

**Qubit Control with a Quadrupole-transition in Cs Atoms and Rydberg
Gate Laser System Design**

by

Jacob Scott

A dissertation submitted in partial fulfillment of
the requirements for the degree of

Doctor of Philosophy

(Physics)

at the

UNIVERSITY OF WISCONSIN–MADISON

2025

Date of final oral examination: 7th of August, 2025

The dissertation is approved by the following members of the Final Oral Committee:

Mark Saffman, Professor, Physics

Jen Choy, Associate Professor, Electrical and Computer Engineering

Deniz Yavuz, Professor, Physics

Matthew Otten, Assistant Professor, Physics

© Copyright by Jacob Scott 2025
All Rights Reserved

Para mi familia.

ACKNOWLEDGMENTS

I'd like to thank my advisor, Mark Saffman, for the opportunity to work on this project and for his guidance throughout my time in the lab. I'm deeply grateful to Trent Graham, Hong Ming Lim, Kais Jooya, Isaac Scott, Abraham Scott, Quinn Meece, Juan Bohorquez, Cody Pool(e), Sam Norrell, Linipun Phuttitarn, Preston Huft, Eunji Oh, and Jake Thompson—for their insight, help, and solidarity through the long days (and nights) of experiments and overheating laboratories.

Thank you to my family and close friends for their steady support over the years. I'm especially grateful to my brother Daniel, who encouraged me, challenged me, and believed in me when I needed it most.

There are people whose support and impact stayed with me and helped shape who I've become. Even if this page doesn't say your name—thank you.

CONTENTS

Contents	iii
List of Figures	vii
Abstract	xvii
I Introduction	1
1 Quantum Computing with Neutral Cesium	2
1.1 The AQuA Experiment	2
2 The $6S_{1/2} \rightarrow 5D_{5/2}$ Quadrupole Interaction in Cesium for Improved Quantum Control and Measurements	4
2.1 Motivation Behind the Quadrupole Interaction	4
2.2 Modeling the Quadrupole Interaction	6
2.3 Matrix Elements of the Quadrupole Interaction	8
2.4 Raman Depumping from Cycling on the Quadrupole Interaction	9
II Experiment	18
3 Experimental Apparatus and Design for Creating Qubits and Per- forming Quadrupole Control	19
3.1 The Experimental Apparatus	19
3.1.1 The 2D MOT	19
3.1.2 The 3D MOT	20
3.1.3 Single Atom Trap Design with Hole Array	24
3.1.4 Imaging single atoms onto the EMCCD	30
3.2 Developing the 685 nm Laser System	31

3.2.1	685 nm Requirements and Design	31
3.2.2	685 nm M2 Laser System	34
3.2.3	685 Injection Lock Amplification Setup	36
3.2.4	685 nm Locking Scheme	43
4	Cooling and Imaging Cs Qubits using the Quadrupole Interaction	53
4.1	Implementing Free-Space Second-Stage Cooling	53
4.1.1	685 nm MOT Experimental Sequence	53
4.1.2	685 nm Free-Space Push Spectroscopy	54
4.1.3	685 nm Free-Space Cooling	56
4.1.4	Comparison to Theory	59
4.1.5	685 nm MOT Atom Cloud Fluorescence	62
4.1.6	685 nm MOT Next Steps	62
4.2	Trapping Single Atoms	64
4.2.1	852 nm Imaging of Single Atoms	64
4.2.2	EMCCD Imaging Light Capture Efficiency	66
4.2.3	Measuring trap frequency	67
4.2.4	Optical Pumping Single Atoms into Qubits	68
4.2.5	Setting up Microwave Qubit Manipulations	73
4.3	685 nm Spectroscopy in the 803 nm Trap	74
4.3.1	Zeeman Effect Contribution	75
4.3.2	AC Stark Shift on the 685 nm Transition and Magic condition of 803 nm trap	76
4.3.3	Single Atom - 685 nm Spectroscopy Measurements	80
4.4	Implementing Trapped Single-Atom Cooling	82
4.4.1	Measuring atom temperature through T_2^*	83
4.4.2	Performing trap drop temperature measurements	85
4.5	Imaging and Quantum State Readout	89
4.5.1	685 nm Qubit Imaging	89
4.5.2	685 nm Raman Depump Rate	100
4.5.3	Paths to Improved Readout with 685 nm Imaging	101

5	Developing Low Noise Lasers for Rydberg Gates	103
5.1	The 459 nm Rydberg Laser	103
5.1.1	Architecture of 459 nm Laser System	103
5.1.2	Locking Scheme for the 459 nm Rydberg Laser	105
5.1.3	Measurements of the 459 nm Spectral Noise	110
5.2	The 1040 nm Rydberg Laser	114
5.2.1	RIN of the 1040 nm Rydberg Laser	115
5.2.2	Measurements of the 1040 nm Spectral Noise	115
5.2.3	S_{dv} of Rydberg Lasers	116
5.3	A small collection of thoughts and experiences on laser locks	118
5.3.1	Tuning up a ULE cavity	118
5.3.2	Lasers That Stay Locked	120
5.3.3	Auto-Relock for Continuous Operation	121
6	Developing Low Noise Electronic Systems for Laser and Current Control	125
6.1	The QPAL Lockbox	125
6.1.1	Design	125
6.1.2	Features	127
6.1.3	The Input Section	129
6.1.4	The Slow Loop	131
6.1.5	Rectifying the Slow Loop Output	136
6.1.6	The Fast Loop	137
6.1.7	Other Sections	139
6.1.8	Auto-Relock Capabilities	140
6.1.9	Future Improvements	141
6.2	The QPAL Coil-Driver	142
6.2.1	Design and Features	142
6.2.2	Determining Maximum Current Draw	145
6.2.3	Noise Performance	147
6.2.4	Future Improvements	148

III Next Steps and Conclusions	149
7 Next Steps and Outlook	150
7.1 Upgrades and Future work on the 685 nm Laser Systems . . .	150
7.1.1 Independent Beam Control for a 685 nm MOT	150
7.1.2 Upgrading the 685 nm Injection Lock Amplification Setup	151
7.1.3 Upgrading the 685 nm Source	152
7.2 Exploration of Quadrupole Cooling Performance	152
7.2.1 Loading a 685 nm MOT from a 2D MOT and Push Beam	153
7.3 Improving Quadrupole Imaging Performance	153
7.4 Outlook	154
References	156

LIST OF FIGURES

2.1	$6S_{1/2} \rightarrow 5D_{5/2} \rightarrow 6P_{3/2} \rightarrow 6S_{1/2}$ cycling and depumping diagram. Solid red lines indicate the primary cycling transition. Dashed red lines indicate paths that lead back to cycling after a Raman event. The dashed lavender lines indicate the Raman events that lead to the $F = 3$ state dark state. Finally, the thin solid orange lines show the decay paths that may lead back to the $F = 4$ cycling state if a future decay returns there (i.e. $5D_{5/2}F'' = 5 \rightarrow 6P_{3/2}F' = 4 \rightarrow 6S_{1/2}F = 4$	11
2.2	Logarithmic plot of the cycling rate over the Raman rate as a function of the detuning, for a few saturation parameters	17
2.3	Logarithmic plot of the cycling rate over the Raman rate as a function of the detuning, for a few saturation parameters	17
3.1	3D MOT Layout for the upper cell. 685 nm and 852 nm beams are combined to make a MOT with diagonal-top/bottom left/right-beams, as well as a pair of beams for the front and back.	20
3.2	An image of the upper cell chamber.	21
3.3	Bottom 685 nm+852 nm MOT light combining optics. Lasers are combined on a dichroic and sent to the bottom mirrors in figure 3.1.	23
3.4	4-f Imaging setup for trap array. Distances not to scale. Each lens is separated from the previous lens (along the beam path) by the sum of their focal lengths for 4-f imaging.	25

3.5	Analysis of trap shape and size in final image plane before the atom image plane. The cross-section plot's horizontal axes are scaled with a factor of $3.45 * 22.8/250$ $\mu\text{m}/\text{pixel}$ to demonstrate the beam size in the atom plane. The vertical axes are scaled to normalize the intensity of the gaussian fit. Top-Left: The raw data measured by the Thorlabs CCD camera with 3.45 μm per pixel. Top-Middle(Right): $x/(y)$ cross-sectional fits of the normalized intensity (peak intensity is equal to 1). Bottom-Left: A zoom in of the Top-Left image, centered around the hole in the trap. Bottom-Middle(Right): $x/(y)$ cross-sectional fits of the gaussian subtracted trap to a parabolic potential. Functional forms of fits are presented in the legends of the cross-section plots.	27
3.6	Trap shape analysis as seen in figure 3.5 but with temperature scaling in μK for the vertical axis, using the total beam power of 125 mW , gaussian beam waist of 20 μm , and the ground state scalar polarizability of -365×10^{-24} cm^3 (cgs units) at 803 nm to calibrate.	29
3.7	685 nm Laser System from a 15W Sprout pump at 532nm , an MSquared SolsTiS lasing at 685 nm , and a ULE cavity for locking the laser.	32
3.8	Optical layout of the 685 nm M2 laser system.	34
3.9	A picture of the 685 nm amplification stage in action.	37
3.10	Lasing wavelength vs temperature for the HL69001DG diode taken from the datasheet. At a case temperature of 0 Celsius, the lasing wavelength is 685 nm and the center wavelength at room temp is 690nm	38
3.11	Dew Point (C) vs air temperature based on the Magnus formula. Reproduced from Easchiff, Wikimedia Commons, CC BY-SA 4.0; [46].	39

3.12	Interference pattern of the injection-locked diode with the seed laser on a PICAM CCD camera for the unlocked laser (left) and locked laser (right). The interference pattern is used to detect an unlock event through Fourier transform image processing. . .	43
3.13	Basic feedback system with labels. A setpoint and a measured process variable are compared to create an error signal. A gain stage acts on the error signal to create a control variable, which is then applied to the process. The process then is measured to create the measured process variable.	44
3.14	Locking scheme for the 685 nm laser. The fast piezo loop is responsible for stabilizing the laser frequency, while the slow piezo loop ensures the fast loop remains locked over longer timescales by monitoring the fast loop control variable after passing through an LPF. The slow loop is effectively a double integrator, but constructed as a PI loop that drives the fast loop control variable to zero. There is also an EOM feedback loop that applies a P-LPF loop to the EOM.	45
3.15	685 nm Self-Heterodyne phase noise data, windowed fit, and extrapolated phase noise from the windowed fit. White noise: $h_0 = 8.4 \text{ Hz}^2/\text{Hz}$. Fit Servo Bumps at $f_g = 20. \text{ kHz}$, $\sigma_g = 920 \text{ Hz}$, $h_g = 15 \times 10^3 \text{ Hz}^2/\text{Hz}$ and $f_g = 57 \text{ kHz}$, $\sigma_g = 5 \text{ kHz}$, $h_g = 46 \text{ Hz}^2/\text{Hz}$. More bumps are clearly seen, but have much less contribution. Analyzed using the same tools and techniques presented in section 5.1.3.	48
4.1	Results of a 685 nm push spectroscopy scan with a 685 nm cooling phase, fit to a Gaussian. The measured atom temperature after the cooling phase was $5.4 \mu\text{K}$ as is discussed in section 4.1.3. The measured center frequency setting for the AOMs is 147.44MHz and the FWHM is 230kHz	55

4.2	Images of the atom cloud with and without the 685 nm push beam activated, as well as a subtraction between the two. a) on resonance, b) off resonance.	56
4.3	685 nm MOT atom cloud images for a 1.5 G/cm magnetic field gradient and a detuning of -1.25Γ . The time of flight is measured at 0, 3, 6, 9, 12, and 15 ms after the cooling phase. Each image is fit to a symmetric 2D Gaussian, with the x and y cross-sections also shown.	58
4.4	Measured atom cloud symmetric Gaussian radius vs. time for a 1.5 G/cm magnetic field gradient and a detuning of -1.25Γ . The points are fit to Equation 4.3 to extract the temperature of the atom cloud, which is $T = 5.3\ \mu\text{K}$	59
4.5	Measured temperature of the atom cloud vs. detuning for a 1.5 G/cm and a 0.25 G/cm (near molasses) magnetic field gradient. The minimum temperature occurs at approximately -1.25Γ detuning for both cases and increases drastically for detunings less than -1.0Γ	60
4.6	Theoretical calculations of the atom cloud temperature vs. detuning for 685 nm quadrupole cooling molasses, shown as a function of the dimensionless parameter $\mathbf{u} = \frac{2\delta\mathbf{s}}{\omega_r} = 77\frac{\delta}{\gamma}\mathbf{s}$ for (a) $\sigma^+ - \sigma^-$ and (b) $\text{lin}\perp\text{lin}$ polarization schemes. Figure reproduced from [71]. Vertical axes are in units of the photon recoil energy, $\hbar\omega_r = 0.154\ \mu\text{K}$. The values of \mathbf{u} for $\mathbf{s} = 1.7$ at detunings $-\gamma/2$, $-\gamma$, and -2γ are 66, 132, and 264, respectively.	61
4.7	Background-free atom cloud fluorescence image using only 685 nm light. The image exposure time was 100 ms to acquire sufficient photons on the camera. Pixel size is $53\ \mu\text{m}/\text{pixel}$. No repump light was used for this image, demonstrating the high pump-depump ratio of the 685 nm light.	63

- 4.8 852 nm imaging statistics for single atoms in the trap array. The images are taken with a -5.6γ detuning, a 70 ms exposure time, and a total intensity of $2.5 I_{\text{sat}}$ across all six beams. A 2x2 charge binning is applied on the EMCCD camera before the analog measurement, followed by a 3x3 digital binning step on the analysis computer. (a) A histogram of counts for shot 0. (b) A scatter plot where each point represents data from a single atom shot; the x-axis shows photoelectron counts in the first shot, and the y-axis shows photoelectron counts in the second shot. The dashed lines indicate threshold locations used to determine whether an atom is present or vacant. 65
- 4.9 Measured trap frequency from the trap drop technique. The time delay between two trap drops is scanned to extract the trap frequency. The data is fit to a sinusoidal function of the form $A \sin(2\omega t + \phi)$, and the trap frequency, ω is directly extracted from the fit (since the fit includes the factor of 2 in the frequency). The measured trap frequency is $\omega = 2\pi \times 10.55$ kHz, which is very close to our calculated values of 10.9 kHz and 9.9 kHz. . . 69
- 4.10 Measured optical pumping (a) and depump (b) rates of the optical pumping process. The $1/e$ time of the pumping is 0.2 ms and the depumping is 80 ms. The ratio of the pump-depump rate is often used as a figure of merit to determine the quality of the optical pumping. In our case, this ratio is 400. 73
- 4.11 AC Stark shift for the $6S_{1/2}, F = 4$ state and the $5D_{5/2}, F = 6$ state at 803 nm. The plot shows the scalar polarizability for both states (solid lines), as well as the total polarizability, $\alpha_{nJF,5D_{5/2}}^{\text{total}}$, for each m_F level (dashed lines). The points where $\alpha_{nJF,5D_{5/2}}^{\text{total}}$ intersects with the $6S_{1/2}$ are the magic conditions on the AC-Stark Shift. At 803nm, we have the $m_F = 6$ magic condition. . . 79

- 4.12 Spectroscopy results for the $6S_{1/2}, F = 4 \rightarrow 5D_{5/2}, F' = 6$ transition in the 803 nm trap. The three scans correspond to different powers of the 803 nm trap light, demonstrating the effects of the AC Stark shift on the 685 nm transition. The "Full power" scan is performed with the full power of the 803 nm trap light, the "1/3 Power" is performed with 41.7 mW of 803 nm trap light and the "Full Power - High Background" is performed with a different trap that has a larger amount of background light in the image plane, causing a large amount of trap shift. As can be seen, there is a strong correlation between the amount of 803 nm light seen by the trapped atom and the broadening of the 685 nm transition. 81
- 4.13 Fit Ramsey oscillations for varying gap times between $\pi/2$ pulses. The oscillations are fit to a cosine function with an amplitude and phase, $\frac{1}{2}(1 + A_t \cos(\phi + \phi_t))$. This is then used to extract the T_2^* time by fitting the amplitude A_t to a decaying exponential function in figure 4.14. 85
- 4.14 Decay of the Ramsey oscillation amplitude A_t as a function of the gap time between $\pi/2$ pulses. The data is fit to a decaying exponential function, $A_t = A_0 e^{-t/T_2^*}$, allowing for the extraction of $T_2^* = 7.6(3)\text{ms}$ 86
- 4.15 Measured retention rate as a function of trap drop time for 685 nm cooling pre- and post-optical pumping, compared with Monte Carlo simulation results for different temperatures. We find that the temperature of the atoms in the trap before optical pumping is around $3.5\mu\text{K}$ by calibrating the simulation with the T_2^* measurement. 88
- 4.16 Average shot of the atom using 685 nm imaging. The 5×5 ROI used to measure the pixel brightness is with a rectangle. This atom is very clearly aberrated and much larger than ideal. . . . 91

- 4.17 Logarithmic histogram of the 685 nm imaging data, with each shot of the 9 images combined into a single histogram, for a total of 45k points. Histogram is fit to the bi-modal distribution described in section 4.5.1. The black dashed line is the result of the bimodal fit, and the blue and green dashed lines are the "no atom" and "atom" modes of the fit, separated. The vertical red line is the location of the threshold determined by the bimodal fit, which is used to classify the qubit state. 93
- 4.18 Scatter plot of the 685 nm imaging data, with each point representing a pair of counts from the first (x-axis) and second (y-axis) shots. The numerically determined threshold (see 4.5.1) is shown as the dashed vertical and horizontal lines. This figure is of the same style as figure 4.8. 94
- 4.19 Scatter plots of the 685 nm imaging data, with each frame presenting an individual shot of the 8 shots after the first shot. The x-axis of the scatter is the number of photoelectrons seen in the first shot and the y-axis is the number of photoelectrons seen in the nth shot. 95
- 4.20 Retention rate of the 685 nm imaging data after each shot. The blue curve presents the total number of atoms retained after each shot with respect to the first shot, while the orange curve presents the number of atoms retained for each shot with respect to the number of atoms in the previous shot. As can be seen, the orange shot has a very constant retention rate of 99.1(1)% across all shots. This rate is not increasing either, which indicates that the atoms are not being heated during the repetitive imaging process. 96

4.21	Classification fidelity of the 685 nm imaging data as a function of the threshold point used for the bimodal fit. The solid blue line is the classification fidelity for the best-fit bimodal distribution, while the shaded blue region is the 1-sigma uncertainty region of the classification fidelity. The dashed red line is the threshold that maximizes the classification fidelity, which is found to be at a photoelectron count of 35.4 photoelectrons. The classification fidelity is found to be 0.9993(4).	98
4.22	Zoomed in version of figure 4.21, allowing for visualization of the classification fidelity and uncertainty.	99
4.23	Diagram of the 685 nm imaging sequence used to measure the Raman depump rate and the vacuum limited lifetime of the atoms in the trap. The blue curve presents the total loss ratio of the atoms, measured at the first shot after the pulse. The red curve presents the ratio of atoms depumped into the $F = 3$ state, based on the ratio of atoms found again in the second shot after the pulse. The orange curve presents the difference between the blue and red curves, which is the ratio of atoms lost due to the vacuum limited lifetime. From fitting these curves, the total loss $1/e$ lifetime is 20(1)s, the Raman depump lifetime is 42(4)s, and the vacuum limited lifetime is 43(5)s.	102
5.1	Image of 459 nm laser system, locked to the SHG cavity.	103
5.2	RIN comparison with the noise eater enabled and disabled for the 459 nm laser. Units are in Fractional Noise Spectral Density, $1/\sqrt{\text{Hz}}$	106
5.3	918 Locking Scheme Functional Diagram.	106
5.4	Moglabs FSC fast loop mod R91 location.	108
5.5	Moglabs FSC slow loop modification R87 location.	109

5.6	918 nm Self-Heterodyne phase noise data, windowed fit, and extrapolated phase noise from the windowed fit. White noise: $0.5 \text{ Hz}^2/\text{Hz}$. Fit Servo Bumps at $f_g = 6.2 \text{ kHz}$, $\sigma_g = 910 \text{ Hz}$, $h_g = 72 \text{ Hz}^2/\text{Hz}$, $f_g = 51 \text{ kHz}$, $\sigma_g = 40 \text{ kHz}$, $h_g = 1.2 \text{ Hz}^2/\text{Hz}$, and $f_g = 145 \text{ kHz}$, $\sigma_g = 5.8 \text{ kHz}$, $h_g = 1.1 \text{ Hz}^2/\text{Hz}$	114
5.7	Ti:Sa absorption and emission spectrum, from [95]. The emission gain at 1040 nm falls off from the peak gain at 800 nm by almost an order of magnitude.	115
5.8	RIN comparison between the 459 nm (with noise eater) and the 1040 nm Rydberg lasers. The free-running noise of the 1040 nm laser is less than that of the 459 nm laser with the noise eater.	116
5.9	1040 nm self-heterodyne phase noise data, windowed fit, and extrapolated phase noise from the windowed fit. White noise: $6.6 \text{ Hz}^2/\text{Hz}$. Fit servo bumps at $f_g = 69.1 \text{ kHz}$, $\sigma_g = 40.6 \text{ kHz}$, $h_g = 11 \text{ Hz}^2/\text{Hz}$ and $f_g = 234 \text{ kHz}$, $\sigma_g = 1.3 \text{ kHz}$, $h_g = 1.7 \times 10^3 \text{ Hz}^2/\text{Hz}$	117
5.10	The S_{dv} of both Rydberg lasers based on fitting white noise and servo bumps to the measured self-heterodyne spectra, following the analysis found in [93]. See Figures 5.6 and 5.9 for fit parameters.	117
5.11	Figures taken from [93], presenting Rabi error vs. laser noise parameters for a 2π rotation: a) Rabi error vs. white phase noise, h_0 . b) Rabi error vs. servo bump frequency (f_g with amplitude of $h_g = 10^3 \text{ Hz}^2/\text{Hz}$). For both situations, the Rabi frequency used in the simulation is $\Omega_0/2\pi = 1 \text{ MHz}$. See the reference for more details.	118
6.1	Image of the populated PCB of a single channel of the QPAL Lockbox.	125
6.2	Schematic diagram of Ch 1 of the QPAL Lockbox.	127
6.3	Input Processing Schematic Portion	129
6.4	Slow Loop Schematic Portion	130

6.5	Fast Loop Schematic Portion	138
6.6	Image of the frontpanel of the QPAL Coil Driver.	142
6.7	Feedback Control Schematic of Coil driver	144
6.8	MOSFET Power Dissipation for a 1 Ohm coil and 8.5V V_{HC} . . .	146
6.9	Current Noise Spectral Density at 1A for the coil driver. The current noise spectral density is less than 1ppm. The noise below 1kHz is dominated by the noise of the input signal, including the spikes on the spectrum.	147

ABSTRACT

Neutral atom quantum computing platforms depend on precise control and measurement of atomic qubits to realize high-fidelity operations at scale. This thesis presents a set of experimental and engineering contributions that address critical requirements on laser systems and light-atom interactions for advancing quantum control of cesium atoms.

It details the development of low-noise electronic feedback systems for laser frequency locking and magnetic field stabilization, engineered to support stable and low-noise experimental operation. It further introduces the design, implementation, and noise characterization of narrow-linewidth laser systems driving Rydberg-level transitions, which underpin high-fidelity two-qubit gates via the Rydberg blockade mechanism.

The work also establishes the electric quadrupole transition in cesium as a powerful tool for state-selective, background-free quantum state readout, achieving a classification fidelity of 0.9993 and an atom survival probability of 0.991. Additionally, it demonstrates that laser cooling on this transition effectively reduces post-optical pumping atom temperatures to 5.4 μK .

Collectively, these results advance the frontier of robust, high-performance control in cesium-based quantum computing and significantly expand the capabilities of neutral atom architectures.

Part I

Introduction

1 QUANTUM COMPUTING WITH NEUTRAL CESIUM

Quantum computing is a rapidly growing field that aims to utilize the fundamentals of quantum mechanics to perform computations that are ultimately useful for the world. There are many different approaches to quantum computing, including superconducting qubits [1, 2], trapped ions [3, 4], quantum dots [5, 6], and neutral atoms [7–10]. While all these approaches have their own advantages and disadvantages, neutral atoms are particularly promising due to their long coherence times, scalability, and the general fun of building a quantum computer out of lasers and individual atoms.

Two major challenges in using the neutral atom approach to quantum computing are improving the fidelity of quantum gates [11–15] and increasing the speed of readout [16–19]. This thesis presents my work and the work of my collaborators on building towards these goals. There are three main components in this thesis: (1) improving the fidelity of Rydberg blockade gates through the development of low noise Rydberg lasers, (2) improving the fidelity and capabilities of quantum state measurement through the implementation of a non-destructive, background-free readout scheme using the $6S_{1/2} \rightarrow 5D_{5/2}$ quadrupole interaction, and (3) developing low noise electronics and systems to accomplish these goals.

1.1 The AQuA Experiment

The AQuA experiment is a neutral atom quantum computing experiment using Cesium atoms at the University of Wisconsin-Madison. This experiment has published many works on our approach to quantum computing [10, 20–23], through which the schemes and components of the quantum computer architecture have been developed and described. Discussing all of these components, architecture, approaches, etc., is beyond the scope of this thesis, since it is not the focus of this work. However, this work does

build closely on the design of laser systems required for Rydberg gates and potential readout schemes, which will be described here.

In particular, the AQuA experiment uses a Rydberg blockade gate [11, 24–26] to perform 2-qubit gates between neutral atoms, necessary for the implementation of quantum algorithms. This gate works by exciting an atom into a Rydberg state, where a neighboring atom in close proximity would experience a perturbative van der Waals interaction that shifts the energy of its Rydberg states. This perturbative potential can be used to impart a phase shift on the neighboring atom, which can be used to perform a controlled phase gate.

However, for this to work, the atoms need to be excited into the Rydberg state. For this, a two-photon excitation is performed, where the $7P_{1/2}$ state is used as the intermediate state before being excited into the $75S_{1/2}$ Rydberg state [10, 22, 24]. These two lasers are at 459 nm ($6S_{1/2} \rightarrow 7P_{1/2}$) and 1040 nm ($7P_{1/2} \rightarrow 75S_{1/2}$). One component of this thesis involves the design and characterization of these lasers, which is described in Chapter 5.

2 THE $6S_{1/2} \rightarrow 5D_{5/2}$ QUADRUPOLE INTERACTION IN CESIUM FOR IMPROVED QUANTUM CONTROL AND MEASUREMENTS

The primary focus of my work in this thesis is the development of a non-destructive, background-free readout and cooling scheme using the $6S_{1/2} \rightarrow 5D_{5/2}$ quadrupole interaction in Cesium. This interaction is a second order electric quadrupole interaction, which allows for excitation to the $5D_{5/2}$ state from the $6S_{1/2}$ state, which is impossible with a standard electric dipole interaction. From a second-order analysis of the light-atom interaction, the quadrupole interaction allows for two units of angular momentum to be exchanged between the atom and the light field.

Because of this second order interaction, it is much weaker than the standard electric dipole interaction, and as such, doesn't primarily decay directly to the $6S_{1/2}$ state, but rather to the $6P_{3/2}$ state first, through standard electric dipole spontaneous emission. As such, a closed loop cycle can be implemented where the atom is excited through $6S_{1/2}, F = 4 \rightarrow 5D_{5/2}, F'' = 6 \rightarrow 6P_{3/2}, F' = 5 \rightarrow 6S_{1/2}, F = 4$, with the $6S_{1/2}, F = 3$ dark state being forbidden due to electric dipole selection rules. Section 2.2 describes and analyzes the theory behind this interaction.

2.1 Motivation Behind the Quadrupole Interaction

The first key point of interest in this interaction is the high pump-depump ratio introduced because of the decay selection rules and narrow effective linewidth ($\gamma = < 2\pi \times 118$ kHz) compared to the hyper-fine splitting (hfs) of the $5D_{5/2}$ state, ($\Delta_{5D_{5/2}, \text{hfs}} = 2\pi \times 126$ MHz between the $F'' = 5$ and $F'' = 6$ states). Section 2.4 takes a deep dive into analyzing and calculating

this ratio, finding it to be of the order 1.6×10^7 . The primary reason for the interest in this is the ability to scatter photons along this line while preventing the information of the atom's state from being lost by decaying into the $F = 3$ dark state.

Another point of interest in this interaction is its capabilities for being utilized for background-free imaging. Standard imaging of trapped Cesium atoms is performed using the D2 line at 852nm, $6S_{1/2}, F = 4 \rightarrow 6P_{3/2}, F' = 5$, which is a standard electric dipole interaction. In fact, this interaction is the primary interaction used for laser cooling of the Cesium atoms as well. However, while imaging, the atoms are illuminated with a large amount of near-resonant light, which always finds itself making its way into the EMCCD cameras, causing large amounts of background noise. Since the $6S_{1/2}, F = 4 \rightarrow 5D_{5/2}, F'' = 6$ quadrupole interaction decays through this state, the photons scattered by the atom will be at a completely different wavelength (852nm) from the cycling pump light (685nm), which can easily be filtered out using standard optical band-pass filters. This allows for a background-free imaging scheme. Background free imaging has been performed on alkali atom clouds [27–31], as well as with single trapped atoms[32]. But these measurements did not also have state-selective readout.

The last key point of interest in this transition is the low Doppler temperature, which is a consequence of its narrow linewidth, $\gamma = 2\pi \times 118$ kHz. The Doppler temperature is given by [33]:

$$T_D = \frac{\hbar\gamma}{2k_B} \quad (2.1)$$

where \hbar is the reduced Planck's constant, k_B is the Boltzmann constant, and γ is the linewidth of the transition. For our case, this gives a Doppler temperature of $T_D = 3$ μ K. The Doppler temperature is a laser cooling limit for a given atomic transition, based on the linewidth of the transition. It is the temperature where an equilibrium is reached where the cooling rate

becomes equal to the heating rate caused by the linewidth of the transition.

In this thesis, we utilize the 685nm transition to implement improved laser cooling of the Cs atoms and a non-destructive background-free imaging scheme that also cools the atoms during the process.

2.2 Modeling the Quadrupole Interaction

The dynamics of the atomic transition $6S_{1/2}, F = 4 \rightarrow 5S_{5/2}, F = 6$ can be modeled through a quadrupole interaction model. In this model, the oscillating electric field ($\mathbf{E} = |E|e^{i(\mathbf{k}\cdot\mathbf{r}-\omega t)}\hat{\mathbf{u}}$), interacts with the atom through the AC Stark Effect ($\mathcal{H} = -\mathbf{d} \cdot \mathbf{E} = e\mathbf{r} \cdot \mathbf{E}$) where \mathbf{d} is the electric dipole moment, e is the electron charge, \mathbf{r} is the position vector of the electron, and $\hat{\mathbf{u}}$ is the polarization vector of the electric field. The interaction Hamiltonian can be written as:

$$\mathcal{H} = e|E|(\mathbf{r} \cdot \hat{\mathbf{u}})e^{i(\mathbf{k}\cdot\mathbf{r}-\omega t)} \quad (2.2)$$

Here, we can expand the complex exponential of $\mathbf{k} \cdot \mathbf{r}$ as a series: $e^{i(\mathbf{k}\cdot\mathbf{r})} \approx 1 + i\mathbf{k} \cdot \mathbf{r}$. In a dipole interaction (E1), this term would be taken to the 0th order. This enforces angular momentum selection rules such that the maximal change in angular momentum is 1. However, if we include the first-order term, we introduce the quadrupole interaction (E2) which makes the Hamiltonian second order in \mathbf{r} , allowing for $\Delta_L = 2$ transitions. Our AC Stark Effect becomes:

$$\mathcal{H} = e(\mathbf{r} \cdot \hat{\mathbf{u}})|E|(1 + i\mathbf{k} \cdot \mathbf{r})e^{-i\omega t} \quad (2.3)$$

with $E = \mathcal{E}e^{-i\omega t}$. We can expand (2.3) as

$$\mathcal{H} = \mathcal{H}_{E1} + \mathcal{H}_{E2} \quad (2.4)$$

with:

$$\begin{aligned}\mathcal{H}_{E1} &= e|E|e^{-i\omega t}(\mathbf{r} \cdot \hat{\mathbf{u}}) \\ \mathcal{H}_{E2} &= \frac{e|E|\omega}{c}e^{-i\omega t}(\mathbf{r} \cdot \hat{\mathbf{u}})(\mathbf{r} \cdot \hat{\mathbf{k}})\end{aligned}\quad (2.5)$$

using $\mathbf{k} = \frac{\omega}{c}\hat{\mathbf{k}}$, where c is the speed of light.

$$\begin{aligned}\mathcal{H}_{E1} &\sim rY_{1q} \\ \mathcal{H}_{E2} &\sim r^2Y_{2q}\end{aligned}\quad (2.6)$$

where Y_{1q} and Y_{2q} are spherical harmonics of rank 1 and 2, respectively. The first term, \mathcal{H}_{E1} , is the electric dipole interaction, while the second term, \mathcal{H}_{E2} , is the electric quadrupole interaction.

Selection rules dictate that \mathcal{H}_{E1} cannot excite the $6S_{1/2} \rightarrow 5D_{5/2}$ transition. Therefore,

$$\langle 6S_{1/2} | \mathcal{H} | 5D_{5/2} \rangle = \langle 6S_{1/2} | \mathcal{H}_{E1} + \mathcal{H}_{E2} | 5D_{5/2} \rangle \quad (2.7)$$

$$= \langle 6S_{1/2} | \mathcal{H}_{E2} | 5D_{5/2} \rangle \quad (2.8)$$

$$(2.9)$$

leaving our Hamiltonian for the quadrupole interaction as:

$$\mathcal{H} = e\frac{\omega|E|}{c}e^{-i\omega t}(\mathbf{r} \cdot \hat{\mathbf{u}})(\mathbf{r} \cdot \hat{\mathbf{k}}) \quad (2.10)$$

Immediately, we can infer that the irreducible spherical tensor component must be of second order, and there is a second order interaction in r , i.e. $\mathcal{H} \sim r^2Y_{2q}$. For now, we can assume an AC Stark effect perturbation from an oscillation electric field will couple two atomic states with a Y_{2q} spherical tensor. The $(\mathbf{r} \cdot \hat{\mathbf{k}})$ component is the essential reason why the quadrupole interaction is feasible.

2.3 Matrix Elements of the Quadrupole Interaction

To calculate the matrix elements of the quadrupole interaction, we need to decompose the interaction Hamiltonian:

$$\mathcal{H} = e \frac{\omega |E|}{c} e^{-i\omega t} (\mathbf{r} \cdot \hat{\mathbf{u}})(\mathbf{r} \cdot \hat{\mathbf{k}}) \quad (2.11)$$

and more specifically, the term $(\mathbf{r} \cdot \hat{\mathbf{u}})(\mathbf{r} \cdot \hat{\mathbf{k}})$, into a radial dependent component and a spherical tensor component. Using the following relation from [34, 35] for commuting vectors A, B, A', B' :

$$(A \cdot B)(A' \cdot B') = \sum_{K=0,1,2} (-1)^K \{A \otimes A'\}_K \cdot \{B \otimes B'\}_K \quad (2.12)$$

where the q component of the tensor product notation is given by:

$$\{A \otimes B\}_{Kq} = \sum_{q_A, q_B} C_{1q_A, 1q_B}^{Kq} A_{q_A} B_{q_B} \quad (2.13)$$

with $C_{1q_A, 1q_B}^{Kq}$ being the Clebsch-Gordan coefficient. In our case of the quadrupole interaction, we can write:

$$(\mathbf{r} \cdot \hat{\mathbf{u}})(\mathbf{r} \cdot \hat{\mathbf{k}}) = \sum_{K=0,1,2} (-1)^K \{\mathbf{r} \otimes \mathbf{r}\}_K \cdot \{\hat{\mathbf{u}} \otimes \hat{\mathbf{k}}\}_K \quad (2.14)$$

However, we know that we only care about the $K = 2$ component since only the Y_{2q} components will couple the quadrupole interaction. This becomes:

$$(\mathbf{r} \cdot \hat{\mathbf{u}})(\mathbf{r} \cdot \hat{\mathbf{k}})_2 = \{\mathbf{r} \otimes \mathbf{r}\}_2 \cdot \{\hat{\mathbf{u}} \otimes \hat{\mathbf{k}}\}_2 \quad (2.15)$$

The $\{\hat{\mathbf{u}} \otimes \hat{\mathbf{k}}\}_2$ component is independent of the quantum state, and instead is just a factor that depends on the polarization and k -vector of the light.

It can be evaluated as:

$$\{\hat{\mathbf{u}} \otimes \hat{\mathbf{k}}\}_{2q} = \sum_{q_u, q_k} C_{1q_u, 1q_k}^{2q} \hat{\mathbf{u}}_{q_u} \hat{\mathbf{k}}_{q_k} \quad (2.16)$$

where q_u and q_k are the spherical components of the polarization and k-vector, respectively. However, since there is an inner product with the $\{\mathbf{r} \otimes \mathbf{r}\}_2$ component, any terms in $\{\hat{\mathbf{u}} \otimes \hat{\mathbf{k}}\}_{2q}$ that are zero will not contribute to the interaction since $\mathbf{A}_{\kappa q} \cdot \mathbf{B}_{\kappa q'} = \delta_{qq'}$.

The $\{\mathbf{r} \otimes \mathbf{r}\}_2$ component can be decomposed to:

$$\{\mathbf{r} \otimes \mathbf{r}\}_{2q} = r^2 \sqrt{\frac{2}{3}} \sqrt{\frac{4\pi}{5}} Y_{2q} \quad (2.17)$$

where Y_{2q} is the spherical harmonic of rank 2 and order q .

All together, the quadrupole interaction Hamiltonian can be written as:

$$\mathcal{H} = e \frac{\omega |E|}{c} e^{-i\omega t} \sqrt{\frac{2}{3}} \sqrt{\frac{4\pi}{5}} r^2 \sum_q Y_{2q} \{\hat{\mathbf{u}} \otimes \hat{\mathbf{k}}\}_{2q} \quad (2.18)$$

We will revisit this in section 2.4.

2.4 Raman Depumping from Cycling on the Quadrupole Interaction

In this section we calculate the pump-depump ratio of this cycling transition, where the dark state is the $6S_{1/2}, F = 3$ state. This ratio is an important figure of merit indicating how many times an average atom will cycle before being depumped to the dark state.

Although the cycle is perfectly closed, there is a possibility of a Raman event where the electron is excited to the $5D_{5/2}, F < 6$ state instead of $F = 6$. Even if the atom is excited to the $5D_{5/2}, F = 5$ state, it can still decay back

down to the $6S_{1/2}, F = 4$ state, completing the cycle since a decay to the $6P_{3/2}, F' = 5$ state is not forbidden. The branching ratios which govern the relative rates of these events are calculated in this section to obtain the pump-depump ratio.

The cycling transition is from the $6S_{1/2}F = 4 \rightarrow 5D_{5/2}F' = 6$ which primarily decays down to the $6P_{3/2}(F'' = 5)$ with a linewidth of $\gamma = 2\pi \times 118$ kHz, and will only decay directly to the $6S_{1/2}F = 4$ via a quadrupole decay with only a few Hz linewidth (i.e. negligible). The $6P_{3/2}(F'' = 5)$ then decays to the $6S_{1/2}(F = 4)$, completing the cycle. Raman events occur if the photon excites the atom to the $F'' = 2, 3, 4, 5$ states (with $F' = 1$ forbidden). An illustration of this can be found in figure 2.1.

To obtain the Raman rate, first, we must consider the rate of our atom exciting to each of the non-cycling F' states, and proceed to evaluating the probability that these states do not eventually decay to $F = 4$. The summation of each of these rates (down to $F = 3$) determines the Raman rate. The cycling rate is simply the rate of the atom exciting to the cycling $F' = 6$, since it will always decay down to $F = 4$, summed with the Raman rates ending in $F = 4$ (which should hopefully be a small contribution if the detuning is small).

To this end, a rate vector, $r_{F''}$, indicating the rate of excitation from $F = 4$ to F'' , will be calculated from the rate equation:

$$r_{F''} = \frac{\gamma}{2} \frac{I/I_{s,eff}}{1 + 4\Delta^2/\gamma^2 + I/I_{s,eff}} \quad (2.19)$$

Where γ is the linewidth, I is the field intensity, $I_{s,eff}$ is the effective saturation intensity (compensating for the degeneracy of the initial F and final F'' states), and Δ is the detuning of the oscillating field from the transition. The effective saturation parameter, $I/I_{s,eff}$, is evaluated below:

$$I/I_{s,eff} = \frac{1}{\sum_{m_g, q}} \sum_{m_e, m_g, q} I/I_{s, m_e, m_g, q}$$

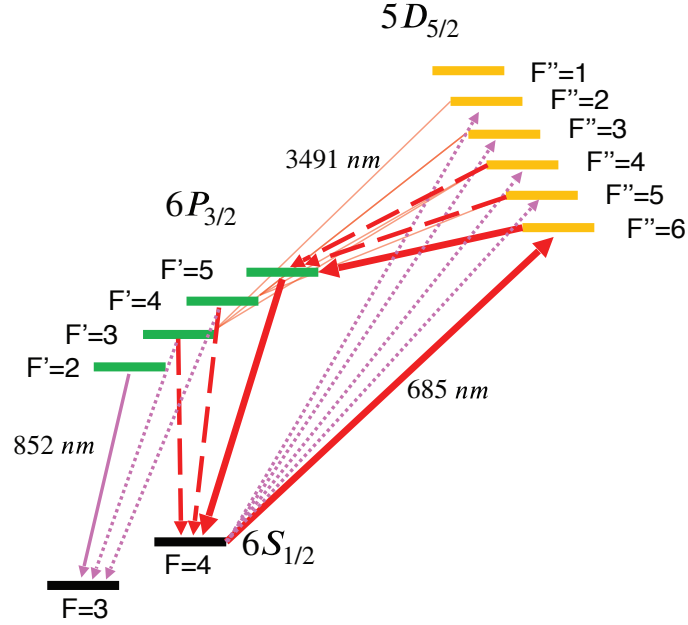


Figure 2.1: $6S_{1/2} \rightarrow 5D_{5/2} \rightarrow 6P_{3/2} \rightarrow 6S_{1/2}$ cycling and depumping diagram. Solid red lines indicate the primary cycling transition. Dashed red lines indicate paths that lead back to cycling after a Raman event. The dashed lavender lines indicate the Raman events that lead to the $F = 3$ state dark state. Finally, the thin solid orange lines show the decay paths that may lead back to the $F = 4$ cycling state if a future decay returns there (i.e. $5D_{5/2}F'' = 5 \rightarrow 6P_{3/2}F' = 4 \rightarrow 6S_{1/2}F = 4$).

which we can write in terms of the Rabi rate and detuning ($I/I_s = \frac{2|\Omega|^2}{\gamma^2}$):

$$I/I_{s,\text{eff}} = \frac{2}{(2F_g + 1)(2\kappa + 1)} \sum_{m_e, m_g, q} \frac{|\Omega_{F_g, m_g \rightarrow F''_e, m_e}|^2}{\gamma^2} \quad (2.20)$$

$$= \frac{I/I_{s,c}}{(2F_g + 1)(2\kappa + 1)} \sum_{m_e, m_g, q} \frac{|\Omega_{F_g, m_g \rightarrow F''_e, m_e}|^2}{|\Omega_c|^2} \quad (2.21)$$

where in the last step, $\gamma^2 = \frac{2|\Omega_c|^2}{I/I_{s,c}}$ has been substituted which will allow cancellation of the reduced matrix element and the radial contributions to

the Rabi frequencies. To this end,

$$\frac{|\Omega_{F_g, m_g \rightarrow F_e'', m_e}|^2}{|\Omega_c|^2} = \frac{|\langle F_g, m_{F,g} | \mathcal{H}_{E2} | F_e'', m_{F'',e} \rangle|^2}{|\langle F_g, m_{F,g} | \mathcal{H}_{E2} | F_g + \kappa, m_{F,g} + \kappa \rangle|^2} \quad (2.22)$$

$$= \frac{|\langle F_g, m_{F,g} | r^2 Y_{2,q} | F_e'', m_{F'',e} \rangle|^2}{|\langle F_g, m_{F,g} | r^2 Y_{2,2} | F_g + 2, m_{F,g} + 2 \rangle|^2} \quad (2.23)$$

where, as indicated in the previous section, we take $\kappa = 2$, since we are considering quadrupole interactions. We can obtain the second equality, where we ignore radial contributions (and any other coefficients), because the radial terms in \mathcal{H} do not interact with angular momentum. Now, we will expand these into reduced matrix elements, and Wigner n-j symbols, to simplify.

The total angular momentum, F , is an addition of the total electronic angular momentum, J , and the nuclear spin, I_n . In the case of the states that we are summing over, J and I_n are the same (i.e., the only non-cycling states we are considering are all in $5d_{5/2}$, and as such, we must only obtain an expansion of the angular momentum coupling in terms of J (which is what is coupled with the perturbation), and I_n (which is orthogonal to the perturbation)). Thus, we expand:

$$\langle I, J', F', m_F' | r^2 Y_{\kappa,q} | I, J, F, m_F \rangle = \frac{(-1)^{2\kappa}}{\sqrt{2F'+1}} C_{F, m_F, \kappa, q}^{F', m_F'} \langle I, J', F' || r^2 Y_{\kappa} || I, J, F \rangle \quad (2.24)$$

$$\begin{aligned} \langle I, J', F' || r^2 Y_{\kappa} || I, J, F \rangle &= (-1)^{F+J+I+J'} \sqrt{(2F+1)(2F'+1)} \\ &\times \left\{ \begin{matrix} J & I & F \\ F' & \kappa & J' \end{matrix} \right\} \langle I, J' || r^2 Y_{\kappa} || I, J \rangle \end{aligned} \quad (2.25)$$

which is now in terms of a common reduced matrix element, $\langle I, J' || r^2 Y_{\kappa} || I, J \rangle$.

Putting it together:

$$\begin{aligned} \langle I, J', F', m'_F | r^2 Y_{\kappa, q} | I, J, F, m_F \rangle &= (-1)^{F+J+I+J'+2\kappa} \sqrt{2F+1} C_{F, m_F, \kappa, q}^{F', m'_F} \\ &\times \left\{ \begin{matrix} J & I & F \\ F' & \kappa & J' \end{matrix} \right\} \langle I, J' || r^2 Y_{\kappa} || I, J \rangle \end{aligned} \quad (2.26)$$

To fully evaluate the matrix element, this process of decomposing the angular momentum states out of the reduced matrix element can be repeated until we can evaluate the $R_{n'L', nL}$ component, using radial integrals as in [36]. However, we can use this for our purposes of calculating the pump-depump ratio since it will cancel out Equation (2.23) becomes:

$$\frac{|\Omega_{F_g, m_g \rightarrow F'_e, m_e}|^2}{|\Omega_c|^2} = \frac{2F'_e + 1}{2F_g + 3} \frac{\left| \left\{ \begin{matrix} J_g & I & F_g \\ F'_e & \kappa & J_e \end{matrix} \right\} \right|^2}{\left| \left\{ \begin{matrix} J_g & I & F_g \\ F_g + 2 & 2 & J_e \end{matrix} \right\} \right|^2} \left| C_{F_g, m_g, \kappa, q}^{F'_e, m_e} \right|^2 \quad (2.27)$$

And so the effective saturation parameter ratio (s_{eff}/s , using Equation 2.20) becomes:

$$s_{eff, F'_e} / s = \frac{I/I_{s, eff}}{I/I_{s, c}} = \frac{(2F'_e + 1) \sum_{m_e, m_g, q} \left| C_{F_g, m_g, \kappa, q}^{F'_e, m_e} \right|^2}{(2F_g + 3)(2F_g + 1)(2\kappa + 1)} \frac{\left| \left\{ \begin{matrix} J_g & I & F_g \\ F'_e & \kappa & J_e \end{matrix} \right\} \right|^2}{\left| \left\{ \begin{matrix} J_g & I & F_g \\ F_g + 2 & 2 & J_e \end{matrix} \right\} \right|^2} \quad (2.28)$$

where the sum of the magnitude squared Clebsch-Gordan coefficients over

\mathbf{m}_g and \mathbf{q} is unity, such that the total sum is $2F_e'' + 1$, leaving:

$$s_{\text{eff}, F_e''} / s = \frac{(2F_e'' + 1)^2}{(2F_g + 3)(2F_g + 1)(2\kappa + 1)} \frac{\left| \begin{pmatrix} J_g & I & F_g \\ F_e'' & 2 & J_e \end{pmatrix} \right|^2}{\left| \begin{pmatrix} J_g & I & F_g \\ F_g + 2 & 2 & J_e \end{pmatrix} \right|^2} \quad (2.29)$$

With this, we can acquire a vector representing the rate that the $F_g = 4$ ground state will excite to each of the physically possible $F_e'' = 2, 3, 4, 5, 6$ states.

Now, to get the "Raman" rate and the Cycling rate, lets look at the probabilities that each of these F_e'' states decay either into either the $F_g = 4$ (cycling) or $F_g = 3$ (not cycling, or "Raman"). To this end, the branching coefficients of each of the states must be calculated. Furthermore, since there is an intermediate state ($J_e' = 3/2$), there will be two sets of branching coefficients to calculate.

The branching coefficient is simply the magnitude squared of the matrix element of a rank 1 spherical tensor operator, $T_{\kappa, q}$ (because this will be an E1 decay) between the two states in question, normalized such that the sum of branching coefficients from the same state is unity. I.e., for the decay $F_e, \mathbf{m}_e \rightarrow F_g, \mathbf{m}_g$:

$$b_{F_e, \mathbf{m}_e \rightarrow F_g, \mathbf{m}_g} = \frac{|\langle F_g, \mathbf{m}_g | T_{1, \mathbf{m}_e - \mathbf{m}_g} | F_e, \mathbf{m}_e \rangle|^2}{\sum_{F_i, \mathbf{m}_j} |\langle F_i, \mathbf{m}_j | T_{1, \mathbf{m}_e - \mathbf{m}_j} | F_e, \mathbf{m}_e \rangle|^2} \quad (2.30)$$

where the sum over i indexes the possible \mathbf{m} values for the given j state, which symbolizes all other possible states (obviously, only the possible states will have a non-zero matrix element, so this sum is over the possible states).

If we combine the degenerate \mathbf{m} states, this becomes:

$$\mathbf{b}_{F_e \rightarrow F_g} = \frac{\sum_{\mathbf{m}_g, \mathbf{q}, \mathbf{m}_e} |\langle F_g, \mathbf{m}_g | T_{1, \mathbf{q}} | F_e, \mathbf{m}_e \rangle|^2}{\sum_{F_i, \mathbf{m}_j, \mathbf{m}_e} |\langle F_i, \mathbf{m}_j | T_{1, \mathbf{m}_e - \mathbf{m}_j} | F_e, \mathbf{m}_e \rangle|^2} \quad (2.31)$$

Using (2.26) again, this becomes:

$$\mathbf{b}_{F_e \rightarrow F_g} = \frac{(2F_g + 1) \sum_{\mathbf{m}_g, \mathbf{q}, \mathbf{m}_e} \left| C_{F_g, \mathbf{m}_g, 1, \mathbf{q}}^{F_e, \mathbf{m}_e} \right|^2 \left| \begin{Bmatrix} J_e & I & F_e \\ F_g & 1 & J_g \end{Bmatrix} \right|^2}{\sum_{F_i} (2F_i + 1) \sum_{\mathbf{m}_j, \mathbf{q}, \mathbf{m}_e} \left| C_{F_i, \mathbf{m}_j, 1, \mathbf{q}}^{F_e, \mathbf{m}_e} \right|^2 \left| \begin{Bmatrix} J_e & I & F_e \\ F_i & 1 & J_g \end{Bmatrix} \right|^2} \quad (2.32)$$

And as before, the sum over CB coefficients will simply equal $2F_e + 1$ and cancel out (if the transition is allowed), and the reduced matrix elements cancel out since $I_g = I_e$ and all $J_i = J_g$ are equal.

$$\mathbf{b}_{F_e \rightarrow F_g} = \frac{(2F_g + 1) \left| \begin{Bmatrix} J_e & I & F_e \\ F_g & 1 & J_g \end{Bmatrix} \right|^2}{\sum_{F_i} (2F_i + 1) \left| \begin{Bmatrix} J_e & I & F_e \\ F_i & 1 & J_g \end{Bmatrix} \right|^2} \quad (2.33)$$

Now, we can define two matrices for these branching coefficients, the first from the $5D_{5/2} \rightarrow 7P_{3/2}$ and the second for the $7P_{3/2} \rightarrow 6S_{1/2}$. The first matrix, will be a 4 by 6 matrix (there are four F'_e states in $7P_{3/2}$ and six F''_e in $5D_{5/2}$), and the second will be 2 by 4 (there are two F_g states in $6S_{1/2}$). These matrices can then be multiplied together to get a total branching coefficient matrix, for coupling each state in $5D_{5/2}$ down to $6S_{1/2}$

through the two photon process with the intermediate $7P_{3/2}$ state.

$$\mathbf{b}_{5D_{5/2} \rightarrow 7P_{3/2}} = \begin{pmatrix} 1 & \frac{3}{5} & \frac{3}{14} & 0 & 0 & 0 \\ 0 & \frac{2}{5} & \frac{5}{8} & \frac{11}{24} & 0 & 0 \\ 0 & 0 & \frac{9}{56} & \frac{99}{200} & \frac{18}{25} & 0 \\ 0 & 0 & 0 & \frac{7}{150} & \frac{7}{25} & 1 \end{pmatrix} \quad (2.34)$$

$$\mathbf{b}_{7P_{3/2} \rightarrow 6S_{1/2}} = \begin{pmatrix} 1 & \frac{3}{4} & \frac{5}{12} & 0 \\ 0 & \frac{1}{4} & \frac{7}{12} & 1 \end{pmatrix} \quad (2.35)$$

$$\mathbf{b}_{5D_{5/2} \rightarrow 6S_{1/2}} = \begin{pmatrix} 1 & \frac{9}{10} & \frac{3}{4} & \frac{11}{20} & \frac{3}{10} & 0 \\ 0 & \frac{1}{10} & \frac{1}{4} & \frac{9}{20} & \frac{7}{10} & 1 \end{pmatrix} \quad (2.36)$$

And so our Cycling and Raman rates become:

$$\begin{pmatrix} r_C \\ r_R \end{pmatrix} = \begin{pmatrix} 1 & \frac{9}{10} & \frac{3}{4} & \frac{11}{20} & \frac{3}{10} & 0 \\ 0 & \frac{1}{10} & \frac{1}{4} & \frac{9}{20} & \frac{7}{10} & 1 \end{pmatrix} \begin{pmatrix} r_{F''=1} \\ r_{F''=2} \\ r_{F''=3} \\ r_{F''=4} \\ r_{F''=5} \\ r_{F''=6} \end{pmatrix} \quad (2.37)$$

where the values for $r_{F''}$ come from (2.19) and the values for the saturation intensities come from (2.29) as a function of the saturation intensity.

These can be evaluated as a function of detuning and/or saturation intensity, as well as a ratio (r_C/r_R), Figure 2.2. Additionally, the cycling rate can be plotted as a function of the detuning and/or saturation intensity, Figure 2.3.

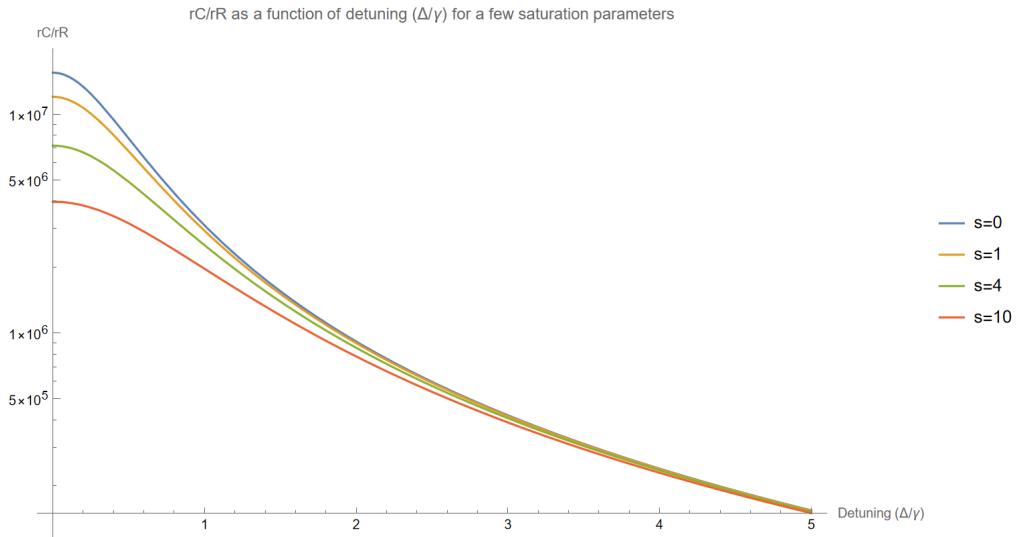


Figure 2.2: Logarithmic plot of the cycling rate over the Raman rate as a function of the detuning, for a few saturation parameters

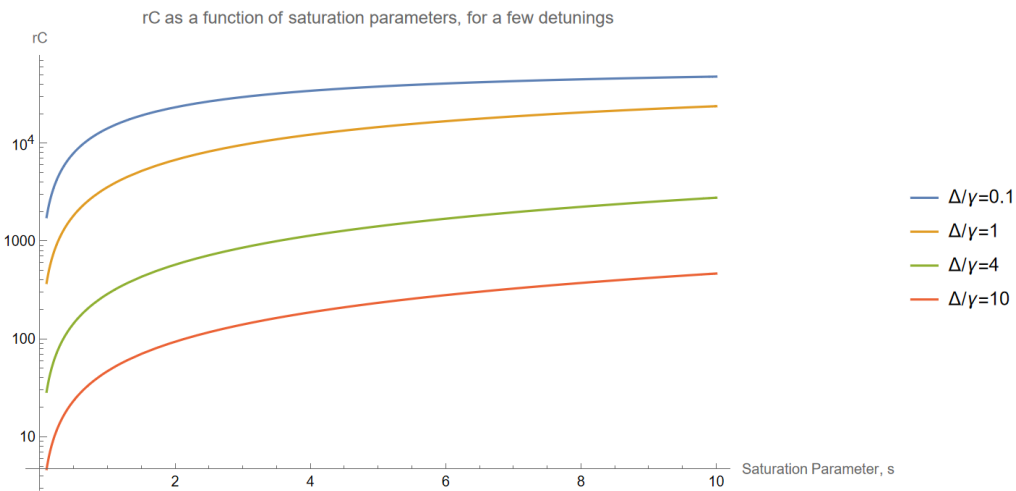


Figure 2.3: Logarithmic plot of the cycling rate over the Raman rate as a function of the detuning, for a few saturation parameters

Part II

Experiment

3 EXPERIMENTAL APPARATUS AND DESIGN FOR CREATING QUBITS AND PERFORMING QUADRUPOLE CONTROL

3.1 The Experimental Apparatus

The experimental setup involves a dual-glass cell vacuum chamber, separated by a pinhole allowing for a differential pressure, where a Cs dispenser is placed in the bottom cell and atoms have to travel through a small pinhole before meeting the ion pump and the upper glass cell where the main experiment is performed.

This section describes the 2D and 3D Magneto-Optical-Trap (MOT) systems in detail, and presents the general design and apparatus of the experiment.

3.1.1 The 2D MOT

The lower cell has a 2D MOT setup with large waist beams and permanent magnets for creating the quadrupole field. These 2D MOTs are created using a pair of retro-reflected 852 nm lasers [37, 38]. To get atoms to the upper chamber, a push-beam (using a sample of the same light for the 2D MOT beams) blows atoms out of the 2D MOT, through the pinhole, and into the upper cell. This creates a column of atoms into the upper chamber, which is used to load a 3D MOT.

The beams for the 2D MOT are composed of -2γ detuned 852 nm light for pumping the $6S_{1/2}, F = 4 \rightarrow 6P_{3/2}, F' = 5$. Since this transition has a much higher Raman depump rate than the quadrupole transition (as discussed in 2.4), a repump laser is necessary. This repump laser is locked to the $6S_{1/2}, F = 3 \rightarrow 6P_{3/2}, F' = 4$, such that atoms depumped into the ground $F = 3$ dark state will be pumped back into the $F = 4$ dark state




Figure 3.1: 3D MOT Layout for the upper cell. 685 nm and 852 nm beams are combined to make a MOT with diagonal-top/bottom left/right-beams, as well as a pair of beams for the front and back.

through the $6P_{3/2}, F' = 4 \rightarrow 6S_{1/2}, F = 4$ decay channel. These lasers are combined on a PBS which splits the power between the push beam and the 2D MOT beams, which are then fiber coupled (with a polarization cleanup stage) to the experiment apparatus.

3.1.2 The 3D MOT

The 3D MOT is setup with 6 independent beams into the square cell. These include 2 pairs of diagonal beams (top/bottom-left/right) as well as a front/back pair. Figure 3.1 presents this layout and figure 3.2 is an actual image of the system. Since the diagonal beams need to make it around the objective, there is a 50 degree angle-of-incidence (AOI) from the normal of the glass for these beams. The beams have a waist of 3.2 mm ($1/e^2$ radius) for the diagonal beams and 3.7 mm for the front/back. The beams are aligned to create a MOT at the center of the glass cell, at an intersection with the push beam from the 2D MOT glass cell below.

To create the proper polarization of light required for the MOT ($\sigma + / -$)




Figure 3.2: An image of the upper cell chamber.

for both wavelengths, achromatic quarter-waveplates are used to rotate the 852 nm and 685 nm light equally. We use protected silver-coated mirrors instead of dielectric mirrors for their improved achromatic polarization performance. These quarter-waveplates are placed as close to the cell wall as possible. However, the cell wall itself is not well coated for 685 nm and is coated for 852 nm. This causes an unknown chromatically dependent retardation that can't be accounted for. The consequences of the effects of the coating issues are discussed further in 3.2.1.

In order to get the bottom diagonal beams into the cell, mirrors have to be placed (see Figure 3.1). Unfortunately, waveplates cannot be fit into this space and they have to be placed before the mirror. This causes some further issues, preventing ideal circular polarization in the glass cell at the atoms for the cooling. In the end, the best that we could do was optimize the waveplates for MOT brightness.

To observe the 3D MOT, an Andor Luca EMCCD camera with a zoom lens images the atom cloud. This camera has an 852 nm line filter to remove background light (like the 685 nm light). This camera is able to also take very short exposure time images (1 ms) allowing for time-of-flight atom temperature measurements (implemented in section 4.1.3).

A layout for the optics before the cell for the bottom beams can be found in figure 3.3. Since the beams have to be very close to the main optical table, this is done on a separate breadboard hanging of the main optical table. The top beams are combined similarly, but with a vertical breadboard that is aiming down into the cell. However, in this case, to improve the independent alignment capabilities of the 685 nm vs 852 nm beams, we add a pair of PBS's to first split the 685 nm light, then recombine before adding with the 852 nm on the dichroic mirror before another PBS.

The front/back 685 nm + 852 nm beam pair is not combined on a dichroic mirror, and is instead a skewed beam. The angle of incidence of the 852 nm light is close to normal, but the 685 nm light has an angle




Figure 3.3: Bottom 685 nm+852 nm MOT light combining optics. Lasers are combined on a dichroic and sent to the bottom mirrors in figure 3.1.

of incidence about 5 degrees from the normal on the x-y plane (see the coordinate system in figure 3.1).

Each counter-propagating beam has a relative frequency of 1 kHz imposed with the switching AOMs. This highly reduces standing wave effects in the beams which improves the MOT stability.

There are a total of 4 fiber launches for the MOT beams for both colors of light (so 8 total). To improve the intensity stability of the beams each of these fiber launches has a photodiode with a trans-impedance amplifier to measure the power of the light for calibration purposes. This voltage is measured on a National-Instruments analog input card, which is measured within the experiment cycle. Since there can be erroneous light leakage from the counter-propagating beams, the analog input measurement cycle for this calibration stage involves a round-robin approach where each beam is independently turned on for a brief moment. This occurs every experiment cycle (i.e. before the MOT is loaded). The feedback of the beam power measurement error is sent to a set of Thorlabs K10CR1 rotators with L/2 waveplates. These allow for feedback to ensure the power is always at a specific setpoint, even as the laser systems' alignment drift.

Coil drivers (see section 6.2) are used to drive x-y-z shim coils and an anti-Helmholtz coil, where x is the front-back axis, y is the axis of the objective, and z is the vertical axis. As such, the diagonal MOT beams are on the y-z plane and thus not on a directly controllable axis with the coils. This makes MOT alignments a little more difficult, but still manageable. The quadrupole magnetic field axis (i.e. anti-Helmholtz axis) is along the x-axis and therefore does have a dedicated beam for it.

3.1.3 Single Atom Trap Design with Hole Array

Atom trapping is achieved by imaging a mask with holes on it, similar to the work done in [39]. In our situation, the mask only has an array of 3x3 holes and the separation between the holes is 1.5 mm. The holes are completely opaque with a diameter of 200 μm .

The laser is an 803 nm laser which is a "blue" Far Off-Resonant Trap (FORT) laser for the $6S_{1/2}$ ground state of Cesium, but a "red" trap for the $6P_{3/2}$ excited state. However, this wavelength is "magic" for atoms in the $5D_{5/2}$ excited state, meaning that the differential AC stark shift can between $6S_{1/2}, F = 4$ and $5D_{5/2}, F = 6, |m_F| = 6$ cancel each other out, which is discussed in much further detail in section 4.3.2 as well as in [40, 41].

As described in [39], the light from the trap has to be imaged into cell from the mask. This involves setting up an series of "4f" imaged lenses. Figure 3.4 shows the technical design of this setup. Figure 3.5 shows an analysis of the trap array by measuring the trap shape on a CCD camera with 3.45 μm pixels. The camera was placed at the final image plane before the atom image plane, and was used to optimize the position and closure of the iris for the trap shape.

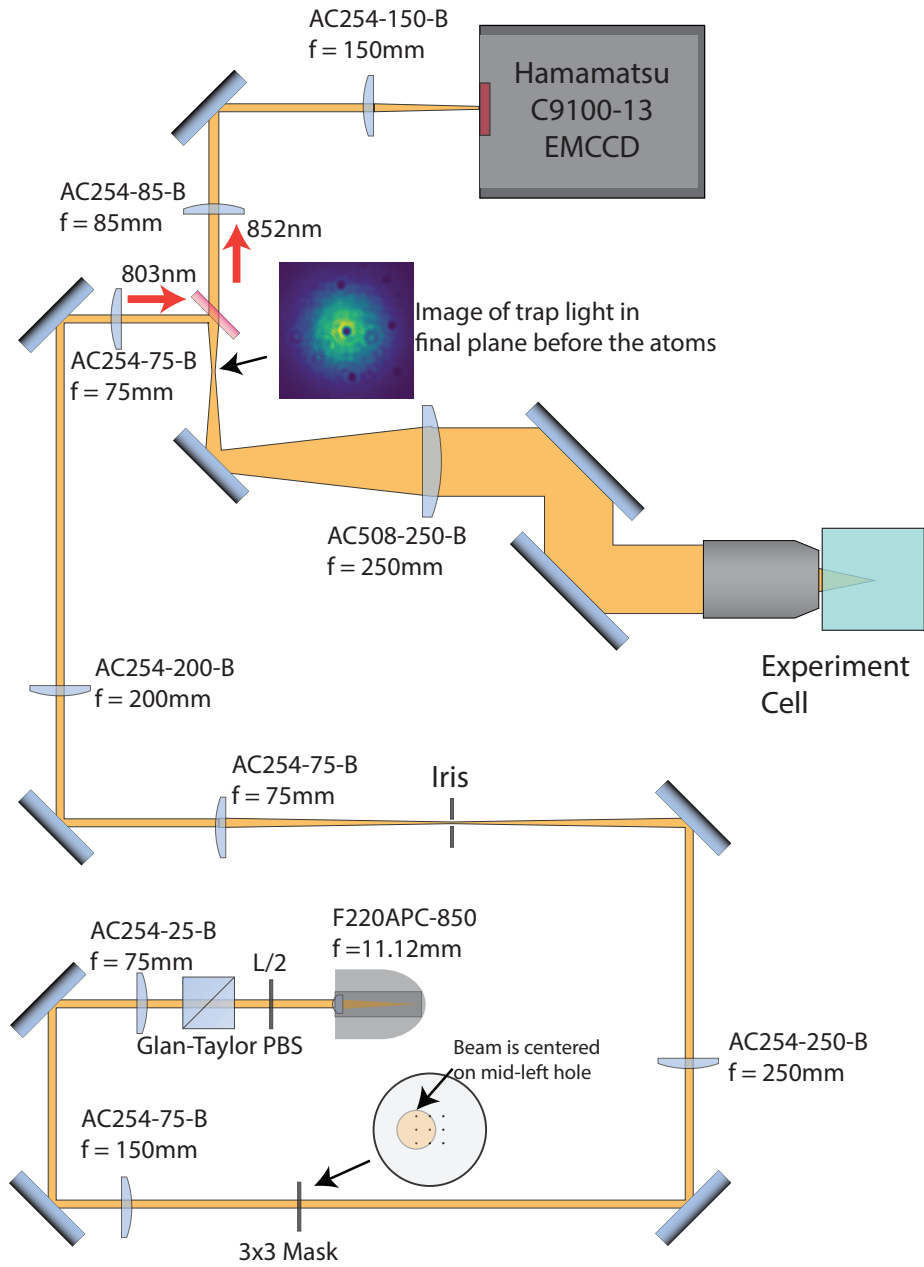


Figure 3.4: 4-f Imaging setup for trap array. Distances not to scale. Each lens is separated from the previous lens (along the beam path) by the sum of their focal lengths for 4-f imaging.

Analyzing the trap shape in final image plane

The trap shape fitting in figure 3.5 involved fitting a gaussian potential ($I_0 e^{-2(x/\sigma_x)^2 - 2(y/\sigma_y)^2}$) to the global beam profile followed by fitting the hole in the trap (the subtraction of the fit gaussian intensity from the actual intensity) to a parabolic function ($\alpha_x x^2 + \alpha_y y^2 - I_{\min}$). Since there is an objective ($f = 22.8\text{mm}$) and achromat ($f = 250\text{mm}$) between the final image plane and the atom image plane, a scaling factor of $3.45 \text{ }\mu\text{m}/\text{pixel} * 22.8/250 = 0.317 \text{ }\mu\text{m}/\text{pixel}$ is used to scale the horizontal axes of the subplots to be in the distance scale at the atoms (this assumes that we are not diffraction limited and the imaging is done perfectly).

The total power of the beam is measured to be 125 mW. Using this, we can extrapolate the peak intensity of the gaussian beam ($I_0 = \frac{2P_0}{\pi\sigma^2}$, where σ is the $1/e^2$ intensity waist). From the fits, we find $\sigma_x = 19.6\text{ }\mu\text{m}$ and $\sigma_y = 20.1\text{ }\mu\text{m}$. With $\sigma = 19.8\text{ }\mu\text{m}$ (the average), the peak intensity of the "gaussian beam" is $I_0 = 0.2\text{mW}/\text{ }\mu\text{m}^2$.

The AC Stark shift on the $6S_{1/2}$ ground state from the trap potential can be calculated from the scalar polarizability of the ground state at 803 nm:

$$\Delta U_{\text{ac}} = -\frac{1}{4}\alpha_{\text{SI}}|\mathcal{E}|^2 \quad (3.1)$$

$$= -\frac{1}{4}\alpha_{\text{SI}} \left(\frac{2I}{\epsilon_0 c} \right) \quad (3.2)$$

where α_{SI} is the scalar polarizability in SI (not cgs), \mathcal{E} is the electric field, I is the beam intensity, ϵ_0 is the vacuum permittivity constant, and c is the speed of light. In with α in cgs units; $\alpha_{\text{SI}} = 4\pi\epsilon_0\alpha_{\text{cgs}}$, equation 3.1 becomes:

$$\Delta U_{\text{ac}} = \frac{-2\pi\alpha_{\text{cgs}}I}{c} \quad (3.3)$$

The scalar polarizability of the $6S_{1/2}$ ground state of Cs is $\alpha_{\text{cgs}} = -365 \times 10^{-24} \text{ cm}^3$ in the cgs unit system. With this, we can compute the

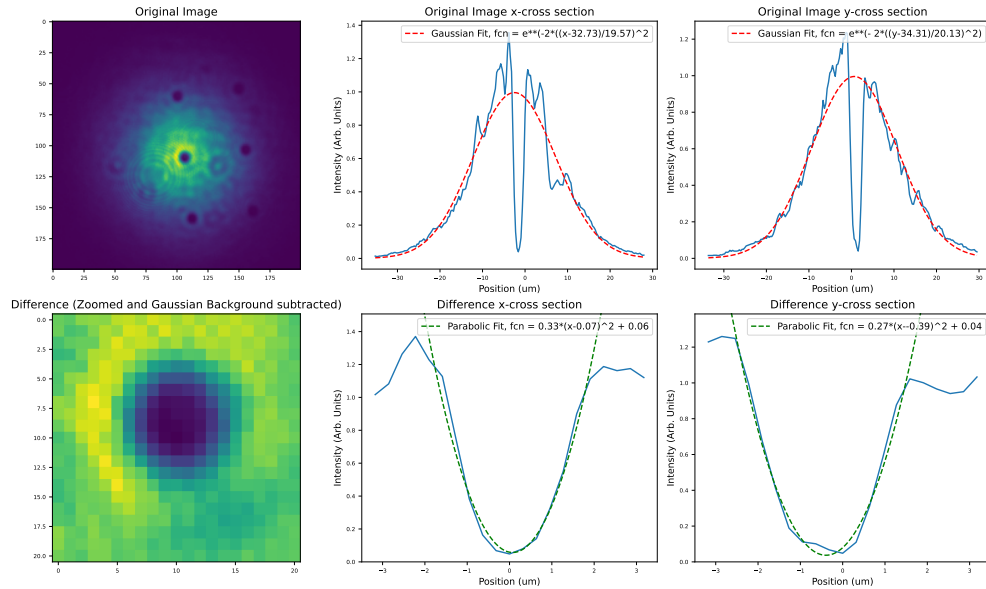


Figure 3.5: Analysis of trap shape and size in final image plane before the atom image plane. The cross-section plot's horizontal axes are scaled with a factor of $3.45 * 22.8/250$ um/pixel to demonstrate the beam size in the atom plane. The vertical axes are scaled to normalize the intensity of the gaussian fit. Top-Left: The raw data measured by the Thorlabs CCD camera with 3.45 um per pixel. Top-Middle(Right): x/(y) cross-sectional fits of the normalized intensity (peak intensity is equal to 1). Bottom-Left: A zoom in of the Top-Left image, centered around the hole in the trap. Bottom-Middle(Right): x/(y) cross-sectional fits of the gaussian subtracted trap to a parabolic potential. Functional forms of fits are presented in the legends of the cross-section plots.

AC stark shift at the peak intensity ($I_0 = 0.2\text{mW}/\mu\text{m}^2$) of our 803 nm trap laser on the ground state to be $\Delta U_{\text{ac}} = 1.56 * 10^{-27}\text{J}$. We can also scale this by the Boltzmann constant, which relates the thermal energy of a particle to the thermal temperature of that particle. With this scaling factor, we can have an idea for what the maximum temperature is for an atom trapped in the potential. Performing this scaling gives us a trap depth of $\Delta U_{\text{ac}} = 113\mu\text{K}$.

Since this is the scaling factor at the peak of the gaussian distribution in figure 3.5, we can also apply this scaling factor to all of the plots in the figure to determine the potential energy in terms of μK at all points. 3.6 presents this data.

We find that the minimum trap potential is at around $7\mu\text{K}$, assuming that there are no diffraction-limited effects causing an increased offset on the bottom of the potential. This also indicates that atoms over $106\mu\text{K}$ will escape the trap.

Calculating the trap frequency from fit

If we model the trap potential as a harmonic oscillator, we can extrapolate the trap frequency. A harmonic oscillator potential,

$$U_{\text{sho}} = \frac{1}{2}m\omega^2x^2 \quad (3.4)$$

where m is the mass of the Cs atom, ω is the angular trap frequency, and x is the spatial displacement. The cross-sectional parabolic fits shown in figure 3.6 are of the form αx^2 with a spatial offset for x and a vertical constant offset. Since the offsets are unimportant for the trap frequency, we can extrapolate $\alpha = m\omega^2/2$, or:

$$\omega = \sqrt{\frac{2\alpha}{m}} \quad (3.5)$$

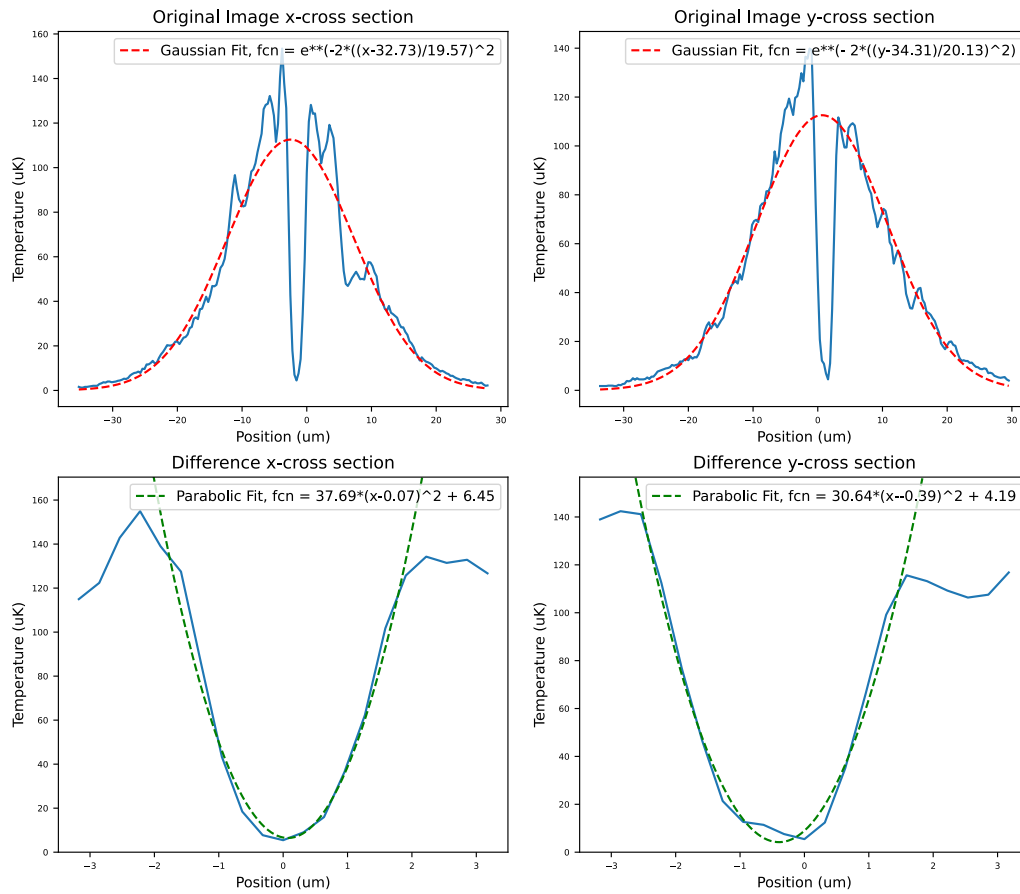


Figure 3.6: Trap shape analysis as seen in figure 3.5 but with temperature scaling in μK for the vertical axis, using the total beam power of 125 mW, gaussian beam waist of 20 μm , and the ground state scalar polarizability of $-365 \times 10^{-24} \text{ cm}^3$ (cgs units) at 803 nm to calibrate.

From the fits, $\alpha_x = 38 \mu\text{K}/\mu\text{m}^2$ and $\alpha_y = 31 \mu\text{K}/\mu\text{m}^2$. The mass of a Cs atom is $2.2 \times 10^{-25} \text{ kg}$. With this, we can solve for the angular trap frequency, $\omega_x = 69 \text{krad/s}$, $\omega_y = 62 \text{krad/s}$. Or in terms of frequency, $f_x = 10.9 \text{ kHz}$, $f_y = 9.9 \text{ kHz}$. These values are experimentally verified in section 4.2.3.

3.1.4 Imaging single atoms onto the EMCCD

As figure 3.1 indicates, the 852 nm atom fluorescence and 803 nm are imaged using the same achromatic high NA objective. This allows for simple alignment in order to be able to see the single atoms in the cell. As seen in figure 3.4, there is a dichroic which splits the 852 nm imaged light and 803 nm trap light. The imaging light is magnified by a 75 mm/85 mm telescope and imaged onto the EMCCD camera. Due to imperfect achromaticity of the lenses, it is expected for there to be a chromatic spatial shift of the atom fluorescence image plane and the trap image plane. This can be compensated for easily, since the placement of the 85 mm lens can be adjusted to image the right plane.

Unfortunately, the dichroic filter is not placed in a good location of the optical train. It is very close to an image plane with a very small waist, which means that the angle of divergence is relatively high. If the light from the atom has a $1 \mu\text{m}$ waist (rough estimate based on trap size), then the size of the atom in that image plane is about $11 \mu\text{m}$. This gives an angle of divergence of $\theta = \frac{\lambda}{\pi w_0} = 0.025 \text{ rad} = 1.5 \text{ deg}$. When this diffracting beam travels through the dichroic, an astigmatic-like aberration is created because of the thickness of the dichroic (which is at a 45 deg angle from the beam k-vector). The axis of the beam that does not see the 45 degree angle of the dichroic has a different phase accumulation than the axis of the light that does see the 45 degree angle of incidence. There is also a spherical aberration caused by the fact that the beam has the 1.5 deg divergence angle through the flat window. Since the beam component that is further

from the beam center will have to travel through the dichroic for a longer distance, it will have an increased phase accumulation than a beam traveling at the center.

An improvement to the imaging train would be to place the dichroic in a collimated plane, near the plane directly after the objective. In this (Fourier) plane, the focal waist is massive, so the beam is "collimated". The astigmatism caused by the angle of the dichroic would become much less problematic because the beam is very large. And since the beam has a very small divergence angle, there would be very little spherical aberration. Unfortunately, we did not have a dichroic large enough for this, and if we had used the dichroic we do have, it would have clipped the beam.

3.2 Developing the 685 nm Laser System

The first step in utilizing the quadrupole transition in Cesium is getting a laser. This section discusses the laser infrastructure design and noise characterization of the 685 nm laser. Figure 3.7 shows an image of the laser source for the system, described in section 3.2.2.

3.2.1 685 nm Requirements and Design

The glass cell used for the experiment was unfortunately not coated for 685 nm, and the cell wall transmission is very bad. Table 3.1 presents the cell wall transmission data, based on measurements performed by shining 685 nm light through the cell and measuring the amount of light coming out of the other side. Taking the square root of the total loss ratio is the loss of a single window, which is the loss to reach the atoms. As indicated, the front and back axes have very low transmission of around 50%, while the diagonal axes have a transmission of around 90%. This is the core reason for why a retro-reflection MOT system isn't feasible, leading to our 6-beam approach. The 90% transmission becomes 70% on a the third pass through




Figure 3.7: 685 nm Laser System from a 15W Sprout pump at 532nm, an MSquared SolsTiS lasing at 685 nm, and a ULE cavity for locking the laser.

the window for the diagonals, and the 50% transmission becomes 13% for the front/back. Additionally, since the transmission of the front and back is so low, a 5 degree angle of incidence was added to the beam to reduce a cavity effect caused by the normal incidence, resulting in spacial interference patterns.

The beam waists are also described in the diagram, the resulting estimated peak intensities, and the power of the laser before entering the cell. As described in section 3.1.2, the 685 nm beams for the top right/left come from the same fiber, and the bottom right/left also come from the same fiber. The front and back each come from independent fibers. These fibers

Table 3.1: 685 nm cell window transmission data measured by square root of total transmission through whole cell (meaning that the loss per cell wall was considered constant.) Diagonal beam's angle of incidence is 50 deg and front/back is 5 degrees.

Beam	Waist (mm)	Transmission (%)	Peak Intensity (W/cm ²)	Power before cell (mW)
Front	1.1	51	0.10	7.5
Back	1.1	56	0.10	6.8
Top-Left	0.58	92	0.23	2.6
Bottom-Right	0.58	85	0.47	5.8
Top-Right	0.58	87	0.58	7.0
Bottom-Left	0.58	91	0.31	3.6
Total			1.80	

are switched with upstream AOMs (before the fiber) in the 685 nm Injection Lock Amplification setup (3.2.3).

The powers of the beam exiting the fibers are around 15 mW for the top and bottom launches, and about 7 mW for the front and back. These powers are stabilized using a photodiode immediately after the fiber cleanup polarizer to measure the power of the light, as described in section 3.1.2. The power is then stabilized using a K10CR1 rotator with an L/2 waveplate, which is controlled by a National Instruments analog input card. Up to 20 mW are coupled out of the top/bottom fibers, and 10 mW for the front/back, but this power is brought down to allow for the power stabilizer to have space to operate.

The design of the 685 nm system first involves a M2 laser system (3.2.2) which has an SolsTiS laser locked to a ULE cavity. This light is coupled to the injection lock setup (3.2.3) which amplifies the power through injection locking. This setup is also responsible for housing the switching AOMs and the fiber launch and rotating stage systems before being sent to the experimental apparatus presented in section 3.1.2.

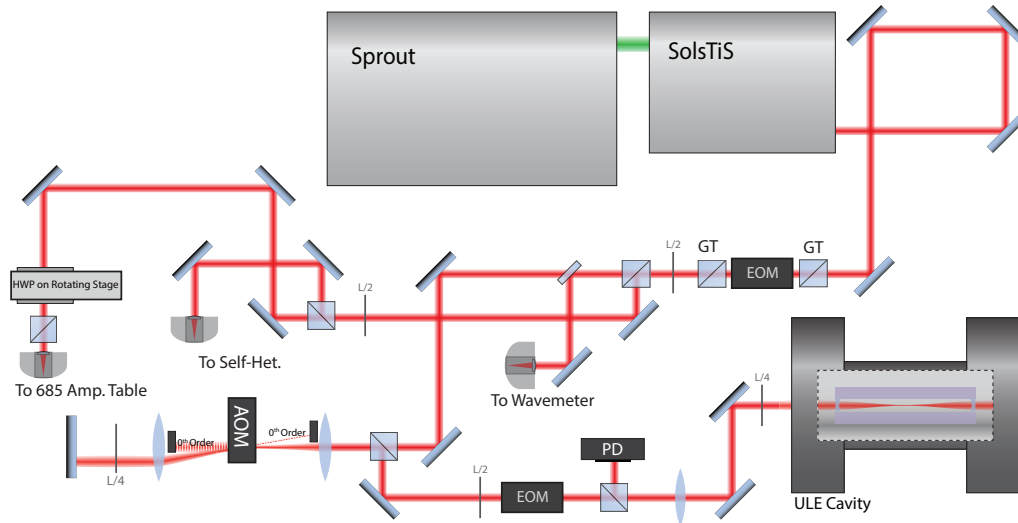


Figure 3.8: Optical layout of the 685 nm M2 laser system.

3.2.2 685 nm M2 Laser System

The 685 nm laser system is based on a 15 W Sprout laser at 532 nm, which pumps an MSquared SolsTiS cavity lasing at 685 nm. The amount of power from the SolsTiS started out at around 600 mW, but has been constantly degrading over the past many years. As of the time of writing, and a lot of the data taken, the power is now down to around 300 mW from 15 W pump, which is just enough to keep the SolsTiS etalon locked.

This is an older SolsTiS, before they added EOMs to the cavity, which means that the only feedback paths are the two piezos (fast and slow). Section 5.1.2 discusses the locking scheme used for the 459 nm laser which also has the more modern, fast EOM feedback path, allowing for a simpler locking scheme. Section 3.2.4 discusses the electronics and feedback design for the 685 nm laser locking scheme.

Figure 3.8 shows the optical layout of the 685 nm laser system. After exiting the SolsTiS, the light is sent through some beam shaping optics. This then enters a Glan-Taylor polarizer and into an out-of-cavity EOM. This

EOM is aligned to apply a phase shift to the light, and if properly setup, can be used to maintain the narrow the spectrum of the laser (see section 3.2.4 for a mathematical analysis of the capabilities of an out of cavity EOM). This light is then cleaned up in polarization again with another Glan-Taylor and sent to a few destinations with beam splitters/samplers. These destinations include the ULE cavity setup, a low power fiber launch to send about 10 uW of power to the wavemeter, an adjustable fiber launch to the self-heterodyne setup (more or less power can be sent here to increase the amount of heterodyne signal), and a fiber launch to the 685 nm injection lock amplification setup (3.2.3).

The light heading to the ULE cavity includes a double pass AOM for allowing a frequency offset of the laser from the cavity mode frequency. The AOM is driven at a 115.1 MHz frequency with a negative first order diffraction mode. This means that the laser frequency out of the SolsTiS is $115.1 \text{ MHz} * 2 = 230.2 \text{ MHz}$ higher than the cavity mode. After the double pass, the light passes through the PDH optics (involving an EOM modulated at 12 MHz) and is coupled to the cavity. The cavity mode used is a TEM_{01} mode and has a PDH feedback based locking scheme described in section 3.2.4. It is important to ensure that very little power is coupled into the cavity, since any thermal loss will cause the cavity to heat up. Since the cavity is in a vacuum cell with large amounts of thermal isolation from the environment, the cavity does not have a way to dissipate heat either.

There is about 250 mW of power coupled into the fiber to the injection lock amplification setup. This includes a polarization cleanup stage to align the polarization of the light to the fiber. However, only 120 mW is coupled out of the fiber, with each PBS transmission causing a 10% loss (a 19% total loss), leaving a fiber coupling efficiency of about 70%. There is a K10CR1 rotator with an L/2 waveplate to stabilize the power of the light before the fiber, but this rotating stage was ultimately never used because it was found to cause oscillations when run in conjunction with the downstream rotators.

3.2.3 685 Injection Lock Amplification Setup

This section describes the amplification setup for the 685 nm laser system, which utilizes injection-locked diodes to increase the amount of power. There is also a set of switching AOMs to allow for all four channels (front, back, top-left/right, bottom-left/right) to be independently switched and frequency shifted. Figure 3.9 shows an image of the setup in action.

As mentioned in section 3.2.2, there are 120 mW of power coupled out of the fiber (and PBS). There are two injection-locked diodes in the setup, one for the top left/right and one for the bottom left/right beams. The front and back beams come directly from the fiber out of the 120 mW of power. Only 5mW are needed for the seed power to the injection-locked diodes [42–44], but this has to travel through lossy optics and about 7 mW are sent to each of the injection lock diodes. Since the light coupled to the amplification table is a very clean gaussian beam (from a single-mode fiber), we can recouple the light back into a fiber after passing through the AOM switch and K10CR1 rotator stages and polarization cleanup optics. There is a maximum of about 20 mW of power coupled into each of the front and back fibers, since the seed laser has to be sent to many locations, and travel through an optical isolator ($T_x \approx 80\%$). So the budget is $120\text{mW} \times 0.8 = 96\text{mW}$ (optical isolator loss), from which 14 mW are sent to the injection-locked diodes, leaving about 80 mW for the front and back fibers, 40 mW for each. With an AOM diffraction efficiency of up to 70%, and other losses from PBS's and lenses, we get 20 mW of power for each of the front and back beams.

Injection-Locked Diode Enclosure and Temperature Control

The injection-locked diodes are HL69001DG, made by Ushio. Figure 3.10 shows the lasing wavelength vs. temperature for this diode, taken from the datasheet. The center wavelength at room temperature is 690 nm, but allows for a tuning range down to 685 nm with a case temperature of 0 °C. We decided to give these diodes a try and attempted to build a




Figure 3.9: A picture of the 685 nm amplification stage in action.




Figure 3.10: Lasing wavelength vs temperature for the HL69001DG diode taken from the datasheet. At a case temperature of 0 Celsius, the lasing wavelength is 685 nm and the center wavelength at room temp is 690nm.

humidity-controlled enclosure for the diode so that the 0 °C temperature would not cause both condensation and subsequent freezing of water onto the optics. This enclosure was built by Quinn Meece and the control system for the temperature and humidity was programmed by Isaac Scott.

The enclosure involved building a small box with active dehumidifier elements from Rosahl that use a "micro-membrane" to perform electrolysis of water molecules passing through. We never really got this system to work well. While it did reduce the humidity, it never quite met the specs. But after over-engineering the enclosure, we managed to safely bring the temperature of the diode mount inside the enclosure to 0 °C while keeping the dew point at a lower temperature, without condensation forming. This also involved lowering the temperature of the ambient air in the enclosure, effectively creating a two-stage cooling scheme.

With the dehumidifier, the humidity inside the enclosure was reduced to around 30% relative humidity. This means that if the air temperature inside




Figure 3.11: Dew Point (C) vs air temperature based on the Magnus formula. Reproduced from Easchiff, Wikimedia Commons, CC BY-SA 4.0; [46].

the box was 15 °C, the dew point would be negative. Figure 3.11 shows the dew point vs. air temperature based on the Magnus formula [45]. Note that this figure is reproduced from Easchiff on Wikimedia Commons, licensed under CC BY-SA 4.0. From the figure, we can see that at an air temperature of 15 °C and relative humidity of 30%, the dew point is around -5 °C, which gives enough room to ensure that the diode is safe. Temperature interlocks were set up to ensure that if the humidity or air temperature were too high for the diode temperature, the laser temperature would be ramped up to a safe level.

¹ However, after building and preparing this entire enclosure, we decided to just try injection locking the diode at 25 C. And to our surprise and disappointment for all the effort to prepare this enclosure, the diode injection worked perfectly at 25 C. So the entire humidity and temperature control system was never used.

¹Image reproduced from Easchiff, *Dewpoint-RH.svg*, Wikimedia Commons, <https://commons.wikimedia.org/wiki/File:Dewpoint-RH.svg>, licensed under CC BY-SA 4.0.

Laser Diode Injection-Lock setup

The injection-locked diodes run at around 160 mA, giving about 180 mW of power out of the diode. 200 mW is the absolute maximum output power, so running at this level improves diode longevity. A 5 mW seed is used to injection lock the diode, which means we get an amplification factor of 36, or about 1.5 orders of magnitude. The injection-locked stability is very good, with an injection current lock range of about 3 mA. Some references that discuss injection locking and other (more complicated) applications include [47, 48].

The optical train for our injection-locking scheme first involves a PBS beam splitter to take the 7 mW sample of the seed laser. This sample is then sent through a beam sampler (for auto-relock, see Section 3.2.3), and then through a PBS, a 45deg Faraday rotator, and another PBS (rotated at 45deg). This is, indeed, the description of a disassembled Faraday isolator.

The seed light passing through the Faraday isolator is output from the 45deg PBS, further rotated by a $\lambda/2$ waveplate, then coupled into the injection-locked diode, with some beam-shaping optics as well. However, the light leaving the injection-locked diode enters the Faraday rotator and gets rejected by the PBS, allowing for a separation of the seed light and the injection-locked light. This is the key to an injection-locked diode, since the seed light cannot itself be injection-locked by the injection-locked diode. (Also, it would be nice for power to be output somehow.)

There is an additional stage of optical isolation after the seed laser fiber launch, as mentioned before, to further attenuate the injection-locked laser source from coupling into the seed laser.

After the Faraday and PBS elements, the injection-locked light passes through a beam sampler (also for the auto-relock scheme, 3.2.3), then to the switching AOM for that beam path (top or bottom), and coupled to the fiber with a rotator and polarization cleanup as has been described above.

685 nm Switching AOMs

The switching AOMs are used to switch the 685 nm light for the front, back, top-left/right, and bottom-left/right beams. Each of these AOMs are driven at 147.5 MHz and are also responsible for applying small frequency shifts to the beam for spectroscopy and dynamic frequency shifts. However, they are setup in a single-pass configuration, meaning that there is a spatial coupling to the frequency shift. To mitigate this effect, we place the AOM at the focus of a $f = 100$ mm telescope with a 1:1 magnification. As such, when the AOM applies a k -vector shift to the diffraction order in the "Fourier" plane, the beam is simply displaced slightly in the image plane. This is better than the alternative where there is only ever a collimated beam that is deflected at an angle. In this situation, the beam continues to be displaced as it propagates to the fiber launch, building up a larger and larger displacement. A double-pass AOM would be better for the frequency shift and shuttering, since this would prevent the spatial coupling of the frequency shift, but there would be a much larger power loss [49].

We align the AOMs at the center frequency of 147.5 MHz, very near the atomic resonance of the $6S_{1/2}, F = 4 \rightarrow 5D_{5/2}, F'' = 6$ transition, based on measurements of the ULE cavity's frequency done by previous members of the group [50, 51]. With the telescope technique, a 5 MHz frequency shift (AOM driven at 152.5 MHz), causes a power loss of about 90%. This isn't a huge issue because we can still perform power balancing and calibration of the beams by adjusting the AOM drive power in the DDS. All of the 685 nm imaging and cooling operations described in this thesis involve calibration of the DDS power to mitigate the uncalibrated frequency dependent power loss of the AOMs.

We also apply a relative 1 kHz frequency shift to the top vs bottom beams and the front vs back beams. This reduces standing wave effects of the light in the cell without having to use piezo electric actuators to vibrate the mirrors [52].

Auto-relocking the injection-locked diode

Re-locking an injection-locked diode is very simple, involving a simple ramp of the injection current up by about 5 mA, then back down to the nominal current. A script that re-locks the diode was implemented in Python, which sends a set of VISA commands to the Arroyo current controller, which performs this ramp. However, this requires an awareness of an unlock event. These events are rather rare, but can easily happen if the seed power has drifted too much, or if the seed laser has been interrupted for a brief instant (for example, if somebody were to place their hand in the beam path briefly).

To detect an unlock event, an interference measurement of the seed laser with the injection-locked laser is performed on a simple PICAM CCD camera, with an image analysis script. To create an interference pattern, a two beam samplers are placed for the seed laser and the injection-locked laser, similar to a Michelson interferometer ([53]). These beams are then combined on the camera with a small angle between them to create a clean interference pattern. Figure 3.12 shows the interference pattern of the injection-locked diode with the seed laser on the camera. The left (unlocked) image is washed out and doesn't show much of an interference pattern while the right image (locked), does show a clear interference pattern. The left image does have a washed out interference pattern, likely caused by some gain of the seed laser in the injection-locked diode, meaning the diode is running multi-mode. As the current is ramped up and back down, the interference suddenly hops into the injected lasing mode.

To analyze the interference pattern, the image is rotated to align the interference fringes with the y-axis, then a spacial Fourier transform [54, 55] is performed on the x-axis. Since the spacial frequency of the interference pattern is primarily dependent on the angle of incidence between two beams [53], the Fourier transform will have a peak at a specific spacial frequency. When this spacial frequency is above a threshold, the diode is considered to be locked. The camera has a frame rate of 30 Hz, so the analysis can




Figure 3.12: Interference pattern of the injection-locked diode with the seed laser on a PICAM CCD camera for the unlocked laser (left) and locked laser (right). The interference pattern is used to detect an unlock event through Fourier transform image processing.

be performed at this rate, and when the diode ever becomes unlocked, it performs the re-locking script mentioned before.

3.2.4 685 nm Locking Scheme

As mentioned many times through this text, the locking scheme of the 685 nm laser to the ULE cavity is a little bit more complex than a standard locking scheme [56, 57]. The SolsTiS is the resonating cavity of the laser and has two internal feedback mechanisms for control. As is standard, these two loops are a fast and a slow piezo loop. As mentioned in section 5.1.2, the bandwidth of the SolsTiS fast piezo is around 100 kHz, while the slow loop has a bandwidth of around 50 Hz (but a very large dynamic range). As such, the fast loop's responsibility is to narrow the frequency noise of the laser while the slow loop is responsible for handling the slow drifts caused by the many environmental effects such as temperature, humidity, and pressure.




Figure 3.13: Basic feedback system with labels. A setpoint and a measured process variable are compared to create an error signal. A gain stage acts on the error signal to create a control variable, which is then applied to the process. The process then is measured to create the measured process variable.

For references and labels, figure 3.13 shows a basic feedback system [58–60]. The setpoint is the voltage offset of the generated PDH signal. The measured process variable is what is generated from the PDH, which may have an offset. The error signal is the difference between these two, which is then amplified by a gain stage (i.e. a PID) to create a control variable. This control variable is then applied to the process/plant, which in this case is the laser cavity. The process/plant is then measured again to create the measured process variable. For a system with two feedback loops (fast and slow), it is common to use a "double integrator" [60]. But as I lengthily discuss in section 6.1.4, double integrators are problematic. To this end, I present a different locking scheme.

Consider another perspective on the purpose of a slow loop: the slow loop is responsible for maintaining the fast laser lock—to keep the fast loop control variable integrator from acquiring a DC offset. This framing of the problem invites a very interesting solution. Instead of feeding the error signal to the slow loop, we can define a new error signal based on the control variable of the fast loop. The goal is to drive the fast loop control variable to zero using an integrator.

This is effectively the same thing as a double integrator, since the error signal is being integrated twice. However, in this scheme, the second integration is entirely about maintaining the fast loop and preventing it

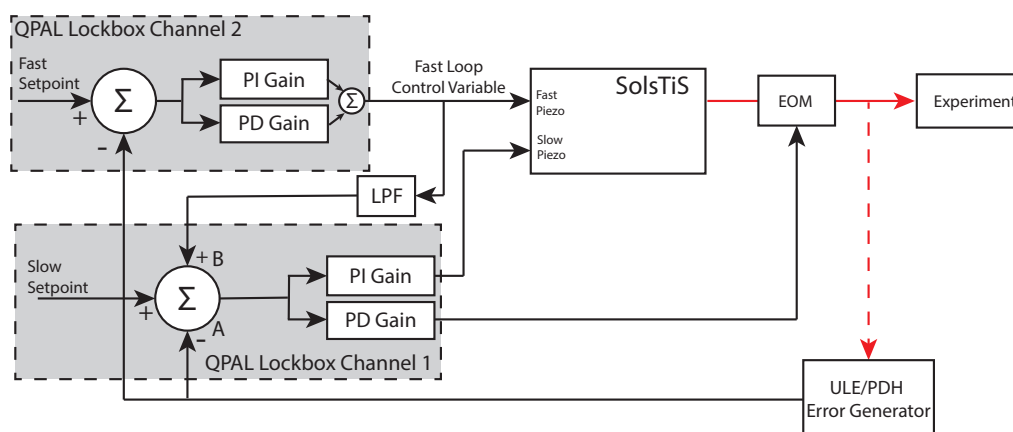


Figure 3.14: Locking scheme for the 685 nm laser. The fast piezo loop is responsible for stabilizing the laser frequency, while the slow piezo loop ensures the fast loop remains locked over longer timescales by monitoring the fast loop control variable after passing through an LPF. The slow loop is effectively a double integrator, but constructed as a PI loop that drives the fast loop control variable to zero. There is also an EOM feedback loop that applies a P-LPF loop to the EOM.

from railing—and it does this without requiring a massive low-frequency gain coefficient. This means we are not limited by the non-linearities of the op-amps: specifically, railing, input offset voltages, drifts, etc.

A detailed depiction of the locking scheme for the 685 nm laser is shown in figure 3.14. To understand this scheme, it is recommended to be familiar with the QPAL lockbox and its features. See section 6.1 for a full description of the lockbox. A PDH signal is generated, as standard. This is sent in parallel to both channels of the QPAL lockbox. For both of these channels, the generated error signal (this is the measured process variable in figure 3.13) is sent to A-in of both channels. The setpoints are generated on-board with trim-pots.

Detailed 685 nm Locking Scheme

Channel 2 (the right channel) of the lockbox is responsible for driving the fast piezo using both of its feedback paths. This is therefore a PI-LPF + PD-LPF gain stage, which is not quite the same as a PID, because it bypasses the nonlinearities of realistic systems introduced by op-amp-based integrators, proportionals, and derivatives. This design allows for clean and independent control of the integrator loop and the proportional loop. The integrator technically has two parameters: a PI corner and a P gain. The same applies for the PD loop. By separating these (and giving them independent low-pass filters), it improves the tuning capabilities of the system. This is described in more detail in the QPAL lockbox section (Section 6.1).

Channel 1 (the left channel) also has two loops utilized, but these are not connected in this case. Here, the fast loop of the lockbox is fed to the EOM. As discussed in Section 3.2.4, the EOM can be modeled to intrinsically act as having a built-in derivative term when applying feedback to the laser. So in this case, a PI loop would actually become a DP loop (i.e., $P \rightarrow D$ and $I \rightarrow P$). However, if the integrator is used and only differentiated, then it is not a good feedback scheme, because there will be nothing to prevent it from wanting to rail. Thus, the ideal gain term for the EOM is a P-LPF loop. A derivative would only become a second-order derivative.

It is also important to note that a DC offset on the EOM will have absolutely no effect on the laser frequency—it will just cause the laser to have an increased phase offset due to a longer effective path length. As such, the low-frequency components of the error signal generated in Channel 1 do not affect this loop.

Channel 1 also has a second loop, which is the PI loop that feeds back to the slow piezo. As mentioned above, the slow piezo is responsible for keeping the loop controlling the fast piezo (Channel 2) from riling, by monitoring the fast piezo control variable. To achieve this, the control variable of Channel 2 is sent through a 1 Hz low-pass filter (LPF) and then sent to the

B-in of Channel 1. The setpoint input is intended to combine with the A-in, so that when the laser is at the lock point, $\text{setpoint} - \text{A-in} = 0$. Then, when the fast piezo control variable is also added in, $\text{setpoint} - \text{A-in} + \text{B-in} = 0$. This way, the slow piezo is both driving the error signal to zero and driving the fast piezo control variable to zero. This is the key to the slow loop: to maintain the fast loop.

This design is great for the longevity of the lock. The lock can survive overnight drifts, and has remained stable for as long as 36 hours at one point. However, the dynamic range of the slow loop is always limited, and eventually, the error signal offset itself drifts, leading to the lock being lost. This is actually a problem with the PDH generation and the cavity. Unfortunately for this 685 nm cavity, the window into the cavity is not optimally coated for 685 nm, so there is an additional cavity: the window into the ULE cavity itself. This causes a second PDH signal to appear, and this signal is much more susceptible to environmental drifts. There isn't much that can be done about this without rebuilding the ULE cavity.

Another reason the laser drifts is the usage of a very old Sprout pump laser, which sometimes has increased amounts of noise, requiring the "self-tuning" mode to be performed every once in a while. In addition, the spatial mode of the pump laser also drifts, requiring a re-coupling to the SolsTiS on occasion.

685 nm Self-Heterodyne Measurements

A self-heterodyne system is also prepared to allow for the continuous monitoring and optimization of the 685 nm laser noise. See section 5.1.3 for a description of the self-heterodyne system used for the 459 nm laser, which is designed similarly to the 685 nm laser, and for the noise fitting performed here. Figure 3.15 shows the analysis of the 685 nm self-heterodyne data to extrapolate the phase noise of the laser. The white noise is around $8.4 \text{ Hz}^2/\text{Hz}$. The servo bumps are at $f_g = 20. \text{ kHz}$, $\sigma_g = 920 \text{ Hz}$, $h_g = 15 \times 10^3 \text{ Hz}^2/\text{Hz}$

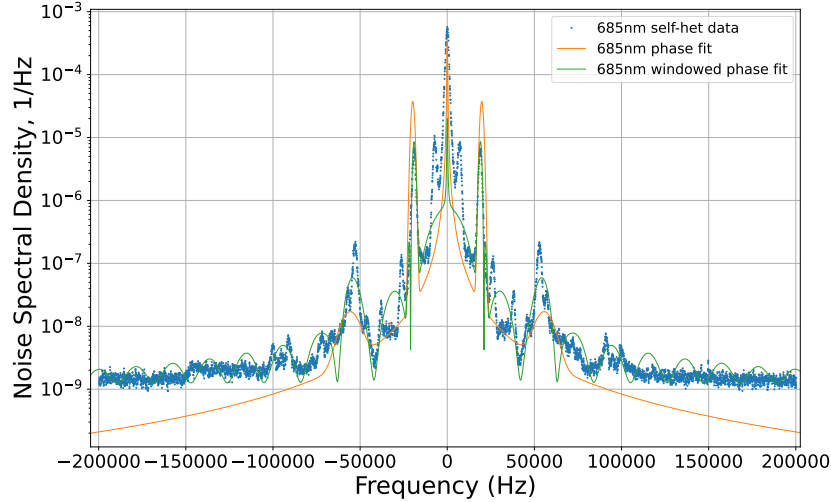


Figure 3.15: 685 nm Self-Heterodyne phase noise data, windowed fit, and extrapolated phase noise from the windowed fit. White noise: $h_0 = 8.4 \text{ Hz}^2/\text{Hz}$. Fit Servo Bumps at $f_g = 20. \text{ kHz}$, $\sigma_g = 920 \text{ Hz}$, $h_g = 15 \times 10^3 \text{ Hz}^2/\text{Hz}$ and $f_g = 57 \text{ kHz}$, $\sigma_g = 5 \text{ kHz}$, $h_g = 46 \text{ Hz}^2/\text{Hz}$. More bumps are clearly seen, but have much less contribution. Analyzed using the same tools and techniques presented in section 5.1.3.

and $f_g = 57 \text{ kHz}$, $\sigma_g = 5 \text{ kHz}$, $h_g = 46 \text{ Hz}^2/\text{Hz}$. More bumps are clearly seen, but have much less contribution. Since the 685 nm laser is not used for phase sensitive operations, the only issues that are of primary concern are servo bumps that are far from the center mode. We do not care if the laser has a fluctuating Rabi rate or phase accumulation from servo bump noise because we are not performing Qubit operations with this laser, we are simply cycling the atoms. As such, the only concerns we have are if there is heating caused by a servo bump exciting the atoms closer to resonance, instead of from a red-detuned beam.

Using an EOM external to the cavity

When an EOM is in a cavity, a DC electric field will cause the laser to observe a longer phase delay for a round trip through the cavity [61–64]. Thus, a DC signal into a cavity based EOM will cause a DC frequency shift of the laser, since the frequency ω depends on the cavity length, L :

$$\omega = \frac{2\pi c}{L} \quad (3.6)$$

where c is the speed of light. Thus, if $L \rightarrow L + dL$ because of the in-cavity EOM, the frequency changes to

$$\omega' = \frac{2\pi c}{L + dL} \approx \frac{2\pi c}{L} \left(1 - \frac{dL}{L}\right) \quad (3.7)$$

with a first order expansion on dL/L . The frequency shift is proportional to the change in length of the cavity, or in this case, the effective cavity length.

However, we do not have an EOM in the cavity. Instead, a DC electric field applied to the EOM will add a DC phase shift to the light passing through, which does not change the laser frequency. But if we apply a time-dependent electric field, then the phase shift applied will also be time-dependent.

Out-of-Cavity EOM as a Derivative Actuator

Lets model the electric field of a laser with

$$E \sim e^{-i(\omega t - kx)} \quad (3.8)$$

with ω the frequency, k the wavenumber, t time, and x displacement. To simplify things, we will be looking at a particular point in space, taking x to be constant and absorbing its contribution to the ignored complex coefficient outside of the exponential (hidden inside of " \sim "). An EOM has

a linear relationship between the electric field applied and the change in index of refraction of the crystal [65]. And since the crystal has a constant length and capacitance, there is a linear relationship between relative phase accumulation and voltage:

$$\delta\phi(t) = \frac{\pi V(t)}{V_\pi(\lambda)} \quad (3.9)$$

where $V_\pi(\lambda)$ is the voltage required to apply a π phase shift for a given wavelength. This is a number often cited in datasheets. The important factor here is that there is a linear relationship between the applied voltage and the relative phase shift of the laser after traveling through the crystal.

A PDH error signal, in the small error linear regime, is a linear relationship between error signal voltage, $e(\omega_{\text{err}})$ and frequency error, ω_{err} [56, 57]:

$$e(\omega_{\text{err}}) \sim \omega_{\text{err}} \quad (3.10)$$

Suppose a PID is applied to this signal and sent to the EOM, $V_{\text{EOM}}(\omega_{\text{err}}) = \text{PID}\{e(\omega_{\text{err}})\}$. Thus, the phase correction is linear to the frequency error (instead of the frequency correction being linear to the frequency error).

If we revisit the electric field of our laser, eq 3.8, but now with phase:

$$\mathbf{E} \sim e^{-i(\omega t + \delta\phi(t))} \quad (3.11)$$

$$\sim e^{-i(\omega + \delta\phi(t)/t)t} \quad (3.12)$$

such that

$$\delta\omega(t) = \phi(t)/t \sim V_{\text{EOM}}(t)/t \quad (3.13)$$

This means that, in order to correct a DC frequency error, a linear voltage ramp has to be applied. This is akin to an integrator. It is the derivative of the voltage change that causes the frequency of the laser to change.

If, in a time-step dt , there is a laser frequency error accumulation, $d\omega_{\text{err}}$

then, from the PID and EOM, the phase adjustment will become:

$$de(\omega_{\text{err}}) \sim d\omega_{\text{err}} \quad (3.14)$$

$$d\phi \sim \text{PID}\{de(\omega_{\text{err}})\} \sim \text{PID}\{d\omega_{\text{err}}\} \quad (3.15)$$

and the frequency correction in this dt will therefore be:

$$\omega_{\text{cor}} dt = d\phi \quad (3.16)$$

$$\omega_{\text{cor}} = \frac{d\phi}{dt} \quad (3.17)$$

$$= \frac{d}{dt}\{\text{PID}\{d\omega_{\text{err}}\}\} \quad (3.18)$$

$$= \text{DPD}^2\{d\omega_{\text{err}}\} \quad (3.19)$$

Where the DPD^2 is a Derivative-Proportional-Double Derivative resulting from a Proportional-Integral-Derivative gain, respectively. As such, a PI loop will become a "DP" loop. However, it is not actually useful to perform an integrator loop for this. As the integrator accumulates phase, the op-amps will quickly rail and no longer be able to contribute to applying a proportional frequency adjustment. The DC accumulation of the integrator would be completely useless. As such, these external EOMs can be useful for implementing a derivative loop for high frequency error correction, but cannot perform proportional or integral gain if a real op-amp is used.

Bandwidth Limitations of the EOM

The actual bandwidth of an EOM can be impressively high, with demonstrations at bandwidths greater than 10GHz [66, 67]. In our case, the bandwidth limitation is the electric limitations caused by the capacitance of the crystal and non-zero output impedance of the lockboxes. The Thorlabs EO-PM-NR-C1 (which is what was used here) has an input capacitance of 14pF. This is a negligible contribution compared to the capacitance of

an RG-58 coaxial cable, which tends to be about 100 pF/m. With a three meter cable, and the maximum short-circuit current output of the AD8429 being 35mA (see datasheet [68]), the maximum rate the QPAL lockbox can drive the voltage of the coils, assuming a 2 meter cable, is

$$\frac{I}{C} = \frac{35 \times 10^{-3} \text{A}}{214 \times 10^{-12} \text{F}} = 160 \text{MV/s}$$

which is greater than the slew rate of 22MV/s (see datasheet). But keep in mind that if the cable length is greatly increased, then the capacitance will become higher and may become the limiting factor.

Either way, the V_{pi} at 685 nm for our EOM is about 220V [69]. With the 22MV/s speed limit of the lockbox, the maximum phase accumulation would be 300 krad/s, or 50kHz. It seems pretty clear that this is a major limitation that I had never considered until I finally did the calculations in writing this thesis. However, these can easily be fixed and improved by using a buffer op amp gain stage after the output of the QPAL lockbox.

The AD825 has a open loop bandwidth of 41 MHz, a slew rate of 125 MV/s, and a short-circuit current of 80mA, according to the datasheet [70]. If this op amp was set up to apply a gain of 10, then the effective slew rate of the QPAL lockbox output can be increased by a factor of 10. With these numbers, a maximum phase accumulation rate of $2\pi \times 300$ kHz could be reached. Care should be taken to minimize stray capacitance by using high quality connectors and short cables.

4 COOLING AND IMAGING CS QUBITS USING THE QUADRUPOLE INTERACTION

4.1 Implementing Free-Space Second-Stage Cooling

In this section, I discuss the results of the free-space cooling of the atoms using the 685 nm laser. This cooling involves 6-axis cooling on the $6S_{1/2}, F = 4$ to $5D_{5/2}, F = 6$ transition with red-detuned 685 nm light. This is studied with and without a quadrupole magnetic field gradient. A brief comparison to theoretical quadrupole cooling models in [71] is also presented.

This is the beginning to part of the work presented in [36].

4.1.1 685 nm MOT Experimental Sequence

We use the experimental apparatus described throughout section 3.1. We start by loading a 3D MOT using the 2D MOT and push beam into an 852 nm MOT, with an anti-Helmholtz magnetic field gradient of about 6 G/cm along the axis of the coils. After loading the atoms into the 852 nm MOT for about 1000ms, the magnetic field gradient and bias fields are ramped to the new parameters for the 685 nm MOT cooling stage. The actual fields depend on if a molasses or MOT is being performed with the 685 nm light field. In addition, different 685 nm axes can be enabled or disabled to allow for different types of experiments to be performed. For example, the experiment sequence for spectroscopy presented in section 4.1.2 involves enabling just one axis of the 685 nm light for a brief amount of time.

Imaging of the MOT cloud is performed using an Andor Luca EM-Gain camera with an 852 nm line filter. To image the resulting atom cloud after the 685 nm cooling phase, the 685 nm light is turned off and the 852 nm light is turned on, with magnetic fields for creating an optical molasses [33],

meaning zero gradient and a bias fields zeroing out the magnetic fields at the atoms. The image of the atoms is typically taken with a 10ms integration time and is triggered by the HSDIO. This integration time is sometimes longer or shorter depending on the experiment. The use of molasses fields makes a less spatially dependent force on the atoms, allowing for images to be taken without deforming the atom cloud, since the shape of the atom cloud is an important characteristic of the measurements required for temperature measurements.

4.1.2 685 nm Free-Space Push Spectroscopy

To perform spectroscopy of the 685 nm light, a single axis of the 685 nm light is enabled, with the other axes disabled. After the 852 nm MOT has been loaded, the magnetic fields are quickly ramped to a molasses level, a single axis of the 685 nm light is turned on for a brief moment, then an image of the resulting cloud is taken. The closer to resonance the 685 nm light is, the more the atom cloud is pushed. The force experienced by the atoms is proportional to the scattering rate [33], which goes as

$$r(\mathbf{t}) = \frac{\Gamma}{2} \frac{s}{1 + s + 4\Delta^2/\Gamma^2} \quad (4.1)$$

where $\Gamma = 2\pi\gamma$ is the natural linewidth of the transition, $s = \frac{I}{I_s}$ is the saturation parameter, and Δ is the detuning from resonance. However, atom temperature is a big factor in the detuning and will artificially broaden this resonance [33]. Regardless, we can still perform a spectroscopy scan and fit to a gaussian to determine the resonant frequency and effective linewidth.

For each iteration of the spectroscopy, the 852 nm MOT and subsequent cooling stage is performed twice. The first time, a 1 ms 685 nm push beam on the front axis is applied. The second time, the atoms are allowed to free-expand for the same 1 ms duration. After the second atom cloud has dispersed, a third shot is taken as a background measurement. To analyze

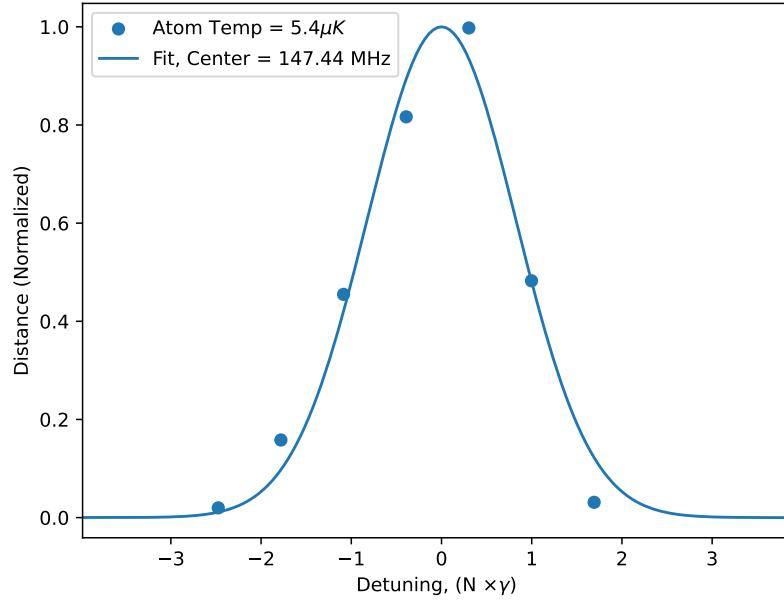


Figure 4.1: Results of a 685 nm push spectroscopy scan with a 685 nm cooling phase, fit to a Gaussian. The measured atom temperature after the cooling phase was $5.4 \mu\text{K}$ as is discussed in section 4.1.3. The measured center frequency setting for the AOMs is 147.44MHz and the FWHM is 230kHz.

the data, the background subtracted images are fit to a 2D gaussian, and the distance between the center of the Gaussians is calculated. This measurement is repeated many times for each frequency setpoint in the spectroscopy. From basic projectile motion, we know that the displacement is proportional to the force, which in turn is a function of the detuning. Figure 4.1 shows the results of a spectroscopy scan of the 685 nm light after a 50 ms 685 nm MOT cooling phase with a -1.0γ detuning (see 4.1.3. The measured atom temperature after the cooling phase was $5.4 \mu\text{K}$.) 4.1.3. During the 685 nm phase, the repump light is turned off, demonstrating the high pump-depump ratio, as mentioned in section 2.4 and further investigated in section 4.5.2.




Figure 4.2: Images of the atom cloud with and without the 685 nm push beam activated, as well as a subtraction between the two. a) on resonance, b) off resonance.

Figure 4.2 shows the images of the atom cloud with and without the 685 nm push beam activated, as well as a subtraction between the two for on- and off-resonant 685 nm light.

This push spectroscopy can also be used to optimize the 685 nm beam pointing. When on resonance, each beam can be aligned to optimize the push distance of the atom cloud.

4.1.3 685 nm Free-Space Cooling

With the resonant frequency measured, we can analyze the performance of the 685 nm cooling with respect to different detunings and magnetic field gradients. As mentioned above, the 685 nm light is applied after loading the 852 nm MOT. In this section, I present the results of the 685 nm cooling

with a 1.5 G/cm magnetic field gradient and a 0.25 G/cm magnetic field gradient (near-PGC). To measure the atom cloud temperature, time-of-flight (TOF) measurements are performed.

TOF measurements involve dropping all light fields for a amount of time, dt , after the cooling phase. Then the atoms are caught in the 852 nm optical molasses field. During this time, the atoms will expand freely as well as fall under the influence of gravity, which can be used to calibrate the size of the atom cloud. References [72–74] have discussed this technique in detail and it has become a common technique for measuring atom cloud temperatures.

The expansion of an atom cloud in free-space can be described by the following equation:

$$f_r(t) = N \left(\frac{2}{\pi w_r^2(t)} \right)^{3/2} \exp \left(-\frac{2r^2}{w_r^2(t)} \right) \quad (4.2)$$

where $f_r(t)$ is the atom density function of distance r and time t , N is the total number of atoms, and $w_r(t)$ is the radius of the atom cloud at time t . The radius of the atom cloud can be described by:

$$w_r(t) = \sqrt{w_{r0}^2 + \frac{4k_b T t^2}{m}} \quad (4.3)$$

with w_{r0} being the initial gaussian radius of the atom cloud, k_b the Boltzmann constant, T the atom cloud temperature, and m the mass of the atoms.

To measure a temperature, the atom cloud is imaged at a series of free-expansion time steps after the cooling phase, and each (averaged) point is fit to a symmetric 2D Gaussian. The Gaussian waist vs. time is then fit to Equation 4.3 to extract the temperature of the atoms. If the atoms evolve too far from the center of the image, asymmetric light fields can distort the shape of the atom cloud during imaging, which can lead to error in the temperature measurement.

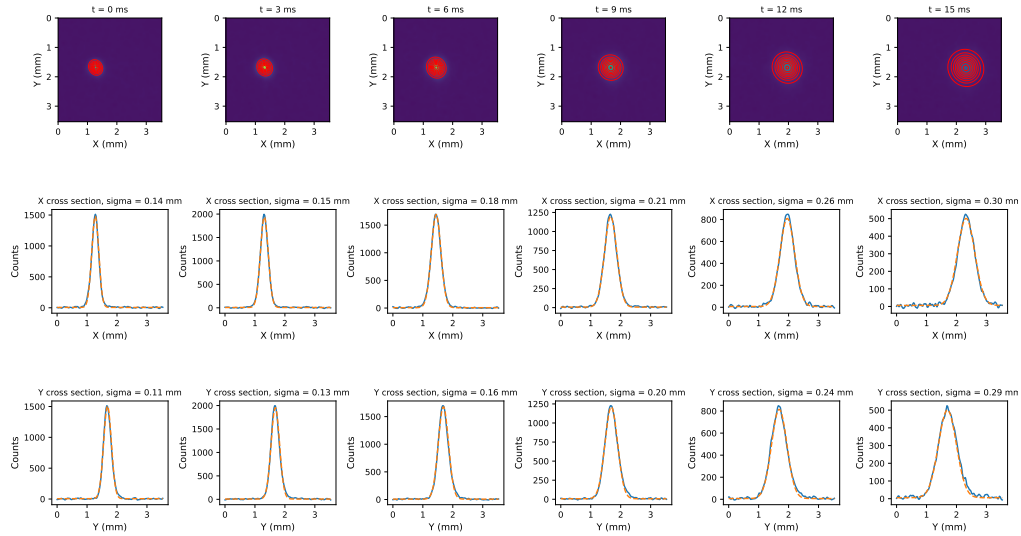


Figure 4.3: 685 nm MOT atom cloud images for a 1.5 G/cm magnetic field gradient and a detuning of -1.25Γ . The time of flight is measured at 0, 3, 6, 9, 12, and 15 ms after the cooling phase. Each image is fit to a symmetric 2D Gaussian, with the x and y cross-sections also shown.

To calibrate the size of a pixel in the image plane of the camera, the central location of the atom cloud and its displacement from the initial location over time is analyzed and fit to a parabola. This is then compared to the known value of gravity to extrapolate the millimeters-per-pixel calibration.

The time of flight is measured at 0, 3, 6, 9, 12, and 15 ms after the 685 nm cooling phase. Figure 4.3 shows each shot of the atom cloud and its fit to a 2D Gaussian for these time steps, for a 1.5 G/cm magnetic field gradient and a detuning of -1.25Γ . The x and y cross-sections are also shown. Figure 4.4 plots the measured atom cloud symmetric Gaussian radius vs. time and includes a fit to Equation 4.3 to extract the temperature. The measured temperature at these parameters is $T = 5.3 \mu\text{K}$.

This data is also taken for a series of different detunings. Unfortunately, too large a detuning causes the atoms to disperse too quickly. Figure 4.5

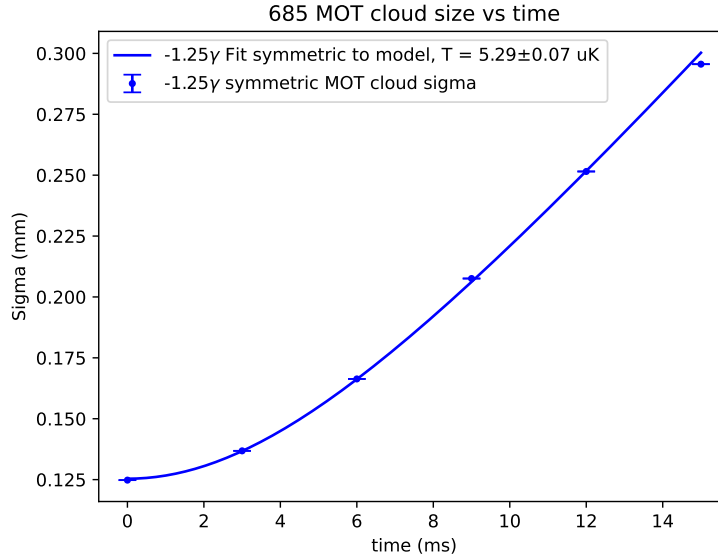


Figure 4.4: Measured atom cloud symmetric Gaussian radius vs. time for a 1.5 G/cm magnetic field gradient and a detuning of -1.25Γ . The points are fit to Equation 4.3 to extract the temperature of the atom cloud, which is $T = 5.3\ \mu\text{K}$.

shows the measured temperature of the atom cloud vs. detuning for a 1.5 G/cm and a 0.25 G/cm (near molasses) magnetic field gradient. The minimum temperature is around -1.25Γ detuning for both cases and begins to increase drastically for detunings below -1.0Γ . Performing power dependence scans is not productive because the MOT cloud disperses very quickly.

4.1.4 Comparison to Theory

As mentioned before, [71] presents a theoretical framework for the feasibility and performance of cooling on a quadrupole transition, specifically discussing the $6S_{1/2} \rightarrow 5D_{5/2}$ transition in Cs. This is analyzed as a molasses, not a MOT, for both $\sigma^+ - \sigma^-$ and lin \perp lin polarization schemes. In our case, the

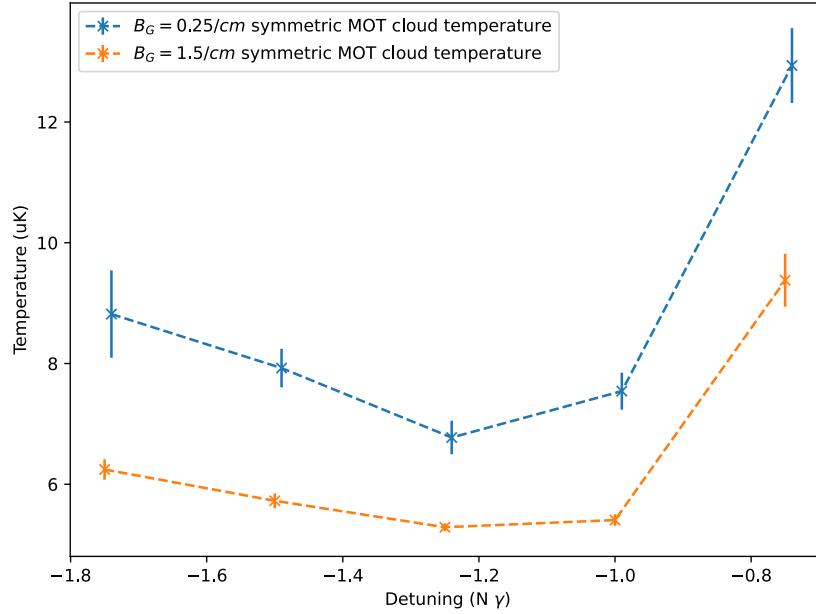


Figure 4.5: Measured temperature of the atom cloud vs. detuning for a 1.5 G/cm and a 0.25 G/cm (near molasses) magnetic field gradient. The minimum temperature occurs at approximately -1.25Γ detuning for both cases and increases drastically for detunings less than -1.0Γ .

MOT is performed with a theoretically circularly polarized scheme, but the cell wall is poorly coated, so this is not necessarily maintained.

Figure 4.6 shows the theoretical calculations of the atom cloud temperature vs. detuning, taken from [71], for a 685 nm quadrupole cooling molasses plotted against the dimensionless parameter \mathbf{u} , defined as

$$\mathbf{u} = \frac{\delta}{\omega_r} \frac{|\Omega|^2}{\delta^2 + \gamma^2/4} = 77 \frac{\delta}{\gamma} s \quad (4.4)$$

with δ the detuning, γ the natural linewidth of the transition, $s = \frac{I}{I_s}$ the saturation parameter, and $|\Omega|$ the Rabi frequency. The figure shows theoretical results for both polarization schemes. The vertical axes are in units of the photon recoil energy, $\hbar\omega_r = 0.154 \mu\text{K}$. The values of \mathbf{u} for

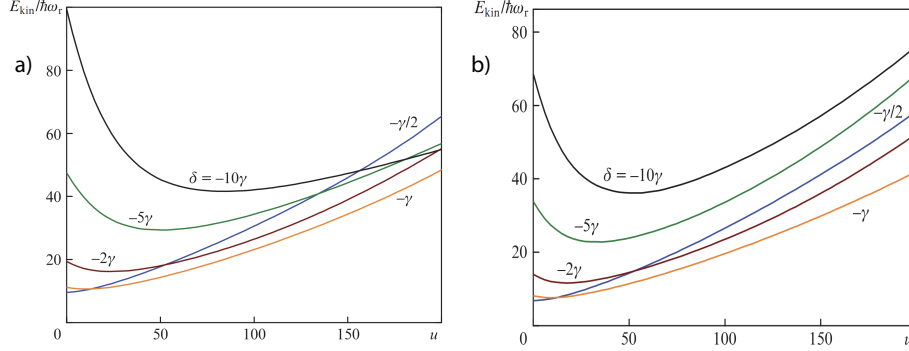


Figure 4.6: Theoretical calculations of the atom cloud temperature vs. detuning for 685 nm quadrupole cooling molasses, shown as a function of the dimensionless parameter $u = \frac{2\delta s}{\omega_r} = 77 \frac{\delta}{\gamma} s$ for (a) $\sigma^+ - \sigma^-$ and (b) $\text{lin} \perp \text{lin}$ polarization schemes. Figure reproduced from [71]. Vertical axes are in units of the photon recoil energy, $\hbar\omega_r = 0.154 \mu\text{K}$. The values of u for $s = 1.7$ at detunings $-\gamma/2$, $-\gamma$, and -2γ are 66, 132, and 264, respectively.

$s = 1.7$ at detunings $-\gamma/2$, $-\gamma$, and -2γ are 66, 132, and 264, respectively.

Extrapolating from the plot, the theoretical temperatures at these detunings are approximately $3 \mu\text{K}$ for $-\gamma/2$, $4.6 \mu\text{K}$ for $-\gamma$, and $11 \mu\text{K}$ for -2γ (this last value requires extrapolating the trend, as $u = 264$ is off the plot). We were unable to run our experiment at -2γ because the atom cloud dispersed too quickly, and reducing the 685 nm laser power significantly made the atom cloud very difficult to maintain.

At low detuning, the measured atom temperature begins to spike, which does not match theoretical predictions. I believe this is due to servo bumps around 50 kHz in the 685 nm laser causing heating of the atoms. If photons are near resonance, they will heat the atoms instead of cooling them through the Doppler effect.

Further study of atom temperature and the performance of 685 nm cooling should be conducted in the future, with more data points and improved beam control (e.g., better polarization). In addition, it would be interesting to attempt to capture atoms directly from the 2D MOT into the

685 nm MOT. When these experiments were performed, the 2D MOT was not well implemented and barely improved our 3D MOT loading rate. As a result, the 852 nm MOT was loaded primarily from room-temperature background gas. If the 2D MOT had been operational at the time, the atoms would have been colder, and the 685 nm cooling could have been more effective.

4.1.5 685 nm MOT Atom Cloud Fluorescence

The last experiment performed in free space with the 685 nm light was to take a picture of a background-free atom cloud by imaging the 852 nm fluorescence of the atoms while illuminating only with 685 nm light. This was also performed without using any repump light.

Figure 4.7 presents the background-free atom cloud fluorescence image. The image exposure time had to be very long to acquire enough photons on the camera. For this shot, the camera zoom lens had an approximately 2-inch clear aperture and was placed about 20 inches away from the cell. The size of the pixels in the image is $53 \mu\text{m}/\text{pixel}$, and the image was taken with a 100 ms exposure time.

4.1.6 685 nm MOT Next Steps

The 685 nm MOT can be a very useful technique for pre-cooling atoms before trapping them in a trap array. It is easy to cool atoms down to $5 \mu\text{K}$ in a low-gradient magnetic field, and the atoms can then be cooled directly into a trap. It would be very interesting to attempt to capture single atoms directly from the 2D MOT into the 685 nm MOT—and even more so if they could be transferred into a molasses. With a molasses, the atoms could also be cooled into a trap array, potentially allowing for loading of single atoms without destroying existing atoms in the array.




Figure 4.7: Background-free atom cloud fluorescence image using only 685 nm light. The image exposure time was 100 ms to acquire sufficient photons on the camera. Pixel size is $53 \mu\text{m}/\text{pixel}$. No repump light was used for this image, demonstrating the high pump-depump ratio of the 685 nm light.

Further analysis of 685 nm cooling should be performed, including more data points and improved beam control (especially polarization). It would be interesting to attempt even colder temperatures. Improving the spectral performance of the 685 nm laser would also be beneficial, as servo bumps at 50 kHz can cause heating of the atoms. One possible improvement—discussed in more detail in Section 7.1—is to set up an injection-lock amplification stage for light transmitted through the ULE cavity, which could attenuate the servo bumps due to its narrow linewidth.

In addition, increasing the available power by implementing six injection-locked diodes—one for each axis—instead of just two could significantly enhance performance. This would require substantial infrastructure upgrades

but could allow for up to twice as much 685 nm power, enabling further study of the dependence of 685 nm cooling performance on laser power.

4.2 Trapping Single Atoms

To trap single atoms in the trap array described in section 3.1.3, a standard 852 nm 3D MOT is loaded (with an improved 2D MOT) for 300ms. This is followed by enabling the trap light and performing a 5ms PGC cooling phase. The parameters (frequency, intensity, time, and Bx/y/z shims) were optimized using the Nelder-Mead optimization algorithm (see [75]) with the cost function maximizing the amount of photons collected during an extra long imaging phase. The imaging magnetic fields were optimized similarly. Finally, an LAC was enabled as an intermediate step to optimize single atom loading, optimized through a similar cost function.

4.2.1 852 nm Imaging of Single Atoms

Figure 4.8 a presents imaging statistics for 852 nm imaging for single atoms in the trap array. These images are taken with a -5.6γ detuning, a 70ms exposure time, and $2.5I_{\text{sat}}$ of intensity across all 6 beams. The images are taken with 2x2 charge binning performed on the EMCCD camera before the analog measurement, and are then digitally binned on the analysis computer in a 3x3 bin. From the figure, we see a separation of about 250 photoelectrons between the two peaks.

For the data shown in Figure 4.8, a second shot is also taken after the first, with 685 nm cooling performed in between the two shots (see Section 4.4). This allows us to determine the probability that the imaging process heats the atom out of the trap. Figure 4.8(b) shows a scatter plot of the number of photoelectrons in the first shot vs. the number in the second shot. Dashed lines on both axes indicate the thresholds used to classify whether an atom is present.

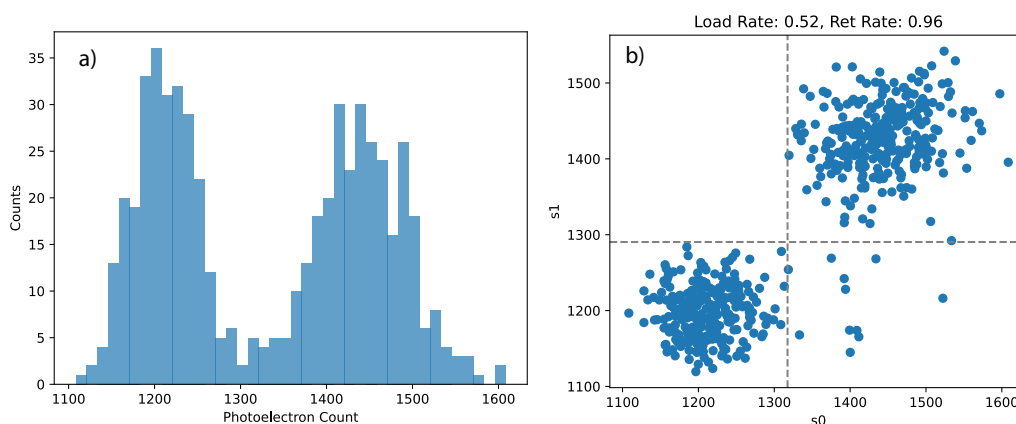


Figure 4.8: 852 nm imaging statistics for single atoms in the trap array. The images are taken with a -5.6γ detuning, a 70 ms exposure time, and a total intensity of $2.5 I_{\text{sat}}$ across all six beams. A 2×2 charge binning is applied on the EMCCD camera before the analog measurement, followed by a 3×3 digital binning step on the analysis computer. (a) A histogram of counts for shot 0. (b) A scatter plot where each point represents data from a single atom shot; the x-axis shows photoelectron counts in the first shot, and the y-axis shows photoelectron counts in the second shot. The dashed lines indicate threshold locations used to determine whether an atom is present or vacant.

Using these thresholds, four types of events can be identified, based on which quadrant a given point falls into:

- **Bottom-left quadrant:** No atom present in either shot.
- **Top-right quadrant:** Atom present in both shots.
- **Bottom-right quadrant:** Atom present in the first shot, but not the second—likely due to heating out of the trap.
- **Top-left quadrant:** Atom present in the second shot, but not the first—should not occur, and would imply that a new atom entered the trap mid-sequence. This outcome could happen if there were multiple

atoms in the trap, but in this experiment, only one atom was loaded. No such cases were observed in this dataset.

The retention rate of this dataset is 96 %, calculated as the fraction of cases where an atom was present in the second shot given that it was present in the first.

Not much effort was put into optimizing the 852 nm imaging parameters, since this imaging beam is primarily used as a diagnostic tool to implement and monitor 685 nm cooling, imaging, and readout. We are not attempting to perform high-fidelity quantum operations with this light—our goal was simply to establish a reliable tool. However, this imaging line is widely used for cesium atom detection in many other experimental platforms. Significant work has been done on optimizing imaging fidelity on this line, such as in [19, 76, 77].

4.2.2 EMCCD Imaging Light Capture Efficiency

These imaging parameters for the data in Figure 4.8 can tell us how many photons are scattered by the atom in 70ms. Using the scattering rate equation,

$$r(t) = \frac{\Gamma}{2} \frac{s}{1 + s + 4\Delta^2/\Gamma^2}$$

where $\Gamma = 2\pi \times 5.2$ MHz is the natural linewidth of the transition, $s = \frac{I}{I_s}$ is the saturation parameter, and Δ is the detuning from resonance. For our imaging parameters, we have $\Delta = -5.6 \gamma$, $s = 2.5$, and $\Gamma = 2\pi \times 5.2$ MHz. In 70ms, this rate gives us a count of 22,000 photons scattered by the atom. Not all of these photons make it to the camera. There is an estimated optical train efficiency of 0.8 due to the many optical components and filters. Additionally, the EMCCD has a datasheet-specified quantum efficiency of 55% at 852 nm, and the $NA = 0.55$ of the objective lens gives a collection efficiency of 8.2%. Together, these give us a total efficiency of $0.8 \times 0.55 \times 0.082 = 0.036$ (3.6%), supposedly giving us 800 photoelectrons on the camera.

Since we only see around 250 photoelectrons, this suggests there is an additional loss of 67%, which is unaccounted for, meaning that the actual efficiency of photoelectrons reaching the camera from the atom is 1.2%. This is likely caused by issues due to the imaging train, (see the imaging setup section for more details, 3.1.4).

However, there could also be issues with the EMCCD camera. For starters, the EM gain of the camera is supposed to be 1200 at the full gain setting. However, upon measuring the EM gain by imaging the same MOT at different EM gain settings, the ratio of counts measured of from the MOT at the two settings is only a factor of 44.5. According to the datasheet, the EM gain at the lowest setting should be 4 and at the highest setting it should be 1200. If we assert that the EM gain at the bottom of our measured range is indeed 4, then the EM gain at the top range would be 178. This is the number that we used to calculate the number of photoelectrons from the camera counts. But since the EM gain is not actually 1200, this indicates that the camera has degraded in sensitivity, and there could be other issues with the camera.

4.2.3 Measuring trap frequency

To ensure the trap is working properly and is not being heavily distorted from diffraction or aberrations, the trap frequency is measured and compared to the values calculated from the trap shape measurements in section 3.1.3 ($f_x = 10.9$ kHz, $f_y = 9.9$ kHz). The trap shape is not symmetric between the transverse axes (orthogonal to the k-vector) and the longitudinal axis and we don't have any good measurements of the longitudinal trap frequency, but it is always expected to be less than half of the transverse trap frequency [39] because the traps are only closed due to the diffraction.

To measure the trap frequency, we utilize the technique in [78, 79]. This technique involves performing a series of two short trap drops, with an evolving time delay between the shots. In our implementation, we also apply

685 nm cooling before the trap drops. The duration of the trap drops are $30\mu\text{s}$ and have a survival rate of about 90%. After the first drop, the atoms are trapped again for a short amount of time before the second drop. The time between the two drops is varied from $20\mu\text{s}$ to $90\mu\text{s}$.

This can be thought of as applying a kick to the atoms with the first drop, imparting kinetic energy to the atoms. The atoms then oscillate in the trap, at the trap frequency, and depending on the phase of the oscillation at the time of the second drop, the probability that the atom will be lost varies. For example, if at the time of the second drop, the atom is at the bottom of the trap, it will be less likely to be lost than if it is at the edge. However, this oscillation occurs at double the trap frequency, since a full round trip occurs when the atom returns to the original position, not at the halfway point.

Figure 4.9 shows the results of the trap frequency measurement. The x-axis is the time between the two drops, and the y-axis is the retention rate. The data is fit to a sinusoidal function of the form $A \sin(2\omega t + \phi)$, and the trap frequency, ω is directly extracted from the fit (since the fit includes the factor of 2 in the frequency). The measured trap frequency is $\omega = 2\pi \times 10.55 \text{ kHz}$, which is very close to our calculated values of 10.9 kHz and 9.9 kHz. Since these values agree, it indicates that our trap shape is indeed parabolic (this was our assumption in the calculations) and is not heavily distorted by diffraction or other aberrations at the bottom of the potential.

The confirmation of the parabolic trap potential is an important factor for how we calculate our atom temperatures in section 4.4.1.

4.2.4 Optical Pumping Single Atoms into Qubits

The next tool we develop and implement for our trapped atoms, for the purpose of measuring the performance of our 685 nm cooling capabilities, is optical pumping. This is important for measuring the atom temperature

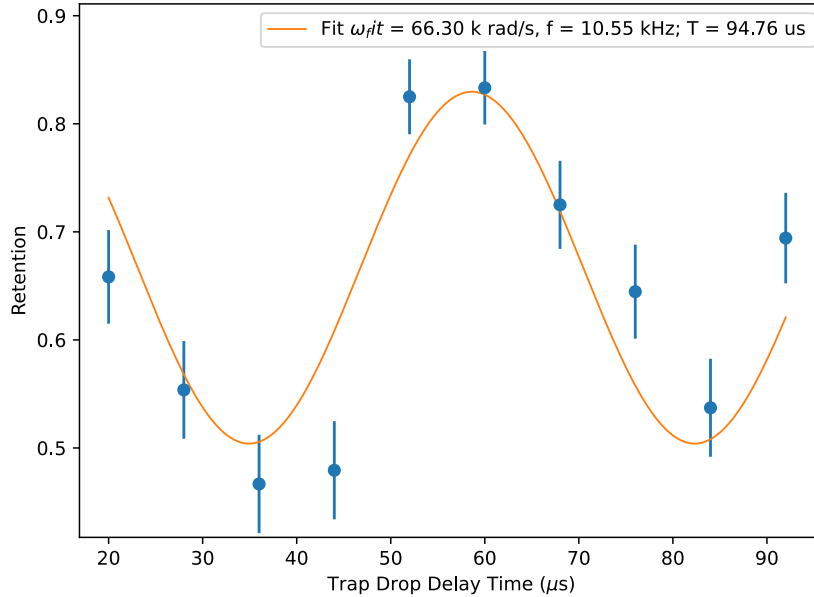


Figure 4.9: Measured trap frequency from the trap drop technique. The time delay between two trap drops is scanned to extract the trap frequency. The data is fit to a sinusoidal function of the form $A \sin(2\omega t + \phi)$, and the trap frequency, ω is directly extracted from the fit (since the fit includes the factor of 2 in the frequency). The measured trap frequency is $\omega = 2\pi \times 10.55 \text{ kHz}$, which is very close to our calculated values of 10.9 kHz and 9.9 kHz.

through Ramsey gap time measurements, as discussed in Section 4.4.1. This section discusses the implementation of optical pumping to prepare the atoms in the $6S_{1/2}, F = 4, m_F = 0$ state, converting our trapped atoms into qubits.

To perform optical pumping (OP), we utilize a laser resonant with the $6S_{1/2}, F = 4 \rightarrow 6P_{1/2}, F' = 4$ transition at 895 nm. This laser is linearly polarized and is applied at about 10° from the normal to the x-axis of the cell (as seen in Figure 3.1). A 2.6 G magnetic field is applied along the y-axis (specifically, orthogonal to the OP laser) of the cell and is used to lift the degeneracy of the m_F states.

Optical Pumping Dark State

The $m_F = 0$ state of the $6S_{1/2}, F = 4$ manifold forms a dark state under linearly polarized light resonant with the $6S_{1/2}, F = 4 \rightarrow 6P_{1/2}, F' = 4$ transition because it does not couple to the $6P_{1/2}, F' = 4, m_{F'} = 0$ state. This is due to the Clebsh-Gordan coefficient $C_{J_1 0 1 0}^{J_2 0}$ being zero. The key reason for this zero is through the orthogonality requirements of the total angular momentum states, $|J_1, J_2; J, m_J\rangle$. From the angular momentum selection rules of an E1 transition ($J_2 = 1$), the possible values of J are $J = J_1 - 1, J_1, J_1 + 1$. Since $|J_1, 1; J, m_J\rangle = \sum |J_1, m_{J1}\rangle|1, m_{11}\rangle$, and if we only look at the state where $m_J = m_{J1} + m_{11} = 0$, then $|J_1, 1; J, 0\rangle = \sum |J_1, m_{J1}\rangle|1, -m_{J1}\rangle$. This sum has three terms: $m_{J1} = 0, 1, -1$. The magnitude of coefficients for the $m_{J1} = \pm 1$ terms must be equal due to symmetry:

$$\langle j_1, m_1; j_2, m_2 | J, m_J \rangle = (-1)^{j_1+j_2-J} \langle j_1, -m_1; j_2, -m_2 | J, -M \rangle \quad (4.5)$$

which, for $m_J = 0$ and $j_2 = 1$:

$$\langle j_1, m_1; j_2, -m_1 | J, 0 \rangle = \langle j_1, -m_1; j_2, -m_1 | J, 0 \rangle \text{ for } J = j_1 \pm 1 \quad (4.6)$$

$$\langle j_1, m_1; j_2, -m_1 | J, 0 \rangle = -\langle j_1, -m_1; j_2, -m_1 | J, 0 \rangle \text{ for } J = j_1 \quad (4.7)$$

Putting this in a tabular form, this can be visualized as:

	$ J + 1, 0\rangle$	$ J, 1\rangle$	$ J - 1, 0\rangle$
$ J, 1; 1, -1\rangle$	A	B	C
$ J, 0; 1, 0\rangle$	D,	E,	F
$ J, -1; 1, 1\rangle$	A,	-B,	C

where

$$A = C_{J,1,1,-1}^{J+1,0} \quad (4.8)$$

$$B = C_{J,1,1,-1}^{J,0} \quad (4.9)$$

$$C = C_{J,1,1,-1}^{J-1,0} \quad (4.10)$$

$$D = C_{J,0,1,0}^{J+1,0} \quad (4.11)$$

$$E = C_{J,0,1,0}^{J,0} \quad (4.12)$$

$$F = C_{J,0,1,0}^{J-1,0} \quad (4.13)$$

$$(4.14)$$

are the Clebsch-Gordan coefficients, $C_{J_1, m_1, J_2, m_2}^{J, m_J}$, for the different states. The last condition we impose is the orthogonality condition:

$$|J, 0; 1, 0\rangle = |J, 1; 1, -1\rangle \times |J, -1; 1, 1\rangle = [2BC, 0, -2BC] \quad (4.15)$$

thus, $E = 0$:

$$C_{J,0,1,0}^{J,0} = 0 \quad (4.16)$$

for all J based on symmetry.

Optical Pumping Results

We optimize the optical pumping (OP) parameters by performing optical pumping with repump for a long time, about 100 ms. The time it takes to optically pump is much shorter than this (measured $1/e$ time is 0.2 ms; see below). The purpose of this is to ensure that the magnetic field is aligned orthogonally to the k -vector of the OP laser, such that the $6S_{1/2}, F = 4, m_F = 0$ state is truly dark. If it isn't, the resonant OP light will eventually heat the atoms out of the trap. Thus, the optical pumping can be optimized by maximizing the retention rate between the two shots (taken before and after the OP phase) as a cost function.

Another method would be to perform optical pumping with repump for a short time (say, 5 ms), then disable repump for around 10 ms, then blow away the atoms in the $6S_{1/2}, F = 4$ state with a $6S_{1/2}, F = 4 \rightarrow 6P_{3/2}, F = 5$ resonant 852 nm blow-away beam (single MOT beam enabled on the +x-axis). Then optimize by minimizing the retention rate between the two shots (before and after this operation). This method works by depumping the atoms from the $6S_{1/2}, F = 4, m_F = 0$ state into the $6S_{1/2}, F = 3$ dark state, which no longer has a repump. This $F = 3$ state is dark to the blow-away, but when imaging is done after the blow-away, repump gets re-enabled and pumps the atoms back into the $F = 4$ state, allowing for imaging. Thus, if fewer atoms are present in the second shot, then the OP was successful.

Both methods were used, but the second method is a little more tricky because there is an additional reason for why no atoms were seen in the second shot—namely, the OP heated the atoms out of the trap while the repump was on. The first method allows for optimization of the OP without the risk of unintentional heating and doesn't require a blow-away beam. Of course, a blow-away beam is likely already going to be implemented for future experiments, so the second method is still useful.

Once it is optimized, we can measure the optical pumping rate of the atoms into the $F = 4, m_F = 0$ state by performing optical pumping with repump for a variable amount of time. Then, a π microwave pulse (see section 4.2.5) is applied to rotate atoms between the $F = 4, m_F = 0$ and $F = 3, m_F = 0$ states. A blow-away beam then heats the atoms in the $F = 4$ state out of the trap, allowing the $F = 3, m_F = 0$ atoms (the atoms that were successfully optically pumped) to be imaged. A short pulse of repump light is applied right before the microwave rotation to ensure that there are no atoms in the $F = 3$ state. The retention rate tends to follow a decaying $1/e$ growth curve, which can be fit to extract the optical pumping rate.

The depump rate from the $F = 4, m_F = 0$ dark state by the OP laser is measured by performing the above "second method" procedure for optimizing




Figure 4.10: Measured optical pumping (a) and depump (b) rates of the optical pumping process. The $1/e$ time of the pumping is 0.2 ms and the depumping is 80 ms . The ratio of the pump-depump rate is often used as a figure of merit to determine the quality of the optical pumping. In our case, this ratio is 400.

the OP, but with a variable amount of time where there is OP light but no repump. The number of atoms that are depumped (found in the $F = 4$ state) also follows a decaying $1/e$ growth curve, which can be fit to extract the depump rate.

Figure 4.10 shows the measured pump (a) and depump (b) rates of the optical pumping process. We find that the $1/e$ time of the pumping is 0.2 ms and the depumping is 80 ms . The ratio of the pump-depump rate is often used as a figure of merit to determine the quality of the optical pumping. In our case, this ratio is 400, which is very good. Although the pump time is 0.2 ms , the atoms are pumped for around 1 ms so that all the atoms are fully pumped into the $F = 4, m_F = 0$ state.

4.2.5 Setting up Microwave Qubit Manipulations

With atoms in the $F = 4, m_F = 0$ state, we now utilize a microwave horn to drive Rabi oscillations between the $F = 4, m_F = 0$ and $F = 3, m_F = 0$ states. This setup also enables measurement of the Zeeman splitting between the $F = 4, m_F = 0$ and $F = 3, m_F = 1$ states, which provides a way to

determine the magnitude of our magnetic field gradient.

The microwave horn is mounted just above the cell. Its large aperture and wide emission pattern make alignment straightforward. To generate the microwave signal, we use a function generator locked to a 10 MHz GPS reference to produce a frequency of $9.192631770\text{GHz} - 200\text{MHz}$. A DDS, also locked to the same GPS reference, generates a tunable 200 MHz signal. This is mixed with the 8.992631770 GHz signal to produce the desired 9.192631770 GHz microwave tone. The output passes through a band-pass filter and is amplified by a 3 W compression point amplifier before being sent to the microwave horn.

The microwave π -pulse time is measured to be $45.1\ \mu\text{s}$ and is optimized by maximizing the retention rate after a microwave π -pulse. The peak retention rate observed is 95%, which is approximately the state measurement fidelity of the 852 nm atom imaging. This π -pulse time corresponds to a Rabi frequency of 11.1 kHz.

The DDS used to generate the 200 MHz signal includes a phase offset feature, allowing us to apply a fixed, coherent phase shift to the microwave tone. However, this phase shift must be programmed prior to the experiment and cannot be changed dynamically within a cycle. A more advanced AWG designed for coherent control would enable this functionality, but for our current temperature measurements, a fixed phase is sufficient.

4.3 685 nm Spectroscopy in the 803 nm Trap

We now begin discussion of utilizing the 685 nm transition for interacting with the trapped atoms. The first step is to perform spectroscopy of the $6S_{1/2}, F = 4 \rightarrow 5D_{5/2}, F' = 6$ transition for atoms in the 803 nm trap potential. We already know approximately what the resonance frequency is based on the push spectroscopy measurements in section 4.1.2, but there are other effects that can shift the resonance frequency. These include the

Zeeman effect and the AC Stark shift from the 803 nm trap light.

4.3.1 Zeeman Effect Contribution

The DC Zeeman effect caused by non-zero magnetic fields can cause splitting of the ground and excited states of the $6S_{1/2}, F = 4 \rightarrow 5D_{5/2}, F = 6$ transition. The energy shift caused by the magnetic field is given by the Zeeman effect:

$$\Delta U_Z = \mu_B g_F m_F B \quad (4.17)$$

where μ_B is the Bohr magneton, g_F is the Landé factor of the state, m_F is the magnetic quantum number, and B is the magnetic field strength. The Landé factor for the $6S_{1/2}, F = 4$ state is $g_F = 1/4$ and for the $5D_{5/2}, F = 6$ state is $g_F = 1/2$. We can also use the Planck constant h to convert the energy shift into a frequency shift. For the states of interest, the frequency shift per magnetic field per m_f is:

$$\frac{\Delta f_{6S_{1/2}, F=4}}{h m_F B} = 350 \text{kHz/G} \quad (4.18)$$

$$\frac{\Delta f_{5D_{5/2}, F=6}}{h m_F B} = 700 \text{kHz/G} \quad (4.19)$$

This will apply a broadening to the 685 nm transition, with a maximal width occurring for the $6S_{1/2} F = 4, m_F = 4 \rightarrow 5D_{5/2} F'' = 6, m_{F''} = 6$ transition, which would have a magnitude of 2.8 MHz/G. This then is multiplied by two for the negative m_F states, giving a total broadening of 5.6 MHz/G or kHz/mG for the 685 nm transition. Thus, for a broadening equal to the linewidth, $\gamma = 2\pi \times 0.118 \text{MHz}$, the magnetic field has to be less than 0.02 G.

To minimize the residual magnetic field seen by the atoms, we scan the applied magnetic field along a single axis while maintaining a relatively large bias field along an orthogonal axis. Rather than measuring the resonant

microwave transition frequency directly (which would require repetitive spectroscopy scans for each field point), we set the microwave frequency slightly red-detuned from the resonance of the $F = 4, m_F = 0 \rightarrow F = 3, m_F = 1$ transition at the current magnetic field. We then apply a calibrated π -pulse and monitor the transfer efficiency. As the scanned magnetic field shifts the transition frequency into resonance, the population transfer is maximized. This provides an indirect but efficient way to identify the zero-crossing point along that axis.

If the scan and bias fields are truly orthogonal, this procedure can be repeated for each axis independently to locate the magnetic zero for all three directions. The resulting field setpoints can then be combined to extrapolate a magnetic field zero. However, since a nonzero bias field must be maintained to lift the degeneracy of the Zeeman levels, the true zero-field condition cannot be directly measured. Based on the sensitivity of this measurement, we estimate an upper bound of the magnetic field to be 10 mG, or roughly 60 kHz.

4.3.2 AC Stark Shift on the 685 nm Transition and Magic condition of 803 nm trap

The other main perturbation on our states is the AC Stark shift on the 685 nm transition caused by the 803 nm trap light. The AC Stark shift has multiple contributions which can also cause hyperfine splitting through the vector and tensor components of the AC Stark shift. Since our trap light is linearly polarized, the vector component of the AC Stark shift is zero [34, 80]. In this section, we discuss the potential contributions to scalar components causing a DC shift and tensor components causing hyperfine m_F level splitting.

As discussed in [34, 80], the scalar and tensor AC Stark shift is given by

$$\Delta E_s = -\frac{1}{4}|\mathcal{E}|^2 \left(\alpha_{nJF}^s - (1 - 3|\mathbf{u}_z|^2) \frac{3m^2 - F(F+1)}{2F(2F-1)} \alpha_{nJF}^T \right) \quad (4.20)$$

where \mathcal{E} is the electric field amplitude, α_{nJF}^s is the scalar polarizability, α_{nJF}^T is the tensor polarizability, \mathbf{u}_z is the component of the polarization vector along the quantization axis, F is the total angular momentum quantum number, and m_F is the magnetic quantum number. In our situation, since the electric field is linearly polarized and there is no quantization axis, we can take $\mathbf{u}_z = 1$. The scalar polarizability is independent of the magnetic quantum number, while the tensor polarizability is dependent on the magnetic quantum number to second order, meaning it causes broadening of the transition.

The value of the scalar and tensor polarizabilities can very easily be calculated using open source tools, like the Alkali.ne Rydberg Calculator (ARC) [81], as I've done here:

$$\alpha_{nJF, 6S_{1/2}}^s = -373 \text{ \AA}^3 \quad (4.21)$$

$$\alpha_{nJF, 5D_{5/2}}^s = -553 \text{ \AA}^3 \quad (4.22)$$

$$\alpha_{nJF, 6S_{1/2}}^T = 0 \text{ \AA}^3 \quad (4.23)$$

$$\alpha_{nJF, 5D_{5/2}}^T = 170 \text{ \AA}^3 \quad (4.24)$$

with cgs units in cubic Angstroms. With these values, we can calculate a total polarizability coefficient, $\alpha_{nJF}^{\text{total}}(m_F)$, (i.e. the combination of scalar and scaled tensor polarizabilities) as a function of the magnetic quantum number m_F , such that:

$$\Delta E_s = -\frac{1}{4}|\mathcal{E}|^2 \alpha_{nJF}^{\text{total}}(m_F) \quad (4.25)$$

Table 4.1 shows the calculated tensor polarizability contribution for a given $5D_{5/2}, F = 6$ state m_F level with the scaling factor at $\mathbf{u}_z = 1$, the state's

total polarizability coefficient ($\alpha_{nJF,5D_{5/2}}^{\text{total}}(\mathbf{m}_F)$), and the difference between the total polarizability of the $5D_{5/2}, F' = 6$ state and the scalar polarizability of the $6S_{1/2}, F = 4$ state (which has no tensor contribution). This difference establishes a magic condition, where the total polarizability of the $5D_{5/2}, F' = 6$ state at $\mathbf{m}_F = 6$ is almost equal to the scalar polarizability of the $6S_{1/2}, F = 4$ state at 803 nm. where each of the polarizabilities is given in cgs units of

$F' = 6, \mathbf{m}_{F'}$	$2 \times \frac{3m^2 - F(F+1)}{2F(2F-1)} \alpha_{nJF}^T$	$\alpha_{nJF,5D_{5/2}}^{\text{total}}$	$\alpha_{nJF,5D_{5/2}}^{\text{total}} - \alpha_{nJF,6S_{1/2}}^{\text{total}}$
0	-108.072	-661.899	-288.225
1	-100.353	-654.180	-280.506
2	-77.194	-631.021	-257.347
3	-38.597	-592.424	-218.750
4	15.439	-538.388	-164.714
5	84.914	-468.913	-95.239
6	169.827	-384.000	-10.326

Table 4.1: Table of calculated polarizabilities at 803 nm for the $5D_{5/2}, F' = 6$ state at multiple \mathbf{m}_F levels. The first column is the magnetic quantum number \mathbf{m}_F , the second column is the tensor polarizability contribution with the scaling factor at $\mathbf{u}_z = 1$, the third column is the total polarizability $\alpha_{nJF,5D_{5/2}}^{\text{total}}(\mathbf{m}_F)$, and the fourth column is the difference between the total polarizability of the $5D_{5/2}, F' = 6$ state and the $6S_{1/2}, F = 4$ state, which establishes a magic condition. The values are calculated using the Alkali.ne Rydberg Calculator tool [81].

cubic Angstroms. As can be seen, for $\mathbf{m}_F = 6$, there is a magic condition, where the total polarizability of the $5D_{5/2}, F = 6, \mathbf{m}_F = 6$ state is equal to the scalar polarizability of the $6S_{1/2}, F = 4$ state.

To visualize this table further, Figure 4.11 plots the AC stark shift vs. frequency for the $6S_{1/2}, F = 4$ state and the $5D_{5/2}, F = 6$ state around 800nm. In addition, it plots the total polarizabilities from $\mathbf{m}_F = 0$ to $\mathbf{m}_F = 6$ as dashed lines. As can be seen, at around 803 nm, the total coefficient of the AC Stark shift at $\mathbf{m}_F = 6$ is equal to the ground state scalar polarizability (as is also indicated in table 4.1).

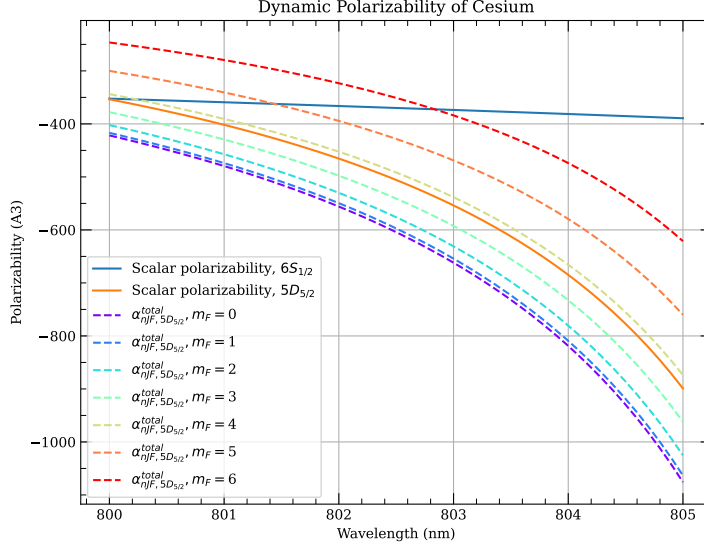


Figure 4.11: AC Stark shift for the $6S_{1/2}$, $F = 4$ state and the $5D_{5/2}$, $F = 6$ state at 803 nm. The plot shows the scalar polarizability for both states (solid lines), as well as the total polarizability, $\alpha_{nJF,5D_{5/2}}^{\text{total}}$, for each m_F level (dashed lines). The points where $\alpha_{nJF,5D_{5/2}}^{\text{total}}$ intersects with the $6S_{1/2}$ are the magic conditions on the AC-Stark Shift. At 803nm, we have the $m_F = 6$ magic condition.

It is important to note that the tensor polarizability scales with m_F^2 (see equation 4.20). As such, the difference in polarizability between the $m_F = 6$ and $m_F = 5$ states is $85A^3$, while the difference between the $m_F = 1$ and $m_F = 0$ states is only $8A^3$. This is an important feature that we see in 4.3.3 when we perform spectroscopy on the $6S_{1/2}$, $F = 4 \rightarrow 5D_{5/2}$, $F = 6$ transition in the 803 nm trap.

The sign of $\alpha_{nJF,5D_{5/2}}^{\text{total}} - \alpha_{nJF,6S_{1/2}}^{\text{total}}$ for all states is negative, meaning that the sign of the AC Stark shift is positive (there is an additional negative sign in equation 4.20 that performs a final inversion). As such, exciting an atom to the $5D_{5/2}$, $F' = 6$, $m_F < 6$ requires a higher frequency than to the

$m_F = 6$ state, and this will scale with the trap light intensity. Meanwhile, from the magic condition, the $m_F = 6$ state resonance frequency will be the lowest frequency of all the m_F states, and should be independent of the trap light intensity if the magic condition is indeed valid.

4.3.3 Single Atom - 685 nm Spectroscopy Measurements

Figure 4.12 shows the results of the 685 nm spectroscopy performed on the $6S_{1/2}, F = 4 \rightarrow 5D_{5/2}, F' = 6$ transition in the 803 nm trap. The spectroscopy is performed by scanning the frequency of the 685 nm laser while measuring the retention rate of the atoms in the trap. The measurements are performed at the near magnetic field zero discussed in 4.3.1. Since all six 685 nm beams are aligned for these measurements, these spectroscopy measurements can and do demonstrate 685 nm cooling.

Three spectroscopy scans are shown, the blue "Full power" scan is performed with the full 125 mW of 803 nm trap light, the orange "1/3 Power" is performed with 41.7 mW of 803 nm trap light and the lavender "Full Power - High Background" is performed with a different trap alignment that had a large amount of background light in the image plane, causing a large amount of trap shift. As can be seen, there is a strong correlation between the amount of 803 nm light seen by the trapped atom ("1/3 Power" < "Full Power" < "Full Power - High Background") and the broadening of the 685 nm transition.

As is most clearly seen in the difference between the "Full Power - High Background" and the "1/3 Power" scans, the unperturbed resonant frequency of the 685 nm transition is at $f \sim 148\text{MHz}$, and appears to converge at that frequency for all three scans. Furthermore, there appears to be 3 clearly distinguishable peaks in the "Full Power - High Background", before the long dip in retention rate. This is consistent with the theory presented in section 4.3.2, where the $m_F = 6$ state is resonant with the 685 nm laser, while the $m_F < 6$ are resonant at a higher frequency. Furthermore, since

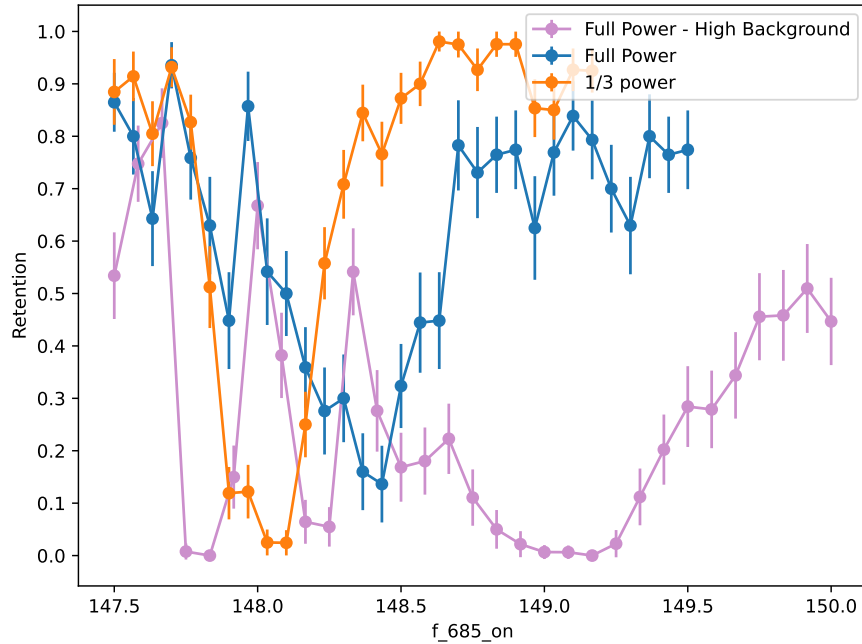


Figure 4.12: Spectroscopy results for the $6S_{1/2}, F = 4 \rightarrow 5D_{5/2}, F' = 6$ transition in the 803 nm trap. The three scans correspond to different powers of the 803 nm trap light, demonstrating the effects of the AC Stark shift on the 685 nm transition. The "Full power" scan is performed with the full power of the 803 nm trap light, the "1/3 Power" is performed with 41.7 mW of 803 nm trap light and the "Full Power - High Background" is performed with a different trap that has a larger amount of background light in the image plane, causing a large amount of trap shift. As can be seen, there is a strong correlation between the amount of 803 nm light seen by the trapped atom and the broadening of the 685 nm transition.

the peaks are not resolved at higher frequency, this is also consistent with the decreasing relative frequency between each m_F level as it approaches $m_F = 0$, since the spacing scales as m_F^2 .

These measurements were taken at multiple times throughout our experiment's life-cycle, and the actual frequency of the 803 nm was never really measured or scanned when these measurements were taken. Future work should involve further investigating the m_F level dependence of these states, as well as empirically demonstrate the magic condition of the 803 nm trap at the $m_F = 6$ state.

Regardless of any magic conditions, these measurements also demonstrate cooling with the 685 nm beams, as mentioned above. The base retention of these measurements is about 90% (about 50% for the "Full Power - High Background") because of a trap drop performed after the heating. This is done to allow an indication of cooling as well as heating. Thus, if the retention rate increases, the atoms are being cooled. Clearly, before the first dip, the atoms are cooled, which is expected for red-detuned traps. We investigate this next!

4.4 Implementing Trapped Single-Atom Cooling

As seen in section 4.3.3 (and referenced in multiple sections), the 685 nm quadrupole transition can be used to cool the atom in the trap. To optimize this 685 nm cooling, the retention rate of the atom after a trap drop is used as the cost function, with the power balancing, magnetic field bias, and detuning of the 685 nm beams as the optimization parameters. As shown in the spectroscopy scans, atoms experience cooling in the trap when the 685 nm light is red-detuned, and this can be further improved through the described optimization scheme. Due to the effectiveness of the 685 nm cooling, it is applied both before the first image and again after the first image, prior to running the experiment (e.g., the Ramsey experiment).

This section discusses temperature measurements of the atom in the trap using two methods: the first is a Ramsey T_2^* method (section 4.4.1), and the second is a trap drop method (section 4.4.2).

4.4.1 Measuring atom temperature through T_2^*

The first method we use to measure the atom temperature in the trap is by measuring the T_2^* time, which is the $1/e$ decay time of the dephasing between an atom in the $F = 4, m_F = 0$ state and an atom in the $F = 3, m_F = 0$ ground states. As presented in [82], the T_2^* time is related to the temperature of the atom in the trap through the relation:

$$T_2^* = 0.97 \frac{2\hbar}{\frac{\omega_{\text{hf}}}{\Delta_{\text{eff}}} k_B T} \quad (4.26)$$

where ω_{hf} is the hyperfine splitting frequency ($\omega_{\text{hf}} = 2\pi \times 9.192\text{GHz}$), Δ_{eff} is the effective detuning of the trap laser taking into account the D1 and D2 lines [83], and k_B is the Boltzmann constant. The effective detuning is given by:

$$\Delta_{\text{eff}} = \frac{3}{\frac{1}{f_{\text{trap}} - f_{\text{D1}}} + \frac{2}{f_{\text{trap}} - f_{\text{D2}}}} \quad (4.27)$$

where $f_{\text{trap}} = 374\text{THz}$, $f_{\text{D1}} = 335\text{THz}$, $f_{\text{D2}} = 352\text{THz}$, and so $\Delta_{\text{eff}} = 25\text{THz}$.

To measure the T_2^* time, a Ramsey gap time experiment [82] is performed, as described here. After the 685 nm cooling phase, optical pumping is performed (4.2.4), preparing the qubits in the $F = 4, m_F = 0$ ground state. This is followed by a microwave $\pi/2$ pulse that rotates the qubits into a superposition of the $F = 4, m_F = 0$ and $F = 3, m_F = 0$ states, $|\psi\rangle = \frac{1}{\sqrt{2}}(|F = 4, m_F = 0\rangle + |F = 3, m_F = 0\rangle)$. This state then evolves for a time t , during which the time evolution causes a phase between the two states, subjected simply to the hfs splitting, $\omega_{\text{hf}} = 2\pi \times 9.192\text{GHz}$. Thus, the state at time t is $|\psi(t)\rangle = \frac{1}{\sqrt{2}}(|F = 4, m_F = 0\rangle + e^{i\omega_{\text{hf}}t}|F = 3, m_F = 0\rangle)$.

However, when considering the perturbation caused by the trap potential, the AC Stark shift (scalar in this case since it is a ground state) changes the relative energy difference between the two states, since the $F = 4, m_F = 0$ state has less detuning from the D1/D2 lines than the $F = 3, m_F = 0$ state. As such, the random motion of the qubit from its temperature causes a random phase accumulation as it traverses the trap potential, leading to a loss of phase coherence between the two states. It should be noted that a random phase accumulation can also be caused by magnetic field noise causing decoherence through a Zeeman effect perturbation, meaning that these measurements give us an upper limit on the temperature of the atom in the trap.

As described in [82], the dephasing rate causes a decay in the amplitude of the Ramsey oscillations as the gap time between a pair of $\pi/2$ pulses is increased. As such, we perform a series of Ramsey experiments with increasing gap times, then fit and measure the $1/e$ decay time of the amplitude of the Ramsey oscillations. The fit Ramsey oscillations are found in Figure 4.13. Each Ramsey oscillation is fit to a cosine function with an amplitude and phase, $\frac{1}{2}(1 + A_t \cos(\phi + \phi_t))$. The amplitude A_t is then fit to a decaying exponential function, $A_t = A_0 e^{-t/T_2^*}$ in figure 4.14. The T_2^* time is then extracted from the fit, which is found to be $T_2^* = 7.6(3)\text{ms}$.

With the T_2^* time measured, we can use equation 4.26 to extract the temperature of the atom in the trap. Plugging in the values, we find a temperature of $T = 5.4\mu\text{K}$.

It is important to note that this temperature calculation is both an upper limit on the temperature, since it is unable to distinguish between the dephasing caused by a random walk in the trap perturbation (temperature) and the dephasing caused by the magnetic field noise. Additionally, optical pumping had to be performed before the Ramsey experiment, which also causes heating of the atoms. To demonstrate this, we can perform trap drop temperature measurements, in section 4.4.2.

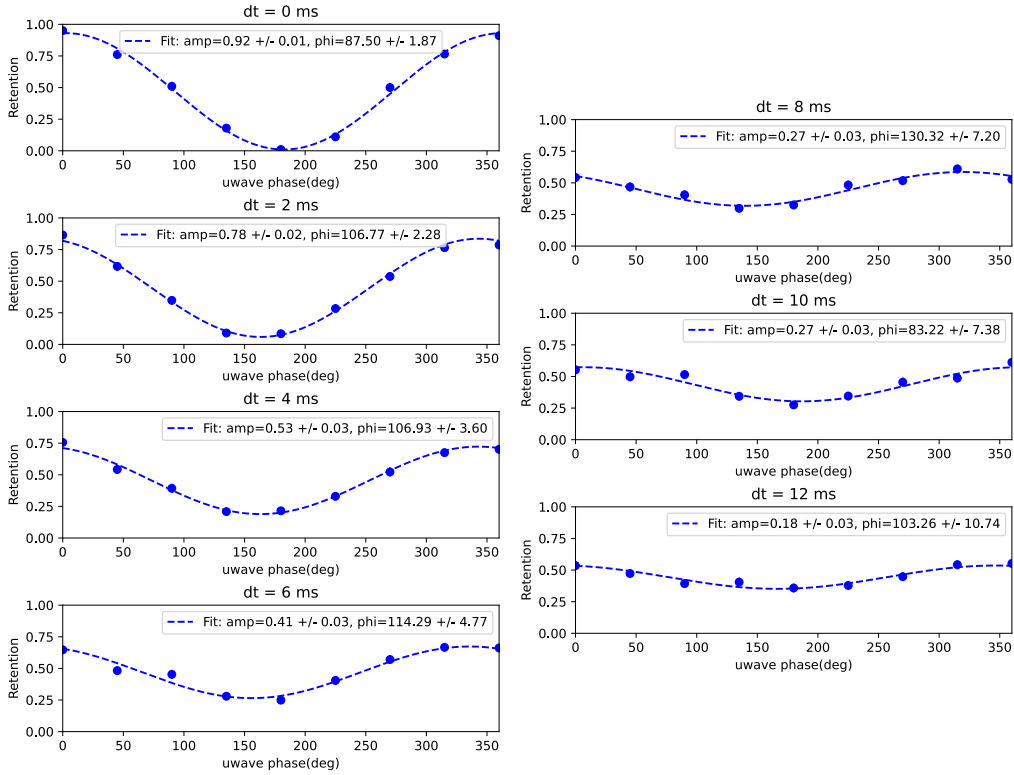


Figure 4.13: Fit Ramsey oscillations for varying gap times between $\pi/2$ pulses. The oscillations are fit to a cosine function with an amplitude and phase, $\frac{1}{2}(1 + A_t \cos(\phi + \phi_t))$. This is then used to extract the T_2^* time by fitting the amplitude A_t to a decaying exponential function in figure 4.14.

4.4.2 Performing trap drop temperature measurements

In this section, we discuss the second method used to measure the atom temperature, the trap drop and recapture measurements [74, 84]. These measurements can be compared to a Monte-Carlo simulation of the trap drop and recapture process to extract the atom temperature. The trap drop and recapture process is performed by dropping the trap light for a short time, then recapturing the atoms by turning the trap light back on. The imaging retention rate is then compared to the Monte-Carlo simulation to extract the atom temperature.

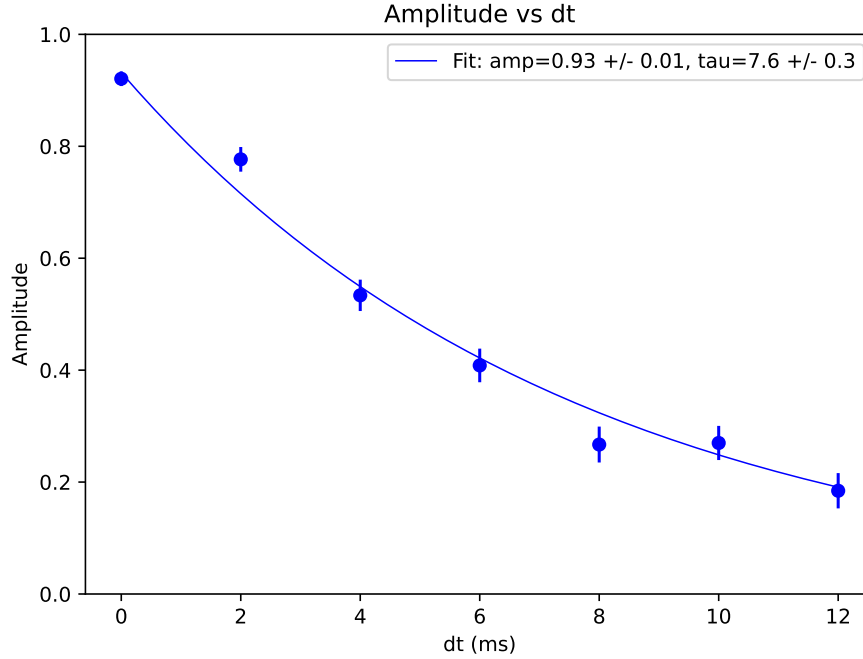


Figure 4.14: Decay of the Ramsey oscillation amplitude A_t as a function of the gap time between $\pi/2$ pulses. The data is fit to a decaying exponential function, $A_t = A_0 e^{-t/T_2^*}$, allowing for the extraction of $T_2^* = 7.6(3)\text{ms}$.

This technique also allows us to measure the temperature of the atoms before and after optical pumping. Since we already know the temperature of the atoms after optical pumping, we can calibrate the Monte-Carlo simulation with this temperature.

To perform the Monte-Carlo simulation, we model the trap potential as a harmonic potential, and iterate through 1000 particles with a random initial position and velocity based on the Boltzmann distribution at a given temperature [84]. We use the trap frequencies measured in section 4.2.3 to describe the trap potential. After allowing these particles to evolve, the atoms are determined to be lost or recapture-able based on their distance from the center of the trap. From the trap characterization in section 3.1.3,

the harmonic potential "hole" intersects with the gaussian beam at around 1.7 μm from the center of the trap. Thus, if a given particle has traveled a distance greater than 1.7 μm from the center of the trap, it is considered lost. The retention rate is then calculated as the ratio of recaptured particles to total particles for a given time step.

With the parameter of 1.7 μm max distance, the simulations led to results that were hotter than the measured temperature from T_2^* measurements ($\sim 8\mu\text{K}$). However, if the maximum distance that the particle is allowed to travel is 1.4 μm , then the 5.5 μK simulation matches with the T_2^* measurements, and can give us a measurement for the temperature of the atoms in the trap before the optical pumping. With this parameter, the temperature of the atoms in the trap before optical pumping is measured to be around $T = 3.5\mu\text{K}$. Figure 4.15 shows the measured retention rate as a function of the trap drop time, as well as the Monte Carlo simulation results for different temperatures.

Saying that the maximum distance is 1.4 μm instead of 1.7 μm is a questionable assumption and may not be entirely valid. However, the estimated temperature it yields for the atoms in the trap before optical pumping is about 2 μK lower, which is consistent with other experiments involving optical pumping. The rationale for using this calibration is that when the trap is turned back on, a significant amount of potential energy is suddenly added to the atoms. This leads to subsequent dynamics in which atoms traverse the trap potential and may escape due to imperfections in the trap. While we have confirmed the harmonicity of the trap potential at low energy (by measuring the trap frequency), the potential deviates from harmonic as it approaches the trap edges. By calibrating the maximum allowable distance an atom can travel before being considered "lost" (using a known atom temperature of 5.4 μK), we can extrapolate an approximate temperature for the colder atoms prior to optical pumping. Ultimately, we are interested in determining the temperature difference of the atoms before




Figure 4.15: Measured retention rate as a function of trap drop time for 685 nm cooling pre- and post-optical pumping, compared with Monte Carlo simulation results for different temperatures. We find that the temperature of the atoms in the trap before optical pumping is around $3.5\mu\text{K}$ by calibrating the simulation with the T_2^* measurement.

and after optical pumping, and this calibration enables that comparison. Further investigation into the validity of this assumption is warranted.

This $3.5\mu\text{K}$ atom temperature is found with a 685 nm detuning of -2.8Γ , which had an intensity of $I = 1.8 \text{ W/cm}^2$ and a saturation parameter of $s = 2$.

4.5 Imaging and Quantum State Readout

Here we present the results of imaging and non-destructive quantum state readout of the Cs qubits using the $6S_{1/2}, F = 4 \rightarrow 5D_{5/2}, F'' = 6 \rightarrow 6P_{3/2}, F' = 5 \rightarrow 6S_{1/2}, F = 4$ cycling transition. As discussed in section 2.4, the depump rate from the $6S_{1/2}, F = 4$ cycling state to the $6S_{1/2}, F = 3$ dark state is very low compared to the cycling rate, allowing for non-destructive readout of the atomic quantum state. This implementation of imaging is background-free and allows atoms to be continuously cooled during imaging, since all six axes can be used to illuminate the atom.

In “destructive” standard neutral Cs qubit experiments [10], the $F = 4$ vs $F = 3$ state is measured by applying a blow-away beam that removes atoms in the $F = 4$ state from the trap, followed by imaging the remaining $F = 3$ atoms with a repump beam. However, this method results in atoms in the $F = 4$ state being lost from the trap. Other non-destructive measurement schemes using 852 nm imaging light involve shelving atoms and using single-axis imaging fields, but these approaches typically lead to significant atom heating [23, 85, 86].

The primary drawback of the quadrupole interaction is the narrow linewidth of the transition, which necessitates a longer exposure time for imaging. However, a theoretical technique involving stimulated emission to quench the $5D_{5/2}, F'' = 6 \rightarrow 6P_{3/2}, F' = 5$ decay has been proposed as a way to increase the photon scattering rate [36]. This technique is not implemented in this work, but presents a promising direction for improving imaging rates in future experiments.

4.5.1 685 nm Qubit Imaging

The 685 nm imaging is performed very similarly to the 685 nm cooling, except with slightly less detuning. The experimental cycle involves loading the atoms into the 803 nm trap followed by 852 nm PGC and LAC, as

has been described in 4.2. The 685 nm cooling and imaging is performed, involving a 20ms cooling phase followed by a 200ms imaging phase. The magnetic fields are held at the same values for both 685 nm imaging and cooling phases, but the detuning of the 685 nm beams is -2.54Γ during imaging and -2.8Γ during cooling.

To measure the imaging fidelity, we assess two key metrics: the classification fidelity, which quantifies how reliably we can determine whether a qubit is in the $F = 4$ state, and the survival rate, which measures whether the qubit remains in the $F = 4$ state after imaging. The experiment consists of a sequence of ten 685 nm cooling-imaging cycles. Each cycle includes a 20 ms cooling period followed by image acquisition. The cooling phase serves two purposes: it allows the EMCCD camera to complete image readout and data transfer to the computer, and it resets the system for the next measurement cycle. During the cooling phase, we red-detune the 685 nm light by an additional -30kHz . This decision is based on earlier measurements showing that the atom temperature is minimized at this detuning, which helps preserve atom confinement and improves overall fidelity.

The intensity of the imaging light summed across all 6 beams is $I = 1.8 \text{ W/cm}^2$ and the saturation intensity of the $6S_{1/2}, F = 4 \rightarrow 5D_{5/2}, F'' = 6$ transition is $I_{\text{sat}} = 0.879 \text{ W/cm}^2$, giving a saturation parameter of $s = 2$. Based on the scattering rate equation, we estimate about 6000 photons scattered by the atom during the 200ms imaging phase, and with the 1.2% photoelectron collection efficiency measured in section 4.2.2, we expect to see around 70 photoelectrons of separation between the dark and bright modes in the EMCCD camera.

Figure 4.16 shows the average shot of the atom using 685 nm imaging. To determine the photoelectron counts, a 5×5 region of interest (ROI) is used to measure the photons collected from the atom, which is drawn in the figure. This atom is abnormally large, and should ideally be composed of a single pixel. Furthermore, the atom is clearly aberrated with a seeming astigmatic




Figure 4.16: Average shot of the atom using 685 nm imaging. The 5x5 ROI used to measure the pixel brightness is with a rectangle. This atom is very clearly aberrated and much larger than ideal.

shape, which is unfortunately anticipated, as discussed in 3.1.4. Regardless, since the image is background free (for the most part, as discussed in 4.5.1), the pixel brightness can be used to determine the number of photoelectrons seen in the atom and the size of the ROI can be made larger without worrying about including much background light. On top of this, the camera has a 2x2 binning mode, meaning that the atom signal is spread across a total of 100 EMCCD pixels (10x10 pixels). This should be improved in future work through the utilization of a much improved imaging train.

Regardless of the size of the atom, we can analyze the photoelectron counts and analyze the fidelities of this imaging approach. Figure 4.17 shows the histogram of the 685 nm imaging data, with each of the 9 shots combined into a single histogram. The 9 shot sequence was repeated 5000 times. As can be seen, there are two distinct modes in the histogram, indicating the presence of a single atom in the $F = 4$ or no atom. As is mentioned before, there is no repump light enabled at any point in this sequence. A bi-modal fit is performed on the histogram to determine the ideal threshold point, based on maximizing the classification fidelity, which is discussed in detail in section 4.5.1. The figure presents the bimodal fit as a black dashed line, with the two modes of the fit as blue and green dashed lines. The vertical red line is the threshold determined by the bimodal fit for all the images.

To see this data in a different perspective that includes information about retention, figures 4.18 and 4.19 show scatter plots of the imaging data, with each point representing a pair of counts from two different shots. Figure 4.18 shows the scatter plot of the first shot counts vs. the second shot counts, while figure 4.19 shows the scatter plots of the first shot counts vs. each of the counts of the 8 shots after the first shot. These scatter plots are similar to the scatter plots shown in figure 4.8, but here with the 685 nm imaging. The horizontal and vertical dashed lines are the threshold determined by the bimodal fit. The plots also include labels which indicate the loading rate of the first shot, and the ratio of atoms retained in the given shot with respect to the first shot.

We can further analyze this retention data by comparing the retention rate as we continue to image the same atom. Figure 4.20 plots the retention rates of the data in figure 4.19 vs the shot number. The blue curve presents the retention rate of the atoms seen in the given shot vs. the first shot, while the orange curve presents the retention rate of the atoms seen in the given shot vs. the previous shot. As can be seen, the orange curve has a very constant retention rate, with an average of 99.1(1)% across all shots.




Figure 4.17: Logarithmic histogram of the 685 nm imaging data, with each shot of the 9 images combined into a single histogram, for a total of 45k points. Histogram is fit to the bi-modal distribution described in section 4.5.1. The black dashed line is the result of the bimodal fit, and the blue and green dashed lines are the "no atom" and "atom" modes of the fit, separated. The vertical red line is the location of the threshold determined by the bimodal fit, which is used to classify the qubit state.

Furthermore, since this retention rate is not decreasing as more shots are taken, the atoms are not being heated during each subsequent shot. Analysis of the cause of atom loss through this imaging process is discussed in section 4.5.2.

Characterizing 685 nm Imaging Classification Fidelity

As described in 4.5.1, a threshold needs to be placed to determine if a given site in a shot either contains an atom or not. This section discusses the method used to pick the ideal threshold through fitting the histogram to a bimodal distribution and determines the resulting classification fidelity,




Figure 4.18: Scatter plot of the 685 nm imaging data, with each point representing a pair of counts from the first (x-axis) and second (y-axis) shots. The numerically determined threshold (see 4.5.1) is shown as the dashed vertical and horizontal lines. This figure is of the same style as figure 4.8.

which follows the method described in [12].

In this work, the bimodal distribution is composed of two separate distributions which are summed to make the overall distribution. For the no-atom, or "dark" shot, the distribution used is an exponentially modified Gaussian distribution. For the atom, or "bright" shot, the distribution used is a skew-normal distribution. These distributions are weighted by the relative number of atoms in each distribution before being summed together.




Figure 4.19: Scatter plots of the 685 nm imaging data, with each frame presenting an individual shot of the 8 shots after the first shot. The x-axis of the scatter is the number of photoelectrons seen in the first shot and the y-axis is the number of photoelectrons seen in the nth shot.

They are:

$$P_{\text{dark}}(x) = \frac{1}{2K\sigma_{\text{dark}}} e^{\left(\frac{1}{2K^2} - \frac{(x-\mu_{\text{dark}})}{K\sigma_{\text{dark}}}\right)} \operatorname{erfc} \left(-\frac{1}{\sqrt{2}} \left(\frac{x - \mu_{\text{dark}}}{\sigma_{\text{dark}}} - \frac{1}{K} \right) \right) \quad (4.28)$$

$$P_{\text{bright}}(x) = \frac{1}{\sqrt{2\pi}\sigma_{\text{bright}}} e^{-(x-\mu_{\text{bright}})^2/(2\sigma_{\text{bright}}^2)} \left(1 + \operatorname{erf} \left(\frac{\alpha(x - \mu_{\text{bright}})}{\sqrt{2}\sigma_{\text{bright}}} \right) \right) \quad (4.29)$$

$$P(x) = A_{\text{dark}}P_{\text{dark}}(x) + (1 - A_{\text{dark}})P_{\text{bright}}(x) \quad (4.30)$$

where K is the exponential modification factor for the dark distribution, σ_{dark} and σ_{bright} are the standard deviations of the dark and bright distributions, μ_{dark} and μ_{bright} are the means of the dark and bright distributions, α is the skewness of the bright distribution, and A_{dark} is the relative number of atoms in the dark vs. bright distributions. The results of the fit (black dashed) and the dark (blue) and bright (green) modes are shown in the histogram figure, 4.17.




Figure 4.20: Retention rate of the 685 nm imaging data after each shot. The blue curve presents the total number of atoms retained after each shot with respect to the first shot, while the orange curve presents the number of atoms retained for each shot with respect to the number of atoms in the previous shot. As can be seen, the orange shot has a very constant retention rate of 99.1(1)% across all shots. This rate is not increasing either, which indicates that the atoms are not being heated during the repetitive imaging process.

Now that the distribution is fit, the threshold can be placed by determining the number of photoelectron counts where the minimum number of classification errors between the two modes occurs. For a dark shot, the classification error occurs when a shot's photoelectron count, \mathbf{n} , is classified as an atom when it is actually a dark shot. As such, we can calculate the error of this as:

$$E_{c, \text{dark}}(\mathbf{n}) = A_{\text{dark}} \int_{\mathbf{n}}^{\infty} dx P_{\text{dark}}(x) \quad (4.31)$$

Since the $P_{\text{dark}}(x)$ function is normalized, this integral gives the probability of misclassification a dark shot as a bright shot if the threshold is placed at

\mathbf{n} . Similarly for the bright shot case, where the error instead occurs when a bright shot is classified as a dark shot:

$$E_{c, \text{ bright}}(\mathbf{n}) = (1 - A_{\text{dark}}) \int_0^{\mathbf{n}} dx P_{\text{bright}}(x) \quad (4.32)$$

This integral gives the probability of misclassification a bright shot as a dark shot if the threshold is placed at \mathbf{n} . The total classification error is then the sum of these two errors:

$$E_c(\mathbf{n}) = A_{\text{dark}} \int_{\mathbf{n}}^{\infty} dx P_{\text{dark}}(x) + (1 - A_{\text{dark}}) \int_0^{\mathbf{n}} dx P_{\text{bright}}(x) \quad (4.33)$$

This classification error will have a minimum at the optimal threshold, \mathbf{n}_{opt} , and the value of this minimum is the classification error rate. The classification fidelity is $1 - E_c(\mathbf{n})$, which is the probability of correctly classifying a shot as either dark or bright given a threshold of \mathbf{n} .

To perform these calculations, the SciPy.stats library [87] has many built-in functions for working with these exact probability distributions, which can also be used to numerically calculate the integrals in equations 4.31, 4.32, and 4.33. The probability density function (PDF) and cumulative distribution function (CDF) of the distributions are already programmed in to the library allowing for quick and easy calculations.

Figure 4.21 (and a zoomed in version, figure 4.22) show the classification fidelity as a function of the threshold point used for our imaging data fits. The solid blue line is the classification fidelity for the best-fit bimodal distribution, while the shaded blue region is the 1-sigma uncertainty region of the classification fidelity. The dashed red line is the threshold that maximizes the classification fidelity. To compute the uncertainty region, bootstrapping was performed, where 1000 samples of the of the bimodal fit parameters were randomly selected based on their uncertainties, and the classification fidelity curve was computed for each sample. Then, at each threshold point, the average and the 1-sigma uncertainty were computed.




Figure 4.21: Classification fidelity of the 685 nm imaging data as a function of the threshold point used for the bimodal fit. The solid blue line is the classification fidelity for the best-fit bimodal distribution, while the shaded blue region is the 1-sigma uncertainty region of the classification fidelity. The dashed red line is the threshold that maximizes the classification fidelity, which is found to be at a photoelectron count of 35.4 photoelectrons. The classification fidelity is found to be 0.9993(4).

The classification fidelity is found to be 0.9993(4) at a photoelectron count of 35.4 photoelectrons. This value is used for the classification of the Qubit state for the retention statistics in section 4.5.1.

Background Light from Trap Array Raman Scattering

In figure 4.16, there is an obvious, yet dim, background light throughout the image. Although the 685 nm scattering is intrinsically background free, and there are physical beam shutters in place during the 685 nm experiments, there still seems to be background light. This light is coming from the 803 nm trap laser. Rather than being 803 nm light, it is indeed 852 nm light

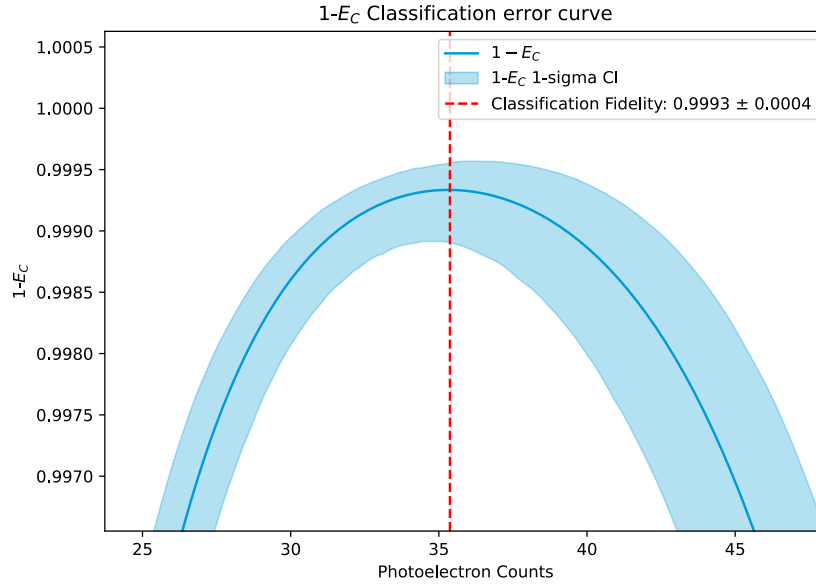


Figure 4.22: Zoomed in version of figure 4.21, allowing for visualization of the classification fidelity and uncertainty.

being generated by stimulated Raman scattering in a focused plane of the trap array, likely in either the glass cell wall or the objective.

Without any 685 nm or 852 nm light entering the optical setup, a 200ms exposure time image is taken with the EMCCD camera and the number of photoelectrons and their distribution is counted both for with and without the 803 nm trap light enabled. In the 5x5 ROI, the standard deviation of the photoelectron counts is 4 photoelectrons without the trap light and 10 photoelectrons with the trap light on. Since these scale in quadrature, the standard deviation of the photoelectrons caused by the trap light is 9 photoelectrons. Meanwhile, the standard deviation of the photoelectron counts for the bright mode was about 20 photoelectrons.

Mitigating the background light caused by the 803 nm trap light can be achieved by applying an iris in the image plane to block the majority of the 803 nm trap light, since the majority of the trapping potential is not the large gaussian, but instead the parabolic potential at the center. The light

in the image plane outside this potential does not contribute to trapping, and thus can be clipped out. This is a future direction for improving the 685 nm imaging results, allowing the imaging to be faster while maintaining the high classification fidelity.

4.5.2 685 nm Raman Depump Rate

In this section, we study the sources of $F = 4$ atom loss during the 685 nm imaging that lead to the 0.991 retention rate measured in 4.5.1. Since we know that atoms are not being heated, the primary sources of atom loss are the Raman depumping rate from the $F = 4$ state to the $F = 3$ dark state, and from the vacuum limited lifetime of the atom trap, independent of the imaging process.

To measure the contributions of these two sources of loss, we perform a 685 nm pulse with the imaging detuning and parameters, scanning the duration of the pulse from 400ms to 5.4s. Since the atoms are scattering light through this entire pulse, we take an image for the first 200ms of the pulse and the last 200ms of the pulse (resulting in the minimum time of 400ms), allowing us to measure the number of atoms lost from the $F = 4$ state during the pulse. This is our total loss rate and it needs to be distinguished between the contributions from the Raman depumping rate vs. the vacuum limited lifetime loss rate. We know that the atoms that were depumped are now in the $F = 3$ state, which we can measure by blowing away the atoms in the $F = 4$ state with a blow-away beam, applying a short 852 nm repump pulse to pump the atoms from $F = 3$ into $F = 4$, and perform a final image of the atoms with the 685 nm light.

Figure 4.23 presents a diagram of this sequence as well as the results of these measurements. The blue curve presents the total loss ratio of the atoms, measured at the first shot after the pulse. The red curve presents the ratio of atoms depumped into the $F = 3$ state, based on the ratio of atoms found again in the second shot after the pulse. The orange curve presents

the difference between the blue and red curves, which is then the ratio of atoms lost due to the vacuum limited lifetime. Each of these curves are fit to a $1/e$ lifetime. We find that the total loss lifetime is $20(1)\text{s}$, the Raman depump lifetime is $42(4)\text{s}$, and the vacuum limited lifetime is $43(5)\text{s}$.

4.5.3 Paths to Improved Readout with 685 nm Imaging

In 200ms , the loss from depumping is $e^{-200\text{ms}/42\text{s}} \approx 0.42\%$ and the loss from the vacuum lifetime is $e^{-200\text{ms}/43\text{s}} \approx 0.46\%$. Together, these loss channels account for the 0.991% retention rate seen in figure 4.20. If more photons could be collected from the atom or the camera and trap light noise could be suppressed, the exposure time necessary for the imaging could be reduced. Both of these loss effects scale linearly with the exposure time (to first order), so reducing the exposure time by a factor of 2 would also decrease the loss rate by a factor of 2.

With modern camera systems, the number of photoelectrons required for a reliable state measurement is around 15 [88]. In our experiment, we are collecting around 75 photoelectrons per atom. With an improved imaging scheme, optics, and a modern camera, the exposure time could be reduced to around 40ms , which would lead to a loss rate of about 0.1% from both depump and vacuum lifetime, for a retention fidelity of around 99.8% .

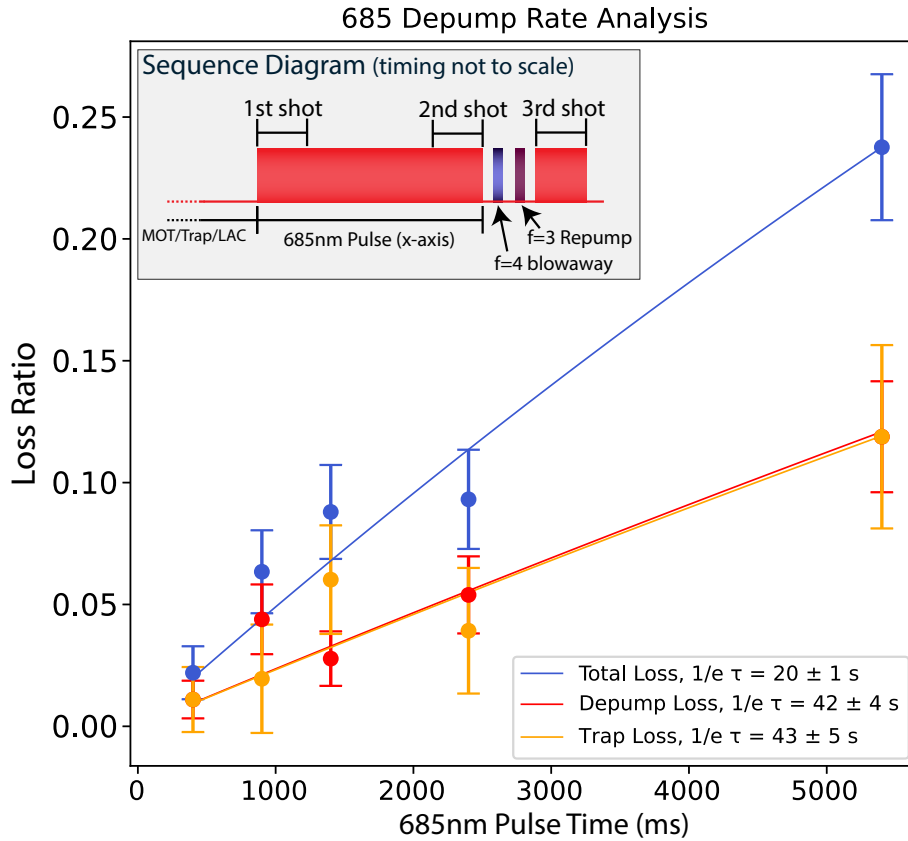


Figure 4.23: Diagram of the 685 nm imaging sequence used to measure the Raman depump rate and the vacuum limited lifetime of the atoms in the trap. The blue curve presents the total loss ratio of the atoms, measured at the first shot after the pulse. The red curve presents the ratio of atoms depumped into the $F = 3$ state, based on the ratio of atoms found again in the second shot after the pulse. The orange curve presents the difference between the blue and red curves, which is the ratio of atoms lost due to the vacuum limited lifetime. From fitting these curves, the total loss $1/e$ lifetime is $20(1)\text{s}$, the Raman depump lifetime is $42(4)\text{s}$, and the vacuum limited lifetime is $43(5)\text{s}$.

5 DEVELOPING LOW NOISE LASERS FOR RYDBERG GATES

5.1 The 459 nm Rydberg Laser

5.1.1 Architecture of 459 nm Laser System

The 459 nm Rydberg laser was implemented using an MSquared SolsTiS Ti:Sa bow-tie cavity lasing at 918 nm and an Equinox pump laser, which is then frequency doubled using an LBO crystal [89] in a bow-tie cavity [90] down to 459 nm. This architecture involves an 18W 532nm pump laser, which pumps the Ti:Sa crystal in the SolsTiS. The bow-tie cavity is tuned with multiple frequency selection schemes including a Birefringent Filter (BRF), an etalon, an optical isolator, a fast+slow pair of piezo-electric actuators (dubbed "piezo" henceforth), and an electro-optic modulator (EOM).

To narrow and lock the 459 nm light, the scheme involves locking the




Figure 5.1: Image of 459 nm laser system, locked to the SHG cavity.

918 nm laser, before frequency doubling, through the utilization of a high finesse Ultra-Low Expansion (ULE) crystal cavity. The cavity is placed in a ultra-high-vacuum (UHV) can, with an ion pump for maintaining vacuum. Although UHV is overkill to mitigate thermal conduction from the crystal to the vacuum can, ion pumps do not impart vibrations and are thus very ideal for this application. In addition to this, the vacuum can is temperature controlled through a thermo-electric coupler (TEC) to maintain a constant temperature. This effort of maintaining an isolated ULE cavity is to be able to use it as a frequency reference instead of using a saturated absorption line or similar.

The FSR of the ULE cavity was measured to be 499.811371 MHz at 960 nm. This means we need the capability to shift the frequency of the light by up to 500 MHz in order to reach any desired lock point. To this end, a double-pass acousto-optic modulator (AOM) setup is used to shift the frequency of the light before entering the cavity. It should be noted that this requires an AOM with a bandwidth of 250 MHz, which in turn requires an even higher center frequency. We use a 650 MHz, 225 MHz BW AOM for this, which has a peak single-pass diffraction efficiency of 50%. Applying this frequency shift to the light going to the cavity instead of the light going to the experimental apparatus allows for simple broad tuning of the light frequency by large amounts without inducing a large cost to the high intensity required for second-harmonic generation (SHG) in the LBO crystal. As will be discussed later, an additional double-pass AOM setup is added for dynamic tuning of the laser frequency after the frequency doubling. But using the 650 MHz double-pass AOM as a dynamic shifter will not work since the laser lock would effectively become responsible for maintaining the frequency shift. Although it easily has the dynamic range capabilities for a small frequency change, it still is a system with finite bandwidth.

The SHG bow-tie cavity with an LBO crystal inside has a finesse of 140 (when the LBO temperature is incorrect, resulting in no SHG). The beam is

focused through the LBO crystal using parabolic mirrors. With 500 mW of 918 nm, we get about 100 mW of 459 nm. The SHG cavity has a piezo on one of the flat mirrors. A Hänsch-Couillaud (HC) locking scheme [91] is utilized for locking the cavity length to the laser (since the laser is locked to the ULE).

After SHG, a "noise-eater" is implemented using an AOM. The light that passes on to the experiment is the 0th order of the AOM. A photodiode (PD) measures the intensity of a sample of that light after the noise eater, and an error signal is generated where the goal is to set the intensity of the laser power to a constant value that is less than the total power transmitted through the AOM. The goal is to discard the excess power by modulating the amplitude of the RF going to the AOM to remove the correct amount of power from the laser to make the intensity equal to a desired setpoint. Thus, the error signal is the intensity measured by the PD subtracted by some constant setpoint intensity. This mechanism allows for random intensity noise (RIN) stabilization of the Rydberg lasers. Figure 5.2 demonstrates the effectiveness of this system, with the noise eater on vs. off.

5.1.2 Locking Scheme for the 459 nm Rydberg Laser

To actually lock the laser to the ULE cavity, a Pound-Drever-Hall (PDH) technique [56] is used to create an error signal with an 80 MHz modulation. There are three feedback channels (and servo loops) for this MSquared system: a slow piezo, which has a high dynamic range (> 1 GHz) but a very low bandwidth (< 50 Hz); a fast piezo, which has a small dynamic range (100 MHz) and a bandwidth of 100 kHz; and an intra-cavity EOM, which has a very low dynamic range but a very high bandwidth (10 MHz). All three of these feedback channels work in unison to change the effective cavity length of the SolsTiS bow-tie cavity: the piezos change the light's path, and the EOM changes the index of refraction, thereby changing the cavity length.




Figure 5.2: RIN comparison with the noise eater enabled and disabled for the 459 nm laser. Units are in Fractional Noise Spectral Density, $1/\sqrt{\text{Hz}}$

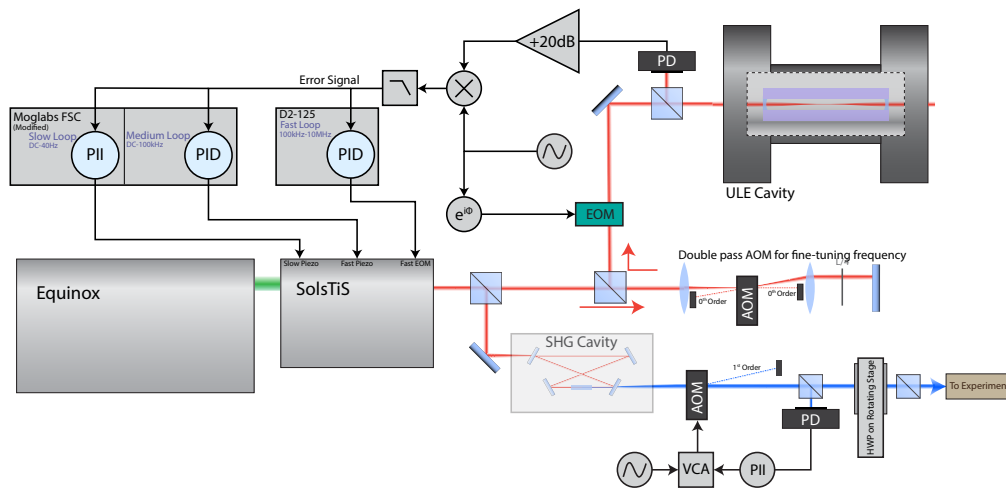


Figure 5.3: 918 Locking Scheme Functional Diagram.

Since the free-running linewidth of the SolsTiS is < 100 kHz, there is very little noise at the 10 MHz level, so only a small feedback signal needs to be applied through the EOM. In addition, an AC-coupled error signal suffices. The servo utilized in the feedback loop for the high-frequency noise is a Vescent D2-125 lockbox. This lockbox has an output voltage span of ± 10 V and a high bandwidth. This loop involves a PI (no derivative), with a PI corner at 200 kHz. This allows for increased gain in the 10 kHz–100 kHz band, reducing the responsibility of the fast piezo in its most fragile range. A proportional gain is also applied for feedback above 200 kHz. Since the signal is AC-coupled, there is a gain limit for the integrator.

For the two piezos, a modified Moglabs FSC servo controller (which has a fast and a slow channel) is used to perform feedback. This feedback system is designed such that the laser can be locked with just the fast piezo and the EOM. However, due to the limited dynamic range of the fast piezo, the DC drifts of the system (thermal, pressure, humidity) will quickly cause the integrators for the fast piezo to rail (on the order of 5 minutes). Since the slow loop has a very large dynamic range, it is enabled after the other loops and therefore assists the fast piezo in obtaining long-term locks. So this loop is just a proportional loop.

The modification to the Moglabs FSC is to increase the dynamic output range of the "fast" loop by replacing an internal limiting resistor. In addition, the dynamic range of the slow-loop integrator is increased as well with a simple modification.

To increase the output range on the Moglabs FSC, R91, which acts as an output voltage limiter, must be replaced with a 500 Ohm 0402 package resistor. The original value of the resistor was measured to be 2.5 kOhms. Contact Moglabs to confirm this before replacing if you want to undo the modification on the board. This upgrade will change the range from about ± 0.625 V up to about ± 2 V. It is highly recommended to label the lockboxes that these changes have been implemented, since this can damage a laser

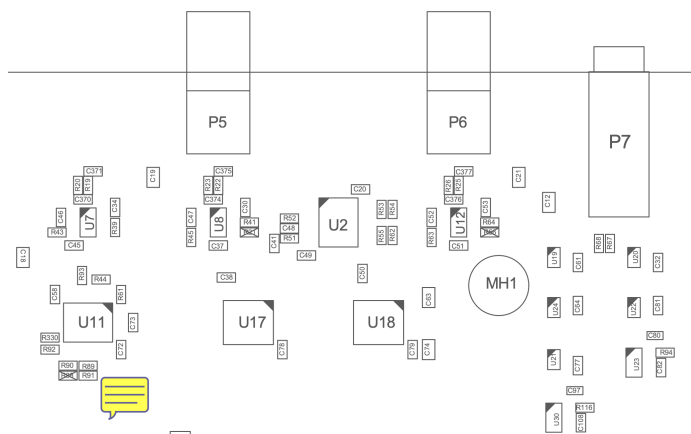


Figure 5.4: Moglabs FSC fast loop mod R91 location.

diode if plugged back into a standard Moglabs laser unit. See Figure 5.4 for R91 location.

The slow loop integrators of the Moglabs FSC have an attenuation stage before being summed with the ramp/offset signals. This is a 200:1 attenuation, which limits the range of the slow output to around ± 50 mV, while the total output range with the offset is around 200 mV to 1000 mV. To decrease this attenuation, R87 (nominally 100 kOhm) must be replaced with a smaller resistance. A 5 kOhm resistor has given an output range of ± 400 mV and should be implemented in most designs with this modification. Again, label the lockbox with the applied change. See Figure 5.5 for R87 location.

The fast piezo is implemented with a PI gain. The PI corner is placed at 50 kHz, and the proportional gain is at a low level. To ensure that the integrator for the fast loop doesn't beat the integrator for the slow loop (the slow loop needs to be responsible for the DC drifts), a gain limit is applied to the fast piezo, which corners off the upper end of the integrator.

Figure 5.3 presents a functional diagram of the locking scheme for the 459 nm Rydberg Laser.

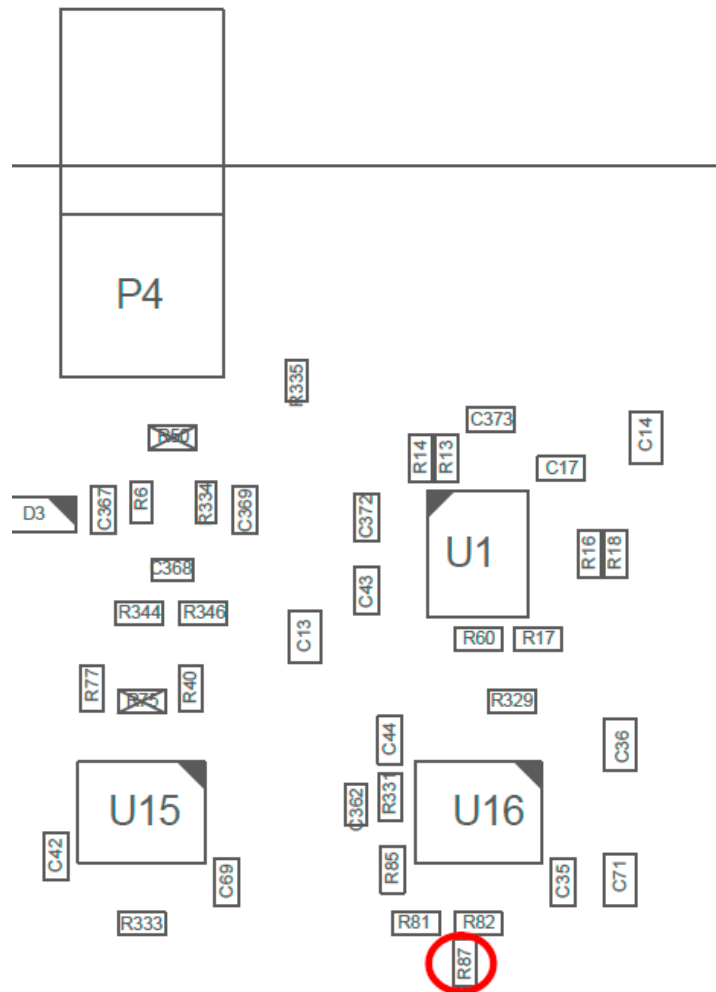


Figure 5.5: Moglabs FSC slow loop modification R87 location.

5.1.3 Measurements of the 459 nm Spectral Noise

The self-heterodyne technique [92] is used to measure the spectral noise of the 459 nm laser system. We measure the spectral noise of the 918 nm laser, since there is more power available and this is the laser that is actually being narrowed.

The 918 nm laser has a permanent self-heterodyne system set up so that we can have continuous monitoring of the noise performance of the laser. This is sent to a spectrum analyzer to constantly monitor and fit the location and amplitude of the servo bumps as well as the white noise. It then uploads this data to the Origin database server. This server then feeds information to the system monitoring system for the AQuA experiment.

The self-heterodyne setup was designed by Sam Norrell and was created using a series of fiber splitters and a fiber-based AOM. The beam is input into the system through a fiber patch cord butt-coupled to a 50-50 fiber non-polarizing beam splitter. One path is sent to a 10 km delay fiber, and the other path is sent to the fiber AOM, which is modulated at 100 MHz and outputs the first-order beam in a fiber. This beam is then combined with the 10 km delayed fiber on a 90-10 beam splitter, where the delayed fiber is combined on the 90% port and the AOM-diffracted path is combined on the 10% port. This is because the delayed fiber exhibits loss. An issue with this design is that the polarization of the light traveling through the split paths may dynamically change due to environmental effects. This can result in a changing amplitude in the self-heterodyne signal, since polarizations need to match for a beat note to be created on the diode. An upgrade to this system would involve a way to rotate the polarization so that it can be matched between the beams.

Modeling the servo bumps

To analyze and perform the fits for monitoring the noise performance of the Rydberg lasers from the spectrum of the self-heterodyne, the analysis

described in [93] is used.

Consider a laser with a pure white frequency noise of amplitude \mathbf{h}_0 (typically measured in Hz^2/Hz):

$$\mathbf{S}_{\text{dv}} = \mathbf{h}_0 \quad (5.1)$$

The phase noise of such a noise spectrum would be:

$$\mathbf{S}_{\phi, \text{white}}(f) = \frac{\mathbf{h}_0}{f^2} \quad (5.2)$$

where the simple relation between phase and frequency noise of $\mathbf{S}_{\phi} = \mathbf{S}_{\text{dv}}/f^2$ has been used.

Similarly, we can model a servo bump to be a symmetric gaussian contribution of peak amplitude, \mathbf{h}_g , in terms of frequency noise:

$$\mathbf{S}_{\text{dv, bump}}(f) = \mathbf{h}_g \left[\exp\left(-\frac{(f - f_g)^2}{2\sigma_g^2}\right) + \exp\left(-\frac{(f + f_g)^2}{2\sigma_g^2}\right) \right] \quad (5.3)$$

and so the phase noise would be:

$$\mathbf{S}_{\phi, \text{bump}}(f) = \frac{\mathbf{h}_g}{f_g^2} \left[\exp\left(-\frac{(f - f_g)^2}{2\sigma_g^2}\right) + \exp\left(-\frac{(f + f_g)^2}{2\sigma_g^2}\right) \right] \quad (5.4)$$

To model our lasers, we can consider them to be superpositions of these noise components, where there is a pure white noise with several sets of servo bumps.

$$\mathbf{S}_{\phi, \text{total}} = \mathbf{S}_{\phi, \text{white}}(f) + \sum_j \mathbf{S}_{\phi, j}(f) \quad (5.5)$$

where $\mathbf{S}_{\phi, j}(f)$ is the j -th servo bump.

Due to the finite time-delay of a self-heterodyne system, a Discrete Fourier Transform "windowing"-like effect (see, for example, the Blackman-Harris window: [94]) occurs when measuring the spectrum with this technique, where a sinusoidal modulation appears. This modulation is caused by points

in the spectrum where the time-delay, t_d , induced has accrued as much phase as an entire period of an oscillation at a given frequency. As is fully analyzed and discussed in [93], the actual measured spectrum of a laser with pure white noise would be:

$$S_i(f) = \frac{2h_0}{f^2 + (2\pi h_0)^2} + e^{-4\pi^2 h_0 t_d} \left(\delta(f) - \frac{2h_0}{f^2 + (2\pi h_0)^2} \left[\cos(2\pi f t_d) + \frac{2\pi h_0}{f} \sin(2\pi f t_d) \right] \right) \quad (5.6)$$

And similarly for a servo bump, this would be:

$$S_i(f) \approx \delta(f) + \frac{4h_g}{f_g^2} \sin^2(\pi f t_d) \left[\exp\left(-\frac{(f - f_g)^2}{2\sigma_g^2}\right) + \exp\left(-\frac{(f + f_g)^2}{2\sigma_g^2}\right) \right] \quad (5.7)$$

Fitting the servo bumps

The first step is to normalize the power measured by the spectrum analyzer by fitting the contribution of the center (highest amplitude) component of the self-heterodyne spectrum. This has to be done each time because the power of the laser continuously changes due to polarization drifts in the fibers. The center of the spectrum is fit to

$$S_{i,\text{peak}} = \frac{s_p \sigma^4}{(f^2 + \pi^2 \sigma^2)^{5/2}} \quad (5.8)$$

which has a total power of $\frac{3\pi^2 s_p}{4}$. This factor is used to normalize the measured data.

Next, the data is fit to a summation of Eqs. 5.6 and 5.7 for the windowed white phase noise, $S_{i,\text{white}}$, and the j -th windowed servo bump phase noise,

$S_{i,\text{bump } j}$, building off of Eq. 5.5:

$$S_{i,\text{fit}} = S_{i,\text{white}}(f) + \sum_j S_{i,\text{bump } j}(f) \quad (5.9)$$

Note that the delta-function component of $S_{i,\text{white}}$ is omitted from the fit since it isn't physically measurable with a spectrum analyzer. The SciPy least squares "curvefit" algorithm is used to perform the fitting [87]. To reduce noise in the measurement, many measurements of the spectrum are averaged down for a single shot before the data is fit.

An ansatz for the number of servo bumps and their general location is used for the servo bump fitting. For each prominent servo bump, an initial guess is given to the spectrum analysis tool based on where the bumps have been known to exist. For the 918 nm laser, a successful fit with initial parameters involving three servo bumps at $f_g = 6.2$ kHz, $\sigma_g = 910$ Hz, $h_g = 72$ Hz²/Hz, $f_g = 51$ kHz, $\sigma_g = 40$ kHz, $h_g = 1.2$ Hz²/Hz, and $f_g = 145$ kHz, $\sigma_g = 5.8$ kHz, $h_g = 1.1$ Hz²/Hz is used, since this is where previous trial and error of the number of servo bumps and their locations placed them. The white noise is also initially guessed, with a similar initial guess assumption (0.5 Hz²/Hz for the 918 nm laser).

Once the accurate strengths (h_g), locations (f_g), and sizes (σ_g) of the servo bumps are fit, as well as the white noise amplitude (h_0), we extrapolate the actual phase noise, without the windowing applied by the self-heterodyne beating mechanism or dark noise of the spectrum analyzer. Equation 5.5 is used here.

The results of one step of this analysis of the laser phase noise is presented in figure 5.6. The measured data (blue) is fit to the windowed phase fit described in Eq. 5.9 (green) and the resulting actual phase noise is presented based on Eq. 5.5 (orange).




Figure 5.6: 918 nm Self-Heterodyne phase noise data, windowed fit, and extrapolated phase noise from the windowed fit. White noise: $0.5 \text{ Hz}^2/\text{Hz}$. Fit Servo Bumps at $f_g = 6.2 \text{ kHz}$, $\sigma_g = 910 \text{ Hz}$, $h_g = 72 \text{ Hz}^2/\text{Hz}$, $f_g = 51 \text{ kHz}$, $\sigma_g = 40 \text{ kHz}$, $h_g = 1.2 \text{ Hz}^2/\text{Hz}$, and $f_g = 145 \text{ kHz}$, $\sigma_g = 5.8 \text{ kHz}$, $h_g = 1.1 \text{ Hz}^2/\text{Hz}$.

5.2 The 1040 nm Rydberg Laser

The 1040 nm Rydberg laser is created in a very similar fashion, except there is no SHG. The gain of the Ti:Sa crystal at this wavelength is much lower [95], so this laser is much more susceptible to mode hops (see Figure 5.7 for the gain/emission spectrum of the Ti:Sa crystal).

In addition, there is no double-pass frequency shifter required for the laser before it is locked to the ULE cavity, since there already is one on the 459 nm laser.

The locking scheme of the 1040 nm is identical to that of the 459 nm, with the same feedback loops and lockboxes, including modifications.




Figure 5.7: Ti:Sa absorption and emission spectrum, from [95]. The emission gain at 1040 nm falls off from the peak gain at 800 nm by almost an order of magnitude.

5.2.1 RIN of the 1040 nm Rydberg Laser

Since the 1040 nm light is coming directly out of the SolsTiS and is not being frequency doubled, the intensity noise is small, even without a noise eater. We did not implement a noise eater for this system. Figure 5.8 presents the RIN for the 459 nm light (with the noise eater on) vs. the free-running 1040 nm light.

5.2.2 Measurements of the 1040 nm Spectral Noise

The 1040 nm spectral noise is also measured using a dedicated self-heterodyne system with a similar monitoring setup to the 918 nm laser. Figure 5.9 presents the 1040 nm self-heterodyne spectrum using the same fitting mechanism described.

For this laser system, there are only two servo bumps (see the figure). One of the servo bumps is at $f_g = 234$ kHz and has an amplitude of $h_g = 1.7 \times 10^3$ Hz²/Hz. This servo bump has a much bigger impact than any of the 918 nm servo bumps because it is at a frequency closer to the




Figure 5.8: RIN comparison between the 459 nm (with noise eater) and the 1040 nm Rydberg lasers. The free-running noise of the 1040 nm laser is less than that of the 459 nm laser with the noise eater.

Rabi frequency. Section 5.2.3 briefly describes the effects of this noise.

5.2.3 S_{dv} of Rydberg Lasers

With the fit and characterized phase noise, the frequency noise in Hz^2/Hz can easily be presented as well. Figure 5.10 plots the frequency noise of the two Rydberg lasers based on the fits of the windowed phase noise from the self-heterodyne measurements. Based on the analysis in [93], we can also extrapolate the error contribution of this noise for our Rydberg gates.

The white noise level for these lasers is very low compared to the amplitude of the servo bumps. The issue with servo bumps is how close they are in frequency to the Rabi frequency. The single-atom Rabi rate for two qubit Rydberg gates in the AQuA experiment is 1.7 MHz (see [10]).

Since the white noise levels are very low, the servo bump with the largest




Figure 5.9: 1040 nm self-heterodyne phase noise data, windowed fit, and extrapolated phase noise from the windowed fit. White noise: $6.6 \text{ Hz}^2/\text{Hz}$. Fit servo bumps at $f_g = 69.1 \text{ kHz}$, $\sigma_g = 40.6 \text{ kHz}$, $h_g = 11 \text{ Hz}^2/\text{Hz}$ and $f_g = 234 \text{ kHz}$, $\sigma_g = 1.3 \text{ kHz}$, $h_g = 1.7 \times 10^3 \text{ Hz}^2/\text{Hz}$.




Figure 5.10: The S_{dv} of both Rydberg lasers based on fitting white noise and servo bumps to the measured self-heterodyne spectra, following the analysis found in [93]. See Figures 5.6 and 5.9 for fit parameters.

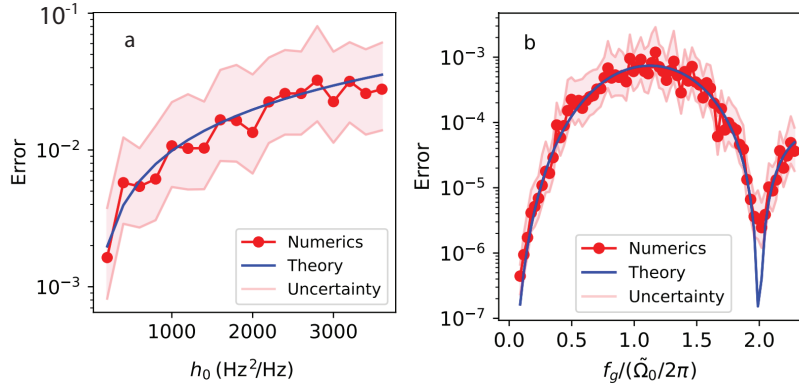


Figure 5.11: Figures taken from [93], presenting Rabi error vs. laser noise parameters for a 2π rotation: a) Rabi error vs. white phase noise, h_0 . b) Rabi error vs. servo bump frequency (f_g with amplitude of $h_g = 10^3 \text{ Hz}^2/\text{Hz}$). For both situations, the Rabi frequency used in the simulation is $\Omega_0/2\pi = 1 \text{ MHz}$. See the reference for more details.

contribution is in the 1040 nm laser at $f_g = 234 \text{ kHz}$ with an amplitude of $1.7 \times 10^3 \text{ Hz}^2/\text{Hz}$. Figure 5.11 presents the results of [93], demonstrating the effects of laser frequency noise on the atoms for a 2π Rydberg rotation.

Using the results shown in the figure, the low white noise levels and $f_g/(\Omega_0/2\pi) = 0.14$ contribute to the Rabi error by less than 10^{-2} and 10^{-5} respectively. This informs that the noise of the Rydberg lasers were not the limiting factor for the 95.5% (SPAM corrected) Bell state achieved [10].

5.3 A small collection of thoughts and experiences on laser locks

5.3.1 Tuning up a ULE cavity

To tune up the lock from a base state—and this is my general approach for all laser systems with multiple feedback loops locked to a ULE cavity—is to first set up a sweep of the laser through the cavity mode. With no feedback,

there should be a very thin, brief pulse of cavity transmission. To ensure the right gain sign is applied, attempt a simple high-gain proportional loop. If the proportional loop can extend this brief cavity pulse, even if it's noisy, then you have the right gain sign. Attempt to increase the gain so that this pulse is extended. There will likely be occasional dips, but the pulse should ultimately be extendable. This loop will likely look like an integrator with finite gain or a low-pass filter with a high proportional component. Derivatives can be useful too, but as the frequency increases, the gain also increases — at 10 dB per decade. So there must be some sort of second-order low-pass filter (or a low-pass filter and a finite-gain derivative) to cancel out this derivative term and apply enough attenuation before the first phase margin. I find that, in general, this isn't ever necessary and can be handled much more simply with just a high proportional term and a low-pass filter.

Once this is done, the next step is to tune up the "medium" loop. The gain sign can be picked in a similar way to the fast loop, where a high proportional gain is applied to extend the lock acquisition window through the ramp. This term can use an integrator, but it should have finite gain (so that the slow loop can ultimately handle all DC drifts). When the integrator (if used) is active and the ramp is disabled, this should be able to keep the laser from coming unlocked for at least a few seconds.

If the laser can last for a few minutes without the integrator railing, this is a good time to try to optimize the gain and low-pass filter parameters of both the fast and slow loops. With a self-heterodyne, you can do a great job optimizing the servo bump performance, but with just the error signal and cavity transmission, you can try to minimize the noise of the transmission.

Once the proper lock performance is acquired, the slow loop can finally be implemented. The reason the slow loop is applied last is that it shouldn't affect the higher frequency dynamics of the laser spectrum. It should just be a signal that is driving the integrator (or near-DC component) of the "medium" loop to zero. This keeps things simple. A very low-frequency low-

pass filter can even be applied in this stage. An interesting application of this loop can be found in the 685 nm laser lock discussion earlier (Section 3.2.4).

If a laser comes unlocked shortly after an integrator is applied, the operator could look into the reason why such a situation is happening. For example, study the dynamics of the output signals (Does the integrator rail? Does it oscillate? etc.). Generally, turning down the gain can also lead to insight.

Ultimately, all laser systems are different and require different forms of care, love, and attention. This guide is highly pointed toward situations where there are three feedback paths, with a clear "fast", "medium", "slow" architecture. The section on the 685 nm laser lock discusses a much more interesting laser lock, which actually utilizes a derivative loop (Section 3.2.4).

5.3.2 Lasers That Stay Locked

Lasers tend to come unlocked for two primary reasons: adiabatic and spontaneous environmental effects. Adiabatic drifts tend to surround temperature, pressure, and humidity fluctuations, causing the actuators to need to accumulate a large amount of actuation. Or maybe a ground loop changes such that an unprotected error signal gets affected. These environmental effects can also cause mode hops in a laser, which are never going to be fixable with a lockbox. Passive mitigation of environmental effects is the first step in having a laser that stays locked. At the same time, ensuring that the integrator has a large enough dynamic range is also a useful tool. If it doesn't, consider a lockbox with a larger-range integrator or a modification.

The other possible reason is spontaneous environmental effects, like somebody dropping a wrench in the lab. Don't drop a wrench in the lab. Locks can be optimized to handle these types of perturbations, but this generally involves applying more proportional or derivative gain such that the sudden high-frequency influences can be handled and are not left to an

unstable lower-frequency loop. But this often leads to servo bumps. Just because it can survive a wrench doesn't mean it has been tuned well.

An interesting technique to find the weak points in your spectrum is using an app on your phone that creates a very quiet sinusoidal tone. If the frequency of the tone is at one of the phase margins, even a very quiet amplitude can cause an unlock. Anyway, it's best to optimize a laser for servo bump performance and not necessarily robustness to a wrench being dropped. You can tune for either, but they are unlikely to be the same point.

5.3.3 Auto-Relock for Continuous Operation

An important feature implemented for the AQuA quantum computing project was auto-relocking of the Rydberg lasers, so that the experiment could continue to run autonomously overnight. Depending on the laser system, this is a task that involves utilizing tools and infrastructure outside of just the lockbox. I will be discussing the auto-relock system implemented for relocking the 1040 nm Rydberg laser, which involves an MSquared SolsTiS laser. The auto-relock scheme for the 459 nm system was similar but includes an extra stage for relocking the SHG cavity. Abraham Scott is primarily responsible for building and implementing the software for these systems, while I assisted him in setting up and designing the hardware.

To auto-relock, there needs to be a way of knowing that the system has become unlocked, a way to find the correct lock point, a way of enabling and disabling the integrator (see Section 6.1.8 for example), and a way to apply tuning to the laser's systems.

The first issue is knowing that the laser has become unlocked. This can't be determined simply from the transmission intensity of the ULE cavity, because a common unlock mode—especially for the 1040 nm laser—involves switching to a different mode. The 1040 nm laser is normally locked to a TEM_{05} mode but will often hop to a TEM_{04} mode during an unlock event.

This means that the cavity transmission does not always drop to zero and may even increase slightly.

One possible solution to this, especially if the mode being locked to is a TEM_{00} mode, is a well-tuned "edge" detection analog measurement device. However, if there is a spontaneous and brief unlock event where the laser ultimately relocks itself within a few milliseconds (e.g., if somebody drops a wrench, causing the laser to jitter but ultimately re-settle), then this would produce false positive readings.

As such, while the ULE cavity transmission is indeed digitized (with a Raspberry Pi-integrated analog-to-digital converter shield) and used in this auto-relock scheme, it is not the primary detection method for an unlock event.

Another detection mode used is image analysis of the transmitted light through the cavity. We already have a Raspberry Pi camera set up to image the mode while locking. Abraham also wrote an image analysis script that identifies the TEM_{0x} cavity mode being transmitted (or no mode, if the laser is unlocked). This software implements contour analysis and similar tools in Python to analyze the mode.

The last measurement tool is a Moglabs "Compact Wavemeter (MWM)," which uses diffractive optics, lenses, and a camera to measure the wavelength of a fiber-coupled laser. This device is highly precise—within 100 MHz according to the datasheet—but due to many possible environmental effects, it is not very accurate. As temperature, humidity, and pressure drift, this device loses its calibration and acquires large, unpredictable systematic errors that change throughout the day. However, the ULE cavity does not drift. So, as long as the laser is properly locked, we can continuously calibrate this device by saving the currently measured frequency as the desired lock point. This allows us to have 100 MHz resolution on our laser frequency using a very cheap wavemeter. The wavemeter communicates over Ethernet with the main Python auto-relocking program, and the aforementioned

"calibration" is performed in software (i.e., the lock goal is saved in code, not on-device). This device can also be used to monitor unlock events—if the frequency suddenly changes, for example. The free spectral range of the 10 cm ULE cavity is $\Delta\nu_{\text{FSR}} = \frac{c}{2L} = 1.5$ GHz, so if the laser jumps from one axial cavity mode to the next while maintaining the same transverse cavity mode, the laser frequency will have changed by at least 1.5 GHz, which is easily measurable by this wavemeter, since all environmental drifts occur at much lower frequencies than an unlock event.

These three analysis modes are used both for detecting an unlock event and for tuning the laser back to the correct frequency after one. The Python program constantly monitors these systems to determine whether the laser is in a proper lock state. Once it detects an error state, it unlocks the integrators—meaning that the integrator capacitors in the lockboxes are drained and shorted.

For these Rydberg lasers, as mentioned above, the lockboxes used are the Moglabs FSC lockboxes. These lockboxes have a lock control input, which allows the lockbox to independently clear both the fast and slow loop integrators and hold them there. The GPIO output of the Raspberry Pi is used to send this signal to the lockbox.

Once the integrators are cleared, the lock point needs to be found again. The SolsTiS digitizes the signal sent to the slow (and fast) feedback lines before sending it to the piezos. It is unclear what architecture they use, but I wouldn't be surprised if they use a high-voltage DAC and skip the need for a high-voltage amplifier. Regardless, the "resonator tuner" setting on the SolsTiS webpage controls the DC offset of the slow piezo voltage (while the slow input is meant as a fine-tuning input). As such, the laser frequency can be ramped using the already existing digital interface of the SolsTiS. To do this, the Python script sends HTML commands to the SolsTiS to scan the frequency of the lock.

Once the scanning script has found the lock point—based on the cavity

transmission mode and frequency measured by the wavemeter—the integrator of the fast loop in the Moglabs FSC is re-enabled. There is always a decent chance that the lock attempt does not succeed, so an immediate check is performed to see whether the laser successfully locked before enabling the slow lock as well. If the laser did not successfully lock, it finds the lock point and tries again. Then it enables the slow loop and performs another check. Rinse and repeat.

For the 459 nm system with SHG, a very similar scheme is used, but this time there is no specific transverse mode that needs to be reached, and the unlock measurement is simply 459 nm power. A scan searches for intensity again and applies the lock.

One of the primary limitations of this auto-relock scheme is if there is a mode hop as a result of the unlock event. Mode hops can be a cause for lasers to unlock, but the SolsTiS architecture tends to be very resilient to these types of situations as long as it is well coupled with a pump laser. However, the etalon lock can be lost if the slow piezo is changed at too high of a rate, so scanning the laser should be performed slowly. If the laser system in use tends to always mode hop when it unlocks—such as in a Vexlum or diode laser situation—more complicated scanning procedures would need to be implemented to find a stable lock point again.

6 DEVELOPING LOW NOISE ELECTRONIC SYSTEMS FOR LASER AND CURRENT CONTROL

6.1 The QPAL Lockbox

6.1.1 Design

The QPAL Lockbox is a multi-loop feedback box designed to control and lock lasers to a feature in an error signal. The design emphasizes an implementation with useful features such as dual balance-able inputs, an emphasis on ground loop suppression through instrumentation amplifiers, low output impedance, interfaceable re-locking, a ramp trigger generator, and many more features.

The QPAL lockbox is the next in a series that the QPAL group has




Figure 6.1: Image of the populated PCB of a single channel of the QPAL Lockbox.

designed, based off of the original famed "JILA Lockbox" which was designed by John Hall. Though, at this point, the lockbox has taken its own form in a fully fleshed out implementation.

These lockboxes were designed to be produced for the many lasers and laser systems around the lab. Once the prototyping stage was completed, we sent the lockbox design to a company which printed and populated the PCB using pick-and-place machines, and allowed for easy and quick implementation of the lockboxes. Unfortunately, the lockboxes (and coil drivers) suffered greatly from the chip shortages of 2021 and some of the parts and components were hard to acquire. Due to this, we had to purchase and populate some of the chips ourselves. Care should be taken to ensure that any lockbox re-orders are properly vetted before just outright purchasing the same boards from the vendor (Macrofab).

The latest version of the QPAL Lockbox authored by me is version 1.2.4. This version utilized the AD8429 instrumentation amplifier (in-amp). The utilization of this chip is at the core of the QPAL Lockbox's ground loop noise suppression. The key idea is that it performs a subtraction of the ground voltage from the signal voltage applied to the lockbox inputs without electrical contact between the grounds. Unfortunately, these chips have a bandwidth limitation of about 1MHz. At the time of design, the AD8429 chips are the highest unitary-gain capable bandwidth in-amps (with other necessary properties for a lockbox such as voltage range, offsets, etc.).

If we didn't buffer the inputs in this manner, we could have higher bandwidth lockboxes. But the design emphasis for these lockboxes is ground loop mitigation, which drove this approach. If a higher bandwidth lockbox is desired, usage of a lockbox such as the Moglabs FSC (lower signal peak-to-peak amplitude and less ground loop suppression, but much higher bandwidth) or the Vescent Lockbox can be used. This lockbox is primarily intended to drive piezoelectric actuators (such as in an MSquared SolsTiS) and laser diode current modulation (through interface with a current driver

allowing for different loop parameters to be implemented, switches, and trim-pots. Unfortunately, due to the attempts to maximize lockbox performance and minimize PCB size, the front panel arrangement is a mess and doesn't follow a logical placement of control locations.

The lockbox has two loops. A "fast" loop and a "slow" loop. Intrinsicly, these loops are not actually intentionally bandwidth limited (unless a lowpass filter is enabled) and the name simply has to do with the way the loops are designed. For the "fast" loop, the architecture involves a Proportional-Derivative-Lowpass. The "slow" loop involves a Proportional-Integrator-Double Integrator-Lowpass. The purpose of separating these loops is in case there are multiple types of feedback actuators in the system. For example, an MSquared SolsTiS has a stack of both a fast (100kHz BW) and a slow (50Hz BW) piezo, or an ECDL may have a slow piezo and a fast current modulation.

But there can also be a desire for a single combined system, where both fast and slow loops combine together at the very end of the loop. This opens the door for even more complicated non-linear PID models to exist, specifically a summation between two PIDs. The lockbox has a capability to internally perform analog addition of the output of the "fast" loop to the output of the "slow" loop. This can add for great tuning power, where a user may want to specifically target high frequency noise by implementing a low strength PDL (Proportional-Derivative-Lowpass) loop that has a high bandwidth on the lowpass, while also performing a stronger PIL (Proportional-Integral-Lowpass) loop for the lower frequencies, allowing the lockbox to drive the error signal to zero. This summation is intrinsicly still linear, but expands on the PID model, which allows the user large amounts of control over the system noise due to its versatility and practicality and ability to bypass the nonlinearities introduced by real systems. It is also an easy enough model to visualize, similar to first order perturbation theory.

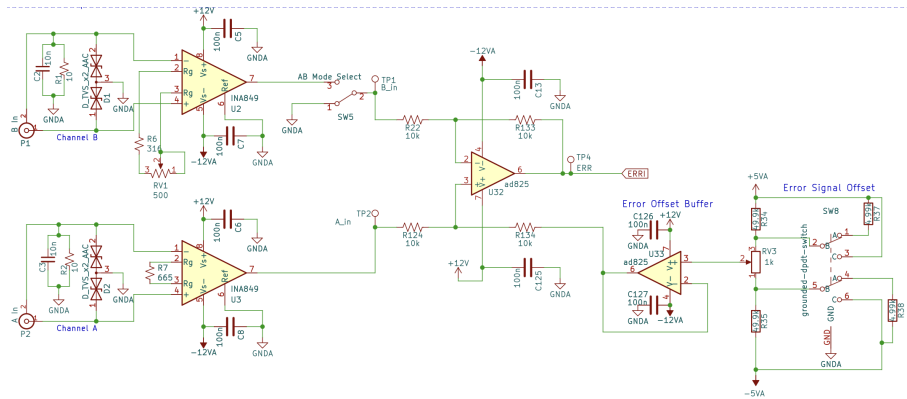


Figure 6.3: Input Processing Schematic Portion

The lockbox has three main sections to it: a shared "input" section, a fast loop, and a slow loop. There are also auxiliary sections such as the power management and filtering section and the ramp input trigger generator. The following sections discuss each in detail.

6.1.3 The Input Section

This section is the first section, where the signals enter the lockbox. The lockbox has two error signal inputs (A and B). As mentioned above, these inputs are heavily buffered for ground loop error suppression (this mitigates issues where ground drifts cause an already very small error signal to drift in offset).

Once buffered, the error signal is either composed of just the A input, or an analog voltage subtraction A-B (selected by the B enable switch). The purpose of having this second input can include background light subtraction, or applying nested integrator loops as demonstrated for the 685nm laser in section 3.2.4. Another possible use of this channel is to allow for an externally controlled error offset, which could even be digitally controlled.

Next, the internally generated error signal offset is applied. This is a constant voltage set manually by the Error Offset CCW 10-turn trim pot

6.1.4 The Slow Loop

The slow loop involves a Proportional, Integrator, second Integrator, and Low pass filter (PIIL). These gains are applied sequentially (in this order). There is a rotary switch for controlling the operating mode of the slow loop, "State Selector". This switch has settings of "NF, P, PI, PI+, PII", and depending on the setting, enables or disables different loops. "NF" stands for "no feedback", "P" is proportional, "PI" is a "Proportional-Integral", "PI+" is a "Proportional-Integral-Integral with limited gain", and "PII" is a "Proportional Double Integral" (with no limit on the integral gain other than the open loop gain of the op-amp).

However, it is important to note that this control mode is deprecated. This switch was implemented for legacy sake, but the best way to implement this lockbox would be with this switch set to PI followed by using other switches to enable or disable features.

The feedback mode rotary switch

NF is the "no feedback" mode of this switch. However, for each of these settings, the lockbox is not actually disabling the loop, rather highly attenuating the particular gain contribution. So even if the switch is set to NF, there is still proportional feedback, but highly attenuated. To truly disable the feedback, there is a physical switch which drives the output to 0V through the utilization of a short.

P mode is simply a "proportional" mode, but the integrators are not fully disabled, simply highly attenuated. It is ultimately not a pure proportional! (This is the theme. Again, I am not proud of the feature, and the reader may be asking "why is it implemented", and I can't remember why I didn't just remove it. Just leave it to PI and don't use it).

PI mode is the main mode that I leave the switch on. It doesn't disable the two important loops (P, I) and it disables the unnecessary loops (the

second integrator). This mode can be turned into a true P mode through utilization of the "integrator clear" switch (more on this later).

The PI+ and PII modes are similar, where they enable the second integrator, but to different degrees. To this day, I have never used a second integrator. The concept of a second order integrator, to my understanding, is to increase the DC gain with respect to a different feedback loop. It makes sense in an analytical analysis of multiple PID systems working in unison on the error signal, in that you can make a particular loop (maybe a loop with the highest dynamic range, like the slow piezo in an MSquared SolsTiS), be responsible for the very low frequency component of the error signal.

However, even in the analytical model, this does have a large draw back where the second order integrator now has a phase shift of 180 degrees, at all points in the frequency spectrum. All feedback from a second integrator has the wrong gain sign, which is the epitome of a servo bump. Second order integrals need to be implemented with large amounts of care, even if electronic systems were perfect. But, even with this analytical issue with second order integrators, there is also an issue with the fact that we are actually implementing these systems with imperfect electronics that don't completely follow the linear model of the "ideal op-amp". There are effects like offset voltages and currents, voltage rails, slew rate, etc. that are very influential on the capabilities of an op-amp in special regimes of operation. In the case of second order integrals, the offset voltages and currents are the primary issue, because a very small DC signal error will propagate through the integral. An improved implementation of utilizing these lockboxes for solving this particular concern (of wanting a slow piezo controlling the DC gain while still utilizing an integrator on the fast piezo) is discussed, implemented, and demonstrated in the 685nm locking section 3.2.4.

Using the PI setting as the primary setting

As has been mentioned profusely, the best way to use the lockbox is to set this switch to PI, and use the linear switches on the lockbox to control the different states. Namely, to enter "no-feedback", set the "Feedback Control" switch to off, which shorts out the feedback signal (while still allowing the ramp and output offset). To disable the integrator, (say, while trying to lock the laser), simply set the integrator clear switch to "Clr" or "(Clr)" and the integration capacitors will be shorted using an analog switch (the parenthetical Clr is a switch notation to indicate that the switch has a spring in it that will push the state out of that mode unless the user is actively pushing it up; its a temporary state. This is very useful if the user wants to quickly re-lock the laser and can't be bothered to flop it twice). Additionally, this analog switch has a "de-bouncer" before it, which is a digital circuit which delays the change of the digital pulse, so that the "shutter" of the switch change doesn't cause the integrator to quickly enter and leave and reenter the integrator mode (which can cause a mode-hop on a laser, which is often considered "annoying").

The proportional gain level is controlled by the Servo Gain rotary switch. This switch is labeled in signal dB, indicating how much gain (or attenuation) is applied to the error signal. After large amounts of usage of this lockbox, I have determined that a higher level of attenuation would have been nice to implement, say -30dB instead of +70dB. This would be a good upgrade for a future revision.

The PI (Hz) rotary switch selects where the PI corner of the integrator is placed, in Hz. This "corner" in the frequency domain is where the contribution of the integrator is of equal amplitude to the proportional contribution. The "INF" means that there is no corner, and the proportional never meets the integrator. This would be where the user would set it if they wanted a pure integral for the gain of the lockbox.

Hidden among the fast loop settings, there is a Slow Invert switch, which

is used to invert the slow signal. This switch has to be placed in the right setting in order to perform positive or negative feedback, depending on the physical properties of the controlled system (does a positive voltage push the error signal in the positive or the negative direction?). It is unfortunately hidden among the fast loop settings so look sharp.

Next is the 9dB corner rotary switch. This controls the I-I corner, for when the second integrator is enabled. I've never wanted to use the second integrator so I haven't tested this out much.

The last component to the feedback, and one of the most important components, is the low pass filter at the very end. This is a first order lowpass filter, implemented with an array of resistors and capacitors, and a trim pot. The range listed on the rotary switch indicates what the RC corner of the filter will be if the "Lowpass RC Adjust (CCW)" trim-pot is fully turned "clockwise" - "counter-clockwise" (for all trim-pots, there is a "CW" or "CCW" written, which informs the user which direction increases the parameter. So in this case, CCW increases the RC corner frequency). The "INF" setting here disables the low pass. It is only recommended if the user wants to apply their own external filtering, or if there is intrinsic, clean filtering in the feedback path after the lockbox. But it should be noted that the lockbox does not have the cleanest phase performance at its peak bandwidth, since it is composed of many op-amps and in-amps. A filter is always necessary for a realistic PID system to mitigate servo bumps at high frequencies.

There is also a "slow fine gain (CCW)" trim pot. I have mixed feelings for this thing. It is implemented as an attenuator after the feedback path, including the integrator. This means that it can limit the full-scale range of the integrators output, which is a primary avenue for lasers to come unlocked after large amounts of time (eventually, slow environmental effects like lab temperature may require a piezo to be actuated at higher and higher levels, but if the integrators output voltage can't reach that level, even with

an external amplifier, then the lockbox will be unable to maintain the lock). However, it is an important feature to be able to reduce the proportional fine gain of the lockbox, so this trim pot attenuator is necessary. It is my recommendation that it be set to the CCW position initially when picking the servo-gain parameter, and brought down slightly as needed.

You now have a feedback signal. However, to utilize this signal, it is also useful to apply an offset voltage, or a ramp. This is where the slow ramp and slow offset knobs come into play. The slow offset knob will add a voltage from 0V (fully clockwise) to +10V (fully counterclockwise) to the output signal. This is useful for scanning for cavity modes or atomic transitions in the error signal. However, the maximum output voltage of all of the op-amps is limited to about 11V, due to the power line noise cleanup with regulators (see the power management in 6.1.7). So if an offset voltage of +10V is applied, the lockbox will have limited dynamic range. Thus, if this knob is being used on a feedback line that is not intended to control the global offset of the system (say a fast piezo or a current modulation port), then rotating the offset knob fully clockwise will disable any offsets.

The "slow ramp CW" knob and "ramp control" are used to control the ramp that is added to the output of the lockbox. The lockbox does not internally generate a ramp. This would have required lots of infrastructure and probably made the lockbox cost twice as much, and be much larger to allow for enough controls to reasonably create a ramp. Modern lockbox instrumentation tends to involve digital ramp generators, but I didn't want to mix digital ground noise into this lockbox. However, an external ramp can be inputted, and the input has strong ground noise buffering to reduce noise from external sources. When a ramp signal with a 0 DC component is inputted into the Ramp In BNC port, a ramp trigger signal is outputted (see 6.1.7), and this signal can be added to the output of the lockbox. The "ramp control" switch shorts the ramp. If it is enabled, there can also be an attenuation applied to the ramp before adding it into the

output, controlled by the slow ramp CW knob. The amount of attenuation applied is infinite (meaning, no signal can pass through) when the knob is at the full counterclockwise position, and decreases non-linearly (the ramp amplitude increases) as it is rotated clockwise. This allows the user to ramp down into a feature in the error signal before activating the integrator. It is recommended, once locked and the ramp knob is fully counterclockwise, to set the "ramp control" switch to "off" to truly set the attenuation of the ramp input to infinity.

The ramp input can be used for other features, such as allowing external control of the output of the lockbox. See the auto-relock section (6.1.8) for more details.

The last component of the slow loop discussed here is output monitor port. The "Monitor Mode Select" switch controls the functionality. This output monitor can be configured to either see the exact voltage being outputted from the lockbox (when "Post-Offset" is selected), or just before the offset voltages (ramp and offset) are added in (when "Pre-Offset" is selected). However, due to the nature of the way the lockbox is implemented, there is an inversion from the actual output when the Pre-Offset mode is enabled, meaning that the actual pre-offset voltage is negative of what is being seen. The monitor output does include the added fast loop output, if that connection is enabled (and this output is not inverted).

When in doubt, look at the circuit diagram to fully understand what is happening.

6.1.5 Rectifying the Slow Loop Output

It may be important to ensure that the output of the slow loop is always a positive voltage. For example, if the lockbox is directly controlling the voltage into a piezo, then a negative voltage can destroy the device. To this end, an internal rectification diode can be enabled by moving the J19 jumper from the bottom two pins to the top two pins. This forces all the

signals to travel through the diode, which means any voltages less than 0.6V are met with a very high impedance and a 10kOhm resistor ties the output to 0V. The consequence of this is that all positive voltages are reduced by 0.6V. But the "Slow Offset CCW" knob can add this 0.6V back very easily with half a turn. Unfortunately, the only way to see and access J19 is to remove the box cover, so it is recommended that the user labels such a change on the cover. A future improvement can involve adding a hole in the cover to allow external maintenance of this jumper.

An important consideration is that this rectifier does not rectify any signals that may have come in through the fast loop. So ensure that the fast loop is not set to output to the slow loop. And since the label on the front panel was labeled incorrectly, this means that the switch labeled "Medium to Slow Out" should be set to "enabled" to ensure that it is actually disabled. For added clarification, when the "Medium to Slow Out" switch is pointing upwards, this is enabled and the fast loop output will be added onto the slow loop output. When it is pointing downwards, it will only output to its own output.

6.1.6 The Fast Loop

The very first thing to understand about the "fast" loop is that it is sometimes labeled as a "medium" loop as well. These words are used interchangeably to describe the same loop, but I will primarily call it the fast loop.

Firstly, there is the rotary switch labeled "Med PD Corner (Hz)". This allows the user to either disable the derivative entirely or choose where it is equal in gain to the proportional feedback in the frequency domain. If D off is selected, the loop will only be proportional.

The next important aspect is the "Fast D Gain Limit (dB)". This rotary switch is used to determine where the Derivative levels off and becomes a flat gain in the frequency domain. This is physically implemented by applying a resistor in series with the derivative capacitor, which is a very

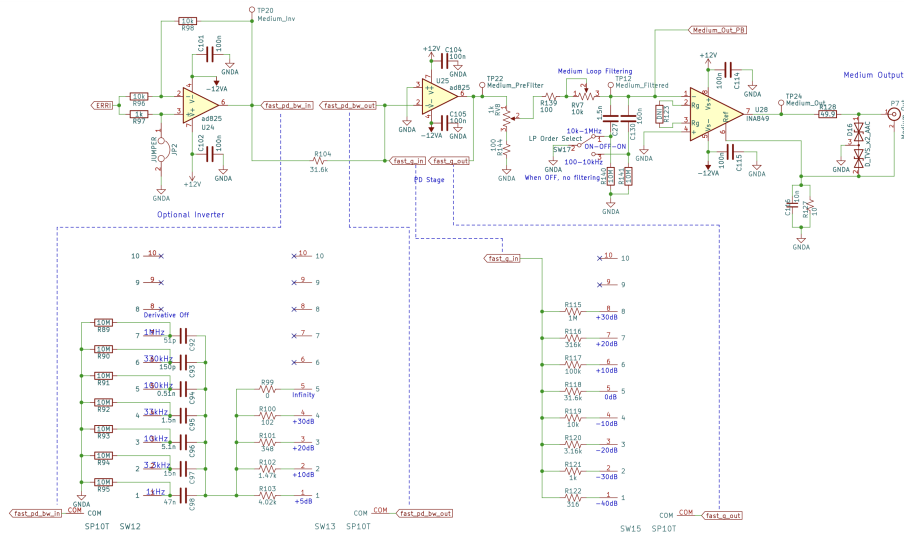


Figure 6.5: Fast Loop Schematic Portion

clean way of maintaining the linearity of the op-amp. However, this has the same mathematical effect as applying a first order low-pass filter after the derivative. This gain limit can be disabled by setting it to "Inf", which will cause the gain limit to be influenced by the other nonlinearities throughout the system.

Below the "Fast D Gain Limit" is an unlabeled rotary switch, which is the course proportional gain setting in signal dB. There is also a medium fine gain trim pot for applying fine gain control. To apply a gain inversion (the gain sign), there is the "Med Inverter" switch which is used to select the inversion setting of the medium loop.

Finally, there is a lowpass filter, with parameters selected by a linear switch and a trim pot, very similar to the slow loop section. Setting the linear switch to the middle position completely disables this low pass filter.

The final important part of the fast loop is the "Medium to Slow Out" switch. The very first thing to know about this switch is that the labels on the frontpanel are backwards, where the "enabled" and "disabled" labels

should be swapped. For added clarification, when the "Medium to Slow Out" switch is pointing upwards, this is enabled and the fast loop output will be added onto the slow loop output. When it is pointing downwards, it will only output to its own output. This won't disable the output from being sent to the "Med Out" port. However, the signal sent to the slow output is an inverted signal from this output. If there is a signal rectification, it is very important to understand that this signal bypasses it. Read section 6.1.5 closely.

Always, when in doubt, take a look at the circuit diagram.

6.1.7 Other Sections

This section discusses the input power filtering and the ramp input trigger generator.

The input power filtering is performed using passive components (capacitors and ferrite beads) as well as active components (linear voltage regulators). The capacitors and ferrite beads allow for energy to be stored locally on the board, so that when a chip needs it, there doesn't have to be a pull leading all the way up to the power supply to acquire the energy needed.

The voltage regulators are intended to add additional power supply noise reduction, on top of that already provided by an op amp or in-amp. It also protects the devices from being supplied a voltage above specifications. Unless modified, these are +/- 12V regulators (as well as a +5V regulator for the simple digital circuits on the board).

The other feature discussed in this section is the ramp input trigger generator. This is a poorly implemented generator, and could really use some improvements. The idea is that it monitors the ramp input and creates a TTL signal that is high when the ramp is positive or 0 when the ramp is negative. This in turn can be plugged into an oscilloscope. This is implemented using a simple comparator. The primary issue found with this

implementation is insufficient jitter control. If the ramp signal is noisy as it passes 0V, there might be multiple TTL swaps, which can be annoying for triggering on the oscilloscope. This can generally be fixed by doing a high frequency reject on the oscilloscope trigger settings. An improvement to the lockbox, where a lowpass filter is added to the ramp signal before the comparator would be able to greatly improve the performance in this regard.

6.1.8 Auto-Relock Capabilities

This section discusses the auto-relock capabilities of this lockbox. A discussion of an implementation of auto-relocking using other lockboxes (before this lockbox was developed) is discussed in section 5.3.3.

The QPAL lockbox also has a "Computer Interface" Molex connector used to allow for TTL signals to change the state of the lockbox. Alongside controlling the state of the lockbox (not tested for aforementioned reasons), the primary control here is to be able to enable or disable the integrators, to allow the user to relock an unlocked laser.

Once a laser has been determined to be unlocked from the desired setpoint, the first thing that occurs is that the integrators either tend to rail because a zero in the error signal cannot be found, or a different lock point is found and the laser is now erroneously locked. To remedy this, the laser needs to be "unlocked", meaning the integrators need to be disabled. This can be done manually (with the switch), but the QPAL lockbox also allows a digital input to override the "unlock" or to enter a "clear" state. Thus, a high ($>2V$) TTL signal inputted on the integrator clear in the computer interface will short the capacitors, allowing the laser to be scanned and relocked. In order for this to work, the lockbox has to be in the uncleared, middle, state of the "Integrator Clear" switch.

Now that the laser is unlocked, the ramp input of the lockbox can be used to scan the offset voltage output to search for the correct error signal. This

would involve an external system that interfaces with external measurement devices such as a wave meter and/or cavity transmission modes. See section 5.3.3 for an example of this.

Once the voltage is ramped to the right setpoint, the integrator clear can be unset and the integrators will resume working to maintain the laser lock. The offset voltage that was applied at the ramp input can now be adiabatically ramped down to zero, which will likely have less noise, depending on the voltage source used.

6.1.9 Future Improvements

There are several areas where the QPAL Lockbox could be improved in future versions. One key upgrade would be replacing the AD8429 instrumentation amplifiers with higher-bandwidth alternatives—if and when they become available—to boost overall responsiveness while still suppressing ground loops. The front panel could also benefit from a clearer layout and better labeling to make tuning more intuitive and reduce user error. On the electronics side, adding a simple low-pass filter to the ramp input trigger would help cut down on noise and reduce unwanted jitter in the generated TTL signal. Mechanically, allowing external access to jumpers like J19 (e.g., through a panel cutout) would make maintenance easier and faster. Finally, expanding the servo gain attenuation range (e.g., down to -30dB rather than topping out at $+70\text{dB}$) would give users more versatile control.

Ultimately, the QPAL lockbox is a very well functioning electronic device that is actively locking many lasers around the lab. From simple situations like current modulation in DFB diodes to new Rydberg laser systems. It has its quirks, but such is the nature of a home-built electronic systems.



Figure 6.6: Image of the frontpanel of the QPAL Coil Driver.

6.2 The QPAL Coil-Driver

6.2.1 Design and Features

The QPAL coil driver is based off of a coil driver design I created during my undergrad at JILA, which was based off of a design built by the on-staff electrical engineers. Danny Wendt contributed in the creation and design of these QPAL coil drivers.

It follows a basic principle where there are two MOSFETS used to drive or pull current through a coil. The current driven is measured through a sense-resistor (1 Ohm) and used as a feedback signal to set the right voltage for the MOSFETs.

To facilitate this, the QPAL coil drivers incorporate over-heat protection, current monitoring, and two separate setpoint inputs. As a theme for all electronic work necessary for high fidelity control of systems in the lab, mitigation of ground loops and noise was an important focus for the coil drivers.

To this end, there are two power inputs to these drivers, the high current

lines (+/- VHC, PWRGND) and the analog power lines (+/- Vcc, AGND). The ground of these inputs are not connected anywhere on the board, since it is understood that the high current "ground" (PWRGND) voltage (with respect to the analog ground) will be linearly dependent on the amount of current being drawn through the system, whereas the analog ground (ANGD) can be considered as the true ground of the system.

To achieve this, an instrumentation amplifier (AD8429) is used to measure the voltage drop across the sense resistor, and reference that voltage to the true, analog ground.

There are bandwidth and damping trim pots used to tune the coil driver to properly drive the current through the inductive coil. Since the coil itself has a high inductance as well as resistance, a properly tuned feedback system is important. This feedback model doesn't really follow a colloquial PID loop.

The bandwidth term limits the rate that feedback can be applied to the coils by measuring and filtering the voltage directly before the coil. This feedback is AC-coupled so it somewhat looks like a derivative term. The damping term observes the actual current through the coil and is observing the low frequency response, contributing similar to a proportional-integrator term in the feedback train. Observing and modeling the circuit diagram can further help a user understand these behaviors.

Figure 6.7 shows a simplified schematic of the feedback and control segments just mentioned.

As mentioned before, there are two inputs for the setpoint. These are a primary setpoint input and a "feedforward" input. These inputted voltages are buffered with an instrumentation amplifier (AD8429) and have an input impedance of 1 MOhm. These inputs have different gains and are summed together. The initial design of the lockbox has the voltage to current gain set to 0.21 Amps per Volt for the primary input and 0.05 Amps per Volt for the "feedforward" input.

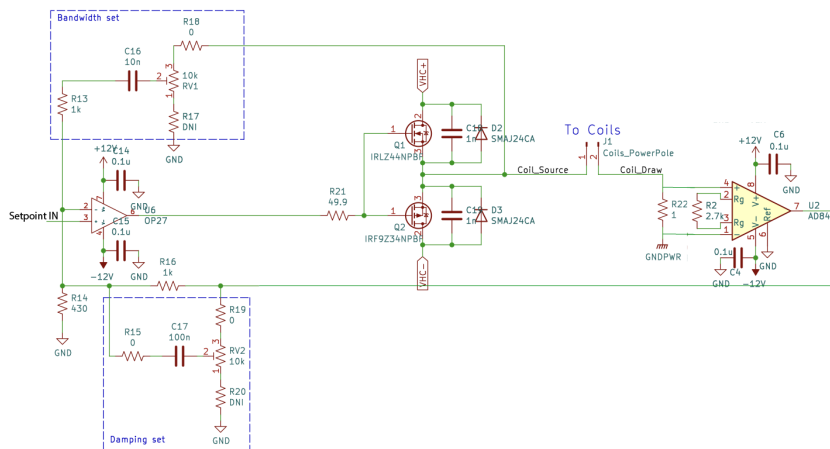


Figure 6.7: Feedback Control Schematic of Coil driver

The feedforward input on the coil driver is designed to accept a corrective current setpoint from an arbitrary waveform generator (AWG), which is triggered by the AC line voltage. This allows for compensating fluctuations in the line voltage by applying a small, synchronized correction to further stabilize the magnetic field near zero. While the coil drivers were designed with this application in mind, the feedforward input is not required for standard operation and may be left floating when unused.

The internal architecture of the coil drivers is designed to run up to 2.5A in a coil that has a resistance on the order of 1 Ohm. Less impedance is fine. But too much resistance can become an issue for the bipolar MOSFET architecture. If a higher current draw is desired, the sense resistor can be reduced in impedance. A further discussion of these limitations is in the next section, 6.2.2.

Over-temperature protection has also been implemented to ensure that the MOSFETs in the coil driver do not overheat. The temperature is measured by a thermistor attached to the MOSFET, and the resistance is measured and compared to a maximum acceptable resistance. If it overheats, the system enters a shutdown state that needs to be reset by physically

pressing the front panel button. With the cooling architecture installed around the MOSFETs, the maximum power dissipation they can handle is around 9W which is further described in the next section, 6.2.2.

Lastly, a monitor port for measuring the temperature of the sense resistor is added. This can be used to determine if there is a systematic error introduced by the resistor changing temperature too much. Although the sense resistor should be picked such that it has a very small thermal coefficient, there can still be a small systematic effect. This allows for countering this effect.

6.2.2 Determining Maximum Current Draw

There are three factors that limit the maximum amount of current the coil drivers can draw. The first is simple, is there enough voltage from the high current power supply? Given a current, will the voltage drop through the coil and sense resistor be too high? This can be described with the relationship $V_{HC} > V_{source} = V_{sense} + V_{coil} = I(R_{sense} + R_{coil})$ (the minimum voltage drop through the MOSFET can be pretty small).

The next factor is considering the voltage at the "source" pin of the MOSFETs vs the maximum voltage that can be driven at the "gate" pin. There is an op amp with +/- 10V rails which drives the gate pin. But if the voltage at the MOSFET source pin is too high, there won't be a high enough voltage drop between the gate and source to activate the MOSFET. This generally has to be up to 3V, but ultimately depends on the actual MOSFETs and the drain voltage as well (check the part datasheet!).

The last consideration has to do with the overheating of the MOSFETs. While a larger V_{HC} can allow for much faster switching speeds of the coil current, it can also lead to MOSFET overheating. For example, the maximum amount of power dissipated in the MOSFETs cannot exceed 9W before the equilibrium temperature exceeds the overheat temperature. We

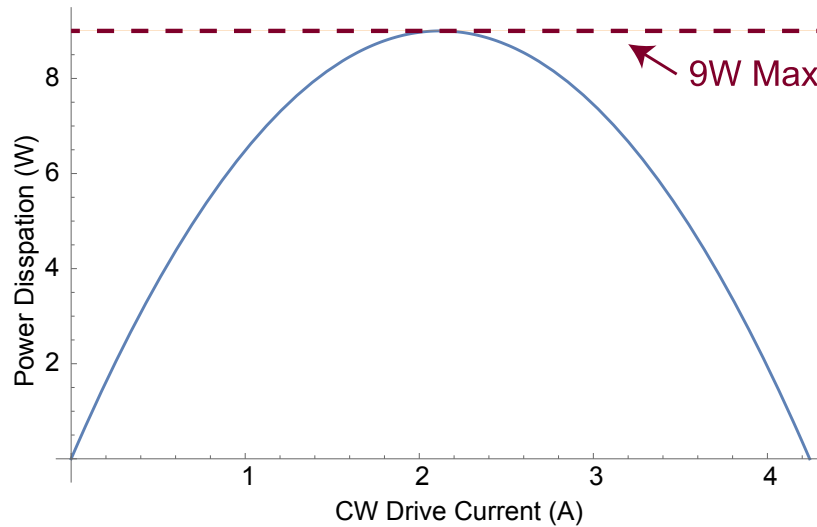


Figure 6.8: MOSFET Power Dissipation for a 1 Ohm coil and 8.5V V_{HC} .

can calculate the power dissipation in the MOSFET:

$$\begin{aligned}
 P_{\text{MOSFET}} &= I \times V_{\text{MOSFET}} \\
 &= I \times (V_{\text{HC}} - V_{\text{sense}} - V_{\text{coil}}) \\
 &= I \times (V_{\text{HC}} - I(R_{\text{sense}} + R_{\text{coil}})) \\
 &= IV_{\text{HC}} - I^2 R_{\text{total}}
 \end{aligned}$$

where R_{total} is the total resistance of the coil and sense resistor and I is the current being driven. This function is a negative parabola, with a maxima at $I_{\text{peak power}} = \frac{V_{\text{HC}}}{2R_{\text{total}}}$. We can use this to solve for the maximum dissipated power given a V_{HC} and R_{total} : $P_{\text{max}} = \frac{V_{\text{HC}}^2}{4R_{\text{total}}}$, or solve for a maximum supply voltage given a maximum allowed power $V_{\text{HC,max}} = 2(P_{\text{max}}R_{\text{total}})^{1/2}$. For a 1 Ohm coil and sense resistor, a 9W maximum power draw, we would want a maximum V_{HC} of about 8.5V. Figure 6.8 demonstrates the relationship between power dissipation in the MOSFET vs current through the coils for this set of parameters.




Figure 6.9: Current Noise Spectral Density at 1A for the coil driver. The current noise spectral density is less than 1ppm. The noise below 1kHz is dominated by the noise of the input signal, including the spikes on the spectrum.

It is important to note that these calculations are assuming constant current. What really matters is the average power dissipation over a small period of time (say, 10 seconds) since 9W is the equilibrium heat dissipation rate.

6.2.3 Noise Performance

The noise performance of the coil drivers is measured by placing a 1 Hz low pass filter on a 4.5V source and inserting this into the coil driver, to drive exactly 1 Amp. The monitor port of the coil driver is plugged into a spectrum analyzer and the current noise spectral density is measured using the 1 A/V gain of the sense resistor. Figure 6.9 presents the noise spectral density of the Coil Driver from 10Hz to 100kHz.

As can be seen in the figure, there are clear spikes at 60Hz, 120Hz, and 180Hz, etc. in this spectrum, caused by AC line noise and its harmonics.

Ideally, these would be much more suppressed, but are introduced by the power supply. A quieter power supply would definitely reduce these. The main points of interest of this spectrum are the broader bumps, at 2kHz and 30kHz, since these are introduced by the feedback mechanisms.

6.2.4 Future Improvements

The coil drivers work rather well for ~ 1 Ohm coils and < 3 Amp currents. But if a user finds themselves outside of this realm, modifications will need to be made. Generally, for experiments with glass cells, a 1 Ohm coil is pretty standard, so this isn't an issue. Modifications can be made to the existing architecture without having to spend lots of effort in redesign.

If mono-polar current is acceptable (say, a quadrupole magnetic field for a MOT), larger currents can be acquired with the same analog architecture by placing the MOSFET after the coils instead of before (meaning $V_{HC} \rightarrow$ Coils \rightarrow MOSFET $\rightarrow R_{sense} \rightarrow$ PWRGND). This way, the resistance of the coils doesn't matter and there will only be R_{sense} to offset the source voltage of the MOSFET.

Another modification, if the user wants to modify the gain of the coil driver to improve noise performance at the expense of dynamic range, the sense resistor can be increased. This way, noise on the input will be attenuated by the gain.

A larger improvement requiring more design effort would be to change the architecture of the current drivers to use a high current op-amp, such as the OPA549 instead of manually implementing the MOSFET architecture by hand. This chip can drive up to 8A CW and has rails within 5V of V_{HC} at this voltage. This is achieved by proper biasing and control of the transistors from professional engineering of these types of systems.

Part III

Next Steps and Conclusions

7 NEXT STEPS AND OUTLOOK

This chapter outlines proposed next steps for advancing experiments with the quadrupole interaction in Cs. These fall into three main categories: upgrades to the 685 nm laser system, further exploration of quadrupole cooling performance, and improvements to quadrupole imaging, including a proposed quenching scheme involving the $6P_{3/2}, F = 5 \leftrightarrow 5D_{5/2}, F = 6$ transition [36].

7.1 Upgrades and Future work on the 685 nm Laser Systems

Most of the upgrades and future work presented in this chapter will require improvements to the 685 nm laser system. This section discusses potential upgrades to the laser system that would allow for higher intensity of 685 nm light at the atoms and better control of the beam polarization and pointing (sections 7.1.1 and 7.1.2), and improvements to the laser's spectral noise (section 7.1.3). Furthermore, as the MSquared laser system continues to degrade, a new laser source will eventually be required. A replacement laser system is also discussed in section 7.1.3.

7.1.1 Independent Beam Control for a 685 nm MOT

The 685 nm light is coupled to the experiment apparatus through four fiber launches, where it is split into a total of six beams, as discussed in section 3.1.2. The front and back beams each have their own fiber launch, while the top and bottom beam paths each share a single launch that splits between the left/right diagonals. The beams are combined with the 852 nm beams on dichroic mirrors and then share achromatic waveplates to create the required σ polarizations.

This design has major drawbacks. It is difficult to achieve independent control of the relative beam pointing and polarization for the 685 nm beams. Furthermore, the glass cell has poor coatings at 685 nm, so the achromatic waveplates are ineffective due to additional chromatically dependent retardance from the cell. Acquiring a new glass cell with coatings optimized for 685 nm would greatly improve system performance.

Additionally, splitting the beams after the fiber launch halves the available power for each beam. A productive upgrade would involve implementing six independent fiber launches (one for each beam) as well as a pair of half- and quarter-waveplates for each beam before combining with the 852 nm beam.

7.1.2 Upgrading the 685 nm Injection Lock Amplification Setup

The proposed upgrade involves using six independent injection-locked diodes instead of just two. Each diode can provide up to 200 mW of power to an independent fiber launch. Assuming an isolator efficiency of 80%, AOM efficiency of 70%, optical train efficiency of 80%, and fiber coupling efficiency of 60%, we estimate about 54 mW of 685 nm light at the table per launch. With a 0.7mm radial $1/e^2$ beam waist, this yields a peak intensity of $7\text{W}/\text{cm}^2$, or $s = 8.5$ per beam, for a total of $s = 51$.

Accounting for 50% reflection losses on the front and back beams, the top and bottom beam sets would need to be decreased by the same factor to maintain balance. This still allows for a total of $s = 25$, significantly higher than the current $s = 2$. Realizing this setup would require substantial infrastructure upgrades.

7.1.3 Upgrading the 685 nm Source

The current 685 nm source is an aging MSquared SolsTiS laser with at least ten thousand lasing hours. With increasing difficulty in sourcing a 532 nm pump laser, a new lockable 685 nm source must be considered.

A 685 nm laser diode in an ECDL configuration is an attractive replacement, using current modulation and a piezo element to lock to the ULE cavity. Although it produces less power than the MSquared laser, it offers modularity and replaceability. The plan is to use an HL6750MG laser diode in an ECDL, coupled to the ULE cavity for PDH locking. The diode has a maximum output of 50 mW.

The transmission of the light from the ULE cavity has a spectral width equal to the linewidth of the cavity, since only the resonant light can be transmitted. With this, the PDH lock can be overdriven to ensure a very stable lock while preventing servo bumps from reaching the atoms. We can use the cavity transmission to injection-lock a second diode [64] using as little as 100 μ W of power. This diode acts as a pre-amplifier before feeding into a six-diode main amplifier stage using HL69001DGs (200 mW each, 5 mW seed required). This method, already demonstrated, offers simple, stable laser sources and is a promising option for the future of the 685 nm system.

7.2 Exploration of Quadrupole Cooling Performance

Further exploration of quadrupole cooling performance has been limited by laser power constraints. With future upgrades, it would be interesting to investigate how cooling performance depends on laser power and beam polarization.

7.2.1 Loading a 685 nm MOT from a 2D MOT and Push Beam

The 685 nm MOT is promising for pre-cooling atoms before trap loading. This work has demonstrated cooling to 5 μK in low-gradient fields and in traps. With higher 685 nm intensity, it may be possible to capture single atoms directly from the 2D MOT and push beam into the 685 nm MOT, enabling steady-state loading.

This approach was briefly attempted but ultimately set aside in favor of investigating cooling and imaging of single atoms in the trap. In Yb, direct trapping into narrow-line MOTs has been demonstrated [96]. It remains an open question whether this can be achieved in Cs. With more power, revisiting this approach would be very interesting.

7.3 Improving Quadrupole Imaging Performance

Improving Imaging Train

As discussed in , trap light introduces background 803 nm light in the images. This can be mitigated by reducing 803 nm in the image plane with an iris and using narrower bandpass filters for 852 nm light. The current setup uses filters with $\pm 3\text{nm}$ and $\pm 0.3\text{nm}$ bandwidths, whereas 852 nm fluorescence has a linewidth of only 0.01 nm. Narrower filters would significantly reduce 803 nm background.

Focusing the atom light onto fewer pixels (e.g., from 100 to 9) would also improve SNR. Since 852 nm background light is generated along the optical path, not in the image plane, it is broadly distributed across the image. Smaller integration regions reduce background contribution. This can be achieved by relocating the dichroic mirror to the Fourier plane of the imaging system, rather than near the image plane as in section 3.1.4.

Quenching the $5D_{5/2}$ Decay for Faster Imaging

Ultimately, for a useful readout scheme, even a 40 ms exposure time is too long, despite the fact that the scheme is non-destructive and cools the atoms. Many qubit experiments [88, 97, 98] are actively working toward sub-100 μs exposure times. Even if 40 ms were achievable with our 685 nm scheme, such a delay is not practical for qubit operations.

One promising approach involves quenching the $5D_{5/2}$ to $6P_{3/2}$ decay using stimulated emission. This technique, discussed in [36], is a future direction for this work. Using two-sided imaging [10] with two NA=0.7 lenses, 80% optical transmission, and 50% quantum efficiency, a total photoelectron collection efficiency of 11% is estimated. Scattering 100 photons would yield about 11 photoelectrons. Simulations suggest that a 60 μs exposure could achieve an imaging fidelity of 0.9995.

7.4 Outlook

The work presented in this thesis demonstrates the practical and conceptual benefits of using the 685 nm quadrupole transition in cesium for high-fidelity, background-free imaging and laser cooling. This transition has proven to be a valuable tool for atomic qubit control and shows strong potential as an important component in scalable neutral atom qubit arrays.

Further work is needed to fully characterize the transition and understand its capabilities. With a more powerful and better-controlled laser platform, it will be possible to probe the detailed dynamics of the interaction and push performance even further. Improvements to imaging optics, beam shaping, and control electronics could also expand the range of applications for this transition within large-scale systems.

In addition to the quadrupole work, this thesis presents a set of tools that improve experimental stability and control—particularly in the context of neutral atom gates. These include low-noise feedback systems Rydberg laser

frequency stabilization and robust and useful electronic systems for feedback and magnetic field control. Together, these developments build toward the next generation of neutral atom quantum computers: more stable, more scalable, and better equipped to meet the fidelity and speed demands of quantum algorithms.

REFERENCES

- [1] Jens Koch, Terri M. Yu, Jay Gambetta, A. A. Houck, D. I. Schuster, J. Majer, Alexandre Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf. “Charge-insensitive qubit design derived from the Cooper pair box”. In: *Phys. Rev. A* 76 (2007), p. 042319.
- [2] Morten Kjaergaard, Mollie E. Schwartz, Jochen Braumüller, Philip Krantz, Joel I.-J. Wang, Simon Gustavsson, and William D. Oliver. “Superconducting Qubits: Current State of Play”. In: *Annual Review of Condensed Matter Physics* 11. Volume 11, 2020 (2020), pp. 369–395. DOI: <https://doi.org/10.1146/annurev-conmatphys-031119-050605>.
- [3] C. Monroe and J. Kim. “Scaling the ion trap quantum processor”. In: *Science* 339 (2013), p. 1164.
- [4] C. Monroe, R. Raussendorf, A. Ruthven, K. R. Brown, P. Maunz, L.-M. Duan, and J. Kim. “Large-scale modular quantum-computer architecture with atomic memory and photonic interconnects”. In: *Phys. Rev. A* 89 (2014), p. 022317.
- [5] John M. Nichol, Lucas A. Orona, Shannon P. Harvey, Saeed Fallahi, Geoffrey C. Gardner, Michael J. Manfra, and Amir Yacoby. “High-fidelity entangling gate for double-quantum-dot spin qubits”. In: *npj Quantum Information* 3.1 (2017), p. 3. DOI: [10.1038/s41534-016-0003-1](https://doi.org/10.1038/s41534-016-0003-1).
- [6] T. F. Watson, S. G. J. Philips, E. Kawakami, D. R. Ward, P. Scarlino, M. Veldhorst, D. E. Savage, M. G. Lagally, Mark Friesen, S. N. Coppersmith, M. A. Eriksson, and L. M. K. Vandersypen. “A programmable two-qubit quantum processor in silicon”. In: *Nature* 555 (2018), pp. 633–637.

- [7] M. Saffman, T. G. Walker, and K. Mølmer. “Quantum information with Rydberg atoms”. In: *Rev. Mod. Phys.* 82 (2010), pp. 2313–2363.
- [8] M. Morgado and S. Whitlock. “Quantum simulation and computing with Rydberg-interacting qubits”. In: *AVS Qu. Sci.* 3 (2021), p. 021101.
- [9] I. S. Madjarov, J. P. Covey, A. L. Shaw, J. Choi, A. Kale, A. Cooper, H. Pichler, V. Schkolnik, J. R. Williams, and M. Endres. “High-fidelity entanglement and detection of alkaline-earth Rydberg atoms”. In: *Nat. Phys.* 16 (2020), p. 857.
- [10] T. M. Graham, Y. Song, J. Scott, C. Poole, L. Phuttitarn, K. Jooya, P. Eichler, X. Jiang, A. Marra, B. Grinkemeyer, M. Kwon, M. Ebert, J. Cherek, M. T. Lichtman, M. Gillette, J. Gilbert, D. Bowman, T. Ballance, C. Campbell, E. D. Dahl, O. Crawford, N. S. Blunt, B. Rogers, T. Noel, and M. Saffman. “Multi-qubit entanglement and algorithms on a neutral-atom quantum computer”. In: *Nature* 604 (2022), pp. 457–462.
- [11] S. J. Evered, D. Bluvstein, M. Kalinowski, S. Ebadi, T. Manovitz, H. Zhou, S. H. Li, A. A. Geim, T. T. Wang, N. Maskara, H. Levine, G. Semeghini, M. Greiner, V. Vuletić, and M. D. Lukin. “High-fidelity parallel entangling gates on a neutral-atom quantum computer”. In: *Nature* 622 (2023), p. 268.
- [12] A. G. Radnaev, W. C. Chung, D. C. Cole, D. Mason, T. G. Ballance, M. J. Bedalov, D. A. Belknap, M. R. Berman, M. Blakely, I. L. Bloomfield, P. D. Buttler, C. Campbell, A. Chopinaud, E. Copenhaver, M. K. Dawes, S. Y. Eubanks, A. J. Friss, D. M. Garcia, J. Gilbert, M. Gillette, P. Goiporia, P. Gokhale, J. Goldwin, D. Goodwin, T. M. Graham, C. J. Guttormsson, G. T. Hickman, L. Hurtley, M. Iliev, E. B. Jones, R. A. Jones, K. W. Kuper, T. B. Lewis, M. T. Lichtman, F. Majdeteimouri, J. J. Mason, J. K. McMaster, J. A. Miles, P. T. Mitchell, J. D. Murphree, N. A. Neff-Mallon, T. Oh, V. Omole, C. Parlo

- Simon, N. Pederson, M. A. Perlin, A. Reiter, R. Rines, P. Romlow, A. M. Scott, D. Stiefvater, J. R. Tanner, A. K. Tucker, I. V. Vinogradov, M. L. Warter, M. Yeo, M. Saffman, and T. W. Noel. “A universal neutral-atom quantum computer with individual optical addressing and non-destructive readout”. In: *arXiv:2408.08288* (2024).
- [13] Richard Bing-Shiun Tsai, Xiangkai Sun, Adam L. Shaw, Ran Finkelstein, and Manuel Endres. “Benchmarking and Fidelity Response Theory of High-Fidelity Rydberg Entangling Gates”. In: *PRX Quantum* 6 (2025), p. 010331.
- [14] Michael Peper, Yiyi Li, Daniel Y. Knapp, Mila Bileska, Shuo Ma, Genyue Liu, Pai Peng, Bichen Zhang, Sebastian P. Horvath, Alex P. Burgers, and Jeff D. Thompson. “Spectroscopy and Modeling of ^{171}Yb Rydberg States for High-Fidelity Two-Qubit Gates”. In: *Phys. Rev. X* 15 (2025), p. 011009.
- [15] J. A. Muniz, M. Stone, D. T. Stack, M. Jaffe, J. M. Kindem, L. Wadleigh, E. Zalus-Geller, X. Zhang, C.-A. Chen, M. A. Norcia, J. Epstein, E. Halperin, F. Hummel, T. Wilkason, M. Li, K. Barnes, P. Battaglino, T. C. Bohdanowicz, G. Booth, A. Brown, M. O. Brown, W. B. Cairncross, K. Cassella, R. Coxe, D. Crow, M. Feldkamp, C. Griger, A. Heinz, A. M. W. Jones, H. Kim, J. King, K. Kotru, J. Lauigan, J. Marjanovic, E. Megidish, M. Meredith, M. McDonald, R. Morshead, S. Narayanaswami, C. Nishiguchi, T. Paule, K. A. Pawlak, K. L. Pudenz, D. Rodríguez Pérez, A. Ryou, J. Simon, A. Smull, M. Urbanek, R. J. M. van de Veerdonk, Z. Vendeiro, T.-Y. Wu, X. Xie, and B. J. Bloom. “High-Fidelity Universal Gates in the ^{171}Yb Ground-State Nuclear-Spin Qubit”. In: *PRX Quantum* 6 (2025), p. 020334.
- [16] A. Fuhrmanek, R. Bourgain, Y. R. P. Sortais, and A. Browaeys. “Free-Space Lossless State Detection of a Single Trapped Atom”. In: *Phys. Rev. Lett.* 106 (2011), p. 133003.

- [17] M. J. Gibbons, C. D. Hamley, C.-Y. Shih, and M. S. Chapman. “Nondestructive Fluorescent State Detection of Single Neutral Atom Qubits”. In: *Phys. Rev. Lett.* 106 (2011), p. 133002.
- [18] Margaret E. Shea, Paul M. Baker, James A. Joseph, Jungsang Kim, and Daniel J. Gauthier. “Submillisecond, nondestructive, time-resolved quantum-state readout of a single, trapped neutral atom”. In: *Phys. Rev. A* 102 (2020), p. 053101.
- [19] Matthew N. H. Chow, Bethany J. Little, and Yuan-Yu Jau. “High-fidelity low-loss state detection of alkali-metal atoms in optical tweezer traps”. In: *Phys. Rev. A* 108 (2023), p. 032407.
- [20] M. T. Lichtman. “Coherent operations, entanglement, and progress towards quantum search in a large 2D array of neutral atom qubits”. In: *PhD thesis, University of Wisconsin-Madison* (2015).
- [21] Trent Graham, Xiaoyu Jiang, Cody Poole, Yuan Sun, Martin Lichtman, and Mark Saffman. “An efficient 2D array of blue-detuned optical traps”. In: *Bull. Am. Phys. Soc.* 63 (2018), T01.85.
- [22] T. Graham, M. Kwon, B. Grinkemeyer, A. Marra, X. Jiang, M. Lichtman, Y. Sun, M. Ebert, and M. Saffman. “Rydberg mediated entanglement in a two-dimensional neutral atom qubit array”. In: *Phys. Rev. Lett.* 123 (2019), p. 230501.
- [23] T. M. Graham, L. Phuttitarn, R. Chinnarasu, Y. Song, C. Poole, K. Jooya, J. Scott, A. Scott, P. Eichler, and M. Saffman. “Mid-circuit measurements on a neutral atom quantum processor”. In: *Phys. Rev. X* 13 (2023), p. 041051.
- [24] M. Saffman and T. G. Walker. “Analysis of a quantum logic device based on dipole-dipole interactions of optically trapped Rydberg atoms”. In: *Phys. Rev. A* 72 (2005), p. 022347.
- [25] T. Graham. “Rydberg gate analysis”. In: (*unpublished*) (2020).

- [26] H. Levine, A. Keesling, G. Semeghini, A. Omran, T. T. Wang, S. Ebadi, H. Bernien, M. Greiner, V. Vuletić, H. Pichler, and M. D. Lukin. “Parallel implementation of high-fidelity multiqubit gates with neutral atoms”. In: *Phys. Rev. Lett.* 123 (2019), p. 170503.
- [27] David V. Sheludko, Simon C. Bell, Russell Anderson, Christoph S. Hofmann, Edgar J. D. Vredenburg, and Robert E. Scholten. “State-selective imaging of cold atoms”. In: *Phys. Rev. A* 77 (3 2008), p. 033401. DOI: 10.1103/PhysRevA.77.033401.
- [28] Hamid Ohadi, Matthew Himsworth, André Xuereb, and Tim Freearge. “Magneto-optical trapping and background-free imaging for atoms near nanostructured surfaces”. In: *Opt. Expr.* 17.25 (2009), pp. 23003–23009.
- [29] Baodong Yang, Qiangbing Liang, Jun He, and Junmin Wang. “Background-free fluorescence detection of cold atoms in a two-color magneto-optical trap”. In: *Opt. Expr.* 20 (2012), p. 11944.
- [30] J. P. McGilligan, K. R. Moore, A. Dellis, G. D. Martinez, E. de Clercq, P. F. Griffin, A. S. Arnold, E. Riis, R. Boudot, and J. Kitching. “Laser cooling in a chip-scale platform”. In: *Appl. Phys. Lett.* 117 (2020), p. 054001.
- [31] Li Li, YIJIA L hou XIAOLONG ZH, DONGYU Huang, ZEMIN Shen, SIJIAN He, JIAN Wang, CHUANFENG Li, and GUANGCAN Guo. “Background-free imaging of cold atoms in optical traps”. In: *Opt. Expr.* 32 (2024), p. 21988.
- [32] Shankar G. Menon, Noah Glachman, Matteo Pompili, Alan Dibos, and Hannes Bernien. “An integrated atom array-nanophotonic chip platform with background-free imaging”. In: *Nature Communications* 15 (2024), p. 6156.
- [33] Harold J. Metcalf and Peter van der Straten. *Laser Cooling and Trapping*. Corrected. Springer, 1999. ISBN: 978-0387987286.

- [34] F. Le Kien, P. Schneeweiss, and A. Rauschenbeutel. “Dynamical polarizability of atoms in arbitrary light fields: general theory and application to cesium”. In: *Eur. Phys. J. D* 67.5 (2013), pp. 1–16.
- [35] D. A. Varshalovich, A. N. Moskalev, and V. K. Khersonskii. *Quantum Theory of Angular Momentum*. World Scientific Publishing, 2008.
- [36] J. Scott, H. M. Lim, U. Singla, Q. Meece, J. T. Choy, S. Kolkowitz, T. M. Graham, and M. Saffman. *Laser cooling and qubit measurements on a forbidden transition in neutral Cs atoms*. 2025. URL: <https://arxiv.org/abs/2507.01720>.
- [37] S. Weyers, E. Aucouturier, C. Valentin, and N. Dimarcq. “A continuous beam of cold cesium atoms extracted from a two-dimensional magneto-optical trap”. In: *Optics Communications* 143.1 (1997), pp. 30–34. DOI: [https://doi.org/10.1016/S0030-4018\(97\)00312-X](https://doi.org/10.1016/S0030-4018(97)00312-X).
- [38] K. Dieckmann, R. J. C. Spreeuw, M. Weidemüller, and J. T. M. Walraven. “Two-dimensional magneto-optical trap as a source of slow atoms”. In: *Phys. Rev. A* 58 (5 1998), pp. 3891–3895. DOI: [10.1103/PhysRevA.58.3891](https://doi.org/10.1103/PhysRevA.58.3891).
- [39] P. Huft, Y. Song, T. M. Graham, K. Jooya, S. Deshpande, C. Fang, M. Kats, and M. Saffman. “Simple, passive design for large optical trap arrays for single atoms”. In: *Phys. Rev. A* 105 (2022), p. 063111.
- [40] A. Sharma, S. Kolkowitz, and M. Saffman. “Analysis of a Cesium lattice optical clock”. In: *arXiv:2203.08708* (2022).
- [41] J. Scott, S. Kolkowitz, and M. Saffman. “A Cesium lattice optical clock (CLOC)”. In: *Bull. Am. Phys. Soc., 54th Annual Meeting of the APS Division of Atomic, Molecular and Optical Physics* (2023), F01.00133.

- [42] N. M. Linke, C. J. Ballance, and D. M. Lucas. “Injection locking of two frequency-doubled lasers with 3.2GHz offset for driving Raman transitions with low photon scattering in Ca^{43+} ”. In: *Opt. Lett.* 38.23 (2013), pp. 5087–5089. DOI: 10.1364/OL.38.005087.
- [43] Tatiana Nazarova, Christian Lisdat, Fritz Riehle, and Uwe Sterr. “Low-frequency-noise diode laser for atom interferometry”. In: *J. Opt. Soc. Am. B* 25.10 (2008), pp. 1632–1638. DOI: 10.1364/JOSAB.25.001632.
- [44] V. M. Schäfer, C. J. Ballance, C. J. Tock, and D. M. Lucas. “Optical injection and spectral filtering of high-power ultraviolet laser diodes”. In: *Opt. Lett.* 40.18 (2015), pp. 4265–4268. DOI: 10.1364/OL.40.004265.
- [45] Oleg A. Alduchov and Robert E. Eskridge. “Improved Magnus Form Approximation of Saturation Vapor Pressure”. In: *Journal of Applied Meteorology and Climatology* 35.4 (1996), pp. 601–609. DOI: 10.1175/1520-0450(1996)035<0601:IMFAOS>2.0.CO;2.
- [46] Easchiff. *Dewpoint-RH.svg: Graph of dew point vs. air temperature at varying relative humidities*. <https://commons.wikimedia.org/wiki/File:Dewpoint-RH.svg>. Own work, based on the Magnus formula; licensed under CC BY-SA 4.0. 2008.
- [47] Ziting Chen, Bojeong Seo, Mingchen Huang, Mithilesh K. Parit, Peng Chen, and Gyu-Boong Jo. “Active control of a diode laser with injection locking using a laser line filter”. In: *Review of Scientific Instruments* 92.12 (2021), p. 123005. DOI: 10.1063/5.0057245.
- [48] M Gertszov and M Rosenbluh. “Injection locking of a diode laser locked to a Zeeman frequency stabilized laser oscillator”. In: *Optics Communications* 170.4 (1999), pp. 269–274. DOI: [https://doi.org/10.1016/S0030-4018\(99\)00470-8](https://doi.org/10.1016/S0030-4018(99)00470-8).

- [49] Carl E. Wieman and Leo Hollberg. “Using diode lasers for atomic physics”. In: *Review of Scientific Instruments* 62.1 (1991), pp. 1–20. DOI: 10.1063/1.1142305.
- [50] Juan Camilo Bohorquez. “Quantum Information Processing with Individually Trapped Neutral Atoms”. Local Identifier: Bohorquez_wisc_0262B_19759. Ph.D. thesis. University of Wisconsin–Madison, 2023, p. 264. URL: <https://digital.library.wisc.edu/1711.d1/CTGGVNG6LNGRC9E>.
- [51] A. W. Carr. “Improving Quantum computation with neutral Cesium: readout and cooling on a quadrupole line, conditions for double magic traps and novel dissipative entanglement scheme”. In: *PhD thesis Univ. Wisconsin-Madison* (2014).
- [52] E. Magnan, J. Maslek, C. Bracamontes, A. Restelli, T. Boulier, and J. V. Porto. “A low-steering piezo-driven mirror”. In: *Review of Scientific Instruments* 89.7 (2018), p. 073110. DOI: 10.1063/1.5035326.
- [53] A. A. Michelson. “The Relative Motion of the Earth and the Luminiferous Ether”. In: *Am. J. Science* 22 (1881), p. 120.
- [54] J. W. Cooley and J. Tukey. “An algorithm for the machine calculation of complex Fourier series”. In: *Math. Comp.* 19 (1965), p. 297.
- [55] Alan V. Oppenheim and Ronald W. Schaffer. *Discrete-Time Signal Processing*. 3rd. Pearson, 2010. ISBN: 9780131988422.
- [56] 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”. In: *Appl. Phys. B* 31 (1983), p. 97.
- [57] Lingze Duan and Kurt Gibble. “Locking lasers with large FM noise to high-Q cavities”. In: *Opt. Lett.* 30.24 (2005), pp. 3317–3319. DOI: 10.1364/OL.30.003317.

- [58] Nicolas Minorsky. “Directional Stability of Automatically Steered Bodies”. In: *Journal of the American Society for Naval Engineers* 34.2 (1922), pp. 280–309.
- [59] Karl Johan Åström and Tore Hägglund. *Advanced PID Control*. ISA - The Instrumentation, Systems, and Automation Society, 2006. ISBN: 978-1-55617-942-6.
- [60] K.J. Åström and T. Hägglund. “The future of PID control”. In: *Control Engineering Practice* 9.11 (2001). PID Control, pp. 1163–1175. DOI: [https://doi.org/10.1016/S0967-0661\(01\)00062-4](https://doi.org/10.1016/S0967-0661(01)00062-4).
- [61] A. Schremer and C. L. Tang. “Single-frequency tunable external-cavity semiconductor laser using an electro-optic birefringent modulator”. In: *Applied Physics Letters* 55.1 (1989), pp. 19–21. DOI: 10.1063/1.101977.
- [62] B. Wacogne, J.-P. Goedgebuer, and H. Porte. “Single lithium niobate crystal for mode selection and phase modulation in a tunable extended-cavity laser diode”. In: *Optics Letters* 19.17 (1994), pp. 1334–1336. DOI: 10.1364/OL.19.001334.
- [63] J. Le Gouët, J. Kim, C. Bourassin-Bouchet, M. Lours, A. Landragin, and F. Pereira Dos Santos. “Wide bandwidth phase-locked diode laser with an intra-cavity electro-optic modulator”. In: *Optics Communications* 282.5 (2009), pp. 977–980. DOI: 10.1016/j.optcom.2008.10.042.
- [64] Megan Yamoah, Boris Braverman, Edwin Pedrozo-Penafiel, Akio Kawasaki, Bojan Zlatkovi, and Vladan Vuleti. “Robust kHz-linewidth distributed Bragg reflector laser with optoelectronic feedback”. In: *Opt. Express* 27.26 (2019), pp. 37714–37720. DOI: 10.1364/OE.27.037714.
- [65] Amnon Yariv. *Quantum Electronics*. 3rd. Wiley, 1989. ISBN: 978-0471609979.

- [66] Dylan Renaud, Daniel Rimoli Assumpcao, Graham Joe, Amirhassan Shams-Ansari, Di Zhu, Yaowen Hu, Neil Sinclair, and Marko Loncar. “Sub-1 Volt and high-bandwidth visible to near-infrared electro-optic modulators”. In: *Nature Communications* 14.1 (2023), p. 1496. DOI: 10.1038/s41467-023-36870-w.
- [67] Peter O. Weigel, Jie Zhao, Kelvin Fang, Hasan Al-Rubaye, Douglas Trotter, Dana Hood, John Mudrick, Christina Dallo, Andrew T. Pomerene, Andrew L. Starbuck, Christopher T. DeRose, Anthony L. Lentine, Gabriel Rebeiz, and Shayan Mookherjea. “Bonded thin film lithium niobate modulator on a silicon photonics platform exceeding 100 GHz 3-dB electrical modulation bandwidth”. In: *Opt. Express* 26.18 (2018), pp. 23728–23739. DOI: 10.1364/OE.26.023728.
- [68] Inc. Analog Devices. *AD8429: 1nV/Hz Low-Noise Instrumentation Amplifier - Data Sheet (Rev. A)*. Analog Devices website. Rev. A; accessed July 2025. 2011. URL: <https://www.analog.com/media/en/technical-documentation/data-sheets/ad8429.pdf>.
- [69] Inc. Thorlabs. *EO-PM-NR-C1: Electro-Optic Phase Modulator (600 900nm) - Data Sheet*. Thorlabs website. Accessed July2025; part no. EO-PM-NR-C1, DC-100MHz non-resonant; LiNbO₃, 14pF, 2mm aperture. 2007. URL: <https://www.thorlabs.com/thorproduct.cfm?partnumber=EO-PM-NR-C1>.
- [70] Inc. Analog Devices. *AD825: Low Cost, General-Purpose, High-Speed FET Input Amplifier - Data Sheet (Rev. G)*. Analog Devices website. Rev. G; accessed July2025; 41MHz bandwidth, 125V/ μ s slew, low-noise J-FET input. 1997. URL: <https://www.analog.com/en/products/ad825.html>.
- [71] A.A. Kirpichnikova, O.N. Prudnikov, and D. Wilkowski. “Investigation of the possibility of ultra-deep laser cooling using a quadrupole transition”. In: *Qu. Electr.* 49 (2019), p. 443.

- [72] P. Lett, R. Watts, C. Westbrook, W. D. Phillips, P. Gould, and H. Metcalf. “Observation of Atoms Laser Cooled below the Doppler Limit”. In: *Phys. Rev. Lett.* 61 (1988), p. 169.
- [73] David S. Weiss, Erling Riis, Yaakov Shevy, P. Jeffrey Ungar, and Steven Chu. “Optical molasses and multilevel atoms: experiment”. In: *J. Opt. Soc. Am. B* 6.11 (1989), pp. 2072–2083. DOI: 10.1364/JOSAB.6.002072.
- [74] S. Pradhan and B. N. Jagatap. “Measurement of temperature of laser cooled atoms by one-dimensional expansion in a magneto-optical trap”. In: *Review of Scientific Instruments* 79.1 (2008), p. 013101. DOI: 10.1063/1.2827517.
- [75] J.A. Nelder and R. Mead. “A Simplex Method for Function Minimization”. In: *Computer Journal* 7.4 (1965), pp. 308–313. DOI: 10.1093/comjnl/7.4.308.
- [76] T.-Y. Wu, A. Kumar, F. Giraldo, and D. S. Weiss. “Stern-Gerlach detection of neutral-atom qubits in a state-dependent optical lattice”. In: *Nat. Phys.* 15 (2019), p. 538.
- [77] Karl N. Blodgett, Saumitra S. Phatak, Meng Raymond Chen, David Peana, Claire Pritts, and Jonathan D. Hood. “Narrow-Line Electric Quadrupole Cooling And Background-Free Imaging Of A Single Cs Atom With Spatially Structured Light”. In: *arXiv:2505.10540* (2025).
- [78] G. Reymond, N. Schlosser, I. Protsenko, and P. Grangier. “Single-atom manipulations in a microscopic dipole trap”. In: *Phil. Trans. R. Soc. Lond. A* 361 (2003), p. 1527.
- [79] Y. R. P. Sortais, H. Marion, C. Tuchendler, A. M. Lance, M. Lamare, P. Fournet, C. Armellin, R. Mercier, G. Messin, A. Browaeys, and P. Grangier. “Diffraction-limited optics for single-atom manipulation”. In: *Phys. Rev. A* 75 (1 2007), p. 013406. DOI: 10.1103/PhysRevA.75.013406.

- [80] P. Rosenbusch, S. Ghezali, V. A. Dzuba, V. V. Flambaum, K. Beloy, and A. Derevianko. “ac Stark shift of the Cs microwave atomic clock transitions”. In: *Phys. Rev. A* 79 (2009), p. 013404.
- [81] N. Šibalić, J.D. Pritchard, C.S. Adams, and K.J. Weatherill. “ARC: An open-source library for calculating properties of alkali Rydberg atoms”. In: *Computer Physics Communications* 220 (2017), pp. 319–331. DOI: <https://doi.org/10.1016/j.cpc.2017.06.015>.
- [82] S. Kuhr, W. Alt, D. Schrader, I. Dotsenko, Y. Miroshnychenko, A. Rauschenbeutel, and D. Meschede. “Analysis of dephasing mechanisms in a standing-wave dipole trap”. In: *Phys. Rev. A* 72 (2005), p. 023406.
- [83] D. Schrader, S. Kuhr, W. Alt, M. Müller, V. Gomer, and D. Meschede. “An optical conveyor belt for single neutral atoms”. In: *Appl. Phys. B* 73 (2001), p. 819.
- [84] C. Tuchendler, A. M. Lance, A. Browaeys, Y. R. P. Sortais, and P. Grangier. “Energy distribution and cooling of a single atom in an optical tweezer”. In: *Phys. Rev. A* 78 (2008), p. 033425.
- [85] B. Grinkemeyer, Elmer Guardado-Sanchez, Ivana Dimitrova, Danilo Shchepanovich, G. Eirini Mandopoulou, Johannes Borregaard, Vladan Vuletić, and Mikhail D. Lukin. “Error-detected quantum operations with neutral atoms mediated by an optical cavity”. In: *Science* 387 (2025), p. 1301.
- [86] B. Nikolov, E. Diamond-Hitchcock, J. Bass, N. L. R. Spong, and J. D. Pritchard. “Randomized Benchmarking Using Nondestructive Readout in a Two-Dimensional Atom Array”. In: *Phys. Rev. Lett.* 131 (2023), p. 030602.
- [87] Pauli Virtanen et al. “SciPy 1.0: fundamental algorithms for scientific computing in Python”. In: *Nature Methods* 17.3 (2020), pp. 261–272. DOI: [10.1038/s41592-019-0686-2](https://doi.org/10.1038/s41592-019-0686-2).

- [88] Alessandro Muzi Falconi, Riccardo Panza, Sara Sbernardori, Riccardo Forti, Ralf Klemt, Omar Abdel Karim, Matteo Marinelli, and Francesco Scazza. *Microsecond-scale high-survival and number-resolved detection of ytterbium atom arrays*. 2025. URL: <https://arxiv.org/abs/2507.01011>.
- [89] David N. Nikogosyan. *Nonlinear optical crystals: a complete survey*. Vol. 64. Springer, 1991, pp. 130–140.
- [90] Marco Pizzocaro, Davide Calonico, Pablo Cancio Pastor, Jacopo Catani, Giovanni A. Costanzo, Filippo Levi, and Luca Lorini. “Efficient frequency doubling at 399  nm”. In: *Appl. Opt.* 53.16 (2014), pp. 3388–3392. DOI: 10.1364/AO.53.003388.
- [91] T. W. Hänsch and B. Couillaud. “Laser frequency stabilization by polarization spectroscopy of a reflecting reference cavity”. In: *Optics Communications* 35.3 (1980), pp. 441–444. DOI: 10.1016/0030-4018(80)90091-3.
- [92] T. Okoshi, K. Kikuchi, and A. Nakayama. “Novel method for high resolution measurement of laser output spectrum”. In: *El. Lett.* 16 (1980), p. 630.
- [93] X. Jiang, J. Scott, M. Friesen, and M. Saffman. “Sensitivity of quantum gate fidelity to laser phase and intensity noise”. In: *Phys. Rev. A* 107 (2023), p. 042611.
- [94] Fredric J. Harris. “On the use of windows for harmonic analysis with the discrete Fourier transform”. In: *Proceedings of the IEEE* 66.1 (1978), pp. 51–83. DOI: 10.1109/PROC.1978.10837.
- [95] Evgeni Sorokin. “Solid-State Materials for Few-Cycle Pulse Generation and Amplification”. In: vol. 95. Topics in applied physics, 2004, pp. 3–73. ISBN: 978-3-540-20115-1. DOI: 10.1007/978-3-540-39849-3_1.

- [96] S. Saskin, J. T. Wilson, B. Grinkemeyer, and J. D. Thompson. “Narrow-Line Cooling and Imaging of Ytterbium Atoms in an Optical Tweezer Array”. In: *Phys. Rev. Lett.* 122 (2019), p. 143002.
- [97] L. Su, A. Douglas, M. Szurek, A. H. Hébert, A. Krahn, R. Groth, G. A. Phelps, O. Marković, and M. Greiner. “Fast single atom imaging for optical lattice arrays”. In: *Nat. Commun.* 16 (2025), p. 1017.
- [98] Jian Wang, Dong-Yu Huang, Xiao-Long Zhou, Ze-Min Shen, Si-Jian He, Qi-Yang Huang, Yi-Jia Liu, Chuan-Feng Li, and Guang-Can Guo. *Ultrafast high-fidelity state readout of single neutral atom*. 2024. URL: <https://arxiv.org/abs/2412.12584>.