method for reducing the fundamental Brownian noise limit is through reducing the coating thermal noise contribution. As demonstrated earlier, coating thermal noise has a $1/w_0^2$ beam size dependence. The beam waist itself cannot be modified once the mirrors are chosen to support a particular radius of curvature, but higher order spatial modes of the cavity can be interrogated to reduce the effective beam size. Recalling that thermal noise is calculated via the fluctuation-dissipation theorem as

$$S_{Brown} = \frac{2k_B T}{\pi^2 f^2} \frac{W_{diss}}{F_0} \quad (5.18)$$

The strain energy dissipated across the coating W_{diss} follows the pressure profile of the beam's amplitude distribution. This will be none other than the spatial profile of the beam which can be described in either the cylindrical Laguerre-Gauss basis, or the cartesian Hermite-Gauss modes. By calculating the effective size increase for a fixed optical power, a merit factor $g_{m,n}$ can be defined to modify the strain energy in the form

$$W_{diss} \propto \frac{g_{m,n}}{w_0^2} \quad (5.19)$$

The merit factor will depend on the mode indices m, n for HG modes, or cylindrical indices ρ, l for LG modes and will have a value less than unity for any non-fundamental modes. For the gaussian 00 mode, $g_{0,0} = 1$ so that $W_{diss}^{m,n \neq 0} < W_{diss}^{0,0}$ and the thermal noise limit is reduced as the mode index increases. Formalism for calculating the merit factor is detailed in [130], and the merit factor differs for coating Brownian and substrate Brownian noises. We only consider the merit factor for coating thermal noise, given that we are only interested in effects originating within the crystalline coating. For coatings, one finds

$$g_{m,n} = \frac{4}{\pi} \int_0^\infty dp \int_0^\infty dq e^{-p^2+q^2} (L_m(p^2)L_n(q^2))^2 \quad (5.20)$$

for the HG basis, where the various $L_m(x)$ are the ordinary Laguerre polynomials of the n -th

order. In the Laguerre-Gauss basis,

$$G_{0,\rho,l} = 2\sqrt{\frac{2}{\pi}} \int_0^\infty dx e^{-2x^2} L_l(x^2)^2 L_{\rho+l}(x^2)^2 \quad (5.21)$$

$$G_{\rho,l} = 2\sqrt{\frac{2}{\pi}} \left(\frac{\rho!}{(\rho+l)!} \right)^2 \int_0^\infty dx e^{-2x^2} x^{4l} L_r^l(x^2)^4 \quad (5.22)$$

For modes with no angular momentum $l = 0$,

$$g_{\rho,l=0} = G_{0,\rho,l} \quad (5.23)$$

For higher angular momentum modes,

$$g_{\rho,l>0} = G_{0,\rho,l} + \frac{1}{2}G_{\rho,l} \quad (5.24)$$




Figure 5.14: Merit factor $g_{m,n}$ for different spatial modes in the rectangular basis is plotted on the top, and $g_{\rho,l}$ for the cylindrical basis is plotted for the bottom. Rectangular modes $HG_{0,i}$ and $HG_{i,0}$ have the same scaling factor due to separability along the two axes. This is not the case for cylindrical modes where increasing the azimuthal mode index leads to a larger effective mode area.

Values of the merit factors for the two bases are displayed in Fig.5.14. In the rectangular basis, separability of the two mode indices implies that $g_{m,n} = g_{n,m}$. A sizeable 25% reduction in thermal noise can be obtained by simply interrogating the first non-fundamental mode, HG_{01} or HG_{10} . This is typically also in the regime where reasonable coupling to a higher order mode can

be achieved simply by adjusting the beam pointing into the cavity. Higher order modes suffer from reduced coupling efficiency, and offer marginal returns in reducing the coating thermal noise. High coupling can be regained by using a mode converter, but alignment of a phase plate is nontrivial for rectangular modes with index larger than 2. One can instead use vortex phase plates to generate LG modes with angular momentum. For mode indices larger than 1, LG modes offer a larger reduction in coating noise when compared to HG modes and the spiral plate also relaxes constraints on alignment. A comparison of phase plates for the HG_{11} and $LG_{3,0}$ mode is displayed in Fig.5.15.




Figure 5.15: Generation of higher order modes with phase plates. The top row shows the phase plate (middle) needed for converting the fundamental mode (left) into the HG_{11} mode (intensity profile on the right). Yellow regions of the mode converter imprint an additional π phase shift when compared to the purple regions. The bottom row shows a similar procedure for generating the LG_{03} mode.

To test mode dependence of the light induced noise, we considered a scheme of generating the LG_{03} with a phase plate. This would in principle easily reveal mode area dependence of birefringent noise since we would be looking for a 50% reduction in S_{meas} . We never ended up implementing this since we had fears that the strong birefringence of the cavity might break the cylindrical symmetry of spatial modes, and the cavity might only support the HG basis. The easiest method would be to couple to the HG_{01} mode by adjusting the pointing, so we opted for this route.

We'll provide a more thorough description in the next section, but our intuitive understanding of the light induced effects paints a picture where the frequency noise is anti-correlated between




Figure 5.16: When using higher order modes, sensitivity to clipping is increased. Not only is the mode area greater, beam size will be larger at all locations along the propagation axis. Here, we compare the beam size for the LG_{00} , LG_{20} , and LG_{02} modes at various locations before the cavity. Locations of various optical apertures are shown as dashed lines at their corresponding distances from the cavity input mirror. A pink dot indicates the size of the 4K active window at its location 9 cm from the cavity. The significantly larger beam radius of the LG_{02} mode makes it much more susceptible to clipping on this optic.

the two birefringent modes of the cavity. We'll consequently call this measured noise "birefringent noise". Given that it significantly dwarfs our technical sources and the projected Brownian thermal noise limit, we can look for dependencies on mode area by directly locking to a single polarization of the cavity. We expect that our measured noise S_{meas} mainly consists of the birefringent noise S_{biref} .

By looking at the corresponding ratio of light induced noise of the two modes $S_{\text{biref}}^{01}/S_{\text{biref}}^{00}$, we can identify any dependence on mode area. Recall that for the higher HG_{01} mode, the scaling factor is calculated as $g_{01} = 0.75$. A direct dependence on the mode area w_0^2 would therefore return

$$\frac{S_{\text{biref}}^{01}}{S_{\text{biref}}^{00}} = \frac{g_{01}}{g_{00}} = \frac{0.75}{1} \quad (5.25)$$

In Fig.5.17 we measure the birefringent noise of the HG_{00} and HG_{01} modes at the same intracavity power and temperature by interrogating a single polarization in the Si6 system. We find a ratio of 0.71 ± 0.02 , similar to the expectation of $g_{01}/g_{00} = 0.75$ for birefringent noise dependent on mode area w_0^2 . Varying the optical power returns consistent results, suggesting that the birefringent noise scales with w_0^2 . We can then return to our traces from Fig.5.13 and rescale them by the corresponding mode areas for the two resonators ($450\mu\text{m}$ for Si5, $300\mu\text{m}$ for Si6). Shown in Fig.5.18,




Figure 5.17: Mode area effect on birefringent noise. We separately measure the birefringent noise of the HG_{00} and HG_{01} modes by addressing a single polarization of the cavity. The ratio of birefringent noise S_{biref}^{01} and S_{biref}^{00} for the two modes reveals the dependence on mode area. If the birefringent noise is independent of the mode area, we expect a ratio of 1. We find a ratio of 0.71, consistent with the expectation of 0.75 if the noise scales with mode area $g_{mn}w_0^2$. A large peak visible at 0.057 mHz is the result of a ground loop that appeared in the S_{biref}^{01} dataset and was omitted in the fit.

both sets of curves appear to be in better agreement.




Figure 5.18: Data from the bottom right of Fig.5.13 is now divided by the mode area of the Si5 and Si6 cavity.

5.2.2 Correlation of light induced noise

For fundamental noise mechanisms, identifying the parameters that set these limits is necessary for devising strategies to reduce their contributions. Both systems see this noise as a limitation, but reducing the intracavity intensity is not practical for either. Ideally we would hope to suppress this effect entirely, but that requires us to coherently cancel this with an identical copy of the noise. We see that this is possible with thermal dissipation noise sources, where thermo-refractive and thermo-elastic terms are engineered to coherently cancel and reduce the overall thermo-optic noise [75, 131]. In this case a second physically distinct mechanism is needed, and our poor understanding of the anomalously high frequency noise in Si6 is already barrier.

However, another distinct noise process is not necessary for coherent cancellation. A clue to the nature of the crystalline coating noise lies in our previous set of photothermal measurements. From the coherent intensity modulation and the transient frequency responses, we saw that the magnitude of the photothermal transfer function was the same for both birefringent modes. However the amplitude was opposite, suggesting that averaging the two could suppress that photothermal effect. It provokes the question of whether the same effect could be present in the steady state noise. The frequency noise we measure for the two polarization modes is identical, but the measurement does not discriminate the relative phase between them. If the birefringent length fluctuations are out of phase, simultaneously probing both modes should suppress this noise. To be explicit, our hypothesis is that length fluctuations of the two modes look like

$$d_{tot}^{E_1}(t) = d_{\text{Brown}}(t) + d_{\text{residual}}(t) + d_{\text{biref}}(t) \quad (5.26)$$

$$d_{tot}^{E_2}(t) = d_{\text{Brown}}(t) + d_{\text{residual}}(t) - d_{\text{biref}}(t) \quad (5.27)$$

We assign a term to length fluctuations originating from different sources: d_{Brown} encapsulates the length noise from the total Brownian thermal noise limit, and technical noise sources (characterized and uncharacterized) are contained in d_{residual} . Both modes share these two common terms. A third, d_{biref} , refers to the birefringent noise that is associated with the orthogonal polarization of the two modes. Notably the magnitude of this term is the same in both expressions, but the amplitude is

opposite between them. Our noise characterization of the cavity suggests that the technical sources are suppressed to a level near the thermal noise, so we expect that

$$d_{tot}^{E_1}(t) \sim d_{\text{biref}}(t) \quad (5.28)$$

$$d_{tot}^{E_2}(t) \sim -d_{\text{biref}}(t) \quad (5.29)$$

If this were true, a simultaneous time series of the length noise $d_{tot}^{E_1}(t)$ and $d_{tot}^{E_2}(t)$ would appear to be largely anticorrelated. Asynchronous measurements would not show this effect, and since phase information is lost when transforming to the frequency domain, measurements of S_y for the two modes would show an identical result:

$$S_y^{E_1} = S_y^{E_2} = S_{\text{Brown}} + S_{\text{residual}} + S_{\text{biref}} \quad (5.30)$$

Only by simultaneously addressing both modes, we would see

$$d_{tot}^{E_1 E_2}(t) = d_{\text{Brown}}(t) + d_{\text{residual}}(t) + d_{\text{biref}}(t) - d_{\text{biref}}(t) \quad (5.31)$$

$$= d_{\text{Brown}}(t) + d_{\text{residual}}(t) \quad (5.32)$$

And correspondingly

$$S_y^{E_1 E_2} = S_{\text{Brown}} + S_{\text{residual}} \quad (5.33)$$

To test this, we need to develop a scheme that can simultaneously interrogate both modes to cancel the birefringent noise.

5.2.3 Dual frequency locking

PTB came up with a unique method that involves two optical setups. One probes the cavity from the top, another from the bottom, and the two are locked on orthogonal modes. A signal stabilized to the sum of the two error signals could be synthesized digitally and fed back to a separate actuator. This proved to be a highly versatile setup, but requires two lasers, two sets of locking and detection electronics, and also requires the two setups to have their individual noise




Figure 5.19: **(a)** Optical layout for probing the Si6 system with active RAM suppression at a modulation frequency (5 MHz) where the cavity error signal is encoded in a traditional PDH locking setup. Dual-tone probing is achieved by driving AOM2 with two RF tones separated by the birefringent mode splitting of Si6 (green and purple sets of tones). A half waveplate before the cavity is aligned to couple an equal mixture of s- and p-polarized light to the two polarization modes of the cavity, labeled $\hat{E}_1$ and $\hat{E}_2$. The collinear beams are then launched upwards with a 45° mirror to couple into the cavity. The reflected cavity error signal has equal contributions from noise of the two polarization modes, and any anti-correlated noise is automatically rejected. **(b)** Frequency landscape of the tones applied to the Si6 cavity. The two tones applied to AOM2 separate the carrier and the two sidebands each into two tones separated by the birefringent mode splitting of the cavity. The purple tones shown address one polarization mode ($\hat{E}_1$) while the green tones simultaneously address the orthogonal mode ($\hat{E}_2$).

budgets suppressed below the level of the birefringent noise. This is an intensive task, so we would like to enact a solution that has less redundancy built into it.

The principle of our desired method is similar to PTB's approach - we can treat the birefringent modes as two distinct cavities, and we want a signal corresponding to their center-of-mass

motion. Our chosen approach is illustrated in Fig.5.19. The optical elements AOM1 and the EOM are used to enact traditional PDH feedback, the EOM generating modulation tones at our typical sideband offset of 5 MHz. A second AOM (AOM2) generates a doublet (two tones generated via DDS) with a splitting between the tones corresponding to the birefringent mode splitting of roughly 770 kHz. These two tones are used to address both modes of the cavity simultaneously in the frequency domain. However, the birefringent modes will only transmit orthogonal polarizations of light, so a half-waveplate after the Faraday element is used to rotate the in-plane p -polarization to project equally along the polarization axes $\tilde{E}_1$ and $\tilde{E}_2$ of the cavity.

A photodiode in cavity transmission can be used to monitor the sum power transmitted for both polarizations, and feedback to the sum (note that this method does not allow us to identify differential amplitude noise of the two tones). In cavity reflection, sidebands and cavity leakage signals along the $\tilde{E}_1$ and $\tilde{E}_2$ and polarizations are recombined on the half waveplate to an identical projection, and then rotated to s -polarization after transmission through the Faraday. Light detected on the PDH photodetector contains the sum of error signals relating to the two distinct modes. If the modes are identical (i.e. we space the two tones by a value corresponding to the cavity free spectral range), we simply recover an error signal twice as large as we would for a single cavity. However if there is any anti-correlated noise between the two modes, we optically average this effect in our proposed scheme. We collect error signals of the two modes corresponding to

$$V_{E_1}(t) = V_{\text{Brown}}(t) + V_{\text{residual}}(t) + V_{\text{biref}}(t) \quad (5.34)$$

$$V_{E_2}(t) = V_{\text{Brown}}(t) + V_{\text{residual}}(t) - V_{\text{biref}}(t) \quad (5.35)$$

So that

$$V_{\text{tot}} = V_{\text{Brown}}(t) + V_{\text{residual}}(t) \quad (5.36)$$

Note that this is identical to the condition from Eq.5.32. There are a few caveats to consider.

First, for the quantity $V_{\text{biref}}(t)$ to be identical in the two equations above, we know that the intracavity power within each birefringent mode needs to be identical. Second, the discrimination

slope for each mode needs to also be similar, i.e. cavity couplings and sideband powers for the two modes must also be the same. Violation of the first condition generates a non-common level of birefringent noise for the two modes, and violation of the second leads to poor rejection of S_{biref} in the optical cancellation.

The advantage of this method is that optical, mechanical, and thermal noise sources are common to both modes. If the half waveplate is set to couple an equal mixture of $\tilde{E}_1$ and $\tilde{E}_2$, cavity couplings will necessarily be the same. Amplitude modulation due to etalon effects will be nearly identical, especially if the etalon free spectral range is significantly larger than the birefringent mode splitting (generally the case). Thus, the technical noise suppression is practically free of extra effort when compared to the dual optical setup scheme.

A disadvantage is that there is no method to directly extract the birefringent noise signal. The optical averaging automatically cancels the antiphase birefringent contributions. One could demodulate the PDH signal at the 770 kHz mode splitting in an attempt to recover the birefringent noise, but we will demonstrate in a later section that this is not in fact possible. A more practical issue is that our scheme generates an additional shot noise contribution due to the fact that a nonresonant tone of light is encoded in each polarization. The electric field projected along $\tilde{E}_1$ sees a copy of the PDH modulation triplet displaced by the mode splitting in the frequency domain, so the effect cavity contrast for each mode is reduced significantly. Avoiding this issue would require us to generate the two polarizations in distinct optical paths and then prepare them in the frequency domain to be resonant with different modes. This negates the rejection of common mode technical noise that we value, so we make do with the reduced signal to noise.

Identifying the mode splitting that separates the dual tones applied to AOM2 can be done by locking to each mode and measuring the beat offset from Si3. An insitu method is also demonstrated in Fig.5.20. Suppose that the AOM2 tones are offset from the actual mode splitting f_{Biref} by the quantity 2Δ and that dual frequency locking is used to lock to the center of the two polarization modes. We can remove one of the tones applied to AOM2 so that the laser is locked to only the E1 mode or the E2 mode. Depending on the mode, the laser frequency will be shifted by $\pm\Delta$ nearly

instantaneously limited by the locking servo's bandwidth. The splitting applied to AOM2 can be adjusted to minimize the frequency jump Δ . A slow frequency change with sign dependent on the intracavity polarization will be visible due to opposite parity of the E1 and E2 photothermal effects.

A natural question is to ask whether a static offset between the tones addressing E1 and E2 is sufficient, i.e. should this be dynamically varied to track the center of both cavity modes. With perfectly balanced phase error contributions from each mode this effect should be negligible, but we know this is not the reality. As mentioned earlier, our method of dual frequency locking does not allow us to access the difference channel for frequency noise between both modes. Without this information we do not have an error signal to allow us to feed back to the mode splitting. Fortunately PTB verified that this is unnecessary, as locking the two modes independently with different laser produces an identical instability as using the dual frequency locking.

5.2.4 Residual noise

Implementing the dual frequency lock provided us with an immediate improvement in the cavity instability. There were several notable features. First, the near 10 dB reduction of the measured frequency noise now displays a low frequency noise spectrum that is consistent with $1/f$ behaviour. There are no additional improvements we made to the technical noise budget indicating that the peculiar $1/f^{1.5}$ components observed on a single polarization mode must have been characteristic of the birefringent noise.

The birefringent noise power dependence of that we had clearly measured when addressing a single polarization is no longer evident in the residual noise. Varying the intracavity power over an order of magnitude consistently returns the same measurements of S_y .

Similar to our studies of the birefringent noise, we can examine the temperature dependence of the residual noise (S_{residual}) by operating the system at 4 K and the lower CTE zero crossing at 16 K. No change in residual noise was observed, indicating that we are unlikely to be limited by some form of thermally driven fluctuations.




Figure 5.20: **Left:** At $t = 0$, the laser is initially dual frequency locked to E1 and E2 modes. Around $t = 10$, the tone addressing E2 is removed so that the laser is only locked to E1. A frequency jump corresponding to Δ is observed. The E2 tone is then reintroduced, and the E1 tone is removed. The laser frequency, distributed after AOM1, is continuously monitored in a beat with Si3. **Right:** A cartoon indicating the location of the AOM2 tones relative to the two polarization modes. Colors correspond to the data on the left. With $\Delta = 0$, there is no frequency jump when switching from the dual mode lock to the single mode lock, indicating that the spacing of the AOM2 tones is equivalent to $f_{\text{Biref.}}$.

A large hump at 100 mHz in the measured frequency noise was identified to be related to some optical scatter within the vacuum. Varying the beam pointing into the cavity could significantly reduce the amplitude of this feature, despite the reduced cavity contrast. By fortune, we discovered that optimized coupling to the HG₀₁ mode via beam pointing would heavily suppress this scattering shelf. Our hypothesis was that some vacuum grease within the cryogenic environment had migrated to one of the 4 K windows. Evidence for this mechanism was observed in the Si4 system, several of the cryogenic windows were observed to have small specks of grease across the top surface. In this case, there was no significant impact observed in the cavity instability but we were not so lucky with Si6.

At the cost of reduced cavity coupling and a higher shot noise contribution, we measured an improved long term instability with the HG₀₁ mode. This test had the additional benefit of allowing us to identify any dependence of this residual noise on mode area. We detected no such difference between S_{residual}^{00} and S_{residual}^{01} , allowing us to conclude that the residual noise in Si6 did not share the same local properties that we measured with the birefringent noise. Whatever noise mechanism is currently limiting the cavity instability will therefore have a correlation length larger than the spot size since we cannot average any frequency fluctuations by increasing the area probed on the coating.

Together with the lack of dependence on spot size and temperature, we can conclude that the residual noise we measure is not the system's thermal noise limit. Though we have not directly measured the coating loss angle, we can place an upper bound on its magnitude at each temperature by equating a loss angle limit to our measured S_y .

In the frequency band of 0.01-0.3 Hz, we identify an instability consistent with $S_y = 3.37(9) \times 10^{-33}/f$ from a continuous 30,000 second measurement against Si3. This performance is slightly better than the thermal noise limited instability measured on Si4 of $S_y = 4.12(5) \times 10^{-33}/f$. Using the formalism of ?? for calculating the coating thermal noise limit, we therefore place constraints




Figure 5.21: Residual noise $S_{\text{residual}} + S_{\text{Brown}}$ after cancellation of the birefringent noise at both 4 K and 16 K shows no temperature dependence.

on the maximum value of the AlGaAs/GaAs coating loss angle ϕ :

$$\phi_{4.7 \text{ K}} < 4.1 \times 10^{-4} \quad (5.37)$$

$$\phi_{16.7 \text{ K}} < 1.2 \times 10^{-4} \quad (5.38)$$

Note that measurements from [18] extract a loss angle of $\text{SiO}_2/\text{Ta}_2\text{O}_5$

$$\phi_{4 \text{ K}} < 5.6 \times 10^{-4} \quad (5.39)$$

$$\phi_{16 \text{ K}} < 3.2 \times 10^{-4} \quad (5.40)$$

These bounds already imply that these crystalline coatings can yield a sensitivity improvement over conventional dielectric coatings at 4 K, and a significant improvement at 16 K.

The origin of the residual noise is not clear. It is natural to ask whether this could be unsuppressed birefringent noise. However, we have already laid out a number of observed properties that are atypical of the birefringent fluctuations. Additional measurements where we introduce a slight imperfect cancellation of birefringence (i.e. improper matching of intracavity powers and discrimination slopes for the two modes) further confirms this, as the residual noise does not increase. From our mode area measurements, we also can conclude that the residual noise is not limited by local fluctuations of the optical coating. It remains possible that the global noise dominating S_y is an




Figure 5.22: Fractional frequency instability of Si6 of birefringent noise (red squares) and residual noise at 16 K (blue circles). The residual noise is measured over a continuous 24000 s dataset.

uncharacterized technical source that we have not yet identified.

5.3 Direct measurement of birefringent noise

We had the motivation to directly measure birefringent noise, independent of technical noise contributions. PTB developed a method for this involving two separate optical setups, one probing the cavity from the top and another from the bottom. Two separate lasers could be locked to the two orthogonal polarizations of the cavity and then independently beat with Si2. A simultaneous time series record of the frequency noise of the two modes then allows one to calculate all linear combinations with the noise of each mode. Noise between the two modes can be calculated directly, and technical noise sources common to both can be strongly suppressed. The power of this method is clear - high frequency information often swamped by vibration noise should now be only limited by optical shot noise. Similarly, thermal fluctuations that move the cavity frequency are common to both modes, so birefringent noise can be measured at temperatures away from the silicon zero crossing.

In principle, measurements made in this fashion on Si6 can identify the temperature depen-

dence of this noise from 4K to nearly 100K. Unfortunately, significant difficulties are associated with this dual optical setup scheme. Technical sources need to be suppressed below the level of the birefringent noise in both setups, and it is clear that we have already struggled with this on a single optical setup. Etalons present only in one setup for example, will appear in the birefringent noise channel when calculating the difference frequency of the lasers stabilized to the orthogonal modes. We therefore attempted a few measurement schemes that I will detail below. Ultimately all were unsuccessful, but they have some interesting ideas associated with them so they are worth recounting.

5.3.1 Two optical paths in one setup

The issue of measuring both polarization modes independently is not a problem unique to birefringence in crystalline coatings. A JILA experiment in 2000 demonstrated a measurement of mirror birefringence between two cavity modes by simultaneously coupling both polarization modes [10]. This, and recent proposals from the LIGO community [132] suggest that measurements of the mode splitting can place constraints on coupling strengths of axion-like dark matter by using the cavity to identify phase shifts between orthogonal light polarizations.

A schematic for this scheme taken from [10] is illustrated in Fig.5.23. While the beam paths share as many optical elements as possible, there is still room to pick up differential phase shifts. It would be prudent to feed back to each EOM for residual amplitude modulation accumulated along each path so a total of four detectors before the cavity are needed in this scheme. An interesting alternative is alluded to in the diagram, with different modulation tones at Ω_1 and Ω_2 fed to EOM1 and EOM2 respectively. Rather than spatially separate the two paths, information from the two cavity modes can be encoded in the frequency domain with different modulation for the two tones. Both tones can be recombined in fiber after EOM1 and EOM2 and outcoupled to share the same optical elements up to the cavity, where a half-waveplate can project the incident light along the two cavity axes. Similar to our current scheme of cancelling birefringent noise, this results in wasted light and higher photon shot noise as each polarization incident on the cavity contains an




Figure 5.23: The optical scheme of [10] splits the beam across two different paths that are prepared in frequency and polarization to address the birefringent modes. Beams are recombined on orthogonal ports of a polarizing beam splitter and coupled into the cavity. A Faraday rotator before the cavity swaps the polarizations of each incoming beams such that reflected light from the cavity fast axis is measured on the same path as incident light addressing the cavity slow axis. A heterodyne beatnote between the birefringent modes can be detected in cavity transmission after interfering the two polarizations.

off resonant tone. Crucially, this no longer locks lasers separately to each mode since the cavity reflected light reaches a photodetector shared for both polarizations. Error signals for the two polarizations are individually obtained by demodulating at their separate sideband frequencies and the sum signal can then yield a signal proportional to the birefringent noise. Though not an ideal method to measure the mode splitting signal, this will save some optics, two photodetectors, and some extra time aligning a second beampath.

We do not implement this on the Si6 system since the use of high-Q resonant photodetectors makes it very difficult to detect the two tones with high gain and low noise. Note that both tones have to be sufficiently spaced apart so that signals encoded on one tone cannot be detected on the other (e.g. in-loop error of the PDH servo).

5.3.2 Asynchronous detection of cavity error signals

The method above requires simultaneous addressing of the cavity with a second frequency tone of opposite polarization, resonant with the other polarization mode.

An analogy can be drawn to the atomic servo that steers a laser to unperturbed clock tran-

sition in an atom. The fundamental picture is identical - a magnetic field is applied to strontium atoms trapped in an optical lattice to define a quantization axis. Spin states in the ground 1S_0 and 3P_0 excited states are perturbed in frequency by the first order Zeeman effect:

$$\Delta\nu_{Z1} = -\frac{\mu_B\delta g}{h}m_FB \quad (5.41)$$

Opposite m_F spin states can be addressed such that the sum channel

$$\delta\nu_{Zeeman} = \Delta\nu_{Z1,m_F=-9/2} + \Delta\nu_{Z1,m_F=+9/2} + \dots \quad (5.42)$$

Is independent of magnetic field to first order. Higher order terms in B^2 persist, but these can be calculated and corrected for since the difference channel $\Delta\nu_{Z1,m_F=-9/2} - \Delta\nu_{Z1,m_F=+9/2}$ directly identifies the magnetic field experienced by the atoms.

Technical complications of simultaneously preparing and interrogating both stretched states $m_F = \pm 9/2$ prevents the experimenter from synchronously cancelling the first order Zeeman shift. Rather, atoms are prepared and interrogated in one spin state, detected, and then prepared and interrogated in the opposite spin state. Key to this method is the assumption that the magnetic field is not drastically drifting between these measurements on a timescale similar to the duty cycle of the experiment, i.e. the measurement is performed faster than the relevant quantity changes. If noise of the clock laser is significantly faster than the fluctuations in magnetic field, this asynchronous measurement is sufficiently fast to identify the measured frequency noise to originate from laser phase noise. The deadtime between these measurements still has consequences in the form of generating additional measurement noise via the Dick effect. The approximate effect of this is to introduce an artificial white noise floor in the power spectrum of the measured signal, and the amplitude of that noise floor will increase with longer deadtimes.

The metaphor between the atomic servo and the birefringent noise is straightforward. Instead of measuring magnetic field noise, we're interested in the noise associated with the splitting of the two cavity polarization modes. The difference channel of interrogating both modes returns this noise, while the sum channel cancels it and reveals the underlying noises. Similar to the atomic




Figure 5.24: **Left:** schematic of a proposed optical setup. Ideally, a switchable polarizing element could rotate the incident light to the fast or slow axis of the cavity in sync with the data acquisition rate. **Right:** A cartoon illustrating the concept of the asynchronous birefringence measurement. An AOM switches the laser frequency to alternate between addressing the two modes. The demodulated signal is recorded on digital multimeter and the corresponding phase error of each cavity mode can be reconstructed to yield the information we want.

servo, the protocol is to measure frequency noise of one polarization mode, jump the laser frequency and polarization to address the other one, measure its noise, and stitch the data together to obtain both measurement channels. We're interested in measuring birefringent noise on timescales of 1 Hz and below, so switching between modes at a 1 kHz rate should be sufficient. There will be a Dick-style aliasing effect but sampling these two polarizations quickly enough will give a faithful representation of the noise at low frequencies.

Maximized coupling to a single cavity polarization requires jumping both the laser frequency and polarization. The former can easily be done, a potential option for the latter is to use a liquid crystal waveplate to rotate the polarization. At visible wavelengths, the rise time can be in the millisecond range but at 1550 nm, all options unfortunately have switching times on the scale of 100 ms. A better option is to use a Pockels cell, which has a significantly higher switching rate but requires a kilovolt drive for a quarter wave phase shift. A proposal is illustrated in 5.24 where the switching element is placed directly before the cavity. We ultimately decided on using a static half waveplate to equally couple to both polarization modes and only jump the laser frequency to address a single mode.

Jumping the laser at a kHz rate to lock to the two modes has the unfortunate consequence that the laser servo will receive a large perturbation at this rate, and attempt to correct for it with a transient that depends on the servo bandwidth. Measurements made in that kHz bandwidth will therefore contain the contribution of the transient and will not be an accurate measure of the cavity frequency. Instead, we develop a method to passively probe the cavity frequency without any servos feeding back information related to it. The Si3 local oscillator has lower frequency noise than a laser locked to Si6 in the entire bandwidth of interest. From the perspective of the Si6 cavity, probing with Si3 stabilized light is akin to probing with a frequency delta function, or performing spectroscopy on Si6 by using Si3.

With the laser locked to Si3, there is no need to encode information from the cavity high in the RF. Clearly we don't want to do a pure homodyne readout of the cavity at DC, but since the phase lock to Si3 is independent of modulation frequency Ω the modulation sidebands can be




Figure 5.25: The Si6-Si3 heterodyne beatnote (blue) where two lasers are locked to the two separate cavities. Probing the Si6 cavity with a laser locked to Si3 recovers the Si6 noise spectrum with the Si3 noise included. This signal is measured in cavity reflection and then converted to frequency noise by calibrating the discrimination slope (pink). Both traces encode the PDH error signal at the same modulation frequency of 5 MHz.

lowered considerably below a MHz. Amplitude modulation picked up at the sideband frequency can be decreased as Ω is lowered. The only constraints are that the modulation needs to be significantly faster than polarization switching rate, and that the frequency is high enough such that power in sidebands doesn't transmit through the cavity.




Figure 5.26: Attenuation factor in transmission of power coupled into the cavity.

A minimum modulation frequency might be 10 kHz for a 1 kHz acquisition rate, but estimation of the required sideband power is a more complicated problem. With the power dependence of the birefringent noise, we want to operate in the ideal situation where nearly all of the power

is concentrated in the sidebands. Unfortunately, more of this is transmitted through the cavity as the modulation frequency is decreased. A worthwhile place to start is by identifying the minimum length noise we can resolve when limited by photon shot noise (in the limit of a weak carrier, strong sideband). Results of a calculation are shown in Fig.5.27.




Figure 5.27: **Left:** Shot noise limited length resolution in a 1 Hz bandwidth for the Si6 cavity. Reducing the optical power in the carrier will decrease the magnitude of the birefringent noise, but also hinder the ability to measure that noise at high frequency. Lowering the measurement bandwidth can mitigate this. **Right:** Johnson noise equivalent for a particular level of shot noise. If the detector noise floor is limited by Johnson noise and other op amp noise sources are below it, this yields the noise equivalent power.

Birefringent noise was measured to have a 1×10^{-16} fractional frequency amplitude at 1 Hz, and if we hope to resolve the noise at this level, the minimum power can detect is around 10 nW. The photodetector used cannot have a transimpedance gain too high, else the optical shot noise of our 10 nW will be hidden under electronic noise of the detector. Supposing we choose a transimpedance gain of 100 k Ω , the feedback resistor's Johnson noise of $40 \text{ nV}/\sqrt{\text{Hz}}$ needs to be pushed below the birefringent noise with a sufficiently large discriminator slope. With a fractional frequency birefringent noise of 2×10^{-32} at 1 Hz, we can calculate the necessary sideband optical power that will push the detector noise well below this quantity:

$$P_s = \left(\frac{\Delta\nu_c}{GP_c(2 \times 10^{-32})} \right) \cdot 40 \text{ nV} \quad (5.43)$$

For transimpedance gain G and cavity linewidth $\Delta\nu_c$. For our minimum choice of carrier power, this returns $P_s \approx 300 \text{ nW}$ so we choose a modulation frequency of 40 kHz to prevent more than a few nW from transmitting through the cavity.




Figure 5.28: Fractional frequency noise of the Si3-Si6 beatnote is computed with an evenly sampled time series of beat (black). The effect of aliasing the data, or introducing deadtime between measurements can be simulated by extracting a time series of every alternate point (e.g. all even timestamps), doubling the effective acquisition period to account for deadtime, and computing the power spectrum (blue). The low frequency noise of the beatnote is fully recovered, but an artificial white noise floor is evident at higher frequencies. Increasing the deadtime or sampling even fewer points has the same effect, but accentuates the magnitude of the white noise. In the asynchronous measurement, datapoints will alternate between addressing the fast and slow axes of the cavity. The degree of common mode rejection for noise present in both modes can be estimated by taking the difference of two signals aliased from the same dataset. In orange, we show the difference channel between an aliased signal composed of all even points of the Si3-Si6 timeseries and an aliased signal composed of all odd points. Low frequency noise in the difference is suppressed to the level of the white noise which is incoherent to the two datasets. Over 20 dB of common noise rejection is visible at the lower range of 100 mHz.

With these parameters, we can begin our study of the Si6 birefringent noise. A compilation of some of these measurements is displayed in Fig.5.30. It is worth noting that the phase lock to Si3 allows us to easily park the probing laser on various parts of the Si6 cavity lineshape. Sensitivity to different parts of the cavity error signal can be extracted by tuning this offset frequency around. In PDH schemes, the error signal is linearized at the point of maximized cavity transmission

meaning power oscillating at the modulation frequency will carry the largest signal to noise ratio. Conversely, the homodyne component will also contain information about the cavity frequency as the laser is detuned from the transmission maximum. A calculation of these different sensitivities in Fig.5.29 indicates that the probe-cavity detuning can in principle be used to extract the same information of the reflected cavity signal at AC and DC. In practice, homodyne measurement is significantly harder since just measuring optical shot noise of the off-resonant probe is a difficult endeavor below a few Hz.




Figure 5.29: Sensitivity to the AC components (blue) from the Pound-Drever-Hall modulation scheme is contrasted to sensitivity obtained from the DC components (green) via the side of fringe interrogation. At a frequency indicated by the dashed magenta line, AC and DC components should yield equivalent sensitivity to cavity noise. The cavity transmission function is overlaid in red.

As a demonstration, we detect cavity noise with our typical scheme of modulating the probe light at 5 MHz, where we can detect an AC PDH signal and simultaneously look at the DC component on the same photodetector. In the left panel of Fig.5.30, we compare the AC PDH signal taken with the probing laser parked at the peak of cavity resonance (red curve) to the DC signal taken with the laser detuned to half the cavity linewidth (indicated with dashed black line in Fig.5.29), where the cavity reflection function provides strong sensitivity but the PDH signal is very weak (blue curve). Comparable features of the cavity noise can be seen in both signals, though a significantly higher white noise floor in the DC signal obscures features with lower frequency noise. This is simply a result of detecting a DC signal while modulating the laser with 5 MHz sidebands,

which provide no frequency noise information at DC and only increase the observed optical shot noise.




Figure 5.30: **Left:** Cavity noise measurements from the AC (modulation at 5 MHz) and DC signals. The DC signal is slightly smaller in amplitude due to a smaller sensitivity to cavity noise at this particular detuning (at the cavity pole). **Right:** Detunings for the AC and DC measurements are now swapped, so that each measurement is insensitive to the amplitude noise converted by the cavity.

We can then reduce the modulation frequency to 40 kHz in accordance with our previously detailed scheme. We were unexpectedly unable to measure cavity noise in a manner similar to what is demonstrated in Fig.5.30 with the modulation at 5 MHz. With the probe off resonant from the cavity, we measure a level of noise consistent with shot noise. Upon tuning this to the cavity resonance, a significant amount of intensity noise was observed at a level higher than any of the noise features of the cavity. Lowering the modulation frequency seemed to have an unintended consequence of generating a significant amount of AM in the measurement. To ensure that this wasn't frequency noise generated in the phase lock to the Si3 light, we directly outcoupled light from the laser stabilized to Si3 and transferred it to the Si6 optical table over a fiber phase noise cancelled link. This has the benefit of removing any contributions of the Si6 laser servo since we are now directly using light from the Si3 stabilized laser.

The same level of excess AM was observed when measuring with the Si3 laser, indicating that it wasn't originating in our control loop. We noticed instead that it was reduced when increasing the gain of the Si3 PDH loop. Since Si3 achieves thermal noise limited performance, clearly the servo

noise was not a limitation so it was strange that a high level of amplitude modulation would only appear when probing the Si6 system on cavity resonance. Additionally, this was only noticeable when detecting in homodyne with the cavity reflected light, or in heterodyne with sidebands placed a few kHz away from the carrier. Modulating the Si3 probe light at a high MHz-level frequency would recover the typical cavity noise we measure when directly locking to Si6. We ultimately discarded this technique to move on to other measurements, but it is useful as a means to probe another system without requiring additional stabilization electronics.

5.3.3 Simultaneous detection

When we realized that we could cancel the birefringent noise with the dual mode - dual polarization scheme, a natural question to ask was whether we could measure the birefringent noise with a similar scheme. We again ended up spending some time to investigate this but realized that it still would not provide us with the same information gained from probing the orthogonal polarization modes from opposite sides of the cavity. Here we'll discuss the main points of the measurement.

In the previous section, we attempted to asynchronously measure the birefringent noise, but discovered that this generates additional white noise due to the Dick effect aliasing. However, this doesn't occur with birefringent cancellation since the two modes are probed simultaneously. Our thought was that in addition to detecting a PDH signal with the cancelled noise, we could also directly measure the birefringent noise by detecting the power oscillating at the 770 kHz mode splitting.

In principle, this would be as easy as taking the photodetector signal and splitting it into two components, one at the sideband modulation frequency and the other at the mode splitting. Unfortunately the narrow resonance of the photodetector's LC circuit provides very low transimpedance gain at frequencies over a MHz away from the PDH modulation. We instead split the signals in the optical domain, employing Terry's wizardry to build another resonant photodetector with a gain resonance near the mode splitting (Fig.5.31).




Figure 5.31: **Top:** Transimpedance gain of detector for measuring birefringent noise measured by Terry Brown. X-axis should be in kHz. **Bottom:** Current noise of the detector. We use the typical Hamamatsu G12180-003 InGaAs PIN with 1.1 A/W responsivity at 1550 nm. At 770 kHz, noise is equivalent to shot noise of 220 nW.

Again, we unfortunately spent some time trying to measure this signal before realizing it was a futile effort. We illustrate this by calculating the sensitivity to noise apparent at the mode splitting. First, consider the electric field incident upon the cavity when we use our dual polarization locking scheme.

$$E_{in} = E_0 \left(c_1 e^{i(\omega+\Delta)t} + c_2 e^{i(\omega-\Delta)t} \right) \hat{x} + E_0 \left(c_1 e^{i(\omega+\Delta)t} + c_2 e^{i(\omega-\Delta)t} \right) \hat{y} \quad (5.44)$$

The reflected field then contains information from the two cavity polarizations

$$E_{ref} = E_0 \left(c_1 R_1(\omega + \Delta) e^{i(\omega+\Delta)t} + c_2 R_1(\omega - \Delta) e^{i(\omega-\Delta)t} \right) \hat{x} \quad (5.45)$$

$$+ E_0 \left(c_1 R_2(\omega + \Delta) e^{i(\omega+\Delta)t} + c_2 R_2(\omega - \Delta) e^{i(\omega-\Delta)t} \right) \hat{y} \quad (5.46)$$

The signal detected at the photodetector signal will be

$$P_{ref} = E_0^2 \left[c_1^2 \left(|R_1(\omega + \Delta)|^2 + |R_1(\omega - \Delta)|^2 \right) \right] \quad (5.47)$$

$$+ c_2^2 \left(|R_2(\omega + \Delta)|^2 + |R_2(\omega - \Delta)|^2 \right) \quad (5.48)$$

$$+ c_1 c_2 \left(R_1(\omega + \Delta) R_1^*(\omega - \Delta) + R_2(\omega + \Delta) R_2^*(\omega - \Delta) \right) e^{2i\Delta t} \quad (5.49)$$

$$+ c_1 c_2 \left(R_1(\omega - \Delta) R_1^*(\omega + \Delta) + R_2(\omega - \Delta) R_2^*(\omega + \Delta) \right) e^{-2i\Delta t} \quad (5.50)$$

The signal oscillating at 2Δ is entirely even, meaning we unfortunately recover no sensitivity to the birefringent noise. This is a natural consequence of the fact that the information of the two cavity modes are encoded in different polarizations and will not interfere naturally at the photodetector. Interference between these modes is necessary. We can place a polarizer at a 45° angle to the $\hat{x}$ and $\hat{y}$ axes to define a new polarization basis $\hat{E}_1$ and $\hat{E}_2$ (rotated 45° from $\hat{x}, \hat{y}$). In this frame, the reflected field becomes

$$E_{ref} = \tilde{E}_0 \left(c_1 R_1(\omega + \Delta) e^{i(\omega + \Delta)t} + c_2 R_1(\omega - \Delta) e^{i(\omega - \Delta)t} \right) \hat{E}_1 \quad (5.51)$$

$$+ \tilde{E}_0 \left(c_1 R_1(\omega + \Delta) e^{i(\omega + \Delta)t} + c_2 R_1(\omega - \Delta) e^{i(\omega - \Delta)t} \right) \hat{E}_2 \quad (5.52)$$

$$+ \tilde{E}_0 \left(c_1 R_2(\omega + \Delta) e^{i(\omega + \Delta)t} + c_2 R_2(\omega - \Delta) e^{i(\omega - \Delta)t} \right) \hat{E}_1 \quad (5.53)$$

$$- \tilde{E}_0 \left(c_1 R_2(\omega + \Delta) e^{i(\omega + \Delta)t} + c_2 R_2(\omega - \Delta) e^{i(\omega - \Delta)t} \right) \hat{E}_2 \quad (5.54)$$

There are two scenarios for generating the electric field interference depending on how we orient the polarizer. First suppose it passes $\tilde{E}_1$. Then the transmitted field is

$$E_{ref} = \tilde{E}_0 \left[c_1 R_1(\omega + \Delta) + c_1 R_2(\omega + \Delta) \right] e^{i(\omega + \Delta)t} \hat{E}_1 \quad (5.55)$$

$$+ \tilde{E}_0 \left[c_2 R_1(\omega - \Delta) + c_2 R_2(\omega - \Delta) \right] e^{i(\omega - \Delta)t} \hat{E}_1 \quad (5.56)$$

Detected power is

$$P_{ref} = c_1^2 \tilde{E}_0^2 \left[R_1(\omega + \Delta) + R_2(\omega + \Delta) \right] \left[R_1^*(\omega + \Delta) + R_2^*(\omega + \Delta) \right] \quad (5.57)$$

$$+ c_2^2 \tilde{E}_0^2 \left[R_1(\omega - \Delta) + R_2(\omega - \Delta) \right] \left[R_1^*(\omega - \Delta) + R_2^*(\omega - \Delta) \right] \quad (5.58)$$

$$+ c_1 c_2 \tilde{E}_0^2 \left[R_1(\Omega + \Delta) R_1^*(\Omega - \Delta) + R_1(\Omega + \Delta) R_2^*(\Omega - \Delta) \right. \quad (5.59)$$

$$\left. + R_2(\Omega + \Delta) R_1^*(\Omega - \Delta) + R_2(\Omega + \Delta) R_2^*(\Omega - \Delta) \right] e^{2i\Delta t} \quad (5.60)$$

$$+ c_1 c_2 \tilde{E}_0^2 \left[R_1(\Omega - \Delta) R_1^*(\Omega + \Delta) + R_1(\Omega - \Delta) R_2^*(\Omega + \Delta) \right. \quad (5.61)$$

$$\left. + R_2(\Omega - \Delta) R_1^*(\Omega + \Delta) + R_2(\Omega - \Delta) R_2^*(\Omega + \Delta) \right] e^{-2i\Delta t} \quad (5.62)$$

Cross terms between $R_1(\omega)$ and $R_2(\omega)$ now appear at both AC and DC indicating interference generated between the two tones. Setting

$$\chi_{\hat{E}_1}(\omega) = R_1(\omega + \Delta) R_1^*(\omega - \Delta) + R_1(\omega + \Delta) R_2^*(\omega - \Delta) \quad (5.63)$$

$$+ R_2(\omega + \Delta) R_1^*(\omega - \Delta) + R_2(\omega + \Delta) R_2^*(\omega - \Delta) \quad (5.64)$$

We can extract the signal oscillating at the mode splitting frequency 2Δ and recover the familiar form of the error signal

$$P_{ref} = c_1 c_2 \tilde{E}_0^2 \left[\text{Re}[\chi(\omega)] \cos(\phi) + \text{Im}[\chi(\omega)] \sin(\phi) \right] \quad (5.65)$$

In the other scenario where the polarizer is oriented to transmit $\tilde{E}_2$, the reflected field changes to

$$E_{ref} = \tilde{E}_0 \left[c_1 R_1(\omega + \Delta) - c_1 R_2(\omega + \Delta) \right] e^{i(\omega + \Delta)t} \hat{E}_2 \quad (5.66)$$

$$+ \tilde{E}_0 \left[c_2 R_1(\omega - \Delta) - c_2 R_2(\omega - \Delta) \right] e^{i(\omega - \Delta)t} \hat{E}_2 \quad (5.67)$$

The equivalent $\chi(\omega)$ quantity will then appear as

$$\chi_{\hat{E}_2}(\omega) = R_1(\omega + \Delta) R_1^*(\omega - \Delta) - R_1(\omega + \Delta) R_2^*(\omega - \Delta) \quad (5.68)$$

$$- R_2(\omega + \Delta) R_1^*(\omega - \Delta) + R_2(\omega + \Delta) R_2^*(\omega - \Delta) \quad (5.69)$$

Plots of P_{ref} for polarizers transmitting $\hat{E}_1, \hat{E}_2$ with varying the laser-cavity detuning ω are displayed in Fig.5.32. Notably there is the presence of a large residual amplitude modulation that can be eliminated by choosing an optimal orientation, but no orientation will recover a dispersive signal that provides maximum sensitivity at zero detuning.




Figure 5.32: Laser-cavity error signals of demodulating the two tone electric field at the birefringent splitting 2Δ . Different curves indicated a different extraction phase ϕ . All curves are shifted horizontally for clarity though there is no vertical offset added. **Left:** Error signals when $\hat{E}_1$ is transmitted. **Right:** Error signals when $\hat{E}_2$ is transmitted. Since components from the cavity modes oscillate out of phase with each other, the large residual amplitude modulation is removed. However, the signal retains its symmetry about zero and is three times smaller. Detuning the laser by 2Δ to measure laser-cavity noise of a single mode no longer has any sensitivity when the polarizer is orientated to transmit $\hat{E}_2$. In both cases, sensitivity to laser noise can be obtained by detuning the probe to the side of both lines, similar to the side-of-line probing demonstrated in the previous section. Alternatively, locking to both modes simultaneously and reducing the splitting of polarization tones to $2\Delta - \Gamma_c$ will naturally shift the absorptive signal by $\Gamma_c/2$ where there is sensitivity to detuning.

These calculations illustrated in Fig.5.32 display the sensitivity to laser-cavity frequency noise when detecting at 770 kHz, but not the noise associated with the birefringence. To calculate that, we are instead interested in measuring the signal that appears as fluctuations of the mode splitting 2Δ . Fig.5.33 displays this by calculating the sensitivity to noise ξ that is anticorrelated for the two modes (i.e. the cavity transfer functions shift to $R_1(\omega + \Delta + \xi)$ and $R_1(\omega - \Delta - \xi)$). The upper row depicts the error signal when probing both polarization modes simultaneously for the two different choices in polarizer orientation. It is immediately apparent that the birefringent signal is not encoded at the splitting 2Δ . Similar to the upper set of curves in Fig.5.32 however, there is

sensitivity to the birefringent noise if we detune the laser by 2Δ , meaning only one of the cavity modes is addressed (bottom row of Fig.5.33). Technical noise affecting the cavity stability can still be a significant contribution to the measurement since we are no longer measuring differentially between the two polarizations. This same information can be obtained much more easily by simply locking to one of the polarization modes with a single frequency tone.

The power of the simultaneous birefringent measurement scheme was that any signal common to the instability of the two polarization modes would be suppressed. Unfortunately it does not appear that there is an easy way to enact this by playing clever tricks with modulation and detection. The optimal scheme may indeed require two separate lasers addressing individual modes, similar to what is proposed in [10]. Birefringent noise independent of common technical sources can then be easily reconstructed.

5.4 Long term instability

So far we've discussed our approach to dealing with the unique issues presented with the length noise of crystalline coatings. However, we were pleased to discover some positive attributes, namely a significantly smaller long term frequency drift. A simultaneous small, highly linear drift and good long term instability could make the system a promising component for an all optical timescale[17]. Our proof of principle study used Si3 as a local oscillator but constant refills of the liquid nitrogen dewar prevented this from operating as a hands-free system. Conversely, the closed cycle environment of Si6 allows us to run it fully autonomously. Choosing to operate Si6 at the 16 K zero crossing relaxes the stringent temperature stabilization environments we encounter on the 21 cm systems. So despite the difficulties of improving the short term instability due to the larger cavity linewidth, the 16 K operation offers some unique advantages for chasing record setting long term stabilities.

Preliminary measurements indicated that the Si3 drift of $-26 \mu\text{Hz/s}$ dwarfed that of Si6. Direct measurement of the Si3-Si6 beatnote required us to actually measure the cavity frequency against Si3 for several days in order to be able to clearly resolve the long term behaviour. New sets




Figure 5.33: **Left column:** Sensitivity to noise on the mode splitting when $\tilde{E}_1$ is transmitted. **Right column:** Sensitivity to noise on the mode splitting when $\tilde{E}_2$ is transmitted. **Top row:** Laser is tuned such that both modes are addressed simultaneously. **Bottom row:** Laser is detuned by 2Δ so that only one mode is addressed. We regain sensitivity to the birefringent noise with the laser resonant with a single mode, but this sensitivity depends strongly on whether $\tilde{E}_1$ or $\tilde{E}_2$ is transmitted.

of technical noise sources which were insignificant in our measurements of the residual noise now became important. These had to be minimized to identify a clear linear behaviour of the cavity frequency. Here, we'll detail some of the principle sources and the methods we used for reducing their impact.

5.4.1 Etalon identification

A significant hurdle to good long term stability of the cavity has been an etalon present in the system. Sudden changes to the thermal environment of the optical setup such as opening the primary enclosure would result in a slow, sinusoidal modulation of the cavity frequency. Eventually this would dissipate, though a diurnal modulation would often still be visible. The amplitude of

this fringe was a few Hz, and even one or two oscillations in a day would be enough to bury the signal of a smaller linear drift less than that of Si3.

It is not immediately obvious that this was an etalon, possible causes could be thermally driven beam pointing, or coupling of the cavity to room temperature. There are systematic ways to eliminate these theories, but the former was invalidated by realigning the entire optical setup several times with no reduction in the amplitude of the frequency modulation. For the latter case, we could measure the cavity coupling to room temperature by heating the vacuum chamber and measuring the corresponding frequency response. Multiplying this transfer function by room temperature fluctuations did not produce any signal coherent with the fringe that we were trying to eliminate on the cavity. Changing the cavity CTE by heating away from 4 or 16 K should also modify this transfer function, and we were unable to detect any difference in the fringe amplitude.




Figure 5.34: Si3-Si6 beatnote with krypton gas flowed into optics enclosure at different times. The blue segment highlights the steady state without any deliberate perturbation to the optical path length. A short burst of gas is released at the beginning of the orange segment. The flow is continuous for the duration of the green segment and then shut off at the end. A clear chirp in the cavity frequency is visible as the gas dissipates.

While trying to measure the sensitivity of the system to air pressure, we performed tests of flowing krypton gas into the optics enclosure. The heavier, inert gas should be less sensitive to ambient turbulence and a lower frequency noise would indicate that Si6 is limited by air pressure fluctuations. While we observed no change in cavity instability, we noticed that flowing the gas into

the Si6 box would drive the same fringe that we were trying to eliminate. This strongly suggests that an etalon was present in the system and that its length was being chirped as the larger index of refraction increased the optical path length.

If the frequency modulation was indeed an etalon, this signal should appear as an amplitude modulation between the carrier and modulation sidebands. We confirmed this in a long measurement of the out of loop RAM, and the next task was to identify the two reflecting surfaces supporting a parasitic cavity. As mentioned earlier, complete realignments of the system were unable to remove this effect so it was unlikely to be formed between two surfaces that we had easy access to. A possibility was that one of the surfaces could be within the cryogenic environment. Intuitively this would not be likely since all vacuum and cryo windows are mounted with a 6° wedge against gravity.

Identifying the troublesome optic is nontrivial since we cannot open the vacuum enclosure to modify its orientation. This is not necessary though, we simply need to deterministically vary the optical path length. This can be done by affecting the temperature of the optic and the nonzero CTE will generate a length change that should allow us to observe a frequency modulation. Recalling that we have heaters on the 300 K, 30 K, and 4K stages, we can step the temperature of each thermal shield and look for a corresponding change in the cavity frequency. Crucially, this change in frequency should occur within the time constant between various stages in order to be associated with the surface that was thermally stepped. For example if we change the setpoint of the 30 K radiation shield and the frequency modulation occurs 15,000 seconds later, the likely scenario is that the heat has diffused to the 4 K stages through the measured 12,000 time constant and is now perturbing the 4 K active window. Conversely if we measure an immediate frequency response, we can pinpoint the 30 K window as the problematic surface.

Iterating through the different thermal stages of the cryostation, we determined the bottom 300 K vacuum window to be one of the etalon surfaces (Fig.5.35). Unfortunately this is the difficult to manage of the cryostation's optical surfaces since it is directly coupled to ambient pressure and temperature fluctuations. It is difficult to access and a heater cannot be placed close enough to the




Figure 5.35: **Top:** Frequency response of the cavity to a step of the vacuum chamber temperature. **Bottom:** Residuals of an exponential fit to the frequency response above display a modulation indicative of an etalon in the system. The lack of a significant time constant between the temperature step and etalon response suggests that one of the etalon surfaces is not thermally isolated from the 300 K stage. Modeling the vacuum chamber temperature as a phase ramp, we can largely reproduce the cavity frequency response by fitting the coupling of temperature to phase. In the orange trace, we correspond a 100 mK vacuum chamber fluctuation to a 2π phase shift. Active suppression of the 2 Hz peak-to-peak etalon amplitude below the 10^{-16} level would require $\sim$ mK stabilization of the window temperature. This is difficult, an easier approach is to attempt to reduce the etalon finesse.

optic to close a loop with our desired bandwidth. We tried shielding it from changes in air pressure with lens tubes but had little success.

Since active stabilization is hard, a better method is to lower the etalon finesse by addressing




Figure 5.36: Sweeping the laser frequency, several transverse modes visible in the scan as well as a modulation on the baseline indicating the presence of an etalon. A fit to this reveals the etalon's free spectral range to be 160 MHz which corresponds to a cavity length of 0.9 m.

the second surface. Precise knowledge of the etalon length is required, but fortunately any etalon seen by the cavity will be present as amplitude modulation at the PDH detector. In Fig.5.36, we scan the laser frequency by nearly 1 GHz and observe a modulation on the baseline amplitude. Note that this modulation amplitude compared to the PDH signal is roughly 400 ppm, fairly consistent with the relative amplitude of the etalon observed in Fig.5.34. Given that these signals roughly match, we have some confidence that the etalon presence is fully captured in the RAM measurement and that we are measuring the same parasitic fringe.

The frequency of the baseline modulation is fitted as 160 MHz, corresponding to an etalon

free spectral range of 0.9 m. Now that we have identified one surface of the parasitic cavity, it is straightforward to identify the next optical element 0.9 away. The free space path is less than 0.4 m long, so somewhat surprisingly it must be the case that the second reflecting surface is a fiber element. We narrowed this down to one of the surfaces of our fiber EOM, presumably the faces are angle cut so its unclear why this was supporting a large reflection. Regardless, we placed a fiber isolator after the EOM and reduced the etalon amplitude by 20 dB.

5.4.2 Pointing stability

With the etalon oscillation suppressed, the next goal was to identify the source of the remaining long term instability. Electronic offsets varying on the timescale of hours or days can be painstaking to pinpoint. Conscientious grounding can help minimize these effects but slowly varying lab temperatures can easily drive voltage offsets or affect the setpoint of voltage references. We determine the sensitivities of our electronics to these effects by locally heating our mixers and servos by a few K (larger than the day-to-day thermal variation within S1B20) and looking at the resulting frequency excursion. With proper thermal insulation and improved grounding, we estimate that we minimize these effects to below the 10^{-16} level.

The cavity circulating power is always stabilized in transmission, but we still observed fluctuations in the DC voltage collected at the PDH photodetector. Correlations between the beatnote and this signal over long time scales (5.37) were evident, suggesting that this could be a reasonable detection channel for identifying the source of the long term beat instability. Our assumption is that there are likely small fluctuations in the beam pointing that either generate some optical scatter or modify the cavity contrast.

While pointing variations likely originate with small changes in the thermal environment, it is easier to actively stabilize the beam position than to temperature stabilize all the optics. A quadrant photodetector can measure horizontal and vertical variations of the beam position at the PDH photodetector location. To quantify our sensitivity to this effect, we can measure the transfer function between changes in the PDH DC voltage and cavity frequency when varying the beam




Figure 5.37: **Top:** Dedrifted Si3-Si6 beatnote overlaid with the PDHDC signal multiplied by the horizontal sensitivity coupling. **Bottom:** Instability of the beat, and PDHDC measurements when propagated through the horizontal and vertical transfer functions. Some sections of the beatnote instability between 100-1000 s appear to be well described by the horizontal pointing deviations. For longer averaging times, the cavity instability might result from a linear combination of vertical and horizontal pointing fluctuations.

pointing. We place a piezo between the pins of one of the steering mirrors used for coupling into the cavity and record the corresponding frequency deviation (Fig.5.38).

We measure vertical and horizontal couplings of the PDHDC to frequency of 300 Hz/V and 100 Hz/V respectively. A desired 0.1 Hz frequency stability (fractionally 5×10^{-16}) requires 300 ppm control in the vertical direction, and 1000 ppm control in the horizontal direction. Based on our measured long term instability, it appears that the passive control is already near this level. A




Figure 5.38: Modulating the voltage applied to the piezo allows us to measure a transfer function between the cavity frequency response and changes in the PDH DC signal. We do this for both axes to measure the system’s sensitivity to beam pointing. **Left:** The piezo is placed in the mirror’s horizontal pin. We drive it with a sawtooth and simultaneously record the frequency deviation. Red trace is the frequency response, and the green trace is the PDH DC voltage converted to frequency via our measured coupling. **Right:** Same data, except the piezo is now driving the mirror’s vertical pin.

feedback loop for further stabilization would only require a gain of 10 dB at an averaging time of 10^4 seconds.

5.4.3 Top setup

The increased sensitivity to beam pointing is largely due to the need to couple to the HG_{01} mode of the cavity. That in turn was necessitated by the observation that the frequency noise from 0.01-0.1 Hz was highly dependent on the input beam angle, and could be reduced by nearly 20 dB when aligning to maximize transmission on HG_{01} . There are two routes to approach this problem - one is to build a quadrant photodetector, measure the beam wander by picking off light before it is focused on to the PDH photodetector and servoing a piezo-controlled mirror. This is relatively straightforward and we designed and assembled a circuit to provide the quadrant sum and difference ports. The quadrant photodiode was unfortunately on backorder for a few months and it wasn’t feasible to wait for this to arrive.

Instead we exercised our second route forward; we decided to construct a setup above the cavity (called the “top breadboard”) that could launch the beat downwards through the top sub-

strate and mirror. The general idea is that we could avoid the bottom optic or substrate that was responsible for the poor RAM signal that we would measure. Of course, this is a significantly larger undertaking and requires constructing an entirely separate setup that still needs to have no additional technical noise contributions.

The majority of measurements made on the top breadboard have been tirelessly carried out by our recent postdoc Zhibin Yao. Already, we have found some promising properties of this scheme. In the bottom setup, we struggled to achieve more than 40% mode coupling to the cavity but from the top we can now achieve closer to 60%. Imaging the beam before it enters the vacuum chamber does not reveal any significant differences in the spatial mode. Rather, this suggests that one of the bottom vacuum optics may be adding a speckle onto the wavefront that decreases its overlap with the cavity's TEM_{00} mode. This is corroborated by images we take of a beam transmitted from both sides of a cavity. A clear gaussian mode is imaged when coupling light from the bottom, whereas coupling from the top results in a highly lopsided transmission profile (Fig.5.39)




Figure 5.39: **Left:** The light is coupled to the cavity from the bottom setup and the transmitted beam is imaged from the top. **Right:** Light is coupled from the top and imaged on the bottom.

We can additionally measure the residual RAM from the top setup after active stabilization on the RAM photodetector. Recall that this was the dominant technical noise contribution measured in Fig.4.38 and was only slightly below our measured instability of Si6 with birefringent noise cancelled. A reduction of this noise source could very well improve the cavity stability. Surprisingly our initial measurements after aligning to cavity showed a dramatic improvement over the out of loop RAM measured in the bottom setup. Without significant optimization the fractional frequency instability

dips below the Si3 thermal noise for over an order of magnitude (Fig.5.40). The fractional frequency noise is also dramatically improved across the entire spectrum when compared to the RAM in the bottom setup. A 10 dB reduction of white noise is observed due to the increased mode coupling and our ability to use the HG_{00} mode.

Our method of mounting the top breadboard on long optical posts above the cryostation opens us up to issues of relative motion between the optical beam and the cavity. Differential displacement noise can appear as refractive index changes while also modulating the cavity coupling. Currently we measure a higher frequency instability from the top breadboard and we suspect that differential motion is the likely cause.

The bottom setup has also been left intact during the installation of the top breadboard. A second set of resonant PDs is required as we simply used the ones from the bottom, but this leaves open the possibility of directly measuring birefringent noise using the PTB technique [20]. Two separate lasers can be locked to the independent birefringent modes of the cavity, and the frequency difference (assuming all technical noise sources are below this level) will yield the direct birefringent noise.

5.5 Drift rate

Low drift rates are coveted in cavities and in ULE spacers, cavity drift was thought to be associated with the spacer slowly sagging with time. The development of cavities with crystalline spacers and substrates held an additional attraction, namely that crystals are expected to have a fixed length dictated by the crystal lattice constant. Drift should be a bygone property, nonexistent in the JILA and PTB silicon cavities.

Surprisingly a nonzero drift was observed; while orders of magnitude smaller than its ULE counterparts, Si3 measured an absolute drift rate of -6 Hz/day, or fractionally $-1.4 \times 10^{-19}/s$. Over time this has exhibited a long term settling from $-30 \mu\text{Hz}/s$ to $-25 \mu\text{Hz}/s$. Intermittent interruptions where the laser falls out of lock have not affected this long term behaviour, suggesting that this might be endemic to the cavity.




Figure 5.40: Fig.4.12 is reproduced but now we include a measurement of RAM from the top setup. The modified allan deviation of RAM instability is now consistently better than the Si3 thermal noise from 0.5-10 s, and the fractional frequency noise shows a considerable 6 dB improvement around 100 mHz. No effort was put into improving the long term instability of RAM in the top setup, so we expect that this can be improved. The top setup couples to the HG_{00} mode while the bottom setup couples to HG_{01} .

Si4 however added some new information, the identification that the drift rate was dependent on the intracavity power. With transmitted powers of 700 nW, the drift rate was over 50 times than what was measured in Si3 at -6×10^{-19} . Significantly reducing the transmitted power to 40 nW resulted in a measured drift of -3×10^{-19} , only two times higher than Si3. Si3 maintains a higher transmitted power of order 100 nW so it's unclear why the behaviour is so different between

the two systems. Power dependent drift was unfortunately not measured in Si3. The Si4 results suggest that some mechanism in the cavity favors few photons circulating between the mirrors. A fit through the Si4 drift measurements suggests a zero drift at zero intracavity power but there is no clear theory as to why this might be the case.

Si6 coming online was exciting from the perspective of adding another data point to these measurements. To our surprise, long term drift measurements of Si6 (after implementing changes from section 5.4 to improve the long term instability) indicated a drift rate substantially lower than anything measured in Si4. Fractional drift rates below $1 \times 10^{-19}/\text{s}$ were measured over several one week periods. A particular measurement is shown in Fig.5.41. The Si3 drift was frequently measured by the SrI lattice clock during an AC stark evaluation so we have accurate knowledge of the average drift over this period. We find a fractional drift rate below $-6 \times 10^{-20}/\text{s}$, a factor three lower than measured in Si3. It is not clear that this is even the true drift rate of the cavity, Fig.5.37 shows the same dataset and suggests that the residual instability could be due to slow variations in beam pointing. The drift innate to the cavity material could be lower.

Regardless, this is significantly smaller than anything measured in Si4. Both cavity spacers and substrates were cut from the same ingot so it follows that the silicon purity should be identical. One conclusion we can infer is that the residual drift rate of Si3 and Si6 could have been due to a drift of the amorphous coatings. In the glassy material, there is no lattice constant to define the intramolecular spacing so these may have been changing length over time. This doesn't to be an unreasonable statement - ULE spacers in our lab are found to have a drift rate of 10 kHz a day. The Si3 drift is over 1000 times smaller. The coating thickness is $10 \mu\text{m}$, so fractionally the coatings are $0.212/(2 \cdot 10 \times 10^{-6}) = 10^4$ smaller than the spacer. The reduction in drift measured in Si3 is 10^3 , not quite commensurate with the switch from a glassy spacer to a crystalline spacer.

This speculation is heuristic and doesn't take into account power dependent effects which haven't been deeply studied in either the Si3 or our ULE systems. Nevertheless, it does appear that the drift rate is significantly suppressed in fully crystalline cavities and a true measurement would require a measurement campaign of several weeks while directly measuring the cavity frequency

with an optical clock.




Figure 5.41: A five day measurement of the Si3-Si6 beatnote with the Si3 drift removed. During this period the Si3 drift was measured by the SrI lattice clock so its sign and magnitude is known. The dotted line is a fit to the entire dataset and finds a fractional drift rate of $-2.3 \times 10^{-20}/\text{s}$, whereas the dashed line fits to the latter two thirds where a more linear behaviour is observed. This segment finds a fractional drift rate of $-5.6 \times 10^{-20}/\text{s}$.

Chapter 6

Novel applications of crystalline cavities

Despite the birefringent noise identified in Si6, the capabilities of the cavity are still powerful. This is not particular to crystalline coatings, but rather the unprecedented long term stability attained with cryogenic silicon cavities. With Si3, several demonstrations highlighted the role of the cavity in future research directions. Here, we will explore some of those experiments. We will identify the role of the cavity as not a technical upgrade, but rather a critical tool that has enabled key advances in frequency metrology, fundamental physics tests, and searches for emergent phenomena [47, 17, 133, 134].

6.1 Ratio Comparison

Searches for physics beyond the Standard Model description have often been associated with a top down approach where massive detectors sensitive to energy at the GeV level attempt direct or indirect observation of new physics. Dark matter searches in particular have often followed this approach; a leading candidate for dark matter, the weakly interacting massive particle (WIMP), has been the subject of intensive experimental efforts at accelerators and other large facilities without conclusive detection [135]. The primary method for direct detection relies on energy deposition, i.e. where the interaction of WIMPs with a large volume of material and a useful signal is obtained from the observation of a few events. With a low event rate, bubble chambers and scintillators will employ 100s to 10000s of liters of a particular material in a sterile, underground environment. The manpower and complexity of these experiments is significant and often prohibitive.

These experiments have succeeded in setting stringent limits on the WIMP candidate mass, but generally in the high energy range. Other candidates for dark matter have been proposed in a significantly lower mass range, including ultralight scalar dark matter (the dilaton). This class of models proposes that dark matter is composed of a light bosonic field where the particle mass could fall in the range of $10^{-24} - 1$ eV [136, 137]. The coupling to Standard Model fields is theorized to lead to a coherent oscillation of the fine-structure α and the electron mass m_e [138]. The oscillation frequency of this scalar field can be directly related to oscillations in the fine structure constant, with the oscillation amplitude dictating the magnitude of the coupling to the electromagnetic gauge with modulus d_e and to the electron mass with d_{m_e} . In particular, the Lagrangian of the dark matter field ϕ is given as [139]

$$\mathcal{L}_\phi = \kappa\phi\left(\frac{d_e}{4}F_{\mu\nu}F^{\mu\nu} - d_{m_e}m_e\bar{\psi}_e\psi_e\right) \quad (6.1)$$

The overall constant κ is static but the first term captures the coupling of the DM field to the electromagnetic tensor $F_{\mu\nu}$ and the second term describes the coupling to the electron mass with electron field ϕ_e .

We've touched on this point several times but optical frequency metrology has demonstrated some of the highest precision in frequency fluctuations. Optical and microwave transitions in atomic clock transitions are known to demonstrate a dependence on the fine structure constant and electron mass. A differential dependence on α is evident between optical and microwave transitions. In optical clocks, the transition frequency exhibits the proportionality

$$f_{optical} \propto \alpha^{2+\epsilon} m_e c^2 \quad (6.2)$$

The factor ϵ originates from the relativistic correction to the atomic energy levels and differs between atoms. In the nonrelativistic approximation, frequency ratios of different electronic transitions between atoms would not exhibit α sensitivity [140, 139].

The enhancement factor ϵ determines the sensitivity to fluctuations in alpha. The variation

of the electronic transition frequency to fine structure oscillations can then be expressed as

$$\frac{\delta E}{E_0} = \epsilon \frac{\delta \alpha}{\alpha} \quad (6.3)$$

A frequency ratio intercomparison between optical clocks with different values of ϵ can therefore be sensitive to scalar DM oscillations. Between two clocks A and B , fluctuations in the frequency ratio can be related to the variations in the fine structure constant via

$$\frac{\delta f_A}{f_A} \bigg/ \frac{\delta f_B}{f_B} \propto (\epsilon_A - \epsilon_B) \frac{\delta \alpha}{\alpha} \quad (6.4)$$

It is therefore the case that precise measurements of frequency ratios between optical clocks can place bounds on the electromagnetic gauge modulus d_e [141, 142, 143]. In the BACON collaboration, we measured three frequency ratios between the Al^+ ion clock, and the Yb and Sr lattice clocks. We use this dataset to perform one of the highest precision measurements of three fundamental constants of nature. In principle if we want to investigate the possibility of fluctuations in α , i.e. fluctuations in these frequency ratios, the same dataset can be used to place bounds on this oscillation. This of course relies on the electronic transition of each clock having a slightly different relativistic correction to α . As it turns out, this is indeed the case and atomic structure calculations reveal sensitivity factors [143]

$$\epsilon_{\text{Al}^+} = 0.008 \quad (6.5)$$

$$\epsilon_{\text{Yb}} = 0.31 \quad (6.6)$$

$$\epsilon_{\text{Sr}} = 0.06 \quad (6.7)$$

From Eq.6.5 we see that the largest sensitivity to α variation is obtained in the $f_{\text{Yb}}/f_{\text{Al}^+}$ frequency intercomparison. However, this does not necessarily imply that frequency intercomparison between these two clocks can lead to the most stringent exclusion bound. To clarify this point, we need to understand how to convert the ratio data to a bound on the electromagnetic gauge modulus d_e .

The ansatz for the coherent DM field is an oscillation [139, 144]

$$\phi(t) = \phi_0 \cos(m_\phi t - m_\phi \langle \mathbf{v} \rangle \cdot \mathbf{r}) \quad (6.8)$$

The first term concerns the time dependent oscillation of the scalar field and we see that the DM mass is equated to the frequency of oscillation. From the perspective of metrology, the dark matter candidate mass m_ϕ is related to the measurement bandwidth:

$$m_\phi = 4.13 \times 10^{-15} f \quad (6.9)$$

Note that this relation implies that a longer data record allows us to place constraints on a lower candidate mass range. In the case of the three-way frequency ratio measurement, the upper end of the mass range is set by the cycle time of the involved clocks. For lattice clocks, a 4 point locking sequence may be used to reject magnetic field noise by probing opposite stretched states, meaning each measurement of the optical transition frequency occurs in ~ 4 s. This corresponds to the upper limit of $m_\phi \sim 10^{-15}$ in our exclusion plot.

The second term identifies a spatial variation of the dark matter field between the measurement locations. The theory considers a slow moving mass of dark matter where $\mathbf{v}$ is the virial velocity of dark matter in our galaxy and is roughly $5 \times 10^{-4}c$. Given that the clocks are not co-located, differential phase shifts due to the spatially varying DM field could be an important contribution.

We can calculate the de Broglie wavelength of the coherent oscillation as $\lambda_{DB} \approx h/(m_\phi v)$. In the mass range of $10^{-21} < m_\phi < 10^{-17}$ eV, the DM wavelength varies from 10-10,000 parsecs, significantly larger than the $\mathcal{O}(1 \text{ km})$ separation between clocks. We can approximate this effect to be spatially homogeneous on the 1 km length scale, so we neglect it in our analysis and only consider the time dependent oscillation. We can rewrite Eq.6.4 in terms of the coupling to the DM field from the Lagrangian as

$$\frac{\delta f_A}{f_A} \bigg/ \frac{\delta f_B}{f_B} = d_e (\epsilon_A - \epsilon_B) \kappa \phi_0 \cos(4.13 \times 10^{-15} ft) \quad (6.10)$$

The quantity ϕ_0 is evaluated from the energy density $\rho_{DM} = \frac{1}{2}m_\phi^2\phi_0^2$, where it is measured that $\rho_{DM} = 0.3 \text{ GeV/cm}^3$. We then identify the scaling

$$\phi_0 = \frac{\sqrt{2\rho_{DM}}}{m_\phi} \quad (6.11)$$

The magnitude of the coupling d_e is then calculated from

$$d_e = \left(\frac{\delta f_A}{f_A} \bigg/ \frac{\delta f_B}{f_B} \right) \frac{1}{(\epsilon_A - \epsilon_B)} \frac{4.13 \times 10^{-15} f}{\kappa \sqrt{2\rho_{DM}}} \quad (6.12)$$

Note the frequency dependence in our calculation of d_e . Measurement of a frequency ratio between two atomic clocks will typically scale with stability $\sigma_y \sim 1/\sqrt{\tau}$ if not limited by any fluctuation systematic effects. The power spectrum of this frequency measurement will look white (f^0), meaning that the d_e exclusion will scale as f^1 . A longer dataset will not only allow for exclusions to be placed at a lower mass range, but will also place a lower limit on the magnitude of d_e . Ultimately, constraints on the low mass range are limited by coherence lengths on the order of galaxies.

The transformation from a ratio dataset to an exclusion limit is therefore conceptually simple. However in practice, the data is often taken in a non-continuous matter. Clocks may be operational for a short period of time, then experience an extended down time, and then continue taking data again. The existence of these time gaps complicates the calculation of the ratio's power spectral density. Traditional estimators for the power spectrum such as the periodogram are not appropriate tools for this situation. We instead use a particular method intended as an improved estimator for gapped data, the Lomb-Scargle periodogram.

The Lomb-Scargle periodogram may be unfamiliar to those primarily concerned with short term clock and oscillator instability since these are typically continuous data records. The periodogram is an appropriate estimator of the continuous Fourier spectrum that is no longer evenly sampled but rather discretely sampled. For long datasets, the latter constraint may not hold and the Lomb-Scargle is an improved estimator for datasets that are not evenly sampled. A confidence interval is required to determine whether the LS periodogram actually contains a periodic signal. Probabilities of these false positives can be obtained by bootstrapping of the data.

We now have all the tools necessary to produce a d_e exclusion limit, and a plot of our calculation is shown in Fig.6.2. Bounds are placed in the mass range of $10^{-21} - 10^{-15}$ eV.

In the ratios involving the Yb lattice clock, a deviation from $d_e \sim f$ is observed. This originates from a deviation in the ratio stability from $\sigma_y \sim 1/\sqrt{\tau}$ at long times. At this point,




Figure 6.1: We simulate a continuous dataset of different types of noise, $S_y \sim f^0, f^{-1}, f^{-2}$. The performance of the periodogram and Lomb-Scargle periodogram estimators is compared. The Lomb-Scargle periodogram consistently underestimates the noise spectrum by a factor of 2 in amplitude. A correction therefore needs to be included in order to obtain a comparable result.

the power spectrum of the Yb/Sr and Al^+/Yb ratios flicker, likely due to a fluctuating systematic offset. The Al^+/Sr ratio has a shorter dataset and is not long enough to resolve this feature.

A more important feature is the comparison between the exclusion limits set by the Yb/Sr

ratio and the Al^+/Yb ratio. From Eq.6.5, we expect

$$\epsilon_{Yb} - \epsilon_{Al^+} = 0.302 \quad (6.13)$$

$$\epsilon_{Yb} - \epsilon_{Sr} = 0.25 \quad (6.14)$$

Sensitivity to a dark matter oscillation is larger in the Al^+/Yb ratio, yet the exclusion bound is not as strong as that between the lattice clock ratios. This is due to the superior stability demonstrated in lattice clocks, largely due to the fact that neutral atom clocks can interrogate several atoms at once and have much smaller quantum projection noise of the measurement. The magnitude of this limit is directly related to the Dick effect - longer interrogation times and shorter deadtimes can improve the d_e exclusion, just as it enables shorter averaging times to reach a particular level of frequency precision. Nevertheless, particular ion-neutral atom frequency comparisons such as $^{171}\text{Yb}^+/\text{}^{87}\text{Sr}$ which are currently operating have such a strong alpha sensitivity that the high instability of the ratio measurement can still lead to superior d_e exclusion limits [145]. Proposed neutral atom comparisons [146] leverage certain transitions that can provide greater sensitivity with the significantly improved instability of frequency comparisons between optical lattice clocks.

A significant variable for improvement is the stability of the local oscillator. Better cavity stabilized lasers with lower thermal noise limited performance will enable a reduction in the d_e limit for all frequency intercomparisons. The cavity performance is therefore a crucial component in the role of clocks probing fundamental physics.

6.2 Dark matter searches with oscillators

Frequency ratio comparisons are clearly a powerful tool for investigating α variations, and the role of a stable oscillator is evident. However, there is another opportunity presented when the oscillator in question is a monocrystalline cavity. In the ratio comparison, sensitivity to d_e is only attained due to the differential dependence on α between electronic transitions in different clock atoms. We saw that this originates from the relativistic correction to the transition energy. In a crystal, the length is classically determined by the number of atoms N and the Bohr radius a_B ;




Figure 6.2: Exclusion limits for d_e measured between the three optical clocks during the BACON frequency intercomparison data campaign. The dataset is the same as that from [11]. The most stringent bounds on d_e are set by the Yb/Sr frequency ratio despite the relatively higher sensitivity of the Al^+/Yb comparison.

there is no room for a significant relativistic contribution. In a cavity with a crystalline spacer, the phase shift defined by the cavity length L is simply $\Phi = \omega L/c$. The frequency ω can be arbitrary but is typically chosen to be the optical transition frequency of the clock atom, so that $\omega \propto e^2/a_B\hbar$. Then,

$$\Phi = \frac{\omega L}{c} \propto \frac{e^2}{a_B\hbar} \frac{Na_B}{c} = N\alpha \quad (6.15)$$

A variation in the fine structure constant then appears as a proportional variation of the laser phase

$$\frac{\delta\Phi}{\Phi} \sim \frac{\delta\alpha}{\alpha} \quad (6.16)$$

The sensitivity of the cavity frequency to fine structure oscillations scales as α^1 . If we now compare this to the electronic transition frequency of the strontium atom, we can obtain a significantly

enhanced alpha sensitivity as

$$\frac{\delta f_{Sr}}{f_{Sr}} \bigg/ \frac{\delta f_{cav}}{f_{cav}} \propto (2.06 - 1) \frac{\delta \alpha}{\alpha} \quad (6.17)$$

Remarkably, measuring the cavity frequency (which is already required in clock operation) can provide a 4 fold enhancement in sensitivity when compared to the Yb/Sr ratio ($\Delta\epsilon = 1.06$ compared to $\Delta\epsilon = 0.25$). The crystalline cavity is no longer a factor in the magnitude of the exclusion limit, but is now directly used as a probe for variations in fundamental constants [47].

We can now perform a calculation of the d_e exclusion limit in the same manner as we did with the ratio data. In both studies, it is important to note that the observation run extends for a time less than the coherence time of the dark matter oscillation. Were this not the case, further collection of data would not improve the signal to noise of the DM signal. The dataset for the f_{Sr}/f_{cav} comparison is a consecutive measurement over 12 days using the 124 K Si3 cavity. Despite the earlier assertion that the ratio dataset could be used for the cavity-Sr atom analysis, technical issues prevent us from doing so. At the time of the three-way ratio data campaign, measurement of the cavity frequency was corrupted by additional path length noise that we had not identified in the distribution of the cavity stabilized light to the Sr atom. Measurement of the cavity frequency needs to be verified by the noise model of the cavity instability. Otherwise, additional noise may be misconstrued as a DM oscillation. The silicon cavity noise spectrum is fit to a model of

$$S_y = 4 \times 10^{-33} f^0 + 1.7 \times 10^{-33} f^{-1} + 9 \times 10^{-38} f^{-2} \quad (6.18)$$

During this measurement campaign, a hydrogen maser at NIST was distributed to JILA across a fiber link, allowing for a direct comparison of the Sr and silicon cavity frequencies to the maser frequency. The maser has a significantly higher instability when compared to the optical oscillator and the optical atomic transition, however it operates on a hyperfine manifold which exhibits a larger sensitivity to α . The transition frequency varies as

$$f_H \propto (\alpha^2 m_e) \alpha^2 \mu_p \left(\frac{m_e}{m_p} \right) \quad (6.19)$$




Figure 6.3: **Left:** Comparison of the data record of the H-maser measured against the silicon cavity frequency, and the 1D Sr lattice clock measured against the silicon cavity. The time series of the maser-cavity frequency ratio is decimated by a factor of 720 for clarity. The short term instability of the atom-cavity measurement is far superior to that of the maser-cavity ratio, but the enhanced sensitivity of the maser to the scalar DM oscillation will enable it to set a competitive bound. **Right:** Frequency noise of the f_H/f_{cav} (blue) and the f_{Sr}/f_{cav} (red) ratios calculated with the Lomb-Scargle method. A noise model for each ratio derived from the frequency noise contributions of the cavity, maser, and 1D Sr lattice is shown in black.




Figure 6.4: A cartoon illustrating the dependence of the H maser, silicon cavity, and strontium clock transition on α and m_e .

Where μ_p is the proton magnetic moment, m_e is the electron mass, and m_p is the proton mass. Comparison with the Sr atom will not provide a significant bound due to the relatively similar $\alpha^{2+\epsilon}$

dependence, but frequency ratio between the H maser and silicon cavity will appear as

$$\frac{\delta f_H}{f_H} \bigg/ \frac{\delta f_{cav}}{f_{cav}} \propto \alpha^4 \mu_p \frac{m_e^2}{m_p} \bigg/ N\alpha \sim \alpha^3 \mu_p \frac{m_e^2}{m_p} \frac{\delta\alpha}{\alpha} \quad (6.20)$$

The scaling with α^3 is significantly more favourable than the $\alpha^{1.06}$ from the cavity-atom comparison. The frequency noise visible in Fig.6.3 indicates an H/Si frequency ratio dominated by the white noise of the maser at nearly all frequencies, but the higher α sensitive will allow the maser comparison to place a competitive d_e bound at low measurement bandwidth.




Figure 6.5: The ratio of the observed DM signal compared to the actual signal amplitude when considering a nonzero phase δ to the DM field.

Before producing an exclusion plot, there are a few points to consider. Unlike the frequency ratio data, both the cavity and H maser experience a linear drift which is removed in the analysis. The impact of this drift on the generated Lomb-Scargle spectrum needs to be carefully quantified. We simulate a DM oscillation over a range of frequencies and compare the peak amplitude recovered through the LS periodogram with the known amplitude. The obtained amplitude is underestimated by a factor of up to 0.7 at our lowest measurement bandwidth of $0.7 \mu\text{Hz}$. A similar correction needs to be applied in consideration of a slowly varying phase of the DM field. Our expression for the scalar field should in fact contain a phase term δ as $\phi(t) \propto \cos(m_\phi t + \delta)$. Particular values of δ can lead to an underestimation of the field amplitude. At high bandwidth, the DM field completes several oscillations over the data campaign so that the observed amplitude is equivalent

to the actual signal amplitude. This is not the case for lower measurement bandwidth, and again a correction needs to be applied. A simulation of the necessary correction is illustrated in Fig.6.5.

With these points in mind, we can generate the d_e exclusion plot which is shown in Fig.6.6. The cavity-atom comparison demonstrates a significant improvement in the exclusion bound within the mass range of $10^{-19} < m_\phi < 10^{-17}$. Below this range, the maser-cavity comparison sets a lower limit due to the improved long term instability of the maser. This is comparable to the constraints placed by a Rb/Cs atom-atom frequency comparison [15].




Figure 6.6: The constraints placed in this work are the blue curve for the maser-cavity frequency measurement, and the red curve denoting the atom-cavity measurement. Purple lines indicated limits set from equivalence principle tests [12, 13]. Black curves are exclusion bounds determined from atom-atom frequency comparisons [14, 15, 16].

It is evident that tabletop atomic physics experiments can provide meaningful contributions to the search for varying fundamental constants. More important is the fact that exclusion bounds

placed by this experiment use an entirely different platform. The cavity stability is directly related to the magnitude of the f_{Sr}/f_{cav} bound that we can set. An improved long term instability or a lower thermal noise limit enables a more stringent exclusion of d_e . The current bound set in Fig.6.6 does not even use the full capability of the local oscillator stability. Uptime of the Sr clock during measurement of the cavity was only 30% of the 11 day data campaign. The dead time introduced by the time gaps is significant; a calculation of the bound that could be set in the f_{Sr}/f_{cav} measurement with a continuous operation is calculated as the dotted green curve of Fig.6.6.




Figure 6.7: Exclusion limit for the electron mass coupling d_{me} obtained from the maser-cavity frequency comparison.

There is one additional point of interest that we have not yet touched on. The second term of the Lagrangian in Eq.6.1 reveals a sensitivity to the electron mass with the electron mass modulus d_{me} . Oscillators or references with differential m_e sensitivity could contribute to an exclusion limit for this coupling. The cavity and Sr atomic frequency both scale with m_e so their frequency ratio will be insensitive to fluctuations of this term. The H maser frequency however scales with m_e^2 so a frequency comparison with either the cavity or Sr atom could constrain the value of d_{me} with sensitivity $\sim m_e$. The H/Si comparison has lower deadtime and probes a larger mass range due to

the longer extent of the data campaign, so we use this to evaluate the d_{me} exclusion.

The exclusion is displayed in Fig.6.7. The maser-cavity measurement results in a significantly lower bound in the range of 2×10^{-20} to 3×10^{-19} . The exclusion magnitude is limited by the H maser instability in the majority of this mass range. Nevertheless, the measurement precision is good enough that an extended data campaign could still place competitive bounds in a lower mass range.

6.3 Timescale

In the previous section we leveraged the excellent long term stability of the silicon cavity to place competitive bounds at the low mass range of d_e and d_{me} . For the d_e limit, a particularly important feature is the lower frequency noise of the silicon cavity when compared to the maser. In Fig.6.3(b) the silicon cavity displays a superior performance at frequencies down to $1 \mu\text{Hz}$. Stated another way, the silicon cavity is a superior oscillator at averaging times up to nearly 10^5 seconds, coincidentally the timescale of a day.

This begs the question of whether the silicon cavity as a free running oscillator can act as a timescale superior to the maser. Over the course of a single day, this is clear from the improved optical local oscillator (OLO) instability. The maser instability is well known from an N-cornered hat analysis of the NIST AT1 maser array. Measurement against the silicon cavity frequency therefore allows us to ascribe features above the maser noise model to the optical cavity instability. In Fig.6.8, we see that measurement of the silicon cavity frequency by the AT1 timescale is only limited by the optical oscillator at 2×10^5 s.

At longer averaging times it may appear that the optical oscillator has already lost the capability to serve as an improved timescale, but instability is not necessarily the most important metric. Eric Oelker and Will Milner vested considerable time with the NIST timescale division to apply the usual predictive model used for microwave standards to the optical oscillator. The details of this study are more thoroughly discussed in [17] so we'll only cover the key conceptual points needed to illustrate the unique space that these silicon cavities occupy. For microwave timescales,




Figure 6.8: Long term instability of the silicon cavity (orange and red points) with our model for the oscillator noise (dashed black line). In a single day’s measurement, particular Si3 datasets have displayed thermal noise limited performance out to 10^4 seconds. This low instability is naturally degraded with the inclusion of significantly more data. A similar effect is apparent in the Si6 long term behaviour, short datasets of 6-8 hours demonstrate instabilities of $\text{mod } \sigma_y < 1 \times 10^{-16}$ for 10^3 s, though a week long dataset shows $\text{mod } \sigma_y = 1 \times 10^{-16}$ at 200 s. Our intent is to point out that length of the dataset is critical when assessing the long term cavity behaviour.

a Kalman filter is typically used to estimate the oscillator frequency by weighting the accumulated timing errors. This is achieved by measuring the oscillator against a frequency standard. In our case we use our 1D Sr optical lattice clock which is unconventional in the benchmarking of microwave timescales. However recall that we established this as an appropriate choice of standard in the three-way optical frequency ratio intercomparison of [11]; measurements of the Yb/Sr optical frequency ratio over a five month period resulted in a series of measurements spread within a 1×10^{-17} range, each with total statistical and systematic uncertainties at the mid 10^{-18} level. Accuracy aside, the capability of the neutral atom Sr clock to consistently measure an optical frequency is evident. More importantly, the ratio evaluation is performed at a level much lower than the OLO and microwave stabilities across all averaging times. This makes the Sr clock suitable as a frequency reference.

It is crucial to note that the prospect of using Si3 as a timescale is inherently tied to the predictable drift rate of crystalline cavities [147]. Predictable not in the sense that the drift rate can be calculated *ab initio*, but rather that it is consistent over several months. To date, no ULE cavity has exhibited such performance and likely never will; changes in the drift rate of our 40 cm ULE cavity already require daily dedrifting for use as an OLO in the 1D optical lattice clock, let alone its use in a timescale. The Si3 fractional drift rate of $1.4 \times 10^{-19} \text{ s}^{-1}$ exceeds that of NISTs older masers ($5 \times 10^{-22} \text{ s}^{-1}$) but is consistently stable over the span of a year, whereas the latter is not stable over intervals of months. A lower drift rate therefore does not establish a superior timescale, a high level of predictability of the oscillator frequency is more important.

Furthermore, a true timescale is generated by *steering* to an atomic standard, not simply evaluating the LO frequency deviations. At this point we've discussed limitations due to deadtime several times over. The situation here is no different - a superior shot term stability will always lead to a lower Dick effect-limited instability. In the context of a timescale, steering a microwave standard with a fixed duty cycle to a precise optical reference will not outperform the steered OLO. This carries the caveat that the reference uptime is not so intermittent that the OLO instability has surpassed the microwave standard, but we see from Fig.6.8 that this simply requires us to operate the clock once per day.

A calculation of this effect is presented in Fig.6.9. The ability of the OLO to maintain low frequency instability is critical for the superior performance of this all-optical timescale. The long term predictive behaviour is equally important in operating this timescale for several months; the first quality has so far only been realized with cryogenic crystalline reference cavities and the second quality is a feature unique to them. A free running optical oscillator oscillator has been used to establish improved timekeeping when compared to conventional microwave standards. Again, we make the point that these crystalline resonators directly enable advances in metrology rather than indirectly aiding clock stability.




Figure 6.9: Figure adapted from [17]. Stability of the optical and microwave timescales is evaluated when measuring against an optical Sr standard for two duty cycles. Measurement of the OLO for 1 hour per day is still superior to the microwave timescale when measured against the strontium clock 12 hours/day. This is directly related to the superior short term thermal noise limited performance of Si3 and the excellent long term control. We additionally display the significant advantage enabled by measuring the microwave timescale against an optical reference compared to a microwave reference.

6.4 Gravitational wave detection

Given the small displacement fluctuations that can be measured with our silicon cavities, it is interesting to consider the strain sensitivity of two orthogonally oriented resonators. This is not a novel concept - resonant mass detection of gravitational waves (GWs) was first investigated in Weber bars [148], and later refined in resonant detectors such as AURIGA [149], miniGRAIL [150], and ALLEGRO [151]. These generally use capacitive transduction to measure the stretching and contraction of the bar detector in response to an incident GW. Crystalline resonators straddle a unique space between GW interferometers and resonant detectors, where optical interferometry is used to read out length variations of the Fabry-Perot but the rigid spacer between the mirrors turns this into a resonant device. The cavity therefore becomes a transducer of mechanical displacement into changes of optical phase, which we are proficient at measuring.

The sensitivity limit to length fluctuations is determined by the measured noise spectrum of the oscillator, and fundamentally limited by the thermal noise of the resonator. We omitted one caveat when discussing this in Chapter 2 primarily because we aim to interrogate the cavity thermal noise at the low bandwidth $\lesssim 1$ Hz regime. In reality, the high frequency thermal noise of each component will be lowpassed above the resonator's resonance; this same principle is present in LIGO's ability to lowpass the suspension thermal noise with a steep $S_y \sim f^{-4}$ character. The total thermal noise of the cavity is more accurately expressed as [152]

$$S_{x,\text{Brownian}} = \frac{4k_B T}{2\pi^2 E_{sp}(R_o^2 - R_i^2)} \frac{f_0^4}{(f_0^2 - f^2)^2 + f_0^4 \phi_{sp}^2} \frac{\phi_{sp}}{f} \quad (6.21)$$

$$+ \frac{4k_B T(1 - \sigma_{sb}^2)}{2\pi^{1.5} E_{sb} w_0} \frac{1}{f} \phi_{sb} \left[1 + \frac{2d(1 - 2\sigma_{sb})}{w_0 \sqrt{\pi}(1 - \sigma_{sb})} \frac{\phi_{coat}}{\phi_{sb}} \right] \quad (6.22)$$

The spacer Brownian noise is lowpassed above the resonant frequency f_0 , and enhanced around this region by the material's quality factor. Note that similar lowpassing of the substrate and coating noise terms also occurs, but the resonant frequency of these components are significantly higher than that of the spacer.

The interferometer's strain sensitivity $S_h(f)$ is determined from the fractional length noise as [153]

$$S_h(f) = \frac{S_y(f)}{|G(f)|^2} \quad (6.23)$$

Here $G(f)$ is the cavity transfer function to a length perturbation. This can be modeled as a damped harmonic oscillator driven by a gravitational wave of frequency f . As we discussed in the thermal noise section, the transfer function of the system displays a lowpassing effect with enhanced sensitivity at the resonance frequency f_0 , with amplitude determined by the mechanical quality factor of the resonating material. The effect of the transfer function is therefore to resonantly amplify a signal at the spring's resonant frequency f_0 . Maximal sensitivity to a GW signal in a local frequency band will occur at f_0 . The efficacy of using cavities for GW detection amounts to monitoring signals around the primary internal mode of vibration.

However there is one important point to note, the driving force, the gravitational wave in this instance, scales with frequency as [154]

$$F_{GW} \propto h(2\pi f)^2 \cos(2\pi ft) \quad (6.24)$$

The effect of this is to scale the typical damped harmonic oscillator transfer function by f^2 , effectively treating it as a highpass filter below the resonance. From Eq.6.23, we can now see that the strain sensitivity of the Fabry-Perot is highly suppressed as the bandwidth decreases. Unfortunately, this is the exact frequency band where these oscillators attain the best length instability.

Note that this measurement, unlike the others presented in this chapter, does not rely on the use of crystalline cavities. Rather, cryogenic silicon cavities provide a measurement advantage due to the lower thermal limit, and the enhanced detection sensitivity at f_0 due to the significantly lower loss angle of silicon compared to ULE. This is important in the context that cavity measurement of gravitational waves does not provide a competitive sensitivity limit in other practical frequency bands. Unlike normal optical interferometry, the resonator geometry design choice is inverted: the mechanical resonance is no longer placed above the bandwidth of interest, but rather is purposely placed at a frequency where the gravitational wave signal power is expected to be high.

We illustrate this effect in Fig.6.10. In Si6, a spacer resonance appears at 10s of kHz, manifesting as a sharp peak in the thermal noise limit of the cavity. In an actual experiment, there will be a white noise limitation due to the shot noise limit which can be reduced by coupling more optical power into the cavity. In our current realization of the Si6 stability, we operate with low levels of light to reduce any birefringent effects and the shot noise contribution at the spacer resonant frequency is much higher than the resonance peak amplitude. In practice, it is not feasible to reduce the shot noise limit by the five orders of magnitude necessary to have the Si6 total thermal noise be dominant at 30 kHz. We consider all cases however, simply for the purpose of providing some intuition.

The cavity's mechanical transfer function also has a resonance at the same mechanical eigenfrequency of the cavity. The resulting strain sensitivity displays a sharp, narrow enhanced sensitivity




Figure 6.10: **Top:** Curves from the top plot are converted to a strain sensitivity by normalization with the cavity's mechanical response to a gravitational wave. The aLIGO strain sensitivity is displayed in red. **Bottom:**

at this region for a cavity with noise spectrum dominated by technical effects, or coating/substrate thermal noise. If instead the spacer thermal noise is dominant, the strain $S_h(f)$ will instead be monotonically decreasing with frequency.

At the highest sensitivity of aLIGO in the 100-1000 Hz band, strain measurements with a thermal noise limited Si6 cavity is nearly 10 orders of magnitude higher in amplitude. State of the art Fabry-Perot cavities have no chance at providing useful bounds on GW detection. Note that this comparison is strictly related to the detection of gravitational waves, meaning that the rigid

spacer of the cavity is a significant inhibitor to detection. If we are instead concerned with the amplitude spectral density of length fluctuations $\sqrt{S_x}$ ($\text{m}/\sqrt{\text{Hz}}$), thermal noise limited cryogenic reference cavities can provide a better resolution. This again requires the shot noise limit to be negligible, but Si6 for example achieves a length noise spectral resolution of $3 \times 10^{-19} \text{ m}/\sqrt{\text{Hz}}$ at 100 Hz, compared to $\sim 2 \times 10^{-20} \text{ m}/\sqrt{\text{Hz}}$ for aLIGO.

The only promising advantage of these systems occurs at high frequencies $> 10 \text{ kHz}$, outside the band of GW interferometer operation. There are of course many technical difficulties associated with achieving this level of strain sensitivity with tabletop optical cavities, and it is only useful if there are interesting GW sources emitting at these frequency bands. Nevertheless it is interesting to consider the whether these systems can serve as useful detectors for astrophysical phenomena.

Chapter 7

Conclusion

We begin our discussion of ultrastable lasers by remarking that the usual and real lasers which we can buy and use in the lab do not in fact emit coherent light. The critical issue here is phase stability, required for any meaningful claim of coherence. We use lasers to allow us to excite transitions in the optical domain, since the higher frequency can lead to a higher line Q. However it also immediately brings us the expectation of phase incoherence within the few millisecond domain due to vibrations imposed on the resonator structure. Our position is that many of the most interesting physical measurements using lasers will only be possible if the full potential coherence time is available – limited only by inescapable fundamental processes. Thus it is interesting work to identify and eliminate non-essential noise sources so that one may begin to appreciate what the real limits may be. - Jan L. Hall

At this point, we've hopefully presented a convincing case as to why developing better local oscillators are imperative to pushing frontiers of optical frequency metrology. However, we've also demonstrated that a tremendous amount of work is necessary to achieve the levels of stability realized in this thesis. Promising candidates for reducing Brownian thermal noise such as AlGaAs/GaAs crystalline coatings have not lead to a straightforward reduction in oscillator instability. In this first investigation of the spectral noise of a silicon cavity with crystalline coatings, we observed several static and dynamic light-induced effects that affect the stability of the cavity's spectral feature. To date, no clear theoretical framework has been presented which will be important for future use of these coatings in ultrastable oscillators, gravitational wave detectors, and quantum optomechanics experiments.

So what is next? It appears that the future of low noise optical interferometry may be at an impasse. Significant research efforts have been dedicated to developing new dielectric coatings that

can exhibit higher mechanical Q than conventionally used $\text{SiO}_2/\text{Ta}_2\text{O}_5$ coatings, but this may not yield the factor of 10 in loss angle that we hope to achieve. Making longer cavities is not prohibitive, but does become technically demanding when placing the resonator in a cryogenic environment.

A more meaningful approach is to consider what levels of stability we may need. It would be a delusion to expect that similar approaches to developing silicon resonators could realize an oscillator with 10^{-19} level instability. The demands of reducing technical noise sources are far too high. Fig.7.1 shows us the optical coherence time T_{coh} that can be expected for different cavity stabilized lasers that we have realized in our lab, or have considered for future systems. Different laser coherence times can be obtained with different measurement methods, here we integrate the Brownian thermal noise limit until the the RMS phase excursion of the laser is 1 rad. This scales as $T_{coh} \sim \tau^{-1}$ so an improved cavity stability will lead to an improved coherence time. What is noteworthy is that particular systems that we could consider (21 cm cavities at 16 K) provide a dramatic improvement in the optical phase coherence. Issues related to crystalline coatings aside, a fully crystalline 21 cm cavity at 16 K could coherently drive a transition for the entire lifetime of the $^{87}\text{Sr } ^3P_0$ clock state.

This is a metrologically useful ultimate goal for oscillator stability, i.e. coherently interrogate the optical clock transition over the maximum possible spectroscopic duration. The technical demands are also not outrageous. In Si6 we reduce all noise sources to the 10^{-17} level or below (with the exception of RAM, which we identify as an optical scatter issue). Electronic contributions will be reduced proportionally with the cavity linewidth, meaning a three times longer cavity will make these sources three times less relevant. Technical effects that couple to the cavity length can be proportionally suppressed with additional thermal mass in the case of temperature fluctuations, or reduced via more aggression active stabilization in the case of vibration noise.

There will certainly be challenges associated with realizing a Brownian thermal noise limited performance of 4×10^{-18} , but this is not out of the realm of imagination, considering what we have already demonstrated in the JILA-PTB collaboration. With the rapid progress still being made with cryogenic reference cavities, we could very well reach the point where a classical resource




Figure 7.1: We plot the total Brownian thermal noise limit for MJM (purple), Si2/Si3 (green), the limit for a 21 cm silicon cavity at 16 K with dielectric coatings using the measured loss angle from [18], and the limit for a 21 cm silicon cavity at 16 K with AlGaAs/GaAs coatings using the measured loss angles from [19, 20]. A dashed black line indicates the instability where an oscillator has experienced a phase excursion of $\phi_{\text{rms}} = 1$ rad.

provides us with all the capabilities we would ever require for an optical oscillator. The record breaking long term instability that we have demonstrated has no immediate competitor so far.

Low noise optical oscillators have come far from the days where Hz-level laser linewidth was a technical feat in the 2000-2010 epoch. We have now afforded ourselves the right to speculate on capabilities of true mHz linewidth oscillators. Future oscillators could remain phase coherent over a significant span of the distance between the earth and the sun. Again, we stress that this will enable new approaches to the way we investigate fundamental physics. Cavity stabilized lasers have already been used to search for Lorentz violations [155, 133], but space-based clocks [156] and interferometers will directly benefit from the enhanced sensitivity afforded by highly phase stable oscillators.

Working in low noise optical interferometry can often be a thankless task, but the importance of its outcomes cannot be disputed. The unofficial patron saint of cavity stabilization summarizes this nicely:

Cheers to all ye who wonder about these kinds of beautiful things - Jan L. Hall

Bibliography

- [1] S. T. Dawkins, J. J. McFerran, and A. N. Luiten, “Considerations on the measurement of the stability of oscillators with frequency counters,” IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control, vol. 54, no. 5, pp. 918–925, 2007.
- [2] A. Gillespie and F. Raab, “Thermally excited vibrations of the mirrors of laser interferometer gravitational-wave detectors,” Phys. Rev. D, vol. 52, pp. 577–585, Jul 1995.
- [3] C. C. Lam and D. H. Douglass, “Observation of oxygen impurities in single-crystal silicon by means of internal friction,” Journal of Low Temperature Physics, vol. 44, pp. 259–264, 1981.
- [4] A. White, “Frequency stabilization of gas lasers,” IEEE Journal of Quantum Electronics, vol. 1, no. 8, pp. 349–357, 1965.
- [5] R. Drever, J. L. Hall, F. Kowalski, J. Hough, G. Ford, A. Munley, and H. Ward, “Laser phase and frequency stabilization using an optical resonator,” Applied Physics B, vol. 31, no. 2, pp. 97–105, 1983.
- [6] T. Hansch and B. Couillaud, “Laser frequency stabilization by polarization spectroscopy of a reflecting reference cavity,” Optics communications, vol. 35, no. 3, pp. 441–444, 1980.
- [7] K. Numata, A. Kemery, and J. Camp, “Thermal-noise limit in the frequency stabilization of lasers with rigid cavities,” Phys. Rev. Lett., vol. 93, p. 250602, Dec 2004.
- [8] G. Harry, T. P. Bodiya, and R. DeSalvo, Optical Coatings and Thermal Noise in Precision Measurement. 2012.
- [9] V. Baglin, “Cryopumping and vacuum systems,” 2020.
- [10] J. L. Hall, J. Ye, and L.-S. Ma, “Measurement of mirror birefringence at the sub-ppm level: Proposed application to a test of qed,” Phys. Rev. A, vol. 62, p. 013815, Jun 2000.
- [11] B. A. C. O. N. B. Collaboration, “Frequency ratio measurements at 18-digit accuracy using an optical clock network,” Nature, vol. 591, no. 7851, pp. 564–569, 2021.
- [12] S. Schlamminger, K.-Y. Choi, T. A. Wagner, J. H. Gundlach, and E. G. Adelberger, “Test of the equivalence principle using a rotating torsion balance,” Phys. Rev. Lett., vol. 100, p. 041101, Jan 2008.
- [13] J. Bergé, P. Brax, G. Métris, M. Pernot-Borràs, P. Touboul, and J.-P. Uzan, “Microscope mission: First constraints on the violation of the weak equivalence principle by a light scalar dilaton,” Phys. Rev. Lett., vol. 120, p. 141101, Apr 2018.

- [14] K. Van Tilburg, N. Leefer, L. Bougas, and D. Budker, “Search for ultralight scalar dark matter with atomic spectroscopy,” *Phys. Rev. Lett.*, vol. 115, p. 011802, Jun 2015.
- [15] A. Hees, J. Guéna, M. Abgrall, S. Bize, and P. Wolf, “Searching for an oscillating massive scalar field as a dark matter candidate using atomic hyperfine frequency comparisons,” *Phys. Rev. Lett.*, vol. 117, p. 061301, Aug 2016.
- [16] P. Wcisło, P. Ablewski, K. Beloy, S. Bilicki, M. Bober, R. Brown, R. Fasano, R. Ciuryło, H. Hachisu, T. Ido, *et al.*, “New bounds on dark matter coupling from a global network of optical atomic clocks,” *Science Advances*, vol. 4, no. 12, p. eaau4869, 2018.
- [17] W. R. Milner, J. M. Robinson, C. J. Kennedy, T. Bothwell, D. Kedar, D. G. Matei, T. Legero, U. Sterr, F. Riehle, H. Leopardi, T. M. Fortier, J. A. Sherman, J. Levine, J. Yao, J. Ye, and E. Oelker, “Demonstration of a timescale based on a stable optical carrier,” *Phys. Rev. Lett.*, vol. 123, p. 173201, Oct 2019.
- [18] J. M. Robinson, E. Oelker, W. R. Milner, D. Kedar, W. Zhang, T. Legero, D. G. Matei, S. Häfner, F. Riehle, U. Sterr, and J. Ye, “Thermal noise and mechanical loss of sio₂/ta₂o₅ optical coatings at cryogenic temperatures,” *Opt. Lett.*, vol. 46, pp. 592–595, Feb 2021.
- [19] G. D. Cole, W. Zhang, M. J. Martin, J. Ye, and M. Aspelmeyer, “Tenfold reduction of brownian noise in high-reflectivity optical coatings,” *Nature Photonics*, vol. 7, no. 8, pp. 644–650, 2013.
- [20] J. Yu, D. Kedar, S. Häfner, T. Legero, F. Riehle, S. Herbers, D. Nicolodi, C. Y. Ma, J. M. Robinson, E. Oelker, J. Ye, and U. Sterr, “Excess noise in highly reflective crystalline mirror coatings,” 2022.
- [21] I. Bloch, J. Dalibard, and S. Nascimbène, “Quantum simulations with ultracold quantum gases,” *Nature Physics*, vol. 8, pp. 267–276, Apr. 2012.
- [22] S. Lloyd, “Universal quantum simulators,” *Science*, vol. 273, no. 5278, pp. 1073–1078, 1996.
- [23] C. Gross and I. Bloch, “Quantum simulations with ultracold atoms in optical lattices,” *Science*, vol. 357, no. 6355, pp. 995–1001, 2017.
- [24] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell, “Observation of bose-einstein condensation in a dilute atomic vapor,” *Science*, vol. 269, no. 5221, pp. 198–201, 1995.
- [25] B. DeMarco and D. S. Jin, “Onset of fermi degeneracy in a trapped atomic gas,” *Science*, vol. 285, no. 5434, pp. 1703–1706, 1999.
- [26] P. Törmä and K. Sengstock, eds., *Quantum Gas Experiments - Exploring Many-Body States*. Cold Atoms Series, Imperial College Press, 2014.
- [27] J. L. Bohn, A. M. Rey, and J. Ye, “Cold molecules: Progress in quantum engineering of chemistry and quantum matter,” *Science*, vol. 357, no. 6355, pp. 1002–1010, 2017.
- [28] W. S. Bakr, J. I. Gillen, A. Peng, S. Fölling, and M. Greiner, “A quantum gas microscope for detecting single atoms in a hubbard-regime optical lattice,” *Nature*, vol. 462, no. 7269, pp. 74–77, 2009.

- [29] L. W. Cheuk, M. A. Nichols, M. Okan, T. Gersdorf, V. V. Ramasesh, W. S. Bakr, T. Lompe, and M. W. Zwierlein, “Quantum-gas microscope for fermionic atoms,” Phys. Rev. Lett., vol. 114, p. 193001, May 2015.
- [30] J. Stewart, J. Gaebler, and D. Jin, “Using photoemission spectroscopy to probe a strongly interacting fermi gas,” Nature, vol. 454, no. 7205, pp. 744–747, 2008.
- [31] M. Bishof, Y. Lin, M. D. Swallows, A. V. Gorshkov, J. Ye, and A. M. Rey, “Resolved atomic interaction sidebands in an optical clock transition,” Phys. Rev. Lett., vol. 106, p. 250801, Jun 2011.
- [32] G. E. Marti, R. B. Hutson, A. Goban, S. L. Campbell, N. Poli, and J. Ye, “Imaging optical frequencies with 100 μ Hz precision and 1.1 μ m resolution,” Phys. Rev. Lett., vol. 120, p. 103201, Mar 2018.
- [33] A. Goban, R. Hutson, G. Marti, S. Campbell, M. Perlin, P. Julienne, J. D’incao, A. Rey, and J. Ye, “Emergence of multi-body interactions in a fermionic lattice clock,” Nature, vol. 563, no. 7731, pp. 369–373, 2018.
- [34] J. L. Hall, “Nobel lecture: Defining and measuring optical frequencies,” Rev. Mod. Phys., vol. 78, pp. 1279–1295, Nov 2006.
- [35] T. W. Hänsch, “Nobel lecture: Passion for precision,” Rev. Mod. Phys., vol. 78, pp. 1297–1309, Nov 2006.
- [36] M. M. Boyd, T. Zelevinsky, A. D. Ludlow, S. M. Foreman, S. Blatt, T. Ido, and J. Ye, “Optical atomic coherence at the 1-second time scale,” Science, vol. 314, no. 5804, pp. 1430–1433, 2006.
- [37] T. Bothwell, C. J. Kennedy, A. Aepli, D. Kedar, J. M. Robinson, E. Oelker, A. Staron, and J. Ye, “Resolving the gravitational redshift across a millimetre-scale atomic sample,” Nature, vol. 602, no. 7897, pp. 420–424, 2022.
- [38] A. W. Young, W. J. Eckner, W. R. Milner, D. Kedar, M. A. Norcia, E. Oelker, N. Schine, J. Ye, and A. M. Kaufman, “Half-minute-scale atomic coherence and high relative stability in a tweezer clock,” Nature, vol. 588, no. 7838, pp. 408–413, 2020.
- [39] D. G. Matei, T. Legero, C. Grebing, S. Häfner, C. Lisdat, R. Weyrich, W. Zhang, L. Sonderhouse, J. M. Robinson, F. Riehle, J. Ye, and U. Sterr, “A second generation of low thermal noise cryogenic silicon resonators,” Journal of Physics: Conference Series, vol. 723, p. 012031, jun 2016.
- [40] D. G. Matei, T. Legero, S. Häfner, C. Grebing, R. Weyrich, W. Zhang, L. Sonderhouse, J. M. Robinson, J. Ye, F. Riehle, and U. Sterr, “1.5 μ m lasers with sub-10 mhz linewidth,” Phys. Rev. Lett., vol. 118, p. 263202, Jun 2017.
- [41] E. Oelker, R. Hutson, C. Kennedy, L. Sonderhouse, T. Bothwell, A. Goban, D. Kedar, C. San-ner, J. Robinson, G. Marti, D. Matei, T. Legero, M. Giunta, R. Holzwarth, F. Riehle, U. Sterr, and J. Ye, “Demonstration of 4.8×10^{-17} stability at 1 s for two independent optical clocks,” Nature Photonics, vol. 13, p. 714–719, 10 2019.

- [42] W. Zhang, J. M. Robinson, L. Sonderhouse, E. Oelker, C. Benko, J. L. Hall, T. Legero, D. G. Matei, F. Riehle, U. Sterr, and J. Ye, “Ultrastable silicon cavity in a continuously operating closed-cycle cryostat at 4 k,” Phys. Rev. Lett., vol. 119, p. 243601, Dec 2017.
- [43] J. M. Robinson, E. Oelker, W. R. Milner, W. Zhang, T. Legero, D. G. Matei, F. Riehle, U. Sterr, and J. Ye, “Crystalline optical cavity at 4 k with thermal-noise-limited instability and ultralow drift,” Optica, vol. 6, pp. 240–243, Feb 2019.
- [44] B. J. Bloom, T. L. Nicholson, J. R. Williams, S. L. Campbell, M. Bishof, X. Zhang, W. Zhang, S. L. Bromley, and J. Ye, “An Optical Lattice Clock with Accuracy and Stability at the 10^{-18} Level,” Nature, vol. 506, pp. 71–75, 2014.
- [45] T. Bothwell, D. Kedar, E. Oelker, J. Robinson, S. Bromley, W. Tew, J. Ye, and C. Kennedy, “Jila sri optical lattice clock with uncertainty of 2.0×10^{-18} ,” Metrologia, vol. 56, p. 065004, 2019-10 2019.
- [46] T. Bothwell, A Wannier-Stark Optical Lattice Clock With Extended Coherence Times. PhD thesis, 2022.
- [47] C. J. Kennedy, E. Oelker, J. M. Robinson, T. Bothwell, D. Kedar, W. R. Milner, G. E. Marti, A. Derevianko, and J. Ye, “Precision metrology meets cosmology: Improved constraints on ultralight dark matter from atom-cavity frequency comparisons,” Phys. Rev. Lett., vol. 125, p. 201302, Nov 2020.
- [48] D. Meiser, J. Ye, D. R. Carlson, and M. J. Holland, “Prospects for a millihertz-linewidth laser,” Phys. Rev. Lett., vol. 102, p. 163601, Apr 2009.
- [49] J. Olson, R. W. Fox, T. M. Fortier, T. F. Sheerin, R. C. Brown, H. Leopardi, R. E. Stoner, C. W. Oates, and A. D. Ludlow, “Ramsey-bordé matter-wave interferometry for laser frequency stabilization at 10^{-16} frequency instability and below,” Phys. Rev. Lett., vol. 123, p. 073202, Aug 2019.
- [50] J. G. Bohnet, Z. Chen, J. M. Weiner, D. Meiser, M. J. Holland, and J. K. Thompson, “A steady-state superradiant laser with less than one intracavity photon,” Nature, vol. 484, pp. 78–81, April 2012.
- [51] M. Thorpe, T. Fortier, M. Kirchner, T. Rosenband, and L. Rippe, “Frequency-stabilization to 6×10^{-16} via spectral-hole burning,” 2011-09-11 2011.
- [52] K. G. D. S., F. F., Z. S., B. S. C., S. S.A., B. S., and S. F., “Towards an active optical clock using an optical conveyor within a ring cavity,” in 2021 Joint Conference of the European Frequency and Time Forum and IEEE International Frequency Control Symposium (EFTF/IFCS), pp. 1–5, 2021.
- [53] The LIGO Scientific Collaboration, “Advanced LIGO,” Classical and Quantum Gravity, vol. 32, p. 074001, mar 2015.
- [54] The Einstein Telescope Collaboration, “The einstein telescope: a third-generation gravitational wave observatory,” Classical and Quantum Gravity, vol. 27, p. 194002, sep 2010.
- [55] The KAGRA Collaboration, “Large-scale cryogenic gravitational-wave telescope in japan: Kagra,” Journal of Physics: Conference Series, vol. 610, no. 1, p. 012016, 2015.

- [56] The LIGO Scientific Collaboration, “Exploring the sensitivity of next generation gravitational wave detectors,” Classical and Quantum Gravity, vol. 34, p. 044001, Feb. 2017.
- [57] J. Cripe, T. Cullen, Y. Chen, P. Heu, D. Follman, G. D. Cole, and T. Corbitt, “Quantum backaction cancellation in the audio band,” Phys. Rev. X, vol. 10, p. 031065, Sep 2020.
- [58] M. T. Nery, J. R. Venneberg, N. Aggarwal, G. D. Cole, T. Corbitt, J. Cripe, R. Lanza, and B. Willke, “Laser power stabilization via radiation pressure,” Opt. Lett., vol. 46, pp. 1946–1949, Apr 2021.
- [59] N. Aggarwal, T. J. Cullen, J. Cripe, G. D. Cole, R. Lanza, A. Libson, D. Follman, P. Heu, T. Corbitt, and N. Mavalvala, “Room-temperature optomechanical squeezing,” Nature Physics, vol. 16, no. 7, pp. 784–788, 2020.
- [60] P. Fulda, PRECISION INTERFEROMETRY IN A NEW SHAPE: HIGHER-ORDER LAGUERRE-GAUSS MODES FOR GRAVITATIONAL WAVE DETECTION. PhD thesis, 2012.
- [61] E. D’Ambrosio, R. O’Shaughnessy, K. Thorne, P. Willems, S. Strigin, and S. Vyatchanin, “Advanced ligo: non-gaussian beams,” Classical and Quantum Gravity, vol. 21, no. 5, p. S867, 2004.
- [62] T. Kessler, C. Hagemann, C. Grebing, T. Legero, U. Sterr, F. Riehle, M. Martin, L. Chen, and J. Ye, “A sub-40-mhz-linewidth laser based on a silicon single-crystal optical cavity,” Nature Photonics, vol. 6, no. 10, pp. 687–692, 2012.
- [63] D. Kedar, J. Yu, E. Oelker, A. Staron, W. R. Milner, J. M. Robinson, T. Legero, F. Riehle, U. Sterr, and J. Ye, “Frequency stability of cryogenic silicon cavities with semiconductor crystalline coatings,” 2022.
- [64] E. Hecht, Optics. Pearson, 2012.
- [65] B. E. A. Saleh and M. C. Teich., Electro-Optics, ch. 18, pp. 696–736. John Wiley Sons, Ltd, 1991.
- [66] H. B. Callen and T. A. Welton, “Irreversibility and generalized noise,” Phys. Rev., vol. 83, pp. 34–40, Jul 1951.
- [67] R. Kubo, “The fluctuation-dissipation theorem,” Reports on Progress in Physics, vol. 29, p. 255, jan 1966.
- [68] P. R. Saulson, “Thermal noise in mechanical experiments,” Phys. Rev. D, vol. 42, pp. 2437–2445, Oct 1990.
- [69] Y. Levin, “Internal thermal noise in the ligo test masses: A direct approach,” Phys. Rev. D, vol. 57, pp. 659–663, Jan 1998.
- [70] F. Bondu, P. Hello, and J.-Y. Vinet, “Thermal noise in mirrors of interferometric gravitational wave antennas,” Physics Letters A, vol. 246, no. 3, pp. 227–236, 1998.
- [71] Y. T. Liu and K. S. Thorne, “Thermoelastic noise and homogeneous thermal noise in finite sized gravitational-wave test masses,” Phys. Rev. D, vol. 62, p. 122002, Nov 2000.

- [72] Optical Coatings and Thermal Noise in Precision Measurement. Cambridge University Press, 2012.
- [73] I. Martin, H. Armandula, C. Comtet, M. M. Fejer, A. Grefarsson, G. Harry, J. Hough, J.-M. M. Mackowski, I. MacLaren, C. Michel, J.-L. Montorio, N. Morgado, R. Nawrodt, S. Penn, S. Reid, A. Remillieux, R. Route, S. Rowan, C. Schwarz, P. Seidel, W. Vodel, and A. Zimmer, “Measurements of a low-temperature mechanical dissipation peak in a single layer of ta_2o_5 doped with tio_2 ,” Classical and Quantum Gravity, vol. 25, p. 055005, feb 2008.
- [74] V. B. Braginsky, V. P. Mitrofanov, V. I. Panov, and R. V. Krotkov, “Systems with small dissipation,” 1986.
- [75] 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 $\text{gaas/al}_{0.92}\text{ga}_{0.08}\text{as}$ bragg mirrors,” Metrologia, vol. 53, p. 860, mar 2016.
- [76] R. V. Pound, “Electronic frequency stabilization of microwave oscillators,” Review of Scientific Instruments, vol. 17, no. 11, pp. 490–505, 1946.
- [77] L. Cutler and C. Searle, “Some aspects of the theory and measurement of frequency fluctuations in frequency standards,” Proceedings of the IEEE, vol. 54, no. 2, pp. 136–154, 1966.
- [78] J. Rutman and F. Walls, “Characterization of frequency stability in precision frequency sources,” Proceedings of the IEEE, vol. 79, no. 7, pp. 952–960, 1991.
- [79] G. J. Dick, “Local oscillator induced instabilities in trapped ion frequency standards,” in Proceedings of the 19th Annual Precise Time and Time Interval Systems and Applications Meeting, pp. 133–147, 1989.
- [80] G. Santarelli, C. Audoin, A. Makdissi, P. Laurent, G. Dick, and A. Clairon, “Frequency stability degradation of an oscillator slaved to a periodically interrogated atomic resonator,” IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control, vol. 45, no. 4, pp. 887–894, 1998.
- [81] G. J. Dick, J. D. Prestage, C. A. Greenhall, and L. Maleki, “Local oscillator induced degradation of medium-term stability in passive atomic frequency standards,” in Proceedings of the 22th Annual Precise Time and Time Interval Systems and Applications Meeting, pp. 487–508, 1990.
- [82] C. Greenhall, “A derivation of the long-term degradation of a pulsed atomic frequency standard from a control-loop model,” IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control, vol. 45, no. 4, pp. 895–898, 1998.
- [83] C. Audoin, G. Santarelli, A. Makdissi, and A. Clairon, “Properties of an oscillator slaved to a periodically interrogated atomic resonator,” IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control, vol. 45, no. 4, pp. 877–886, 1998.
- [84] M. L. Day, P. J. Low, B. White, R. Islam, and C. Senko, “Limits on atomic qubit control from laser noise,” npj Quantum Information, vol. 8, no. 1, pp. 1–10, 2022.
- [85] R. W. Fox, “Temperature analysis of low-expansion fabry-perot cavities,” Opt. Express, vol. 17, pp. 15023–15031, Aug 2009.

- [86] E. Wiens, Q.-F. Chen, I. Ernsting, H. Luckmann, U. Rosowski, A. Nevsky, and S. Schiller, “Silicon single-crystal cryogenic optical resonator,” *Opt. Lett.*, vol. 39, pp. 3242–3245, Jun 2014.
- [87] R. Nawrodt, A. Zimmer, T. Koettig, C. Schwarz, D. Heinert, M. Hudl, R. Neubert, M. Thürk, S. Nietzsche, W. Vodel, P. Seidel, and A. Tünnermann, “High mechanical q-factor measurements on silicon bulk samples,” *Journal of Physics: Conference Series*, vol. 122, p. 012008, jul 2008.
- [88] J. J. Wortman and R. A. Evans, “Young’s modulus, shear modulus, and poisson’s ratio in silicon and germanium,” *Journal of Applied Physics*, vol. 36, no. 1, pp. 153–156, 1965.
- [89] S. D. Penn, M. M. Kinley-Hanlon, I. A. O. MacMillan, P. Heu, D. Follman, C. Deutsch, G. D. Cole, and G. M. Harry, “Mechanical ringdown studies of large-area substrate-transferred gaas/algaas crystalline coatings,” *J. Opt. Soc. Am. B*, vol. 36, pp. C15–C21, Apr 2019.
- [90] A. Schroeter, R. Nawrodt, R. Schnabel, S. Reid, I. Martin, S. Rowan, C. Schwarz, T. Koettig, R. Neubert, M. Thürk, W. Vodel, A. Tünnermann, K. Danzmann, and P. Seidel, “On the mechanical quality factors of cryogenic test masses from fused silica and crystalline quartz,” 2007.
- [91] S. D. Penn, P. H. Sneddon, H. Armandula, J. C. Betzwieser, G. Cagnoli, J. Camp, D. R. M. Crooks, M. M. Fejer, A. M. Gretarsson, G. M. Harry, J. Hough, S. E. Kittelberger, M. J. Mortonson, R. Route, S. Rowan, and C. C. Vassiliou, “Mechanical loss in tantala/silica dielectric mirror coatings,” *Classical and Quantum Gravity*, vol. 20, p. 2917, jun 2003.
- [92] S. D. Penn, P. H. Sneddon, H. Armandula, J. C. Betzwieser, G. Cagnoli, J. Camp, D. R. M. Crooks, M. M. Fejer, A. M. Gretarsson, G. M. Harry, J. Hough, S. E. Kittelberger, M. J. Mortonson, R. Route, S. Rowan, and C. C. Vassiliou, “Mechanical loss in tantala/silica dielectric mirror coatings,” *Classical and Quantum Gravity*, vol. 20, p. 2917, jun 2003.
- [93] D. R. M. Crooks, G. Cagnoli, M. M. Fejer, A. Gretarsson, G. Harry, J. Hough, N. Nakagawa, S. Penn, R. Route, S. Rowan, and P. H. Sneddon, “Experimental measurements of coating mechanical loss factors,” *Classical and Quantum Gravity*, vol. 21, p. S1059, feb 2004.
- [94] G. M. Harry, M. R. Abernathy, A. E. Becerra-Toledo, H. Armandula, E. Black, K. Dooley, M. Eichenfield, C. Nwabugwu, A. Villar, D. R. M. Crooks, G. Cagnoli, J. Hough, C. R. How, I. MacLaren, P. Murray, S. Reid, S. Rowan, P. H. Sneddon, M. M. Fejer, R. Route, S. D. Penn, P. Ganau, J.-M. Mackowski, C. Michel, L. Pinard, and A. Remillieux, “Titania-doped tantala/silica coatings for gravitational-wave detection,” *Classical and Quantum Gravity*, vol. 24, p. 405, dec 2006.
- [95] P. G. Murray, *Measurement of the mechanical loss of test mass materials for advanced gravitational wave detectors*. PhD thesis, 2008.
- [96] M. Granata, E. Coillet, V. Martinez, V. Dolique, A. Amato, M. Canepa, J. Margueritat, C. Martinet, A. Mermet, C. Michel, L. Pinard, B. Sassolas, and G. Cagnoli, “Correlated evolution of structure and mechanical loss of a sputtered silica film,” *Phys. Rev. Mater.*, vol. 2, p. 053607, May 2018.

- [97] G. Vajente, L. Yang, A. Davenport, M. Fazio, A. Ananyeva, L. Zhang, G. Billingsley, K. Prasai, A. Markosyan, R. Bassiri, M. M. Fejer, M. Chicoine, F. m. c. Schiettekatte, and C. S. Menoni, “Low mechanical loss tio_2 : geo_2 coatings for reduced thermal noise in gravitational wave interferometers,” Phys. Rev. Lett., vol. 127, p. 071101, Aug 2021.
- [98] M. Granata, A. Amato, L. Balzarini, M. Canepa, J. Degallaix, D. Forest, V. Dolique, L. Mereni, C. Michel, L. Pinard, et al., “Amorphous optical coatings of present gravitational-wave interferometers,” Classical and Quantum Gravity, vol. 37, no. 9, p. 095004, 2020.
- [99] M. Evans et al., “A Horizon Study for Cosmic Explorer: Science, Observatories, and Community,” 9 2021.
- [100] B. Shyam, K. H. Stone, R. Bassiri, M. M. Fejer, M. F. Toney, and A. Mehta, “Measurement and modeling of short and medium range order in amorphous ta_2o_5 thin films,” Scientific reports, vol. 6, no. 1, pp. 1–7, 2016.
- [101] L. Yang, M. Fazio, G. Vajente, A. Ananyeva, G. Billingsley, A. Markosyan, R. Bassiri, M. M. Fejer, and C. S. Menoni, “Structural evolution that affects the room-temperature internal friction of binary oxide nanolaminates: implications for ultrastable optical cavities,” ACS Applied Nano Materials, vol. 3, no. 12, pp. 12308–12313, 2020.
- [102] K. Prasai, J. Jiang, A. Mishkin, B. Shyam, S. Angelova, R. Birney, D. Drabold, M. Fazio, E. Gustafson, G. Harry, et al., “High precision detection of change in intermediate range order of amorphous zirconia-doped tantala thin films due to annealing,” Physical review letters, vol. 123, no. 4, p. 045501, 2019.
- [103] G. Cole. private communication, Dec. 2022.
- [104] M. Aspelmeyer, T. J. Kippenberg, and F. Marquardt, “Cavity optomechanics,” Rev. Mod. Phys., vol. 86, pp. 1391–1452, Dec 2014.
- [105] B. Slagmolen, D. Shaddock, M. Gray, and D. McClelland, “Frequency stability of spatial mode interference (tilt) locking,” IEEE Journal of Quantum Electronics, vol. 38, no. 11, pp. 1521–1528, 2002.
- [106] F. Diorico, A. Zhutov, and O. Hosten, “Laser-cavity locking at the 10^{-7} instability scale utilizing beam ellipticity,” 2022.
- [107] M. J. Martin, Quantum Metrology and Many-Body Physics: Pushing the Frontier of the Optical Lattice Clock. PhD thesis, 2013.
- [108] N. Courjal, M.-P. Bernal, A. Caspar, G. Ulliac, F. Bassignot, L. Gauthier-Manuel, and M. Suarez, “Lithium niobate optical waveguides and microwaveguides,” in Emerging Waveguide Technology (K. Y. You, ed.), ch. 8, Rijeka: IntechOpen, 2018.
- [109] 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, vol. 2, pp. 1527–1533, Sep 1985.
- [110] W. Zhang, M. J. Martin, C. Benko, J. L. Hall, J. Ye, C. Hagemann, T. Legero, U. Sterr, F. Riehle, G. D. Cole, and M. Aspelmeyer, “Reduction of residual amplitude modulation to 1×10^{-6} for frequency modulation and laser stabilization,” Opt. Lett., vol. 39, pp. 1980–1983, Apr 2014.

- [111] H. Shen, L. Li, J. Bi, J. Wang, and L. Chen, “Systematic and quantitative analysis of residual amplitude modulation in pound-drever-hall frequency stabilization,” Phys. Rev. A, vol. 92, p. 063809, Dec 2015.
- [112] J. Bi, Y. Zhi, L. Li, and L. Chen, “Suppressing residual amplitude modulation to the 1×10^{-7} level in optical phase modulation,” Appl. Opt., vol. 58, pp. 690–694, Jan 2019.
- [113] M. A. Hopcroft, W. D. Nix, and T. W. Kenny, “What is the young’s modulus of silicon?,” Journal of Microelectromechanical Systems, vol. 19, no. 2, pp. 229–238, 2010.
- [114] M. Notcutt, L.-S. Ma, J. Ye, and J. L. Hall, “Simple and compact 1-hz laser system via an improved mounting configuration of a reference cavity,” Opt. Lett., vol. 30, pp. 1815–1817, Jul 2005.
- [115] A. D. Ludlow, The Strontium Optical Lattice Clock. PhD thesis, 2008.
- [116] M. D. Swallows, M. J. Martin, M. Bishof, C. Benko, Y. Lin, S. Blatt, A. M. Rey, and J. Ye, “Operating a 87sr optical lattice clock with high precision and at high density,” IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control, vol. 59, no. 3, pp. 416–425, 2012.
- [117] S. A. Webster, M. Oxborrow, and P. Gill, “Vibration insensitive optical cavity,” Phys. Rev. A, vol. 75, p. 011801, Jan 2007.
- [118] T. Nazarova, F. Riehle, and U. Sterr, “Vibration-insensitive reference cavity for an ultra-narrow-linewidth laser,” Applied Physics B, vol. 83, no. 4, p. 531, 2006.
- [119] J. Keller, S. Ignatovich, S. A. Webster, and T. E. Mehlstäubler, “Simple vibration-insensitive cavity for laser stabilization at the 10^{16} level,” Applied Physics B, vol. 116, no. 1, pp. 203–210, 2014.
- [120] M. F. Modest, “Chapter 4 - view factors,” in Radiative Heat Transfer (Second Edition) (M. F. Modest, ed.), pp. 131–161, Burlington: Academic Press, second edition ed., 2003.
- [121] X. Zhi, R. Cao, C. Huang, K. Wang, and L. Qiu, “Feasibility analysis of hocu₂ and gd₂o₂s as regenerative materials around 4–20 k,” Applied Thermal Engineering, vol. 192, p. 116921, 2021.
- [122] K. Shinozaki, T. Mizutani, T. Fujii, T. Onaka, T. Nakagawa, and H. Sugita, “Thermal property measurements of critical materials for spica payload module,” Physics Procedia, vol. 67, pp. 270–275, 2015. Proceedings of the 25th International Cryogenic Engineering Conference and International Cryogenic Materials Conference 2014.
- [123] S. Pucic, “Diffusion of copper into gold plating,” in 1993 IEEE Instrumentation and Measurement Technology Conference, pp. 114–117, 1993.
- [124] K. G. Lyon, G. L. Salinger, C. A. Swenson, and G. K. White, “Linear thermal expansion measurements on silicon from 6 to 340 k,” Journal of Applied Physics, vol. 48, no. 3, pp. 865–868, 1977.
- [125] B. Edlén, “The refractive index of air,” Metrologia, vol. 2, no. 2, p. 71, 1966.

- [126] K. P. Birch and M. J. Downs, “Correction to the updated edlén equation for the refractive index of air,” Metrologia, vol. 31, p. 315, jan 1994.
- [127] G. Cole, S. Ballmer, G. Billingsley, S. Cataño-Lopez, M. Fejer, P. Fritschel, A. Gretarsson, G. Harry, D. Kedar, T. Legero, C. Makarem, S. Penn, D. Reitze, J. Steinlechner, U. Sterr, S. Tanioka, G.-W. Truong, J. Ye, , and J. Yu, “Substrate-transferred gaas/algaas crystalline coatings for gravitational-wave detectors: A review of the state of the art.” 2022.
- [128] V. Braginsky, M. Gorodetsky, and S. Vyatchanin, “Thermodynamical fluctuations and photo-thermal shot noise in gravitational wave antennae,” Physics Letters A, vol. 264, no. 1, pp. 1–10, 1999.
- [129] A. Farsi, M. Siciliani de Cumis, F. Marino, and F. Marin, “Photothermal and thermo-refractive effects in high reflectivity mirrors at room and cryogenic temperature,” Journal of Applied Physics, vol. 111, no. 4, p. 043101, 2012.
- [130] J.-Y. Vinet, “Thermal noise in advanced gravitational wave interferometric antennas: A comparison between arbitrary order hermite and laguerre gaussian modes,” Phys. Rev. D, vol. 82, p. 042003, Aug 2010.
- [131] C. Panuski, J. Goldstein, D. Englund, and R. Hamerly, “Coherent thermo-optic noise cancellation in an optical microcavity,” in Frontiers in Optics + Laser Science 2021, p. FTh6B.3, Optica Publishing Group, 2021.
- [132] H. Liu, B. D. Elwood, M. Evans, and J. Thaler, “Searching for axion dark matter with birefringent cavities,” Phys. Rev. D, vol. 100, p. 023548, Jul 2019.
- [133] H. Müller, S. Herrmann, A. Saenz, A. Peters, and C. Lämmerzahl, “Optical cavity tests of lorentz invariance for the electron,” Phys. Rev. D, vol. 68, p. 116006, Dec 2003.
- [134] E. Wiens, A. Y. Nevsky, and S. Schiller, “Resonator with ultrahigh length stability as a probe for equivalence-principle-violating physics,” Phys. Rev. Lett., vol. 117, p. 271102, Dec 2016.
- [135] M. Aaboud, G. Aad, B. Abbott, B. Abeloos, S. Abidi, O. AbouZeid, N. L. Abraham, H. Abramowicz, H. Abreu, R. Abreu, et al., “Search for dark matter and other new phenomena in events with an energetic jet and large missing transverse momentum using the atlas detector,” Journal of High Energy Physics, vol. 2018, no. 1, pp. 1–53, 2018.
- [136] E. G. Ferreira, “Ultra-light dark matter,” The Astronomy and Astrophysics Review, vol. 29, no. 1, pp. 1–186, 2021.
- [137] D. Antypas, A. Banerjee, C. Bartram, M. Baryakhtar, J. Betz, J. J. Bollinger, C. Boutan, D. Bowring, D. Budker, D. Carney, G. Carosi, S. Chaudhuri, S. Cheong, A. Chou, M. D. Chowdhury, R. T. Co, J. R. C. López-Urrutia, M. Demarteau, N. DePorzio, A. V. Derbin, T. Deshpande, M. D. Chowdhury, L. Di Luzio, A. Diaz-Morcillo, J. M. Doyle, A. Drlica-Wagner, A. Droster, N. Du, B. Döbrich, J. Eby, R. Essig, G. S. Farren, N. L. Figueroa, J. T. Fry, S. Gardner, A. A. Geraci, A. Ghalsasi, S. Ghosh, M. Giannotti, B. Gimeno, S. M. Griffin, D. Grin, D. Grin, H. Grote, J. H. Gundlach, M. Guzzetti, D. Hanneke, R. Harnik, R. Henning, V. Irsic, H. Jackson, D. F. J. Kimball, J. Jaeckel, M. Kagan, D. Kedar, R. Khatiwada, S. Knirck, S. Kolkowitz, T. Kovachy, S. E. Kuenstner, Z. Lasner, A. F. Leder, R. Lehnert, D. R. Leibbrandt, E. Lentz, S. M. Lewis, Z. Liu, J. Manley, R. H. Maruyama, A. J. Millar, V. N.

- Muratova, N. Musoke, S. Nagaitsev, O. Noroozian, C. A. J. O'Hare, J. L. Ouellet, K. M. W. Pappas, E. Peik, G. Perez, A. Phipps, N. M. Rapidis, J. M. Robinson, V. H. Robles, K. K. Rogers, J. Rudolph, G. Rybka, M. Safdari, M. Safdari, M. S. Safronova, C. P. Salemi, P. O. Schmidt, T. Schumm, A. Schwartzman, J. Shu, M. Simanovskaia, J. Singh, S. Singh, M. S. Smith, W. M. Snow, Y. V. Stadnik, C. Sun, A. O. Sushkov, T. M. P. Tait, V. Takhistov, D. B. Tanner, D. J. Temples, P. G. Thirolf, J. H. Thomas, M. E. Tobar, O. Tretiak, Y. D. Tsai, J. A. Tyson, M. Vandegar, S. Vermeulen, L. Visinelli, E. Vitagliano, Z. Wang, D. J. Wilson, L. Winslow, S. Withington, M. Wooten, J. Yang, J. Ye, B. A. Young, F. Yu, M. H. Zaheer, T. Zelevinsky, Y. Zhao, and K. Zhou, "New horizons: Scalar and vector ultralight dark matter," 2022.
- [138] A. Derevianko and M. Pospelov, "Hunting for topological dark matter with atomic clocks," *Nature Physics*, vol. 10, no. 12, pp. 933–936, 2014.
- [139] A. Arvanitaki, J. Huang, and K. Van Tilburg, "Searching for dilaton dark matter with atomic clocks," *Phys. Rev. D*, vol. 91, p. 015015, Jan 2015.
- [140] M. S. Safronova, D. Budker, D. DeMille, D. F. J. Kimball, A. Derevianko, and C. W. Clark, "Search for new physics with atoms and molecules," *Rev. Mod. Phys.*, vol. 90, p. 025008, Jun 2018.
- [141] V. A. Dzuba, V. V. Flambaum, and J. K. Webb, "Calculations of the relativistic effects in many-electron atoms and space-time variation of fundamental constants," *Phys. Rev. A*, vol. 59, pp. 230–237, Jan 1999.
- [142] V. A. Dzuba, V. V. Flambaum, and J. K. Webb, "Space-time variation of physical constants and relativistic corrections in atoms," *Phys. Rev. Lett.*, vol. 82, pp. 888–891, Feb 1999.
- [143] V. Flambaum and V. Dzuba, "Search for variation of the fundamental constants in atomic, molecular, and nuclear spectra," *Canadian Journal of Physics*, vol. 87, no. 1, pp. 25–33, 2009.
- [144] Y. V. Stadnik and V. V. Flambaum, "Enhanced effects of variation of the fundamental constants in laser interferometers and application to dark-matter detection," *Phys. Rev. A*, vol. 93, p. 063630, Jun 2016.
- [145] M. Filzinger, S. Dörscher, R. Lange, J. Klose, M. Steinle, E. Benkler, E. Peik, C. Lisdat, and N. Huntemann, "Improved limits on the coupling of ultralight bosonic dark matter to photons from optical atomic clock comparisons," 2023.
- [146] M. S. Safronova, S. G. Porsev, C. Sanner, and J. Ye, "Two clock transitions in neutral yb for the highest sensitivity to variations of the fine-structure constant," *Phys. Rev. Lett.*, vol. 120, p. 173001, Apr 2018.
- [147] C. Hagemann, C. Grebing, C. Lisdat, S. Falke, T. Legero, U. Sterr, F. Riehle, M. J. Martin, and J. Ye, "Ultrastable laser with average fractional frequency drift rate below $5 \times 10^{-19}/\text{s}$," *Opt. Lett.*, vol. 39, pp. 5102–5105, Sep 2014.
- [148] J. Weber, "Gravitational-wave-detector events," *Phys. Rev. Lett.*, vol. 20, pp. 1307–1308, Jun 1968.
- [149] Prodi, G. A. and others, "The ultracryogenic gravitational wave detector auriga," in *International Conference on Gravitational Waves: Sources and Detectors*, 1996.

- [150] A. de Waard, L. Gottardi, J. van Houwelingen, A. Shumack, and G. Frossati, “Minigrail, the first spherical detector,” *Classical and Quantum Gravity*, vol. 20, p. S143, apr 2003.
- [151] B. Abbott, R. Abbott, R. Adhikari, J. Agresti, P. Ajith, B. Allen, R. Amin, S. B. Anderson, W. G. Anderson, M. Arain, M. Araya, H. Armandula, M. Ashley, S. Aston, P. Aufmuth, C. Aulbert, S. Babak, S. Ballmer, H. Bantilan, B. C. Barish, C. Barker, D. Barker, B. Barr, P. Barriga, M. A. Barton, K. Bayer, K. Belczynski, J. Betzwieser, P. T. Beyersdorf, B. Bhawal, I. A. Bilenko, G. Billingsley, R. Biswas, E. Black, K. Blackburn, L. Blackburn, D. Blair, B. Bland, J. Bogenstahl, L. Bogue, R. Bork, V. Boschi, S. Bose, P. R. Brady, V. B. Braginsky, J. E. Brau, M. Brinkmann, A. Brooks, D. A. Brown, A. Bullington, A. Bunkowski, A. Buonanno, M. Burgamy, O. Burmeister, D. Busby, R. L. Byer, L. Cadonati, G. Cagnoli, J. B. Camp, J. Cannizzo, K. Cannon, C. A. Cantley, J. Cao, L. Cardenas, M. M. Casey, G. Castaldi, C. Cepeda, E. Chalkey, P. Charlton, S. Chatterji, S. Chelkowski, Y. Chen, F. Chiadini, D. Chin, E. Chin, J. Chow, N. Christensen, J. Clark, P. Cochrane, T. Cokelaer, C. N. Colacino, R. Coldwell, R. Conte, D. Cook, T. Corbitt, D. Coward, D. Coyne, J. D. E. Creighton, T. D. Creighton, R. P. Croce, D. R. M. Crooks, A. M. Cruise, A. Cumming, J. Dalrymple, E. D’Ambrosio, K. Danzmann, G. Davies, D. DeBra, J. Degallaix, M. Degree, T. Demma, V. Dergachev, S. Desai, R. DeSalvo, S. Dhurandhar, M. Díaz, J. Dickson, A. Di Credico, G. Diederichs, A. Dietz, E. E. Doomes, R. W. P. Drever, J.-C. Dumas, R. J. Dupuis, J. G. Dwyer, P. Ehrens, E. Espinoza, T. Etzel, M. Evans, T. Evans, S. Fairhurst, Y. Fan, D. Fazi, M. M. Fejer, L. S. Finn, V. Fiumara, N. Fotopoulos, A. Franzen, K. Y. Franzen, A. Freise, R. Frey, T. Fricke, P. Fritschel, V. V. Frolov, M. Fyffe, V. Galdi, J. Garofoli, I. Gholami, J. A. Giaime, S. Giampanis, K. D. Giardina, K. Goda, E. Goetz, L. Goggin, G. González, S. Gossler, A. Grant, S. Gras, C. Gray, M. Gray, J. Greenhalgh, A. M. Gretarsson, R. Grosso, H. Grote, S. Grunewald, M. Guenther, R. Gustafson, B. Hage, W. O. Hamilton, D. Hammer, C. Hanna, J. Hanson, J. Harms, G. Harry, E. Harstad, T. Hayler, J. Heefner, I. S. Heng, A. Heptonstall, M. Heurs, M. Hewitson, S. Hild, E. Hirose, D. Hoak, D. Hosken, J. Hough, E. Howell, D. Hoyland, S. H. Huttner, D. Ingram, E. Innerhofer, M. Ito, Y. Itoh, A. Ivanov, D. Jackrel, B. Johnson, W. W. Johnson, D. I. Jones, G. Jones, R. Jones, L. Ju, P. Kalmus, V. Kalogera, D. Kasprzyk, E. Katsavounidis, K. Kawabe, S. Kawamura, F. Kawazoe, W. Kells, D. G. Keppel, F. Y. Khalili, C. Kim, P. King, J. S. Kissel, S. Klimenko, K. Kokeyama, V. Kondrashov, R. K. Kopparapu, D. Kozak, B. Krishnan, P. Kwee, P. K. Lam, M. Landry, B. Lantz, A. Lazzarini, B. Lee, M. Lei, J. Leiner, V. Leonhardt, I. Leonor, K. Libbrecht, P. Lindquist, N. A. Lockerbie, M. Longo, M. Lormand, M. Lubinski, H. Lück, B. Machenschalk, M. MacInnis, M. Mageswaran, K. Mailand, M. Malec, V. Mandic, S. Marano, S. Márka, J. Markowitz, E. Maros, I. Martin, J. N. Marx, K. Mason, L. Matone, V. Matta, N. Mavalvala, R. McCarthy, B. J. McCauley, D. E. McClelland, S. C. McGuire, M. McHugh, K. McKenzie, J. W. C. McNabb, S. McWilliams, T. Meier, A. Melissinos, G. Mendell, R. A. Mercer, S. Meshkov, E. Messaritaki, C. J. Messenger, D. Meyers, E. Mikhailov, P. Miller, S. Mitra, V. P. Mitrofanov, G. Mitselmakher, R. Mittleman, O. Miyakawa, S. Mohanty, V. Moody, G. Moreno, K. Mossavi, C. MowLowry, A. Moylan, D. Mudge, G. Mueller, S. Mukherjee, H. Müller-Ebhardt, J. Munch, P. Murray, E. Myers, J. Myers, T. Nash, D. Nettles, G. Newton, A. Nishizawa, K. Numata, B. O’Reilly, R. O’Shaughnessy, D. J. Ottaway, H. Overmier, B. J. Owen, H.-J. Paik, Y. Pan, M. A. Papa, V. Parameshwaraiah, P. Patel, M. Pedraza, S. Penn, V. Pierro, I. M. Pinto, M. Pitkin, H. Pletsch, M. V. Plissi, F. Postiglione, R. Prix, V. Quetschke, F. Raab, D. Rabeling, H. Radkins, R. Rahkola, N. Rainer, M. Rakhmanov, M. Ramsunder, K. Rawlins, S. Ray-Majumder,

- V. Re, H. Rehbein, S. Reid, D. H. Reitze, L. Ribichini, R. Riesen, K. Riles, B. Rivera, N. A. Robertson, C. Robinson, E. L. Robinson, S. Roddy, A. Rodriguez, A. M. Rogan, J. Rollins, J. D. Romano, J. Romie, R. Route, S. Rowan, A. Rüdiger, L. Ruet, P. Russell, K. Ryan, S. Sakata, M. Samidi, L. Sancho de la Jordana, V. Sandberg, V. Sannibale, S. Saraf, P. Sarin, B. S. Sathyaprakash, S. Sato, P. R. Saulson, R. Savage, P. Savov, S. Schediwy, R. Schilling, R. Schnabel, R. Schofield, B. F. Schutz, P. Schwinberg, S. M. Scott, A. C. Searle, B. Sears, F. Seifert, D. Sellers, A. S. Sengupta, P. Shawhan, D. H. Shoemaker, A. Sibley, J. A. Sidles, X. Siemens, D. Sigg, S. Sinha, A. M. Sintes, B. J. J. Slagmolen, J. Slutsky, J. R. Smith, M. R. Smith, K. Somiya, K. A. Strain, D. M. Strom, A. Stuver, T. Z. Summerscales, K.-X. Sun, M. Sung, P. J. Sutton, H. Takahashi, D. B. Tanner, M. Tarallo, R. Taylor, R. Taylor, J. Thacker, K. A. Thorne, K. S. Thorne, A. Thüring, K. V. Tokmakov, C. Torres, C. Torrie, G. Traylor, M. Trias, W. Tyler, D. Ugolini, C. Ungarelli, K. Urbanek, H. Vahlbruch, M. Vallisneri, C. Van Den Broeck, M. Varvella, S. Vass, A. Vecchio, J. Veitch, P. Veitch, A. Villar, C. Vorvick, S. P. Vyachanin, S. J. Waldman, L. Wallace, H. Ward, R. Ward, K. Watts, J. Weaver, D. Webber, A. Weber, A. Weidner, M. Weinert, A. Weinstein, R. Weiss, S. Wen, K. Wette, J. T. Whelan, D. M. Whitbeck, S. E. Whitcomb, B. F. Whiting, C. Wilkinson, P. A. Willems, L. Williams, B. Willke, I. Wilmot, W. Winkler, C. C. Wipf, S. Wise, A. G. Wiseman, G. Woan, D. Woods, R. Wooley, J. Worden, W. Wu, I. Yakushin, H. Yamamoto, Z. Yan, S. Yoshida, N. Yunes, M. Zanolin, J. Zhang, L. Zhang, P. Zhang, C. Zhao, N. Zotov, M. Zucker, H. zur Mühlen, and J. Zweizig, “First cross-correlation analysis of interferometric and resonant-bar gravitational-wave data for stochastic backgrounds,” *Phys. Rev. D*, vol. 76, p. 022001, Jul 2007.
- [152] G. González, “Suspensions thermal noise in the ligo gravitational wave detector,” *Classical and Quantum Gravity*, vol. 17, p. 4409, nov 2000.
- [153] L. A. Wainstein and V. D. Zubakov, “Extraction of signals from noise,” *American Journal of Physics*, vol. 31, no. 1, pp. 68–68, 1963.
- [154] *Gravitational-Wave Detectors*, ch. 6, pp. 197–267. John Wiley & Sons, Ltd, 2011.
- [155] C. Braxmaier, H. Müller, O. Pradl, J. Mlynek, A. Peters, and S. Schiller, “Tests of relativity using a cryogenic optical resonator,” *Phys. Rev. Lett.*, vol. 88, p. 010401, Dec 2001.
- [156] S. Kolkowitz, I. Pikovski, N. Langellier, M. D. Lukin, R. L. Walsworth, and J. Ye, “Gravitational wave detection with optical lattice atomic clocks,” *Phys. Rev. D*, vol. 94, p. 124043, Dec 2016.
- [157] R. Nawrodt, C. Schwarz, S. Kroker, I. W. Martin, R. Bassiri, F. Brückner, L. Cunningham, G. D. Hammond, D. Heinert, J. Hough, T. Käsebier, E.-B. Kley, R. Neubert, S. Reid, S. Rowan, P. Seidel, and A. Tünnermann, “Investigation of mechanical losses of thin silicon flexures at low temperatures,” *Classical and Quantum Gravity*, vol. 30, p. 115008, may 2013.
- [158] K. Craig, J. Steinlechner, P. G. Murray, A. S. Bell, R. Birney, K. Haughian, J. Hough, I. MacLaren, S. Penn, S. Reid, R. Robie, S. Rowan, and I. W. Martin, “Mirror coating solution for the cryogenic einstein telescope,” *Phys. Rev. Lett.*, vol. 122, p. 231102, Jun 2019.
- [159] J. T. VanderPlas, “Understanding the lomb–scargle periodogram,” *The Astrophysical Journal Supplement Series*, vol. 236, p. 16, may 2018.

- [160] P. G. Westergaard, J. Lodewyck, and P. Lemonde, “Minimizing the dick effect in an optical lattice clock,” IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control, vol. 57, no. 3, pp. 623–628, 2010.

Appendix A

Dick effect limited instability and laser noise

In the context of Dick effect limited stability, a distinction between Ramsey and Rabi spectroscopy is the unique feature that increased spectroscopic time does not always result in an improved clock stability. As we discussed in chapter 2, this is related to the spectroscopic sensitivity function of Rabi interrogation where laser phase evolution at the interrogation time $\tau/2$ is weighted more heavily than the beginning and end of spectroscopy. This concept is familiar to us in lab, and it is not difficult to measure a reduction in stability for longer Rabi interrogation times. However, it is a misconception to think that this is a universal feature of Rabi spectroscopy. This feature is instead dependent on the spectrum of the local oscillator. For particular types of frequency noise (e.g. white PM), sampling laser noise with the Rabi spectroscopic sensitivity function has the feature where increased spectroscopic time will result in an improved stability. Here we will show the dependence of clock stability for different T_c and τ when the local oscillator has different spectrums $S_y(f) \sim h_\alpha f^\alpha$.

Each figure displays a contour plot for the 1 s Dick-effect limited instability for a given spectroscopy time and dead time. We show the power spectral density of the laser's fractional frequency noise as well as the mod $\sigma_y(\tau)$. In optical clocks, we are most familiar with the thermal noise limited laser spectra of Fig.A.3. We can see the clear feature where increased Rabi pulse time results in a reduced clock stability, whereas this is not the case with Ramsey interrogation. However, lasers with excess high frequency noise show the opposite behaviour. Fig.A.1 calculates stabilities for a laser with a white phase noise spectrum, $S_y \sim f^2$. The relative insensitivity of

Rabi spectroscopy to high frequency noise (when compared to Ramsey), results in a superior clock stability for equivalent deadtimes and spectroscopic times. The point here is that optimization of clock stability requires us to give careful thought to the spectroscopic sequence and depends strongly on the interrogation oscillator's frequency noise spectrum.




Figure A.1: The laser has a white phase noise spectrum where $S_y(f) \sim f^2$. Top row shows the laser $S_y(f)$ and the corresponding modified Allan variance. The bottom row shows the Dick effect limited 1 s clock stability that can be expected for different dead times and pulse times. Different behaviour will be visible for Rabi spectroscopy (**lower left**) and Ramsey spectroscopy (**lower right**).




Figure A.2: The laser has a white frequency noise spectrum where $S_y(f) \sim f^0$.




Figure A.3: The laser has a flicker frequency noise spectrum where $S_y(f) \sim f^{-1}$.




Figure A.4: The laser has a random walk frequency noise spectrum where $S_y(f) \sim f^{-2}$.

Appendix B

Analog locking

We've discussed a few cases where reduced optical power within the cavity has associated benefits, e.g. lower drift rate in Si4 and lower birefringent noise in Si6. An attractive scenario would be to operate the cavity with optical power just above the shot noise limited length resolution without any detrimental effects. Lowering the overall power incident on the cavity can be appropriate method at the cost of a reduced cavity signal. Photodetector noise and other technical noise sources will often become consequential at these levels. A useful method employed in Si6 is to increase the modulation index of the EOM sidebands. Power in the carrier, the important metric for the aforementioned effects, is reduced while power in the sidebands increases. However, this is limited to a practical sideband/carrier ratio of ≈ 10 as illustrated in Fig.B.1. The Bessel function dictating sideband amplitude restricts the maximum power we can store in the first order modulation tones. Driving $\beta > 3$ generates stronger second order harmonics which can be harmful in terms of RAM control.

The ideal situation would be to couple only a few photons into the cavity (again, just above the length resolution limited by optical shot noise) while maintaining massive sidebands for a strong heterodyne beat. This would ensure that technical issues related to the servo lock and photodetection never become significant problems. Rather than relying on the phase modulation generated via an EOM, we can try to synthesize the PDH triplet by applying a set of independently generated tones to an AOM. Decoupled control to each tone can allow us to arbitrarily adjust their phase and amplitudes without any effect on the other two. Generally the most significant source




Figure B.1: Amplitude of different sidebands generated via phase modulation on an EOM.

of measured RAM is generated by the EOM. Triplet synthesis within the AOM is insensitive to polarization fluctuations, so it possible that the free running RAM introduced with this method is much lower.

This sounds like a novel scheme but we've already been playing around with this concept in our dual frequency locking. In that case there is no dynamic readjustment of the electric field phasors but we are able to independently control the amplitude and phase. To generate PDH

- The two sidebands need to maintain a relative phase of 180°
- We still want to implement active RAM cancellation so need to be able to feed back to the

phase and amplitude of the sidebands

- As discussed in the RAM section of chapter 3, this requires applying common phase offsets $\phi_+ = \phi_i$ to the two tones, and inverse amplitude offsets $a_+ = -a_-$

Triplet synthesis via a DDS allows us direct access to the amplitude and phase of the electric fields, a unique capability that we do not enjoy when using active RAM stabilization on EOM generated sidebands. These are truly orthogonal quadratures but it is worth considering whether our in-phase and quadrature detection channels are aligned to the parameters we can modify (i.e. does the in-phase error signal solely depend on the sideband amplitude). We'll revisit a derivation of the error signals, and also demonstrate that this is indeed the case when we apply a correction signal.

In section 3.2.2 we found that the generalized electric field generated by the EOM is

$$E_i = E_0 \left(J_0(\beta) e^{i\omega t} + \alpha_+ J_1(\beta) e^{i(\omega+\Omega)t} e^{i\phi_+} - \alpha_- J_1(\beta) e^{i(\omega-\Omega)t} e^{i\phi_-} \right) \quad (\text{B.1})$$

and the corresponding detected RAM signal in the presence of AM and PM is

$$V \propto \text{Re} \left[\alpha_+ e^{-i\phi_+} - \alpha_- e^{i\phi_-} \right] \cos(\varphi) + \text{Im} \left[\alpha_+ e^{-i\phi_+} - \alpha_- e^{i\phi_-} \right] \sin(\varphi) \quad (\text{B.2})$$

A particular issue is that we will not be identify all four α_i, ϕ_i components. In AM for example, a single stage of demodulation at the sideband frequency Ω will not yield information about the individual components α_+ and α_- . The correction we apply will contain some combination of the two and we should consider how bad of an approximation this is - can this suppress the AM (and similarly the PM) to necessary levels?

First, let's consider quadrature demodulation of the error signal. The in-phase component is found to be

$$V_I \propto \text{Re} \left[\alpha_+ e^{-i\phi_+} - \alpha_- e^{i\phi_-} \right] \approx \alpha_+ \left(1 + \frac{\phi_+^2}{2} \right) - \alpha_- \left(1 + \frac{\phi_-^2}{2} \right) \quad (\text{B.3})$$

$$= (\alpha_+ - \alpha_-) - \frac{1}{2} (\alpha_+ \phi_+^2 - \alpha_- \phi_-^2) \quad (\text{B.4})$$

We can expand the overall field amplitudes of the two sidebands as

$$\alpha_+ = 1 + \epsilon_+ \quad (\text{B.5})$$

$$\alpha_- = 1 + \epsilon_- \quad (\text{B.6})$$

The corresponding in-phase signal is then

$$V_I \propto (\epsilon_+ - \epsilon_-) - \frac{1}{2}(\phi_+^2 - \phi_-^2) + \frac{1}{2}(\phi_+^2 \epsilon_+ - \phi_-^2 \epsilon_-) \quad (\text{B.7})$$

$$\approx \epsilon_+ - \epsilon_- \quad (\text{B.8})$$

For the quadrature signal, we find

$$V_Q \propto \text{Im}[\alpha_+ e^{-i\phi_+} - \alpha_- e^{i\phi_-}] \approx \alpha_+ \phi_+ + \alpha_- \phi_- \quad (\text{B.9})$$

$$\approx (\phi_+ + \phi_-) - (\phi_+ \epsilon_+ + \phi_- \epsilon_-) \quad (\text{B.10})$$

$$\approx \phi_+ + \phi_- \quad (\text{B.11})$$

Crucially, this tells us how to properly apply feedback to the RAM signal to suppress both components. Recall that $\phi_{\pm}$ are defined as common mode rotations of the two sideband phasors, i.e. will result in upper and lower sidebands with respective phases of ϕ_+ and $-\pi + \phi_-$. In our first order approximation in the presence of zero residual offset, $\phi_+ = -\phi_-$ will cause no quadrature RAM signal. Unless these terms are zero, there will still be a contribution apparent at the in-phase channel V_I .

By directly modifying the phase of the DDS-synthesized sidebands, we can apply a phase correction to the upper and lower sideband electric fields as $\phi_c = -(\phi_+ + \phi_-)/2$. Each sideband therefore acquires an additional phase shift as

$$E_{upper,lower} = J_1(\beta) e^{i(\omega \pm \Omega)t} e^{i(\phi_c + \phi_{\pm})} \quad (\text{B.12})$$

Consider how this modifies the in-phase signal. We can substitute

$$\phi_+ \rightarrow \phi_+ + \phi_c = (\phi_+ - \phi_-)/2 \quad (\text{B.13})$$

$$\phi_- \rightarrow \phi_- + \phi_c = (\phi_- - \phi_+)/2 \quad (\text{B.14})$$

Then this results in

$$V_Q = 0 \quad (\text{B.15})$$

Substitution into the in-phase signal, we find

$$V_I \propto (\epsilon_+ - \epsilon_-) - 0 + \mathcal{O}(\epsilon_{\pm})\mathcal{O}(\phi_{\pm}^2) \quad (\text{B.16})$$

$$= \epsilon_+ - \epsilon_- \quad (\text{B.17})$$

Now the first order approximation we made is an exact reduction to the AM components. In practice, it may be easier to apply the correction phase shift only to a single sideband, such that $\phi_+ \rightarrow \phi_+ + \phi_c$, ϕ_- is unchanged, and $\phi_c = -(\phi_+ - \phi_-)$. It's therefore instructive to consider whether actuating on a single sideband produces an appreciable difference from actuating on both sidebands. We find that the quadrature signal V_Q will be nulled but the second order in-phase signal will be modified as

$$\mathcal{O}(\epsilon_{\pm})\mathcal{O}(\phi_{\pm}^2) = (\phi_+^2 - 2\phi_+\phi_- + \phi_-^2) \left(\frac{\epsilon_+ - \epsilon_-}{2} \right) \rightarrow (\phi_+^2 - 2\phi_+\phi_- + \phi_-^2) \epsilon_+ - \phi_-^2 \epsilon_- \quad (\text{B.18})$$

For small amounts of AM and PM, this is a negligible correction.

Unlike PM, a common mode amplitude change to both sidebands $\epsilon_+ = \epsilon_-$ results in a null RAM signal for the in-phase quadrature (to first order). An AM correction can be made by scaling the two amplitudes. Given our error signal, we apply a correction

$$\alpha_c = -(\epsilon_+ - \epsilon_-) \quad (\text{B.19})$$

Actuating on a single sideband, we can then perform the transformation

$$\alpha_+ \rightarrow \alpha_c + \alpha_+ = 1 + \epsilon_- \quad (\text{B.20})$$

Such that

$$V_I = 0 + \frac{1}{2}\alpha_- (\phi_+^2 - \phi_-^2). \quad (\text{B.21})$$

Notably, this section is titled 'Analog locking' rather than 'Digital locking'. The initial realization of a DDS-generated triplet found that analog feedback to the sideband amplitudes through variable voltage attenuators was more straightforward than writing digital words to the DDS.