

**Enhancing optical clocks with ultrastable lasers and
spin-squeezing**

by

John M. Robinson

B.S., University of Nevada, Las Vegas 2015

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 Science
2023

Committee Members:

Jun Ye, Chair

James K. Thompson

Ana Maria Rey

Konrad Lehnert

Joshua Combes

Robinson, John M. (Ph.D., Physics)

Enhancing optical clocks with ultrastable lasers and spin-squeezing

Thesis directed by Prof. Jun Ye

Traditional optical atomic clocks are limited in their performance by laser frequency noise and the intrinsic quantum noise of uncorrelated atoms. In this thesis, we advance the field of optical clocks on both of these fronts. By developing the next generation of ultrastable laser technology, we enable clock comparisons that have approached the quantum projection noise limit. To go beyond this limit, we build and operate an optical clock with the capability of spin-squeezing. Employing conditional spin squeezing via quantum nondemolition measurements based on cavity QED, we produce a spin squeezed state that yields a spectroscopic enhancement of 1.7 dB beyond the standard quantum limit. We then run a clock comparison between two spin squeezed clock ensembles, making use of a movable optical lattice to individually squeeze and readout the ensembles with cavity QED. This differential comparison between the two squeezed clocks directly verifies enhanced clock stability of 1.9 dB beyond the quantum projection noise limit, and reaches a measurement precision level of 10^{-17} . This constitutes the first direct demonstration of quantum enhanced measurement in an operational optical atomic clock.

Dedication

To my parents.

Acknowledgements

First off, huge thanks to Jun for being the best research advisor imaginable. When I first showed up, I was intimidated by Jun's quick thinking and questions, he seemed to be moving at 100 miles per hour. Over time, I realized that Jun's tough questions were the embodiment of his dedication to his students, he was often thinking about our problems, both technical and fundamental. He simply wants to help - both in science and in life. I'll always remember the lessons I've learned working in the Sr lab, and I'm eternally grateful for the opportunities Jun has given me.

When I first joined the lab, I spent the summer helping out Sr1, and I thank Shimon Kolkowitz, Sarah Bromley, and Toby Bothwell for being supportive in those early days. Soon enough, Jun asked me if I was interested in stable lasers, and I began my journey over in S1B20. That is when I started working with Wei Zhang. Wei is perhaps the most generous and kind person I've ever met. His eagerness to help young students, sitting for hours on end until things were clear for everyone, is truly admirable. Lindsay Sonderhouse was working on the comb and cavities at the time, and I will always appreciate her careful thinking on all sorts of issues in the lab and her eternal kindness. Eric Oelker soon joined us, and brought with him some demystifying of the stable laser dark magic. Eric somehow also knew when I needed to go on with some issue/topic on my own to grow, giving those little nudges in the right direction along the way. The more recent members of Sr1 - Alex Aeppli and Kyungtae Kim were always eager to help and have been great labmates. The newly joined member Will Warfield will surely help that experiment continue to flourish.

One wonderful aspect to stable lasers is the collaboration with our PTB colleagues. When Lindsay and I traveled to PTB to help work on the silicon cavity project at PTB, Dan Matei, Thomas Legero, Fritz Riehle and Uwe Sterr were all so very kind and helpful. Their careful attention to detail continues to be inspiring - always taking systematic steps, and somehow *always* in the right direction.

Sr2 was always a very friendly bunch. Sara Campbell became a close friend, and time spent in the lab with her was always fun. She, along with other Sr2 folks: Ross Hutson, Will Milner, Ed Marti, and Aki Goban, all became truly amazing colleagues over the years. Ross is one of the labs wizards, somehow always either knowing the answer to your question, or knowing how to figure it out. Will is incredibly giving with his time in fixing technical issues in the lab, and you can be sure to have a good conversation if you ask him a physics question. Ed and Aki were amazing post-docs in the Sr lab, each always being helpful to young graduate students. When Christian Sanner showed up, he came over to stable lasers to help out on the Si4 upgrades. Christian's helpfulness and absolute technical wizardry made for some serious fun in the stable laser room, and this carried on throughout his time at JILA. Working on Sr1 for a while was a great introduction of using stable lasers with the atoms. Colin and I worked closely on the Sr1 buildup cavity for many weeks, and during periods of setbacks, he always came into the lab with a positive attitude nonetheless. Toby is a careful scientist, and was often extremely helpful during the building of Sr3, using his experience on the Sr1 experiment to help guide us. Dhruv is a very skilled physicist, and his work on Si6 took the next steps in stable laser research in exactly the right direction.

Some of the initial Sr3 effort involved all of us on Sr1, early cavity and chamber designs were made by Colin and Toby. When Sr1 split into two experiments, I took on Sr3 and was first joined by Josie Meyer. Josie's positive attitude was always infectious, and some of the Sr3 machine would not be there without her. Shortly after, we were joined by Maya Miklos, an extremely sharp physicist, and without her, many critical aspects of Sr3 would be vastly weaker. Yee Ming Tso joined the project, and dove right into many critical aspects of the experiment. I'm certain the future of Sr3 is in great hands.

The work in this thesis would not be possible without support from many others, including my Mom, Dad, Tom, Connor, Jake, Keene, Mark, Selina, Luna. Lastly, thank you to Raven Larson, for your unwavering support and love.

Contents

Chapter	
1 Introduction	1
1.1 Characterizing clocks	3
1.2 Limits to clock performance	4
1.3 Overcoming the limits	5
1.4 Outline of this thesis	6
2 Ultrastable lasers based on optical cavities	8
2.1 Motivation and outline	8
2.2 Lasers stabilized to optical cavities	9
2.2.1 Fundamental limit - thermal noise	10
2.3 Technical noise budget of Si4	13
2.3.1 Temperature	13
2.3.2 Vibration noise	17
2.3.3 Residual amplitude modulation	20
2.3.4 Photothermal effects	23
2.3.5 Path length noise	25
2.3.6 Electronics noise	26
2.4 Stability measurements	27
2.5 Linear frequency drift : power dependence	30

2.6	Inferring loss angle of $\text{SiO}_2/\text{Ta}_2\text{O}_5$ coating	36
2.7	Si_3 : zero crossing of CTE	43
3	Classical clock demonstrations and applications	46
3.1	Classical clock comparisons at the quantum limit	46
3.2	Dark matter limits from precision spectroscopy	50
3.2.1	Sensitivity to fundamental constants	50
3.2.2	Dark matter couplings	51
3.2.3	Dark matter Limit	54
3.2.4	Extra details on our analysis	54
3.3	Boulder Atomic Clock Optical Network	58
4	Spin-squeezed clock : background and considerations	62
4.1	Basic theory and definitions	62
4.2	Historical background	64
4.3	How do we design a cQED system?	65
4.4	How should we probe the atom-cavity system?	68
4.4.1	Atomic motion and probing time scales	72
5	Sr_3 : a new spin-squeezed clock apparatus	77
5.1	UHV System	78
5.1.1	Design	78
5.1.2	Assembly and bake-out procedure	79
5.2	MOT bucket design	81
5.3	Accuracy considerations	82
5.3.1	Faraday shielding - DC Stark	82
5.3.2	In-vacuum calibrated thermal sensors	84
5.4	Laser systems and cooling	84

5.4.1	461 nm laser system and blue MOT	84
5.4.2	Repump lasers	86
5.4.3	689 nm lasers	86
5.4.4	Optical lattice and transport	87
5.5	Clock laser	91
5.5.1	Fiber noise cancellation	91
5.6	Atomic state preparation	92
5.7	Cavity-stabilization 813nm laser	95
5.8	cQED cavity	96
5.8.1	Cavity design and construction	96
5.8.2	Cavity stabilization	102
5.9	Homodyne interferometer	108
5.9.1	Path length stabilization	110
6	Generating and measuring entanglement on the Sr clock transition	113
6.1	Clock spectroscopy	113
6.2	Entering the collective strong-coupling regime	119
6.3	QPN and the effective atom-cavity coupling	122
6.3.1	QPN in slightly detuned regime	122
6.3.2	Measuring QPN	123
6.3.3	Noise in QPN measurement	125
6.4	Conditional spin squeezing	127
6.4.1	Spectroscopic enhancement	127
6.4.2	Calibrating incident photon number	128
6.4.3	Spin noise reduction	130
6.4.4	Ramsey contrast	132
6.4.5	Measured Wineland parameter	133

7	Squeezed clock comparisons	138
7.1	QPN and SQL	139
7.2	Self-synchronous clock comparison with cQED	140
7.3	Optimal estimators	145
7.4	Vertical extent of cloud	146
7.5	Common-mode suppression of laser noise	149
7.6	Quantum efficiency	151
7.7	Spin state tomography	152
7.8	Independence of the two ensembles	152
8	Conclusion and outlook	156
	Bibliography	158
	Appendix	
A	Vacuum noise limits of cavity frequency determination	170
A.0.1	Heterodyne detection	171
A.0.2	Homodyne detection	172
B	Calculation of effective atom-cavity coupling	173
C	Calculation of free space scattering	175

Tables

Table

5.1	Key parameters of optical cavity	96
7.1	Quantum efficiency contributions	151

Figures

Figure

- 2.1 We illustrate the major contributions to the thermal noise for three optical cavities at various temperatures and with length ranging from 6 cm to 40 cm. We use the coating loss angle ϕ_c from our extraction [54] such that the noise values better match our measured frequency noise. Fractional noise associated with coating and substrate scales inversely with the cavity length. Substrate thermo-elastic and spacer Brownian noise are below 10^{-36} Hz^{-1} for the cryogenic systems. 11
- 2.2 Time domain step response of the Si4 thermal damping system. A step function is applied to the active plate/shield, and we measure the response on the passive shield, and the corresponding frequency change of the cavity measured via a heterodyne beatnote against the MJM cavity. 16
- 2.3 (a) The mechanical layout of the Si4 system. The top view shows how the split plate is mechanically isolated from the main plate, and the only real physical connection is a flexible thermal connection. (b) Measured coherence between the accelerometer on the AVI and the beat note of Si4 against a reference laser. (c) Vertical vibration sensitivity in fractional frequency units per g of acceleration, versus the azimuthal angle between the silicon crystal axis and the three-axis supports. 18

2.4	(a) Vertical vibration sensitivity of Si4 measured versus the relative azimuthal angle between the tripod support and the crystal axis of the silicon crystal. (b) Improved high frequency noise performance of Si4 by sensitivity improvement, and optimization of AVI position.	18
2.5	(a) Measured Si3-Si4 beat note. (b) Velocity waveform measured at the Briggsdale seismic station. We can see the clear arrival of the different types of seismic waves, and the correlations with our optical frequency beat note measurement.	21
2.6	(A) Amplitude response of the photothermal effect in Si4. Dashed lines are phenomenological fits, and the solid black line is a theoretical estimate. (B) Phase response of the measured transfer functions.	24
2.7	Technical noise contributions for the 4 K system, along with the sum of the technical noise (pink line). The sum of the technical noise sources is well below the coating Brownian noise (red dashed line).	28
2.8	Setup for measuring the beatnote between Si4 and Si3. Si4's radiation shields are shown. The Si3 stability is checked using the Sr atomic clock. Comparison against the UTC(NIST) timescale allows for calibration of the Si3 drift when the Sr clock is not running.	29
2.9	(A) Modified Allan deviation for the Si4 cavity. (B) Fractional frequency noise power spectral density (S_y) of the Si4 cavity. In both panels, the shaded green region indicates the spread of the estimated thermal noise floor given two previously published values for the coating loss angle [65, 66].	31
2.10	(A) FFT of the beat measured at 1542 nm (black circles), fit to a Lorentzian lineshape (red line). (B) Histogram of the measured Si4 linewidths for 100 measurements. The Si4 linewidths are obtained by multiplying the beat linewidth by 0.85.	32
2.11	Fractional frequency drift in units of 10^{-18} for Si4 at different operating temperature. The Si3 drift versus intensity is shown as the open red squares, exhibiting a far smaller sensitivity to optical power.	34

2.12	Fractional frequency drift in units of 10^{-18} for Si4 at different operating temperature. The Si3 drift versus intensity is shown as the open red squares, exhibiting a far smaller sensitivity to optical power.	35
2.13	The measured noise PSD at various temperatures (solid curves). The dashed lines are the fits to $S_y(f) = a/f^b$	39
2.14	Frequency-dependent mechanical loss angle for the $\text{SiO}_2/\text{Ta}_2\text{O}_5$ coatings at different operating temperature. The shaded bands indicate the 1σ uncertainty.	41
2.15	Mechanical loss of the individual coating constituents (SiO_2 , Ta_2O_5), and $\text{SiO}_2/\text{Ta}_2\text{O}_5$ coatings measured by ringdown (orange diamonds [75], blue triangles [76], green squares [77]). Other direct measurements are also included (black triangle [78], open black square [25], and purple diamond [79]). Our results from direct CTN measurements are shown as the red circles and the red arrow.	42
2.16	Si3 zero crossing measurement of the coefficient of thermal expansion. Measured both against Si4 and the hydrogen maser. The fitted zero-crossings are consistent, with $T_0 = 123.345542(2)$ K from Si3-Si4, and $T_0 = 123.34594(7)$ K from the Si3-maser.	45
3.1	(a) Schematic of the cross-correlation method. Each laser is fiber noise cancelled to the frequency comb table, and we measure the beat between Si4 and Si3, and MJM vs Si3. (b) The black data is the cross-correlation spectrum for the Si3 system (21 cm Si). The red line is a laser model that we will use for Dick effect calculations and dark matter limit simulations.	47
3.2	(a) Anti-synchronized clock comparison between Sr1 and Sr2. The solid gray line is the calculated QPN, blue line is calculated Dick effect, orange dashed-dot is the quadrature sum of QPN and the Dick effect laser noise, and the dashed green line is the fitted stability, giving $5.5 \times 10^{-17}/\sqrt{\tau}$. (b) Synchronous comparison, showing a single clock stability of $3.5 \times 10^{-17}/\sqrt{\tau}$	49

- 3.3 (a) Fractional frequency deviation of the Si - Sr frequency comparison. The data has been dedrifted by fitting the data to Eq. 3.13. (b) Power spectral density of the Si-Sr data. The model for the Si cavity is shown as the dashed line. (c) Fractional deviation of the Si - maser frequency difference. (d) PSD of the Si - maser data. The maser model has a white noise term, and a broadened resonance to account for the microwave link noise. The silicon model is the same as for (b). 53
- 3.4 (a) 95% confidence limits placed on the dimensionless couplings to the electromagnetic field d_e , showing the improved limits set by $f_{\text{Sr}}/f_{\text{Si}}$ (red) and $f_{\text{H}}/f_{\text{Si}}$ (blue) ratios in the mass range above 1×10^{-19} eV. The maximum projected sensitivity for a search of the same 11-day duration without data gaps is included (light green), highlighting the potential of our technique. Limits derived from previous spectroscopic searches (black lines), with the network [95], Dy-Dy comparison [96], and Rb/Cs [97]. Limits from equivalence principle tests (purple lines), with MICROSCOPE [98] and UW [99]. (b) Demonstration of a significantly improved limit on the electron mass modulus (d_{m_e}) derived from the $f_{\text{H}}/f_{\text{Si}}$ (blue) ratio. Limits from equivalence principle tests (purple lines), including MICROSCOPE [98] and measurements from UW [99]. Shaded regions in both (a) and (b) are excluded at the 95% confidence level given the observed signal and noise models. 55
- 3.5 Corrections to the d_{m_e} limit from dedrifted and the unknown dark matter phase. The vertical dashed line indicates the lowest Fourier frequency for which a dark matter limit is computed. 57

- 3.6 (a) Schematic for the out-of-loop evaluation of the BRAN fiber link for the 2017/2018 JILA-NIST clock comparison. A stable laser (Si3) is fiber noise cancelled, with the noise detected between the mirrors labelled as “ref #1” and “ref #2”. The out-of-loop beat is measured between the local light and the “remote” light. (b) Overlapping Allan deviation of the inloop error for the fiber noise cancellation (red) and the out-of-loop beat (black points). (c) Histogram of the out of loop frequency measurements, with the residuals to the Gaussian fit shown below. 60
- 3.7 Allan deviation of the three ratios on a given day. Also shown are noise traces for various components of the optical network, including the frequency combs, fiber link, free-space link and a network loopback measurement. We can see good agreement between the free space and fiber measurements of the Yb/Sr ratio. 61
- 4.1 (a) Schematic of the detection scheme, where laser light is split into a probe path and LO path, and these are combined on a beamsplitter. (b) By sweeping the probe laser across resonance we can map out the $\mathcal{I} - \mathcal{Q}$ response of the cavity response. (c) Normalized \mathcal{I} quadrature. (d) Normalized \mathcal{Q} quadrature. 70
- 4.2 (a) The atom-cavity coupling $g^2(x, y)$, which follows the intensity of the cavity mode standing wave is shown as the contour plot, normalized to a peak value of one. The lines indicate the $1/e^2$ radius. Overlaid on top of g^2 is 100 different atomic trajectories simulated for the particle motion in a single pancake of the optical lattice. (b) Corresponding traces of $g^2(t)/\delta_c$ for all of the different particles, showing that most atoms sample many standing waves well within 1 ms. (c) Fractional noise of the cavity shift for a measurement time of 2 ms, versus the dwell time between frequency measurements. We show the magnitude of the oscillations versus atomic temperature (d) Same as (c), with a measurement time of 40 ms. We can see that the longer measurement time offers improved averaging of the radial motion. 74

4.3	Simulations of cavity frequency shift noise arising from thermal atomic motion. With a fixed separation time between measurements of 10 ms, we predict fluctuations in the shifted cavity frequency on the percent level. Both low atomic temperatures and longer measurement time T_m are critical for minimizing this effect.	76
5.1	(a) Angle-view of the MOT bucket. The smaller MOT bucket pieces are welded together on a 10 inch conflat flange that is bolted to the bottom of the chamber. (b) Side cut-view of the MOT bucket. The MOT coils and the cooling fixtures inside the MOT bucket are labelled. In total, there are 5 in-vacuum welds.	81
5.2	COMSOL simulations of the DC electric field at the location of the atoms. We simulate the effect of having 1 nC of charge distributed uniformly on the surface of one of the mirrors. Also shown is the effect of an applied voltage on that same PZT. Even at the highest operating voltage, the overall shift should be on the order of 1×10^{-18}	83
5.3	(a) Top view of the optical cavity. Shown are the two thermal sensors that are installed in the cavity spacer. (b) Calibration of R_{grey} in a dry-block comparison with three reference sensors. (c) Calibration uncertainty for the two sensors as a function of operating temperature.	85
5.4	Lattice phase noise measurements. The free running lattice phase noise is the red trace, and the locked inloop error is the black trace. For comparison, the laser noise models for the two clock lasers MJM and Si3 are shown as the purple and cyan dashed curves. The photodiode noise floor is the gray trace.	89
5.5	The left axis indicates the relative frequency detuning of the optical lattice beams for transport, and the right axis is the corresponding velocity of the lattice. The total transport distance is approximately 42 mm.	90

5.6	Differential fiber noise between the 0 th order and 1 st order beams. The black trace indicates when we place the 1 st order retroreflection near the 0 th order retro, and red is placing the retro as close to the vacuum chamber as possible, just before the viewport. For reference, we plot the Si3 noise model, showing reasonably adequate noise cancellation.	93
5.7	The top panel shows the experimental sequence roughly laid out in time. Below we graphically indicate the location of the MOTs, and the subsequent loading of the lattice and transport into the cavity. We also note the magnitude and direction of the relevant magnetic field for each step.	94
5.8	Measurement of the cavity finesse at 689 nm by monitoring the transmission after abruptly shutting off the light. The black data and fit are before the bake out, giving $\mathcal{F} = 20,000$. The red data and fit are after the chamber bakeout, giving $\mathcal{F} = 16,000$. As a check, we temporarily move the photodetector to before the cavity (gray data), to confirm our light shut-off time is sufficiently short.	97
5.9	Up close photo of the mirror assembly with the PZT stack. You can also see two of the Ceramabond epoxy joints holding the mirror inside the mirror cup. The whole epoxied assembly is bolted to the cavity spacer with four vented screws, of which two are visible here. The other mirror, is identical apart from the overall height of the PZT.	99
5.10	Full view of the cavity, as it would be seen from below after installation. In clear view are the two mirrors. Also visible is the slot we cut in the cavity, to reduce Eddy currents. The cavity is mounted onto an aluminum jig that was used to secure the cavity during installation of the various components.	101

- 5.11 Response of the magnetic field when rapidly switching off the current through the MOT coils. The magnetometer is placed inside a test cavity spacer having roughly the same geometry as our actual spacer. We measure the response without the spacer (orange), with the spacer (blue) and with the spacer but with a slot cut into it (green). We show a dramatic effect of including the slot in the spacer, showing a faster damping time constant of induced Eddy currents in the spacer. 103
- 5.12 The optics layout and stabilization scheme of the cavity QED setup. The 813nm laser is pre-stabilized, and the cavity is locked to that pre-stabilized laser via the cavity PDH lock. The 689nm laser is pre-stabilized to the red cavity system. The 689 ECDL is split into the LO and probe path. The leakage light from the cavity is detected via homodyne detection. 104
- 5.13 Frequency diagram of the cavity modes and of the different laser light tones. The bare cavity frequency ω_c is detuned from the bare atomic resonance ω_a by $\delta_c/2\pi = 1$ MHz. The cavity is shifted by the presence of atoms in the ground state $|\downarrow\rangle$. The relevant frequencies are shown. The carrier is synthesized by frequency shift of AOM-2, and the sideband that probes the cavity is generated by the EOM drive. The LO is correspondingly frequency shifted to be on resonance with the sideband for homodyne detection. 105
- 5.14 Noise measurements of cavity QED stabilization scheme. The blue line is the inloop error of the 813 ECDL lock to the comb. The gray line is the inloop error of the PDH lock of the cavity to the 813 laser. The 689nm pre-stabilized laser then probes the cavity - serving as the out of loop noise measurement. We can see the frequency bands where we are still limited by the inloop lock of the cavity. 106

- 5.15 (a) Measured output of the homodyne voltage (blue) and the accelerometer (orange) placed on the optical table near the homodyne interferometer. Both signals have been bandpass filtered in software to make the correlations clearly visible. (b) The same signals in the frequency domain. After we subtract the accelerations from the homodyne voltage, we can see strong cancellation of the noise peaks at 20 Hz and 30 Hz, without impacting the homodyne noise at other frequencies. We note here that the photodiode voltage is amplified by a voltage factor of 25, on top of the 10^6 gain of the transimpedance amplifier. 111
- 6.1 Sideband scans of the clock transition. Without any ramping of the lattice or additional cooling (green), with sideband cooling (red), and with the lattice ramp along (black). The lattice ramp gives a sample prepared primarily in the ground motional band $\langle n_z \rangle = 0$, and a narrower blue sideband (radial temperature). The carrier signal is heavily saturated in this clock scan. 114
- 6.2 Radial Doppler scans of the clock transition. Red data points are without the lattice ramp, black points are with the lattice ramp. The radial temperature extracted from the Gaussian fits (corresponding dashed lines) are $T=290(10)$ nK for with lattice ramp, and $T = 447(20)$ nK without the lattice ramp. 116
- 6.3 Rabi flops (a) Measured excitation versus the Rabi pulse time in seconds, with a drive Rabi frequency of $\Omega_R/2\pi = 111$ Hz. (b) Same as (a), with the drive Rabi frequency set to $\Omega_R/2\pi = 12$ Hz. 117
- 6.4 Ramsey contrast for increasingly complex pulse sequences. The pulses are labelled with a sub-script, indicating the relative rotation angle. $\pi/2_{90}$ refers to an “align” $\pi/2$ pulse, required for squeezing. The introduction of the π_0 is the π pulse we need to perform the pre-measurement in our squeezing experiments. 118

- 6.5 Center : we show the eigenvalue spectrum of the atom-cavity coupled system versus the atom-cavity detuning $\delta_c/2\pi$. Left : When we tune the cavity onto the bare atomic resonance, we observe the vacuum Rabi splitting, shown in the left plot. This demonstrates entering the collective strong-coupling regime. Right : For the QND measurements, we operate in the slightly detuned regime. 120
- 6.6 (a) Pulse sequence for measuring QPN. Black open pulses correspond to clock rotations, and the red pulses are cavity frequency shift measurements. (b) QPN measurement. The red datapoints have no rotation noise subtracted, and the model including a linear term is fit (red line) giving $g/2\pi = 5230(213)$ Hz. The open blue circles have an estimated clock rotation noise contribution subtracted, and is fit without a linear term giving $g/2\pi = 5127(192)$ Hz. 124
- 6.7 (a) Pulse sequence used for measuring spin noise reduction. After an initial $\pi/2$, we measure $J_{z,p}$, wait a dwell time τ , and then measure $J_{z,f}$. (b) The high-degree of correlations are shown between $J_{z,p}$ and $J_{z,f}$ 131
- 6.8 The left panel shows the pre-measurement (red points) and the final-measurement (open cyan points). The incident photon number per measurement is 4.2×10^4 . The optimal subtraction $J_{z,f} - \beta J_{z,p}$ is shown in the right panel. The dashed lines indicate the calculated QPN, and the solid lines indicate the calculated standard deviation of the measured data. The spin noise reduction here is $R = -6.3(4)$ dB. 132
- 6.9 (a) Pulse sequence used for measuring the initial contrast with no QND probing applied. (b) Pulse sequence for measuring final contrast C_f with the QND probing. (c) Open green squares is the Ramsey fringe with no probing, black points is with 2.3×10^4 photons per measurement, and red triangles is with 4.2×10^4 photons per measurement. We do not observe any significant phase shift due to the probing. 134
- 6.10 Measured fractional contrast C_f/C_i versus the incident photon number per measurement. The solid line indicates the calculated expected loss of contrast from free-space scattering, described in Appendix C. 135

- 6.11 Relative spin noise reduction R (cyan circles), fractional contrast loss C_i/C_f^2 (red circles), and the corresponding Wineland parameter ξ (black circles). At the optimal photon number, we directly measure a Wineland parameter of -1.8(7) dB. Expected R given our estimated quantum efficiency of $Q = 0.28$ (cyan line), expected contrast loss via free-space scattering (red line), and the corresponding expected Wineland parameter (black line, see Methods). 137
- 7.1 (a) Clock pulses are the black pulses, measurements of ensemble A are the red pulses, and the transports are shown as the green and purple pulses. The Bloch spheres depict the spin state distribution at various points during the sequence. (b) Ramsey fringe measured by varying the phase of the final $\pi/2$ pulse. (c) Pre and final measurements of ensemble A and B. (d) The final measurements of ensemble A and B show strong correlations, allowing for the subtraction of the common-mode laser phase noise. 141
- 7.2 (a) Schematic showing the transported atomic cloud for squeezing/readout of each ensemble. The clock laser propagates vertically, addressing both ensembles in a synchronous fashion. (b) Fractional frequency stability given by the total Allan deviation for the CSS - CSS comparison (green open squares) and the SSS - SSS comparison (black circles). The green dashed line is the estimated QPN limit, and the dotted blue line is the SQL. We note that we plot the observed comparison stability, so we have not divided by $\sqrt{2}$ to estimate single clock stability. 143
- 7.3 The optimal estimators calculated for the spin-squeezed clock comparison data. (a) Optimal estimator β_A for the pre-measurement of A. (b) Optimal estimator β_B for the pre-measurement of B. (c) Optimal estimator β_D , which is the differential scale factor between the two ensembles. The shaded region indicates the ± 1 standard deviation for the datapoints beyond experimental shot 150, and the dashed line is the corresponding mean value. 147

7.4	Measured shift of the cavity-like mode versus vertical displacement of the atomic cloud. The green dashed line is the Gaussian fit of the form of Eq. 7.8, giving a fitted $\sigma_z = 130\mu\text{m}$	148
7.5	Measured cavity shift noise with artificial white frequency noise imposed on the clock laser. Each ensemble (red and blue circles) becomes noisier as the magnitude of the added noise is increased, but taking the difference (open green squares) gives a 10 dB common mode noise reduction factor.	150
7.6	(a) Pulse sequence used for performing spin state tomography. (b) Relative spin noise after a rotation angle of Ψ . We infer the spin noise just after the pre-measurement by assuming the noise in the pre and final measurement is uncorrelated. The black dashed line is a unitary spin squeezed state, and the dashed-dot line indicates our expected anti-squeezing for a quantum efficiency of 0.28.	153
7.7	Independence of atomic ensembles (a) Measured correlation coefficient between $J_{z,A}$ and $J_{z,B}$ versus the separation between the ensembles. (black circles). The blue line is a Monte Carlo simulation. (b) Corresponding change of the QPN due to the finite overlap of the ensembles, with numerical Monte carlo simulation (blue) and analytical calculation (orange). At our operating ensemble separation, the change to QPN is 0.05 dB.	155

Chapter 1

Introduction

One common thread in fundamental physics is the continual “zooming in” on the universe. It also allows one to periodically “zoom out” and see the world as a whole. It is this simple pursuit that has led us through developments of all branches of physics. The development and refinement of the tools used for investigating the universe aids this journey in a critical way. We will highlight a few aspects of this journey that leads to the current world of optical atomic clocks.

First, we need atoms. In 1909, Rutherford instructed his colleagues Hans Geiger and Ernest Marsden to fire alpha particles (later known to be helium nuclei) at very thin sheet of metal foils like gold and silver, and non-intuitively, look for diffuse reflections of the alpha particles at angles greater than 90° [1]. To their astonishment, sitting in the dark and counting scintillations by eye, they found that these alpha particles were indeed being scattered at large angles. Rutherford sat down for a few years to think about the ramifications of this experiment, and eventually developed the picture of the modern atom [2]. This new model of the atom had a central nucleus with net positive charge¹, and electrons orbiting outside that nucleus. We then have Niels Bohr to thank for explaining that these electrons occupied quantized orbits around the atom [3]. It will be the transition of atoms between these quantized energy levels that we will take advantage of in building an atomic clock.

The question is - how can we probe these orbital levels inside the atom? I.I. Rabi showed the way forward by establishing the magnetic resonance method [4]. This involves tuning an oscillating

¹ Neutrons were not known to exist at this time, but Rutherford was careful to simply state that the nucleus must have a net positive charge.

magnetic field until it hits the same frequency as the Larmor precession frequency of the electron inside the atom. When the oscillating field is on resonance, the radiation induces the nuclear spins to transition to another spin state. One could use this technique to peer into the spectrum of different atoms and molecules, and measure the energy spectrum with a new level of precision. Surely this method would lead to some fundamental discoveries about the nature of the atom. Willis Lamb and Robert Retherford continued to zoom in when they employed microwave spectroscopy and found a shift between the $^2S_{1/2}$ and $^2P_{1/2}$ levels in the hydrogen atom [5]. This discovery spurred the exploration and further understanding of quantum electrodynamics (QED). Then, in 1950, Norman Ramsey introduced a new method - later named Ramsey spectroscopy - where one uses two phase-coherent pulses separated in time (or space) to probe the atomic resonance [6]. It was clear that continued work in spectroscopy would yield further paths to explore physics.

I.I. Rabi first suggested using the methods of magnetic resonance techniques to build an atomic clock in a 1945 lecture [7]. The idea is that one uses two energy levels of a particular atom as a frequency reference. The microwave radiation would be used as the pendulum, and one can count the waves of this radiation to provide the ticks of the clock. Sources of microwave generation on their own are imperfect, possessing drift and random frequency deviations. However, by performing spectroscopy on the atoms, we can actively tune the microwave radiation to be on resonance, thus fully stabilizing the light and giving us an atomic clock. The early development of atomic clocks moved very rapidly. The first atomic clock was an ammonia maser, with the construction led by Harold Lyons [8]. Further improvement in maser technology came in 1958 with the development of the hydrogen maser in Ramsey's group (led by Dan Kleppner) [9]. Around the same time, scientists at the National Physical Laboratory developed the first cesium beam clock [10]. At what seems like a shockingly fast pace to us now, in 1967 the cesium hyperfine transition frequency was chosen to redefine the second in the International System of Units (SI).

While the early development of microwave atomic clocks was revolutionary - we should continue to seek more refined tools to investigate nature. It soon became clear that these were waiting for us in the optical domain. The invention of the laser and related optical technologies was another

explosion of ingenuity and fast progress [11]. Venturing into the world of visible wavelength lasers with femto-second timing resolution allows us to sub-divide the second into $\approx 10^6$ more cycles of radiation per second as compared to microwave clocks. This leads to an intrinsic improvement - missing one cycle of the radiation corresponds to a greatly reduced timing error. Shortly after their first realization in the lab, optical clocks began to surpass their microwave counterparts in all aspects of their performance [12, 13]. Progress since then has been steady, with reported accuracies of $\Delta\nu/\nu = 10^{-18}$ [14, 15, 16]. State-of-the-art comparisons between independent clocks are now reaching the 10^{-17} level at an averaging time of 1 second [17, 18]. Further, by comparing sub-ensembles within a single chamber one can leverage the incredible precision and measure the gravitational redshift over a distance of just 1 mm [19].

1.1 Characterizing clocks

It is important to establish metrics for clocks, to evaluate the performance of a given clock by asking the question - how well can a given clock tell time? Viewing this question from another perspective - how sensitive is a given clock to new physics and perturbations? The more sensitive it is - the more we can learn.

There are two broad metrics for atomic clocks - accuracy and stability. Accuracy refers to the fact that the atomic energy levels chosen for the clock transition are not completely immune to systematic effects from the environment. For example, the largest systematic shift in a room-temperature strontium atom clock is the dc Stark shift arising from black body radiation of the surrounding environment. The fact that the thermal environment radiates microwave photons means there is a finite electric field due to that radiation. Due to differential polarizability of the ground and excited state, that electric field leads to a shift in the apparent clock frequency. In other words, if you make a clock in the lab and someone across the world makes a similar clock, but puts the atom in a very cold environment, the two clocks will disagree. It is the job of the clock scientist to carefully calibrate and understand the systematic shift. If we both write down what we believe to be the shift, and our uncertainty in that shift, we can then compare clocks to

see if we still agree.

My thesis will focus more on aspects of *stability*, which is effectively a statement of how noisy a clock (or laser) is. Put in other words, stability refers to how long it takes to reach a given statistical uncertainty of the clock frequency. A clock that has more shot-to-shot noise will have a worse stability - it will take longer to reach a given statistical uncertainty. We wish to be quantitative about stability. Natural experience in measurement science is that when you take datapoints close together in time, they end up more clustered together with some shot-to-shot noise. This shot-to-shot noise could have several sources, such as technical electronic noise, or more fundamental noise. But if we go home for the night, and come back the next day, we might find that the relative frequency had changed. There can easily be environmental conditions that influence noise on longer time scales. We desire a statistical quantity that can capture this behavior - and by far the best for clocks is the *Allan deviation*,

$$\sigma_y(\tau) = \sqrt{\left\langle \frac{(\bar{y}_{i+1} - \bar{y}_i)^2}{2} \right\rangle} \quad (1.1)$$

where $y = \nu/\nu_0$ is the fractional frequency, and the bar indicates that we have averaged over a time interval τ . This statistical operator allows us to separate and understand the influence of noise on different time scales. White frequency noise in the Fourier domain will result in an Allan deviation with a slope of $1/\sqrt{\tau}$, whereas flicker frequency noise ($1/f$) will yield a plateau in the Allan deviation, indicating that more data is not giving a more precise measurement. We will use the metric of the Allan deviation when looking at the stability of ultrastable lasers as well as the optical clock. In fact, the fundamental noise floor of optical reference cavities comes in the form of a “flicker floor” in the Allan deviation. For clocks, the limit is usually arising from quantum projection noise or aliased frequency noise, which we discuss in the following section.

1.2 Limits to clock performance

Now that we are acquainted with how we characterize clock performance, we ask what are the limits to the performance? The finite signal to noise ratio of a given clock measurement will

ultimately set the limitations to the clock performance. There are two general sources of noise in the clock - one coming from the interrogating clock laser and the other coming from the atoms themselves.

The fundamental noise that arises from the atoms of the clock is quantum projection noise (QPN) [20]. This noise is associated with the projective measurement of a system placed into a superposition of two states $|\downarrow\rangle$ and $|\uparrow\rangle$. The standard deviation (denoted in this thesis by Δ) of the atom number in $|\downarrow\rangle$ and $|\uparrow\rangle$ respectively is $\Delta N_{\downarrow} = \Delta N_{\uparrow} = \sqrt{N}/2$. The fluctuations in the two clock states are perfectly anti-correlated, so the standard deviation of the difference is $\Delta(N_{\downarrow} - N_{\uparrow}) = \sqrt{N}$. These population fluctuations are converted to fluctuations of the clock frequency by the slope of the Ramsey fringe. The highest performing microwave clocks - the fountain clock - operate at the limit of QPN [21]. Today, state of the part optical clocks approach the QPN limit [17, 18].

A primary roadblock to observing a clock operating at the noise limit given by QPN is laser noise. Due to finite deadtime in the clock, we do not track the laser phase at all times. This means that the stroboscopic probing of the atoms leads to aliased high frequency noise, known in the clock community as the Dick effect [22, 23]. The particular pulse sequence (Rabi or Ramsey) and the dead time of the clock dictates the magnitude of the laser noise impact on the clock [24].

1.3 Overcoming the limits

Viewed from a practical perspective, both the classical and quantum limitations to the noise performance are critical to understand, optimize and improve and subsequently improve. For the classical laser noise - we take a two-pronged approach. First, we push the boundary of stable lasers down to the 10 mHz laser linewidth by developing state-of-the-art lasers based on cryogenic silicon cavities [25, 26, 27]. This laser stabilization pushed our clock performance up against the QPN limit [17]. With the motivation of removing laser noise from clock comparisons, the Ye group developed the “self-synchronous” measurement technique, by comparing two parts of the same ensemble [28]. This was the technique employed in the measurement of the gravitational redshift with a single atomic cloud [19].

The other fundamental limit - the quantum projection noise - can be a tricky noise source to improve upon. Of course, one way forward is to make the QPN small by operating at large atom numbers. However, there will always be some technical limit to atom number, so it is interesting to wonder if we can obtain clock performance beyond. This can be done by introducing spin-squeezing. We can prepare states with reduced noise along a particular direction of the Bloch vector, at the expense of increased noise along the orthogonal axis (anti-squeezed quadrature). Since Ramsey spectroscopy is primarily sensitive to one quadrature of the spin state, we can align this spin-squeezed state along the phase-sensitive axis. There are a multitude of ways to generate such non-classical spin states [29]. One involves performing quantum non-demolition measurements of the joint atom-cavity system [30, 31, 32, 33, 34]. With correlated quantum noise between the measurements, taking the difference allows one to go beyond the QPN limit. This will be the method we use when generating entanglement and using that in direct optical clock comparisons. There has been much work on integrating spin squeezing in optical clocks [35, 36]. However, until this work, there has not been an observation of an optical clock comparison operating below the QPN limit. By combining the best clock laser technology with clock comparisons with cavity QED squeezing and readout of sub-ensembles, we achieve this long-standing goal.

1.4 Outline of this thesis

This thesis discusses efforts to push the frontier of optical clocks on both the local oscillator and atomic reference side.

Chapter 2 describes the details of the Si4 system, the first cavity to reach the thermal noise floor below an operating temperature of 124 K, achieving this at both 4 K and 16 K. We present a thorough noise budget of this system and stability measurements. Using the measured thermal noise floor, we infer the loss angle of the high reflectivity dielectric coatings, providing valuable information for both reference cavity and gravitational wave detector research.

In Chapter 3, we describe experiments involving classical atomic clocks that push the frontier of precision measurement science. These include demonstrating clock comparisons approaching the

QPN limit, placing limits on dilaton dark matter, and optical frequency ratio measurements at the 10^{-18} level for the first time.

In Chapter 4, we discuss various aspects of a spin-squeezed clock, including design choices, probing atom-cavity system, and the influence of atomic motion on considerations of this probing.

In Chapter 5 we describe the construction of “Sr3”, the spin-squeezed clock apparatus.

Chapter 6 focuses on the first result coming out of Sr3, where we generate entanglement directly on the Sr clock transition. We measure a Wineland parameter of -1.8 dB.

In Chapter 7, we apply spin-squeezing to a direct optical clock comparison. We directly observe an enhancement in the clock performance of 2.0 dB beyond the QPN limit. This constitutes the first direct observation of a spin-squeezed neutral atom optical clock operating below QPN.

Chapter 2

Ultrastable lasers based on optical cavities

2.1 Motivation and outline

Ultrastable lasers are at the core of most of the world's best precision measurements, including optical atomic clocks [37], tests of relativity [38, 39], and gravitational wave detectors [40]. Improved optical coherence will open the door for more precise optical clocks [41, 42]. These lasers will further studies in fundamental physics in several aspects, including the search for dark matter [43], atom-based gravitational wave detectors [44] and many-body physics [45]. Furthermore, optical clocks will play a defining role in the next generation of optical timescales [46, 47, 48]. In our lab, over the past few years we have demonstrated numerous specific applications of ultrastable lasers. For example, one involved comparing the silicon optical cavity against the Sr clock transition, setting a bound on various coupling of dark matter to the Standard Model [49]. This experiment leveraged the fact that this cavity-atom comparison is linearly sensitive to the fine-structure constant. A second example was the first realization of an all-optical timescale [50]. In that work, the flywheel for the timescale was the optical local oscillator instead of a microwave source. This all optical timescale out-performed timekeeping that relies on microwave technology, showing the clear advantage of optical clock technology.

All of these applications greatly benefit from improved short, mid, and long-term laser frequency stability. Looking forward to our spin-squeezed optical clock discussed in the later half of this thesis, we show that possessing a low noise optical oscillator is a crucial ingredient to pushing the quantum frontier of clocks. Without a pristine optical oscillator, it would become difficult or

impossible to debug various sources of noise, and achieve optical coherence at state-of-the-art levels.

We are thus motivated to explore the fundamental limits of ultrastable lasers, as progress in their stability directly advances the performance of optical atomic clock technology. Other approaches for generating ultrastable laser light are being pursued with promising results [51, 52]. However, to this date, stabilization to optical reference cavities results in the lowest noise laser light. In particular, the highest performing systems are based on cryogenic silicon cavities, taking advantage of the reduced thermal noise floor at lower temperature, as we discuss in Section 2.2.1.

We begin with a short introduction to ultrastable lasers based on optical cavities, and the fundamental limit of thermal noise of these systems. We then describe in more detail a optical cavity referred to as "Si4", which was a 6 cm cavity operating near 4 K cooled by a closed-cycle cryocooler¹. The noise budget, stability measurements, drift characterization, and inferring loss angle will all be discussed. At the end we will also describe some of the details of the Si3 cavity, which has been the work-horse optical cavity for Sr3 and all of the other Sr experiments at JILA, including Sr1, Sr2 in Jun's lab, and the Sr experiments in labs of Adam Kaufman and James Thompson.

2.2 Lasers stabilized to optical cavities

All lasers start as free running, with their frequency dictated by the lasing conditions inside the laser cavity. These free running lasers will have both short-term and long-term fluctuations, due to a myriad of environmental perturbations through temperature, pressure, pump characteristics, etc. Following the footsteps of our wiser predecessors, we know that one can make the frequency of a laser more stable by locking to an optical cavity. From the outside a cavity looks like a simple object - two mirrors facing one another, both fixed to a common cavity spacer. If the light enters the cavity exactly on resonance, it will experience a phase shift of modulo 2π , and thus exits the cavity with no relative phase shift. The basic idea of using the optical cavity as a frequency discriminator

¹ We use the past tense because since this work, Si4 has been upgraded to Si6 lead by Dhruv Kedar. This new Si6 system implements crystalline optical coatings, with the motivation of further reducing the fundamental thermal noise floor from the coatings.

is that when the laser is detuned slightly from the cavity resonance, it will pick up a phase shift. We can imagine that the magnitude of this phase shift for a given frequency detuning will depend on the quality factor of the cavity resonance. Cavities with a longer length or a higher finesse have narrower cavity resonance linewidths. After undergoing many bounces in the cavity, some of the light will leak out of the input mirror - available to detect and measure. The leakage light coming out of the cavity can be described in $\mathcal{I} - \mathcal{Q}$ space, which is just fancy language for saying the leakage light has a real (\mathcal{I}) and imaginary (\mathcal{Q}) component. For example, if we can probe an optical cavity and sense the corresponding \mathcal{Q} -quadrature, or the dispersive response, then we can use this as a feedback signal to the laser, to keep it on resonance.

This is the underlying idea of the Pound-Drever-Hall technique [53]. In the PDH scheme, one applies sidebands to the laser well outside the linewidth of the optical resonator. When the laser is incident on the cavity, the sidebands will simply bounce off of the front mirror, and they simply reflect the present phase of the laser. Meanwhile, the carrier enters the cavity, and will “buildup” in the cavity due to the high finesse resulting in many bounces per photon. This light experiences some phase shift due to detuning between the laser and the cavity, and will then leak back out of the cavity carrying that information with it. At low Fourier frequency below the cavity pole of the cavity, the PDH error signal is a frequency discriminator. At Fourier frequencies above the cavity pole, the frequency detuning is integrated and becomes phase information. Sophisticated servo loops will take this cavity pole in the transfer function into account, although we have found this is not always necessary. With sufficient bandwidth, we can lock the laser to the cavity resonance down to the mHz level.

2.2.1 Fundamental limit - thermal noise

State-of-the-art optical cavities are limited in their length stability by coupling to the thermal environment. The presence of mechanical loss in the cavity constituents (spacer, substrates, and mirror coatings) leads to mechanical displacements which directly impact the frequency stability of the laser locked to the cavity [55]. A great overview of the different sources of thermal noise for

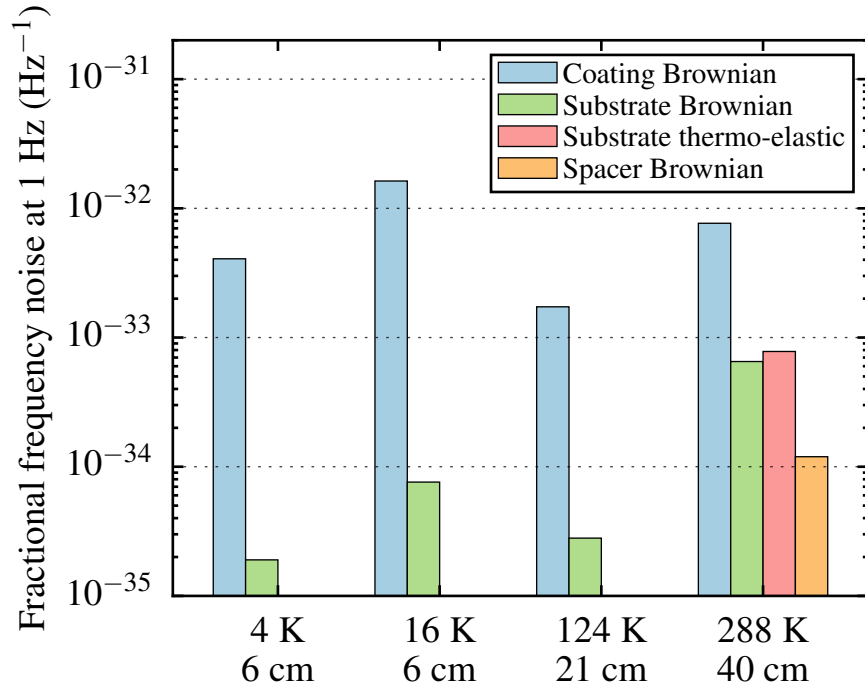


Figure 2.1: We illustrate the major contributions to the thermal noise for three optical cavities at various temperatures and with length ranging from 6 cm to 40 cm. We use the coating loss angle ϕ_c from our extraction [54] such that the noise values better match our measured frequency noise. Fractional noise associated with coating and substrate scales inversely with the cavity length. Substrate thermo-elastic and spacer Brownian noise are below 10^{-36} Hz^{-1} for the cryogenic systems.

rigid cavities can be found in both Chapter 3 of Michael J. Martin’s thesis and the book chapter within the thermal noise book referenced here [24, 56]. We will briefly overview a few sources of thermal noise.

When Einstein first considered Brownian motion, he thought about the random motion of particles suspended in a liquid. Einstein’s work is considered to be an antecedent to the *fluctuation-dissipation theorem*. The fluctuation-dissipation theorem states that the thermally-driven fluctuations of a mechanical system can be directly tied to sources of friction or damping. The components of the optical cavity are bulk mechanical objects, and we can think of the system as responding according to a modified Hooke’s law [57],

$$F(f) = -(1 + i\phi) kX(f) \tag{2.1}$$

where $F(f)$ is an applied force in the Fourier domain, k is the spring-constant, $X(f)$ is the displacement in Fourier domain, and ϕ is the loss angle. We refer to ϕ as the *loss angle*, because the fractional energy loss per cycle is $2\pi\phi$.

The task of calculating the Brownian motion for a macroscopic optical interferometer was made much easier in 1997 when Yuri Levin published a direct approach [58]. Before Levin, the method for calculating the thermal noise involved a computationally expensive normal-mode analysis. The conceptual idea of Levin’s approach is to apply an oscillating pressure (or force) to the observed surface, and the displacement of the surface responds according to the fluctuation-dissipation theorem. One very important takeaway from Levin’s approach is the role of a lossy mirror coating. Levin’s analysis results in a thermal noise contribution due to localized dissipation on the surface (from the mirror coating) that scales as $1/r_0^2$, where r_0 is the beam radius on the mirror. This means that that the loss angle of the coating can be critical for small beam sizes (like in our optical cavities). Since the previously used normal-mode analysis assumes homogenous mechanical dissipation and thus ignored the possibility of localized dissipation from the coating, Levin’s approach was pivotal for the development and understanding of precision optical interferometers.

Besides Brownian motion of all of the cavity constituents, there are several other types of thermal noise. One that is important for ULE glass cavities that employ fused silica substrates is substrate thermo-elastic noise. Thermo-elastic thermal noise refers to the fact that there are fundamental temperature fluctuations, and through the finite coefficient of thermal expansion, will lead to fundamental length fluctuations. A similar source of thermal noise is coating thermo-optic noise, where the length fluctuations arise through the refractive index of the coating. These sources of noise have much longer correlation length, making it fundamentally different from Brownian thermal noise.

Calculated thermal noise contributions for three different cavity systems in our lab are shown in Fig. 2.1. We note that for the cryogenic silicon systems, by far the dominate source is coating Brownian noise. Due to the small mechanical loss of crystalline silicon, both the substrate and cavity spacer Brownian terms are negligible. For the room temperature ULE cavity with fused silica substrates, you can see the significance of the other sources of thermal noise.

2.3 Technical noise budget of Si4

To build a cavity-stabilized laser that is limited by the thermal noise, one has to characterize and control all of technical sources of noise. In this section, we outline various sources of technical noise for the Si4 cavity system.

2.3.1 Temperature

Temperature fluctuations couple to the cavity length via the coefficient of thermal expansion (CTE). The expansion of the cavity and the corresponding change in fractional frequency can be simply expressed in a Taylor expansion

$$\frac{\Delta L}{L} = \frac{\Delta \nu}{\nu} = \alpha(T_0)\Delta T + \frac{1}{2}\alpha'(T_0)\Delta T^2 + \mathcal{O}(\Delta T^3), \quad (2.2)$$

where ΔT is the temperature difference with respect to T_0 , $\alpha(T_0)$ is the coefficient of thermal expansion evaluated at T_0 , and $\alpha'(T_0)$ is the slope. This equation is what motivates some of the

more technical design decisions when considering what material to use and what temperature to operate at. For silicon, there are two zero-crossing of $\alpha(T)$, one near 123 K and another near 16 K. As the temperature approaches zero, the CTE tends to zero, and so another strategy is to operate at the lowest temperature possible. However, cooling down to such low temperature can introduce detrimental vibrations induced by the cryogenics system used. This explains the natural decision to explore the 124K operating point first, where an open-cycle cryogenics system using boiled liquid nitrogen for cooling. Wei Zhang, the post-doc who was leading the initial Si4 effort, always said that he would feel the cryocooler moving with his fingers and remark “that’s no good home for a cavity”.

The Si4 system mainly operated near 4 K, and for a short time, near 16 K. The system is cooled by a commercial cryostat system from Montana Instruments. At the time, this system was custom built, designed in collaboration between Montana Instruments and the stable lasers team, which at the time was Wei and Jun. The cavity operates at 4 K using a two-stage Gifford-McMahon closed-cycle cryocooler [59, 60]. The temperature fluctuations of the cryocooler are typically 20 mK at an averaging time of 6 s. Given a CTE of $\approx 2 \times 10^{-11}/\text{K}$ at 4 K, the corresponding frequency instability is on the 10^{-13} level, which is several orders of magnitude above the thermal noise floor. To suppress temperature fluctuations from the cryocooler, we use a custom, multi-stage thermal damping system [26]. The shields are designed to maximize their heat capacity, while the connections between the shields have a low thermal conductance.

The first shield is coupled to the 30 K stage of the cryocooler, and initially had no active temperature control or monitoring. The cooling power for the next shield, called the active shield, comes from the second stage of the cryocooler (2.2 K) through a mechanically flexible connection called the cold finger. Between the cold finger and the bottom of the active shield, a cylindrical plate made of holmium copper (HoCu_2), which has a specific heat 2500 times larger than copper at 4 K, is inserted to passively suppress temperature fluctuations from the cold finger. The temperature of the active shield is controlled using a resistive heater. For additional thermal damping, we insert a passive shield between the active shield and the cavity, which is mounted on a cylindrical ring

made of G10 fiber glass (not shown). The low thermal conductivity of G10 and the ring's small cross-sectional area (96 mm^2) minimize the thermal conductance between the two shields. The bottom of the passive shield is constructed out of HoCu_2 . A cup-shaped structure made of G10 is attached to the base of the passive shield, upon which the cavity is supported. The support structure has three nested layers to increase its effective thermal path, and has three 1 mm-long G10 rods on top to support the cavity. The contact area of each rod is 0.75 mm^2 to minimize the thermal conductivity.

In order to evaluate the impact of temperature fluctuations on the cavity we directly measured the thermal response transfer function (Figure 2.2). We apply a step of the active shield temperature of 0.5 K, and observe the low-passed response of the passive shield by plotting the temperature (dashed red). The optical cavity response is measured by a heterodyne beatnote against the MJM cavity (blue). The thermal model is equivalent to an electrical second-order resistor-capacitor low-pass filter. The time constant is $\approx 1400\text{ s}$ between the active and passive shields, and is $\approx 600\text{ s}$ between the passive shield and the cavity.

This was the state of the Si4 cryogenic system as of the end of 2017, after our first result. In order to push this system to the fundamental thermal noise limit, we sought to improve various aspects of the system. One important improvement was made regarding the first radiation shield. Initially, the first radiation shield was a relatively thin aluminum radiation shield coupled to the 30 K stage of the cryocooler. On top of this, it was not actively controlled, only the 30 K stage of the cryocooler was controlled, which is far from the shield itself. We discovered a significant coupling between the temperature of this supposed 30 K shield and the frequency of the cavity. After discovering this coupling between the 30K stage and the cavity, we wanted to redesign and upgrade the radiation shield to a copper radiation shield with more thermal mass. Further, initially we had no thermal sensors on the Al radiation shield. After opening up the chamber to begin our upgrades, we attached a thermal sensor to this Al radiation shield and cooled back down (without the cavity), and to our surprise we found that the shield was closer to 80-90K, with a massive 10 K gradient along the vertical extent of the shield; this made us happy that we already planned to

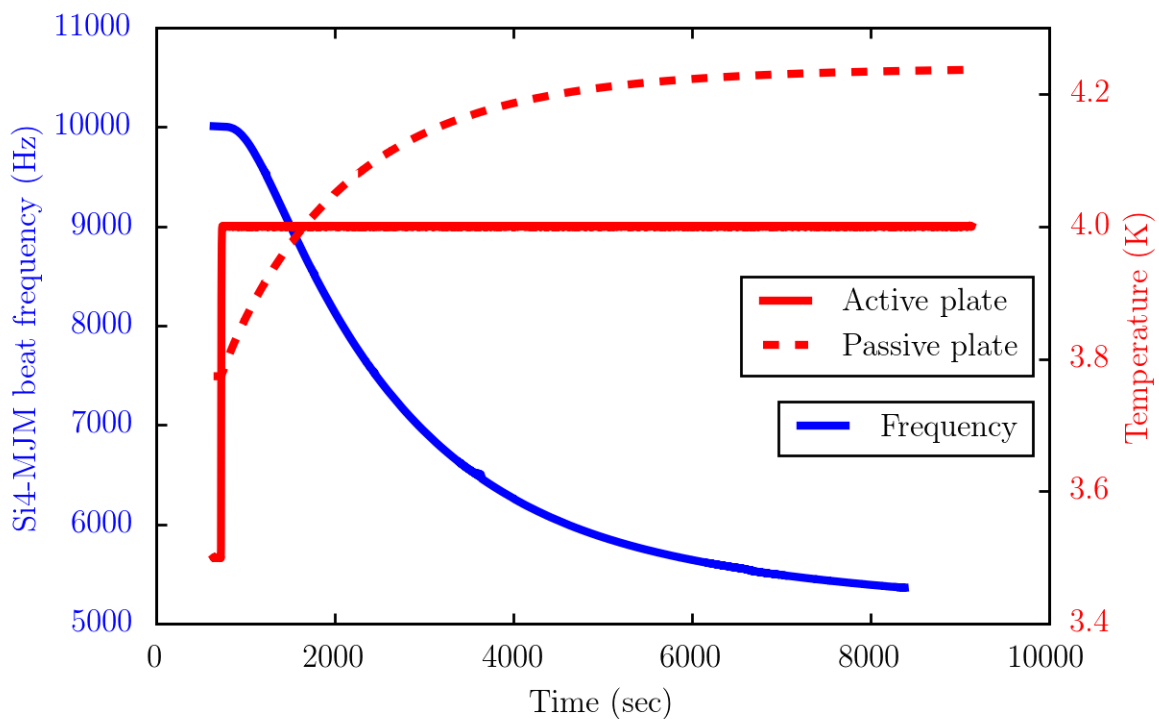


Figure 2.2: Time domain step response of the Si4 thermal damping system. A step function is applied to the active plate/shield, and we measure the response on the passive shield, and the corresponding frequency change of the cavity measured via a heterodyne beatnote against the MJM cavity.

replace this shield! In hindsight - aluminum is not a good choice for the low temperature cryogenic shields as compared to copper given the thermal conductivity, and the thin 1 mm thick walls both gave a small thermal mass, and a small cross-section for propagation of heat to the top of the shield. After replacing this Al shield for the Cu shield - we were able to cool that shield down to 40K with relative ease, and achieved a sub 1K gradient along the entire extent of the shield.

2.3.2 Vibration noise

Vibrations are a significant source of short-term instability for the Si4 system because of the closed-cycle cryostat. We minimize vibrations coming from the cryocooler by carefully designing the mechanical layout of the system [26]. The mechanical design (Fig. 2.3a) isolates the cavity from the vibrations of the cryocooler. The cryocooler sits on a main plate that is firmly fixed to a monolithic aluminum pillar that is bolted to the concrete foundation of the laboratory. There are three mandatory mechanical links between the cryocooler and the chamber: a cold finger to cool the active shield, a thermal link connecting the first stage of the cryocooler to the 30 K shield, and a vacuum bellows to isolate them from the environment. All other mechanical connections are eliminated by placing the vacuum chamber on a split plate, which is mechanically separated from the main plate (see the top view in Fig. 2.3a). The gap between the main plate and the split plate is 1.5 mm, and is exaggerated in the figure. The split plate is supported by three rods that are bolted onto an active vibration isolation (AVI) table. The AVI rests on an optical table. The optical table is supported by four aluminum blocks, roughly 1 foot tall, giving the setup an overall low profile for minimizing tilting motion.

Our initial characterization of the vibration sensitivity of the Si4 system was obtained by using ambient vibrations as the drive. We simultaneously measured the vibration noise and a heterodyne beat between the silicon cavity and a stable laser at 698 nm based on a 40 cm-long ultra-low expansion glass (ULE) cavity (called the MJM cavity throughout this thesis) using a Yb frequency comb. We found that the beat frequency against the MJM cavity was highly correlated with the vibrations measured by an accelerometer placed on the AVI table at Fourier frequencies

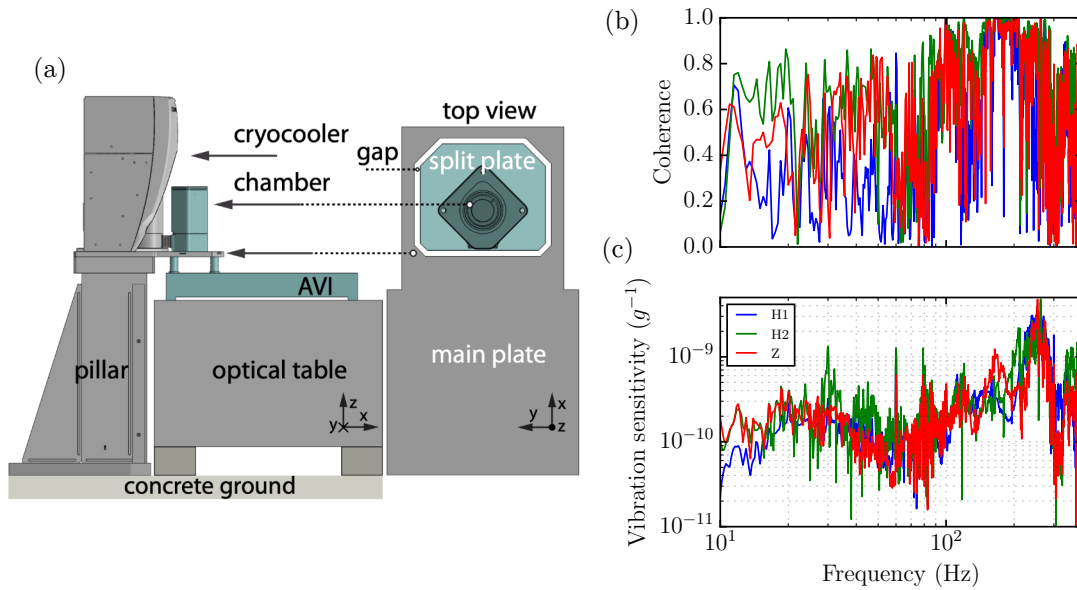


Figure 2.3: (a) The mechanical layout of the Si4 system. The top view shows how the split plate is mechanically isolated from the main plate, and the only real physical connection is a flexible thermal connection. (b) Measured coherence between the accelerometer on the AVI and the beat note of Si4 against a reference laser. (c) Vertical vibration sensitivity in fractional frequency units per g of acceleration, versus the azimuthal angle between the silicon crystal axis and the three-axis supports.

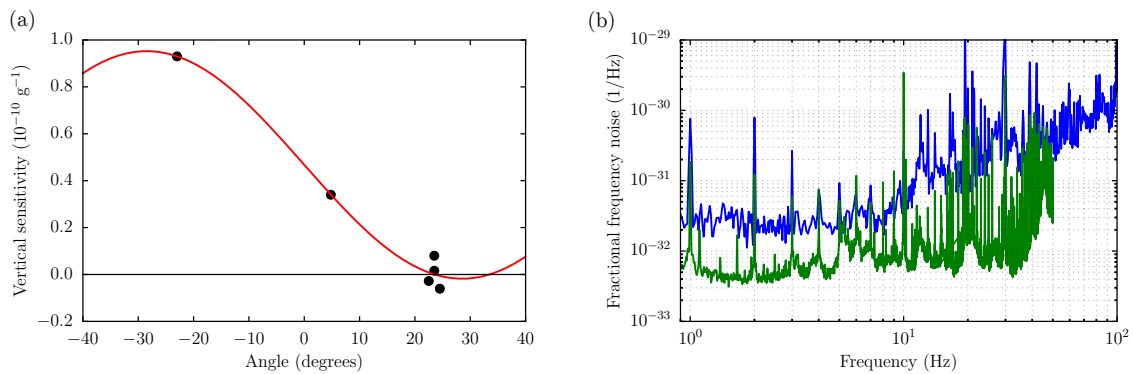


Figure 2.4: (a) Vertical vibration sensitivity of Si4 measured versus the relative azimuthal angle between the tripod support and the crystal axis of the silicon crystal. (b) Improved high frequency noise performance of Si4 by sensitivity improvement, and optimization of AVI position.

above 10 Hz. To quantify the correlation between vibrations measured on the AVI table and the frequency noise of the beat note we introduce the *coherence* : $C_{a_i\nu}(f) = |S_{a_i\nu}(f)|^2/[S_{\nu\nu}(f)S_{a_ia_i}(f)]$. Here, $S_{a_i\nu}(f)$ is the cross-spectral density between frequency fluctuations of the beat signal and acceleration noise in the i th direction ($i = x, y, z$), and $S_{\nu\nu}(f)$ and $S_{a_ia_i}(f)$ are the autospectral densities for frequency noise and accelerations, respectively. Three accelerometers are mounted to the AVI table, one for each direction labelled as H_1 , H_2 and z . The coherence between each accelerometer and the phase noise of the beatnote are shown in Figure 2.3b. A nonzero coherence value at a given f indicates that the laser noise and vibrations are correlated at that spectral frequency. This is an important prerequisite for measuring transfer functions from accelerations to frequency noise, $H_{a_i}(f) = S_{a_i\nu}(f)/S_{a_ia_i}(f)$, since phase estimates in the cross spectrum are only trustworthy where significant frequency-domain correlation exists. From 10 Hz to 400 Hz, we measure appreciable coherence between frequency noise and vibrations along all three directions. Within this frequency range, we compute transfer functions from each accelerometer on the AVI table to the beat frequency spectrum by averaging $H_{a_i}(f)$ over an hour-long dataset. The measured transfer functions along each direction are shown in Figure 2.3c. Si4 reached the level of $10^{-10} g^{-1}$ for 10 - 100 Hz, with some resonant features at higher frequencies. We have since made significant improvement in lowering the vibration sensitivity by paying attention to the mechanical support structure and mounting in the design of Si6 (an effort led by Dhruv Kedar).

After our initial result of [26], we sought to improve the short-term stability. Due to the anisotropic nature of the silicon crystal [61], the vertical vibration sensitivity is sinusoidally modulated, with a period of 120 degrees. When the cavity was initially installed, the vertical sensitivity versus azimuthal mounting angle had not yet been experimentally checked, although Wei had placed the cavity to the simulated zero-crossing of the sensitivity. We had reason to suspect how close we really were to the zero-crossing, since the simulated zero-crossing can be modified by imperfections in the mechanical model of the cavity and in the finite-element-modeling. Since we were planning to open the vacuum chamber in order to make the upgrades to the thermal system as discussed in Ch. 2.3.1, we planned to measure this. We directly measured the vertical vibration sensitivity versus the

relative angle between the mounting points and the crystal axis orientation of the silicon. We were able to reduce the vertical vibration sensitivity to $(5 \pm 2) \times 10^{-12}/g$ at a driving frequency of 9.5 Hz. The horizontal vibration sensitivity in each direction was measured to be $(2 \pm 1) \times 10^{-10}/g$. A reduction of vibrations at the cavity was obtained by fine tuning the relative position of the vacuum chamber and the cryostat. The combined improvements in sensitivity and noise provide a tenfold reduction in the frequency noise power spectral density (PSD) for Fourier frequencies of 10-50 Hz compared to previous work [26]. The horizontal sensitivity depends critically on the transverse positions of the cavity mirrors with respect to the optical axis of the cavity, and thus a more precise centering of the mirrors with respect to the symmetry axis of the cavity will provide a reduced horizontal sensitivity.

One other aspect related to vibration sensitivity is the impact of earthquakes. During a JILA/NIST clock comparison data point, suddenly all of the clocks came unlock at the same time. By chance, I went to the NYtimes website and saw that an earthquake had just occurred in Mexico (February 16, 2018). We looked up the local seismometer station, and found that one can download the ground velocity data, and this is plotted in Figure 2.5a. Looking at the beatnote between the Si3 and Si4 cavity systems (Figure 2.5), we see clear correlation between the two signals. One can see the arrival of the different components of the seismic waves produced by the earthquake. We have a remarkably sensitive optical interferometer, that we have attempted to make as insensitive to vibrations as possible. The fact that we can still see such a clear signal is indicative of the relatively large ground motion coming from earthquakes, even far away. It also reminds you that low frequency vibrations are not cancelled by our AVI table. This means that using a lower frequency sensor of motion, like a tiltmeter, could give enhanced control over the low frequency vibration environment.

2.3.3 Residual amplitude modulation

One of the most important technical noise sources for Si4 is residual amplitude modulation (RAM). RAM is a systematic effect in frequency modulation based stabilization schemes, which




Figure 2.5: (a) Measured Si3-Si4 beat note. (b) Velocity waveform measured at the Briggsdale seismic station. We can see the clear arrival of the different types of seismic waves, and the correlations with our optical frequency beat note measurement.

arises when the two modulation sidebands are unequal in magnitude, not exactly 90° out of phase, or a mixture of the two effects. From an operational point of view, RAM that is completely stationary is not a real problem, but in reality, the RAM always has some time-dependent component. This time-dependent RAM manifests itself as time-varying baseline of the PDH error signal. Since the PDH servo loop can not distinguish this change in the error signal from real frequency deviations between the laser and the cavity, this RAM will be written as frequency drift of the locked laser. There can be various sources of RAM, including time-varying RF pickup, changes in the EOM temperature, and optical etalons. The most nefarious sources of RAM is time-varying optical interference, or etalons. If you have some sort of standing wave in your optical system, that means the carrier and sideband can experience differential phase shifts or amplitude attenuations. Then, as the lab temperature or pressure slowly varies, then this relative phase shift or amplitude attenuation will appear as a slowly changing DC offset of the error signal. We point the interested reader to [62] for an in-depth analysis of the role of etalons in different sections of the optics layout.

We employ active control of the RAM, following the strategy described in [63], with a few minor differences. In our case, the RAM is detected by picking off the beam with a 90/10 non-polarizing beamsplitter before the PDH PBS. We demodulate the RAM, with the phase set such that we are detecting the \mathcal{I} quadrature. The relative phase for the \mathcal{I} quadrature is typically set by maximizing the RAM error signal. This demodulated signal is then sent through a servo loop filter, and we apply the feedback as a DC voltage onto the EOM. The \mathcal{I} quadrature dependence on voltage is a sinusoidal function of the EOM temperature, and so to optimize the gain and dynamic range, we tune the temperature so that this slope is maximized. For our setup, the \mathcal{Q} quadrature was always much too small to play a significant role.

The way to diagnose the out-of-loop RAM is by tuning the laser off of the cavity resonance, and observe the demodulated PDH error signal. If this baseline is fluctuating, then you can be sure that this same fluctuation is impacting the PDH error signal when you tune on resonance. By tilting optics, controlling temperature and pressure of the overall optics setup, and adding optical isolation in the correct spots, one can reduce this RAM to the parts-per-million level or better.

The measured out-of-loop RAM for Si4 is shown as the red trace in Figure 2.7. We can see a broad peak in the 10-40 mHz region, which we ascribe to fluctuating etalons in the system.

Of course, this off-resonance test is not perfect, and things could degrade further when you are on resonance, since you are allowing the beam to enter the cavity. For example, the beam could strike a piece of dust on a window in transmission of the cavity. This scattered light could make its way back to the PDH error signal. It will appear as a fluctuating phase shift between the carrier and the sidebands, and the PDH servo will apply feedback based on this signal. So while checking the out-of-loop RAM as we have discussed is a first step, one must always check the stability of the stabilized laser by forming a three-cornered comparison with two other lasers.

2.3.4 Photothermal effects

Photothermal effects are those that change the frequency of the cavity depending on the amount of laser light circulating in the cavity. Photothermal effects can arise from different physical mechanisms. At very low frequency, the absorbed power in the mirror coatings propagates and changes the length of the spacer. At higher frequencies, the response tends to be dominated by the coating, since the substrate contribution with silicon is negligible. Unfortunately there are not great references for the different critical parameters. From the perspective of just trying to make the best laser possible - it is sufficient to measure the transfer function, and then stabilize the intensity to a high enough degree. In Figure 2.6a, we show the measured photothermal amplitude in units of Hz/W, where W refers to the amplitude of the modulated power measured in transmission. We measured the amplitude and phase response at different overall transmitted power, and found that the photothermal transfer function strongly depended on this transmitted power.

It is natural to ask what is limiting the transfer function in the different frequency ranges in Figure 2.6. In order to attempt to understand the frequency dependence of the photothermal response, we calculate the expected transfer function based on the model developed in [64]. At low frequency the dominate contribution is the thermal response of the cavity spacer. At Fourier frequencies above 10^{-2} Hz, the photothermal response is dominated by the coating. There are

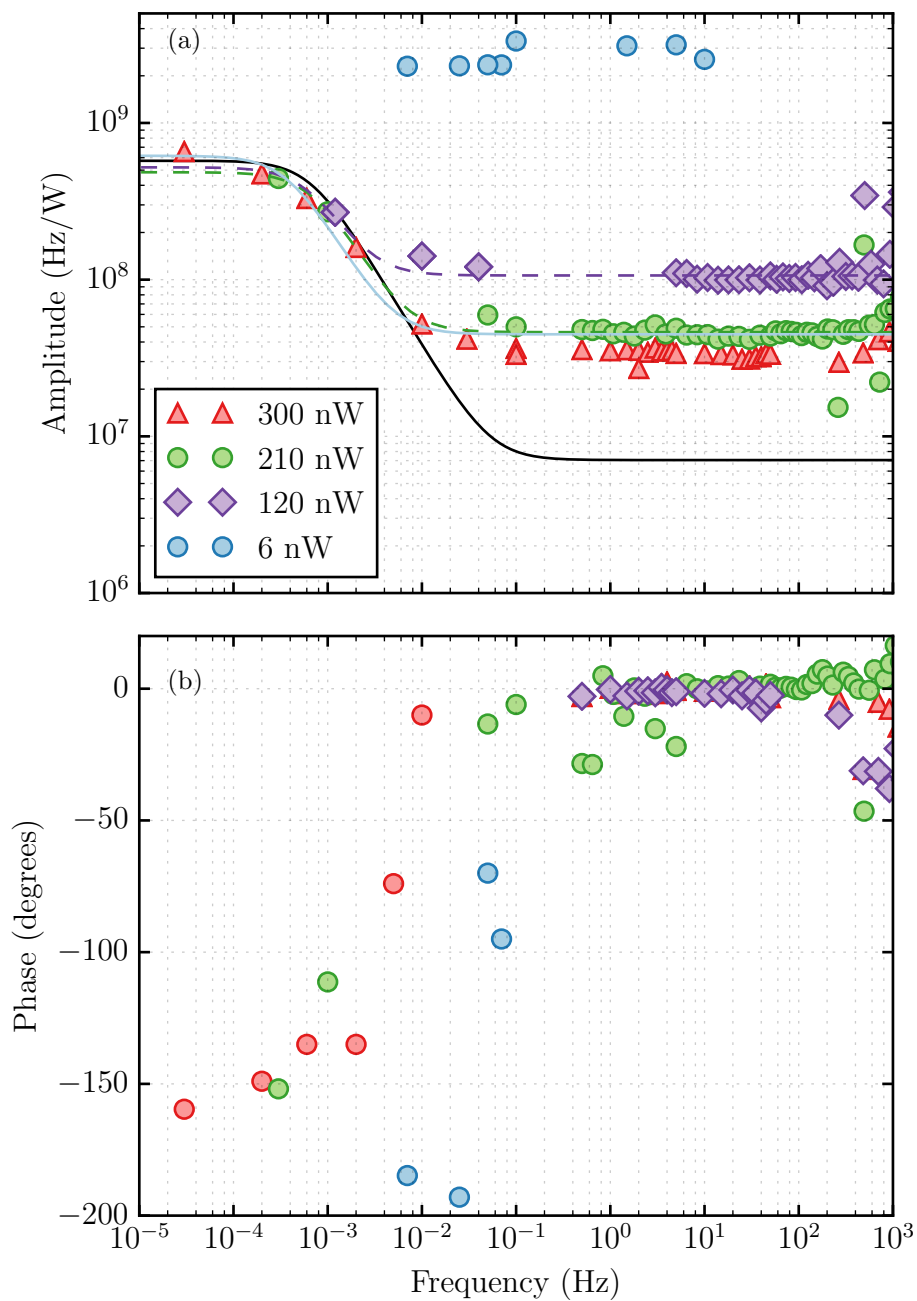


Figure 2.6: (A) Amplitude response of the photothermal effect in Si4. Dashed lines are phenomenological fits, and the solid black line is a theoretical estimate. (B) Phase response of the measured transfer functions.

two contributions to the coating photothermal response - the thermo-elastic and photo-thermo-refractive. The thermo-elastic effect refers to the fact that a temperature change results in a change in the coating thickness through the coefficient of thermal expansion. The photo thermo-refractive effect arises due to the finite dependence of the refractive index on temperature dn/dT . To calculate these effects, we need to plug in various material properties at cryogenic temperatures, which can be difficult to assign. Nonetheless, we use the parameters from [64], and plot the result of the calculation as the black line in Figure 2.6.

We find it difficult to assign the exact physical effect dominating the transfer function at higher frequency. Photo thermo-refractive noise seems unlikely, since dn/dT tends to get small at cryogenic temperatures. It is also interesting to note that the phase of the photothermal response flips by 180° as we go from the low frequency cavity response to the higher frequency coating contribution. Furthermore, the dependence on this photothermal transfer function on the optical power might suggest the presence of another light-matter effect.

From a practical perspective, we minimize the impact of the photothermal effect in Si4 by actively stabilizing the transmitted optical power, down to 40 nW. We can thus reduce the impact of intensity fluctuations to be well below the Brownian thermal noise floor (see Figure 2.7). Further efforts in our group on coating photo-thermal effect would reveal other interesting effects associated with different mirror coatings (in our case, crystalline coatings).

2.3.5 Path length noise

Another critical source of noise is path length fluctuation. When we are considering stabilizing a laser to a cavity, the critical junction is when you split the light in front of the cavity, where some of the light heads toward the cavity and some heads toward either a beat note or toward a particular optical clock experiment. Since any noise written onto the path length after the split point is differential, one must be careful that the phase noise introduced in that path is sufficiently small. For the Si4 cavity, we characterized this noise by performing a loop-back measurement. We split the light in fiber before the output coupler, and send part through 5 m fiber across the

lab, and then loop the fiber back to the beamsplitter. We close the loop on the path length noise of the fiber, and beat the cancelled light against the local light. The relative frequency noise is well below the thermal noise limit of the the cavity. In actual operation, there is additional free space path that is differential. In order to minimize the impact of noise in this free space path, we built a brass enclosure that surrounds the entire optical setup. On top of this, we added an additional layer of isolation around the optics using blackened aluminum foil. Around the optical table we installed heavy-duty vinyl curtains to further suppress both room temperature fluctuations and pressure variations.

2.3.6 Electronics noise

There are various sources of electronics noise in a cavity-stabilized laser. One source of noise is from the photodetector (photodiode + transimpedance amplifier) used for PDH detection. In most optical cavities, this noise source is dealt with by simply turning up the optical power such that the signal to noise of the detected error signal is limited by photon shot noise. For Si4, due to significant photothermal effects, both in noise and drift, we desire a low operating power. To this end, we employ a resonant photodetector that has a noise equivalent power (NEP) of ≈ 60 nW, allowing operation at the 100 nW level without significant contribution from the photodiode. Another source of electronics noise can come from the servo used for PDH locking. We used a standard JILA low delay servo, which is a circuit that has been around JILA for two decades, having undergone several upgrades. For use in a stable laser system, we found that the low frequency noise had to be improved. We achieved this by changing the input op-amp from a very high speed op-amp to a low bandwidth one, with a lower $1/f$ corner. The final PDH servo noise is shown in Figure 2.7 as the dotted curve. While this was low enough for Si4, Eric Oelker designed a further improved servo for Si3 and Si6, to achieve an even lower noise floor.

The full technical noise budget of the 6 cm system is shown in Figure 2.7. We show fits to the measured noise arising from temperature, Pound-Drever-Hall (PDH) servo, and the PDH photodetector (PD). The dominant technical noise source is residual amplitude modulation (RAM),

shown as the red data. The measured out of loop RAM has a prominent peak around 20 mHz, and such a broad peak can easily arise from time-dependent etalon effects. The sum of the technical noise is well below the predicted thermal noise floor (dashed red line).

2.4 Stability measurements

To determine the instability of the Si4 cavity, we measure a beat between Si4 and Si3 (see Figure 2.8). Si4 and Si3 are both fiber-noise-cancelled to the comb table located within the same room as the cavities. Si3 consists of a $1.5 \mu\text{m}$ laser stabilized to a silicon cavity which operates at 124 K with a thermal noise floor of 4×10^{-17} [61, 25]. The short-term instability (averaging times of 0.1 to 10 s) of Si3 is determined by a three-cornered comparison with Si4 and a ULE clock laser at 698 nm. The long-term instability (>10 s) is directly measured by a strontium optical lattice clock. These measurements show that Si3 is at its thermal noise floor for averaging times from 0.1 to 1000 s [17].

Measuring the noise of a beatnote at the 10^{-16} level and better is no easy task. We often make use of frequency counters to measure beatnotes. However, while frequency counters are great devices, they have broad bandwidth, and so they can be very sensitive to other spurious tones in the RF neighborhood. Furthermore, the shortest gate time of the counter is 1 ms, which can lead to significant aliasing of noise above the Nyquist frequency of 500 Hz. Thus, we found that in order to obtain reliable noise measurements with the counter, we needed to heavily band-pass the beatnote signal. We employed an 8-pole active band-pass filter, with a 3 dB bandwidth of 1 kHz, and a center frequency of 10 kHz. At these lower frequencies, the counter can become sensitive to amplitude fluctuations, so our active filter is designed to give a square-wave output.

The modified Allan deviation of this beat, after subtracting the reference laser instability of 4×10^{-17} in quadrature, is displayed in Figure 2.9(A). The modified Allan deviation is calculated from a 24,000 second long measurement record made with a dead-time free lambda-type counter. We compute the instability after removing the linear drift of the beat with a magnitude of $\sim 3 \times 10^{-18}/\text{s}$. The Si4 instability reaches 6.5×10^{-17} for averaging times of $0.8 < \tau < 80$ s, which is

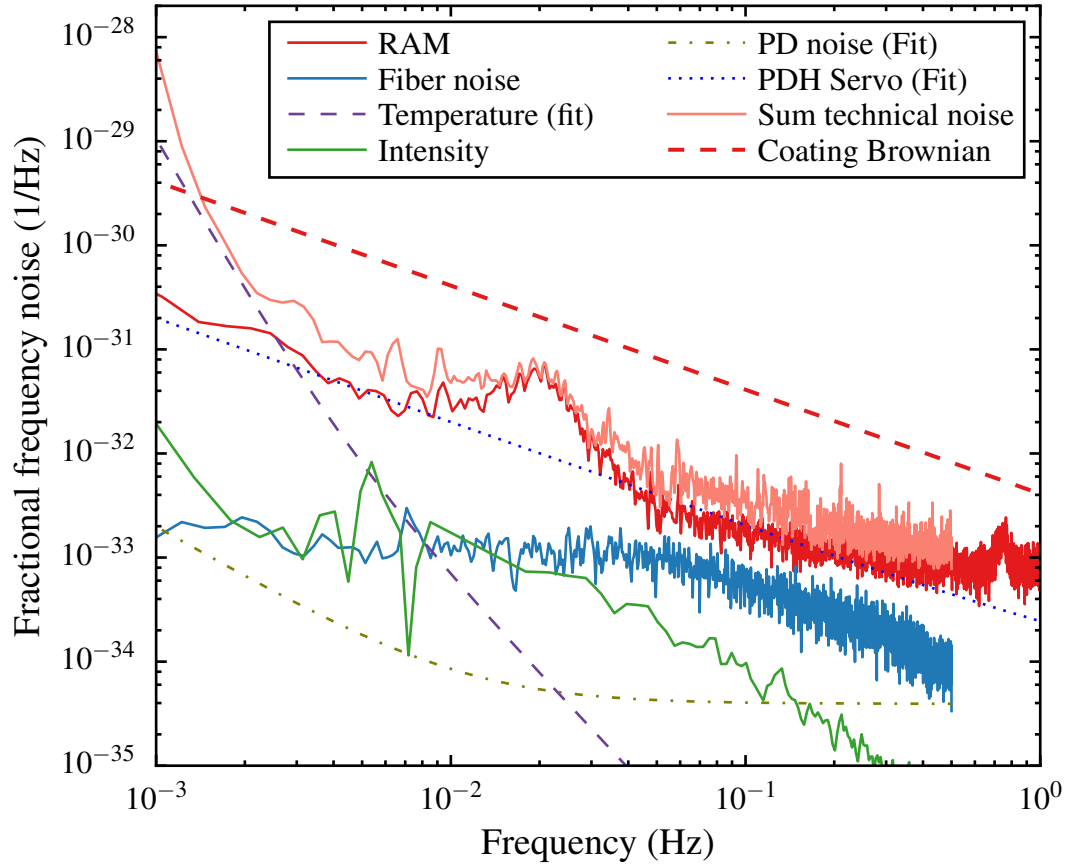


Figure 2.7: Technical noise contributions for the 4 K system, along with the sum of the technical noise (pink line). The sum of the technical noise sources is well below the coating Brownian noise (red dashed line).




Figure 2.8: Setup for measuring the beatnote between Si4 and Si3. Si4's radiation shields are shown. The Si3 stability is checked using the Sr atomic clock. Comparison against the UTC(NIST) timescale allows for calibration of the Si3 drift when the Sr clock is not running.

consistent with the predicted thermal noise floor (green shaded region). At the time of this work, we had limited knowledge of the actual loss angle of the dielectric coatings at cryogenic temperatures. The uncertainty in the thermal noise floor is estimated from the two published values of the loss angle at 4 K. [65, 66].

The corresponding frequency noise PSD for Si4 is shown in Figure 2.9(B). The PSD is calculated from the time-series of the beat obtained by the frequency counter. The thermal noise floor of the Si3 cavity ($S_y = 1.7 \times 10^{-33}/f$) is subtracted from the beat PSD [25]. The Si4 laser is limited by the thermal noise floor for Fourier frequencies over nearly three decades, from 5 mHz to 2 Hz. We fit the measured PSD to a function $S_y = af^{-1}$ and obtain the fit parameter $a = 4.12(5) \times 10^{-33}$. The thin noise spikes at 1 Hz and higher harmonics come from the cryocooler vibrations. The laser deviates from thermal noise at Fourier frequencies below 0.5 mHz, potentially due to etalons or temperature fluctuations.

We experimentally determine the laser linewidth from a Fast Fourier Transform (FFT) of the Si3-Si4 beat. The beat is mixed down to 10 Hz and digitized with an analog-to-digital converter. One example of such a measurement is shown in Figure 2.10(a). We use a measurement time of 128 seconds and employ a Hanning window, corresponding to a Fourier limit of 10.9 mHz. The expected linewidth for $1/f$ frequency noise is given by a statistical distribution [25]. The distribution is multiplied by the ratio $\sigma_{Si4}/(\sqrt{\sigma_{Si3}^2 + \sigma_{Si4}^2}) = 0.85$ in order to estimate the relative contribution of the Si4 laser to the beat linewidth. Here, $\sigma_{Si4(Si3)}$ refers to the thermal noise floor of Si4(Si3). We repeat this measurement 100 times and plot the histogram of results in Fig. 2.10(B). The median laser linewidth for the distribution in Fig. 2.10(b) is 16 mHz, representing by far the lowest observed to date for a laser locked to an optical cavity placed inside a closed-cycle cryocooler.

2.5 Linear frequency drift : power dependence

We measure the drift of the Si4 system by counting the beat Si3-Si4 as shown in Fig. 2.8. This requires careful calibration of the drift of the Si3 system. The Si3 laser is used as the clock laser for a strontium optical lattice clock, giving a direct measurement of the drift [17]. As an

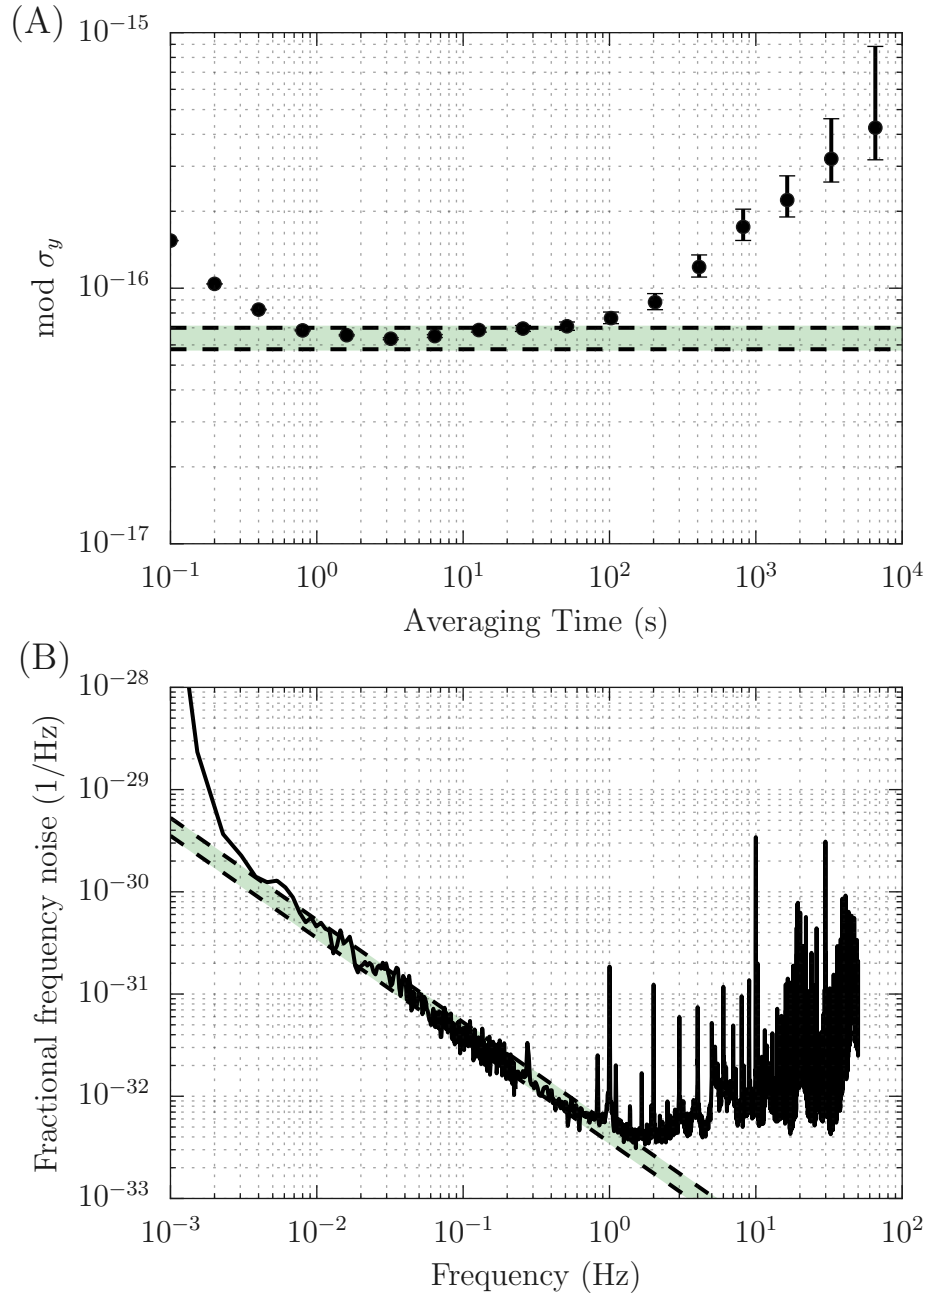


Figure 2.9: (A) Modified Allan deviation for the Si4 cavity. (B) Fractional frequency noise power spectral density (S_y) of the Si4 cavity. In both panels, the shaded green region indicates the spread of the estimated thermal noise floor given two previously published values for the coating loss angle [65, 66].

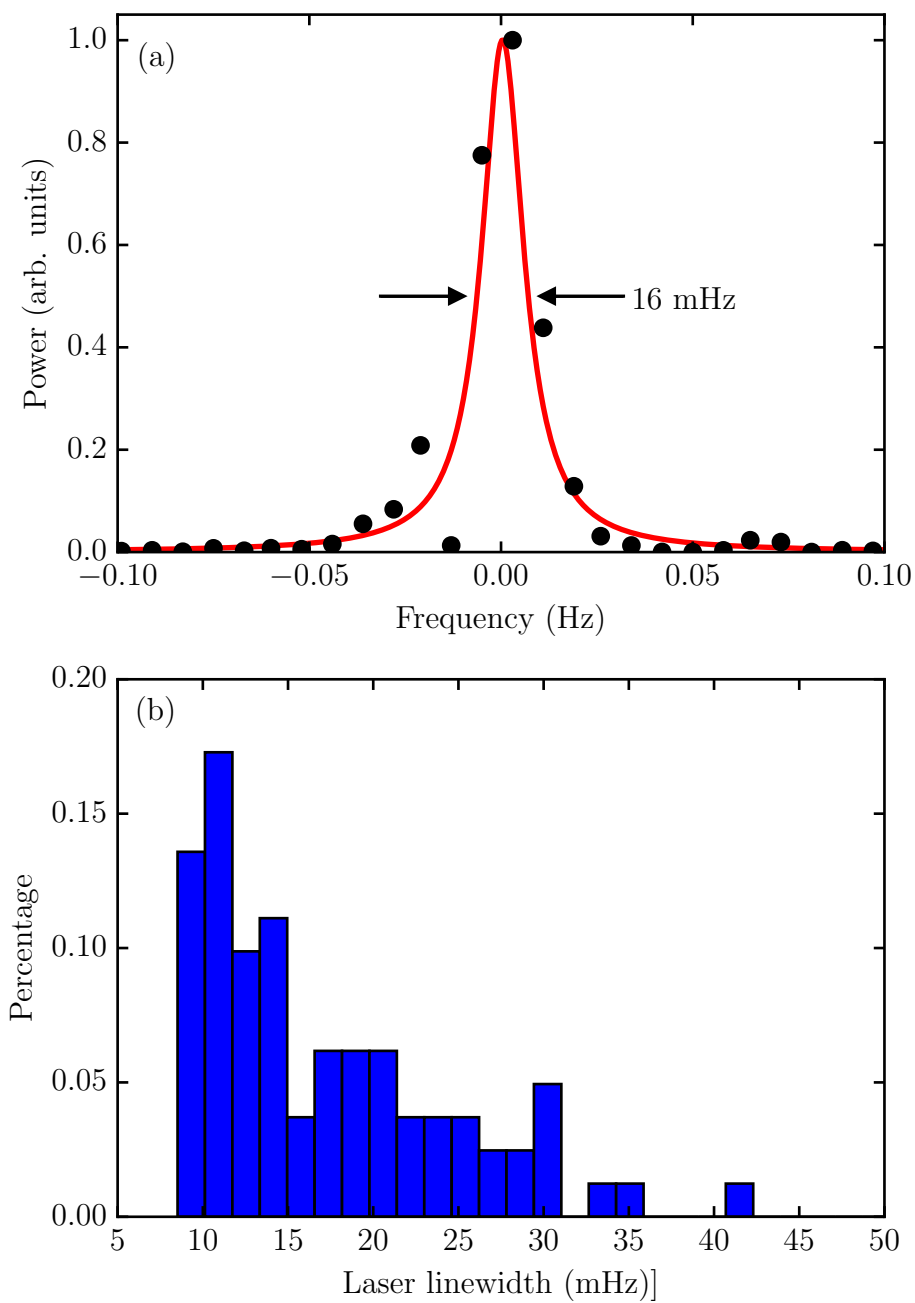


Figure 2.10: (A) FFT of the beat measured at 1542 nm (black circles), fit to a Lorentzian lineshape (red line). (B) Histogram of the measured Si4 linewidths for 100 measurements. The Si4 linewidths are obtained by multiplying the beat linewidth by 0.85.

independent check, the drift of the Si3 system is continuously monitored against a hydrogen maser from NIST via an optical frequency comb. This maser is then calibrated against UTC(NIST) as depicted in Fig. 2.8. The long-term linear frequency drift of Si3 is $-3 \times 10^{-19}/\text{s}$ with $2.8 \mu\text{W}$ of transmitted power. The measured linear drift of Si3 is removed from the Si3-Si4 beat, thus giving the drift of Si4.

The linear frequency drift of the Si4 cavity is dependent on the transmitted optical power as shown in Fig. 2.11. We vary the incident power and stabilize the cavity transmission at various levels as shown in Fig. 2.8. With a cavity finesse of $\mathcal{F} = 500,000$, total transmission of $T = 2$ ppm, and total loss of $A = 4$ ppm, a transmitted power of 40 nW corresponds to a circulating power of 2 mW. Each time the optical power in the cavity is changed, a frequency transient is observed. In order to extract the linear frequency drift, we wait several time constants for the transient to decay away. Interestingly, the observed time-constant for the settling of the cavity to its asymptotic drift rate depends on the final power setting, as shown in Fig. 2.12. This means that at the lowest optical power (40 nW in transmission), we needed to take data for about 1 month to obtain a reliable estimate of the frequency drift. We note that in order to achieve high performance at low optical power by employing resonant photodetectors for both the PDH and the RAM detection, providing a shot-noise limited signal-to-noise ratio at 68 nW (input referred noise of $0.15 \text{ pA}/\sqrt{\text{Hz}}$).

The linear power dependence of the drift is striking evidence for a new mechanism of length drift of an optical cavity at low temperatures. The sign of the frequency drift is always negative, meaning the physical length of the cavity is getting longer over time. The slope of the power dependence is roughly $-7 \times 10^{-21}/\text{s/nW}$. One potential explanation is thermal-induced mechanical creep of the mirror coating, where the mismatch in the coefficient of thermal expansion for the substrate and the coating gives a temperature-dependent creep. To reduce the impact of optical power on the long-term drift, Wiens et al. minimized the irradiation of their mirrors by periodically scanning the laser across the cavity resonance to measure the cavity frequency [39]. We present the first rigorous characterization of a power-dependent frequency drift in an optical cavity. The lowest operating power we have achieved is 40 nW in transmission, giving a fractional frequency

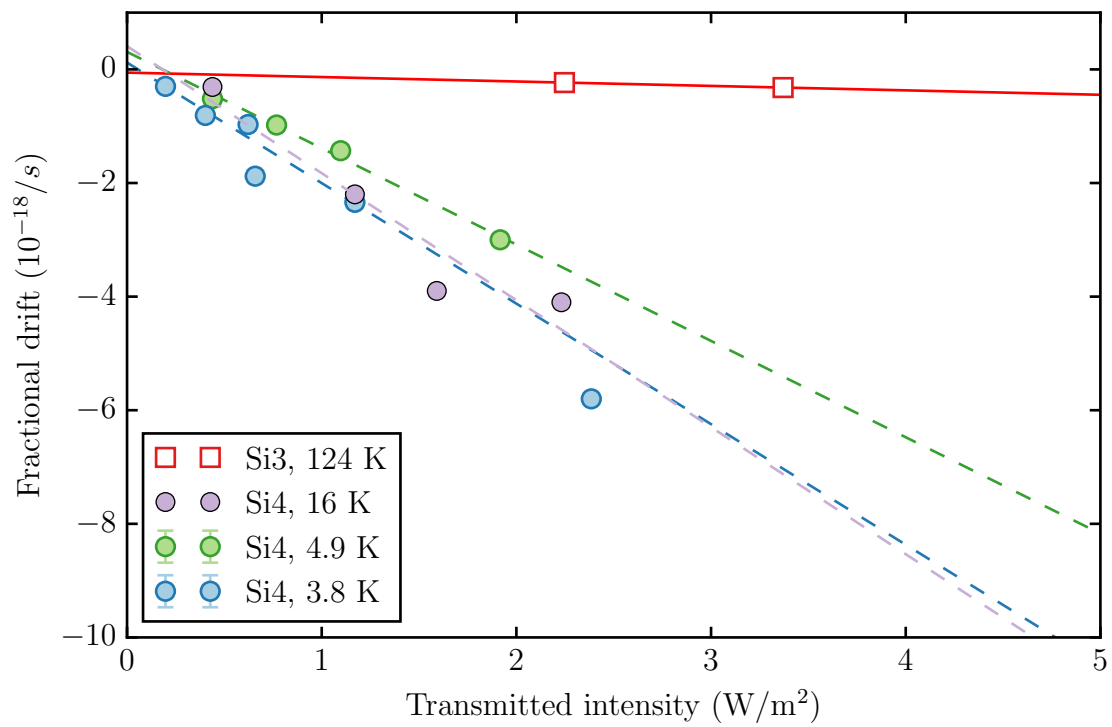


Figure 2.11: Fractional frequency drift in units of 10^{-18} for Si4 at different operating temperature. The Si3 drift versus intensity is shown as the open red squares, exhibiting a far smaller sensitivity to optical power.

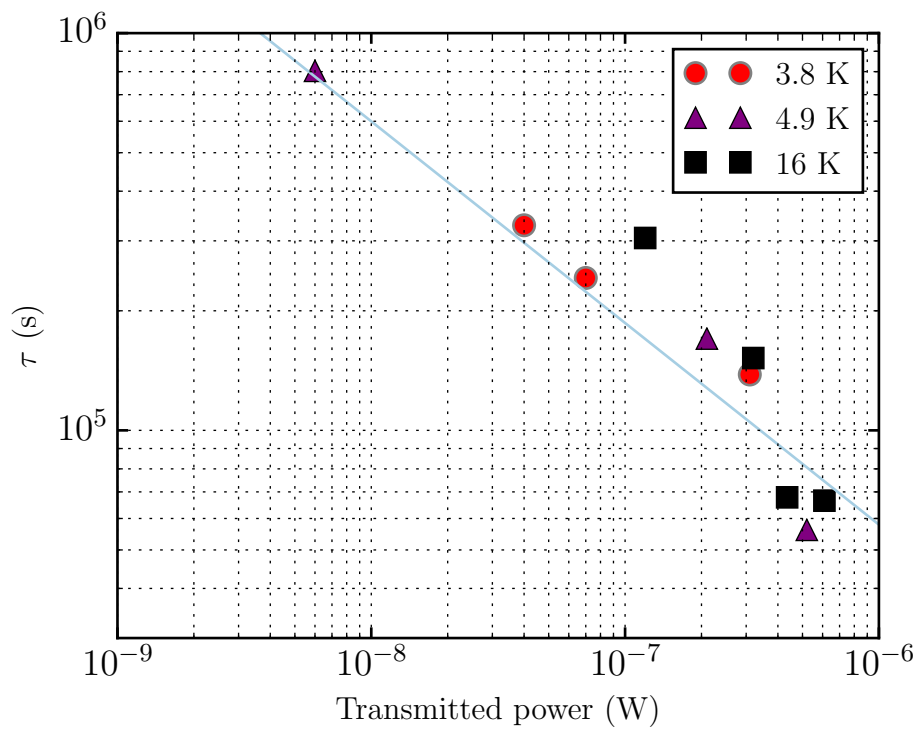


Figure 2.12: Fractional frequency drift in units of 10^{-18} for Si4 at different operating temperature. The Si3 drift versus intensity is shown as the open red squares, exhibiting a far smaller sensitivity to optical power.

drift of -3×10^{-19} /s. This frequency drift is comparable to the previous state-of-the-art obtained from a 124 K silicon cavity [67].

However, the implication of the current finding is tantalizing in that as we continue to reduce the incident power, we can access an extremely low value of cavity drift, making it possible that such a cavity alone could be useful as a potential time scale. At this low power, the fractional noise of the laser is higher than that showed in Fig. 2.9 by about a factor of two. The extra noise is due to the photodetector, and will be mitigated with an improved design. We also note that since this work, crystalline mirrors have been implemented in the "Si6" system. There, the overall magnitude of the drift is much smaller [68]. This supports our understanding that the drift in the Si4 system comes primarily from some physical effect in the coating, or in some light-matter interaction in the coating. It is interesting to note that a coating with more order (crystalline), yet with generally higher loss (defects), still gives an overall smaller drift.

The advances presented here point to a clear direction for ultrastable lasers. To operate with a minimal frequency drift, optical cavities at low temperature will have to operate at very low optical power. Reduction of the thermal noise will be possible by replacing the conventional $\text{SiO}_2/\text{Ta}_2\text{O}_5$ mirrors with crystalline mirrors [69, 70]. Such crystalline mirrors have been shown to exhibit a factor of 10 lower loss angle at room temperature [70]. Increasing the cavity length further will also reduce the fractional frequency noise. We can now foresee a strong possibility of achieving an ultrastable cavity with fractional instability $< 1 \times 10^{-17}$ using a continuously-running closed-cycle cryocooler at 4 K.

2.6 Inferring loss angle of $\text{SiO}_2/\text{Ta}_2\text{O}_5$ coating

As we discussed in Section 2.2.1, mechanical dissipation in the coating often dictates the amplitude of length fluctuations. These fluctuations are referred to as Brownian coating thermal noise (CTN), and can limit the precision of interferometric measurements. Thus, there has been considerable experimental effort to characterize the level of CTN present in current optical coatings.

There are two distinct methods for determining the Brownian CTN of an optical coating. The first is referred to as the "mechanical ringdown approach". This involves measuring the mechanical quality factor, Poisson ratio, and Young's modulus for each coating material. These mechanical properties for the coating and substrate are then used to calculate the CTN for a given mirror. This method of calculating the CTN may not account for all multilayer phenomena in the coating. The second method, referred to as a "direct measurement", determines the CTN of the coating by measuring the frequency stability of optical cavities. This approach has the challenge of extracting a coating loss angle from a single measurement over a broad frequency range, requiring input on other coating and substrate properties. Due to the challenges in each approach, it is vital that both are undertaken as independent and complementary research efforts. In this Section, we present direct CTN measurements of $\text{SiO}_2/\text{Ta}_2\text{O}_5$ using the most stable cavities in operation today at cryogenic temperatures up to room temperature. To make calculations of the fundamental noise floor of a given optical resonator, including both reference cavities or gravitational wave detectors, one must know the "loss angle" of the coating material. As of right now, all gravitational wave detectors make use of amorphous coatings, mainly made up of $\text{SiO}_2/\text{Ta}_2\text{O}_5$ dielectric stacks.

To extract a mechanical loss angle from the measured noise spectrum, we must perform a characterization of all other noise sources in the reference cavity systems. The 124 K system has been thoroughly characterized as described in [61]. Since that work we have made several improvements to the setup, including active stabilization of the transmitted optical power, active temperature control of the outermost vacuum chamber, and improved optical quality of the optics. All of these improvements are primarily focused on the lower frequency noise, allowing us to characterize the $1/f$ thermal noise over a wider frequency range.

The technical noise budget of the Si4 system is shown in Fig. 2.7. We show fits to the measured noise arising from temperature, Pound-Drever-Hall (PDH) servo, and the PDH photodiode (PD). The dominant technical noise source is residual amplitude modulation (RAM), shown as the red data. The sum of the technical noise is well below the measured noise. Figure 2.7B shows the intrinsic thermal noise of the cavity, including Brownian noise from the spacer (dotted), substrate

(dash-dot) and coating (dashed). The measured noise is shown (black) along with the sum of technical noise and thermal noise (green). The CTN is the dominant noise source from 5 mHz to 0.8 Hz.

The MJM cavity has routinely performed at a fractional frequency stability of 1×10^{-16} . Since we have not developed a thorough noise budget for this system, we report only an upper limit on the loss angle.

With the technical noise on the cryogenic cavities being sufficiently low, we can extract a mechanical loss angle at each respective temperature. We assume that the mechanical loss in the parallel and perpendicular directions are identical ($\phi_{\parallel} = \phi_{\perp}$). Our measurements presented in this manuscript do not distinguish these two loss components. Without loss of generality we lump them together as a total loss angle. We also note a gradual shift away from the loss angle convention of ϕ_{\parallel} and ϕ_{\perp} in the LIGO community, and the loss angle can instead be defined in terms of bulk loss ϕ_B and shear loss angles ϕ_S [71]. Approximately, the assumption of $\phi_B = \phi_S$ is equivalent to that of $\phi_{\parallel} = \phi_{\perp}$, without detailed treatment of light penetration into the coating layers. However, the reference [71] argue that using ϕ_{\parallel} and ϕ_{\perp} can lead to negative strain energies for a non-zero Poisson ratio $\sigma \neq 0$.

For the i^{th} mirror, the expression for the fractional frequency PSD arising from coating Brownian noise is then [72]

$$S_y^i(f) = \frac{2k_B T d}{\pi^2 f L^2} \frac{1 - \sigma_{sub}^2}{\omega_i Y_{sub}^2} \frac{\phi_c(f)}{Y_c (1 - \sigma_c^2) (1 - \sigma_{sub}^2)} \times [Y_c^2 (1 + \sigma_{sub})^2 (1 - 2\sigma_{sub})^2 + Y_{sub}^2 (1 + \sigma_c)^2 (1 - 2\sigma_c)] \quad (2.3)$$

Here, "c" labels the coating and "sub" labels the substrate. $\sigma_{c(sub)}$ is the Poisson's ratio, $Y_{c(sub)}$ is Young's modulus, d is the coating thickness, ω_i is the $1/e^2$ beam radius at the i^{th} mirror, $\phi_{c(sub)}$ is the loss angle.

Note that we are allowing for the coating loss angle to have frequency dependence. Several authors have reported a frequency-dependent loss angle for SiO₂/Ta₂O₅ coatings [73, 74]. We fit the measured spectra to the functional form of $S_y(f) = a/f^b$ in order to allow for any potential

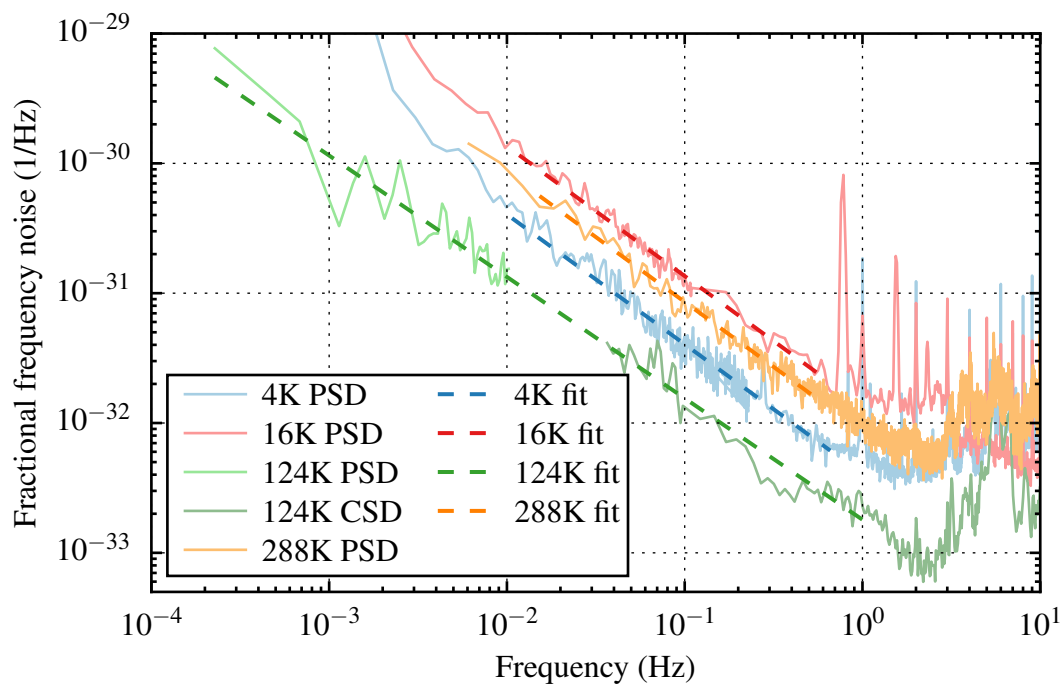


Figure 2.13: The measured noise PSD at various temperatures (solid curves). The dashed lines are the fits to $S_y(f) = a/f^b$.

frequency dependence. A deviation from $b = 1$ for the CTN indicates a frequency-dependent loss angle. The fits to the measured spectra are shown as dashed lines in Figure 2.13. We find weak dependencies on frequency as summarized in Figure 2.14.

With the fits from Fig. 2.13 and Eq. 2.3, we derive $\phi_c(f)$ and plot the results in Fig. 2.14.

The results are

$$\phi_{4\text{ K}}(f) = (5.6 \pm 0.9) \times 10^{-4} \times f^{-0.05 \pm 0.01} \quad (2.4)$$

$$\phi_{16\text{ K}}(f) = (3.2 \pm 0.3) \times 10^{-4} \times f^{-0.11 \pm 0.02} \quad (2.5)$$

$$\phi_{124\text{ K}}(f) = (2.4 \pm 0.3) \times 10^{-4} \times f^{0.06 \pm 0.02} \quad (2.6)$$

The shaded band indicates the 1σ uncertainty. We find that the frequency dependence, although rather weak, has the opposite sign at 4 and 16 K compared with the 124 K result. The 4 and 16 K results use mirrors from a different coating run than the 124 K cavity, and so this difference could in principle be coating-dependent. However, the other measurements of the frequency dependence to the loss angle all have positive coefficients (see [74, 73]). It might be that there is a real change in the coating loss properties with the drastic change in temperature.

For comparison with other results published in the literature, we plot ϕ_c as a function of temperature in Fig. 2.15. The orange diamonds [75], green squares [77], and blue triangles [76] are ringdown measurements that do not explicitly consider ϕ_{\perp} and ϕ_{\parallel} . The pink square [25], black triangle [78] and purple diamond [79] are all from direct thermal noise measurements which assume $\phi_{\perp} = \phi_{\parallel}$. Our results are plotted at 1 Hz, and no frequency dependence has been assumed for the referenced values, which are largely measured in the kHz frequency range. We measured CTN at several temperatures near 4 K, and average these values together. The x-axis error bar for that data point indicates the spread in temperature explored, while the y-axis error bar is the standard deviation of the extracted ϕ_c . We do not see a strong temperature dependence between 4 to 16 K.

We have presented direct measurements of the thermal noise for $\text{SiO}_2/\text{Ta}_2\text{O}_5$ HR coatings. These measurements are complementary to those attained by ringdown spectroscopy. For the design of future gravitational wave detectors, independent measurements of coating thermal noise




Figure 2.14: Frequency-dependent mechanical loss angle for the SiO₂/Ta₂O₅ coatings at different operating temperature. The shaded bands indicate the 1σ uncertainty.




Figure 2.15: Mechanical loss of the individual coating constituents (SiO_2 , Ta_2O_5), and $\text{SiO}_2/\text{Ta}_2\text{O}_5$ coatings measured by ringdown (orange diamonds [75], blue triangles [76], green squares [77]). Other direct measurements are also included (black triangle [78], open black square [25], and purple diamond [79]). Our results from direct CTN measurements are shown as the red circles and the red arrow.

are important. For ultrastable reference cavities, the use GaAs/AlGaAs crystalline coatings at cryogenic temperatures is important to explore. Other potentially low thermal noise optical coatings could also be tested using cryogenic ultrastable silicon cavities.

2.7 Si3 : zero crossing of CTE

So far, we have focused on the Si4 cavity system, but for the optical clock experiments in the Ye lab, we have only used Si3, the 124K cavity system. In this section, we will describe the measurement of the zero crossing of the CTE of Si3. After the optical timescale data campaign [50], we were curious how much of the long-term instability was coming from temperature fluctuations. When the Si3 cavity was commissioned, we performed a broad scan of the cavity temperature with the help of PTB colleagues who had come to visit JILA. At that time, we beat the Si3 laser against Si4, which did not have its power-dependent drift worked out yet. Thus, we decided to remeasure the zero-crossing, since we weren't exactly sure how far off we were. During the Christmas break of 2017, a quiet lab allowed for taking the Si3 offline for some characterization work.

To measure the zero-crossing of the CTE, we want to sweep the cavity temperature over a small range, say around 10 mK. We then record the cavity frequency (length) versus the temperature of the cavity, and should observe a parabola when we cross through zero. Sweeping the cavity temperature over a reasonable time scale can be tricky, since the time constant from the active shield to the cavity is 12 days. Thus, one has to use “thermal inertia” in order to sweep the temperature. We first make a large step to a lower temperature, wait some time, and then jump up above the initial setpoint. By using a detailed thermal model of the cryogenic system, and continuously checking the parameters versus the measured temperatures, one can predict how the thermal system will respond, and fine-tuning of the step amplitudes and time can be performed. The cavity temperature is read by a thermocouple, which can exhibit jumps in its reading. For this reason, this sensor is only meant as a somewhat rough out-of-loop monitor, and the inloop sensors are precision PT-100 thermistors installed on the radiation shields. So during the sweep, one has to keep track of any inadvertant noise jumps in the data. This becomes especially important

when you have determined the zero-crossing point, but now need to steer the system to the correct temperature reading.

With the cavity temperature now sweeping linearly, we measure the Si3 frequency against Si4 and the hydrogen maser reference from NIST (Figure 2.16). The measurement against the maser has much more short-term frequency noise, but serves as a consistency check for the fact that we measured against Si4, another drifting system. We have measured the zero crossing to the μK level, and can easily steer to this temperature to within 1 mK.




Figure 2.16: Si3 zero crossing measurement of the coefficient of thermal expansion. Measured both against Si4 and the hydrogen maser. The fitted zero-crossings are consistent, with $T_0 = 123.345542(2)$ K from Si3-Si4, and $T_0 = 123.34594(7)$ K from the Si3-maser.

Chapter 3

Classical clock demonstrations and applications

In this chapter, we briefly overview various clock results using classical, uncorrelated atomic ensembles. These experiments take direct advantage of the advances in ultrastable laser technology. We first describe the clock comparisons between Sr1 and Sr2 that were undertaken during the second half of 2018 [17]. Then, we discuss placing limits on ultralight bosonic dark matter, coming from precision comparisons between the Si cavity, Sr atoms and a hydrogen maser [49]. After this, we switch gears toward time-keeping, discussing a measurement campaign made by the Boulder Area Clock Optical Network (BACON), where in 2018 we measured three optical clock ratios at the 18th digit. This lays the groundwork for the eventual redefinition of the SI second, and also provides further opportunities for testing fundamental physics. These measurements and applications push the frontier of what one can achieve state-of-the-art optical clock technology, without introducing quantum entanglement. This chapter sets the stage for the rest of the thesis, where we report our effort to push a further boundary by introducing quantum correlations and spin-squeezing.

3.1 Classical clock comparisons at the quantum limit

After many years of hard work spanning multiple PhD timescales [80, 24, 81], our group developed the first cryogenic silicon cavity to reach the thermal noise floor [25]. This new laser had unprecedented frequency stability - reaching a laser linewidth below 10 mHz. Two of these systems had been built at PTB in Germany, and it was time to transport one of the systems to JILA. Eric Oelker and Lindsay Sonderhouse both travelled to Braunschweig Germany, assisted in packing up

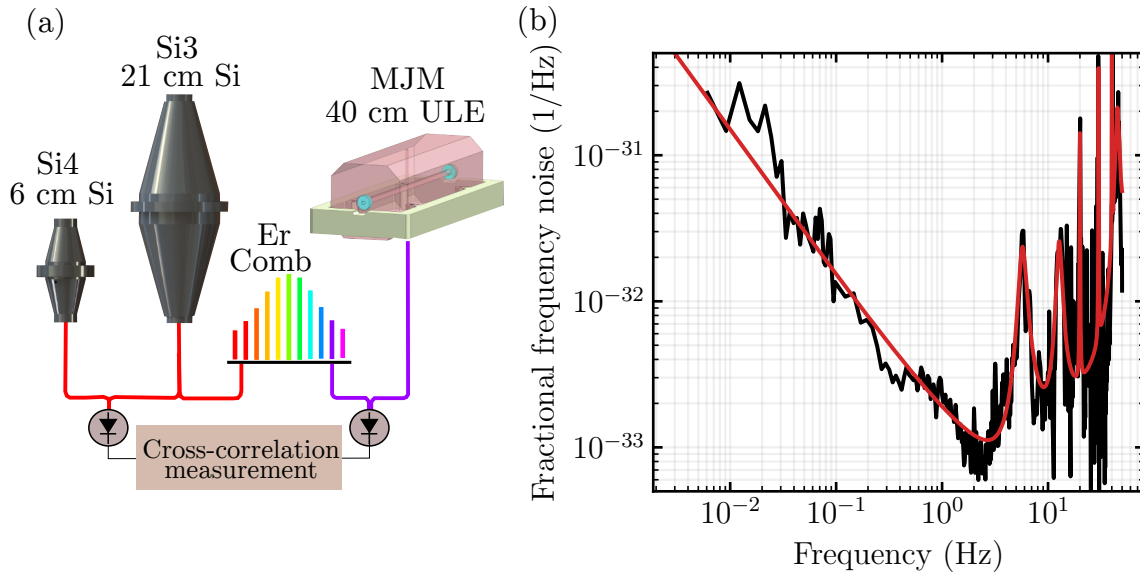


Figure 3.1: (a) Schematic of the cross-correlation method. Each laser is fiber noise cancelled to the frequency comb table, and we measure the beat between Si4 and Si3, and MJM vs Si3. (b) The black data is the cross-correlation spectrum for the Si3 system (21 cm Si). The red line is a laser model that we will use for Dick effect calculations and dark matter limit simulations.

the systems, and hand-carried the cavity to Boulder, Colorado. Eric led the effort to rebuild the system, and very quickly we were very happy to see that the Si3 system was dipping below the Si4 stability within just a few weeks of operation at 124 K. It seemed like a prime time to conduct some new clock comparisons with this new clock laser.

As we mentioned in Chapter 1, neutral atom clocks make use of large ensembles in order to make the influence of QPN negligible. However, it is quite difficult to operate a clock at the QPN limit, due to frequency noise on the clock laser. Aliased high frequency noise, known as the Dick effect, can contribute significant amounts of clock noise. Until this work, the laser noise had limited the direct observation of anti-synchronous clock comparisons to the low $10^{-16}/\sqrt{\tau}$ level, where τ is the averaging time in seconds [41, 82, 83]. By employing the improved optical local oscillator, we sought to push single clock stability into the $10^{-17}/\sqrt{\tau}$ level for the first time. To this end, we performed clock comparisons between the 1D Sr1 clock and the 3D Sr2 clock.

We start by characterizing the clock laser performance by using a cross-correlation measurement (Figure 3.1a). The measured cross correlation for Si3 is shown as the black data in Figure

3.6b. The noise model includes a thermal noise term of $1.5 \times 10^{-33}/f$. There are several peaks in the spectrum. The peak near 5 Hz is a vibration peak that shows up on all three lasers, and at a larger magnitude on both the MJM and Si4 systems, making us suspect that the true magnitude on the Si3 laser system could easily be smaller than observed and modeled here. Two peaks at 20 Hz and 30 Hz are likely to arise from HVAC motor vibrations. Also included are white frequency and white phase noise terms.

For this comparison, the Sr1 1D clock (before the rebuild described in [84]) operated with 1,500 - 2,000 atoms, probing the $|^1S_0, m_F = \pm \frac{9}{2}\rangle \rightarrow |^3P_0, m_F = \pm \frac{9}{2}\rangle$ transitions. The \pm label refers to the averaging of the two transitions to become first order insensitive to magnetic field fluctuations. The Sr2 system is a 3D optical lattice, and for this experiment was run without their typical evaporation cooling stage [85]. The Sr2 system also operated on the magnetically insensitive σ transition, $|^1S_0, m_F = \pm \frac{3}{2}\rangle \rightarrow |^3P_0, m_F = \pm \frac{5}{2}\rangle$, which is 22 times less sensitive to the magnetic field noise than the $\pm 9/2$ transition [86]. Since this work, several other groups/experiments have begun using this particular clock transition [87, 18].

Furthermore, to achieve the best clock performance, we had to hunt for various sources of noise in the systems. We initially stabilized the 0th order of an AOM, and probed the atoms with the 1st order. We found that differential noise between these two orders gave noise beyond our desired level of precision. With some ingenuity (a combined effort of Christian Sanner, Ross Hutson, Eric Oelker, Colin Kennedy and others on the Sr team), a simple solution was found for Rabi spectroscopy. We retroreflect the 1st order beam for stabilization, but need to deal with the turning on and off of the pulse accordingly, without introducing significant phase transients when the laser is on. To work around this, we first detune the laser from the clock transition by 1 MHz, turn on the light to a low level, and engage the fiber noise servo. The laser is then swept onto resonance and turned to the clock drive power level. This solution works great for Rabi spectroscopy because the time when the laser is swept onto resonance (ms level) is short compared to the entire Rabi pulse duration of 550 ms. Furthermore, we characterized the overall frequency shift uncertainty of this probing scheme to very high precision [14]. Ramsey spectroscopy is further

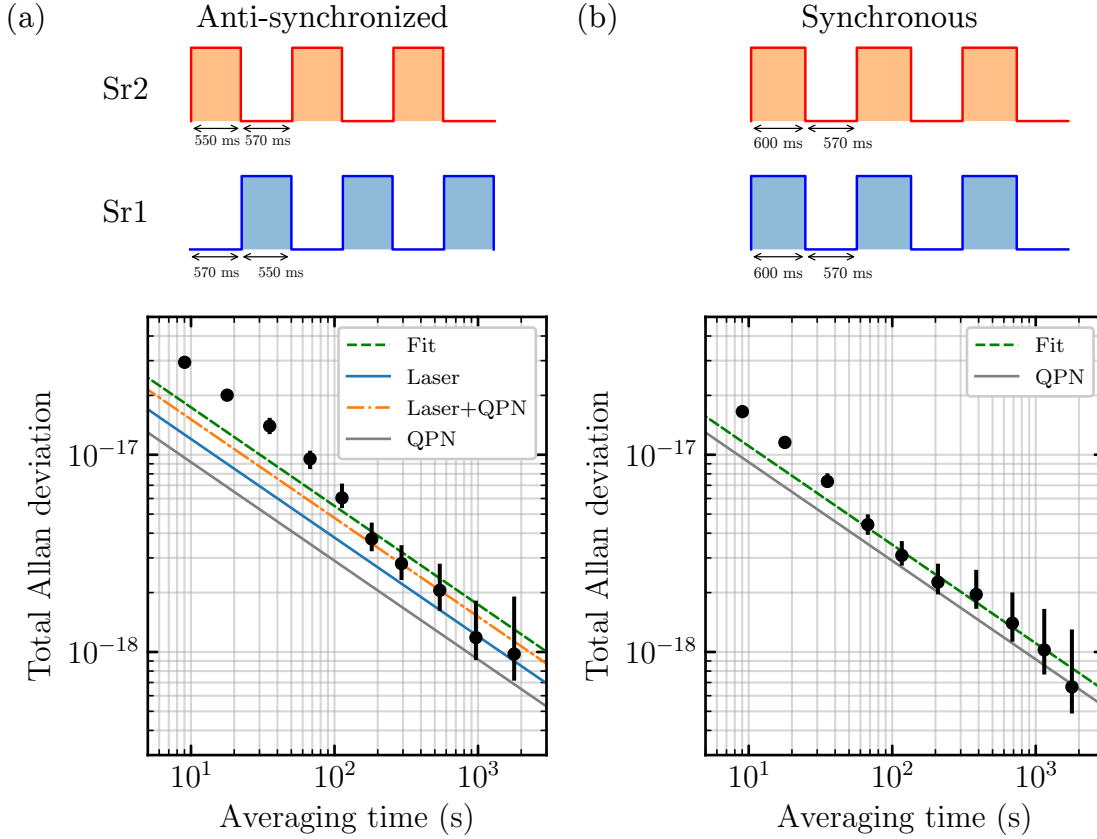


Figure 3.2: (a) Anti-synchronized clock comparison between Sr1 and Sr2. The solid gray line is the calculated QPN, blue line is calculated Dick effect, orange dashed-dotted is the quadrature sum of QPN and the Dick effect laser noise, and the dashed green line is the fitted stability, giving $5.5 \times 10^{-17}/\sqrt{\tau}$. (b) Synchronous comparison, showing a single clock stability of $3.5 \times 10^{-17}/\sqrt{\tau}$.

complicated when you want to stabilize the 1st order, but has been implemented with clever tricks.

The comparison stabilities are shown in Figure 3.2. The anti-synchronized comparison is run so that the two clocks have no temporal overlap of the clock pulses, with a pulse time of 550 ms, and the dead time tuned so there is a buffer of 10 ms on either side of the pulse. The Dick effect is calculated to be $3.8 \times 10^{-17}/\sqrt{\tau}$ using the laser noise model for the Si3 laser, plotted as the solid blue line. Shown as the gray line is the anticipated QPN level is $2.9 \times 10^{-17}/\sqrt{\tau}$, and we combine the Dick effect and QPN for the orange dashed-dotted line. The fitted stability of $5.5(3) \times 10^{-17}$ is in good agreement with the anticipated stability of $5.7 \times 10^{-17}/\sqrt{\tau}$.

To remove the common-mode laser noise, we perform synchronous comparisons (Figure 3.2b).

Here, the observed stability is $3.5(2) \times 10^{-17}/\sqrt{\tau}$, just 20% (1.6 dB) above the predicted QPN of $2.9 \times 10^{-17}/\sqrt{\tau}$. This is the closest to the QPN level for any independent, neutral atom optical clock comparison to date. It is this very fact that motivates the development of a spin-squeezed clock in our lab. State-of-the-art clocks are running into the fundamental QPN limit.

3.2 Dark matter limits from precision spectroscopy

Overwhelming evidence from astrophysics points to the fact that dark matter makes up the majority of the observable matter density in the universe [88]. However, direct detection of dark matter has so far alluded the scientific community, despite significant efforts spanning generations of experimental astrophysicists [89]. Ultralight bosonic dark matter is one proposed type of dark matter out of a great soup of theoretical proposals [90]. Using our precision measurement techniques, we can make measurements that are sensitive to fundamental constants.

3.2.1 Sensitivity to fundamental constants

We first consider the frequency ratio of our optical cavity to the Sr atom f_{Sr}/f_{Si} , and show that it is sensitive to the fine-structure constant, $\alpha = \frac{e^2}{hc}$. To see how this scaling comes about, we make some simple considerations. The strontium optical transition frequency scales as

$$f_{Sr} \sim \frac{R_y}{h} \quad (3.1)$$

where R_y is the usual Rydberg constant. Then, let's suppose the silicon cavity frequency is given by $f_{Si} \sim \frac{c}{\lambda}$, and that λ must be an integer multiple of the Bohr radius,

$$f_{Si} \sim \frac{c}{Na_B} \quad (3.2)$$

where c is the speed of light, N is the number of silicon atoms making the length of the cavity, and the a_B is the Bohr radius. Then, in our measurements we look at the frequency ratio of the strontium frequency and the silicon cavity frequency,

$$\frac{f_{Sr}}{f_{Si}} \sim \frac{R_y Na_B}{hc} \quad (3.3)$$

$$\frac{f_{Sr}}{f_{Si}} \sim \frac{\frac{m_e e^4}{h^2} N \frac{h^2}{m_e e^2}}{hc} \quad (3.4)$$

$$\frac{f_{Sr}}{f_{Si}} \sim \frac{(2\pi)^2 e^2 N}{hc} = b\alpha \quad (3.5)$$

where I've absorbed a constant factor $b = (2\pi)^2 N$. We see that the linear dependence on alpha naturally comes out of this simple derivation. The next step is to take into account relativistic corrections, which can be written quite clearly as an additive correction $f_{Sr} \sim R_y(1 + k\alpha^2)$.

Sensitivity to the electron mass is realized by comparing the silicon cavity against a hydrogen maser. The hydrogen maser operates on the hyperfine ground state transition, which has the usual scaling of

$$f_H \propto \left(\frac{e^2}{a_B \hbar} \right) [\alpha^2 F_{rel}(Z\alpha)] \mu_p \left[\frac{m_e}{m_p} \right] \quad (3.6)$$

where α is the fine-structure constant, μ_p is the magnetic moment of the proton, m_e and m_p are the electron mass and proton mass respectively, and $F_{rel}(Z\alpha)$ is a relativistic factor depending on Z . The frequency of the silicon cavity is $f_{Si} = c/\lambda$. Using the silicon cavity frequency from Eq. 3.2, we take the ratio,

$$\frac{f_H}{f_{Si}} \propto N [\alpha^3 F_{rel}(Z\alpha)] \mu_p \left[\frac{m_e}{m_p} \right] \quad (3.7)$$

Thus, the silicon cavity - hydrogen maser comparison is linearly dependent on the electron mass.

For an optical timescale to be established, we need to run the silicon cavity against strontium over an extended (ideally forever) period of time. The comparison does not have to be continuous, but it is repeated many times nonetheless. We also naturally run the silicon cavity against hydrogen masers in order to connect to the more conventional timescales. Establishing the optical timescale thus gives us “free” access to this search of ultralight dark matter.

3.2.2 Dark matter couplings

The standard way to search for dilaton dark matter is to consider phenomenological couplings of the dark matter field $\phi(t)$ to various terms in the Standard Model Lagrangian. Following [91, 92],

we consider scalar couplings of the dark matter field ϕ to the Standard Model. If we consider couplings to the electromagnetic term and the fermionic mass term, we get the following dilaton Lagrangian (with $\hbar = c = 1$),

$$\mathcal{L}_\phi = \kappa\phi \left(\frac{d_e}{4e^2} F_{\mu\nu} F^{\mu\nu} - d_{m_e} m_e \bar{p} \bar{s} i_e \psi_e \right) \quad (3.8)$$

where d_{m_e} and d_e are the dimensionless coupling constants, d_{m_e} is the coupling of the dark matter to electron mass m_e , and ψ_e is the fermion spinor [92]. The dark matter field oscillates with the form

$$\phi(t) = \phi_0 \cos(m_\phi t + \delta) \quad (3.9)$$

The overall normalization is mass dependent

$$\kappa\phi_0 = \frac{6.4 \times 10^{-31} \text{ eV}}{m_\phi}, \quad (3.10)$$

where here (and only here), $\kappa = \frac{4\pi}{m_{Pl}}$ and m_{Pl} is the Planck mass. The $1/m_\phi$ scaling means that the amplitude of the expected dark matter increases for lower mass, such that the overall energy density of the dark matter is held fixed. The fermionic term in the Lagrangian gives a time-dependence to the mass of the electron (for details see Appendix A of [93],

$$m_e(\phi(t)) = m_e (1 + \kappa\phi_0 d_{m_e} \cos(m_\phi t + \delta)). \quad (3.11)$$

and similarly for the electromagnetic field term, we obtain time dependence of α ,

$$\alpha(\phi(t)) = \alpha (1 + \kappa\phi_0 d_e \cos(m_\phi t + \delta)). \quad (3.12)$$

Thus, if one can find a measurement that is sensitive to the electron mass then a limit on the magnitude of d_{m_e} may be set.

For the silicon-maser data, we use the Si3 cavity locked to stabilize an Er: fiber optical frequency comb (same setup described in [17]). A hydrogen maser (ST-14) signal at 1 GHz is transferred from NIST to JILA over a phase-stabilized optical fiber to the strontium lab [94]. The maser signal is mixed with the repetition rate of the frequency comb to a frequency of 1 MHz and then counted with a lambda-type frequency counter. This gives us the direct frequency comparison between the




Figure 3.3: (a) Fractional frequency deviation of the Si - Sr frequency comparison. The data has been dedrifted by fitting the data to Eq. 3.13. (b) Power spectral density of the Si-Sr data. The model for the Si cavity is shown as the dashed line. (c) Fractional deviation of the Si - maser frequency difference. (d) PSD of the Si - maser data. The maser model has a white noise term, and a broadened resonance to account for the microwave link noise. The silicon model is the same as for (b).

silicon cavity and the hydrogen maser. For the optical data, silicon cavity is locked to the Sr1 clock [14]. The control signal of the atomic clock lock yields the silicon cavity - strontium relative frequency.

The data campaign spans from Modified Julian Date (MJD) 58443 - 58477. The raw data for the two datasets are shown in Figure 3.3a and Figure 3.3c, in terms of fractional frequency deviation versus the elapsed time. The data has been dedrifted by fitting to a model of the form,

$$f(t) = a + bt + ce^{-t/d}, \quad (3.13)$$

where (a,b,c,d) are the fit parameters. The power spectral density (PSD) of the residuals is computed using the Lomb-Scargle periodogram, shown as the red data points in Figure 3.3b for the optical data and Figure 3.3d for the microwave data. The microwave PSD is calculated for Fourier frequencies from $0.7\mu\text{Hz}$ to 1 mHz. We leave out Fourier frequencies above 1 mHz for two reasons. One is that due to the noise level, the bound is considerably higher than currently set limits. Another is to simplify the analysis, since the noise above 1 mHz is dominated by the microwave link noise.

In order to set a dark matter bound, we require the 95% confidence limit of the PSD. This is achieved by establishing a noise model for the measurement and simulating many time-domain traces that have the same noise as the model. In Figure 3.3d, we plot traces of the Si3 noise model (dashed green). Also plotted is the white noise level of the Hydrogen maser (dashed red) of

$1.1 \times 10^{-25} \text{ Hz}^{-1}$, which corresponds to an Allan deviation of $1.5 \times 10^{-13}/\sqrt{\tau}$. The measured noise is largely dominated by the maser and some link noise at around $2 \times 10^{-4} \text{ Hz}$. The link noise is modeled as a sinusoidal signal at $2 \times 10^{-4} \text{ Hz}$ with phase noise introduced to approximate the shape of the link noise. We then simulate 100 different time-domain traces, compute the Lomb-Scargle for each generated trace and build a statistical distribution. The 95% confidence limit is then calculated from the distribution.

3.2.3 Dark matter Limit

With the 95% confidence limit for the PSD in hand, one can calculate the dark matter limit. We use the same normalization. The experimental bound for the silicon-hydrogen maser is shown as the black curve in Figure 3.4. This bound includes corrections coming from the impact of dedrifting the frequency measurement and any potential degradation in sensitivity due to the unknown phase of the dark matter. The discussion of these effects may be found in the following section.

In Figure 3.4b, we plot various other experimental bounds on the d_{m_e} coupling. The torsion pendulum bounds are obtained by searching for a fifth force, since the existence of a scalar field ϕ would result in the exchange of virtual ϕ quanta [100]. The MICROSCOPE experiment is a space-based torsion-pendulum [101, 98]). We note that in the 2018 PRL, the d_{m_e} result is not directly plotted, but another paper shows the limit on d_{m_e} [93].

3.2.4 Extra details on our analysis

In this sub-section, we expand on a few details regarding our data analysis. Some other details may also be found in the supplemental materials of [49]. To check the normalization of the Lomb-Scargle periodogram, we took two independent approaches. The first, most straightforward check is by checking the amplitude of a sine wave. We generate a sine wave with a given amplitude of 1×10^{-12} . In order to verify the amplitudes of the spectral densities, we simply integrate each amplitude spectral density (ASD). The Scipy (a commonly used Python package) ASD gives us the correct rms value of 7×10^{-13} , whereas the Lomb-Scargle gives us an amplitude that is exactly a

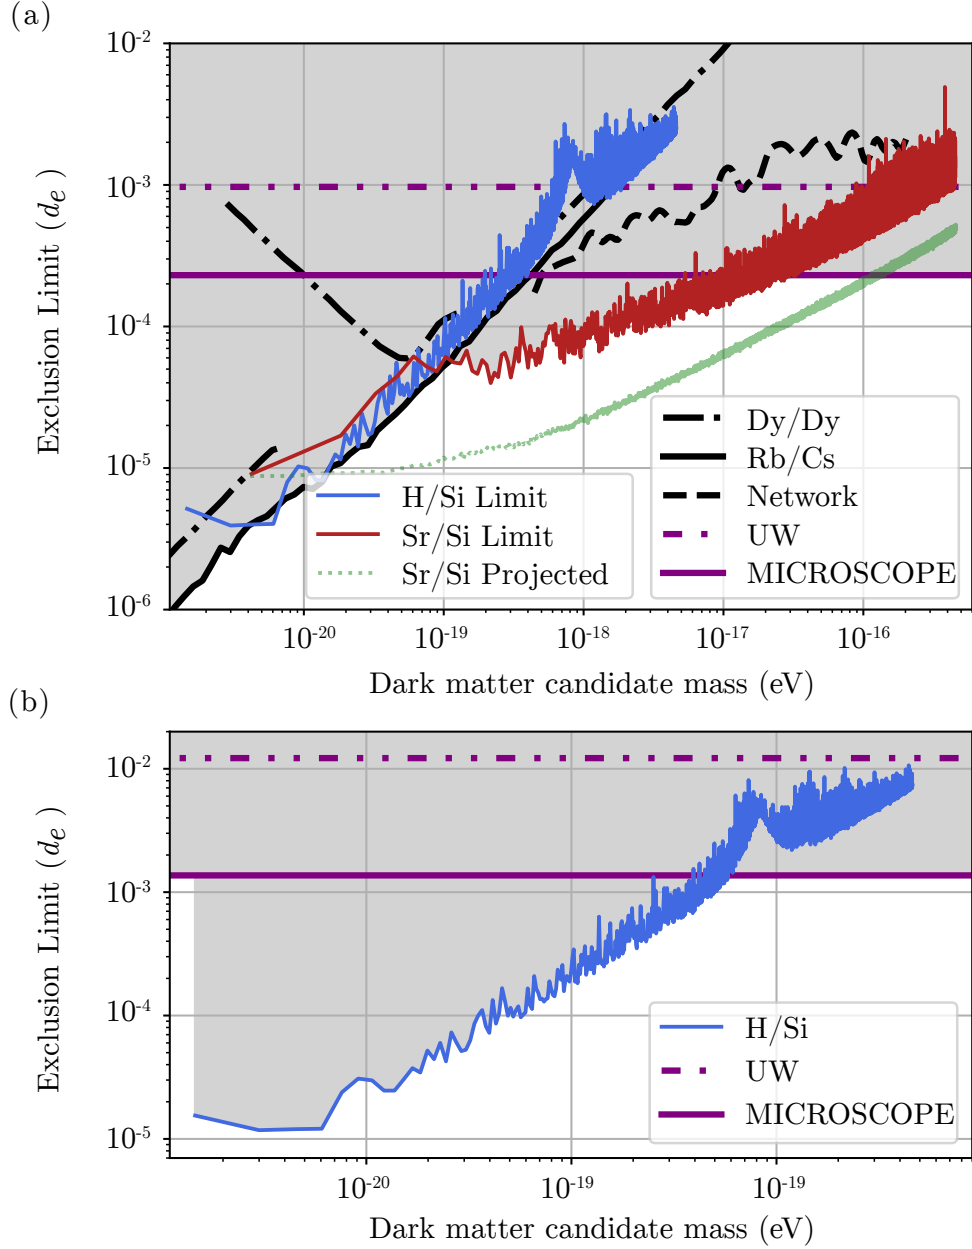


Figure 3.4: (a) 95% confidence limits placed on the dimensionless couplings to the electromagnetic field d_e , showing the improved limits set by $f_{\text{Sr}}/f_{\text{Si}}$ (red) and $f_{\text{H}}/f_{\text{Si}}$ (blue) ratios in the mass range above 1×10^{-19} eV. The maximum projected sensitivity for a search of the same 11-day duration without data gaps is included (light green), highlighting the potential of our technique. Limits derived from previous spectroscopic searches (black lines), with the network [95], Dy-Dy comparison [96], and Rb/Cs [97]. Limits from equivalence principle tests (purple lines), with MICROSCOPE [98] and UW [99]. (b) Demonstration of a significantly improved limit on the electron mass modulus (d_{m_e}) derived from the $f_{\text{H}}/f_{\text{Si}}$ (blue) ratio. Limits from equivalence principle tests (purple lines), including MICROSCOPE [98] and measurements from UW [99]. Shaded regions in both (a) and (b) are excluded at the 95% confidence level given the observed signal and noise models.

factor of two smaller. One possibility is that the Scipy ASD is giving the one-sided spectral density whereas the Lomb-Scargle is giving the two-sided. In any case, we will multiply the resulting Lomb-Scargle periodograms by a factor of two for amplitude, and by four for power spectral density (PSD). The second check we did was to generate time-domain data of white noise, flicker frequency noise and random walk frequency noise with known amplitude, and checking the Lomb-Scargle vs FFT.

Since our dark matter limit for d_{m_e} extends down to a Fourier frequency of $0.7 \mu\text{Hz}$, we were curious to check the effect of dedrifting the signal, as well as a random phase. In order to explore this effect, we add a hypothetical sine-wave signal to the raw data (before dedrifting) with a known amplitude of 1×10^{-13} . We then dedrift the data+sine by fitting the drift model given by Eq 3.13. Then we can take a Lomb-Scargle periodogram of the data and check the RMS value of the signal by integrating over the relevant frequency bins, and compare with the value we know we should obtain for such a sine wave. The result of this simulation is plotted in Figure 3.5. We plot the ratio of the obtained dark matter signal amplitude to the expected amplitude. Simulating the time domain traces with a length of 30 days is computationally expensive, and so to reduce the time required for the simulations we only calculate the correction at the points plotted as the orange circles. We then fit the correction to a $1 - e^{-x/f_0}$ model so we can interpolate the correction to all frequencies for the dark matter bound. We note that these corrections are more relevant for the Si - maser dataset since it is three times longer, allowing us to reach lower Fourier frequencies.

The scalar dark matter described by Equation 3.9 involves an arbitrary phase δ that could in principle take on any value from 0 to 2π . Since the data is not perfectly sampled, we will have a non-uniform sensitivity function for dark matter fields possessing different phase δ . We have no way of knowing what the correct phase of the signal could be, so we evaluate this effect by injecting dark matter signals of known amplitude with varying phase δ . We then compare the obtained amplitude from our least-squares fit of the data with the actual injected amplitude, and plot this ratio as the blue datapoints in Figure 3.5. Similar to the dedrift correction, we fit the ratio to obtain the corrections at our final Fourier frequencies. For the final corrections to the dark matter limit, we multiply both the dedrift impact and the phase impact, and this is plotted as the solid

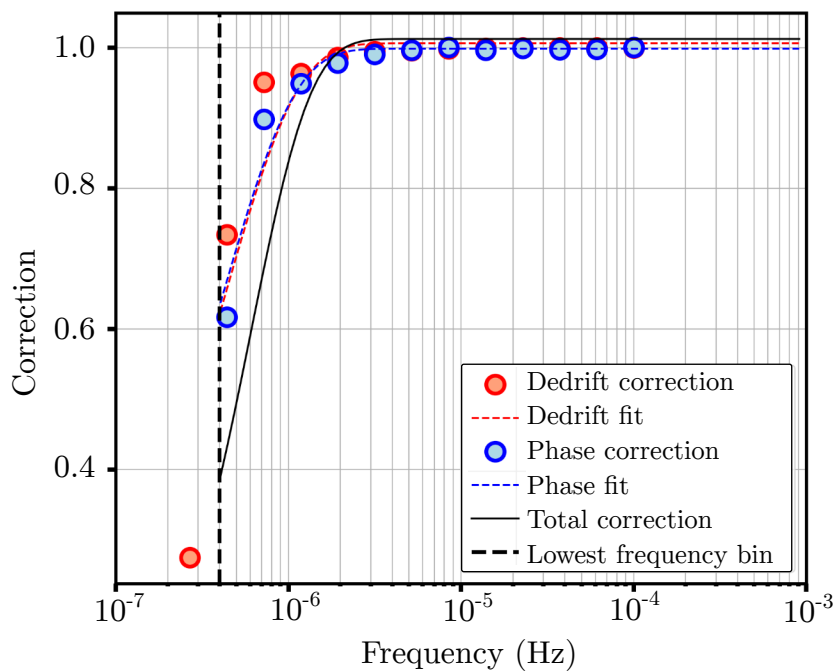


Figure 3.5: Corrections to the d_{m_e} limit from dedrifting and the unknown dark matter phase. The vertical dashed line indicates the lowest Fourier frequency for which a dark matter limit is computed.

black line in Figure 3.5.

3.3 Boulder Atomic Clock Optical Network

The Boulder Atomic Clock Optical Network (BACON) is a collaboration and corresponding network between the Sr lab at JILA and the Yb and Al⁺ ion groups at NIST. The relative density of high performing clocks in Boulder, Colorado is truly impressive. At the time of writing this Thesis, the Yb clock has a systematic uncertainty of 1.4×10^{-18} [15], the Al⁺ ion clock is 9.4×10^{-19} [16], and the Sr clock is at 2.0×10^{-18} [14]. In 2017, we embarked to measure the frequency ratios between these three atomic clocks. This optical clock frequency ratio measurement serves as a critical step toward the redefinition of the SI second. Optical frequency ratios are the best way to check consistency between optical clocks because they do not rely on cesium fountain clocks that currently define the second, which would limit such a comparison to the 10^{-16} level. By using an optical frequency comb as the flywheel, we can connect various optical standards across a wide range of frequencies [102, 103]. Since the optical frequency combs introduce minimal noise, we can measure the ratio with the utmost precision [104].

One of the first steps in the comparison was establishing a link between the JILA Sr clock and the clock/frequency comb labs at NIST. There exists a long network of optical fiber interconnecting various labs and institutes around the city of Boulder called the Boulder Research and Administrative Network (BRAN). Before the clock comparison, we revived the BRAN fiberlink between JILA and NIST. The servo loop had to be adjusted to prevent instability, since the finite time-delay leads to phase shifts that can give π phase shifts and corresponding instabilities [105]. The bandwidth of the servo ended up being approximately 5 kHz. We set up the fiber link in a loop-back configuration to evaluate the out-of-loop performance (Figure 3.6a). We install AOM-2 in the retroreflection of the light from ref #2, so that the light is marked in frequency and is distinguished from all of the other back reflections in the long fiber and interconnections. The inloop error is measured by out-coupling the error signal and counting it with a frequency counter (lambda-type, deadtime free), and is plotted as the red points in Figure 3.6b. We also measure the out-of-loop

noise, shown as the black circles. We see a two distinct bumps in the Allan deviation, likely due to temperature/pressure fluctuations of the out-of-loop fiber. The Allan deviation averages down to 5×10^{-19} , which we take to be the uncertainty of the fiber link. In Figure 3.6c, we plot a histogram of the out-of-loop frequency measurements, and show the residuals to the Gaussian fit in the bottom panel.

An example Allan deviation of a daily ratio measurement is shown in Figure 3.7. Within a single day, the neutral atom clock ratio of Yb/Sr averages down to the low 10^{-18} level, while the ratio with the single ion clock averages to 10^{-17} . The contributions to the instability are shown as the dashed lines, showing that they are low enough noise as to not contribute significant noise to the ratio. The ratios were measured over several months, and in the end the Yb/Sr data was over-scattered, partially due to this outstanding single-day precision. To treat the data from this campaign - we had to introduce the idea of “dark uncertainty”, which would account for unknown fluctuating systematics (this data analysis was done in collaboration with NIST statistics expert Amanda Koepke). The final ratios were measured to be

$$\begin{aligned}\nu_{\text{Al}^+}/\nu_{\text{Yb}} &= 2.162887127516663703(13) \\ \nu_{\text{Al}^+}/\nu_{\text{Sr}} &= 2.611701431781463025(21) \\ \nu_{\text{Sr}}/\nu_{\text{Yb}} &= 1.2075070393433378482(82),\end{aligned}\tag{3.14}$$

representing some of the most accurate measured quantities to date.

Using the same data from the clock comparison we can also measure the absolute frequency of the strontium clock transition. A detailed analysis of the connection between the optical clock and the international definition of the SI second determined the Sr absolute frequency to be [106]

$$\nu_{\text{Sr}} = 429228004229873.19(0.15) \text{ Hz}.\tag{3.15}$$

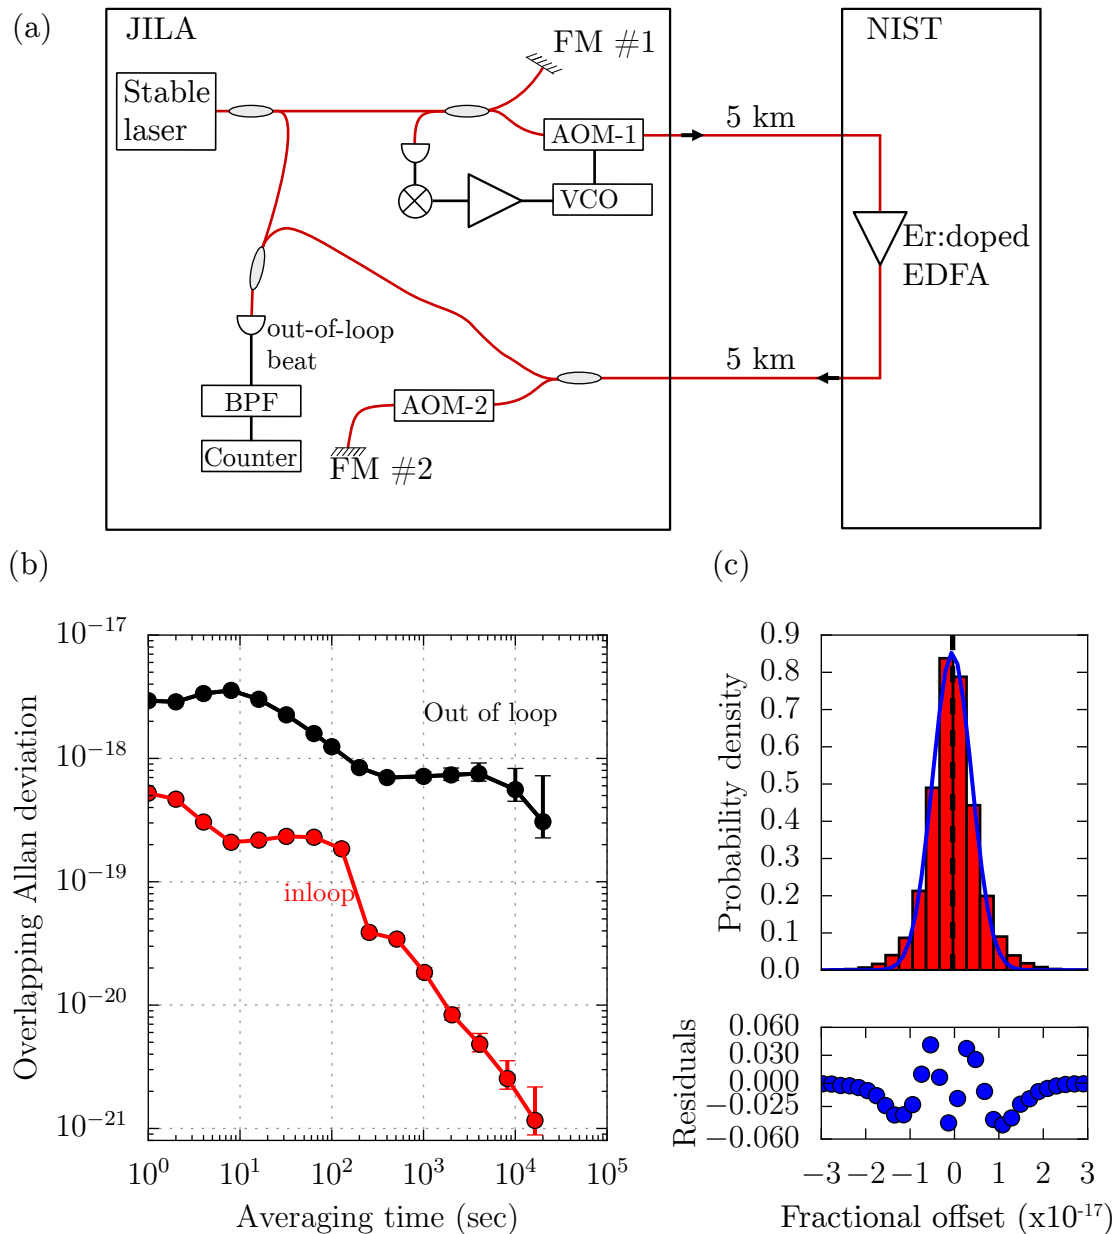


Figure 3.6: (a) Schematic for the out-of-loop evaluation of the BRAN fiber link for the 2017/2018 JILA-NIST clock comparison. A stable laser (Si3) is fiber noise cancelled, with the noise detected between the mirrors labelled as “ref #1” and “ref #2”. The out-of-loop beat is measured between the local light and the “remote” light. (b) Overlapping Allan deviation of the inloop error for the fiber noise cancellation (red) and the out-of-loop beat (black points). (c) Histogram of the out of loop frequency measurements, with the residuals to the Gaussian fit shown below.

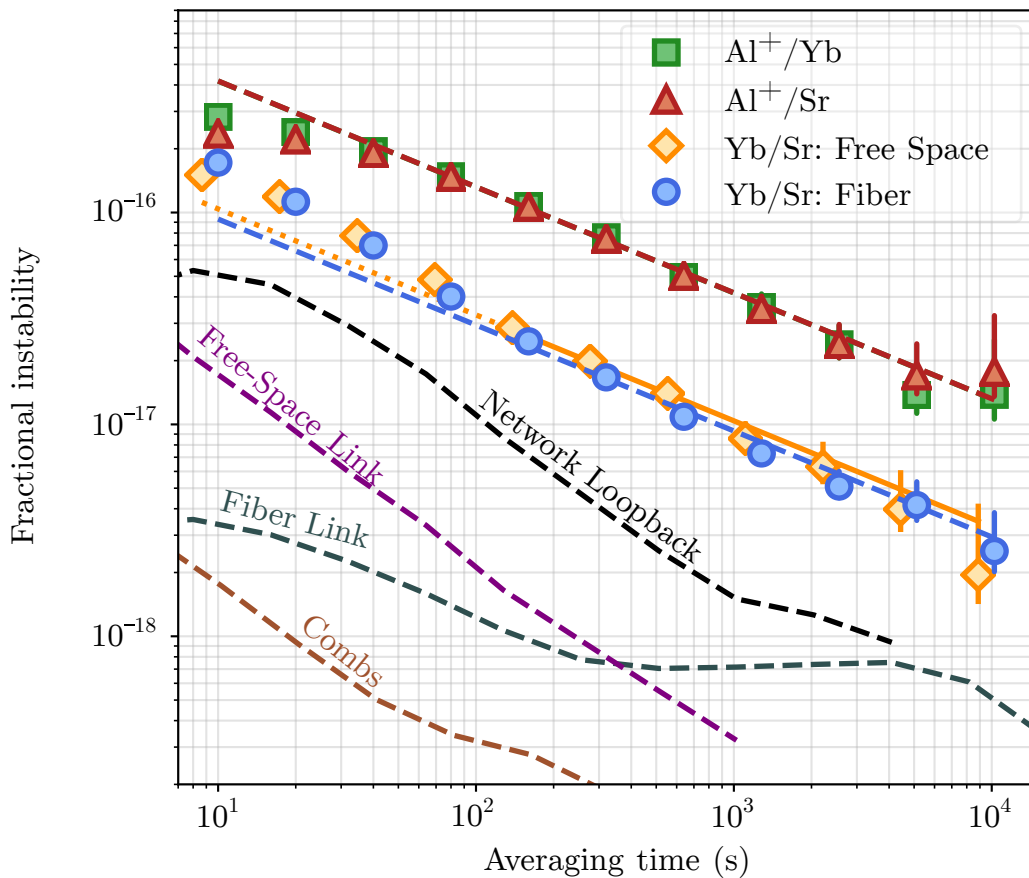


Figure 3.7: Allan deviation of the three ratios on a given day. Also shown are noise traces for various components of the optical network, including the frequency combs, fiber link, free-space link and a network loopback measurement. We can see good agreement between the free space and fiber measurements of the Yb/Sr ratio.

Chapter 4

Spin-squeezed clock : background and considerations

4.1 Basic theory and definitions

Neutral atom clocks that employ uncorrelated independent atoms are fundamentally limited by QPN. We describe the state of this atomic ensemble using the collective spin operator $\hat{\mathbf{J}} = \hat{J}_x \hat{x} + \hat{J}_y \hat{y} + \hat{J}_z \hat{z}$, which is formed by simply taking the sum over all of the individual spin vectors. The $\hat{x}, \hat{y}, \hat{z}$ are unit vectors that form the basis for the collective state vector. The z -projection of the collective spin operator can be expressed as population differences between the two spin states as

$$J_z = \frac{N_\downarrow - N_\uparrow}{2}, \quad (4.1)$$

where N_\downarrow and N_\uparrow are the atoms in the ground and excited clock states, respectively. We will concern ourselves with J_z , because it contains our clock discriminator. Ramsey spectroscopy converts phase accumulation in the azimuthal plane to a tipping of the Bloch vector along the z direction on the Bloch sphere (J_z).

Coherent spin states (CSS) describe the non-classical state formed by N identical spin-1/2 particles. We can express the CSS as a product state [24, 107]

$$|\theta, \phi\rangle = \prod_{i=1}^N \cos \frac{\theta}{2} |\uparrow_i\rangle + e^{i\theta} \sin \frac{\theta}{2} |\downarrow_i\rangle \quad (4.2)$$

This state is represented by a simple arrow with the endpoint living on the Bloch sphere. It responds to optical rotations in a classical fashion. We will consider a state pointing along the y direction on the Bloch sphere after atoms are initially prepared in a coherent superposition of $|\downarrow\rangle$ and $|\uparrow\rangle$.

The noise associated with this state is along the orthogonal directions z and x , and is obtained by applying Heisenberg’s uncertainty principle,

$$\Delta J_z \Delta J_x \geq |J_y|/2 \quad (4.3)$$

where Δ denotes the standard deviation. We visualize the noise of this state as a Gaussian distribution on the surface of the Bloch sphere. This noise arises from quantum projection, which is the noise associated with the projective measurement of a state placed into a superposition.

The ultimate phase resolution potentially achievable by a CSS is set by the QPN and is defined as the standard quantum limit. We express QPN in terms of the uncertainty of the z -projection of the Bloch vector

$$\Delta J_{z,CSS} = \sqrt{N}/2 \quad (4.4)$$

QPN arises from the projective measurement of an ensemble in a superposition state, when the atoms “collapse” into either the $|\downarrow\rangle$ or $|\uparrow\rangle$ state. This QPN will lead to a limit in the phase-estimation of a Bloch vector with length N to

$$\Delta\theta_{SQL} = \frac{\sqrt{N}/2}{N/2} = \frac{1}{\sqrt{N}} \quad (4.5)$$

We note that we have assumed the contrast associated with the state is $C_i = 1$, and we treat the case of $C_i < 1$ in Section 7.1. Equation 4.5 makes it very clear that operating with a higher atom number is advantageous to push the SQL to the lowest phase resolution possible. It is difficult to run a clock that operates at the SQL, and even the highest performing synchronous clock comparisons tend to operate a few dB above this level, often due to the presence of other technical noise in the clock [17, 19].

The idea of spin-squeezing is to engineer entanglement in the atomic ensemble in order to surpass the SQL. Since the clock signal is captured in J_z , the simple idea is to shuffle the uncertainty to the anti-squeezed quadrature J_x , and improve the phase resolution while still satisfying Equation 4.3. Experiments that involve spin-squeezing greatly increase the technical complexity of the optical clock. One has to engineer interactions between the atoms in order to obtain non-

classical spin states. When introducing these interactions, one must preserve the coherence of the atomic ensemble.

4.2 Historical background

The development of quantum entanglement has provided an exciting new direction for reducing the impact of QPN in quantum sensors, offering the opportunity to greatly advance upon this state-of-the-art performance. Spin squeezing, in particular, was proposed early on to utilize quantum correlation to conceal noise from individual atoms and thus achieve improved measurement precision and bandwidth [108, 109]. The creation of entanglement for metrology has been explored in a wide variety of atomic quantum sensors including microwave clocks [110, 111, 112, 30, 31, 34, 33, 29], ion clocks [113, 114], magnetometers [115], and matterwave interferometers [116].

The first step toward implementing spin-squeezing in a cavity QED system is achieving strong coupling between the atom (or atoms) and the electromagnetic cavity mode. This strong coupling regime is also referred to the resolved vacuum Rabi splitting limit, where the vacuum Rabi frequency,

$$\Omega = 2g\sqrt{N}, \tag{4.6}$$

is much larger than any other time scales in the system ($\Omega \gg \kappa, \Gamma$). There are two general regimes for strong coupling: microwave domain [117, 118], and optical domain [119]. In the earliest work, both of these general approaches were pursued by coupling either single atoms, or beams of atoms to the microwave/optical cavity. The approach in the Haroche and Walther group was to couple a beam of circular Rydberg atoms to a superconducting microwave cavity. In that case, the dipole moment of the atomic transition was enhanced by using Rydberg atoms, and a very high quality factor microwave cavity minimized κ . This enabled the realization of the strong-coupling regime for the atomic beam passing through the microwave cavity. The second approach was pioneered in the group of Jeff Kimble, where alkali atoms were coupled to a high finesse optical cavity via dipole-allowed transitions. In this regime, strong-coupling can be attained for just single atoms

[119].

In later work in the groups of Vladan Vuletic, James Thompson, Mark Kasevich, and others, ensembles of atoms were trapped and coupled to optical cavities, enabling the inspiring work on implementing spin-squeezing in microwave clocks [30, 31, 33, 29]. By coupling ensembles of many atoms, one can enhance the collective coupling between the atoms and the cavity mode. We will follow along this line, but instead run the clock in the optical domain, both for the cavity QED coupling and the clock states.

Despite tremendous progress, spin-squeezing in atomic clocks has yet to yield enhancement at state-of-the-art stability levels. A spin-squeezed microwave clock has observed 11 dB of enhancement [33] at the 10^{-10} stability level at 1 s, in contrast with microwave fountain clocks at 10^{-14} [120]. We note that the spin-squeezed clock in [33] operated at a considerably shorter Ramsey dark time than the state-of-the-art atomic fountain clocks. For optical clocks, which operate at much higher stability, generation of entanglement has been demonstrated by a measured Wineland parameter of -4.6 dB [35]. After subtraction of a laser noise model, an optical clock employing a SSS was inferred to operate -4.4 dB below the SQL at a fractional frequency stability of $1.3 \times 10^{-13} \tau^{-1/2}$ (where τ is the averaging time in seconds) [35]. However, metrological applications require direct observation of clock stability enhancement without post-processed removal of technical noise. It remains to be seen whether spin-squeezing will enhance the operation of state-of-the-art atomic clocks. The primary objection of the Sr3 machine is to demonstrate a competitive atomic clock with cavity QED based squeezing, and directly observe metrological enhancement.

4.3 How do we design a cQED system?

What should we consider when designing a cavity QED system for an optical clock? To answer this, we will examine some simple calculations of various physical quantities, and discuss how different design choices impact various parameters through fundamental and technical scalings.

The Jaynes-Cummings Hamiltonian describes a single atom interacting with the quantized

electromagnetic field of the cavity mode,

$$\hat{H} = \hbar\omega_a\hat{\sigma}_+\hat{\sigma}_- + \hbar\omega\hat{a}^\dagger\hat{a} + \hbar g\left(\hat{\sigma}_-\hat{a}^\dagger + \hat{\sigma}_+\hat{a}\right), \quad (4.7)$$

where the first term is the energy of the atom and ω_a is the atomic transition frequency between the energy levels $^1S_0 \rightarrow ^3P_1 = |e\rangle$. The second term is the energy of the photons, where \hat{a} (\hat{a}^\dagger) is the photon annihilation (creation) operator, so $\hat{a}\hat{a}^\dagger$ corresponds to the total photon number. ω is the optical cavity resonance frequency. The last term describes the interaction between the photon and the atom. $\hat{\sigma}_-\hat{a}^\dagger$ represents an atom going from the excited to ground state and emitting a photon, while $\hat{\sigma}_+\hat{a}$ describes an atom absorbing a photon and going from the ground to excited state.

As usual, we go into the rotating frame of the atom,

$$\hat{H} = \hbar\delta_c\hat{a}^\dagger\hat{a} + \hbar g\left(\hat{\sigma}_-\hat{a}^\dagger + \hat{\sigma}_+\hat{a}\right), \quad (4.8)$$

where $\delta_c = \omega - \omega_a$ is the relative detuning of the cavity from the atomic transition. A key parameter here is the single atom cavity coupling [121],

$$g = \mu\sqrt{\frac{\omega_a}{2\hbar\epsilon_0 V}}, \quad (4.9)$$

where μ is the atomic dipole matrix element, ω_a is the atomic transition frequency, and V is the cavity mode volume. The atomic dipole moment μ is

$$\mu = \left(\frac{3\epsilon_0\hbar\lambda^3}{8\pi^2}\Gamma\right)^{1/2} \quad (4.10)$$

We note that g only depends on properties of the atomic transition and the volume of the cavity mode - which will depend on the cavity parameters.

To extend this Hamiltonian to describe N atoms interacting with the cavity mode, we replace the atomic operators with collective raising and lowering operators $\hat{J}_\pm = \frac{1}{2}\sum_{i=1}^N\hat{\sigma}_\pm^i$. The Tavis-Cummings Hamiltonian that describes the many-atom ensemble coupled to a quantized mode of electromagnetic radiation is then

$$\hat{H} = \hbar\delta_c\hat{a}^\dagger\hat{a} + \hbar g\left(\hat{J}_-\hat{a}^\dagger + \hat{J}_+\hat{a}\right). \quad (4.11)$$

Assuming we are in the weak-excitation limit, where the number of atoms excited to $|e\rangle$, then we can make an approximation to this Hamiltonian [32, 122],

$$\hat{H} = \hbar\delta_c\hat{a}^\dagger\hat{a} + \hbar g\sqrt{N}\left(\hat{b}\hat{a}^\dagger + \hat{b}^\dagger\hat{a}\right). \quad (4.12)$$

where the collective operators \hat{J}_\pm are replaced with new creation and annihilation operators $\hat{b} = J_+/\sqrt{N}$ and $\hat{b}^\dagger = J_-/\sqrt{N}$. We have finally arrived at the enhancement of atom-cavity coupling by using an ensemble of N atoms, increasing the frequency by a factor of \sqrt{N} .

One might think that choosing a transition with the largest dipole moment μ leading to a larger g will give the maximum achievable squeezing. This is not true, however, and instead the key parameter that dictates the maximum achievable squeezing through QND measurements is the single-atom cooperativity \mathcal{C} , expressed as

$$\mathcal{C} = \frac{(2g)^2}{\kappa\Gamma}. \quad (4.13)$$

Examining this expression, we can see that the cooperativity \mathcal{C} does not depend on Γ (or the strength of the atomic transition dipole moment). We can understand this intuitively: a transition with a larger dipole moment will increase the numerator of Equation 4.13, but that increased dipole moment simultaneously increases the rate of incoherent scattering of photons into free space. The spirit of Equation 4.13 is that what matters is the tradeoff between photons leaving through the cavity mirror as useful information and the leakage into the universe via free space scattering. This means that one can either use a broad transition ($^1S_0 \rightarrow ^1P_1$, $\Gamma = 2\pi \times 35$ MHz), or a narrower transition ($^1S_0 \rightarrow ^3P_1$, $\Gamma = 2\pi \times 7.5$ kHz). Reasons for choosing one transition over another are thus more technical rather than fundamental. For example, working high finesse cavities at shorter wavelengths can be problematic in terms of mirror coating damage [123].

The spectroscopic enhancement using QND measurements scale with the spin noise reduction and the shortening of the Bloch vector (loss of Ramsey fringe contrast). We assume that the contrast loss is limited by free space scattering, so we can write the spectroscopic enhancement as

$$\xi = \left(\frac{\Delta J'_z}{\Delta J_{z,QPN}}\right)^2 e^{-m_s} \quad (4.14)$$

where m_s is the number of scattered photons and $\Delta J'_z$ is noise of a differential J_z measurement. In the far-detuned limit where the atom-cavity detuning is much greater than the vacuum Rabi splitting $\delta_c \gg \Omega$ and assuming negligible readout noise of the final measurement, and assuming the loss of coherence is dominated by free-space scattering, the potential spectroscopic enhancement scales as [32]

$$\xi^{-1} = \frac{QNC}{e}. \quad (4.15)$$

This means that the maximum achievable squeezing only depends on the collective cooperativity NC and the total quantum efficiency Q . This result is independent of the atomic transition linewidth Γ .

One important assumption of the above equation is that we are working with a cycling transition. This is more complicated when working with the fermionic isotope of Sr, where we need to consider the complicated hyperfine levels. Transitions that take the atom out of the initial ground-state will lead to loss of coherence and an increase of noise via diffusion [32]. For these reasons, we choose to operate on the $|^1S_0, F = 9/2, m_F = 9/2\rangle \rightarrow |^3P_1, F = 11/2, m_F = 11/2\rangle$. We apply a bias magnetic field along the direction of the cavity, and probe the cavity with circularly polarized light. Barring any impure polarization or relative misalignment of the magnetic field and the k-vector, we should realize a cycling transition with relatively small probability of a spin flip.

4.4 How should we probe the atom-cavity system?

We want to probe the atom-cavity system such that we optimize the information gained about the collective atomic state versus the incoherent collapse of the single-particle wave function. This means we need the most precise information about the state of light leaving the optical cavity.

In QND measurements, we aim to utilize the phase of the light leaking from the cavity as the witness for the ground state population of the atomic ensemble in the cavity. The leakage light that comes out of the cavity through the input mirror can be written relative to the incident field

E_i , [32]

$$\frac{E_r}{E_i} = \mathcal{I} + i\mathcal{Q}, \quad (4.16)$$

where the \mathcal{I} quadrature response is

$$\mathcal{I} = \left(1 - \frac{\beta_r}{1 + \left(\frac{\delta}{\kappa_p/2}\right)^2} \right) \quad (4.17)$$

and the \mathcal{Q} quadrature is

$$\mathcal{Q} = \frac{\beta_r \left(\frac{\delta}{\kappa_p/2}\right)}{1 + \left(\frac{\delta}{\kappa_p/2}\right)^2} \quad (4.18)$$

where $\beta_r = \sqrt{\kappa_1/\kappa}$ is the radius of the $\mathcal{I} - \mathcal{Q}$ circle for the light coming back through the input mirror. Here, κ_1 is the rate of coupling out of the input mirror, κ_2 is the rate of coupling out of the second mirror, $\kappa = \kappa_1 + \kappa_2 + \kappa_L$ is the total rate of power decay for the cavity mode. The detuning between the probe laser and the cavity mode is δ . To show a direct example of this, we measure the $\mathcal{I} - \mathcal{Q}$ response of our cavity with the setup in Figure 4.1. We scan our probe laser across the cavity resonance and record a heterodyne beat between the probe and the LO beam. By directly digitizing this signal, we can demodulate the RF signal in software. The normalized $\mathcal{I} - \mathcal{Q}$ response is plotted parametrically in Figure 4.1b. The individual \mathcal{I} and \mathcal{Q} responses are plotted versus the probe laser frequency in Figure 4.1b and c. The red line in each panel is the result of a simultaneous fit to Eq. 4.17 and Eq. 4.18. We can see some slight deviation from the theoretical response, as seen in the baseline drift in the \mathcal{I} quadrature. This can come from a slight miscalibration of the demodulation waveform, and does not significantly impact the extraction of the cavity frequency, as we can restrict the fitting window to the most sensitive parts of the frequency sweep (closer to resonance).

The QND measurement can be distilled down to a measurement of the leakage light to extract the phase. It is this light that carries information about the atomic state inside the cavity. However, in cavity QED experiments, this probe power must be kept at an extraordinarily small level, in our case around 10^{-15} W. Since it is such a small amount of probe power, we opt to measure the phase shift against a stronger local oscillator beam, in order to make the photon shot noise dominant

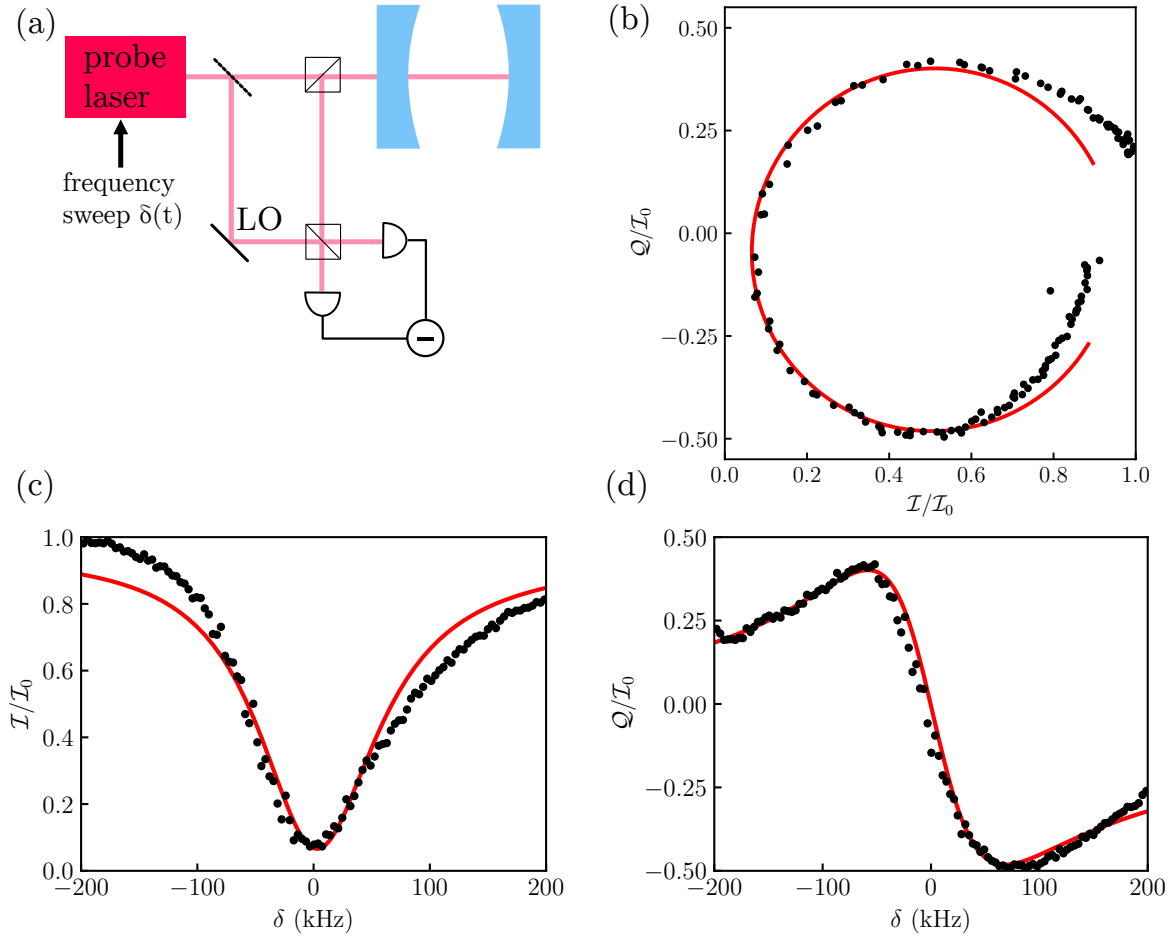


Figure 4.1: (a) Schematic of the detection scheme, where laser light is split into a probe path and LO path, and these are combined on a beamsplitter. (b) By sweeping the probe laser across resonance we can map out the $\mathcal{I} - \mathcal{Q}$ response of the cavity response. (c) Normalized \mathcal{I} quadrature. (d) Normalized \mathcal{Q} quadrature.

over the detector noise, giving a shot noise limited signal-to-noise. The amplitude of the resulting interference term is proportional to the product of the electric fields, or equivalently the square root of the optical powers. For sufficient LO power, the noise of the photodetector and measurement system is negligible compared to the LO photon shot noise. In this case, the signal-to-noise ratio is simply proportional to the signal beam,

$$SNR \propto \frac{\sqrt{P_s P_{LO}}}{\sqrt{P_{LO}}} = \sqrt{P_s} \quad (4.19)$$

To achieve the photon shot noise limit, we also have to be careful about classical intensity noise of the LO, which is more severe at greater LO powers. We combat this by employing balanced detection in addition to active control of the LO intensity.

While this simple scaling gives good intuition on quantum optical measurements, there is the choice between “heterodyne” and “homodyne” detection schemes. This refers to the fact that one can tune the LO beam to be at a different frequency from the cavity probe and look at the AC signal, or tune to the same frequency and extract the information at DC. The heterodyne method can be simpler to implement, but suffers from a loss of 3 dB of signal to noise ratio. This is seen by considering the influence of the “image frequencies”. When you form a heterodyne beatnote, you are looking at the interference of two distinct optical frequencies, which produce a beatnote at the difference frequency $f_{het} = \nu_{LO} - \nu_s$, where ν_{LO} is the optical frequency of the LO and signal respectively. However, on the other side of the LO is not a tone, but shotnoise, that gets mixed to the same beat frequency f_{het} . Hence the shot noise contribution gets doubled, resulting in the loss of 3 dB in signal to noise. The most straightforward way around the 3 dB loss is by performing homodyne detection. There, the LO is tuned to the same frequency as the signal beam. By adjusting the relative phase of the LO beam with respect to the signal beam, one can choose the particular quadrature of detection of I or Q . Making the signal beam and the LO beam to be the same frequency, there is no image shot noise, and hence no extra loss of 3 dB of signal to noise.

While homodyne detection optimizes the signal-to-noise, it also allows for further optimization by performing “fixed tone measurements”, which refers to fixing a probe tone on the resonance

of the cavity mode as opposed to sweeping the laser frequency across the resonance. For example, if you sit exactly on resonance with heterodyne, the signal will spend half of the time in the \mathcal{Q} quadrature, and thus half of the probe will not carry useful information. However, if we perform homodyne detection, then the entire probe field will carry the most sensitive phase information. It's also clear that sweeping will not provide the optimal measurement. This is because the optimal measurement quadrature varies from the \mathcal{I} to \mathcal{Q} and back the \mathcal{I} again during the sweep. So unless one varies the detection phase during the sweep, one loses the 3 dB gain.

With all of these considerations - it is clear that the best QND measurement will be made by performing homodyne detection and probing with a fixed frequency probe on resonance. This is the primary motivation for operating our system in a detuned regime. In this regime, we can park a probe tone near the shifted cavity-like mode, and the atom number fluctuations are small enough such that we stay in the linear response of the homodyne signal. The one con of this approach is that we lose the common mode suppression of laser noise that one enjoys by probing the vacuum Rabi splitting. We combat this by building a low noise cavity stabilization setup, discussed in Chapter 5.8.1.

4.4.1 Atomic motion and probing time scales

In the experiments described in this Thesis, we keep all intra-cavity light to a minimum, and thus do not confine the atoms along the cavity axis. The atomic motion is dictated by their temperature and the trapping potential of the vertical trapping optical lattice. At higher order, the optomechanical effects from the probing light may play a role [34], but for now we will ignore this effect. We treat each pancake of the vertical optical lattice as a 2D harmonic oscillator. The atoms occupy the radial modes of the 2D harmonic oscillator according to the corresponding Boltzmann distribution. To estimate the order of magnitude for the spatial spread, the thermal radius for our cloud can be estimated from $\frac{1}{2}k_B T = \frac{1}{2}m_{87}(2\pi\nu_R)^2\sigma_y^2$, where m_{87} is the mass of ^{87}Sr , ν_R is the radial trap frequency. From this, we get a thermal cloud radius of $\sigma_y = 26 \mu\text{m}$. We note that this definition is consistent with our density distribution defined in Appendix B. The cavity mode gives

rise to the spatial dependence of the atom-cavity coupling

$$g_i^2(X, Y, Z) = g_0^2 e^{-(Y^2+Z^2)/w_0^2} \sin^2(2\pi X/\lambda_{689}) \quad (4.20)$$

So as the atoms move, they will experience a time-dependent atom-cavity coupling, $g(X(t), Y(t), Z(t))$.

In the far-detuned regime ($\delta_c \gg \Omega$), the time-dependent frequency shift for the total atomic ensemble can be decomposed into the dispersive shifts of each of the individual atoms,

$$\delta\nu_i(t) = \frac{g_i^2(t)}{\delta_c}. \quad (4.21)$$

We seek to simulate the atomic motion of atoms in the 2D harmonic oscillator, and use those trajectories to calculate $g_i^2(t)$ for each i^{th} atom. We can then calculate the mean cavity shift due to each atom, and simulate how the cavity shift depends on probing time scales or the atomic temperature.

Simulating the atomic trajectories is carried out as follows. We randomly sample initial X and Y positions for each atoms according to the normal distribution with $\sigma_Y = 26 \mu\text{m}$. With the spatial coordinate, we then pull from Boltzmann distributions to assign the initial momentum for p_x and p_y . The distributions along the two directions are taken to be independent, each having a standard deviation of $\sigma_p = \sqrt{2m_{87}k_B T_r}$. We then iteratively solve Hamilton's equation of motion to obtain new values of position and momentum. These trajectories are plotted in Figure 4.2a as the white traces. They are plotted on top of the intensity distribution of the standing wave inside the cavity. The solid lines indicated the $1/e^2$ radius of the cavity mode, $w_0 = 71 \mu\text{m}$.

For each atom, we track the motion and plot the quantity $g_i^2(t)/((2\pi)^2\delta_c)$, giving units of Hz/atom. We can see that after a very short amount of time, most of the atoms traverse several standing waves of the probing light. This is consistent with a rough estimate made by considering a 1D Boltzmann distribution and calculating the fraction of atoms that traverse 5 standing waves to be around 0.1 ms. Thus, we can be rest assured that sufficient thermal averaging along the cavity mode is achieved after a relatively short amount of probing time.

However, the story for the radial motion is not as straightforward. This is because the thermal cloud radius $\sigma_Y = 26 \mu\text{m}$ is not that much smaller than the cavity mode waist of $w_0 = 71 \mu\text{m}$.

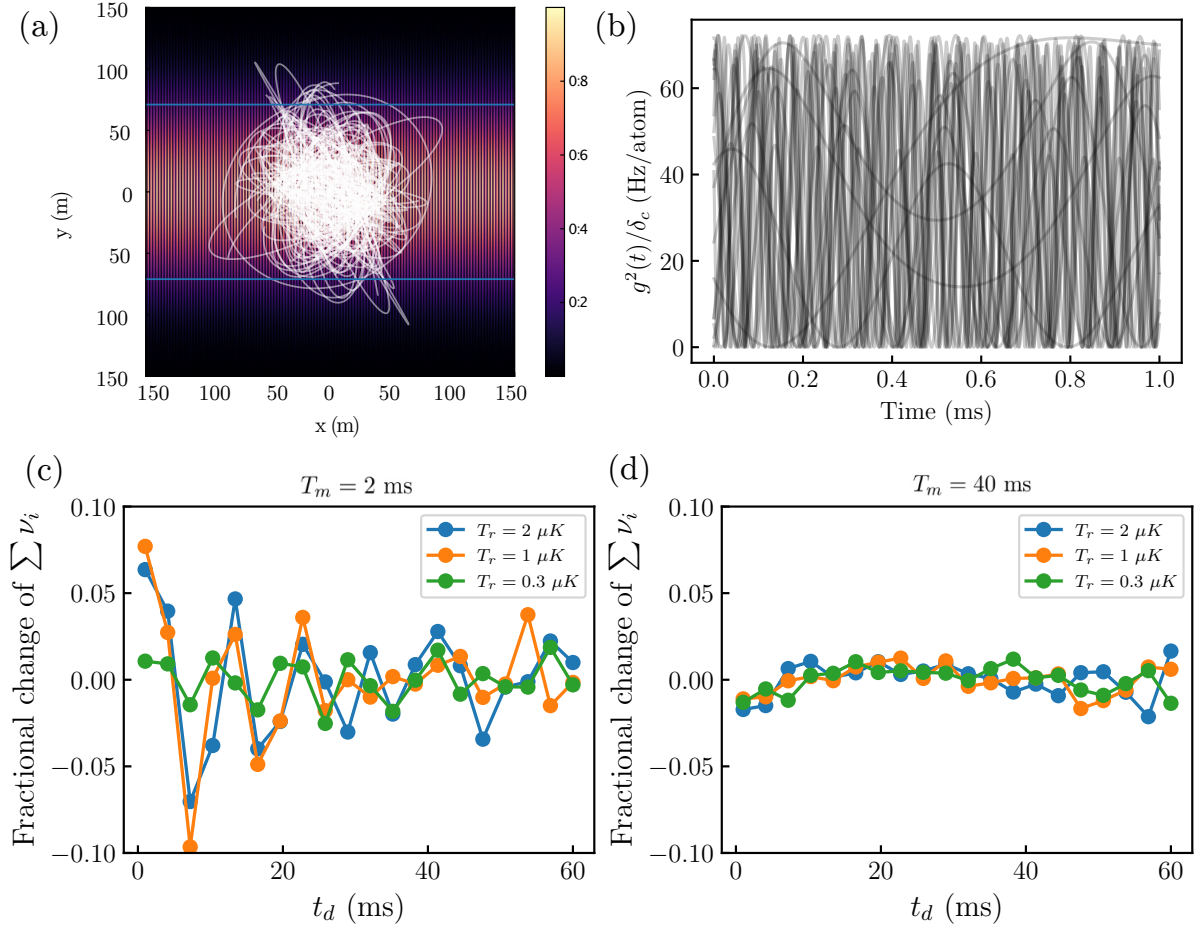


Figure 4.2: (a) The atom-cavity coupling $g^2(x, y)$, which follows the intensity of the cavity mode standing wave is shown as the contour plot, normalized to a peak value of one. The lines indicate the $1/e^2$ radius. Overlaid on top of g^2 is 100 different atomic trajectories simulated for the particle motion in a single pancake of the optical lattice. (b) Corresponding traces of $g^2(t)/\delta_c$ for all of the different particles, showing that most atoms sample many standing waves well within 1 ms. (c) Fractional noise of the cavity shift for a measurement time of 2 ms, versus the dwell time between frequency measurements. We show the magnitude of the oscillations versus atomic temperature (d) Same as (c), with a measurement time of 40 ms. We can see that the longer measurement time offers improved averaging of the radial motion.

This means that the radial motion will lead to significant changes in the atom-cavity coupling. We consider the impact of radial motion on a two window measurement of the cavity frequency shift, similar to what we implement for our QND probing of J_z in Chapter 6 and 7. In Figure 4.2c, we plot the total fractional shift versus the dwell time between the two measurements t_d for a short measurement window of $T_m = 2$ ms. We can see the radial oscillations are damped out as the radial motion is not perfectly coherent and there exists a spread of radial trapping frequencies. If we set the measurement time to $T_m = 40$ ms (Figure 4.2d), we can see that there is further averaging of the radial motion giving a drastically smaller effect.

Now we want to simulate the magnitude of noise in dressed atom-cavity frequency measurements that can occur due to the atomic motion. We fix the dwell time $t_d = 10$ ms, chosen to represent the time it takes to perform a π -pulse. We then vary the measurement time, and look at the standard deviation of the fractional cavity shift difference $(\nu_p - \nu_f)/\nu_p$ (Figure 4.3). Further, by varying the atomic temperature, we can see that lower temperature atoms will give smaller noise due to the atomic motion - matching intuition. We can also see that a measurement time of 20 ms or longer is a safe choice for mitigating motional effects. As we will discuss further, experimentally we use a measurement time of $T_m = 40$ ms. This choice is primarily influenced by considering two effects: the atomic motion discussed in this Section, and the influence of cavity - laser technical noise, discussed in Chapter 5.




Figure 4.3: Simulations of cavity frequency shift noise arising from thermal atomic motion. With a fixed separation time between measurements of 10 ms, we predict fluctuations in the shifted cavity frequency on the percent level. Both low atomic temperatures and longer measurement time T_m are critical for minimizing this effect.

Chapter 5

Sr3 : a new spin-squeezed clock apparatus

In this Chapter, we describe the design and construction of a new experimental apparatus that integrates a high finesse cavity with a state-of-the-art optical lattice clock. We have named this system as “Sr3”, since it is the third atomic clock system being operated in the Ye lab. The first goal of this experiment was to directly observe improvement in clock performance due to spin squeezing. To achieve this, we require the preparation of 1,000 - 10,000 strontium atoms inside a high finesse cavity and in a well controlled environment. The high finesse optical cavity and supporting structure is designed for low noise frequency measurements, allowing us to make non-destructive measurements of the atomic population with resolution below the quantum projection noise. We also implement a movable optical lattice, in order to apply squeezing and read out independent atomic ensembles with the cavity probe. With all of this in mind, we took one step further and designed a system that would enable future performance at the state-of-the-art level for accuracy, including black-body radiation, DC Stark, ac Stark, and other systematic shifts.

We split the experiment into two main regions inside our vacuum chamber, referred to as the “MOT bucket” and the “science cavity”. The MOT bucket sits below the cavity, and serves as the location where we prepare the initial blue and red MOTs. These atoms are loaded into the movable optical lattice, which transports the atoms ≈ 40 mm vertically into the science cavity. The cavity enables us to engineer squeezed states of several thousand atoms, while simultaneously providing a clean environment in terms of clock systematics. The following will describe the design of the ultrahigh vacuum (UHV) system, the MOT bucket, the science cavity, and some of the surrounding

hardware.

5.1 UHV System

One of the requirements for a cavity-QED optical clock system is achieving UHV. A poor vacuum would impact our experiment in a number of different ways. Broadly speaking, a higher background pressure means more collisions of these gas molecules with Sr, leading to loss, or in the case of a grazing collision, leading to a systematic clock frequency shift. Thinking practically about the impact of a finite vacuum lifetime for Sr3, we consider a potential future experiment where we will use the QND nature of the probing to recycle the atomic ensemble instead of reloading atoms each cycle of the experiment. This type of experiment would have the advantage of reducing the dead time of the clock, leading to improved stability due to mitigation of the Dick effect. Of course, a short vacuum lifetime would be a limitation to such an experiment, due to the eventual loss of atoms and the corresponding rise of the QPN. When designing Sr3, we knew we would be placing a lot of different materials in vacuum, including: the copper cavity spacer and supporting structure, mirrors, epoxy, temperature sensors, kapton wires, viton. Thus, we could not expect the same vacuum quality as the Sr2 machine described in [124]. We ended up setting the goal of $< 10^{-10}$ Torr, which should lead to a > 10 s vacuum-limited lifetime.

Each component that was placed in vacuum have different levels of handling, and different cleaning procedures altogether. In the following we describe the UHV system, the pre-bake and cleaning procedures we implemented, and finally the bakeout of the entire system.

5.1.1 Design

We employ an AOSense strontium cold beam source, which includes the oven crucible, Zeeman slower and two 2D MOTS. After transporting the chamber from the auxillary lab to the main Sr room, we found that one of the heater pins for the oven was shorted to the vacuum chamber. After some moments of anxiety, we realized we could actually run the heater in this configuration. The one worry is that if we flow current directly into the vacuum chamber, it could find its way back to

Earth ground through an undesirable path (for example through the breadboard \rightarrow photodiode). We thus attached a thick tin-coated copper braid (the type typically used for electronic cable shielding) to serve as the return for the current, and by checking with a clamp current reader, we confirmed the current was making its way back to the power supply through this attached braid.

The oven is attached to the chamber via a long skinny vacuum tube. The differential pumping speed from the outlet of the all-metal gate valve to our main chamber is calculated to be 0.01 L/s. For vacuum pumping, we employ a 150 L/s Starcell ion pump from Varian, a titanium sublimation pump (TSP), and a 200 L/s NEG. The TSP gives us the largest pumping speed, in excess of 1000 L/s for hydrogen in the chimney area of the chamber. We note that when one calculates capacity for the TSP, one would expect the pump to saturate on the year timescale at 10^{-11} Torr level. However, we have not noticed a rise in pressure from the ion gauge, and so we have not refired the TSP. Interestingly, the story is the same with Sr2. Since the TSP does not provide efficient pumping of inert gases (argon, xenon, etc), we add the ion pump. This is especially important since anytime we replace the oven, one dominant gas will be argon from the back-filling process.

5.1.2 Assembly and bake-out procedure

Given careful cleaning of the in-vacuum components, our chamber should mostly be limited by the level of hydrogen outgassing of the chamber walls and the achieved pumping speed. Based on an initial 150 C degree vacuum bakeout of the empty chamber (all steel blanks, no payload), we estimated our initial steel outgassing rate to be 3×10^{-12} mbar L s⁻¹ cm⁻². It has been shown that vacuum baking steel vacuum chambers at a temperature greater than 400 C for a long period of time can reduce the hydrogen outgassing rate by as much as a factor of ten [125]. We removed all viewports and vacuum pumps except for the turbo-molecular pump (TMP), and proceeded to bake the chamber at 400-450 C for 25 days. At this elevated temperature, we observed the partial pressure of hydrogen to decrease from 1.5×10^{-5} Torr down to 3×10^{-6} Torr, indicating an overall reduction of the hydrogen outgassing rate of nearly a factor of three.

With the steel in the main chamber thoroughly baked, next we needed to ensure that all

objects placed in vacuum were properly cleaned. The surface area of the MOT bucket is not negligible, and all of the inner walls that are exposed to vacuum are steel. Since these parts had been fabricated in the shop, we opted to sonicate the MOT bucket upside down in an Alconox water bath for 30 minutes, followed by a solvent rinse. Since the MOT bucket has several in-vacuum welds, we combined a vacuum test with a pre bake of the MOT bucket. The MOT bucket was attached to a brand new 10 inch diameter nipple, and was subsequently vacuum baked at 150C for 7 days. The final pressure reading of this MOT bucket prebake reached the low 10^{-10} Torr level, limited by the RGA noise floor.

Our science cavity design employs viton for vibration decoupling. The LIGO collaboration has found that viton can outgas a significant amount of hydrocarbons under vacuum, especially above a temperature of 180° C. Since this viton is relatively close to our high finesse mirrors, we took extra care in preparing these pieces for our system. For vacuum cleaning the viton, we followed the LIGO cleaning procedures with some modifications :

1. Ultrasonic bath in Alconox for 15 minutes
2. Ultrasonic bath in de-ionized water for 15 minutes
3. Vacuum bake at 180 C for 72 hours

The cavity pieces that are machined out of copper were cleaned ultrasonically with Alconox, and then in the “usual three solvents” of isoproponal, acetone and methanol. In order to ensure cleanliness and to increase the emissivity of the copper, we then air-baked at 150C for approximately 3 hours. This air-baking made the cavity spacer a reddish-brown color - indicating the formation of a Cuprous Oxide coating, which has a higher emissivity than polished copper. This will increase the relative contribution of the copper to the black-body radiation environment of the atoms, thus giving higher control and knowledge of the thermal environment. The final assembly procedure was an involved effort of stringing various kapton wires through the cavity assembly. The electrical feedthrough we used is part number : IFTRG327013, which uses a MIL-C-26482 style connector with 32 total pins. Of these 32 pins, we use 4 for the two PZTs, and 10 total for the in-vacuum thermal sensors.

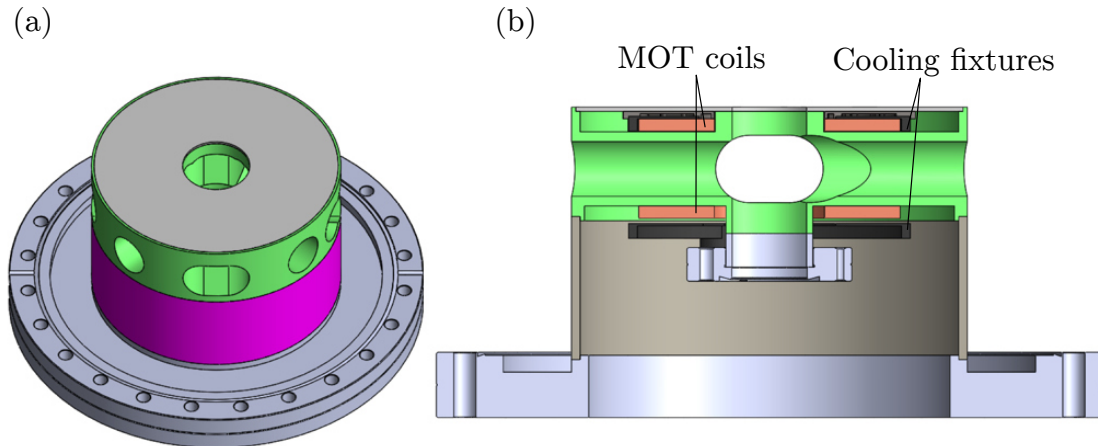


Figure 5.1: (a) Angle-view of the MOT bucket. The smaller MOT bucket pieces are welded together on a 10 inch conflat flange that is bolted to the bottom of the chamber. (b) Side cut-view of the MOT bucket. The MOT coils and the cooling fixtures inside the MOT bucket are labelled. In total, there are 5 in-vacuum welds.

5.2 MOT bucket design

The MOT bucket was designed to provide a space inside the vacuum chamber to produce the ultracold atomic sample for cavity QED. The MOT coils are placed 1.1 inches apart with an overall diameter of 2.5 inches. We used solid-core, kapton coated 13 AWG wire from S&W wire company, with a specified operating temperature up to 250 C. The MOT coils consist of 13 turns of wire in two vertical layers, potted in a high-temperature non-ferromagnetic epoxy. This wire configuration was chosen to make the coils compact, while simultaneously maximizing the surface area contact with the watercooling fixtures that are affixed to one side of each coil (see Figure 5.1b). Four roughly calibrated sensors are installed inside the cooling fixtures, so in the future the thermal environment can be monitored for accuracy studies.

Installing magnetic field coils with so much surrounding metal, especially the copper cooling flanges can be worrisome in terms of generating significant Eddy currents. We measured the responding magnetic fields inside the MOT bucket when we rapidly switch off the MOT coils, and

confirmed that we could switch the magnetic fields on the ms level. The presence of the copper cavity further complicates the issue of Eddy currents, which we discuss and present measurements of in section 5.8.1.

5.3 Accuracy considerations

5.3.1 Faraday shielding - DC Stark

One of the potential drawbacks of implementing a cavity QED clock is the electric field produced by the cavity PZTs. In order to shield the atoms from the PZT voltage, we make use of the cavity itself as a Faraday shield. The cavity spacer is directly grounded to the vacuum chamber through the thermal braid, ensuring a good $\Delta V = 0$ ground plane relative to the chamber, ensuring that the cavity should not build up any significant electrical charge. The aperture cut through the cavity for the cavity mode is the same diameter as the inner diameter of the PZT (6 mm). In order to evaluate the efficacy of this design, we perform a COMSOL simulation of the electric field distribution. The boundary conditions for the voltage are set by grounding the cavity spacer, and also by an applied voltage to the PZT electrodes. The simulated electric field along the axis of the cavity is shown in Figure 5.2. We consider the possibility of a surface charge on the surface of the mirror of 1 nC. This amount of charge is considered to be a worst case scenario, as this was the magnitude of charge seen by SYRTE that gave a large 10^{-13} shift with essentially no Faraday shielding [126].

This same approach was taken in the recently built Sr1 system (described in [84]), where the Faraday shield had the same inner diameter as the PZT. In Sr1, the DC Stark effect was measured to be less than 10^{-18} , and could be easily compensated down to the 10^{-20} level by applying a small field. While we expect success with the Faraday shielding in Sr3, experiments will have to be performed to confirm this. By mounting electrodes outside of vacuum, one should be able to apply a larger enough electric field to check the dc Stark shift in Sr3.

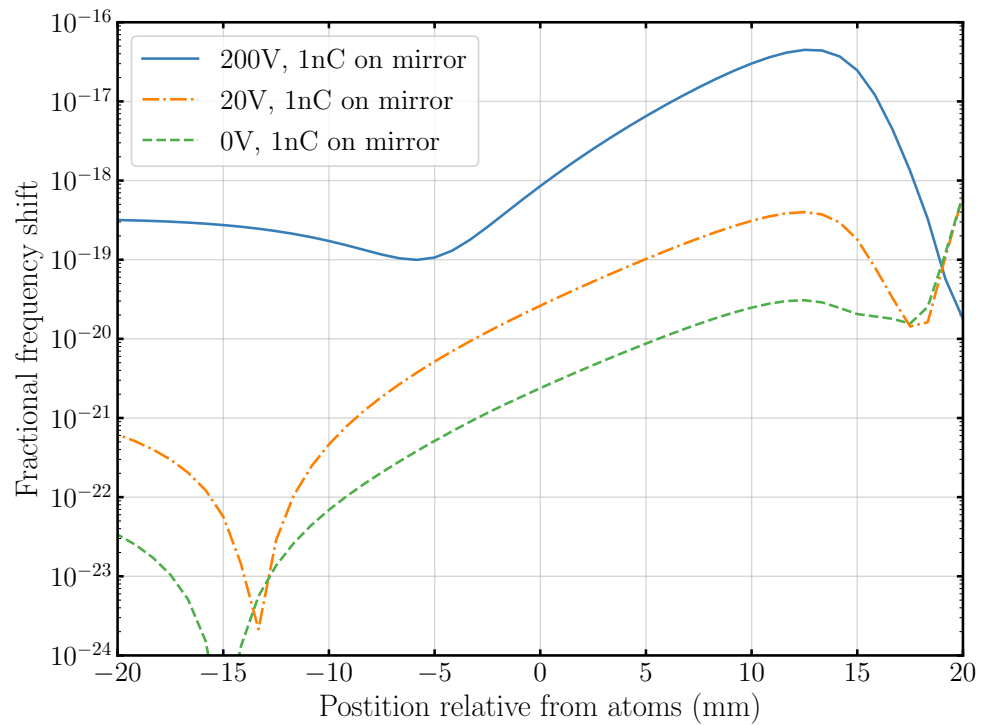


Figure 5.2: COMSOL simulations of the DC electric field at the location of the atoms. We simulate the effect of having 1 nC of charge distributed uniformly on the surface of one of the mirrors. Also shown is the effect of an applied voltage on that same PZT. Even at the highest operating voltage, the overall shift should be on the order of 1×10^{-18} .

5.3.2 In-vacuum calibrated thermal sensors

Looking ahead, the hope is that Sr3 becomes a well-calibrated and accurate clock that can make use of entanglement. In order to achieve the desired accuracy, we require good knowledge of the thermal environment. For this reason, we calibrated four PT100 thermal sensors and installed them in the cavity space and cavity support mount. You can see the two sensors mounted in the cavity spacer in Figure 5.3a. The sensors are housed in a ceramic rod, where the wires are glued to the ceramic rod, but the sensor itself is free floating to avoid strain related shifts of the resistance calibration.

The sensors were first culled by performing repeated thermal cycles up to 150 C and back down to room temperature. After each thermal cycle, the sensors were then checked by realizing an ice melting point, by mixing double distilled water with doubly distilled ice (frozen water). The sensors that exhibited small jumps in the ice melting point value were subsequently calibrated. We used a dry-block comparison method, where we couple the NIST calibrated sensors to a copper block, and then couple our PT100 to that same copper block. We evaluated temperature gradients across the block by comparing the temperature read by each reference sensors, and made sure that the gradient was below 1 mK. Example calibration data for R_{grey} is shown in Fig. 5.3b, with the fit to the resistance data given by a quadratic function. The corresponding uncertainty resulting from the fits, converted to temperature, is shown in Fig. 5.3c.

5.4 Laser systems and cooling

5.4.1 461 nm laser system and blue MOT

The blue laser system is a combination of extended cavity diode lasers (ECDLs) and injection-locked lasers. Over time, the blue system in the Ye lab has become more centralized. The common source is a blue laser locked to the atomic spectrometer. From there, a secondary Toptica 461 nm system serves as seed light for both Sr1 and Sr3. This secondary laser is offset locked to the spectrometer laser. We then seed injection locked lasers. In the past, our lab has routinely used

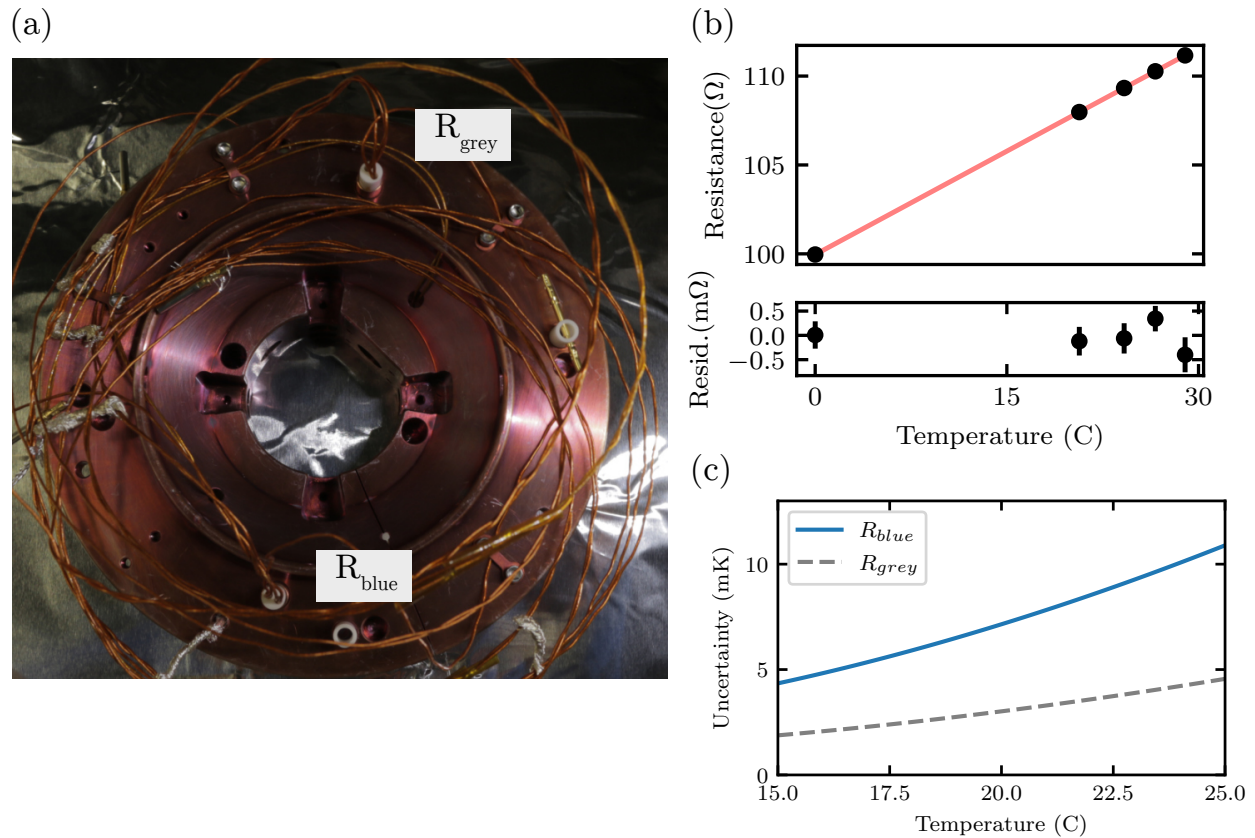


Figure 5.3: (a) Top view of the optical cavity. Shown are the two thermal sensors that are installed in the cavity spacer. (b) Calibration of R_{grey} in a dry-block comparison with three reference sensors. (c) Calibration uncertainty for the two sensors as a function of operating temperature.

injection locked lasers from Nichia operating near 100 mW. For Sr3, we opted to try out a new 500 mW blue laser diode as the injection locked laser [127]. By doing so, we could simplify the blue laser system to only require two injection locked lasers to operate the experiment. The first 500 mW diode provides light for the Zeeman slowing. We find diminishing returns with increased power in the Zeeman slowing beam and that 75 mW at the experiment is a healthy and reasonable amount. Thus, to avoid the damage of the optical fiber, we attenuate the laser before fiber coupling to the experiment. The second diode is split into two overall paths, one serving the 2D MOT and the other for the 3D MOT. On the experiment side, the 2D MOT has about 20 mW total. Each 3D MOT beam has between 5-10 mW of power.

5.4.2 Repump lasers

Initially we utilized two home-built ECDLs for the 679 nm and 707 nm transition. These lasers utilized anti-reflection coated laser diodes, neither of which were close to the desired wavelength at room temperature, thus necessitating heating and cooling near the maximum rated temperatures. We have since replaced these lasers with commercial ECDLs from AOSense. In either case, the lasers are roughly set at the correct frequency using a wavemeter, and then frequency modulated by about 1 GHz in order to address all of the hyperfine states. As a technical note - we generally find that the blue MOT atom number is most critically dependent on the tuning of the repump lasers. This indicates that a better scheme for the repump lasers would be to stabilize them either to an optical cavity or an atomic reference (Sr hollow cathode lamp [128]). Then, one would directly modulate the light with an EOM placing sidebands at precisely the correct frequencies, making much more efficient use of the optical power.

5.4.3 689 nm lasers

To address the narrow, dipole forbidden transition of $^1S_0 \rightarrow ^3P_1$, we require reasonable quiet lasers with minimal frequency drift. We use two 689nm lasers, both homebuilt interference filter ECDLs, which were originally designed by Ben Johnson in Adam Kaufman's group, though the

rough design for interference filter lasers have been around for quite some years. Josie Meyer helped us out by building a few of these ECDLs for the 689 nm laser system. We have one laser each for the $F=9/2$ and $F=11/2$ transitions of $^1S_0 \rightarrow ^3P_1$ in ^{87}Sr . Each laser is phase-locked to the 689 nm cavity system, which is an old ‘‘Ludlow’’ cavity that has an expected performance on the 1-10 Hz level, well below any of our noise requirements for cQED experiments. To control the frequency of each 689nm laser, we put the reference synthesizer for the phase lock into frequency modulation mode. Then, we can either drive the FM port with an arbitrary waveform generator (AWG), to achieve arbitrary modulation. For simplicity, we use an RF switch to change the modulation source to a DAC channel, so we perform more simple linear ramps without having to reprogram the AWG. Just before transferring the atoms from the blue MOT to red MOT, we ramp down the blue MOT laser intensity to about $1/3$ the initial value. This results in a colder blue MOT and optimizes the transfer efficiency to the red MOT.

To capture atoms from the blue MOT to the red MOT, we need to broaden the laser linewidth of the red trapping lasers in order to match the Doppler broadening of the \approx mK temperature blue MOT atoms. For this initial ‘‘broadband red mot’’, we frequency modulate both lasers with sine-wave modulation, and the modulation depth is ramped linearly along with the intensity of the red MOT beams in order to adiabatically transfer atoms from a broadband red MOT to the single frequency red MOT. In order to reach the lowest possible red MOT temperature, the final intensity setpoint is $\approx 0.1 \times I_{sat}$. Our final red MOT temperature is approximately $1 \mu\text{K}$, as measured by time of flight. One special note is that our $F = 11/2$ laser also serves as our cavity probe laser. The detailed optics layout for the implementation of that laser as the cavity probe is shown in Fig. 5.12.

5.4.4 Optical lattice and transport

We trap the atoms in a vertically-oriented optical lattice formed by the interference of two counter-propagating 813 nm beams. The lattice beams are split at the laser breadboard. From there, each beam goes through their own AOM and then independently fiber-coupled to the exper-

iment. After the fiber outcoupler, the laser is picked off and intensity servoed. In order to stabilize the relative phase of the two lattice beams, we form a Mach-Zender interferometer, using a dichroic on the top breadboard as one of the beamsplitters. When the lattice beams are well-aligned, we observe an interference fringe, and can use this fringe for locking the relative phase of the two lattice beams. One issue is that the observed interference phase is not exactly the same phase that the atoms will experience in the vacuum chamber. Through geometric arguments, one can see that there is a differential phase. Thus, despite stabilizing the interference fringe phase, we must be careful to shield the open path length as best we can so that the atoms experience a stable lattice. As a bare minimum, we evaluate the inloop error of the phase-stabilization, shown in Figure 5.4. We compare the inloop error of the phase stabilization to the laser noise model of Si3 and of MJM. For the synthetic laser that probes the cavity, we adopt the Si3 stability below ≈ 100 Hz, and follow MJM above ≈ 100 Hz. We approach the MJM noise model around 1 kHz, and are otherwise safely below both laser noise models [129, 17].

Further improvements to this stabilization scheme can be considered. One drawback of detecting the relative phase of the lattice at DC is that laser intensity noise is indistinguishable from phase noise. This means that we need to do a careful job at controlling the laser intensity (something we want to do anyway, for lattice heating considerations). While we don't show the traces here, the out-of-loop laser intensity dominates over the phase noise for Fourier frequencies above 10 kHz. This is the primary reason for not pushing the bandwidth of the phase stabilization with this setup, as we will actually write the laser intensity noise onto the phase of the optical lattice. Experimentally, we find that if we continue to increase the bandwidth, we start to suffer from atom loss. Using an EOM to sense the phase difference at an RF frequency in a similar configuration as described in [130] might allow for tighter phase control. Another consideration is that there can be out-of-loop phase fluctuations. For this, shortening the out-of-loop path (although it is already as short as possible given the optics layout), and better shielding of the environment from pressure and temperature fluctuations is always the way to go.

After the atoms are loaded into the lattice, we transport the atoms in this movable lattice

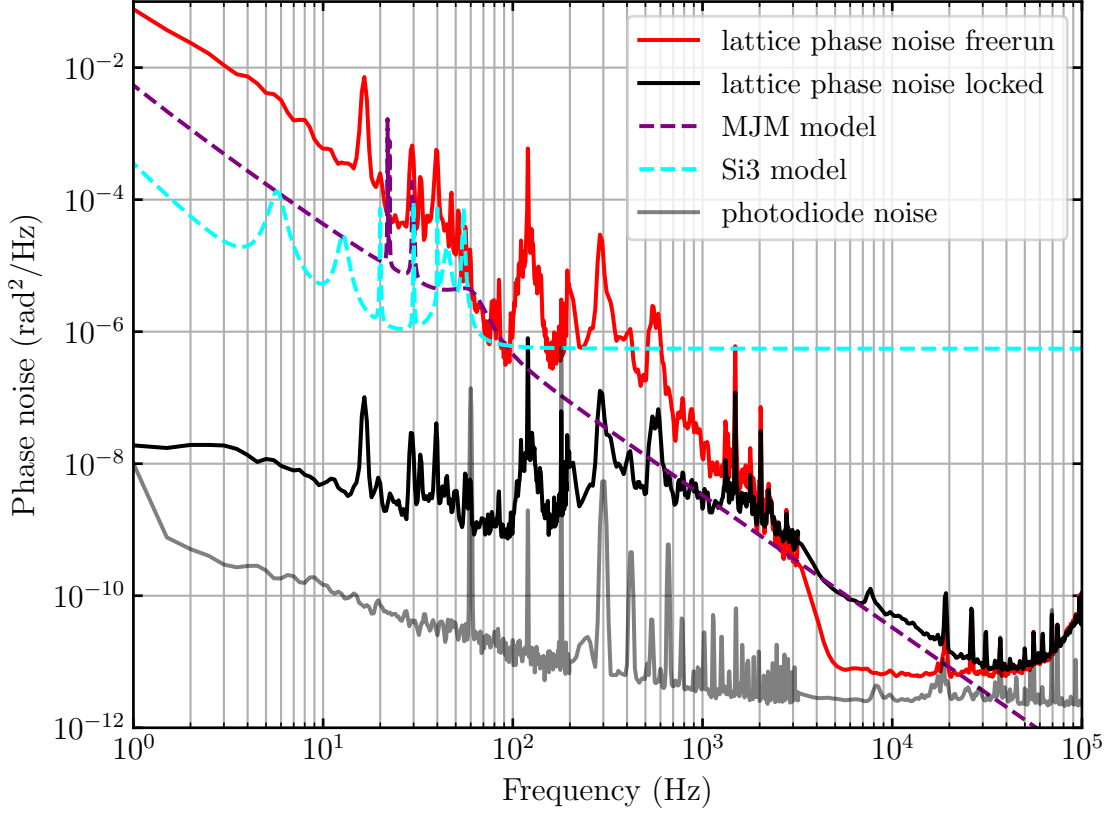


Figure 5.4: Lattice phase noise measurements. The free running lattice phase noise is the red trace, and the locked inloop error is the black trace. For comparison, the laser noise models for the two clock lasers MJM and Si3 are shown as the purple and cyan dashed curves. The photodiode noise floor is the gray trace.

by ramping the frequency of the AOM in the bottom-up lattice beam. The motion of the lattice is induced by changing the relative frequency of the two lattice beams and is given by

$$\Delta \vec{v}_{lat} = \frac{\lambda_{lat}}{2} \Delta \nu, \quad (5.1)$$

where $\Delta \vec{v}_{lat}$ is the velocity of the lattice and $\Delta \nu$ is the relative frequency detuning of the lattice beams. The transport waveform is shown in Figure 5.5. One important detail is that we mix three frequency synthesizers together to form the AOM drive for the bottom-up lattice beam. In this way, we can trigger three independent frequency ramp profiles corresponding to three transport distances. This allows us to do smaller vertical manipulations of the cloud position relative to the cavity, enabling the probing of multiple ensembles with the cavity. Future upgrades to the




Figure 5.5: The left axis indicates the relative frequency detuning of the optical lattice beams for transport, and the right axis is the corresponding velocity of the lattice. The total transport distance is approximately 42 mm.

RF synthesis will allow for more arbitrary configuration of ramp profiles, including more gentle Gaussian ramps instead of the linear ramps we employ here. Another consideration for future upgrades is the fact that during our transports, we allow the lattice phase to free run. One worries about this, especially in the differential clock comparisons of Chapter 7. One can imagine that small phase fluctuations lead to displacement noise of the atomic cloud, and if a measurement is sensitive to this position (i.e. with respect to the clock laser phase), this can be a source of unwanted noise. More sophisticated servo control can keep the lattice phase servoed during the transport operations. We also wish to have a focus-tunable lens so that the tightest confinement can be applied to the atoms in both the MOT bucket and the science cavity. This would allow for high radial trap frequencies, more easily driving the clock transition in the resolved radial sideband regime $\Omega \ll \omega_R$.

5.5 Clock laser

The clock laser is critical to the operation of the Sr3 experiment. We discussed in detail the Si4 and Si6 ultrastable laser systems in Chapter 2. Here we will discuss some of the critical aspects of the clock laser as it applies to the Sr3 experiment.

5.5.1 Fiber noise cancellation

In order to drive the clock transition in a low noise fashion, we need to transfer the laser light from the distribution center to the experimental table without adding appreciable phase noise. The distribution center houses the injection locked laser, which receives its seed from the fully stabilized clock light. We employ fiber noise cancellation (FNC), where we have a reference plane (mirror) at the distribution center, and another reference plane at the experiment where the light is retro-reflected through the fiber. By interfering the light from these two reference planes, we directly measure the relative phase noise introduced by the fiber. The 0th order of the second AOM at the experimental table is retro-reflected as the reference light. We feedback to the AOM at the distribution center, with a bandwidth of $\approx 100\text{kHz}$.

Due to the fact that we stabilize the 0th order, and the 1st order goes to the atoms, any differential noise between these two paths will be seen by the atoms. To evaluate the differential noise, we temporarily insert a retro-reflector in the path of the 1st order beam, while stabilizing the 0th order beam as usual. The 1st order beam will produce a beatnote at the photodetector that is distinguished as a particular RF tone. If we move the retro-reflector and measure the phase noise of this 1st order beat will tell us about the differential noise. The measured noise of this 1st order beat is shown in Figure 5.6. When we place the 1st order retro close to the 0th order reference, we observe the black trace, showing noise well below the laser noise model of Si3 (dashed line). Moving the 1st order retro to just before the chamber, we see the noise is slightly larger, but still at an acceptable level. We note that in order to achieve this noise performance, we cover the clock optics with a crude enclosure of blackened aluminum foil, and close the enclosure around our entire

optics table.

Furthermore, any differential noise between the 1st order clock light and the optical lattice will also appear as clock noise. We place the clock retro for fiber noise cancellation as close to the optical lattice phase reference as possible. To minimize any potential relative path length noise, we cover the optics in a secondary enclosure of blackened foil.

5.6 Atomic state preparation

Now it's time to put all of the above-described laser systems into action. The overall experimental sequence is displayed in Figure 5.7. The sequence starts with cooling and trapping the atoms in a blue MOT. We send 13 A of current through our MOT coils, corresponding to an anti-Helmholtz magnetic field gradient of about 30 G/cm. A rough estimate of the blue MOT atom number is 10 million. At the end of the blue MOT, we linearly ramp the intensity of the 3D MOT beams down in an out-of-loop fashion, by about a factor of 10. This leads to increased transfer efficiency from the blue MOT to the broadband red MOT, giving $\approx 2x$ more atoms.

The red MOT operates on the dipole-forbidden transition of $^1S_0 \rightarrow ^3P_1$ [131]. We transfer the atoms from the blue MOT to a broadband red MOT by simultaneously turning off the blue MOT beams and turning on the red MOT lasers at full intensity ($\approx 1000I/I_{sat}$ each). During the initial broadband red MOT, the 689 nm trapping lasers are frequency modulated to ≈ 1 MHz wide for the 11/2 laser and ≈ 500 kHz for the 9/2 laser, respectively. The broadband to single frequency transition is realized by linearly ramping the modulation depth of the frequency modulation amplitude down to zero over 250 ms. Simultaneous with the frequency modulation ramp, we linearly ramp the laser intensity down to less than I_{sat} . Going to this low intensity results in a red MOT temperature of about 1 μ K, as measured by time-of-flight spectroscopy.

During nominal operation, the vertical optical lattice is kept on during the entire blue and red MOT sequence, meaning that we are continually loading the lattice. We note that because the 813 nm trapping lattice causes an ac Stark shift of the $^1S_0 \rightarrow ^3P_1$ transition, we optimize with this inhomogenous shift always present. We suspect that we get away with this given the size of the




Figure 5.6: Differential fiber noise between the 0th order and 1st order beams. The black trace indicates when we place the 1st order retroreflection near the 0th order retro, and red is placing the retro as close to the vacuum chamber as possible, just before the viewport. For reference, we plot the Si3 noise model, showing reasonably adequate noise cancellation.

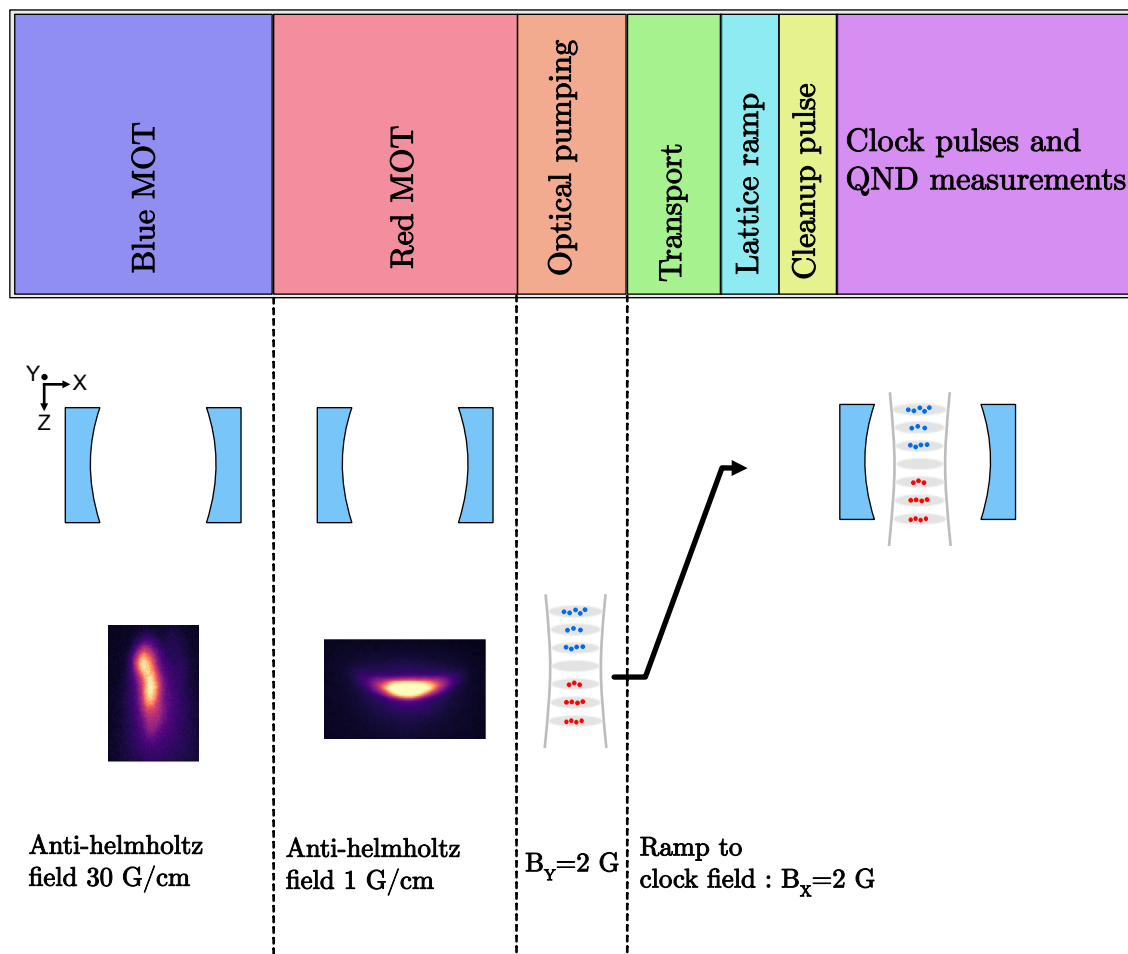


Figure 5.7: The top panel shows the experimental sequence roughly laid out in time. Below we graphically indicate the location of the MOTs, and the subsequent loading of the lattice and transport into the cavity. We also note the magnitude and direction of the relevant magnetic field for each step.

optical lattice is $\approx 150 \mu\text{m}$ diameter at the MOT location, keeping the inhomogeneous ac Stark shift to an acceptable level. However, it might be possible to improve the red MOT by instead ramping the optical lattice up at the end of the MOT sequence. This can involve frequency ramping the red cooling lasers to compensate for the ac Stark shift, to optimize cooling of the atoms inside the lattice.

After extinguishing the red MOT light, we perform optical pumping into the ground-state hyperfine level $|^1S_0, m_F = -9/2\rangle$. We apply a bias magnetic field of about 2 G along the Y direction, and shine on an optical pumping beam tuned near the $|^1S_0\rangle \rightarrow |^3P_1, F = +9/2\rangle$ transition with circularly polarized light, in order to pump the atoms into $m_F = -9/2$. This optical pumping is optimized by adjusting the frequency, intensity, and duration of the pumping, to achieve approximately 70% of the atoms in $m_F = -9/2$. We then move the cloud up into the cavity by detuning the bottom-up lattice beam, as described in Section 5.4.4. Once there, we begin clock state manipulation and QND probing. Future improvements will implement optical pumping in tandem with sideband/radial cooling.

5.7 Cavity-stabilization 813nm laser

In the very early stages of the Sr3 experiment, we stabilized the cQED cavity to our M Squared Ti:S laser. Over time we realized that the performance of this Ti:S was much noisier than what we required for the cQED cavity stabilization. We would often see large amounts of frequency noise manifest itself as intensity noise in the transmission of the cavity, and this signal was actually a good diagnostic for when the Ti:S was misbehaving. We knew we could build a lower noise laser than the Ti:S, albeit at lower power, so we decided to build an independent ECDL for the cavity stabilization. This ECDL is based on the same design as our 689 nm lasers described in section 5.4.3. The 813 nm laser uses an anti-reflection coated laser diode from Eagleyard (EYP-RWE-0840-06010-1500-SOT02-0000), and the interference filter is a 3 nm FWHM bandpass filter from Iridian (MI000002). The free-running Lorentzian linewidth of this laser is ≈ 20 kHz, which is about a factor of 10 better than the M-Squared Ti:S. One unfortunate note on this laser - we have had

Cavity parameters		
	689 nm	813 nm
\mathcal{C}	0.3	NA
$g_0/2\pi$ (kHz)	8.6	NA
L (cm)	6.9	6.9
finesse, \mathcal{F}	16,000	8,000
κ (Hz)	157,000 (measured)	330,000
mode waist (μm)	71	75
T_1 (ppm)	280 (specified)	200
T_2 (ppm)	20 (specified)	200
total loss L	100 (estimated)	100 (estimated)

Table 5.1: Key parameters of optical cavity

the PZT fail on this laser twice during power outages. We therefore warn any graduate student reading this thesis to not use the cylindrical PZTs from Thorlabs in any critical application.

5.8 cQED cavity

5.8.1 Cavity design and construction

Perhaps the most unique and critical aspect to the Sr3 experiment is the optical cavity. The cavity parameters are shown in Table 5.1. We designed the cavity to operate on the $\lambda = 689$ nm $^1S_0 \rightarrow ^3P_1$ transition. The cavity geometry and the atomic transition properties determine the peak atom-cavity coupling, which we calculate to be $g_0/2\pi = 8.6$ kHz. Correspondingly, the peak single-atom cooperativity is given by $\mathcal{C} = 4g_0^2/(\kappa\Gamma) = 0.25$. This enables us to enter the collective strong-coupling regime ($N\mathcal{C} \gg 1$) with a few thousands of atoms.

The cavity is designed as a single-ended cavity at 689 nm, with specified mirror transimission values of $T_1 = 280$ ppm and $T_2 = 20$ ppm. This means that intra-cavity light will preferentially exit through the input mirror, making that the optimal measurement port for quantum efficiency. Before baking the system, we measured a cavity finesse of $\mathcal{F} = 20,000$, which was consistent with the specified transmission values of the mirrors and negligible optical loss. We knew the bake out




Figure 5.8: Measurement of the cavity finesse at 689 nm by monitoring the transmission after abruptly shutting off the light. The black data and fit are before the bake out, giving $\mathcal{F} = 20,000$. The red data and fit are after the chamber bakeout, giving $\mathcal{F} = 16,000$. As a check, we temporarily move the photodetector to before the cavity (gray data), to confirm our light shut-off time is sufficiently short.

process can be a risky time for such high reflectivity mirrors. To try to minimize the thermal gradients in the chamber, we heated up the chamber over the course of ≈ 24 hours. Although the bakeout was successful in reaching the 10^{-11} Torr level, the cavity finesse ended up degrading during the bakeout. We calculate that the cavity obtained an additional loss of $L = 100$ ppm. We suspect that this increase of loss could have come from a deposition of residual hydrocarbons in the chamber. The primary consequence of the increase in loss is a decrease in the first mirror transmission rate κ_1 relative to the total cavity power decay rate κ . This is quantified and referred to as the quantum efficiency for the cavity $Q_{cav} = \kappa_1/\kappa$, which went from $Q_{cav} = 0.85$ to $Q_{cav} = 0.75$. This results in a change in the overall quantum efficiency of just 5%. Given the nonzero chance that cleaning the mirrors could actually make things worse, we opted to move forward with the slightly degraded mirrors.

For cavity stabilization, the cavity mirrors also have high reflectivity at 813nm, with the cavity parameters at this wavelength shown in the second column. The cavity linewidth at 813nm is 330 kHz. This relatively narrow linewidth at 813 nm of $\kappa_{813} = 300$ kHz was chosen to make cavity stabilization less susceptible to technical effects like RAM. However, the narrower linewidth comes with the tradeoff of needing to deal with the problem of FM to AM conversion.

A closeup photograph of one of the cavity mirror assemblies is shown in Figure 5.9, and a full view of the cavity is shown in Fig. 5.10. The cavity employs two PZTs supplied by Noliac, one having $13.2 \mu\text{m}$ travel (NAC2123-H10-A01) and the other having $3.3 \mu\text{m}$ travel (NAC2123-A01). The idea is that the longer travel PZT could allow us to attempt to make the 813nm mode closer to the magic wavelength, but for the experiments in this thesis we have not used the longer travel PZT. We note one consideration to PZT stacks is that the conversion from voltage to displacement is multiplied by the number of PZTs glued together, in our case 4x. For the 100 Hz linewidth goal of our stabilization, the high voltage PZT driver can be a significant source of noise. The PZTs are glued to the copper plate using a very thin layer of EpoTek 353ND epoxy, and then a copper “mirror cup” is glued to the top of the PZT with the same epoxy. The mirror cup has three extrusions to gently center the mirror on the cup. The mirror is glued in the midplane of the mirror




Figure 5.9: Up close photo of the mirror assembly with the PZT stack. You can also see two of the Ceramabond epoxy joints holding the mirror inside the mirror cup. The whole epoxied assembly is bolted to the cavity spacer with four vented screws, of which two are visible here. The other mirror, is identical apart from the overall height of the PZT.

using a ceramic based adhesive Ceramabond 853. Prior to choosing these epoxies for use in our vacuum chamber, both epoxies were vacuum tested using a residual gas analyzer, demonstrating outgassing limited in the test chamber at to the 10^{-10} Torr level. We note that when gluing cavity mirrors for optical cavities - it's critical to choose an epoxy that does not induce a large amount of stress on the optic. During the first gluing of the mirrors for the Sr1 buildup cavity, we found that 353ND epoxy could go as far as to shatter the fused silica optic, which drove us to the more gentle ceramic-based adhesives.

The mechanical support for the cavity is a double pendulum isolation scheme. We suspend a heavy intermediate piece of copper from the vacuum chamber with three short rods of viton serving as the lossy spring. The optical cavity is then suspended from this intermediate piece with three more pieces of viton. The idea of this double pendulum is that it will greatly attenuate vibrations coupled to the cavity above the resonant frequencies. By simulating the coupled system in COMSOL, we predicted that the resonant frequencies of this mechanical system would be clustered between 10 - 50 Hz. In hindsight, while these mechanical resonances do indeed improve the higher frequency noise of the optical cavity, it poses significant challenges when our cavity frequency measurement window lasts 40 ms. Future experiments that further reduce the atomic motion might permit measurements with shorter duration, taking further advantage of this mechanical isolation scheme.

While copper has nice properties for vacuum and BBR, there is the downside of being capable of supporting significant Eddy currents. The largest change in magnetic flux will occur when we switch the MOT field and when we ramp on the clock bias field. These changes in B field can and will induce currents in the copper spacer, leading to time-dependent forces on the cavity. Furthermore, these Eddy currents could take considerable time to damp out. To mitigate this effect, we used a wire-EDM to cut a very small slot in the optical cavity and the supporting copper piece. We measured the effect of including such a slot using a test spacer. We suspended the cavity test spacer from the upper copper piece, and mounted these pieces above the MOT bucket. A 3-axis magnetometer was placed inside the cavity spacer, and we rapidly switched off the current




Figure 5.10: Full view of the cavity, as it would be seen from below after installation. In clear view are the two mirrors. Also visible is the slot we cut in the cavity, to reduce Eddy currents. The cavity is mounted onto an aluminum jig that was used to secure the cavity during installation of the various components.

in the MOT coil and measured the time response of the magnetic field (Figure 5.11). We checked the response with no cavity present (orange), with the cavity but without the slot cut (blue), and the cavity in with the slot cut (green).

5.8.2 Cavity stabilization

In this section we will describe the cavity stabilization scheme that we used for most of our early experiments in Sr3. The requirements for the frequency noise performance depends on several factors. The primary one is the magnitude of the cavity frequency fluctuations due to quantum projection noise. When we are operating in the slightly detuned regime, the QPN fluctuations are roughly 2 kHz for approximately 10,000 atoms. If we want the technical noise floor coming from relative laser - cavity frequency noise to be 20 dB below QPN, we need to reach a frequency resolution of 200 Hz. As we will discuss, the timescale of the probing will change what Fourier frequency we are sensitive to, and this can impact the strategy for stabilization. To put this in context for our experiment, we probe the cavity with a measurement time of 40 ms, meaning that we need to focus on frequency noise in the range of 1 - 100 Hz.

The simple idea is to make the cavity length as stable as possible for cQED probing of the atom-cavity system. To this end, we first stabilize the 813 nm ECDL laser described in section 5.7 to the optical frequency comb by feeding back to the current at high frequency, and to the PZT of that ECDL at lower frequencies. We achieve a locking bandwidth of ≈ 300 kHz. This allows us to stabilize this laser down to the noise floor given by Si3+comb in addition to the uncanceled fiber noise floor between the ECDL and the comb. We then send this 813 nm light to the cQED cavity, and stabilize the cavity length to the ECDL via a PDH lock, feeding back to the cavity PZT. For the experiments in this thesis, we keep the intra-cavity lattice at a low level so as not to significantly trap the atoms along the cavity. To this end, we turn the incident power down to $\approx 10 \mu\text{W}$, and actively stabilize the power in transmission of the cavity. After the bakeout the cavity has a finesse of ≈ 6300 at 813 nm, leading to a power buildup of ≈ 1200 . We estimate that our operational trap depth in the cavity is $\approx 0.1E_r$. The relative large finesse for 813nm allows us to more easily




Figure 5.11: Response of the magnetic field when rapidly switching off the current through the MOT coils. The magnetometer is placed inside a test cavity spacer having roughly the same geometry as our actual spacer. We measure the response without the spacer (orange), with the spacer (blue) and with the spacer but with a slot cut into it (green). We show a dramatic effect of including the slot in the spacer, showing a faster damping time constant of induced Eddy currents in the spacer.

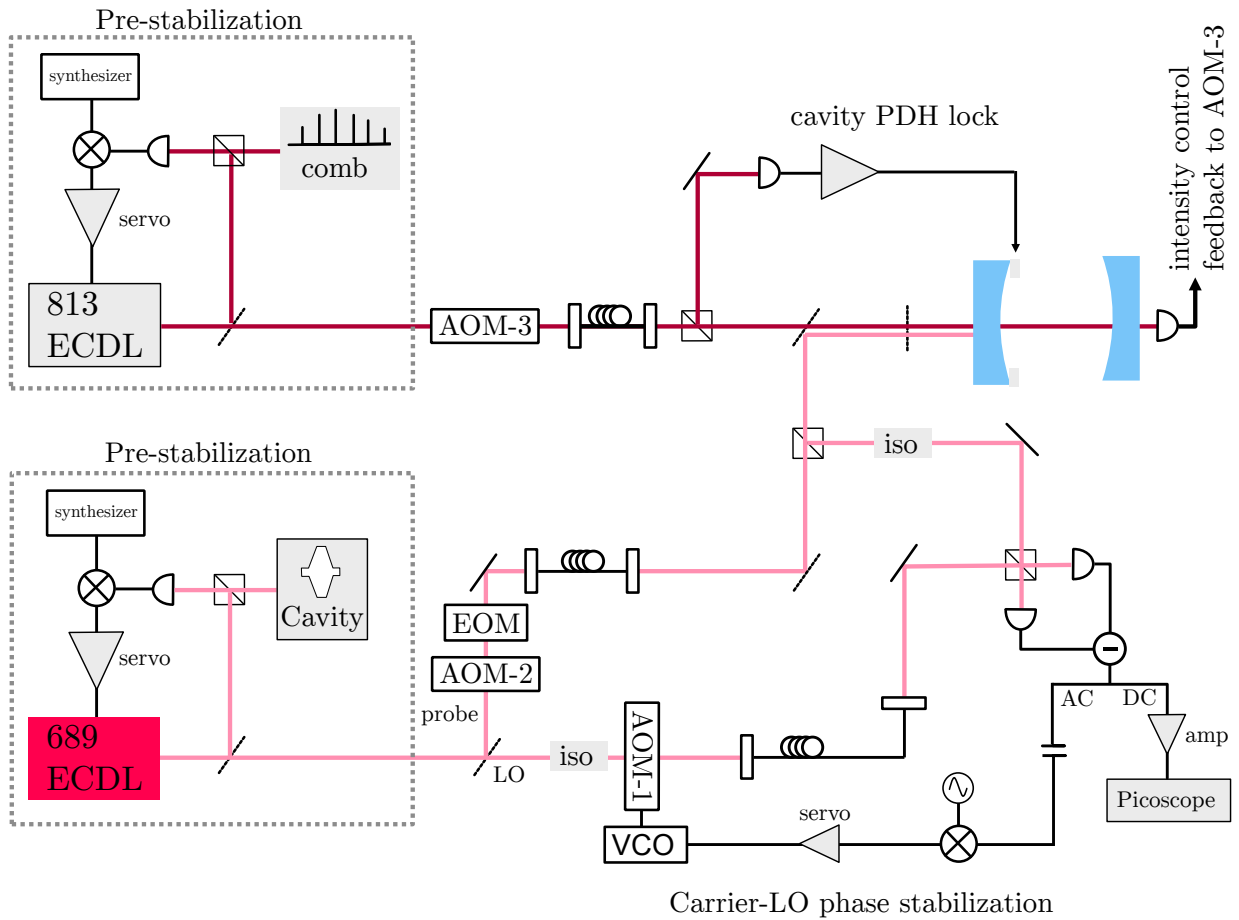


Figure 5.12: The optics layout and stabilization scheme of the cavity QED setup. The 813nm laser is pre-stabilized, and the cavity is locked to that pre-stabilized laser via the cavity PDH lock. The 689nm laser is pre-stabilized to the red cavity system. The 689 ECDL is split into the LO and probe path. The leakage light from the cavity is detected via homodyne detection.




Figure 5.13: Frequency diagram of the cavity modes and of the different laser light tones. The bare cavity frequency ω_c is detuned from the bare atomic resonance ω_a by $\delta_c/2\pi = 1$ MHz. The cavity is shifted by the presence of atoms in the ground state $|\downarrow\rangle$. The relevant frequencies are shown. The carrier is synthesized by frequency shift of AOM-2, and the sideband that probes the cavity is generated by the EOM drive. The LO is correspondingly frequency shifted to be on resonance with the sideband for homodyne detection.

stabilize the cavity length, since a narrow cavity resonance makes issues like RAM less worrisome. The tradeoff is with FM \rightarrow AM conversion, making the implementation of an intra-cavity lattice more challenging.

To measure the noise floor set by the relative fluctuations of the laser and the cavity, we probe the cavity with a relatively large optical power (nW or so). We do this so that the photon shot noise limit is negligible, and we can unveil the technical frequency noise floor that does not depend on optical power. The relevant noise measurements are shown in Figure 5.14. For reference, the free-running cavity noise is the purple trace, showing enormous noise at lower frequencies. The inloop error of the 813nm lock to the comb is the blue trace, showing more than sufficient stabilization. We note that this is an optimistic trace, as there is a 30 m optical fiber that is currently not fiber

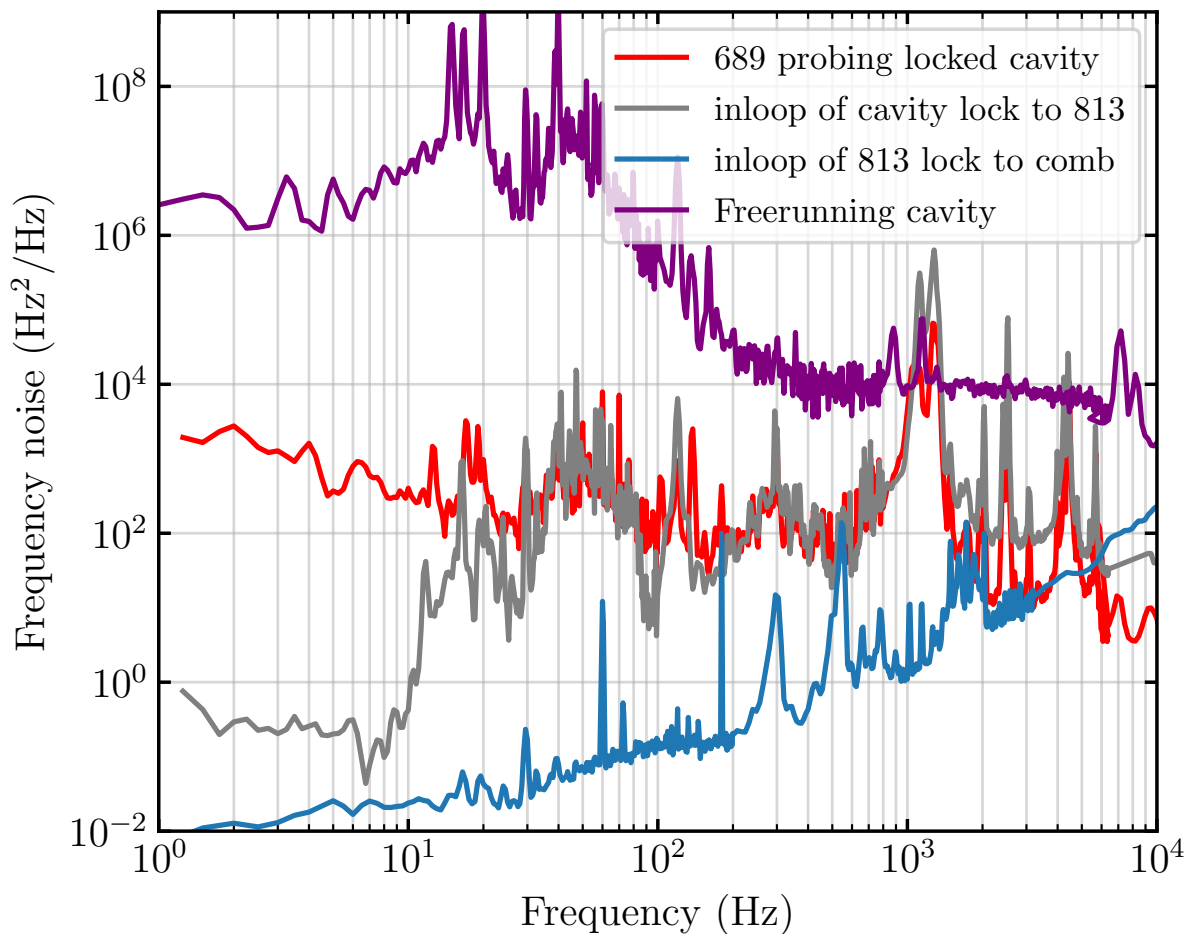


Figure 5.14: Noise measurements of cavity QED stabilization scheme. The blue line is the inloop error of the 813 ECDL lock to the comb. The gray line is the inloop error of the PDH lock of the cavity to the 813 laser. The 689nm pre-stabilized laser then probes the cavity - serving as the out of loop noise measurement. We can see the frequency bands where we are still limited by the inloop lock of the cavity.

noise cancelled, and can easily be a contributor to noise at the $100 \text{ Hz}^2/\text{Hz}$ level. The inloop error of the cavity lock to the 813 nm laser is shown as the gray trace. When we then probe the cavity with the pre-stabilized 689nm laser, we measure the frequency noise spectrum shown as the red curve. We can see the servo bump at around 1 kHz, which is limited by the ringing up of mechanical resonances starting at 1.2 kHz and 2.5 kHz.

To measure the cavity frequency, we employ homodyne detection of the probe laser. For the cavity QED experiments described later in this thesis, we will be doing repeated cavity frequency shift measurements, where the cavity resonance has been shifted by the presence of atom in the ground state. We will refer to the repeated probe of the cavity resonance as “two window probing”. The transfer function for a two-window measurement is

$$T(f) = \frac{4 \sin^2(\pi f(T_m + \tau_m)) \sin^2(\pi f T_m)}{(\pi f T_m)^2}. \quad (5.2)$$

We can multiply the measured noise spectrum (red trace from Figure 5.14) with this transfer function and predict the noise of the two window probing. When we perform this calculation, we find that we can push the technical noise of two repeated cavity frequency measurements to approximately 200 Hz rms for a measurement time of $T_m = 40 \text{ ms}$.

Looking forward, one can imagine various ways to improve the cavity stabilization system. The relative frequency noise between the 689 nm laser and the locked optical cavity is largely limited by the inloop stabilization of the cavity in the band from 30 Hz to 1 kHz. Since there are already strong correlations between the frequency noise at 689 nm and the inloop PDH error signal at 813 nm from 20 Hz - 10 kHz, one could imagine implementing feedforward with no need for any changes to the hardware. As an alternative to feedforward, one could directly improve the relative noise by implementing a more sophisticated loop transfer function that compensates for the mechanical resonances, allowing for a higher bandwidth. Either of these approaches would improve the ultimate technical noise floor of the cavity frequency measurement system. Further improvement to the low frequency noise below 20 Hz would also be beneficial. A primary suspect for the excess low frequency noise is residual amplitude modulation in the 813 nm cavity stabilization

lock.

5.9 Homodyne interferometer

The homodyne interferometer is shown in the lower right corner of Figure 5.12. This is where we will measure the leakage light from the atom-cavity coupled system to perform QND measurements. The primary goal is to achieve a photon shot noise limited homodyne detection of the optical tone probing the atom-cavity system.

Performing homodyne detection combined with our somewhat long probe time of 40 ms means we are increasingly sensitive to low frequency noise. We found that pointing instability of the LO beam on its own (with the signal port blocked) could give rise to significant low frequency noise in the output of the homodyne photodetector. We confirmed the source of the low frequency noise by picking off a small fraction of the beam near the homodyne PD and detect it with a quadrant PD. There were significant correlations between the homodyne noise and the beam position as measured by the quadrant PD. We then reconfigured the interferometer to keep the path as short as was reasonably possible, and intentionally tilted optics to prevent optical etalons. Furthermore, we enclosed the optics surrounding the cavity in a secondary box, although with an imperfect seal, and found that this greatly improved our low frequency noise performance.

The homodyne fringe is detected using a home built balanced photodetector. The photodetector operates in current-subtracting mode, where the two photodiodes are wired back to back such that the sensed current is the difference between the two photocurrents. The technical dark noise of the photodetector is ≈ 31 dB below the LO photon shot noise for an LO power of 80 μ W. This photodetector design has advantages in simplification of the electronics, but means that there is not a convenient way to electronically balance the two ports for optimal intensity noise suppression. While one can fine tune the balance of the beamsplitter with polarization/alignment, we add a safety cushion by adding active intensity control to both the LO and cavity probe beams, The combination of active stabilization of the LO intensity and the common-mode rejection of the LO intensity noise allows the homodyne detection to be *nearly* photon shot noise limited.

There are two primary sources of technical noise on top of the photon shot noise. We send the LO beam through the homodyne interferometer, and measure the output noise (Figure ??). Note here that we do not add any extra gain at the output of the photodetector, so the photocurrent noise is amplified by the transimpedance gain of $10^6 \Omega$. In Figure ??, we can see the noise rises with a $1/f$ scaling, and fit a noise model to the data shown as the orange line. The expected photon shot noise is the red dashed line, and the dark noise is the solid blue line.

If one looks at the noise of the LO alone (signal port blocked) through the homodyne beamsplitter and detection, there are significant peaks in the noise spectrum (Figure 5.15). We placed an accelerometer on the table near the homodyne interferometer, and discovered strong correlations with the homodyne output voltage at certain frequencies. The distinct frequencies at ≈ 20 Hz and ≈ 30 Hz come from air conditioning motor vibrations. There is another cluster of peaks in the 300-400 Hz range that we identified to be the mechanical resonances of the optical breadboard that houses all homodyne optics. Our understanding of the coupling between vibrations of the optical breadboard and the homodyne output voltage is that the vibrations drive differential pointing noise of the LO beam through the interferometer. There can easily be differential etalons between the two paths after the beamsplitter to each photodiode. Such a differential etalon could give a varying intensity due to mechanical fluctuations. Interestingly, we have found that the amplitude of the vibration induced peaks in the homodyne spectrum depends on the precise alignment of the LO beam. This needs to be kept in mind when optimizing noise performance of the homodyne detection. Thinking toward the future, it will be advantageous to make the homodyne interferometer even more compact. An improved design would involve mounting all of the optics to a common chunk of metal in a very compact fashion. The optics would be completely enclosed so as to prevent air currents from giving time-varying etalon effects. The one downside is the compactness of the optics can make it difficult to avoid spurious reflections, but should still be manageable.

When we are making measurements of the cavity frequency, we synchronously record the homodyne signal and the accelerometer output on an oscilloscope (Picoscope), and perform subtraction of the vibration trace in order to reach the photon shot noise limit. The only free parameter

in the subtraction is the arbitrary scale factor between the homodyne voltage and the accelerometer voltage - there is no significant phase shift between the signals.

5.9.1 Path length stabilization

As seen in Figure 5.12, we split the 689 nm laser into the probing path and the LO path. Any relative path length fluctuation between the two paths will lead to relative phase noise between the probe light and the LO light. When probing with heterodyne, this phase noise can impact the measurement, but slow overall relative phase drifts are less of a concern, since the heterodyne signal spends equal time in the \mathcal{I} and the \mathcal{Q} quadrature. However, when probing the cavity on resonance with homodyne detection, we need to actively control and set the relative phase. This is because the most sensitive quadrature on resonance is \mathcal{Q} , and any deviation will distort the homodyne signal and reduce the signal to noise ratio. We thus implement a path length servo, shown in the bottom of Figure 5.12. We detect the beatnote between the carrier and the LO beam, which are shown in frequency space in Figure 5.13. This beatnote is stabilized to a reference synthesizer by feeding back to AOM-1 in the LO path. We note in passing that this path-length beatnote appears at the same exact frequency as the EOM drive.

We mention one minor technical detail regarding the RF driving. While one could simply split the RF drive for the EOM and use that synthesizer as the phase lock reference, we opt to use two independent synthesizers that are phase locked to the same 10 MHz microwave reference. This way, one can change the relative phase by adjusting the synthesizer settings, instead of dealing with delay lines, or other ways of setting the relative phase. The downside is that the relative phase noise is not common-mode at higher frequency. Therefore we opt to use a low phase noise synthesizer SRS SG382, such that the contributing phase noise of the synthesizer is well below the technical noise floor of the laser-cavity system. Another consideration for the path length servo is the photon shot noise. The power in the carrier is approximately 150 nW. This relatively high carrier power compared to the probe sideband means the photon shot noise limit will not degrade the signal to noise of the detection of the sideband. One could imagine turning down the carrier




Figure 5.15: (a) Measured output of the homodyne voltage (blue) and the accelerometer (orange) placed on the optical table near the homodyne interferometer. Both signals have been bandpass filtered in software to make the correlations clearly visible. (b) The same signals in the frequency domain. After we subtract the accelerations from the homodyne voltage, we can see strong cancellation of the noise peaks at 20 Hz and 30 Hz, without impacting the homodyne noise at other frequencies. We note here that the photodiode voltage is amplified by a voltage factor of 25, on top of the 10^6 gain of the transimpedance amplifier.

power by a factor of 10 without facing serious consequences.

Chapter 6

Generating and measuring entanglement on the Sr clock transition

With this newly built Sr3 experiment, we can now ask: can we directly observe an optical clock ticking with fluctuations below the fundamental limits set by the quantum projection noise?

Before even dreaming of obtaining that result, there were many aspects of the system to implement and iron out. In this Chapter, we discuss performing state-of-the-art clock spectroscopy in our system, and use it to understand the limitations imposed by atomic motion. In Section 6.2, we establish coupling between the atoms and the cavity mode, entering the strong collective coupling regime. In Section 6.3, we directly measure the QPN, allowing us to carefully calibrate the effective atom cavity coupling g_{eff} . We then demonstrate QND measurements with measured spin noise reduction of -4.7 dB below QPN. Together with measurements of the Ramsey fringe contrast loss, we measure a spin squeezed state that has phase resolution 1.7 dB below the Standard Quantum Limit, verifying the presence of entanglement.

6.1 Clock spectroscopy

We first explore the temperature of our atomic sample by probing the atoms with clock spectroscopy. As the simulations from Chapter 3 suggest, and by intuition, we know that we are going to be happier with a colder atomic sample. We take a brute force approach to prepare a gas with temperatures less than $1 \mu\text{K}$. After the initial state preparation described in Section 5.6, we remove the higher temperature atoms by applying a “lattice ramp”, which involves ramping the optical lattice to $\approx 7E_r$ over 10 ms, hold there for 20 ms, and then back up to $20E_r$ in 10

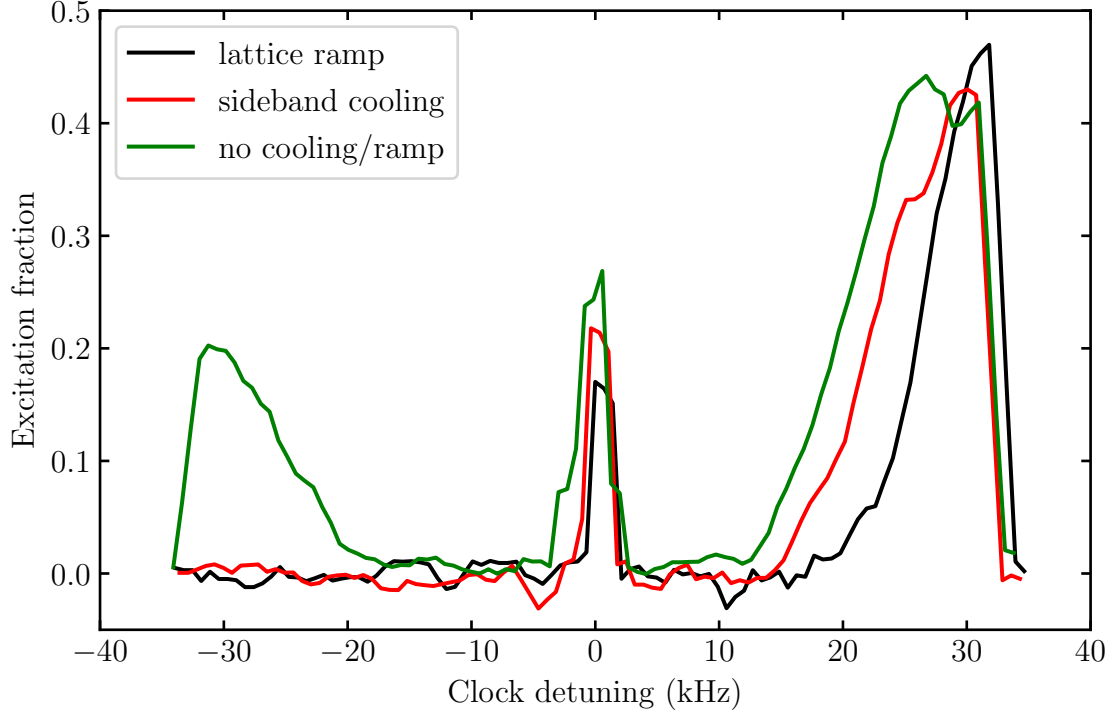


Figure 6.1: Sideband scans of the clock transition. Without any ramping of the lattice or additional cooling (green), with sideband cooling (red), and with the lattice ramp along (black). The lattice ramp gives a sample prepared primarily in the ground motional band $\langle n_z \rangle = 0$, and a narrower blue sideband (radial temperature). The carrier signal is heavily saturated in this clock scan.

ms. The ramp is adiabatic with respect to the fast motion along the tightly confined direction, and quasi-adiabatic with respect to the radial motion. There is no phase-space compression in this lattice ramp, it instead lets the hotter atoms to simply leave the trap. It is a bit of a shortcut compared to implementing efficient radial and sideband cooling. At the expense of throwing away about half of our atoms, it easily leads to a cold sample of 20,000 atoms.

We perform sideband clock spectroscopy in order to look at the population of higher motional bands (Figure 6.1). The sideband spectrum just after transport and optical pumping but without further cooling strategies is the green trace, giving a temperature of $1.2 \mu\text{K}$. We compare just sideband cooling vs just the lattice ramp alone in the red and black traces, respectively. Both lead to $\langle n_z \rangle \rightarrow 0$, but the lattice ramp also leads to a colder radial sample as seen by the width of the blue sideband.

Extracting the radial temperature from the blue sideband is not a reliable method. To carefully measure the radial atomic temperature, we perform radial scans of the clock transition. For this, we send the clock laser along the radial direction of the trapped pancakes in the 1D lattice. In this scenario the clock laser probes the atoms well outside the Lamb-Dicke regime, and will be sensitive to the Doppler shifts of the radially moving atoms. By scanning the frequency of the clock laser, we directly access the Doppler limited linewidth. This measurement is shown in Figure 6.2, where the red data points is without the lattice ramp, and the black data points are with the lattice ramp. The extracted radial temperature is 447(20) nK without lattice ramp, and 291(22) nK with the lattice ramp.

With the atomic temperature nice and cold, we proceed by zooming in on the optical transition. The background magnetic field is nulled to zero at the location of the cavity by performing clock spectroscopy with all 10 ground-state spin states present. When the applied bias field exactly cancels the background magnetic field, all 10 π -transitions will be degenerate. By adjusting the fields in the X, Y, and Z direction and iteratively zooming in on the carrier by probing with longer π -pulses, we achieve a Fourier limited linewidth of 5 Hz. We do not push this further as we do not plan a long interrogation time at the first stage of spin squeezing experiments.

To remove the atoms in hyperfine spin states other than $m_F = -9/2$, we apply a “cleanup pulse”. This involves driving a clock pulse on the magnetic sub-level transition ($m_F = -9/2 \rightarrow m_F = -9/2$), and then blowing away all of the remaining ground-state atoms that were either not promoted to the excited state using a pulse of 461 nm light. With the pure initial excited state, we are finally ready for spectroscopy. In the same bias magnetic field, we can drive the atoms and observe the Rabi flopping between the ground and excited clock states.

For example, we probe the atoms via Rabi flopping with a Rabi frequency of 12 Hz in Figure 6.3a and 111 Hz in Figure 6.3b. We observe dephasing of the Rabi flopping. This is particularly relevant for squeezing experiments where we need to apply several pulses to manipulate the atomic state. To attempt to understand our loss of contrast, we explore the model presented in [132]. In that work, a finite misalignment angle between the clock laser and the optical lattice is considered.

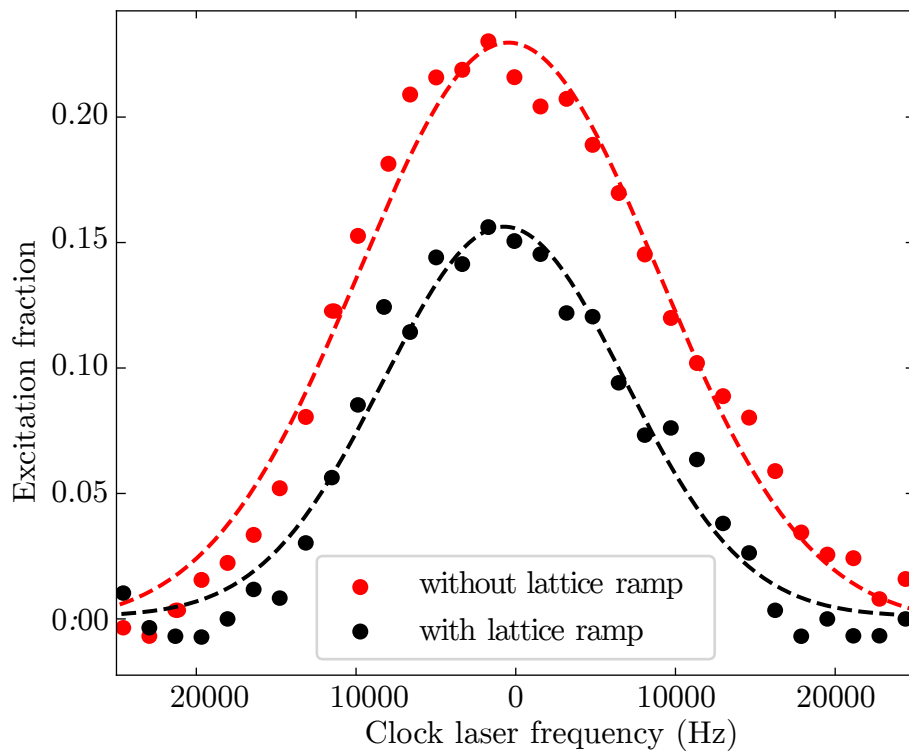


Figure 6.2: Radial Doppler scans of the clock transition. Red data points are without the lattice ramp, black points are with the lattice ramp. The radial temperature extracted from the Gaussian fits (corresponding dashed lines) are $T=290(10)$ nK for with lattice ramp, and $T = 447(20)$ nK without the lattice ramp.




Figure 6.3: Rabi flops (a) Measured excitation versus the Rabi pulse time in seconds, with a drive Rabi frequency of $\Omega_R/2\pi = 111$ Hz. (b) Same as (a), with the drive Rabi frequency set to $\Omega_R/2\pi = 12$ Hz.

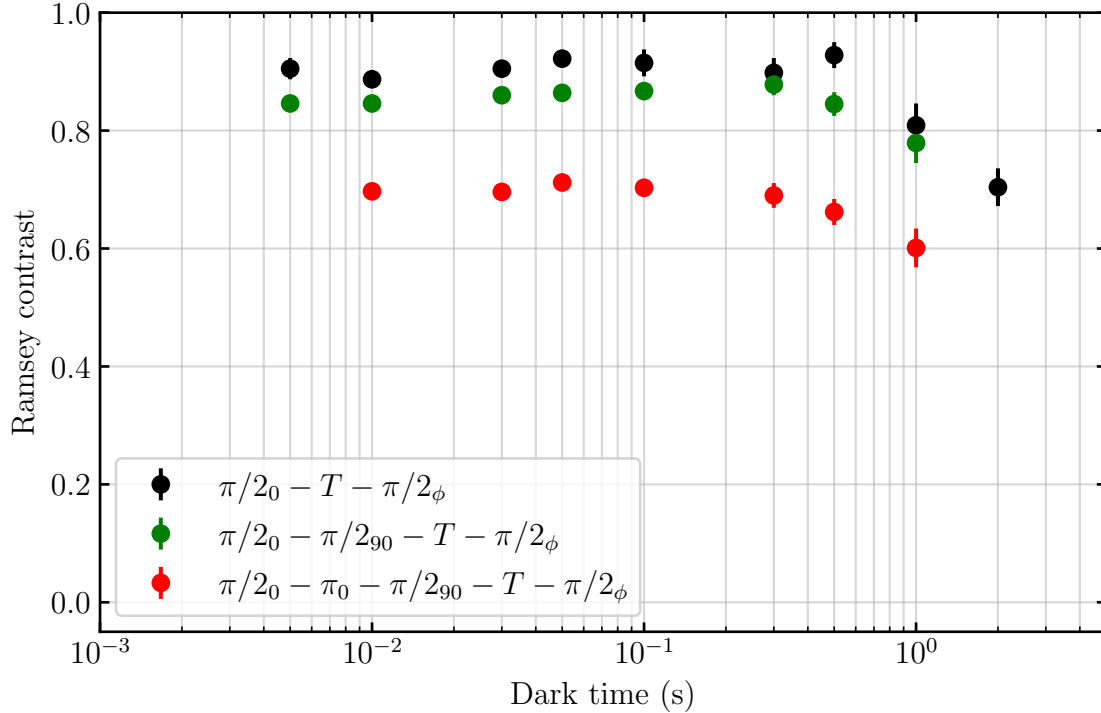


Figure 6.4: Ramsey contrast for increasingly complex pulse sequences. The pulses are labelled with a sub-script, indicating the relative rotation angle. $\pi/2_{90}$ refers to an “align” $\pi/2$ pulse, required for squeezing. The introduction of the π_0 is the π pulse we need to perform the pre-measurement in our squeezing experiments.

Due to this finite angle, there is a projection of the radial motion along the axial clock, modifying and reducing the Lamb-Dicke parameter. We fit our measured data to a model from [132], plugging in our measured radial temperature, and those fits are shown as the solid red lines in Figure 6.3. These fits suggest an overall angle between the clock and lattice of 2.2 mrad.

We also perform Ramsey spectroscopy experiments, and explore the Ramsey contrast versus the dark time T for different pulse sequences. For the simplest Ramsey type experiments, we measure Ramsey contrast of 90 to 95%, as seen by the black points in Figure 6.4. However, for our spin squeezing experiments, we will be required to apply several clock pulses to manipulate the atomic state in the clock manifold. These clock pulses will be separated in time according to the time duration of the population measurements. One can imagine that radial motion will again play an important role in these pulse sequences. As shown in Figure 6.4, we can see by adding

the necessary clock pulses reduce our contrast to the 70% level. This is further reduced by adding dwell time between the $\pi/2_0$ and the π_0 , which is necessary for the QND measurements.

One other crucial takeaway for Ramsey contrast is in relation to our previous discussions of atomic temperature. We have found that the Ramsey contrast can be negatively impacted with hotter atoms. This again suggests that atomic motion is playing an important role in the atomic coherence. The clear way forward is to freeze out the motion by implementing more atomic confinement with a 2D or 3D optical lattice.

6.2 Entering the collective strong-coupling regime

To realize the coupled atom-cavity system, we first need to align the lattice beam to the cavity mode. Most cavity QED experiments trap the atoms in the cavity mode, but we have a more unconstrained geometry without strong trapping in the cavity. Rough alignment is obtained by imaging the atoms through the second mirror ($T \approx 0.5$ at 461 nm) onto a CCD camera. By observing the transmitted cavity mode on the same CCD, rough alignment can be obtained. However, if the cavity mode is not well centered on the imaging lens, there can be small errors in this technique, and thus we look directly for the vacuum Rabi splitting as a measure of successful alignment. To more easily obtain the first signal, we operated with all 10 spin states in a nominally zero magnetic field, so all of the different spin states added together in the vacuum Rabi splitting, and we can take advantage of the higher atom number.

Since we wish to avoid Raman spin flips, we choose to operate on the nominally cycling transition $|^1S_0, F = 9/2, m_F = -9/2\rangle \rightarrow |^3P_1, F = 11/2, m_F = -11/2\rangle$. To couple to this transition, we apply a 2 G bias magnetic field along the cavity optical axis (X direction), and probe the system with σ_- polarization. With the bare cavity tuned on resonance, we measure a typical vacuum Rabi splitting of 750 kHz, as shown in Figure 6.5. Measuring the ground state atom number amounts to measuring the vacuum Rabi splitting $\Omega_\downarrow = 2g\sqrt{N_\downarrow}$. Operating on resonance offers several technical advantages. By simultaneously probing the two vacuum Rabi modes, and taking the difference, one becomes insensitive to relative fluctuations between the laser and the




Figure 6.5: Center : we show the eigenvalue spectrum of the atom-cavity coupled system versus the atom-cavity detuning $\delta_c/2\pi$. Left : When we tune the cavity onto the bare atomic resonance, we observe the vacuum Rabi splitting, shown in the left plot. This demonstrates entering the collective strong-coupling regime. Right : For the QND measurements, we operate in the slightly detuned regime.

cavity. This is because laser-cavity fluctuations will lead to an overall translation of the eigenvalue spectrum in Figure 6.5a. One technical tradeoff for operating in this regime is the signal-to-noise versus the loss of atomic coherence from the probing. This is because atom number fluctuations cause the coupled modes to vary more than $\kappa/2$, which forced us to perform frequency sweeps instead of fixed-tone probing. This could in principle be avoided by implementing an active servo of the probe tones onto the coupled modes, but due to complexity we opted to not pursue that as a first experiment.

Instead, we opt to operate in the slightly detuned regime. There are multiple trade-offs when deciding on an operating detuning. As the operating atom number increases, one can afford to detune the cavity further from the bare atomic resonance, if you want to fix the technical noise floor relative to the QPN. For our experiment we set the atom - cavity detuning to $\delta_c/2\pi = 1$ MHz. We then observe the dispersive shift of the cavity mode due to atoms in the ground-state. At this detuning, the operational cavity frequency shift is anywhere between 100 - 400 kHz depending on the exact atom number. A fixed-frequency laser is then parked on resonance of the cavity-like mode to measure the dispersive shift ω_\downarrow . For our highest atom numbers, the far-detuned approximation for the frequency shift begins to break down, and we need to include relatively significant higher order corrections to the usual g^2/δ_c scaling. We can use the exact expression of the eigenvalues to obtain an expression for the number of atoms in $|\downarrow\rangle$ in terms of ω_\downarrow as

$$N_\downarrow = \omega_\downarrow \frac{\delta_c}{g^2} \left(1 + \frac{\omega_\downarrow}{\delta_c} \right). \quad (6.1)$$

An optical π -pulse is then applied to swap the population between $|\uparrow\rangle$ and $|\downarrow\rangle$. The frequency shift is measured again to determine the excited state population N_\uparrow . To zoom ahead, it is by measuring these frequency shifts in a non-destructive way that we will generate entanglement via conditional spin squeezing. After an initial $\pi/2$, we can make two repeated measurements of J_z , and since the two measurements will contain highly correlated QPN, taking their difference allows us to subtract the common quantum noise [111, 31]. By using that differential information, we will project the atoms into a spin-squeezed-state (SSS).

6.3 QPN and the effective atom-cavity coupling

With the strong-coupling regime in our wheelhouse, we now seek to perform conditional spin squeezing. For that, we will make repeated measurements of the Bloch vector projection along z , and take direct advantage of the correlated QPN in those two measurements. Here, we will make full use of the low-noise cavity probing system we described in Chapter 5.8.2. It is therefore critical to know the level of QPN so we can properly benchmark our QND measurements. It turns out, and we will show derivations below, that the atom-cavity coupling constant g serves as the bridge between QPN in terms of population fluctuations and cavity frequency shift fluctuations. We can see this with a few lines of math and considerations. In the far-detuned limit, the shift of the cavity-like mode is

$$\omega_{\downarrow} = \frac{g^2}{\delta_c} N_{\downarrow} \quad (6.2)$$

The noise from projection noise is $\Delta N_{\downarrow} = \sqrt{N}/2$. Since we infer the atom number from the shift, one can show that the noise of the frequency shift is

$$\Delta\omega_{\downarrow} = g\sqrt{\frac{\omega_{\downarrow}}{2\delta_c}}. \quad (6.3)$$

We can then see in a straightforward way that measuring the QPN using our cavity QED system is directly connected to the atom-cavity coupling g . A stronger atom-cavity coupling means that QPN converts to larger frequency fluctuations, which makes good physical sense. We note that we write g , but we really mean the effective quantity.

6.3.1 QPN in slightly detuned regime

We operate in the slightly detuned regime, where we have set a detuning of $\delta_c/2\pi = 1$ MHz. Even at our lower atom numbers, the shift is 100 kHz or greater. To fit g to our measured noise of the differential cavity frequency shift $\Delta(\omega_{\downarrow} - \omega_{\uparrow})$, we require an expression for this noise in terms of the sum shift $\omega_{sum} = \omega_{\downarrow} + \omega_{\uparrow}$. The eigenvalues for the $|\downarrow\rangle$ and $|\uparrow\rangle$ states are

$$\omega_{\downarrow,\uparrow} = \frac{\delta_c + \sqrt{\delta_c^2 + \Omega_{\downarrow,\uparrow}^2}}{2}, \quad (6.4)$$

where the vacuum Rabi splitting for each spin state are $\Omega_{\downarrow} = 2g\sqrt{N_{\downarrow}}$ and $\Omega_{\uparrow} = 2g\sqrt{N_{\uparrow}}$. The QPN fluctuation of the frequency shift is obtained by propagating the QPN noise ΔN_{\downarrow} through the expression

$$\Delta\omega_{\downarrow} = \left| \frac{d\omega_{\downarrow}}{dN_{\downarrow}} \right| \Delta N_{\downarrow}. \quad (6.5)$$

For a CSS of N total atoms prepared by a $\pi/2$ pulse on the equator of the Bloch sphere with state population $N_{\downarrow} = N_{\uparrow} = N/2$, one can express N in terms of the sum of measured frequency shifts

$$\omega_{sum} = \omega_{\downarrow} + \omega_{\uparrow}$$

$$N = \omega_{sum} \frac{\delta_c}{g^2} \left(1 + \frac{1}{2} \frac{\omega_{sum}}{\delta_c} \right). \quad (6.6)$$

QPN for the two spin states is $\Delta N_{\downarrow} = \Delta N_{\uparrow} = \sqrt{N}/2$. We can calculate the projection noise fluctuations of the frequency shift using the derivative of the eigenvalue expression. Projection noise fluctuations of N_{\downarrow} and N_{\uparrow} are perfectly anti-correlated, and therefore $\Delta(\omega_{\uparrow} - \omega_{\downarrow})$ is twice of the fluctuations of ω_{\downarrow} ,

$$\Delta(\omega_{\uparrow} - \omega_{\downarrow}) = 2\Delta\omega_{\downarrow} = \frac{g^2\sqrt{N}}{\sqrt{\delta_c^2 + \Omega_{\downarrow}^2}}. \quad (6.7)$$

Using Eq. 6.6 and Eq. 6.7, the expression for characterizing g based on the measurement on QPN fluctuations of the initial CSS is

$$\Delta(\omega_{\uparrow} - \omega_{\downarrow}) = g \sqrt{\frac{\omega_{sum}^2/2 + \delta_c\omega_{sum}}{(\omega_{sum} + \delta_c)^2}} \quad (6.8)$$

6.3.2 Measuring QPN

We will now measure the QPN of our ensemble by looking at the fluctuations of $(\omega_{\downarrow} - \omega_{\uparrow})$ versus the total shift $\omega_{sum} = \omega_{\downarrow} + \omega_{\uparrow}$. The pulse sequence for measuring QPN is relatively simple as shown in Figure 6.6a. We perform an initial $\pi/2$ pulse to prepare a CSS on the equator of the Bloch sphere. We measure the shift of the cavity-like mode for the ground state atoms, indicated by ω_{\downarrow} , perform a π pulse to swap populations, and measure ω_{\uparrow} . As we will now go into detail - we need to be careful about the influence of technical noise in the clock rotations or the cavity




Figure 6.6: (a) Pulse sequence for measuring QPN. Black open pulses correspond to clock rotations, and the red pulses are cavity frequency shift measurements. (b) QPN measurement. The red datapoints have no rotation noise subtracted, and the model including a linear term is fit (red line) giving $g/2\pi = 5230(213)$ Hz. The open blue circles have an estimated clock rotation noise contribution subtracted, and is fit without a linear term giving $g/2\pi = 5127(192)$ Hz.

frequency shift measurements, as these can both mimic QPN, and give incorrect normalizations for our conditional spin squeezing.

6.3.3 Noise in QPN measurement

When we perform a measurement of QPN, we have to be aware of systematic or technical effects that can influence our determination of g . There is noise from both rotation noise and the cavity shift measurements. At low atom number where the influence of QPN is small due to the \sqrt{N} scaling, we run into our technical noise floor of the homodyne measurement. When measuring QPN, we can feel free to turn up the optical power so that we can minimize the impact of photon shot noise. Nonetheless, the combination of photon shot noise, backaction on the atoms, and technical laser - cavity noise will all contribute, as seen by the finite offset at $\omega_{\downarrow} + \omega_{\uparrow} = 0$. At higher atom number, as the length of the Bloch vector gets longer, rotation noise can start to become noticeable in the measurements. In principle, the noise can come from amplitude or phase noise of the clock local oscillator.

The pulse sequence for measuring QPN is relatively simple, as shown in Figure 6.6a. Let's first consider amplitude noise during the clock pulses in the QPN measurement sequence. We need to consider the noise introduced in the $\pi/2$ pulse and the π pulse, since this added noise would add in quadrature with the technical measurement noise and the true QPN. Let's consider shot-to-shot fluctuations ϵ in the clock drive Rabi frequency, and assume it is correlated between the $\pi/2$ and π pulse for each shot of the experiment. If we tune the Rabi frequency to Ω_0 to nominally give a $\pi/2$ pulse with a pulse time of $t_{\pi/2}$, we can write the modified Rabi frequency as

$$\Omega_1 = \Omega_0(1 + \epsilon) \tag{6.9}$$

A Bloch vector undergoing evolution with a Rabi frequency Ω_1 gives a Bloch vector angle of $\theta_1 = \pi/2 + \epsilon\pi/2$. Since we are treating the ϵ error as correlated, the π -pulse will result in a further angular deviation of $\pi\epsilon$, and so the two pulses together give the angular error of this pulse sequence

is

$$\Delta\theta = \frac{3\epsilon\pi}{2} \quad (6.10)$$

This can be understood intuitively since we are treating the error as correlated, so the total pulse area of the combined $\pi/2 + \pi$ pulse is $3\pi/2$.

Now we consider the pulse sequence we utilize to benchmark rotation noise, which is $\pi/2 - \omega_{\downarrow} - 2\pi - \omega_{\downarrow}$. Since we always take the difference between the two ω_{\downarrow} measurements, and the initial $\pi/2$ pulse is common to both measurements, this pulse sequence is sensitive only to the noise of the 2π pulse. Again, given the shot-to-shot error in the Rabi frequency of ϵ , the angular error of this 2π pulse is

$$\Delta\omega = 2\pi\epsilon \quad (6.11)$$

Thus, to estimate the amount of rotation noise present in the QPN pulse sequence, we multiply the benchmarked noise by the ratio (1.5/2) and subtract that in quadrature (taking the linear atom number dependence into account).

We also take another approach where we allow for a linear term in our non-linear fit, and see how large that term is, and how it impacts the extraction of g_{eff} . This is potentially an appealing approach, since the rotation noise subtraction can be sensitive to your assumed model of the rotation noise (amplitude versus phase, correlated versus uncorrelated noise). By allowing for a linear term in the model, we allow for a generic classical rotation noise source. Fitting for g this way gives $g/2\pi = 5230(213)$ Hz, which is in good agreement with the method of subtracting the measured rotation noise. We treat this as a consistency check, since the rotation noise term is not perfectly linear in the shift. Our final result is taken as the value obtained by the subtraction method, giving $g = 2\pi \times 5.1(2)$ kHz as shown in Figure 6.6b, and the bare cavity noise offset of $2\pi \times 0.76(5)$ kHz.

6.4 Conditional spin squeezing

6.4.1 Spectroscopic enhancement

The first method of generating useful spin-squeezing we pursued is conditional spin squeezing. This involves making two repeated measurements of $J_z = (N_\downarrow - N_\uparrow)/2$. In principle, we can make this differential measurement with precision well below the level of QPN [34], and how far below QPN is referred to as spin noise reduction R ,

$$R = \left(\frac{\Delta(J_{z,f} - \beta J_{z,p})}{\Delta J_{z,QPN}} \right)^2, \quad (6.12)$$

where the numerator is the difference between the pre and final measurement, with β acting as a Bayesian optimal estimator. The β deweights the pre-measurement, and will typically take on a value from 0 - 1, depending on the strength of the pre-measurement. An important concept to remember is that we have prior measurement about the state we prepare ($J_z \approx 0$), and in principle the noise of that prior distribution is set by the QPN. If the noise of the pre-measurement was higher than the QPN level, and you imagined getting a result for J_z that was far from the expected value near the equator of $J_z = 0$, you would clearly want to deweight that measurement, since you know it is exceedingly unlikely that the actual state is far from $J_z = 0$ (ignoring real outliers due to issues on the experiment). A further reason for a non-unity β is the fact that the pre-measurement is taken at a lower probe power than the final measurement due to the backaction. One therefore should deweight the pre-measurement accordingly.

To certify that the measured spin noise reduction arises from entanglement, one must weigh R against the loss of coherence induced by the probe as measured by the Ramsey fringe contrast. In other words - if the information leaks out to the universe via free space scattering into 4π steradian, then the system loses its entanglement.

The Wineland parameter, ξ , serves as a criterion for the generation of entanglement, and was originally expressed as [108]

$$\xi = R \frac{1}{C_f^2}, \quad (6.13)$$

where we weigh the spin noise reduction below QPN against the loss of signal induced by the squeezing operation. This formulation of the Wineland parameter assumes an initial contrast of $C_i = 1$. We further note that we have defined the QPN assuming unity contrast (all of the atoms contribute to the cavity signal and corresponding noise), $\Delta J_{z,QPN} = \sqrt{N}/2$. For our experiment, we want to express the Wineland parameter taking into account an initial contrast $C_i < 1$. When we have $C_i < 1$, the initial length of the Bloch vector is shorter, making the $\Delta\theta_{SQL}$ worse by a factor of $\sqrt{C_i}$. The Wineland parameter may be rewritten with the initial contrast taken into account as

$$\xi = \left(\frac{\Delta(J_{z,f} - \beta J_{z,p})}{\Delta J_{z,QPN}} \right)^2 \frac{C_i}{C_f^2} \quad (6.14)$$

$$= R \frac{C_i}{C_f^2}, \quad (6.15)$$

where C_i and C_f are the initial and final Ramsey fringe contrast, respectively. Another equivalent way to think about ξ is by looking at the measured phase resolution relative to the SQL limit of phase resolution,

$$\xi = \left(\frac{\Delta\theta_m}{\Delta\theta_{SQL}} \right)^2 \quad (6.16)$$

where $\Delta\theta_m$ is the measured phase resolution and $\Delta\theta_{SQL} = 1/\sqrt{C_i N}$ with C_i being the initial Ramsey fringe contrast. This formulation has the advantage that it shows us exactly what we need to measure in our experiment, which is the phase resolution of our state. The phase resolution can then be easily mapped to Ramsey spectroscopy and how one gains enhancement of the clock performance. Further, it is this relative phase resolution that intuitively shows why ξ is a witness for entanglement.

6.4.2 Calibrating incident photon number

Here we include a short aside on how we calibrate and estimate the number of incident photons on the cavity for a given QND measurement. While the practical method of benchmarking the QND strength is by the loss of Ramsey fringe contrast, determining the incident photon number will allow us to compare our measured spin noise reduction and contrast loss against theoretical

predictions. To get our initial bearings, we first want to independently determine the optical power in the sideband. For this we turn up the EOM drive such that we put at least 10% of the optical power into the first order sideband. We can then measure the total optical power in the beam incident on the cavity (in front of the viewport just before the cavity, loss in the viewport is on the % level) with a power meter. By observing the relative size of the reflection dip on the cavity mode for the carrier and the sideband, we can reliably determine the total amount of sideband optical power incident on the cavity in units of Watts, labelled as $P_{inc,tot,cav}$. We label with “tot” because this includes the part of the beam that is not actually modematched to the cavity. So we know how much optical power is in that sideband, and we can track the optical loss from the incident viewport on the chamber to the homodyne photodiode, and the fraction of power that makes it to the photodetector is the path efficiency. We then know how much sideband optical power is incident on the photodiode, and we label this as $P_{inc,PD}$. The power in the LO incident on the PD is labelled as P_{LO} . The expected peak-to-peak voltage assuming unity mode matching between the reflected beam off of the cavity mirror and the LO is

$$V_{p2p,ideal} = 4S_{PD}G\sqrt{P_{LO}P_{inc,tot,PD}} \quad (6.17)$$

where S_{PD} is the sensitivity of the photodiode ($S_{PD} = 0.48$ A/W), G is the transimpedance gain of the photodetector ($G = 1$ MOhm). However, the finite mode matching between the reflected beam and the LO leads to a smaller observed peak-to-peak $V_{p2p,obs}$. The ratio $F = V_{p2p,obs}/V_{p2p,ideal}$ depends on the quality of various mode-matching and alignment, and can typically be 0.6. To be explicit, we rewrite the signal as

$$V_{p2p,obs} = FV_{p2p,ideal} = 4FS_{PD}G\sqrt{P_{LO}P_{inc,tot,PD}} \quad (6.18)$$

and use this expression to infer $P_{inc,tot,PD}$. Using the path efficiency we can determine the incident power on the cavity, but only a fraction of that incident beam is actually mode matched to the TEM00 mode. We measure this mode matching by back-propagating a beam through the cavity, and observe the back-coupling efficiency through the launch fiber. In the experiments shown in

this thesis, the mode matching of the incident beam to the TEM00 mode was 0.75. With those values, we can deduce the amount of optical power incident on the cavity in the sideband that is actually mode matched to the cavity mode. This allows us to map measured homodyne signal sizes to incident photon number, and use them to compare against predictions of photon scattering into free space.

6.4.3 Spin noise reduction

We now demonstrate repeated measurements of J_z with differential resolution well below QPN. After an initial $\pi/2$, we make a pre-measurement denoted by $J_{z,p}$, wait a dwell time of 20 ms, and perform the final measurement $J_{z,f}$ (Figure 6.7a). While the pre-measurement fluctuates with standard deviation $\Delta J_{z,QPN} = \sqrt{N}/2$, the final measurement contains highly correlated QPN, as seen by the strong correlations between the two measurements (Figure 6.7). We turn up the photon number for the measurement of $J_{z,f}$ to optimize the signal to noise of that final readout measurement. The effect of β is to de-weight the pre-measurement so that we optimize the resulting standard deviation of the difference. We directly observe spin-noise reduction $R = -4.6(6)$ dB relative to QPN at the optimal squeezing photon number of 2.3×10^4 per measurement window.

We can ask the question - after taking into account the technical noise of the final-measurement, what is the full spin noise reduction R that we can infer? To answer this, we assume that the readout noise in the final measurement is uncorrelated with the pre-measurement. Due to our measurement noise roughly matching our expected photon shot noise (see Figure 6.11), this is a somewhat reasonable assumption to make. We then calculate the photon shot noise present in the final two measurements, and subtract this off from our measured spin noise reduction. For example, for the spin noise reduction shown in Figure 6.8, both measurements are taken at the highest probe power of 4.2×10^4 photons per measurement. In this case the final measurement would contribute the same amount of noise as the pre-measurement, and the inferred spin noise reduction is $R = -8.5$ dB. From a practical perspective, this inferred R is not what we could make use of in a real metrological measurement. However, it tells us how much we could improve our conditional squeezing if

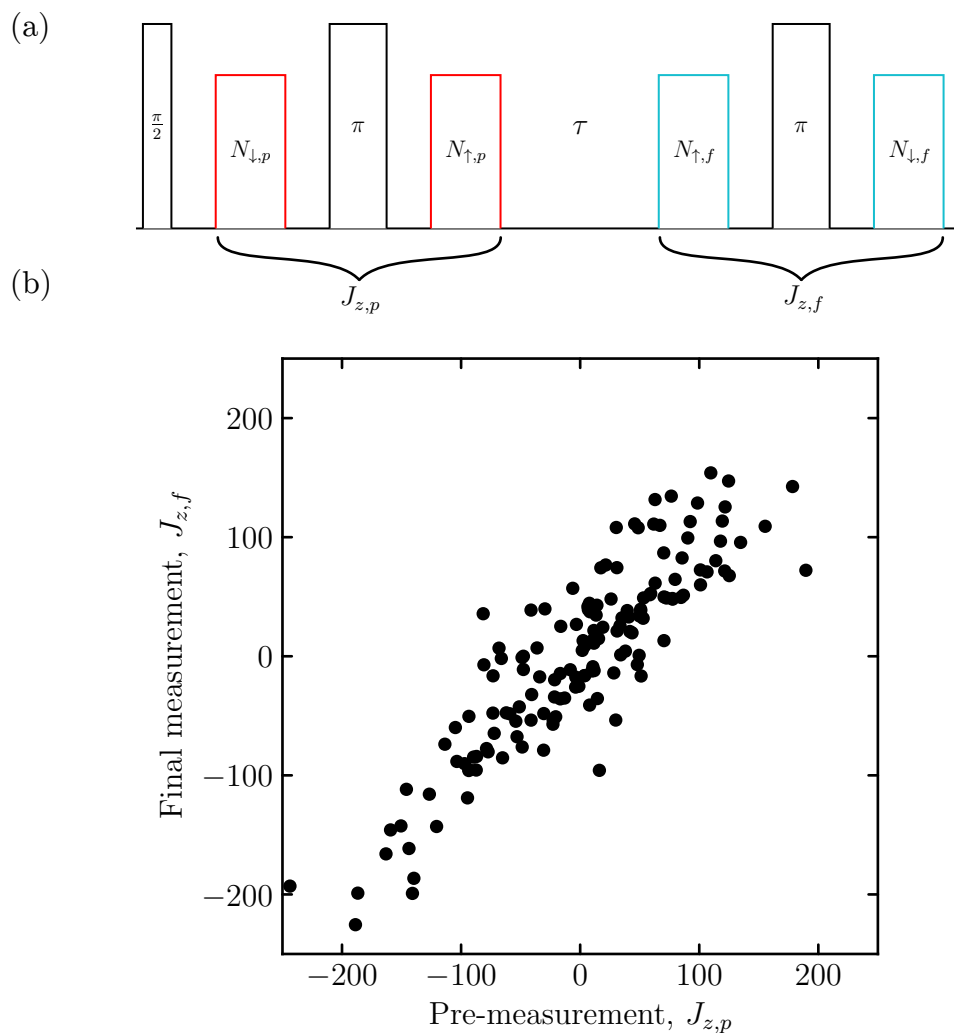


Figure 6.7: (a) Pulse sequence used for measuring spin noise reduction. After an initial $\pi/2$, we measure $J_{z,p}$, wait a dwell time τ , and then measure $J_{z,f}$. (b) The high-degree of correlations are shown between $J_{z,p}$ and $J_{z,f}$.




Figure 6.8: The left panel shows the pre-measurement (red points) and the final-measurement (open cyan points). The incident photon number per measurement is 4.2×10^4 . The optimal subtraction $J_{z,f} - \beta J_{z,p}$ is shown in the right panel. The dashed lines indicate the calculated QPN, and the solid lines indicate the calculated standard deviation of the measured data. The spin noise reduction here is $R = -6.3(4)$ dB.

we could turn up our probe power arbitrarily high to make the final measurement noise completely negligible. At the moment, we observe an increase of the noise at even higher power - indicating extra noise from the light action on the atoms.

6.4.4 Ramsey contrast

For our conditional spin squeezing we measure the initial Ramsey contrast C_i without QND probing using the pulse sequence shown in Figure 6.9a. We then add QND probing and measure the final contrast C_f using the sequence in Figure 6.9b. To map out the fringe, we vary the phase of the final $\pi/2$ pulse and record the excitation fraction which is defined as

$$f_e = \frac{N_\uparrow}{N_\uparrow + N_\downarrow}. \quad (6.19)$$

Examples for the measured Ramsey fringe is shown in Figure 6.9c. The initial contrast with no QND probing is the open green squares. We note that our initial contrast is degraded primarily by the optical manipulations of the π pulse and the $\pi/2$ align pulse. When we add the QND probe,

we see the contrast come down accordingly (black and red points). We note that we do not see an appreciable phase shift of the Ramsey fringe. This leads one to believe that we have sufficiently echoed away the ac Stark shift induced by the QND probe. The maximum phase shift that we observe is 50 mrad.

We plot the measured fractional loss of contrast as C_f/C_i versus the incident photon number per population measurement as the black points in Figure 6.10. The expected loss of coherence from free-space scattering is

$$C_f = C_i e^{-M_s/N} \quad (6.20)$$

where M_s is the number of scattered photons. We estimate the expected number of scattered photons by solving for the steady-state atomic excitation $|\langle b \rangle|^2$. The rate of scattered photons is then $\dot{M}_s = \Gamma |\langle b \rangle|^2$ (see Appendix C for more details). We clearly observe more contrast loss than our free space scattering estimate. In principle this could be explained by a miscalibration of the photon number, but the discrepancy seems too large to be explained by this. It is also possible that there is a second-order in M_s effect playing a role to the contrast loss. We need to further explore our loss of contrast - and whether it depends on any other parameters in the system, like the duration of the measurement window, π -pulse time for the clock rotations, etc. Another possible excess loss of contrast could come from additional inhomogeneous effects induced on the atoms due to the QND probe. Better atomic confinement could lead to improved Ramsey contrast in both the initial state and the final probed atoms.

6.4.5 Measured Wineland parameter

We systematically vary the number of photons incident on the cavity per measurement window, and measure the spin noise reduction R and the contrast loss C_i/C_f^2 (Figure 6.11). For example, turning on a QND probe with 2.3×10^4 photons per population measurement reduces the contrast to $C_f = 0.60(1)$. At this optimal probe power, the measured Wineland parameter reaches -1.7(7) dB (Figure 6.11), a direct verification of spin entanglement in our system. If we want to infer the amount of entanglement in the system before the final measurement, we need to calculate

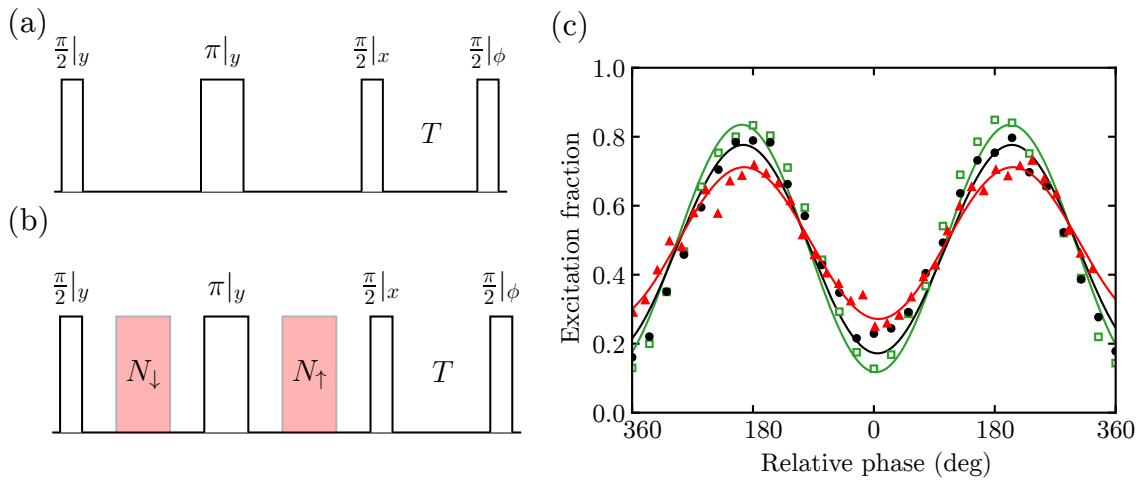


Figure 6.9: (a) Pulse sequence used for measuring the initial contrast with no QND probing applied. (b) Pulse sequence for measuring final contrast C_f with the QND probing. (c) Open green squares is the Ramsey fringe with no probing, black points is with 2.3×10^4 photons per measurement, and red triangles is with 4.2×10^4 photons per measurement. We do not observe any significant phase shift due to the probing.




Figure 6.10: Measured fractional contrast C_f/C_i versus the incident photon number per measurement. The solid line indicates the calculated expected loss of contrast from free-space scattering, described in Appendix C.

the noise of the final measurement to subtract in quadrature (assuming uncorrelated detection noise). Using the model given by the cyan line in Figure 6.11, we subtract the final measurement noise, and infer a Wineland parameter of $-3.7(7)$ dB. This inferred value is not the useful amount of entanglement, as real clock operation relies on the final measurement. It is insightful however, since we can infer how much more squeezing we could expect if we could improve the cavity shift noise at higher photon number.

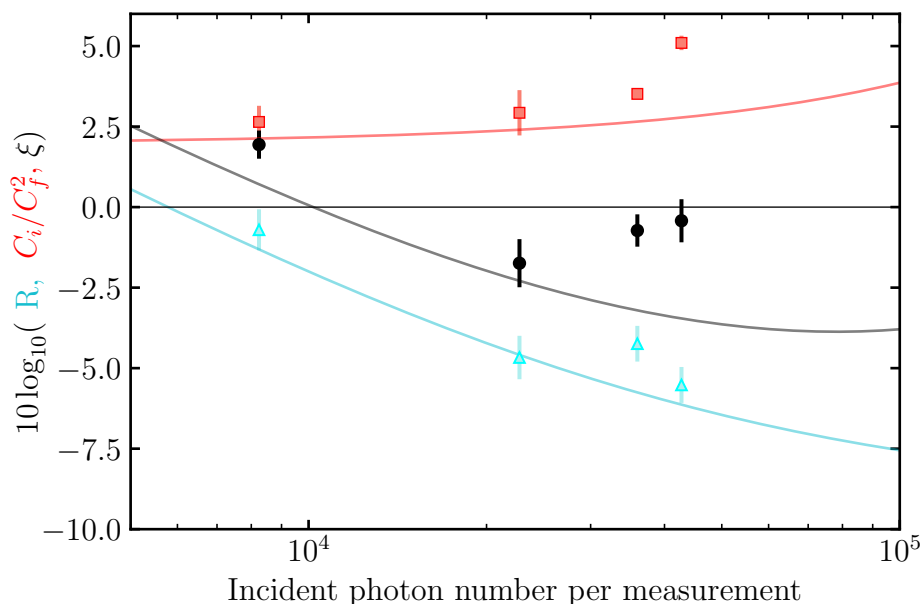


Figure 6.11: Relative spin noise reduction R (cyan circles), fractional contrast loss C_i/C_f^2 (red circles), and the corresponding Wineland parameter ξ (black circles). At the optimal photon number, we directly measure a Wineland parameter of $-1.7(7)$ dB. Expected R given our estimated quantum efficiency of $Q = 0.28$ (cyan line), expected contrast loss via free-space scattering (red line), and the corresponding expected Wineland parameter (black line, see Methods).

Chapter 7

Squeezed clock comparisons

To quickly recap, we have built a new apparatus that is capable of state-of-the-art clock operation with cQED. We then demonstrated that we can generate entanglement on the Sr clock transition. The next question is - can we use this spin squeezing for a real clock measurement, and directly observe an enhancement beyond the QPN limit?

In this Chapter we describe our squeezed clock comparisons that we perform in order to demonstrate the direct advantage of implementing spin squeezing. We perform a differential clock measurement, where we compare two sub-ensembles within the same atomic cloud. Original demonstrations of this technique by the Sr2 team made use of a high resolution imaging system in order to divide their Fermi-degenerate gas into two sub-ensembles [28]. Instead of using imaging, we use the selective nature of the cavity probing to sub-divide our ensemble. By using the movable optical lattice, we can shuffle different parts of the cloud in and out of the cavity mode on demand. In Section 7.1, we first define the QPN and SQL limit to a clock comparison with initial contrast $C_i < 1$. Section 7.2 will provide a description of the differential clock measurements, and present our results for both a classical clock and squeezed clock comparison. In the later sections, we highlight various details that provide further experimental measurements that are relevant to the spin squeezing observation.

7.1 QPN and SQL

Neutral atom optical clocks probe many atoms to obtain the best signal-to-noise ratio. The fundamental noise source in atomic clocks is the quantum projection noise that stems from the inherent population fluctuations associated with the projective measurement of N uncorrelated atoms [20]. Operating at higher atom number is clearly advantageous from a signal-to-noise point of view. However, technical noise from imperfect state readout, intrinsic atom-atom interactions, or aliased frequency noise of the interrogating clock laser pose challenges for observing clock performance at the QPN limit [133]. With precise engineering of quantum states and control of atomic interactions [134], and using laser-noise mitigation techniques such as synchronous comparisons [28, 17], state-of-the-art optical clocks are currently approaching QPN-limited stability with up to 10^5 atoms [19]. It is an exciting prospect to imagine pushing these state-of-the-art clocks past the QPN-limit, making use of quantum correlations to improve precision measurements.

We want to establish a clear definition for the standard quantum limit for a clock with a finite initial contrast $C_i < 1$. The primary loss of atomic coherence in our squeezed clock experiments comes from the multiple optical rotations needed in our sequence. Conceptually, we treat the system as having two components - a pure state with $|J_z| = C_i N/2$, and treating other $(1 - C_i)$ atoms as no longer participating in the same *CSS*. We thus need to calculate the standard quantum limit (SQL) for the C_i atoms that participate in the *CSS*. In this case, we have

$$\Delta\theta_{\text{SQL}} = \frac{\sqrt{C_i N}/2}{C_i N/2} = \frac{1}{\sqrt{C_i N}} \quad (7.1)$$

We find it useful to define another limit, which we call the ‘‘QPN’’ limit. In our situation, the sequence of optical rotations leads to dephasing, but the atoms are not lost from the trap. The dephased atoms still couple to the cavity mode, and will still contribute quantum projection noise. Thus, the QPN limit is the ratio of the projection noise for all of the atoms divided by the shortened Bloch vector,

$$\Delta\theta_{\text{QPN}} = \frac{\sqrt{N}/2}{C_i N/2} = \frac{1}{C_i \sqrt{N}} \quad (7.2)$$

We note here that operation of a clock with phase resolution below $\Delta\theta_{QPN}$ is a real enhancement, in that no classically operating clock could go below this level. However, going below $\Delta\theta_{SQL}$ goes further as a metric for entanglement in the clock.

7.2 Self-synchronous clock comparison with cQED

The timing sequence for the differential clock comparison includes clock rotations, transports, and QND measurements (Figure 7.1a). Clock rotations are shown as black pulses, with each pulse labelled by the pulse area and axis of rotation. The atomic cloud is moved vertically by detuning one of the lattice beams, and each transport operation is indicated by either a green or purple pulse. The waveforms for the optical lattice detuning are linear ramps of the frequency over a total of 5 ms, as shown below the pulse sequence. All clock pulses are applied with the lattice at the same vertical location, so we do not have to take into account the varying clock laser phase. After the pre-measurements, the $\frac{\pi}{2}|_x$ clock pulse rotates the spin-squeezed axis to the phase-sensitive axis. Our Ramsey dark evolution time is $T = 14$ ms, and the final $\pi/2$ pulse converts an accumulated phase to a tipping of the Bloch vector in the z direction. To minimize the impact of photon shot noise, the final measurements use a higher probe photon number of $\approx 4.5 \times 10^4$ per measurement window. This is safe since the final measurement is outside of the Ramsey interferometer, and a shortening of the Bloch vector at this point is harmless. The CSS - CSS comparison uses the exact same pulse sequence, with the only exception that we do not apply any probe light during the pre-measurement windows. Running the two comparisons with the same pulse sequence allows us to see and understand the role of spin-squeezing.

The relative separation between the two ensembles is set such that they give the same shift of $(\omega_\downarrow + \omega_\uparrow) \approx 2\pi \times 215$ kHz, corresponding to ≈ 8500 atoms per independent ensemble. We first measure the Ramsey fringe contrast for each ensemble, shown in Figure 7.1b. We plot $J_z = (N_\downarrow - N_\uparrow)/2$ versus the phase ϕ of the final $\pi/2$ pulse. Red is ensemble A and blue is ensemble B. The open circles correspond to running the sequence with no probe light applied during the pre-measurement windows. The solid circles are the Ramsey fringes with probing. The loss of contrast

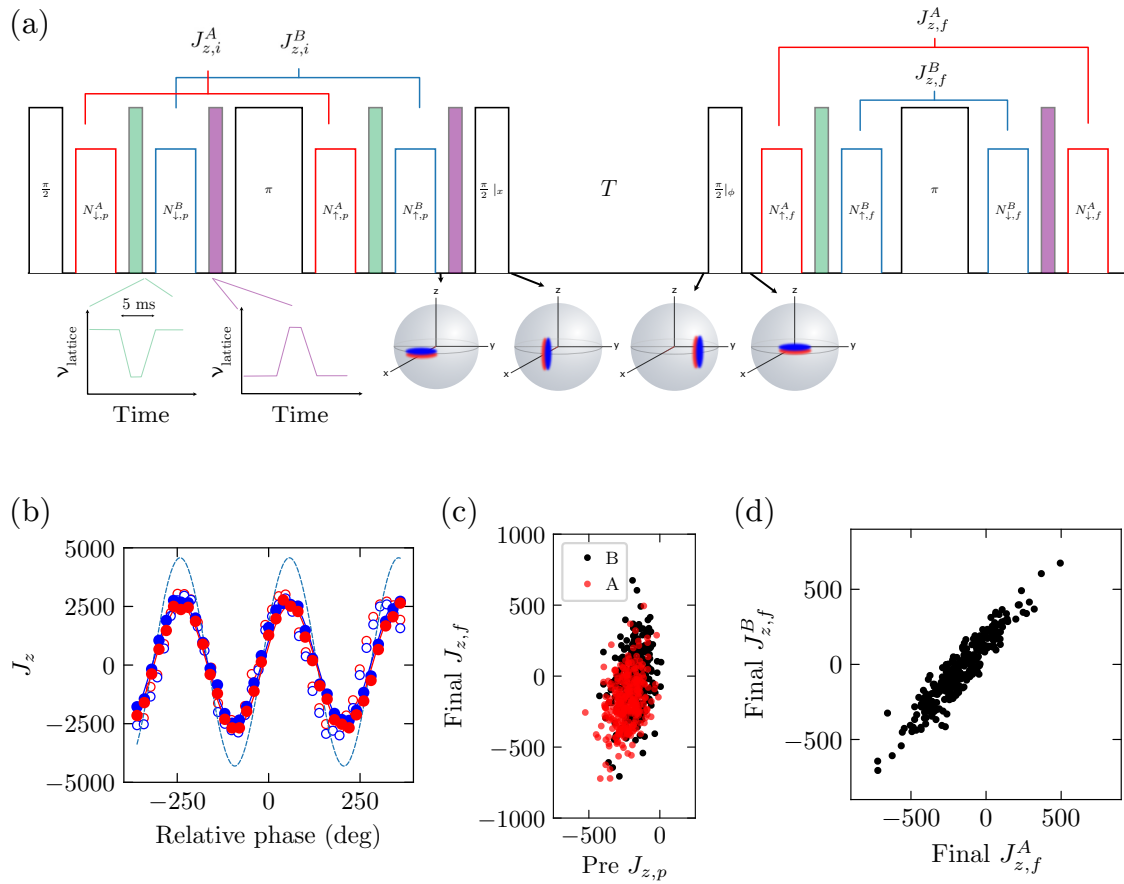


Figure 7.1: (a) Clock pulses are the black pulses, measurements of ensemble A are the red pulses, and the transports are shown as the green and purple pulses. The Bloch spheres depict the spin state distribution at various points during the sequence. (b) Ramsey fringe measured by varying the phase of the final $\pi/2$ pulse. (c) Pre and final measurements of ensemble A and B. (d) The final measurements of ensemble A and B show strong correlations, allowing for the subtraction of the common-mode laser phase noise.

at our operating probe photon number is $\approx 5\%$. An ideal Ramsey fringe with unity contrast is shown as the dotted curve, reaching the maximal values of $J_z = \pm N/2$. The initial C_i and final contrast C_f are given by the ratio of the measured fringe amplitude and the ideal Ramsey fringe amplitude.

For each cycle of the experiment we simply observe the final tipping of each ensemble's Bloch vector. We do not apply feedback to the laser to steer onto the resonance. The phase of the final $\pi/2$ pulse is set such that the final Bloch vector for each ensemble is near the equator of the Bloch sphere, corresponding to $J_z \approx 0$. This puts the Bloch vectors at the point of maximal phase sensitivity for Ramsey spectroscopy. The raw pre and final measurements are plotted for each ensemble in Figure 7.1c (red is ensemble A and blue is ensemble B). One can see the correlations by eye, and the correlation coefficient between the pre and final measurement is ≈ 0.3 . However, it is clear that the laser noise is playing a significant role, as the correlation between the pre and final measurement is far worse than what we see for the spin noise reduction measurements (Figure 6.7). This is due to the addition of the Ramsey pulse sequence. We can look at the final measurements for A and B (Figure 7.1d), which clearly shows rejection of the common-mode laser frequency noise. The correlation coefficient between $J_{z,f}^A$ and $J_{z,f}^B$ for both classical and spin squeezed comparisons is ≈ 0.9 . In Section 7.5 we measure the degree of common mode rejection, and estimate the residual laser noise to contribute a noise floor 30 dB below QPN.

For each experimental shot, we need to combine the measurements to give the best estimate of the differential phase between the two ensembles. To be explicit, we will use the measured amplitude of the Ramsey fringe α (Figure 7.1b) to convert J_z measurements to phase. Assuming the small angle limit, this conversion is expressed as

$$\theta_A - \theta_B = \frac{(J_{z,f}^A - \beta_A J_{z,p}^A) - \beta_D (J_{z,f}^B - \beta_B J_{z,p}^B)}{\alpha} \quad (7.3)$$

Note that we use θ , because the final Ramsey $\pi/2$ pulse converts phase evolution ϕ to tipping in the polar angle θ . The phase difference can then be converted to frequency units using the simple relationship between phase and frequency $\Delta\theta = \Delta\omega T = 2\pi\Delta\nu T$.

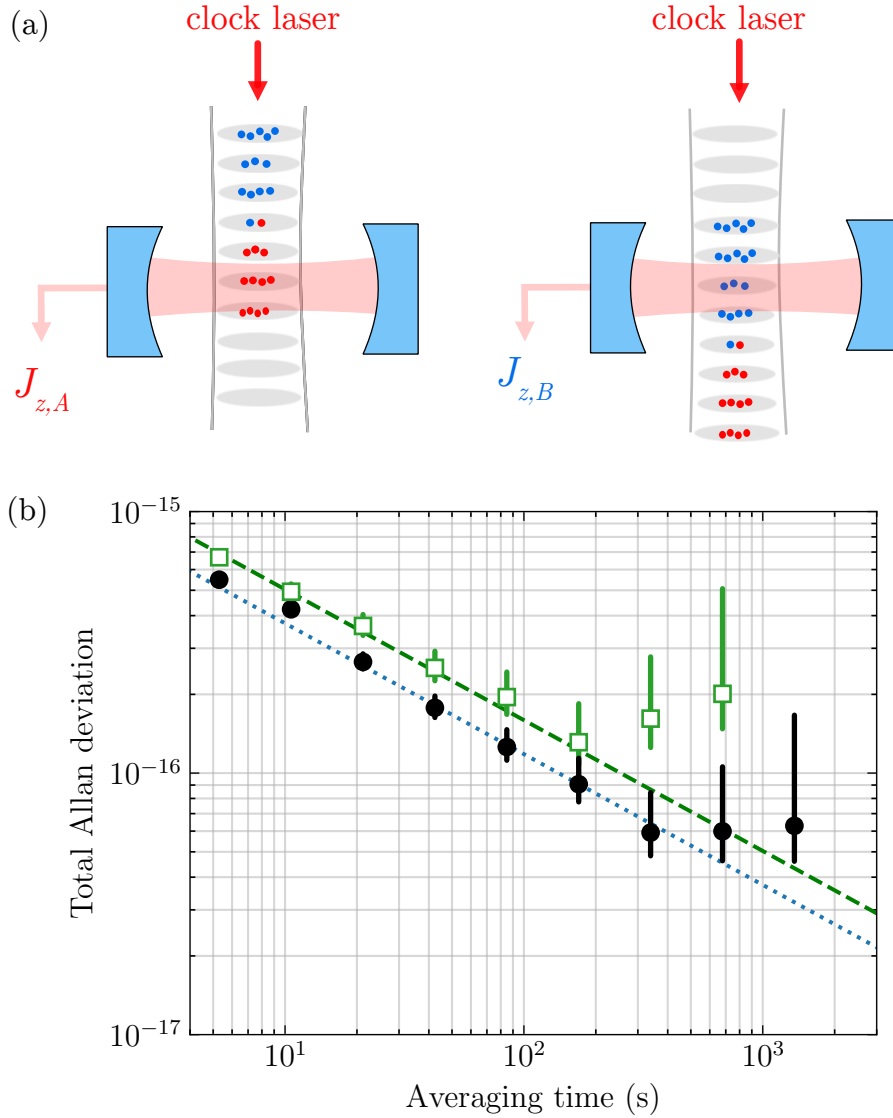


Figure 7.2: (a) Schematic showing the transported atomic cloud for squeezing/readout of each ensemble. The clock laser propagates vertically, addressing both ensembles in a synchronous fashion. (b) Fractional frequency stability given by the total Allan deviation for the CSS - CSS comparison (green open squares) and the SSS - SSS comparison (black circles). The green dashed line is the estimated QPN limit, and the dotted blue line is the SQL. We note that we plot the observed comparison stability, so we have not divided by $\sqrt{2}$ to estimate single clock stability.

We first compare classical clocks (CSS - CSS), where we run the pulse sequence without QND probing. Making use of the eigenvalue expression for the atom-cavity system, we convert these frequency differences to atom number, and then combine the measurements to attain the corresponding J_z . To convert measured J_z for each ensemble directly to phase, we use the amplitude of the Ramsey fringes (Figure 7.1b). The SSS - SSS comparison uses the same pulse sequence, with the QND probe pulses applied.

We determine the stability for each comparison by fitting the first 5 points of the Allan deviations plotted in Figure 7.2. The observed stability is $1.58(3) \times 10^{-15} \tau^{-1/2}$ and $1.25(2) \times 10^{-15} \tau^{-1/2}$ for the CSS - CSS and SSS - SSS comparison, respectively. We directly observe an enhancement of stability by 2.0(2) dB in the SSS - SSS comparison over that of CSS - CSS. We observe a deviation from $1/\sqrt{\tau}$ averaging at longer averaging times. This arises due to a drift in the atom number which changes the conversion from measured homodyne voltage to cavity frequency shift. In future experiments we will be able to eliminate this source of flicker noise by tracking the change of the homodyne fringe slope. For this reason, we fit the Allan deviation for the first four data points. The deviation is likely due to a drift in the slope of the homodyne signal, as the atom number slowly drifts.

To put our results in the proper context, we seek to benchmark the observed SSS-SSS stability to both the practically achievable limit set by QPN and fundamental limit set by SQL. The measurement of $(J_{z,A} - J_{z,B})$ is limited by the quadrature sum of QPN arising from each sub-ensemble. We first define the fundamental limits in the absence of any scale factor between the ensembles ($\beta_D = 1$). The QPN-limited phase estimation between the ensembles is (assuming the same initial contrast C_i for each ensemble)

$$\Delta\theta_{\text{QPN}} = \frac{1}{C_i \sqrt{N_A + N_B}} \quad (7.4)$$

using C_i as the slope of the Ramsey fringe to convert QPN into fractional frequency noise. However, the ultimate bound on the performance of an unentangled ensemble is the SQL. This strictest bound treats the $(1 - C_i)$ fraction of atoms as no longer participating in the pure CSS, thus reducing the

magnitude of the projection noise, resulting in the SQL-limited phase sensitivity of

$$\Delta\theta_{\text{SQL}} = \frac{1}{\sqrt{C_i(N_A + N_B)}} \quad (7.5)$$

The above limits must be adjusted for the optimal estimator β_D . For example, the SQL for the clock comparison is

$$\Delta(\phi_A - \phi_B)_{\text{SQL}} = \sqrt{\frac{1}{C_i N_A} + \beta_D^2 \frac{1}{C_i N_B}}. \quad (7.6)$$

The observed stability of the CSS - CSS comparison is consistent with the QPN-limited stability, and 2.6(3) dB above SQL. Implementing QND-based squeezing operation and accounting for the final contrast of $C_f = 0.50(1)$, the SSS - SSS comparison shows a 2.0(3) dB gain over the QPN-limited stability, demonstrating practical metrological enhancement from the squeezing operation. This result is above SQL by 0.6(3) dB, indicating a Wineland parameter near unity. We therefore can not

This is the first direct observation of an operating optical clock below the QPN limit. The direct observation of the clock comparison below the QPN limit with measurement precision averaging down to the 10^{-17} level is a crucial step towards improving the performance of the best optical lattice clocks via entanglement. Such improvement in differential clock stability translates directly into increased sensitivity for many applications of interest, including the measurement of the gravitational redshift at ever-shorter length scale and the future development of clock networks for fundamental physics [135, 113, 136, 44]. Further improvement will come from enhanced control of atomic motion, which will yield improved spin rotation fidelity and increased coupling to the cavity mode. Larger atom number, improved single-atom cooperativity, and better quantum efficiency will all lead to stronger spin squeezing.

7.3 Optimal estimators

In this Section we will discuss the optimal estimators used in the differential clock comparison in more detail. To attain the ideal clock comparison given the measured quantities, we have three optimal estimators in the estimated phase difference $\phi_A - \phi_B$ in Equation 7.3. The optimal

estimators ϕ_A and ϕ_B represent the weighting of the pre-measurement. Due to the finite signal-to-noise ratio, the pre-measurement contains imperfect information about the collective spin state. This means that the optimal value of β_A and β_B are less than unity. A more subtle optimal estimator for our comparison is β_D . During a particular measurement, we can look at the individual noise of each ensemble on their own $\Delta(J_{z,A,f})$ and $\Delta(J_{z,B,f})$. For example, in the data presented in this chapter, ensemble B was noisier. We found that the relative noise between the two ensembles was not exactly unity. This led us to include an additional optimal estimator β_D to account for the differential technical noise.

The simplest analysis would be to simply simultaneously optimize the three estimators over the whole dataset. We then asked several questions regarding these β values. What is the uncertainty of these optimal estimators; and how does it propagate to the final squeezing result? Furthermore, how does the length of the dataset influence these parameters. To answer these questions, we temporarily break the data up into a “training set” and a “clock run”. We increase the length of the training set to see how the optimal estimators evolve and eventually converge in Figure 7.3. Taking the second half of the data, we calculate the mean values (dashed line) and plot the standard deviation (shaded region). The values are $\beta_A = 0.49(1)$, $\beta_B = 0.48(1)$ and $\beta_D = 0.907(5)$. We take these values to be the optimal estimators for analyzing the whole clock run. To propagate the uncertainty of the optimal estimators, we use a simple numerical method. We pull from normal distributions for $\beta_A, \beta_B, \beta_D$ and g_{eff} using the 1σ standard deviation for each quantity. For each iteration, we calculate the phase resolution compared to the QPN and the SQL limits, and plot these results as histograms.

7.4 Vertical extent of cloud

In the differential clock measurements, we squeeze and compare two ensembles within the same atomic cloud. The final Z position of the atomic cloud can be precisely set by adjusting the transport waveform. We vary the final position of the cloud and measure the mean shift of the cavity-like mode, and plot this in Figure 7.4. Using our knowledge of the cavity mode distribution,

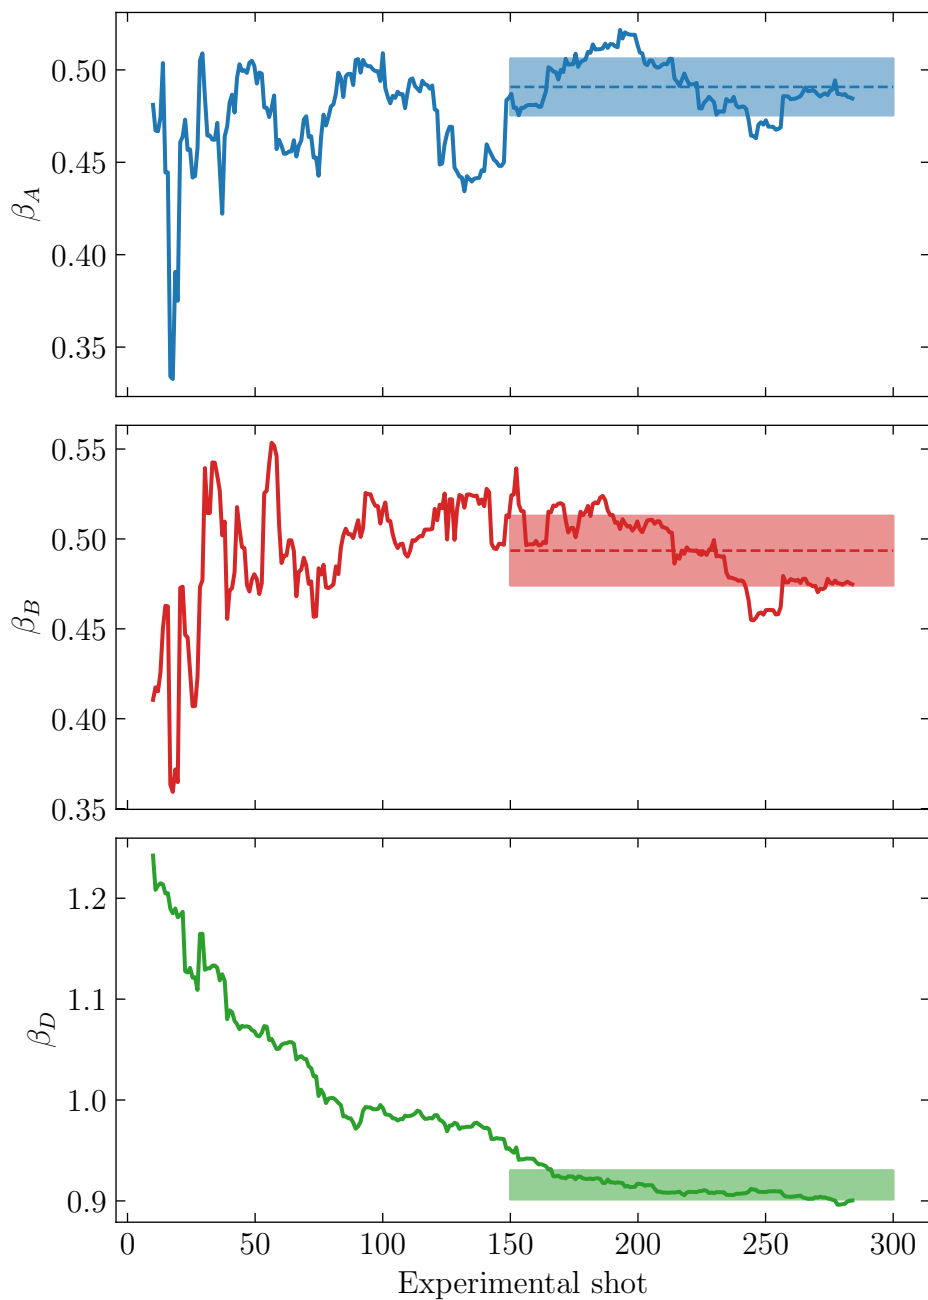


Figure 7.3: The optimal estimators calculated for the spin-squeezed clock comparison data. (a) Optimal estimator β_A for the pre-measurement of A. (b) Optimal estimator β_B for the pre-measurement of B. (c) Optimal estimator β_D , which is the differential scale factor between the two ensembles. The shaded region indicates the ± 1 standard deviation for the datapoints beyond experimental shot 150, and the dashed line is the corresponding mean value.




Figure 7.4: Measured shift of the cavity-like mode versus vertical displacement of the atomic cloud. The green dashed line is the Gaussian fit of the form of Eq. 7.8, giving a fitted $\sigma_z = 130\mu\text{m}$.

we calculate the frequency shift using Eq. B.1 and Eq. B.2

$$\omega_{\downarrow} = \frac{g_{eff}^2 N_{eff}}{\delta_c} = \frac{N_{tot} \langle g_i^2 \rangle}{\delta_c} \quad (7.7)$$

Plugging in the appropriate distributions in $\langle g_i^2 \rangle$, and considering a varying offset of the cloud in the vertical direction Z_0 , we find that the shift has the form,

$$\frac{\omega_{\downarrow}(Z_0)}{\omega_{\downarrow}(0)} = e^{\frac{-2Z_0^2}{4\sigma_Z^2 + w_0^2}}, \quad (7.8)$$

where σ_Z is the rms extent of the cloud along the vertical direction. We then fit our measured cavity frequency shift to Eq. 7.8. After plugging in our cavity mode waist of $w_0 = 71 \mu\text{m}$, we extract a vertical cloud size of $\sigma_Z = 130 \mu\text{m}$. This cloud size is in good agreement with images of the cloud using fluorescence imaging.

One can imagine various different routes for improvement. If we could get an atomic cloud with both a larger atom number and larger extend σ_z , we could easily probe more than 2 ensembles within the same cloud. Another approach would be to create spatially separated ensembles [87], each having $\sigma_z < w_0$. These more compressed clouds could give both an improved effective values for atom number N and the atom-cavity coupling g .

7.5 Common-mode suppression of laser noise

Probing the two sub-ensembles with the same clock laser has the great advantage of allowing the laser phase noise to be common between the evolution of the two Bloch vectors. In order to check our level of common-mode suppression, we operate a two-ensemble Ramsey experiment and add in artificial laser noise. We raise the artificial noise level by adding a white frequency noise term into the fiber noise cancellation reference synthesizer. This is achieved using an SRS SG382 synthesizer, and using the internal noise modulation feature. The measured cavity shift noise versus added white noise is shown in Figure 7.5. It is clear the the added noise shows up in each ensemble (red and blue circles), but when we take the difference (open green squares), we see a good common-mode rejection of noise. Looking at the largest noise, we can see a direct suppression




Figure 7.5: Measured cavity shift noise with artificial white frequency noise imposed on the clock laser. Each ensemble (red and blue circles) becomes noisier as the magnitude of the added noise is increased, but taking the difference (open green squares) gives a 10 dB common mode noise reduction factor.

Table 7.1: Quantum efficiency contributions

Parameter	Value
QE of cavity (κ_1/κ)	0.68
Photodiode QE	0.88
Homodyne-cavity mode overlap	0.75
Path efficiency	0.62
Photodiode technical	0.99
Total	0.28

of noise corresponding to 10 dB. The Si3 clock laser has a frequency noise power spectral density of about 10^{-2} Hz²/Hz. Our estimated contribution in terms of cavity frequency shift noise would be ≈ 50 rms Hz, or 30 dB below QPN. We comment here that only 10 dB of this cushion comes from the common-mode suppression. The full 30 dB is only possible because of our advanced clock laser technology.

7.6 Quantum efficiency

The overall quantum efficiency of the measurement plays a key role in QND-based spin squeezing. In thinking about probing the atom-cavity coupled system, we define the overall quantum efficiency as the product of all of the independent sources of loss of photons, or decrease of signal to noise ratio for the homodyne detection. The amount of attainable spin-squeezing is linearly proportional to the quantum efficiency [32]. Furthermore, $Q < 1$ leads to excess noise in the anti-squeezed quadrature. We first estimate the overall quantum efficiency by considering the different contributions, including the cavity quantum efficiency $\kappa_1/\kappa = 0.68$, mode overlap of the cavity leakage light and the homodyne LO beam of 0.75, quantum efficiency of the photodiode of 0.88, and finite path efficiency of 0.62, and other negligible sources. Multiplying these together gives a quantum efficiency of $Q = 0.28$. We use this value when estimating the expected R (cyan line in Figure 6.11), showing reasonable agreement.

7.7 Spin state tomography

We perform state tomography to evaluate the amount of anti-squeezing introduced by the QND measurement. The pulse sequence used for measuring the anti-squeezed quadrature is shown in Figure 7.6a. We make a pre-measurement of J_z , jump the clock laser phase by 90° to rotate the state by an Ψ , and then perform a final measurement. For the rotation angles larger than 5 degrees, we use a swept readout for the final measurement, since the measurement noise is large enough to depart from the linear regime of the Q response of the homodyne signal. To our surprise, the observed anti-squeezing is well above the expected level given our estimated quantum efficiency of $Q = 0.28$. This indicates additional technical noise in the anti-squeezing quadrature, or in the measurement itself. The measured noise level at the anti-squeezed quadrature is 9 dB above what is expected from the estimated quantum efficiency (Figure 7.6b). This indicates an additional noise source contributing to the anti-squeezing, but it does not preclude observing the benefit of spin squeezing. This is clearly an avenue for further investigation.

7.8 Independence of the two ensembles

There is an inherent limitation in the total atom number per sub-ensemble for a given total atom number and vertical extent of the cloud. For the differential clock comparisons performed in this thesis, we tune the transport distances such that the cavity frequency shift is equal for the two ensembles (to within 10 - 20%). When tuning up the experiment and the relative distances, one wants to confirm the independence of the sub-ensembles and determine the magnitude of the correction factor due to finite overlap between the ensembles and the cavity mode. To explore this, we vary the vertical separation between the ensembles and evaluate the Pearson correlation coefficient between the measured J_z for each ensemble. For no separation, we probe the same ensemble and expect high correlations between the measured QPN to within the measurement noise. Then as we increase the separation distance, the correlation coefficient will decrease accordingly to a noise floor set by correlated noise (classical rotation noise). We first measure $J_{z,A}$ for ensemble A,

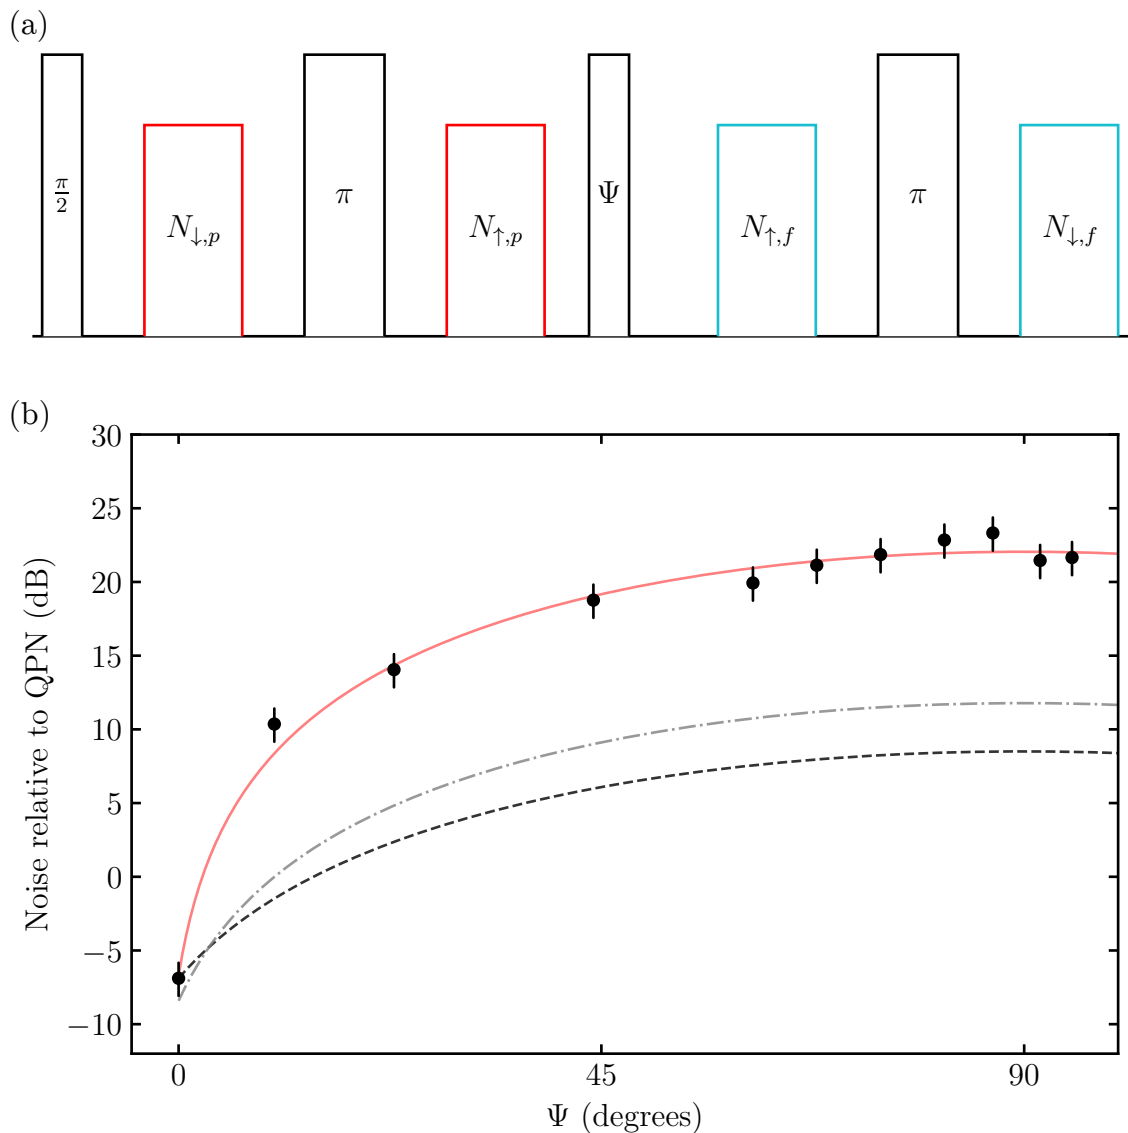


Figure 7.6: (a) Pulse sequence used for performing spin state tomography. (b) Relative spin noise after a rotation angle of Ψ . We infer the spin noise just after the pre-measurement by assuming the noise in the pre and final measurement is uncorrelated. The black dashed line is a unitary spin squeezed state, and the dashed-dot line indicates our expected anti-squeezing for a quantum efficiency of 0.28.

apply a vertical displacement of the cloud via the movable optical lattice to put ensemble B in the cavity, and subsequently measure $J_{z,B}$. The measured correlation coefficient between $J_{z,A}$ and $J_{z,B}$ is shown versus the separation distance between the ensembles in units of the waist of the cavity mode w_0 (Figure 7.7a). The blue line is a numerical Monte Carlo simulation of the correlation coefficient for two ensembles with varying mean separation. Our operational ensemble separation Z_0 is denoted as the dashed grey line, at $Z_0 = 150 \mu\text{m}$.

While the correlation coefficient is a relatively easy quantity to calculate from experimental data, it is not the quantity we really would like to know. Based on the finite overlap of the ensemble, we would like to calculate the change of the combined QPN versus the separation distance. For this we use both analytical and numerical methods, with both shown in Figure 7.7b. In the limit of a flat density distribution, the analytical correction to the QPN is straightforwardly calculated as

$$d(\Delta\theta_{\text{QPN}}) = 1 - e^{-Z_0^2/w_0^2} \quad (7.9)$$

We can calculate the full analytical expression taking into account the finite σ_Z for the ensemble. From this, we perform a series expansion so we can see the second order term,

$$d(\Delta\theta_{\text{QPN}}) = 1 - e^{-Z_0^2/w_0^2} + \left(\frac{Z_0}{2\sqrt{2}\sigma_Z}\right)^2 e^{-Z_0^2/w_0^2}. \quad (7.10)$$

At our vertical cloud size of $\sigma_z = 130 \mu\text{m}$, and a separation distance of $Z_0 = 150 \mu\text{m}$, this higher order term contributes a 0.2% correction. For the differential clock comparisons of Figure 7.2, a spatial separation of $150 \mu\text{m}$ (dashed grey line) is chosen to guarantee sufficient independence of the ensembles. At this separation, the adjustment to the QPN is 0.045 dB.




Figure 7.7: Independence of atomic ensembles (a) Measured correlation coefficient between $J_{z,A}$ and $J_{z,B}$ versus the separation between the ensembles. (black circles). The blue line is a Monte Carlo simulation. (b) Corresponding change of the QPN due to the finite overlap of the ensembles, with numerical Monte Carlo simulation (blue) and analytical calculation (orange). At our operating ensemble separation, the change to QPN is 0.05 dB.

Chapter 8

Conclusion and outlook

This Thesis has covered the improvement of optical clock technology on multiple fronts. We developed some of the best stable lasers using cryogenic silicon cavities. This involved building high finesse cavities that reduced the Brownian thermal noise from the optical coating, which is the primary source of the noise floor of these optical interferometers. We then built a new optical clock apparatus that implements spin-squeezing approaching the state-of-the-art level. For the first time, we were able to directly observe an optical clock comparison operating below the QPN limit.

One obvious goal is to continue to push the performance below the SQL. There are clear improvements and aspects of the system to further understand. The aim is to directly leverage and highlight entanglement for enhanced clock performance. The initial contrast of the atomic ensembles will be further improved by optimizing aspects of the optical lattice and the associated atomic motion. Adding more confinement either along the cavity axis or out of the plane of the cavity could improve the atomic coherence, allowing for immediate improvements in the Wineland parameter. Another approach would be to generate entanglement in the ground-state manifold of ^{87}Sr similar to [35], which could avoid some of the detrimental clock pulses.

Another priority is to further understand the magnitude of our measured anti-squeezed quadrature. Excess anti-squeezing can make us more sensitive to amplitude errors in the clock rotations. Further, as we push the Ramsey dark time to longer intervals, the leakage of the anti-squeezed quadrature will start to play a role [137].

Improving the amount of squeezing involves improving the readout noise of our QND mea-

surement. Since the measurement noise increases with high photon number, we could win 2-3 dB of spin squeezing by understanding and mitigating this effect. We could also explore other ways of generating entanglement, for example using one-axis twisting [109, 111]. Depending on the exact experimental parameters, one might find OAT or QND the favorable approach.

In the longer term, the system demonstrated in this thesis will serve as a blue print for building more advanced versions of entangled clocks, with prospects of two-mode squeezing for clock comparisons and integrating squeezing into the best performing precision to enhance measurement bandwidth.

Bibliography

- [1] H Gegier and Ernest Marsden. On a diffuse reflection of the α -particles. Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character, 82(557):495–500, 1909.
- [2] Ernest Rutherford. The scattering of α and β particles by matter and the structure of the atom. The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science, 21(125):669–688, 1911.
- [3] Niels Bohr. I. on the constitution of atoms and molecules. The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science, 26(151):1–25, 1913.
- [4] I. I. Rabi, J. R. Zacharias, S. Millman, and P. Kusch. A new method of measuring nuclear magnetic moment. Phys. Rev., 53:318–318, Feb 1938.
- [5] Willis E. Lamb and Robert C. Retherford. Fine structure of the hydrogen atom by a microwave method. Phys. Rev., 72:241–243, Aug 1947.
- [6] Norman F. Ramsey. A molecular beam resonance method with separated oscillating fields. Phys. Rev., 78:695–699, Jun 1950.
- [7] Norman F. Ramsey. Obituary : I.I. Rabi, 1988.
- [8] Harold Lyons. The atomic clock: A universal standard of frequency and time. The American Scholar, 19(2):159–168, 1950.
- [9] Norman F. Ramsey. Experiments with separated oscillatory fields and hydrogen masers. Rev. Mod. Phys., 62:541–552, Jul 1990.
- [10] L. Essen and J. V. L. Parry. An atomic standard of frequency and time interval: A caesium resonator. Nature, 176(4476):280–282, 1955.
- [11] John L. Hall. Why it took so long for the laser and the optical comb to be invented: the unmarked trail from concept to experimental reality. J. Opt. Soc. Am. B, 34(2):338–346, Feb 2017.
- [12] Martin M. Boyd, Tanya Zelevinsky, Andrew D. Ludlow, Seth M. Foreman, Sebastian Blatt, Tetsuya Ido, and Jun Ye. Optical atomic coherence at the 1-second time scale. Science, 314(5804):1430–1433, 2006.

- [13] A. D. Ludlow, T. Zelevinsky, G. K. Campbell, S. Blatt, M. M. Boyd, M. H. G. de Miranda, M. J. Martin, J. W. Thomsen, S. M. Foreman, J. Ye, T. M. Fortier, J. E. Stalnaker, S. A. Diddams, Y. Le Coq, Z. W. Barber, N. Poli, N. D. Lemke, K. M. Beck, and C. W. Oates. Sr lattice clock at 1×10^{-16} fractional uncertainty by remote optical evaluation with a Ca clock. *Science*, 319(5871):1805–1808, 2008.
- [14] Tobias Bothwell, Dhruv Kedar, Eric Oelker, John M Robinson, Sarah L Bromley, Weston L Tew, Jun Ye, and Colin J Kennedy. Jila Sr1 optical lattice clock with uncertainty of 2.0×10^{-18} . *Metrologia*, 56(6):065004, 2019.
- [15] W. F. McGrew, X. Zhang, R. J. Fasano, S. A. Schäffer, K. Beloy, D. Nicolodi, R. C. Brown, N. Hinkley, G. Milani, M. Schioppo, T. H. Yoon, and A. D. Ludlow. Atomic clock performance enabling geodesy below the centimetre level. *Nature*, 564(7734):87–90, 2018.
- [16] S. M. Brewer, J.-S. Chen, A. M. Hankin, E. R. Clements, C. W. Chou, D. J. Wineland, D. B. Hume, and D. R. Leibbrandt. $^{27}\text{Al}^+$ quantum-logic clock with a systematic uncertainty below 10^{-18} . *Phys. Rev. Lett.*, 123:033201, Jul 2019.
- [17] E. Oelker, R. B. Hutson, C. J. Kennedy, L. Sonderhouse, T. , A. Goban, D. Kedar, C. Sanner, J. M. Robinson, G. E. Marti, D. G. 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*, 13(10):714–719, 2019.
- [18] William Bowden, Alvise Vianello, Ian R. Hill, Marco Schioppo, and Richard Hobson. Improving the q factor of an optical atomic clock using quantum nondemolition measurement. *Phys. Rev. X*, 10:041052, Dec 2020.
- [19] Tobias Bothwell, Colin J. Kennedy, Alexander Aeppli, Dhruv Kedar, John M. Robinson, Eric Oelker, Alexander Staron, and Jun Ye. Resolving the gravitational redshift across a millimetre-scale atomic sample. *Nature*, 602(7897):420–424, 2022.
- [20] W. M. Itano, J. C. Bergquist, J. J. Bollinger, J. M. Gilligan, D. J. Heinzen, F. L. Moore, M. G. Raizen, and D. J. Wineland. Quantum projection noise: Population fluctuations in two-level systems. *Phys. Rev. A*, 47:3554–3570, May 1993.
- [21] Giorgio Santarelli, Ph Laurent, Pierre Lemonde, André Clairon, Anthony G Mann, S Chang, Andre N Luiten, and Christophe Salomon. Quantum projection noise in an atomic fountain: A high stability cesium frequency standard. *Physical Review Letters*, 82(23):4619, 1999.
- [22] G John 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*, pages 133–147, 1989.
- [23] G. Santarelli, C. Audoin, A. Makdissi, P. Laurent, G.J. 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*, 45(4):887–894, 1998.
- [24] M. Martin. *Quantum Metrology and Many-Body Physics: Pushing the Frontier of the Optical Lattice Clock*. PhD thesis, Boulder, CO 80309-0440, 2013-04 2013.

- [25] 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 μm lasers with sub-10 mHz linewidth. Phys. Rev. Lett., 118:263202, Jun 2017.
- [26] 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., 119:243601, Dec 2017.
- [27] John M. Robinson, Eric Oelker, William R. Milner, Wei Zhang, Thomas Legero, Dan G. Matei, Fritz Riehle, Uwe Sterr, and Jun Ye. Crystalline optical cavity at 4 K with thermal-noise-limited instability and ultralow drift. Optica, 6(2):240–243, Feb 2019.
- [28] G. Edward Marti, Ross B. Hutson, Akihisa Goban, Sara L. Campbell, Nicola Poli, and Jun Ye. Imaging optical frequencies with 100 μHz precision and 1.1 μm resolution. Phys. Rev. Lett., 120:103201, Mar 2018.
- [29] Luca Pezzè, Augusto Smerzi, Markus K. Oberthaler, Roman Schmied, and Philipp Treutlein. Quantum metrology with nonclassical states of atomic ensembles. Rev. Mod. Phys., 90:035005, Sep 2018.
- [30] Monika H. Schleier-Smith, Ian D. Leroux, and Vladan Vuletić. States of an ensemble of two-level atoms with reduced quantum uncertainty. Phys. Rev. Lett., 104:073604, Feb 2010.
- [31] Zilong Chen, Justin G. Bohnet, Shannon R. Sankar, Jiayan Dai, and James K. Thompson. Conditional spin squeezing of a large ensemble via the vacuum Rabi splitting. Phys. Rev. Lett., 106:133601, Mar 2011.
- [32] Zilong Chen, Justin G. Bohnet, Joshua M. Weiner, Kevin C. Cox, and James K. Thompson. Cavity-aided nondemolition measurements for atom counting and spin squeezing. Phys. Rev. A, 89:043837, Apr 2014.
- [33] Onur Hosten, Nils J. Engelsen, Rajiv Krishnakumar, and Mark A. Kasevich. Measurement noise 100 times lower than the quantum-projection limit using entangled atoms. Nature, 529(7587):505–508, 2016.
- [34] Kevin C. Cox, Graham P. Greve, Joshua M. Weiner, and James K. Thompson. Deterministic squeezed states with collective measurements and feedback. Phys. Rev. Lett., 116:093602, Mar 2016.
- [35] Edwin Pedrozo-Penafiel, Simone Colombo, Chi Shu, Albert F. Adiyatullin, Zeyang Li, Enrique Mendez, Boris Braverman, Akio Kawasaki, Daisuke Akamatsu, Yanhong Xiao, and Vladan Vuletic. Entanglement on an optical atomic-clock transition. Nature, 588(7838):414–418, 2020.
- [36] G Vallet, E Bookjans, U Eismann, S Bilicki, R Le Targat, and J Lodewyck. A noise-immune cavity-assisted non-destructive detection for an optical lattice clock in the quantum regime. New Journal of Physics, 19(8):083002, aug 2017.
- [37] Andrew D. Ludlow, Martin M. Boyd, Jun Ye, E. Peik, and P. O. Schmidt. Optical atomic clocks. Rev. Mod. Phys., 87:637–701, Jun 2015.

- [38] Dieter Hils and J. L. Hall. Improved Kennedy-Thorndike experiment to test special relativity. Phys. Rev. Lett., 64:1697–1700, Apr 1990.
- [39] E. Wiens, A. Yu. Nevsky, and S. Schiller. Resonator with ultrahigh length stability as a probe for equivalence-principle-violating physics. Phys. Rev. Lett., 117:271102, Dec 2016.
- [40] The LIGO Scientific Collaboration. LIGO: the laser interferometer gravitational-wave observatory. Reports on Progress in Physics, 72(7):076901, 2009.
- [41] T. L. Nicholson, M. J. Martin, J. R. Williams, B. J. Bloom, M. Bishof, M. D. Swallows, S. L. Campbell, and J. Ye. Comparison of two independent Sr optical clocks with 1×10^{-17} stability at 10^3 s. Phys. Rev. Lett., 109:230801, Dec 2012.
- [42] Y. Y. Jiang, A. D. Ludlow, N. D. Lemke, R. W. Fox, J. A. Sherman, L. S. Ma, and C. W. Oates. Making optical atomic clocks more stable with 10^{-16} level laser stabilization. Nature Photonics, 5:158, 01 2011.
- [43] 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, 93:063630, Jun 2016.
- [44] 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, 94:124043, Dec 2016.
- [45] M. J. Martin, M. Bishof, M. D. Swallows, X. Zhang, C. Benko, J. von Stecher, A. V. Gorshkov, A. M. Rey, and Jun Ye. A quantum many-body spin system in an optical lattice clock. Science, 341(6146):632–636, 2013.
- [46] Fritz Riehle. Towards a redefinition of the second based on optical atomic clocks. Comptes Rendus Physique, 16(5):506 – 515, 2015.
- [47] Hidekazu Hachisu, Fumimaru Nakagawa, Yuko Hanado, and Tetsuya Ido. Months-long real-time generation of a time scale based on an optical clock. Scientific Reports, 8(1):4243, 2018.
- [48] Fritz Riehle, Patrick Gill, Felicitas Arias, and Lennart Robertsson. The CIPM list of recommended frequency standard values: guidelines and procedures. Metrologia, 55(2):188, 2018.
- [49] Colin J. Kennedy, Eric Oelker, John M. Robinson, Tobias Bothwell, Dhruv Kedar, William R. Milner, G. Edward Marti, Andrei Derevianko, and Jun Ye. Precision metrology meets cosmology: Improved constraints on ultralight dark matter from atom-cavity frequency comparisons. Phys. Rev. Lett., 125:201302, Nov 2020.
- [50] William R. Milner, John M. Robinson, Colin J. Kennedy, Tobias Bothwell, Dhruv Kedar, Dan G. Matei, Thomas Legero, Uwe Sterr, Fritz Riehle, Holly Leopardi, Tara M. Fortier, Jeffrey A. Sherman, Judah Levine, Jian Yao, Jun Ye, and Eric Oelker. Demonstration of a timescale based on a stable optical carrier. Phys. Rev. Lett., 123:173201, Oct 2019.
- [51] Matthew A. Norcia, Julia R. K. Cline, Juan A. Muniz, John M. Robinson, Ross B. Hutson, Akihisa Goban, G. Edward Marti, Jun Ye, and James K. Thompson. Frequency measurements of superradiance from the strontium clock transition. Phys. Rev. X, 8:021036, May 2018.

- [52] Judith Olson, Richard W. Fox, Tara M. Fortier, Todd F. Sheerin, Roger C. Brown, Holly Leopardi, Richard E. Stoner, Chris W. Oates, and Andrew D. Ludlow. Ramsey-bordé matter-wave interferometry for laser frequency stabilization at 10^{-16} frequency instability and below. Phys. Rev. Lett., 123:073202, Aug 2019.
- [53] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward. Laser phase and frequency stabilization using an optical resonator. Applied Physics B, 31(2):97–105, 1983.
- [54] John M. Robinson, Eric Oelker, William R. Milner, Dhruv Kedar, Wei Zhang, Thomas Legero, Dan G. Matei, Sebastian Häfner, Fritz Riehle, Uwe Sterr, and Jun Ye. Thermal noise and mechanical loss of SiO₂/Ta₂O₅ optical coatings at cryogenic temperatures. Opt. Lett., 46(3):592–595, Feb 2021.
- [55] Kenji Numata, Amy Kemery, and Jordan Camp. Thermal-noise limit in the frequency stabilization of lasers with rigid cavities. Phys. Rev. Lett., 93:250602, Dec 2004.
- [56] M J Martin and J Ye. High-precision laser stabilization via optical cavities. In Optical Coatings and Thermal Noise in Precision Measurement. Cambridge University Press, 2012.
- [57] Peter R. Saulson. Thermal noise in mechanical experiments. Phys. Rev. D, 42:2437–2445, Oct 1990.
- [58] Yu. Levin. Internal thermal noise in the LIGO test masses: A direct approach. Phys. Rev. D, 57:659–663, Jan 1998.
- [59] W. E. Gifford. The Gifford-McMahon cycle. In K. D. Timmerhaus, editor, Advances in Cryogenic Engineering, pages 152–159, Boston, MA, 1966. Springer US.
- [60] <http://montanainstruments.com/>. Commercial products are identified for technical clarity. Such identification does not imply endorsement by authors.
- [61] DG Matei, T Legero, Ch Grebing, S Häfner, Ch Lisdat, R Weyrich, W Zhang, L Sonderhouse, JM Robinson, F Riehle, et al. A second generation of low thermal noise cryogenic silicon resonators. In Journal of Physics: Conference Series, volume 723, page 012031. IOP Publishing, 2016.
- [62] Hui Shen, Liufeng Li, Jin Bi, Jia Wang, and Lisheng Chen. Systematic and quantitative analysis of residual amplitude modulation in Pound-Drever-Hall frequency stabilization. Phys. Rev. A, 92:063809, Dec 2015.
- [63] 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., 39(7):1980–1983, Apr 2014.
- [64] Alessandro Farsi, Mario Siciliani de Cumis, Francesco Marino, and Francesco Marin. Photothermal and thermo-refractive effects in high reflectivity mirrors at room and cryogenic temperature. Journal of Applied Physics, 111(4):043101, 2012.

- [65] Kazuhiro Yamamoto, Shinji Miyoki, Takashi Uchiyama, Hideki Ishitsuka, Masatake Ohashi, Kazuaki Kuroda, Takayuki Tomaru, Nobuaki Sato, Toshikazu Suzuki, Tomiyoshi Haruyama, Akira Yamamoto, Takakazu Shintomi, Kenji Numata, Koichi Waseda, Kazuhiko Ito, and Koji Watanabe. Measurement of the mechanical loss of a cooled reflective coating for gravitational wave detection. Phys. Rev. D, 74:022002, Jul 2006.
- [66] Eiichi Hirose, Kieran Craig, Hideki Ishitsuka, Iain W. Martin, Norikatsu Mio, Shigenori Moriwaki, Peter G. Murray, Masatake Ohashi, Sheila Rowan, Yusuke Sakakibara, Toshikazu Suzuki, Kouichi Waseda, Kyohei Watanabe, and Kazuhiro Yamamoto. Mechanical loss of a multilayer tantala/silica coating on a sapphire disk at cryogenic temperatures: Toward the kagra gravitational wave detector. Phys. Rev. D, 90:102004, Nov 2014.
- [67] Christian Hagemann, Christian Grebing, Christian Lisdat, Stephan Falke, Thomas Legero, Uwe Sterr, Fritz Riehle, Michael J. Martin, and Jun Ye. Ultrastable laser with average fractional frequency drift rate below 5×10^{-19} /s. Opt. Lett., 39(17):5102–5105, Sep 2014.
- [68] Dhruv Kedar, Jialiang Yu, Eric Oelker, Alexander Staron, William R. Milner, John M. Robinson, Thomas Legero, Fritz Riehle, Uwe Sterr, and Jun Ye. Frequency stability of cryogenic silicon cavities with semiconductor crystalline coatings, 2022.
- [69] Garrett D. Cole, Wei Zhang, Bryce J. Bjork, David Follman, Paula Heu, Christoph Deutsch, Lindsay Sonderhouse, John Robinson, Chris Franz, Alexei Alexandrovski, Mark Notcutt, Oliver H. Heckl, Jun Ye, and Markus Aspelmeyer. High-performance near- and mid-infrared crystalline coatings. Optica, 3(6):647–656, Jun 2016.
- [70] Garrett D Cole, Wei Zhang, Michael J Martin, Jun Ye, and Markus Aspelmeyer. Tenfold reduction of brownian noise in high-reflectivity optical coatings. Nature Photonics, 7(8):644–650, 2013.
- [71] Ting Hong, Huan Yang, Eric K. Gustafson, Rana X. Adhikari, and Yanbei Chen. Brownian thermal noise in multilayer coated mirrors. Phys. Rev. D, 87:082001, Apr 2013.
- [72] Gregory M Harry, Andri M Gretarsson, Peter R Saulson, Scott E Kittelberger, Steven D Penn, William J Startin, Sheila Rowan, Martin M Fejer, D R M Crooks, Gianpietro Cagnoli, Jim Hough, and Norio Nakagawa. Thermal noise in interferometric gravitational wave detectors due to dielectric optical coatings. Classical and Quantum Gravity, 19(5):897–917, feb 2002.
- [73] Alex Amato, Gianpietro Cagnoli, Maurizio Canepa, Elodie Coillet, Jerome Degallaix, Vincent Dolique, Daniele Forest, Massimo Granata, Valérie Martinez, Christophe Michel, Laurent Pinard, Benoit Sassolas, and Julien Teillon. High-reflection coatings for gravitational-wave detectors:state of the art and future developments. Journal of Physics: Conference Series, 957:012006, feb 2018.
- [74] S. Gras and M. Evans. Direct measurement of coating thermal noise in optical resonators. Phys. Rev. D, 98:122001, Dec 2018.
- [75] Kieran Craig, Jessica Steinlechner, Peter G. Murray, Angus S. Bell, Ross Birney, Karen Haughian, Jim Hough, Ian MacLaren, Steve Penn, Stuart Reid, Raymond Robie, Sheila Rowan, and Iain W. Martin. Mirror coating solution for the cryogenic Einstein telescope. Phys. Rev. Lett., 122:231102, Jun 2019.

- [76] Kazuhiro Yamamoto, Shinji Miyoki, Takashi Uchiyama, Hideki Ishitsuka, Masatake Ohashi, Kazuaki Kuroda, Takayuki Tomaru, Nobuaki Sato, Toshikazu Suzuki, Tomiyoshi Haruyama, Akira Yamamoto, Takakazu Shintomi, Kenji Numata, Koichi Waseda, Kazuhiko Ito, and Koji Watanabe. Measurement of the mechanical loss of a cooled reflective coating for gravitational wave detection. *Phys. Rev. D*, 74:022002, Jul 2006.
- [77] Raymond Robie. Characterisation of the mechanical properties of thin-film mirror coating materials for use in future interferometric gravitational wave detectors. PhD thesis, University of Glasgow, 2018.
- [78] Tara Chalermongsak, Frank Seifert, Evan D Hall, Koji Arai, Eric K Gustafson, and Rana X Adhikari. Broadband measurement of coating thermal noise in rigid Fabry–Pérot cavities. *Metrologia*, 52(1):17–30, dec 2014.
- [79] Sebastian Häfner, Sofia Herbers, Stefan Vogt, Christian Lisdat, and Uwe Sterr. Transportable interrogation laser system with an instability of mod $\sigma_y = 3 \times 10^{-16}$. *Optics Express*, 28(11):16407–16416, 2020.
- [80] Thomas Kessler, Christian Hagemann, C Grebing, T Legero, Uwe Sterr, Fritz Riehle, MJ Martin, L Chen, and J Ye. A sub-40-mhz-linewidth laser based on a silicon single-crystal optical cavity. *Nature Photonics*, 6(10):687, 2012.
- [81] C. Hagemann. Ultra-stable laser based on a cryogenic single-crystal silicon cavity. PhD thesis, 2013-04 2013.
- [82] 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*, 506(7486):71–75, 2014.
- [83] M. Schioppo, R. C. Brown, W. F. McGrew, N. Hinkley, R. J. Fasano, K. Beloy, T. H. Yoon, G. Milani, D. Nicolodi, J. A. Sherman, N. B. Phillips, C. W. Oates, and A. D. Ludlow. Ultrastable optical clock with two cold-atom ensembles. *Nature Photonics*, 11(1):48–52, 2017.
- [84] Tobias Bothwell. A Wannier-Stark Optical Lattice Clock With Extended Coherence Times. PhD thesis, Boulder, 2022-04 2022.
- [85] S. L. Campbell, R. B. Hutson, G. E. Marti, A. Goban, N. Darkwah Oppong, R. L. McNally, L. Sonderhouse, J. M. Robinson, W. Zhang, B. J. Bloom, and J. Ye. A Fermi-degenerate three-dimensional optical lattice clock. *Science*, 358(6359):90–94, 2017.
- [86] Martin M. Boyd, Tanya Zelevinsky, Andrew D. Ludlow, Sebastian Blatt, Thomas Zanon-Willette, Seth M. Foreman, and Jun Ye. Nuclear spin effects in optical lattice clocks. *Phys. Rev. A*, 76:022510, Aug 2007.
- [87] Xin Zheng, Jonathan Dolde, Varun Lochab, Brett N. Merriman, Haoran Li, and Shimon Kolkowitz. Differential clock comparisons with a multiplexed optical lattice clock. *Nature*, 602(7897):425–430, 2022.
- [88] David N. Schramm and Michael S. Turner. Big-bang nucleosynthesis enters the precision era. *Rev. Mod. Phys.*, 70:303–318, Jan 1998.

- [89] Gianfranco Bertone and Dan Hooper. History of dark matter. *Rev. Mod. Phys.*, 90:045002, Oct 2018.
- [90] 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. Crespo Lopez-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. Dbrich, 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. Jackson 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.
- [91] Thibault Damour and John F. Donoghue. Equivalence principle violations and couplings of a light dilaton. *Phys. Rev. D*, 82:084033, Oct 2010.
- [92] Asimina Arvanitaki, Junwu Huang, and Ken Van Tilburg. Searching for dilaton dark matter with atomic clocks. *Phys. Rev. D*, 91:015015, Jan 2015.
- [93] Aurélien Hees, Olivier Minazzoli, Etienne Savalle, Yevgeny V. Stadnik, and Peter Wolf. Violation of the equivalence principle from light scalar dark matter. *Phys. Rev. D*, 98:064051, Sep 2018.
- [94] Gretchen K Campbell, Andrew D Ludlow, Sebastian Blatt, Jan W Thomsen, Michael J Martin, Marcio H G de Miranda, Tanya Zelevinsky, Martin M Boyd, Jun Ye, Scott A Diddams, Thomas P Heavner, Thomas E Parker, and Steven R Jefferts. The absolute frequency of the ^{87}Sr optical clock transition. *Metrologia*, 45(5):539–548, sep 2008.
- [95] 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*, 4(12):eaau4869, 2018.
- [96] Ken Van Tilburg, Nathan Leefer, Lykourgos Bougas, and Dmitry Budker. Search for ultralight scalar dark matter with atomic spectroscopy. *Phys. Rev. Lett.*, 115(1):011802, 2015.
- [97] 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.*, 117(6):061301, 2016.

- [98] Joel Bergé, Philippe Brax, Gilles Métris, Martin Pernot-Borràs, Pierre Touboul, and Jean-Philippe Uzan. MICROSCOPE mission: First constraints on the violation of the weak equivalence principle by a light scalar dilaton. *Phys. Rev. Lett.*, 120(14):141101, 2018.
- [99] Stephan Schlamminger, K-Y Choi, Todd A Wagner, Jens H Gundlach, and Eric G Adelberger. Test of the equivalence principle using a rotating torsion balance. *Phys. Rev. Lett.*, 100(4):041101, 2008.
- [100] N. Leefler, A. Gerhardus, D. Budker, V. V. Flambaum, and Y. V. Stadnik. Search for the effect of massive bodies on atomic spectra and constraints on yukawa-type interactions of scalar particles. *Phys. Rev. Lett.*, 117:271601, Dec 2016.
- [101] Pierre Touboul, Gilles Métris, Manuel Rodrigues, Yves André, Quentin Baghi, Joël Bergé, Damien Boulanger, Stefanie Bremer, Patrice Carle, Ratana Chhun, Bruno Christophe, Valerio Cipolla, Thibault Damour, Pascale Danto, Hansjoerg Dittus, Pierre Fayet, Bernard Foulon, Claude Gageant, Pierre-Yves Guidotti, Daniel Hagedorn, Emilie Hardy, Phuong-Anh Huynh, Henri Inchauspe, Patrick Kayser, Stéphanie Lala, Claus Lämmerzahl, Vincent Lebat, Pierre Leseur, Françoise Liorzou, Meike List, Frank Löffler, Isabelle Panet, Benjamin Pouilloux, Pascal Prieur, Alexandre Rebray, Serge Reynaud, Benny Rievers, Alain Robert, Hanns Selig, Laura Serron, Timothy Sumner, Nicolas Tanguy, and Pieter Visser. MICROSCOPE mission: First results of a space test of the equivalence principle. *Phys. Rev. Lett.*, 119:231101, Dec 2017.
- [102] T. Rosenband, D. B. Hume, P. O. Schmidt, C. W. Chou, A. Brusch, L. Lorini, W. H. Oskay, R. E. Drullinger, T. M. Fortier, J. E. Stalnaker, S. A. Diddams, W. C. Swann, N. R. Newbury, W. M. Itano, D. J. Wineland, and J. C. Bergquist. Frequency ratio of Al^+ and Hg^+ single-ion optical clocks; metrology at the 17th decimal place. *Science*, 319(5871):1808–1812, 2008.
- [103] Steven T Cundiff and Jun Ye. Colloquium: Femtosecond optical frequency combs. *Reviews of Modern Physics*, 75(1):325, 2003.
- [104] Holly Leopardi, Josue Davila-Rodriguez, Franklyn Quinlan, Judith Olson, Jeff A. Sherman, Scott A. Diddams, and Tara M. Fortier. Single-branch Er: fiber frequency comb for precision optical metrology with 10^{-18} fractional instability. *Optica*, 4(8):879–885, Aug 2017.
- [105] N. R. Newbury, P. A. Williams, and W. C. Swann. Coherent transfer of an optical carrier over 251 km. *Opt. Lett.*, 32(21):3056–3058, Nov 2007.
- [106] Holly Leopardi, Kyle Beloy, Tobias Bothwell, Samuel M Brewer, Sarah L Bromley, Jwo-Sy Chen, Scott A Diddams, Robert J Fasano, Youssef S Hassan, David B Hume, Dhruv Kedar, Colin J Kennedy, David R Leibbrandt, Andrew D Ludlow, William F McGrew, William R Milner, Daniele Nicolodi, Eric Oelker, Thomas E Parker, John M Robinson, Stefania Romisch, Jeff A Sherman, Lindsay Sonderhouse, Jian Yao, Jun Ye, Xiaogang Zhang, and Tara M Fortier. Measurement of the 27Al^+ and 87Sr absolute optical frequencies. *Metrologia*, 58(1):015017, jan 2021.
- [107] Z. Chen. Breaking Quantum Limits with Collective Cavity-QED: Generation of Spin Squeezed States via Quantum Non-Demolition Measurements. PhD thesis, Boulder, CO 80309-0440, 2013.

- [108] D. J. Wineland, J. J. Bollinger, W. M. Itano, F. L. Moore, and D. J. Heinzen. Spin squeezing and reduced quantum noise in spectroscopy. *Phys. Rev. A*, 46:R6797–R6800, Dec 1992.
- [109] Masahiro Kitagawa and Masahito Ueda. Squeezed spin states. *Phys. Rev. A*, 47:5138–5143, Jun 1993.
- [110] Jürgen Appel, Patrick Joachim Windpassinger, Daniel Oblak, U Busk Hoff, Niels Kjærgaard, and Eugene Simon Polzik. Mesoscopic atomic entanglement for precision measurements beyond the standard quantum limit. *Proceedings of the National Academy of Sciences*, 106(27):10960–10965, 2009.
- [111] Ian D. Leroux, Monika H. Schleier-Smith, and Vladan Vuletić. Implementation of cavity squeezing of a collective atomic spin. *Phys. Rev. Lett.*, 104:073602, Feb 2010.
- [112] Ian D. Leroux, Monika H. Schleier-Smith, and Vladan Vuletić. Orientation-dependent entanglement lifetime in a squeezed atomic clock. *Phys. Rev. Lett.*, 104:250801, Jun 2010.
- [113] B. C. Nichol, R. Srinivas, D. P. Nadlinger, P. Drmota, D. Main, G. Araneda, C. J. Ballance, and D. M. Lucas. An elementary quantum network of entangled optical atomic clocks. *Nature*, 609(7928):689–694, 2022.
- [114] Christian D. Marciniak, Thomas Feldker, Ivan Pogorelov, Raphael Kaubruegger, Denis V. Vasilyev, Rick van Bijnen, Philipp Schindler, Peter Zoller, Rainer Blatt, and Thomas Monz. Optimal metrology with programmable quantum sensors. *Nature*, 603(7902):604–609, 2022.
- [115] W. Wasilewski, K. Jensen, H. Krauter, J. J. Renema, M. V. Balabas, and E. S. Polzik. Quantum noise limited and entanglement-assisted magnetometry. *Phys. Rev. Lett.*, 104:133601, Mar 2010.
- [116] Graham P. Greve, Chengyi Luo, Baochen Wu, and James K. Thompson. Entanglement-enhanced matter-wave interferometry in a high-finesse cavity. *Nature*, 610(7932):472–477, 2022.
- [117] Serge Haroche. Nobel lecture: Controlling photons in a box and exploring the quantum to classical boundary. *Rev. Mod. Phys.*, 85:1083–1102, Jul 2013.
- [118] Gerhard Rempe, Herbert Walther, and Norbert Klein. Observation of quantum collapse and revival in a one-atom maser. *Phys. Rev. Lett.*, 58:353–356, Jan 1987.
- [119] C. J. Hood, M. S. Chapman, T. W. Lynn, and H. J. Kimble. Real-time cavity QED with single atoms. *Phys. Rev. Lett.*, 80:4157–4160, May 1998.
- [120] S Weyers, V Gerginov, M Kazda, J Rahm, B Lipphardt, G Dobrev, and K Gibble. Advances in the accuracy, stability, and reliability of the PTB primary fountain clocks. *Metrologia*, 55(6):789–805, oct 2018.
- [121] H J Kimble. Strong interactions of single atoms and photons in cavity QED. *Physica Scripta*, 1998(T76):127, jan 1998.
- [122] T. Holstein and H. Primakoff. Field dependence of the intrinsic domain magnetization of a ferromagnet. *Phys. Rev.*, 58:1098–1113, Dec 1940.

- [123] Dorian Gangloff, Molu Shi, Tailin Wu, Alexei Bylinskii, Boris Braverman, Michael Gutierrez, Rosanna Nichols, Junru Li, Kai Aichholz, Marko Cetina, Leon Karpa, Branislav Jelenković, Isaac Chuang, and Vladan Vuletić. Preventing and reversing vacuum-induced optical losses in high-finesse tantalum (v) oxide mirror coatings. *Opt. Express*, 23(14):18014–18028, Jul 2015.
- [124] Sara Campbell. A Fermi-degenerate three-dimensional optical lattice clock. PhD thesis, Boulder, 2017-05 2017.
- [125] Makfir Sefa, James A. Fedchak, and Julia Scherschligt. Investigations of medium-temperature heat treatments to achieve low outgassing rates in stainless steel ultrahigh vacuum chambers. *Journal of Vacuum Science & Technology A*, 35(4):041601, 2017.
- [126] Jerome Lodewyck, Michal Zawada, Luca Lorini, Mikhail Gurov, and Pierre Lemonde. Observation and cancellation of a perturbing dc stark shift in strontium optical lattice clocks. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 59(3):411–415, 2012.
- [127] Vladimir Schkolnik, Jason R. Williams, and Nan Yu. Generating 500 mW for laser cooling of strontium atoms by injection locking a high power laser diode, 2020.
- [128] Matthew A. Norcia and James K. Thompson. Simple laser stabilization to the strontium 88sr transition at 707 nm. *Review of Scientific Instruments*, 87(2):023110, 2016.
- [129] M. Bishof, X. Zhang, M. J. Martin, and Jun Ye. Optical spectrum analyzer with quantum-limited noise floor. *Phys. Rev. Lett.*, 111:093604, Aug 2013.
- [130] Ichiro Ushijima, Masao Takamoto, Manoj Das, Takuya Ohkubo, and Hidetoshi Katori. Cryogenic optical lattice clocks. *Nature Photonics*, 9(3):185–189, 2015.
- [131] Thomas H. Loftus, Tetsuya Ido, Martin M. Boyd, Andrew D. Ludlow, and Jun Ye. Narrow line cooling and momentum-space crystals. *Phys. Rev. A*, 70:063413, Dec 2004.
- [132] S. Blatt, J. W. Thomsen, G. K. Campbell, A. D. Ludlow, M. D. Swallows, M. J. Martin, M. M. Boyd, and J. Ye. Rabi spectroscopy and excitation inhomogeneity in a one-dimensional optical lattice clock. *Phys. Rev. A*, 80:052703, Nov 2009.
- [133] Ali Al-Masoudi, Sören Dörscher, Sebastian Häfner, Uwe Sterr, and Christian Lisdat. Noise and instability of an optical lattice clock. *Phys. Rev. A*, 92:063814, Dec 2015.
- [134] Alexander Aeppli, Anjun Chu, Tobias Bothwell, Colin J. Kennedy, Dhruv Kedar, Peiru He, Ana Maria Rey, and Jun Ye. Hamiltonian engineering of spin-orbit-coupled fermions in a Wannier-Stark optical lattice clock. *Science Advances*, 8(41):eadc9242, 2022.
- [135] P. Kómár, E. M. Kessler, M. Bishof, L. Jiang, A. S. Sørensen, J. Ye, and M. D. Lukin. A quantum network of clocks. *Nature Physics*, 10(8):582–587, 2014.
- [136] B M Roberts, P Delva, A Al-Masoudi, A Amy-Klein, C Bærentsen, C F A Baynham, E Benkler, S Bilicki, S Bize, W Bowden, J Calvert, V Cambier, E Cantin, E A Curtis, S Drscher, M Favier, F Frank, P Gill, R M Godun, G Grosche, C Guo, A Hees, I R Hill, R Hobson, N Huntemann, J Kronjger, S Koke, A Kuhl, R Lange, T Legero, B Lipphardt, C Lisdat, J Lodewyck, O Lopez, H S Margolis, H Álvarez-Martínez, F Meynadier, F Ozimek, E Peik,

- P-E Pottie, N Quintin, C Sanner, L De Sarlo, M Schioppo, R Schwarz, A Silva, U Sterr, Chr Tamm, R Le Targat, P Tuckey, G Vallet, T Waterholter, D Xu, and P Wolf. Search for transient variations of the fine structure constant and dark matter using fiber-linked optical atomic clocks. New Journal of Physics, 22(9):093010, sep 2020.
- [137] Boris Braverman, Akio Kawasaki, and Vladan Vuleti. Impact of non-unitary spin squeezing on atomic clock performance. New Journal of Physics, 20(10):103019, oct 2018.
- [138] Jiazhong Hu, Wenlan Chen, Zachary Vendeiro, Hao Zhang, and Vladan Vuletić. Entangled collective-spin states of atomic ensembles under nonuniform atom-light interaction. Phys. Rev. A, 92:063816, Dec 2015.
- [139] C. W. Gardiner and M. J. Collett. Input and output in damped quantum systems: Quantum stochastic differential equations and the master equation. Phys. Rev. A, 31:3761–3774, Jun 1985.

Appendix A

Vacuum noise limits of cavity frequency determination

In this thesis, we have estimated the frequency of our probe laser with respect to the cQED cavity in various ways. One method involves sweeping the probe tone over the cavity and measuring the full IQ response using heterodyne detection. When possible, we have also parked on resonance and measured the \mathcal{Q} response using heterodyne. In the latest work, we switched to homodyne detection of the probe laser and performed both sweeps and fixed-tone probing, where the homodyne detection phase has been chosen to detect the \mathcal{Q} quadrature. In the following, we examine the full $\mathcal{I} - \mathcal{Q}$ response of the cavity leakage field, and use it to calculate the vacuum noise limit on determining the frequency of the cavity (or cavity-like) mode.

The field that leaks from the cavity has the form

$$I_0 = \sqrt{R_{inc}} \left(1 - \frac{\beta_r}{1 + \left(\frac{\delta}{\kappa_p/2}\right)^2} \right) \quad (\text{A.1})$$

$$Q_0 = \sqrt{R_{inc}} \left(\frac{\beta_r \left(\frac{\delta}{\kappa_p/2}\right)}{1 + \left(\frac{\delta}{\kappa_p/2}\right)^2} \right) \quad (\text{A.2})$$

Here R_{inc} is the rate of incident photons, κ_p is the probed FWHM of the cavity mode, and δ is the detuning of the laser with respect to the cavity mode. It is convenient to shift the \mathcal{I} quadrature such that the phasor is centered on the origin. We define our shifted \mathcal{I} quadrature as

$$I_p = I_0 - \sqrt{R_{inc}} \left(1 - \frac{\beta_r}{2} \right) \quad (\text{A.3})$$

Note that the length of the phasor from the center of the centered circle to the edge is $\frac{\beta_r}{2}\sqrt{R_{inc}T}$. Also note that β_r is the diameter of the circle, but for estimating the signal to noise, we require the radius. Estimating the detuning of the cavity probe with respect to the cavity mode amounts to estimating the angle of the phasor

$$\Psi = \arctan\left(\frac{Q_0}{I_p}\right). \quad (\text{A.4})$$

A.0.1 Heterodyne detection

We first consider performing heterodyne detection, and perform a sweep over the cavity resonance. We consider the local sensitivity to how Ψ changes with a small change in the detuning δ by calculating the derivative,

$$\frac{d\Psi}{d\delta} = -\frac{4\delta\kappa_p}{4\delta^2 - \kappa_p^2}. \quad (\text{A.5})$$

If the total number of photons incident is $M_{inc} = R_{inc}dt$, then the radius of the circle is $\frac{\beta_r}{2}\sqrt{R_{inc}dt}$. Quantum optics teaches us that the vacuum noise creates fluctuations that will degrade our ability to estimate Ψ , having magnitude 1/2. So the uncertainty in Ψ is given by the ratio of the noise to the overall length of the phasor,

$$\Delta\Psi = \frac{1/2}{\frac{\beta_r}{2}\sqrt{R_{inc}dt}}. \quad (\text{A.6})$$

We can then calculate the uncertainty of the detuning as

$$\Delta\delta = -\frac{\Delta\Psi}{d\Psi/d\delta} = \frac{4\delta^2 + \kappa_p^2}{4\beta_r\kappa_p\sqrt{R_{inc}dt}}. \quad (\text{A.7})$$

We will think about sweeps of the probe tone across the cavity resonance as taking weighted averages. If we have several measurements, each with uncertainties $\sigma_1, \sigma_2, \dots, \sigma_n$, then the weighted averages is given by the inverse-squared sum $\frac{1}{\sigma_{tot}^2} = \frac{1}{\sigma_1^2} + \frac{1}{\sigma_2^2} + \dots + \frac{1}{\sigma_n^2}$. The swept measurement is thought of as spending a dwell time dt and making a frequency estimate, then we step the frequency and repeat. The final estimate is taken as the weighted sum of all of the measurements. We then express the final estimate as an integral, where we take the sweep rate as $\delta = \alpha t$,

$$\Delta\delta_{\text{sweep}} = \left(\int_{-\infty}^{\infty} \left(\frac{4(\alpha t)^2 + \kappa_p^2}{4\beta_r\kappa_p\sqrt{R_{inc}}} \right)^{-2} dt \right)^{-1/2} = \frac{\kappa_p^2}{2\kappa_1} \sqrt{\frac{\alpha}{\pi R_{inc}\kappa_p}}. \quad (\text{A.8})$$

We can use this expression to estimate the frequency noise corresponding to a single sweep over the cavity resonance. It requires careful calibration of the rate of incident photons R_{inc} . If we then want to take into account a finite quantum efficiency, one simply includes a $1/\sqrt{Q}$ term, meaning that a poor quantum efficiency will lead to a higher $\Delta\delta_{sweep}$.

A.0.2 Homodyne detection

Now let's consider a homodyne measurement of the cavity frequency. We will get the most information if we park directly on the resonance and measure the Q -response of the cavity leakage field. Taking the local derivative $\frac{dQ}{d\delta}$ to lowest order in δ gives

$$\frac{dQ}{d\delta} = \frac{\kappa_p^2}{4\kappa_1\sqrt{M_{inc}}}, \quad (\text{A.9})$$

where M_{inc} is the number of photons in a measurement window. The corresponding frequency uncertainty is given by

$$\Delta\delta_{resonant} = \frac{1/2}{dQ/d\delta} = \frac{\kappa^2}{4\sqrt{M_{inc}}\kappa_1}. \quad (\text{A.10})$$

Appendix B

Calculation of effective atom-cavity coupling

The effective atom-cavity coupling g is independently estimated as a consistency check on our experimentally determined value from Figure 6.6a. We follow the convention from [138], where the effective g is

$$g^2 = \frac{\langle g_i^4 \rangle}{\langle g_i^2 \rangle}, \quad (\text{B.1})$$

where g_i is the atom-cavity coupling for the i^{th} atom. The effective atom number is

$$N = N_{\text{tot}} \frac{\langle g_i^2 \rangle^2}{\langle g_i^4 \rangle} \quad (\text{B.2})$$

where N_{tot} is the total atom number. In this work, when we refer to g and N , we refer to the effective quantities. Our coordinate system is defined such that X is the direction along the cavity axis, Y is the other horizontal axis orthogonal to X , and Z is the vertical direction along gravity. The atomic density distribution in Z is modeled as a Gaussian, $\rho_Z(Z) = \frac{N}{\sqrt{2\pi}\sigma_Z} e^{-Z^2/(2\sigma_Z^2)}$, with standard deviation σ_Z . The probability distribution along Y is $P_Y(Y) = \frac{1}{\sqrt{2\pi}\sigma_Y} e^{-Y^2/(2\sigma_Y^2)}$ is a Gaussian with a standard deviation σ_Y set by the thermal cloud radius. We calculate σ_Y from the radial temperature of $T_r = 290(10)$ nK from a radial Doppler scan of the clock transition and the radial trap frequency of $34(3)$ Hz, giving a $\sigma_Y \approx 26\mu\text{m}$. If we allow the atoms to sufficiently time-average the standing wave along x , we have

$$g_i^2(Y, Z) = \frac{g_0^2}{2} e^{-2(Y^2+Z^2)/w_0^2} \quad (\text{B.3})$$

with peak coupling $g_0 = d_0 \sqrt{\frac{\omega_p}{2\epsilon_0 \hbar V}} = 2\pi \times 8.6$ kHz. Here, d_0 is the dipole matrix element and ω_p is the angular frequency of the $|\downarrow\rangle \rightarrow |e\rangle$ transition, and $V = \frac{1}{4}\pi\omega_0 L$ is the effective cavity mode

volume. The ensemble averages are evaluated as

$$\langle g_i^2 \rangle = \frac{1}{N} \int g_i^2(\mathbf{Y}, \mathbf{Z}) \rho_{\mathbf{Z}}(\mathbf{Z}) P_{\mathbf{Y}}(\mathbf{Y}) d\mathbf{Y} d\mathbf{Z} \quad (\text{B.4})$$

and

$$\langle g_i^4 \rangle = \frac{1}{N} \int g_i^4(\mathbf{Y}, \mathbf{Z}) \rho_{\mathbf{Z}}(\mathbf{Z}) P_{\mathbf{Y}}(\mathbf{Y}) d\mathbf{Y} d\mathbf{Z}, \quad (\text{B.5})$$

and when combined using Eq. B.1 gives the estimated value of $g = 2\pi \times 4.8(2)$ kHz.

Appendix C

Calculation of free space scattering

This section describes the calculation of free space scattering in the atom-cavity coupled system.

To describe the atom-cavity coupled system, we use an input-out formalism [139]. The Heisenberg-Langevin equations of motion are [32]

$$\frac{d\langle a \rangle}{dt} = -(i\delta_c + \frac{\kappa}{2})\langle \hat{a} \rangle - i\frac{\Omega_{\downarrow}}{2}\langle \hat{c} \rangle + \sqrt{\kappa_1}a_{\text{in}} \quad (\text{C.1})$$

$$\frac{d\langle b \rangle}{dt} = -\frac{\Gamma}{2} - i\sqrt{N}g\langle \hat{b} \rangle \quad (\text{C.2})$$

where $\langle a \rangle$ is the cavity photon operator and $\langle b \rangle$ is the atom operator. a_{in} is the drive amplitude, with units of $\sqrt{\text{photons/sec}}$. We note that the cavity operator has time-dependence as $\langle a \rangle = |a|e^{i\delta t}$, and similar for the atomic operator $\langle b \rangle = |b|e^{i\delta t}$. Here, δ is the detuning of the probe light with respect to the bare atomic transition frequency. We wish to solve the equations of motion for the steady-state values, and so take the derivative of each operator and obtain

$$i\langle a \rangle\delta = -(i\delta_c + \frac{\kappa}{2})\langle \hat{a} \rangle - i\frac{\Omega_{\downarrow}}{2}\langle \hat{c} \rangle + \sqrt{\kappa_1}a_{\text{in}} \quad (\text{C.3})$$

$$i\langle b \rangle\delta = -\frac{\Gamma}{2} - i\sqrt{N}g\langle \hat{b} \rangle \quad (\text{C.4})$$

We solve these coupled equations to determine the steady-state excited state population a . It then follows that the rate of scattered photons will be

$$\dot{M}_s = \Gamma|\langle a \rangle|^2 \quad (\text{C.5})$$

The solution has the form

$$|\langle a \rangle|^2 = \frac{4c_i^2 \kappa_1 \Omega_{\downarrow}}{(\Gamma^2 + 4\delta^2)(4(\delta_c + \delta)^2 + \kappa^2) + 2\Omega_{\downarrow}^2(\Gamma\kappa - 4\delta(\delta_c + \delta)) + \Omega_{\downarrow}^4} \quad (\text{C.6})$$

We can then plug in the appropriate vacuum Rabi splitting Ω_{\downarrow} , κ_1 , δ_c . One then plugs in the correct detuning δ to the eigenvalue, assuming we are probing on resonance. We require a good estimate of the incident photon number c_i , which we estimate by detuning the laser from the cavity resonance, and measuring the resulting homodyne fringe amplitude. We can calibrate the sideband power, and determine the mode overlap between the reflected mode and the LO mode.