

**New tools for precision measurement and quantum science
with narrow linewidth optical transitions**

by

Matthew A. Norcia

B.S., Stanford University, 2011

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

2017

ProQuest Number: 10634089

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



ProQuest 10634089

Published by ProQuest LLC (2017). Copyright of the Dissertation is held by the Author.

All rights reserved.

This work is protected against unauthorized copying under Title 17, United States Code
Microform Edition © ProQuest LLC.

ProQuest LLC.
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106 – 1346

This thesis entitled:
New tools for precision measurement and quantum science with narrow linewidth optical
transitions
written by Matthew A. Norcia
has been approved for the Department of Physics

Prof. James K. Thompson

Prof. Jun Ye

Date _____

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Norcia, Matthew A. (Ph.D., Physics)

New tools for precision measurement and quantum science with narrow linewidth optical transitions

Thesis directed by Prof. James K. Thompson

In my thesis work, I have explored a variety of ways in which narrow linewidth optical transitions can be utilized to extend the capabilities of precision measurement and quantum science. These tools include non-destructive atom counting techniques that may enhance the performance of optical lattice clocks, a novel form of cavity enhanced atomic spectroscopy appropriate for laser frequency stabilization, a new form of laser cooling with reduced reliance on spontaneous emission, several methods for enhancing the sensitivity of proposed atomic gravitational wave detectors, a newly observed form of cavity-mediated spin-spin interactions, and a new class of laser based on optical superradiance from narrow and ultra-narrow linewidth transitions. As part of the superradiant laser project, we have demonstrated the most precise active absolute frequency reference realized to date, which still has great potential for improvement.

Many of these tools involve coupling atoms to an optical cavity, a field known as cavity quantum electrodynamics, or cavity QED. In my thesis work, a key focus was the extension of cavity QED techniques into the new regime of coupling large ensembles of atoms to a cavity using narrow linewidth optical transitions (in contrast to the previously explored regimes of microwave-frequency transitions or broad-linewidth optical transitions). To this end, I constructed a new experiment to couple an ensemble of atoms to an optical cavity using two optical transitions in atomic strontium, one with a narrow 7.5 kHz linewidth, and another with an ultra-narrow 1 mHz linewidth. This new system has allowed us to access a unique new regime of cavity QED, and a wealth of interesting and useful applications.

Dedication

To my parents.

Acknowledgements

Here I am, about to graduate the 22nd grade and looking for ways to stick around. For this, I have many people to thank.

First, I would like to thank my parents Pat and Tony. They have always been extremely encouraging and tolerant of my projects, from taking up much of the street and yard with skateboard ramps, to the five or so project cars that cycled through the driveway, to the ultimate entire takeover of the back yard for the construction of an electric vehicle. I would also like to thank my brother Michael (who has currently taken over my parents house as well as the yard for airplane construction) for undertaking these technical adventures with me. These projects provided a great background for my later transition into engineering and physics. My high-school auto-shop teacher Mike Camicia had a lot to do with this as well.

In undergrad, I experienced just enough of physics research after my Freshman year to realize that I was not mentally mature enough for it yet, so I focused my energies on engineering projects with more immediate payoffs. Fortunately, an Engineering Physics degree had recently been introduced (thanks, Pat), so I could do this while leaving open the option of a physics career later on (once my frontal lobes had developed, as Tony would say).

That brings me to grad school where I was lucky enough to join James Thompson's group. The combination of physics, engineering and emphasis on simple visual explanations have made this a far better choice than I could have known. James has been a terrific advisor, with excellent physical intuition and a good sense of priorities both in science and in general.

During my time in the group, I have had the opportunity to work with some fantastic students,

including Zilong Chen, Justin Bohnet, Josh Weiner, Kevin Cox, Graham Greve and Baochen Wu on the Rubidium experiment, and Karl Mayer, Drew Morrill, Julia Cline, Matthew Winchester and postdoc Juan Muniz on the Strontium experiment. Thank you all for the years of great discussions.

Finally, I would like to thank my friends and housemates here in Boulder who have made the last five years here in Boulder a time that I will look back on with great fondness. Perhaps a new high-point?

Contents

Chapter

1	Introduction	1
1.1	Overview of Applications	5
1.1.1	Appliction to Clocks: Non-Destructive Atom Counting	5
1.1.2	Appliction to Laser Stabilization: Magnetically Induced Transparency	6
1.1.3	Superradiance from Narrow Linewidth Optical Transitions	7
1.1.4	Cavity-Mediated Spin-Spin Interactions	11
1.1.5	A Novel Cooling Mechanism	12
1.1.6	Enhanced Gravitational Wave Detection with Atoms	13
1.2	Outline of Thesis	14
2	An Introduction to Atoms in Cavities and Narrow-Linewidth Transitions	15
2.1	Narrow-Linewidth Transitions in Strontium	15
2.2	Atoms Interacting with a Cavity	16
2.3	Cavity QED on Dipole-Forbidden Transitions	23
2.4	Cavity QED with Bloch Spheres and Pendula	26
3	Experimental Apparatus	31
3.1	Experimental Overview	32
3.2	Vacuum System	34
3.3	Atomic Source	36

3.4	MOT Coils and Driver	36
3.5	Laser Systems and Laser Cooling	39
3.5.1	461 nm Laser System	40
3.5.2	Initial Slowing and Cooling	43
3.5.3	Repump Lasers at 707 and 679 nm	45
3.5.4	689 nm System and Narrow-Line Cooling	46
3.5.5	698 nm Laser System	48
3.5.6	Reference Cavity	50
3.6	The Optical Lattice	51
3.7	The Optical Cavity	55
4	Strong Collective Coupling on 7.5 kHz Linewidth Transition and Nondestructive Atom Counting	60
4.1	Demonstration of Collective Strong Coupling	62
4.2	Nondestructive Atom Counting	65
4.3	Conclusions	70
5	Magnetically Induced Transparency	71
5.1	Basic Mechanism	72
5.2	Experimental Characterizations of MIT	77
5.3	Mapping to EIT	81
5.4	Limits on Photon Retrieval Efficiency	82
5.5	Conclusion	84
6	Lasing on the 7.5 kHz Linewidth Transition: the Superradiant Crossover Regime	85
6.1	Pulsed Lasing in the Vacuum Rabi Splitting Limit	87
6.2	Quasi Steady State Operation	88
6.3	Lasing Threshold Behavior	91

6.4	Cavity Pulling, and Tuning Through the Superradiant Crossover Regime	92
7	Pulsed Lasing on the milliHertz Linewidth Clock Transition	95
7.1	Experimental System	95
7.2	Observation of Superradiant Pulses	97
7.3	Seeding Atomic Coherence	102
7.4	Simultaneous Lasing on Multiple Transitions	104
7.5	Simulations of Superradiance	106
7.5.1	Effects of Atomic Decoherence and Loss	107
7.5.2	Effects of Inhomogeneous Coupling	108
8	Frequency Stability Characterization of Pulsed Lasing	111
8.1	Introduction	111
8.2	Basic Techniques	112
8.3	Stability Characterization	115
8.4	Cavity Pulling	118
8.5	Magnetic Field Insensitivity	120
8.6	Absolute Accuracy and Other Perturbations	121
8.6.1	Lattice Shifts	122
8.6.2	Atomic Collisions	125
8.6.3	Other Shifts	126
9	Cavity-Mediated Spin-Spin Interactions	130
9.1	Quantum Description	131
9.2	Classical Description of OAT	134
9.3	Experimental Characterization of OAT Term	138
9.4	The χJ^2 Term	140
9.5	Experimental Characterization of Effects from the χJ^2 Term	142

9.6	Conclusion	150
10	Sawtooth-Wave Adiabatic Passage (SWAP) Cooling	151
10.1	How it Works	152
10.1.1	Adiabatic Passage	152
10.1.2	Basic Cooling Mechanism in One Dimension	154
10.1.3	Where Does the Entropy Go?	156
10.1.4	A Dressed-State Picture	157
10.1.5	Regimes of Applicability	159
10.1.6	Simulated Trajectories	161
10.2	Experimental Demonstration	162
10.2.1	Forces in 1D	162
10.2.2	Cooling and Phase-Space Compression in 1D	167
10.2.3	Practical Advantages	169
10.3	Creating a Restoring Force	170
10.3.1	How it Works	171
10.3.2	Demonstration in One Dimension	173
10.3.3	Application to a Real Experiment	173
10.4	Extensions and Applications	176
11	A new perspective on gravitational wave detection with cold atoms in space	178
11.1	Optical Path Length Changes due to Gravitational Waves	180
11.2	Detecting Changes in Optical Path Length	182
11.3	Clock-like Detectors	185
11.4	Atom Interferometer Detectors	186
11.5	Enhanced sensitivity analogous to large-momentum-transfer	186
11.6	Evolution Times Beyond the Excited State Lifetime	190
11.7	Conclusions	193

11.8 Accounting for Time Delays	195
12 Conclusion and future outlook	197
12.1 Spin squeezing in an optical lattice clock	197
12.2 Towards a steady-state superradiant laser	199
Bibliography	203

Tables

Table

2.1	Key parameters and scalings for cavity QED.	20
3.1	Overview of laser systems.	40
3.2	Cooling and trapping beam parameters.	45
3.3	Typical lattice parameters.	52
3.4	Key parameters of optical cavity.	58
7.1	Summary of key parameters and effects of inhomogeneous coupling for pulsed super-radiance.	109
8.1	Summary of quantum efficiencies in detection chain.	118
8.2	Summary of shifts and uncertainties in frequency comparison.	123

Figures

Figure

- 1.1 (a) Strontium has two valence electrons, whose spins can either be aligned (spin triplet) or anti-aligned (spin singlet). Photons are very bad at flipping spins, so we call transitions between these two types of states “forbidden,” meaning that they occur slowly if at all. (b) A basic overview of our experimental platform. Strontium atoms are cooled to roughly $10\text{ }\mu\text{K}$, and tightly confined in an optical lattice inside of an optical cavity. The wavelength of the lattice is carefully chosen so as not to perturb atomic transition frequencies of interest — the “magic wavelength.” The atoms interact with the cavity via one of several forbidden transitions. 3
- 2.1 A cavity QED system. Atoms exchange excitations with the cavity mode at a rate $\Omega = 2g\sqrt{N}$. Atoms can decay into free-space via single-particle spontaneous emission at a rate γ , leading to loss of coherence. Photons in the cavity leak out through the mirrors at rate κ . This channel of dissipation is easy to detect, and provides only collective information about the state of the atomic ensemble. 17
- 2.2 One of the key manifestations of the strong coupling regime defined by $\Omega \gg \kappa, \gamma$ is the vacuum Rabi splitting. As a weak probe tone coupled to the cavity is swept over resonance with the cavity (here, $\delta_c = 0$), a single transmission peak splits into two. The splitting between these peaks is set by Ω , and the width of each peak is a weighted average of atomic and cavity decay rates. 22

- 2.3 (a) The state of a two-level system with states $|e\rangle$ and $|g\rangle$ can be represented by a Bloch vector that lives on a Bloch sphere. The figure shows a general state $|\psi\rangle = \sin(\theta/2)|g\rangle + \cos(\theta/2)e^{i\phi}|e\rangle$. The vertical projection of the Bloch vector tells us the population inversion of the state (the relative likelihood of finding the atom in the ground or excited state), while the azimuthal angle ϕ tells us the relative phase between the $|e\rangle$ and $|g\rangle$ components. (b) An ensemble of many two-level systems can be represented on a collective Bloch sphere by simply adding up all of the individual Bloch vectors associated with the two-level systems to form a collective Bloch vector \vec{J} . The inversion of the collective Bloch vector is labelled \vec{J}_z , and the atomic coherence (projection onto the equatorial plane) is labelled \vec{J}^+ 26
- 2.4 We can describe the dynamic of an ensemble of atoms interacting with the cavity field using a Bloch sphere picture. The cavity field is represented by an axis \check{C} about which the Bloch vector rotates at a rate proportional to $g|\check{C}|$. As pictured, the rotation is leading to collective decay of the atoms from $|e\rangle$ to $|g\rangle$ 27
- 2.5 With the substitutions given here and in the text, the dynamics of an ensemble of atoms interacting with the mode of a resonant cavity can be mapped onto the motion of a damped pendulum under gravity. The angular velocity of the pendulum takes the role of the cavity field while the decay of photons from the cavity is equivalent to viscous damping of the motion of the pendulum. The length of the pendulum scales as $(g^2J)^{-1}$, which leads to faster dynamics for many atoms or large coupling g . . . 29

- 3.1 A simplified atom loading sequence. Atoms begin as a roughly 500 C collimated beam exiting the oven. Some of the atoms are slowed to around 50 m/s by a Zeeman slower, then redirected by two 2D MOTs into the main experimental chamber. They are then further cooled and spatially compressed by a 3D MOT, which we typically load for 500 ms. The Zeeman slower and both MOT stages also use 461 nm (blue) light. The atoms are then transferred from the blue MOT to a narrow-line red MOT over roughly 50 ms. The magnetic field gradient (represented by the slope of the black line) is reduced by a factor of roughly ten during this step. An asymmetric sawtooth frequency sweep applied to the red MOT beams greatly improves the efficiency of this transfer (See Chapter 10). Atoms are further cooled, and loaded into the optical lattice with a typical temperature of around 10 μ K. 32
- 3.2 Pictures of the experiment from July 2017, to give a sense of scale. Upper left: four-bore reference cavity, with associated optics. Upper right: main vacuum chamber end of experiment. Lower left: the experiment spans two optics tables. Vacuum system is in the foreground, with lasers and laser stabilization in background. Lower right: enclosure houses all red and infrared lasers. 34
- 3.3 The vacuum system consists of three main sections: the atom source, which contains an oven, Zeeman slower, and two 2D MOTs, the science chamber, which contains the optical cavity, and the pumping appendage. The system is pumped by a 75 liter/second ion pump and a titanium sublimation pump. 35
- 3.4 Closeup of MOT and bias coils. MOT coils (large circular coils) are wound from hollow square wire through which water flows. The vertical bias coils are wound onto the MOT coils for support and cooling. The two sets of horizontal bias coils are wound around square frames. Additional freely wound coil in the foreground creates an additional gradient along the cavity axis, but is not typically used. 37

3.5	Left: diagram of circuit used to control current through MOT coils. Right: actual circuit used to control current through MOT coils. All high-power circuit elements are mounted to a water-cooled plate. The MOSFET is the large white component in the top of the image. The shunt resistor is the large black component below it, and the diodes are below that.	38
3.6	Relevant energy levels of strontium. Level structure of ^{88}Sr , which lacks nuclear spin and therefore hyperfine structure, are denoted with thick lines. When present, additional states resulting from the hyperfine structure of ^{87}Sr are shown with thin lines. Zeeman sublevels are neglected.	41
3.7	The 461 nm laser system consists of two master diode lasers and three slave diodes. The frequencies of the two diode lasers are stabilized to a polarization-rotation spectroscopy setup, and by a beatnote lock between the two lasers. The three slave diodes provide power for the Zeeman slower, 2D MOTs, and 3D MOTs respectively.	42
3.8	Closeup of AOSense source, with optics for Zeeman slower (far left) and 2D MOTs. The thin tube with bellows visible in the lower left corner feeds atoms into the main vacuum chamber.	44
3.9	689 nm laser system. An interference-filter ECDL from AOSense provides the master light for our 689 nm laser system. The laser is stabilized to our stable reference cavity by a Pound-Drever-Hall lock on a sideband that provides a tunable offset. Power for the MOT beams is increased using a tapered amplifier (TA). An AOM both switches the cooling light on and off, and allows us to sweep its frequency in a sawtooth ramp for efficient cooling. The 689 nm laser also provides optical pumping for state preparation, among other functions.	47

- 3.10 Fiber-noise cancellation system for clock reference light. This architecture involves two AOMs — one in the Ye lab (AOM1) and the other in the Thompson lab (AOM2). Light that is deflected by AOM2 (giving it a frequency offset f) is reflected off the back mirror of the science cavity, and sent back to the Ye lab. This light is overlapped with light that has not gone through the long fiber, and detected on a fast photodiode. The resulting beatnote (near $2f$ in frequency) is used to lock the frequency of a voltage controlled oscillator (VCO). The output of the VCO is divided by two and used to drive AOM1 at a frequency near f . Any noise present in the fiber leads to a perturbation of the frequency of the beatnote. Because we use opposite diffraction orders on the two AOMs, this frequency offset is cancelled by AOM1. Compared to simpler configurations that use a single AOM, this configuration has the advantage of being immune to unwanted reflections before AOM2. 49
- 3.11 Transfer cavity from Stable Lasers Systems with four sets of mirrors sharing the same spacer. 50
- 3.12 Optical cavity assembly and vibration isolation. The cavity consists of two mirrors separated by 4 cm. The axis of the cavity is horizontal. The mirrors are mounted to piezoelectric actuators (PZTs), which allow for control of the cavity length. To decouple the cavity from acoustic vibrations, the cavity spacer sits on viton balls on top of a support block, which also sits on viton balls. 56

3.13	Mirror assembly. The cavity mirrors are mounted to piezo actuator tubes (PZTs) by an intermediate macor hat. This configuration allowed us to recess the mirrors into the cavity spacer in order to avoid coating them with strontium, and provided an easy way to align the mirrors. The two holes at either end of the cavity spacer in which the mirrors sit were drilled at the same time, ensuring that they share a common axis. All other alignments are referenced to those holes. All connections are made by very thin beads of Torrseal epoxy. Electrical connections to the inner and outer surfaces of the PZT tubes are made by custom clips (not shown) that attach through a slit in the macor hat (bottom, in image).	57
4.1	(a) Experimental diagram. ^{88}Sr atoms are confined by 1D optical lattice in high-finesse cavity. π polarized probe light is coupled through cavity and is detected on a single-photon counting module. This probe interacts only with the $^1\text{S}_0$ to $^3\text{P}_1$, $m_J = 0$ transition. A magnetic field B shifts the $m_J = \pm 1$ Zeeman sub-levels by ± 15 MHz. (b) Relevant level structure of ^{88}Sr , showing dipole forbidden transitions at 689 nm and 698 nm, and dipole allowed transition at 461 nm.	61
4.2	Observation of the collective strong coupling regime on the forbidden $^1\text{S}_0$ to $^3\text{P}_1$ transition. (a) We characterize the vacuum Rabi Splitting by recording transmitted power as probe frequency is swept. Traces correspond to different cavity detunings, δ_c . (b) Detected fluorescence versus vacuum Rabi splitting Ω . (c) Fitted linewidths κ'_\pm of the two resonances at ω_+ (red) and ω_- (blue) versus δ_c , with prediction based on known atomic linewidth γ , measured Ω and cavity linewidth κ	64
4.3	(a) To probe the Vacuum Rabi Splitting, two probe tones at $\omega_{p\pm}$ are simultaneously swept across the two normal mode resonances at ω_\pm while total transmitted power is recorded. (b) The noise power spectra relative to photon shot noise for cavity probe with a single tone (red trace) and two tones (black trace), demonstrate noise cancellation of two-tone probing.	67

- 4.4 (a) Black points show measurement noise relative to hypothetical projection noise in adjacent windows R , versus length of a measurement window. The lowest two points represent $R = -2.4_{1.1}^{0.9}$ dB. Red line shows the the number of scattered photons per atom m_s in a single measurement window. (b) Fractional change in $\Omega^2 = 4g^2N$ due to additional sweep $\epsilon \equiv (\Omega_0^2 - \Omega_s^2)/\Omega_0^2$, versus M_t , the number of photons in the additional sweep. Ω_s and Ω_0 are the measured vacuum Rabi splittings with and without the additional scattering sweep. Solid red line is fit to data, with shaded region representing 1 sigma error on the fit. Dashed line represents a prediction based on the theoretically predicted free space scattering. The number of photons used in two typical 100 μ s windows is represented by the thickness of the left axis. . 69
- 5.1 (a) Simplified experimental diagram. The system is probed with horizontally polarized probe light. The light can be coherently absorbed by the atoms (brown ovals) and reemitted into the cavity at collective vacuum Rabi frequency Ω . The transmitted power is detected on a single photon counting module (SPCM). A magnetic field is applied along the vertical direction. (b) The atomic energy level diagram of the ground 1S_0 and excited states $^3P_1 |m_j\rangle$ with quantization axis \hat{q} parallel to the applied magnetic field \vec{B} . This is the viewpoint adopted in this letter. Here, the applied magnetic field creates a Zeeman splitting Δ between the excited states $|\pm 1\rangle$. Both of these transitions interact equally with the horizontally polarized light inside the cavity with collectively enhanced Rabi frequency $\Omega/\sqrt{2}$. The transmitted probe light is measured as the probe's detuning δ_p is swept. The $|0\rangle$ state is shown, but it does not interact with the horizontally polarized cavity-field. 75

- 5.2 (a) The transmitted power through the cavity versus the probe detuning δ_p , with $\delta_c = 0$. Each trace was taken for different applied magnetic fields, creating different Zeeman splittings Δ labeled on the vertical. The central red trace is taken for $\Delta = 0$ and displays a collective vacuum Rabi splitting $\Omega/2\pi = 5$ MHz. When a magnetic field is applied perpendicular to the probe polarization, inducing a Zeeman splitting Δ , a new transmission feature appears in between the two original resonances of the vacuum Rabi splitting. (b) Linearized theory showing the power P_T and phase ψ of the transmitted light, plotted here for $\Omega/2\pi = 5$ MHz and $\Delta/2\pi = 1$ MHz. 76
- 5.3 (a) The measured linewidth of the central MIT transmission feature versus the induced Zeeman splitting between excited states. The traces are taken for three different collective vacuum Rabi frequencies $\Omega/2\pi = 4.6(5)$ (red), $10(1)$ (blue), and $16(1)$ (green) MHz, with values set by changing the total atom number N . The upper dashed line is the empty cavity's linewidth κ , and the lower dashed line is the atomic transition's linewidth γ . The minimum observed linewidth was 11 kHz. The shaded regions are no-free parameter predictions from the linearized model introduced in the text, indicating the ± 1 standard deviation uncertainty bands based on independent measurements of Ω . (b) The measured peak transmitted power of the central MIT transmission feature for the same collective Rabi frequencies. Here, the transmitted power is normalized to the peak transmitted power when the cavity is empty. Again the shaded regions indicate the ± 1 standard deviation uncertainty bands for the predictions. 78
- 5.4 The pulling coefficient P versus Zeeman splitting Δ for several collective vacuum Rabi frequencies $\Omega/2\pi = 5(1)$ (red), $10(1)$ (blue), and $17(1)$ (green) MHz. The prediction from the linearized theory is shown with ± 1 standard deviation bands. . . 80

- 5.5 In the ‘EIT picture’, the quantization axis is perpendicular the magnetic field and and parallel to the probe light’s polarization. (left) The cavity now couples the ground and excited Zeeman state $|m_{j\perp} = 0\rangle$. The magnetic field couples $|0_{\perp}\rangle$ to neighboring excited Zeeman states $|m_{j\perp} = \pm 1\rangle$. (center) Equivalently, the magnetic field couples to a symmetric combination of Zeeman states $|S\rangle$ while the antisymmetric combination $|D\rangle$ does not participate. (right) For comparison, a traditional Raman EIT system is shown for comparison with the equivalent coupling strengths labeled. 82
- 6.1 Energy level and experimental diagrams. The lasing transition is the 7.5 kHz wide $^3P_1, m_J = 0$ to 1S_0 transition in ^{88}Sr . A magnetic field B is oriented perpendicular to the cavity axis, and pump lasers at 688 and 689 nm are applied with directions and polarizations as shown. Collectively enhanced emission into the TEM00 mode of the cavity is collected in an optical fiber and sent to one of various detectors. . . 86
- 6.2 (a) Laser power emitted from the cavity during a collectively enhanced pulse, with $N = 36\text{k}$ atoms. The atoms are optically pumped into $|e\rangle$ at $t = 0$. (b) Numerical simulation of (a) with no free parameters. (c) The power emitted from the cavity during continuously repumped lasing with $N = 60\text{k}$ atoms. Each atom emits around 35 photons into the cavity before repump-induced heating causes laser operation to cease. For scale, 1 nW corresponds to 3500 intracavity photons. (d) Repumping scheme. Atoms are incoherently repumped through $^3P_1, m_J = -1$ to 3S_1 , where they decay into the 3P manifold. Atoms that fall into 3P_0 or 3P_2 are then repumped by additional lasers via 3S_1 . For display purposes, $^3P_{0,2}$ and 1P_1 states are shown at arbitrary vertical positions. 87

6.3	Threshold behavior of laser. Laser output power in quasi steady state is plotted versus repumping rate w for no added broadening $\gamma'_{\perp} = 0$, (red points) and added broadening, ($2\gamma'_{\perp} \simeq 2\pi \times 3$ MHz $\gg \gamma, \kappa$, black points). Atomic decay rate γ is displayed as vertical blue line, with shaded blue region representing uncertainty in calibration of w . For both conditions, the measured threshold repumping rate w_t is consistent with $w_t = \gamma$	90
6.4	(a) Pulling coefficient P versus the total homogeneous broadening of the lasing transition, along with prediction (blue). The vertical line at 160 kHz represents the cavity decay rate κ . The inset displays the left five points on a linear scale. (b) Averaged heterodyne power spectral density (PSD) between output light and 689 nm pump laser, with recentering of each individual trial before averaging. Lorentzian (Gaussian) fits, shown as red (blue) lines indicate FWHM linewidths of 6.0(3) (4.7(3)) kHz.	91

- 7.1 (a) The atoms can be represented by a collective Bloch vector on a Bloch sphere. The cavity field $C \propto \sqrt{M_c}$, where M_c is the number of photons in the cavity, is represented by a vector that lies in the equatorial plane perpendicular to the atomic coherence J_\perp (when the cavity is on resonance with the atomic transition). The Bloch vector rotates about the axis that represents the cavity field. (b) The Bloch vector behaves like a highly damped pendulum that starts inverted at the north pole of the Bloch sphere (all atoms in $|e\rangle$). Quantum fluctuations disturb the system from its unstable equilibrium position, causing the Bloch vector to swing down the Bloch sphere, emitting peak radiation at the equator and ultimately relaxing to the south pole (all atoms in $|g\rangle$) as inversion is lost. The radiated electric field (red trace) is proportional to the perpendicular projection of the Bloch vector, J_\perp , which at its peak is proportional to N . The radiated power (black trace) is proportional to the square of the radiated electric field, and at its peak is therefore proportional to N^2 . This is one way to understand the origin of the collective enhancement in emission rate. The black output power trace on the projection plane is actual data. 96
- 7.2 Spontaneously generated superradiant pulses. Representative single time traces of photon output rate R for pulses at different atom number $N \approx 100 \times 10^3$ (green), 125×10^3 (blue), 150×10^3 (red), 200×10^3 (black). The equivalent average intracavity photon number is calculated on the right as $M_c = R/\kappa$ 98

- 7.3 (a) Peak photon output rate, R_{peak} , versus initial total atom number. The black line is a quadratic fit to the output power. We observe a horizontal offset in the fit, indicating a threshold atom number N_t (black vertical line in b,c,d). The threshold results from decoherence and atom loss and is a signature of lasing that would not be present for single-atom emission. (b) FWHM pulse duration t_w (blue) and delay of peak power t_d (red) versus initial total atom number. The blue line is a fit to the predicted functional form for the pulse duration, with N_t determined from the fit to R_{peak} . The dashed red line is a fit to the pulse delay, assuming threshold is due to atomic homogeneous broadening without atom loss. The solid red line is a fit to the pulse delay, assuming the threshold is set only by atom loss from the lattice. (c) The ratio of emitted photons M_{tot} to the number of atoms in excess of threshold N_x , plotted versus atom number. The dashed line is the average ratio, showing that 48% of the atoms in excess of threshold participate in the superradiant pulse, largely independent of N 100
- 7.4 Seeded superradiant pulses. (a) By terminating adiabatic transfer with atoms in a superposition of $|e\rangle$ and $|g\rangle$, we seed coherence in the atomic ensemble. Here, the atomic Bloch vector is rotated to just above the equator (i.e. a small amount of initial inversion.) Seeding leads to the immediate onset of superradiant emission, in contrast to the non-seeded pulses shown in Fig. 7.2 for which quantum noise seeds the coherence. (b) Peak output power for seeded pulses exhibits N^2 scaling. In contrast to results of Figure 2, seeded pulses exhibit a peak photon output rate consistent with no threshold: $N_t = 0$ 103

- 7.5 Lasing on multiple transitions at the same time leads to beating in output power
- (a) Atoms prepared in the $m_f = 9/2$ and $m_f = 7/2$ Zeeman sub-levels radiate simultaneously into the cavity. (b) Interference between the electric fields radiated from the two transitions leads to a modulation of the output power. An average of 20 time-traces is plotted, showing that the modulation has the same phase between trials of the experiment. (c) We compute a Fourier transform of an averaged time trace, showing a peak at the frequency splitting of adjacent Zeeman sub-levels (≈ 200 Hz). The peak at ≈ 400 Hz indicates that a smaller number of atoms have been prepared in $m_f = 5/2$. A Lorentzian fit to the average power spectrum (red line) returns a FWHM linewidth of 11 Hz, due to the finite length of the pulse. The inset of (b) shows the center frequency of the Lorentzian fit versus applied magnetic field. The blue line shows the predicted splitting between adjacent Zeeman sub-levels. . . 104
- 8.1 Superradiant pulses in the time and frequency domains for different configurations of spin states. (a) We can populate a single nuclear spin state to observe lasing on a single transition. In the time domain (center column) this leads to a simple superradiant pulse. In the frequency domain (right column), we see a single peak. (b) When two nuclear spin states are populated, we observe a beating of the two transitions in the time domain, and two peaks in the frequency domain. (c) When we populate all ten nuclear spin states, we observe an interesting interference pattern in the time domain from simultaneous lasing on the ten transitions. In the frequency domain, we observe peaks associated with the different transitions, with heights determined by the Clebsch Gordan coefficients of the different transitions. . . . 113

- 8.2 Stability of superradiant pulses expressed as a fractional Allan deviation $\sigma(\tau)$ (see text for definition). Two stretched ($m_f = \pm 9/2$) transitions lase simultaneously. The black points correspond to the average frequency between the two transitions, while the red points correspond to half the difference. Lines represent fits of the small- τ points to $\sigma(\tau) \propto 1/\sqrt{\tau}$, the expected scaling for uncorrelated noise between measurements. 116
- 8.3 Sensitivity of laser output frequency ω_ℓ to changes in the cavity detuning Δ_c for different initial atomic inversions. Red points correspond to an inversion empirically tuned to minimize sensitivity to cavity detuning. Black and blue points correspond to slightly higher and lower inversions. The slope of the red points is 2×10^{-6} 119
- 8.4 A magnetic-field insensitive configuration. The frequencies of the $m_f = 9/2$ and $m_f = -9/2$ transitions (labelled f_{upper} and f_{lower} here) respond in equal and opposite manners to an applied magnetic field, making the average frequency of the two peaks highly insensitive to magnetic fields. 120
- 8.5 Frequency difference between our laser and Ye lab optical lattice clock frequency at different lattice laser detunings and power levels. Different colors of points correspond to different lattice depths. A simultaneous fit (see text) to all data returns a nominal frequency offset consistent with zero, as represented by the horizontal dashed line, with the uncertainty of the fit represented by the blue band. The vertical dashed line represents the fitted magic wavelength (relative to the expected value based on our wavemeter) with fit uncertainty given by the green band. 124
- 8.6 Measured shift to superradiant output frequency associated with density-dependent collisions versus fraction of atoms prepared in 3P_0 at beginning of superradiant pulse. Data taken with roughly 5×10^5 atoms in lattice at beginning of superradiance. . . 126

- 8.7 Calibration of DC Stark shifts. Measured frequency is plotted against voltage applied to electrode on large window of vacuum chamber. Frequency shifts due to background fields in the direction of the applied field from the electrode manifest as an offset between the minimum of the parabola and its value at zero electrode voltage. Shifts here are consistent with zero. 128
- 9.1 (a) The cavity field mediates interactions between the effective spins corresponding to the optical $^1\text{S}_0(g)$ - $^3\text{P}_0(e)$ clock transition. This leads to both a collectively enhanced decay at rate Γ which leads to light leaking out of the cavity, and to interesting unitary dynamics that lead to shifts in the frequency of the emitted light ω_ℓ . (b) The resonance of the optical cavity is detuned from the atomic transition by a frequency Δ 131
- 9.2 Illustration of the effects of the two terms in $H = \hbar\chi[\hat{J}^2 - \hat{J}_z^2]$ (a) The $\chi\hat{J}_z^2$ term causes a twisting of the Bloch sphere, referred to as one-axis-twisting (OAT). A quantum noise distribution (represented by the dashed circle) will be deformed by this twisting, leading to a potentially useful entangled state. (b) The $\chi\hat{J}^2$ term leads to a so-called many-body gap protection effect, where dephasing of the atoms due to inhomogeneous noise sources is suppressed. 133
- 9.3 A Bloch sphere representation of the atomic ensemble radiating into a detuned cavity. The phase of the cavity field C (ϕ_C) is locked to the phase of the atomic coherence J^+ (ϕ) up to a fixed offset $\phi_r = \arctan(2\Delta/\kappa) + \pi/2$. The Bloch vector precesses about the cavity field axis, which is the source of both collectively enhanced decay and of shifts in the output frequency ω_ℓ 134

- 9.4 (a) When the cavity is on resonance with the atomic transition ($\Delta = 0$) the cavity field C is perpendicular to J^+ . This causes a rotation of the Bloch vector downwards, and is the physical origin of superradiant emission. (b) When the cavity is tuned far off resonance from the atomic transition ($\Delta \gg \kappa$), the cavity field C is nearly aligned with J^+ . Now, the rotation of the Bloch vector about C is primarily sideways. Because C follows the Bloch vector as it rotates, this leads to a frequency shift of the Bloch vector precession, and of the emitted light. 135
- 9.5 Shifts in frequency of emitted light ω_ℓ versus cavity detuning Δ , showing expected dispersive behavior. Blue line is a fit with cavity linewidth held to its independently measured value. 138
- 9.6 Main: Pulling coefficient P (black points) and atomic inversion J_z (blue points) versus state prep pulse duration. Blue line is a fit to the blue points with the expected functional form of inversion versus state preparation angle, allowing for a fixed fraction of atoms that remain in the ground state. Black line is predicted value for pulling coefficient given the fit parameters of the blue line, up to an overall scale factor. Inset: the same data can be viewed by plotting P against the value of $J'_z(0)$ inferred from the blue fit line, showing the expected linear scaling. 139
- 9.7 Illustration of the effects of the χJ^2 term. We consider two sub-ensembles, represented by the green and orange Bloch vectors. (a) When the two sub-ensembles have the same transition frequency, the χJ^2 interactions simply lead to a precession about the total Bloch vector J (blue arrow). (b) When the transition frequencies of the two sub-ensembles are split by δ , the two Bloch vectors rotate in opposite directions along the equator when the χJ^2 interactions are absent. (c) When both δ and χ are finite, and $\chi N \gg \delta$, the two sub-ensembles follow small circles and never develop a large opening angle. This is the origin of the many-body gap protection. 141

- 9.8 (a) We introduce controlled inhomogeneity by simultaneously populating the $m_F = \pm 9/2$ states and applying a magnetic field to split the Zeeman sublevels by a frequency δ . (b) Output power spectra for lasing on the two transitions, with atoms prepared near the south pole of the Bloch sphere. In the presence of a large detuning δ , the two branches correspond to light emitted independently on the $m_F = \pm 9/2$ transitions. When the cavity mode is detuned from atomic resonance by $\Delta/2\pi = 150$ kHz $\simeq \kappa/2\pi$ (top row), we experimentally observe an asymmetry in the two branches (left panel). This asymmetry is reversed for opposite cavity detuning (middle row). A simulation with the full Hamiltonian $\hat{H} = \hbar\chi[\hat{J}^2 - \hat{J}_z^2]$ qualitatively reproduces these features (center column), while a simulation that includes only $\hat{H} = \hbar\chi\hat{J}_z^2$ does not (right column). This indicates the important role of the term χJ^2 , which is associated with many-body gap protection. Power spectra at fixed δ (bottom row) highlight this comparison, taken at a detuning δ indicated by the horizontal green lines. (c) Cartoon projections of the Bloch vectors associated with the two spin ensembles onto the xy plane for pure OAT Hamiltonian (upper) and full Hamiltonian (lower) in reference frame of total Bloch vector. For pure OAT, the two ensembles precess independently, leading to a large opening angle. For full dynamics, the two Bloch vectors synchronize, preventing a large opening angle from developing. 144
- 9.9 (a) We introduce controlled, reversible dephasing by illuminating the atoms with light near resonance with the 1S_0 to 3P_1 transition. The beam is offset from center with the atomic cloud, so the atoms experience a spatially varying intensity, and therefore a spatially varying light shift. (b) Applying the beam for roughly 1 ms scrambles the phases of the atoms, reducing atomic coherence and superradiant emission. Applying the laser again, this time with the opposite detuning from the 1S_0 to 3P_1 transition unscrambles the atomic phases, restoring atomic coherence and superradiance. 146

9.10	We measure the change in the difference between optical phase measurements before dephasing (ψ_1) and after rephasing (ψ_2), as the cavity is toggled between opposite detunings $\Delta/2\pi = \pm 30$ kHz: $\Delta\psi = (\psi_2 - \psi_1) _{30 \text{ kHz}} - (\psi_2 - \psi_1) _{-30 \text{ kHz}}$. Red (black) points correspond to trials where the atoms were dephased (not dephased) during T_{hold} , corresponding to the bright (dark) normal modes. $\Delta\psi$ shows no dependence on T_{hold} for the dephased case, and a phase shift that scales linearly with T_{hold} when the atoms are not dephased, confirming a frequency shift between the bright and dark portions of the atomic coherence.	147
9.11	The connection between our observed phase shift and an energy gap can be understood explicitly by considering the effect of dephasing in a Dicke basis for states prepared near $J_z = -N/2$, subject to a pure χJ^2 Hamiltonian. Atoms begin in a superposition of maximally symmetric states with $J = N/2$ (far left states). Dephasing (red arrows) transfers the atomic ensemble to states of minimal coherence for a given J_z . Because the difference in J between the states of maximal and minimal coherence for a given J_z is $J_z + N/2$ for states with $J_z \simeq -N/2$, the energy difference between these states is approximately $\chi(J_z + N/2)$. When coherence is restored in the rephasing step (blue arrows), the different J_z states will have accumulated a phase proportional to this energy shift. Because a linear phase shift between adjacent J_z states corresponds to an azimuthal rotation of the Bloch vector, the measured phase shifts of (b) provide a direct measurement of the energy splitting between manifolds of different J , which is the origin of many-body gap protection.	148
9.12	We may vary the power in the dephasing beams to only partially dephase the atoms. We can infer the residual coherence from the amplitude of the superradiant emission. We observe that the pulling coefficient is unaffected by this process, maintaining the same value as the fully coherent case (to within our experimental uncertainty). . . .	149

10.1	Illustration of Landau-Zener crossing in a dressed state picture. Black lines represent energy eigenstates in the presence of drive coupling Ω . An atom that begins in $ g\rangle$ when the drive has a large negative detuning δ_d will be transferred to $ e\rangle$ as the drive sweeps to large positive detuning (red arrows).	152
10.2	Illustration of basic cooling mechanism. (a) Atoms moving in one dimension with velocity v are illuminated by counter-propagating laser beams. (b) Atoms have a single relevant atomic transition at frequency ω_a . In our case, this transition is dipole-forbidden, so $ e\rangle$ has a relatively long lifetime. (c) Representation of frequencies in lab frame. Because of Doppler shifts resulting from atomic motion, the co-propagating (counter-propagating) laser beams are resonant with $ g\rangle$ to $ e\rangle$ transition when the laser frequency $\omega = \omega_a - kv$ ($\omega_a + kv$). Because ω increases with time, the counter-propagating laser sweeps over the transition before the co-propagating laser. If the atom starts in $ g\rangle$, the counter-propagating laser transfers the atom from $ g\rangle$ to $ e\rangle$ as it sweeps over resonance, and the co-propagating laser transfers the atom back to $ g\rangle$. This results in the transfer of two photon recoils of momenta to the atom in the correct direction to cool.	154
10.3	Eigenstate diagram of coupled atom-light modes versus detuning of the lasers from atomic resonance. Uncoupled eigenstates are labeled as $ m, n, a\rangle$ (defined in text). Atoms start at point (1), then are adiabatically transferred to (2) by the laser sweep. When the laser detuning jumps, they non-adiabatically project to (3), having given up $2\hbar k$ of momentum by transferring a photon from the left-going to right-going modes. Higher-order crossings, which lead to the exchange of multiple photons become significant when $kv \sim \Omega$	159
10.4	Simulated trajectories for atoms with different initial momenta. Produced by a simulation based on a Monte Carlo wave function trajectory method written by Murray Holland and John Bartolotta.	161

- 10.5 Cooling into a moving reference frame. (a) We apply a frequency offset in the lab reference frame between the two beams, which creates a moving reference frame in which the beams have equal frequency. (b) If laser frequency is swept upwards, (green points) the atoms settle into the moving reference frame, which has a velocity indicated by the dashed black line. If frequency is swept downwards (orange points), the atoms accelerate in the other direction, without cooling. If the frequency sweep is symmetric (purple points), atoms experience a small acceleration due to radiation-pressure imbalance. 163
- 10.6 Acceleration versus polarization angle between two beams. Acceleration is enhanced when the polarizations of the two beams are aligned. The blue line represents the maximum acceleration for a saturated two-level atom without stimulated emission. The dashed blue line represents the measured acceleration with a single beam. Black points represent acceleration with both beams applied. The dashed black line represents the acceleration that two photon recoils per sweep of the laser frequency would produce. 164
- 10.7 Measured acceleration versus the sweep center frequency. As long as the lasers sweep over both atomic resonance (blue region), the acceleration is roughly constant. The sweep range here was 4 MHz. 165
- 10.8 Force versus velocity of moving reference frame, as set by the relative detuning $/\Delta$ between counterpropagating cooling beams. Red points correspond to an upwards frequency sweep, while black points correspond to a downwards frequency sweep. Blue line represents half of the sum of the red and black points, while the green line represents half of the difference. We attribute the contributions to the acceleration represented by the blue line to transient effects not directly associated with our cooling mechanism. 166

- 10.9 Experimental study of an atomic ensemble using one-dimensional cooling by adiabatic transfer. (a) An ensemble of atoms, precooled to around $600\ \mu\text{K}$, is illuminated by frequency-swept counterpropagating beams for a variable amount of time. The one-dimensional temperature is decreased to a steady-state value of $45\ \mu\text{K}$ by the cooling lasers. (b) The phase space density in one dimension (grey trace) and three dimensions (black trace) is increased during the cooling process. 167
- 10.10 Temperature versus Rabi frequency for a sweep range of 7 MHz and repetition rate of 15 kHz. Red line represents the results of our simulations, and black points are experimental data. The red band indicates uncertainty of prediction due to statistical uncertainty in simulation results, experimental calibration of Ω and variations in Rabi frequency during the sweep. The vertical line at 130 kHz represents the approximate point at which we expect adiabaticity to break down: $\Omega^2 = \alpha$ 169
- 10.11 Mechanism for realizing a confining force with swept lasers. Counter-propagating beams with opposite circular polarizations illuminate the atoms in the presence of a magnetic field gradient, as in a standard MOT (left). The atoms are located near the zero of the magnetic field. The MOT lasers are swept upwards in frequency in a sawtooth manner, as before. In the simplest system, each laser connects the ground state to a unique excited state, which shift in opposite directions in an applied magnetic field. The lasers sweep over the transition to the state that is shifted to lower energy first, causing acceleration in the direction of the beam that addresses that state, and depleting the population in the ground state (right). Because the excited state is relatively long lived, the atom is less likely to be in the ground state when the other laser sweeps over its resonance. This leads to an imbalance in the number of photons absorbed from each beam, and causes a magnetic field dependent force that can be used to spatially confine the atoms near a magnetic field zero. . . . 171

10.12	Atomic acceleration in the configuration described in Fig. 10.11. The direction of the force switches sign depending on the sign of the magnetic field applied along the direction of the counter-propagating beams, and depending on the direction of the sweep (blue and red points). This is the necessary criteria for generating a restoring force toward a zero in a spatially dependent magnetic field. A symmetric frequency sweep generates a smaller force, which does not switch sign with the magnetic field (black points), likely due to an intensity imbalance between the two beams.	173
10.13	Fraction of atoms transferred from the broad-line MOT at 461 nm to the narrow-line MOT at 689 nm versus center frequency of the sweep for ^{87}Sr and ^{88}Sr for different sweep profiles. Red points represent an upwards-swept asymmetric sawtooth ramp, which we refer to as the SWAP MOT configuration. Black points represent the more commonly used symmetric triangle ramp. Blue points represent a downwards-swept asymmetric sawtooth ramp, which captures almost no atoms.	174
11.1	Optical lattice clock and atom interferometer detectors for gravitational waves. Each type of proposal relies on two ensembles of atoms, one in each satellite. Lasers transmitted between the satellites encode phase shifts due to changing optical path lengths caused by a gravitational wave. The key difference between clock and atom interferometer proposals is that in clocks, atomic recoils due to the momentum of absorbed/emitted photons are suppressed by tightly confining the atoms, while in atom interferometers the atoms are free to recoil.	180

- 11.2 A toy model for gravitational wave detection. (a) Two atoms (a and b) are addressed by a single laser resonant with a transition between the ground state $|g\rangle$ and a long-lived excited state $|e\rangle$. A phase modulator mimics the effect of a gravitational wave by modifying the optical path length between the two atoms. (b) For tightly confined atoms (clocks), a simple Ramsey sequence can be used to detect changes in optical path length. A first $\pi/2$ pulse imprints the laser phase on the two atoms when the path length is shifted by L_1 , and a second $\pi/2$ pulse acts when the path length is shifted by L_2 . Random variations in the laser phase (ϕ_1 and ϕ_2) are common to the two atoms, while phase shifts resulting from changes in path length are not. The signal from the phase modulator manifests as phase shift $\Phi_s = k(L_2 - L_1)$ between the excited state probability oscillations of the two atoms. (c) For unconfined atoms (atom interferometers), transitions between $|g\rangle$ and $|e\rangle$ are accompanied by momentum kicks, which necessitate an additional laser π pulse in the middle of the sequence. The phase shift between excited state oscillations of the two atoms is now $\Phi_s = k(L_3 - 2L_2 + L_1)$ 181
- 11.3 LMT-like enhancement with confined atoms. (a) A toy model to demonstrate enhancement mechanism. Atoms are addressed by counter-propagating lasers, both of which pass through the phase modulator that sits between the atoms. (b) Two blocks of M pairs of π pulses (blue boxes) are inserted between the $\pi/2$ pulses of the standard Ramsey sequence. Each pair of π pulses contains a pulse originating from each direction. The duration of each block is assumed here to be short relative to the GW period, while a long evolution time (grey box) between the two blocks may be comparable to the GW period. These additional pulses lead to a factor of $4M + 1$ enhancement in the signal phase Φ_s compared to simple Ramsey sequence. 187

- 11.4 Use of three-state system to allow atoms to stay in ground state for much of evolution time. (a) Our procedure utilizes two ground states, $|g_1\rangle$ and $|g_2\rangle$ and a single long-lived optically excited state $|e\rangle$. $|g_1\rangle$ and $|g_2\rangle$ can independently be coupled to $|e\rangle$ using laser pulses with different frequency and/or polarization (red and blue arrows). (b) After preparing an atomic superposition of $|g_1\rangle$ and $|e\rangle$ using a $\pi/2$ pulse from the left (atoms begin in $|g_1\rangle$), the portion of the atoms in $|e\rangle$ is transferred to $|g_2\rangle$ using a π pulse from the right. The atoms are then in a superposition of $|g_1\rangle$ and $|g_2\rangle$ with phase dictated by the lasers and the path length L_1 . The sequence of pulses is reversed after a long evolution time, during which time the atoms are not susceptible to excited state decay. The phase shift in the excited state probabilities of the two atoms $\Phi_s = 2k(L_2 - L_1)$ now encodes changes in the optical path length that occurred while the atoms were in the ground states. 191
- 12.1 A proposed architecture for spin squeezing in an optical lattice clock. Atoms are tightly confined in two dimensions by intersecting lattices supported by two crossed optical cavities. One of these cavities is used to perform atom counting measurements on the 7.5 kHz linewidth 3P_1 transition. Clock rotations are performed by a 698 nm laser applied in the plane defined by the two cavity modes, along which direction the atoms are tightly confined and Doppler-shift free. 199
- 12.2 A proposed architecture for a continuously loaded superradiant laser. Atoms emitted from oven are slowed by a Zeeman slower, then captured in a 2D MOT formed on the 461 nm transition. Some atoms that fall into 3P_2 are magnetically trapped and fall guided along the magnetic field zero. They are then depumped to the ground state and cooled transversely with 689 nm light before falling into the mode of the optical cavity. Here, they are loaded and cooled into a moving optical lattice, which transports them to the lasing region, where each atom emits up to one photon. . . 202

Chapter 1

Introduction

The exquisite control of cold atoms, ions, and molecules in tabletop experiments has opened new windows in our understanding of the universe at its most fundamental levels. This has occurred through several primary avenues, notably precision measurement as well as quantum information science and simulation.

Precision measurements have provided experimental values or bounds for fundamental and universal quantities with critical implications for our understanding of the most foundational aspects of physics. These quantities include the electron electric dipole moment [1, 2], the anomalous magnetic moment of the electron and fine structure constant [3], and precision tests of key predictions of relativity [4, 5], to name a few. These studies complement and sometimes extend the results obtained by more established, and far more expensive techniques such as high energy (accelerator) physics.

The same kinds of tools can be applied in a completely different context, exploiting the quantum nature of the experimental platform provided by cold atoms, ions, and molecules. These experiments provide direct experimental confirmation and great intuitive insight into some of the puzzling aspects of quantum mechanics [6, 7]. The current cutting edge of this sub-field involves the use of quantum systems for information processing and simulation tasks, and the related quest to understand the highly non-classical states that may arise in systems of large numbers of interacting quantum particles. It is conceivable that the new tools for computation and simulation will upend our current information and economic infrastructure [8], and would lead to a greater understanding

of fascinating phenomena that arise in condensed matter systems [9].

This continuing bounty of insights and applications has been made possible by the ever-developing tools used to prepare, manipulate, and detect useful states of quantum particles. These advances include continuous improvements in laser cooling and trapping, as well as other methods of creating low-entropy quantum systems such as evaporative cooling and quantum gas assembly [10, 11, 12, 13]. Advances in techniques for the coherent manipulation of the quantum states of particles using microwave and laser radiation have provided crucial tools for enabling precision measurement and information processing, and advances in detection techniques have given experimentalists the capability to observe increasingly interesting states [14, 15, 16]. Finally, creating controllable interactions between the quantum particles provides the potential to study complex quantum systems and for the enhancement of precision measurements [17, 18, 19].

In this thesis, I present a collection of new tools for further advancing precision measurement and quantum science, based on narrow-linewidth optical transitions. These tools include a non-destructive atom counting technique that may enhance the performance of optical lattice clocks and other quantum sensors [20], a novel form of cavity-enhanced atomic spectroscopy appropriate for laser frequency stabilization for portable sensing applications [21], a new form of laser cooling with reduced reliance on spontaneous emission that may improve and extend the use of laser cooling in new experimental systems [22], a method of enhancing the sensitivity of proposed atomic gravitational wave detectors [23], cavity-mediated spin-spin interactions with implications both for precision quantum sensors and for engineering interesting interactions between particles, and a new class of laser (a “superradiant laser”) based on optical superradiance from narrow and ultra-narrow linewidth transitions [24, 25]. As part of the superradiant laser project, we have demonstrated the most precise active absolute frequency standard realized to date, providing an encouraging indication that a superradiant laser may become a key experimental tool for precision measurement in the near future.

With the exception of the gravitational wave detection project (which was purely theoretical), this work was performed in a new experimental system that I began building in the fall of 2013. In




Figure 1.1: (a) Strontium has two valence electrons, whose spins can either be aligned (spin triplet) or anti-aligned (spin singlet). Photons are very bad at flipping spins, so we call transitions between these two types of states “forbidden,” meaning that they occur slowly if at all. (b) A basic overview of our experimental platform. Strontium atoms are cooled to roughly $10\ \mu\text{K}$, and tightly confined in an optical lattice inside of an optical cavity. The wavelength of the lattice is carefully chosen so as not to perturb atomic transition frequencies of interest — the “magic wavelength.” The atoms interact with the cavity via one of several forbidden transitions.

this system, an ensemble of strontium atoms interact in a collective manner with an optical cavity via narrow-linewidth optical transitions.

The important property of strontium for this work is the presence of two valence electrons. The spins (intrinsic magnetic moments) of these two electrons can either be aligned in a configuration known as a “triplet” state, or anti-aligned in a configuration known as a “singlet” state (See Fig. 1.1a). To go from a singlet state to a triplet state or vice-versa requires flipping the spin of one of the electrons. Photons are very bad at doing this, so these transitions happen very slowly if at all, and are thus known as “forbidden transitions” or “dipole-forbidden transitions.” The work here utilizes two forbidden transitions in atomic strontium: the somewhat forbidden 7.5 kHz linewidth $^3\text{P}_1$ to $^1\text{S}_0$ transition and the far more forbidden 1 mHz linewidth $^3\text{P}_0$ to $^1\text{S}_0$ clock transition. These

transition linewidths correspond to the inverse of the excited state lifetime, which is $20\ \mu\text{s}$ for $^3\text{P}_1$ and $160\ \text{s}$ for $^3\text{P}_0$. Both of these transitions involve the absorption or emission of a red photon. For scale, an “allowed” transition (which does not involve flipping a spin, or any other processes that a photon cannot do) of this color would have a linewidth of order $10\ \text{MHz}$, or a decay rate one thousand and ten billion times faster than the two forbidden transitions. These narrow-linewidth transitions in strontium and similar elements have already proved very useful, and have enabled the development of the most precise and accurate clocks to date [26, 27, 28, 29]. The tools that I describe in this thesis have the potential to improve the performance of these clocks even further, and to extend the impact of forbidden optical transitions to new types of systems and sensors.

The key enabling technology for much of my work is the coupling of strontium atoms to an optical cavity by these narrow linewidth transitions (See Fig. 1.1b). The optical cavity consists of two highly reflective mirrors that light can bounce back and forth between many times. The cavity provides distinct resonant optical modes, defined by their spatial profile and frequency, with which the atoms can interact through the exchange of photons. The strontium atoms are cooled using laser light to a temperature of order 10^{-5} Kelvin, and suspended within the optical cavity by an optical lattice — a standing wave of light that is resonant with the optical cavity. Atoms are attracted to regions of high intensity of the lattice light, which pins the cold atoms in place during the experiment, greatly reducing a variety of problems associated with atomic motion. The wavelength of the lattice light is set near a “magic” value, for which it does not significantly perturb the frequency of the narrow linewidth clock transition.

The study of atoms coupled to an optical cavity constitutes a field known as cavity quantum electrodynamics, or cavity QED. Historically, cavity QED with atoms has typically focused on coupling the atoms to the cavity using either microwave transitions between Rydberg states [30], or strong allowed optical transitions [31]. The same kind of ideas have been extended to superconducting circuits in the microwave domain [32]. A key focus of my thesis work is the extension of cavity QED techniques into a new regime by instead coupling the atoms to the cavity via narrow-linewidth optical atomic transitions. The slow interaction rates associated with these

narrow-linewidth transitions required substantial care in order for the atom-cavity interactions to dominate over technical sources of decoherence.

Before beginning the construction of the strontium experiment, I contributed to several experiments that I will mention only briefly here. These experiments were performed on the preexisting rubidium cavity QED apparatus with fellow graduate students Justin Bohnet, Josh Weiner and Kevin Cox. The first was a spin-squeezing experiment where we demonstrated for the first time an order of magnitude improvement in sensor performance (10 dB of directly observed spin squeezing) using entanglement generated from collective measurements of atoms inside an optical cavity [33]. In the second experiment, we demonstrated a technique in which we imparted controlled, reversible disorder to an ensemble of atoms in order to suppress technical errors in coherent atomic manipulations [34].

I summarize experiments performed in the strontium cavity QED system next, then discuss two sets of additional tools that are enabled by narrow linewidth optical transitions but do not involve an optical cavity — a new form of laser cooling that relies on narrow-linewidth transitions, and a method for enhancing the performance of proposed gravitational wave detectors.

1.1 Overview of Applications

1.1.1 Application to Clocks: Non-Destructive Atom Counting

The coupling of ensembles of atoms to a cavity via a narrow linewidth optical transition holds promise for improving the performance of optical lattice clocks, which are currently the most precise sensors of time and frequency in existence [26, 27, 28, 29]. Like many clocks, optical lattice clocks rely on state-dependent atom number measurements to infer frequency. Typically, these measurements are made through fluorescence detection, where atom number is inferred by scattering light off of an atomic transition. This technique has the advantage of simplicity and good signal-to-noise, but has the disadvantage that the scattered photons heat the atoms to the point where a new ensemble must be prepared. During the time that new atoms are being prepared, the

frequency of the clock may wander, which can degrade its performance [35].

Another alternative approach to fluorescence detection is to couple the atoms to an optical cavity, and count the atoms by measuring perturbations to the resonance frequency of the cavity. In addition to excellent precision, this method has the advantage that it may be performed in a non-destructive manner, eliminating the need to prepare a new ensemble of atoms, and reducing problematic dead-time in the clock sequence. Further, by making collective measurements of atomic population one can generate useful entanglement or “spin squeezing” in the ensemble, which can further improve the performance of the clock [36].

Proof-of-principle demonstrations of non-destructive atom counting and spin squeezing have been performed using dipole allowed transitions in alkali atoms coupled to optical cavities [14, 37, 33, 15, 17]. An outstanding goal is to apply these techniques to a state-of-the-art optical lattice clock based on an ultra-narrow optical transition in an alkaline-earth or alkaline-earth-like atom such as strontium or ytterbium. A promising path towards this goal is to couple the atoms to a cavity via the 7.5 kHz linewidth dipole-forbidden transition to perform atom counting and spin squeezing.

As a first application of the new strontium system, we demonstrated that this transition could indeed be used to perform non-destructive atom counting at a level compatible with spin squeezing in a system applicable to state-of-the-art optical clocks. We showed that we could count atoms with precision compatible with quantum-limited clock operation, and inferred that the measurement was sufficiently non-destructive to enable spin squeezing.

1.1.2 Application to Laser Stabilization: Magnetically Induced Transparency

We have also observed and characterized an interesting and potentially useful effect that arises when an atom is coupled to the optical cavity by multiple narrow-linewidth transitions. For most applications of cavity QED, each atom is viewed as a two-level system with a single transition. By carefully preparing the initial state of the atoms interacting with a cavity mode of well-defined frequency and polarization, it is often possible to neglect additional atomic structure,

which greatly simplifies the understanding of the system. In this work, we probe a cavity mode whose polarization is orthogonal to an applied magnetic field, and thus interacts with two atomic transitions that share a ground state. This allows us to observe a feature that we call “magnetically induced transparency” (MIT) in analogy with a well-studied phenomena called electromagnetically induced transparency (EIT) [38, 39, 40]. This feature manifests as a spectrally narrow transmission window through the cavity that appears when the applied magnetic field is sufficiently large. The center frequency of this window has low sensitivity to changes in the cavity resonance frequency, making it an appealing candidate for laser stabilization. Further, the ability to tune the nature of this feature by changing the magnitude of the applied magnetic field could enable its use for storing and retrieving photons in a manner similar to EIT [41, 42, 43, 44, 45].

1.1.3 Superradiance from Narrow Linewidth Optical Transitions

Much of my thesis focused on the study of superradiance from narrow-linewidth optical transitions, with the ultimate goal of creating a narrow linewidth superradiant laser. Superradiance refers to a phenomena in which the rate that photons are emitted from an ensemble of atoms greatly exceeds the sum of the rates at which photons would be emitted from the individual atoms. Superradiance itself has been studied in detail as a curiosity for decades [46, 47, 48, 49]. More recently it has been proposed [50, 51] that one could take advantage of an ultra-narrow linewidth clock transition like the one in strontium to realize a superradiant laser with linewidth of order 1 mHz (surpassing the current records for laser stability [52]).

Like any laser, a superradiant laser would consist of a gain medium that sits within the mode of an optical buildup cavity. In contrast to a conventional laser, where the gain medium provides broad-band amplification of power and the lasing frequency is set by the resonance frequency of the optical cavity, a superradiant laser would operate in a regime where the gain medium is by far the most spectrally selective element in the system, and the optical cavity serves to enhance emission in a comparatively broadband manner. The output frequency of the laser is then determined by the center frequency of the gain medium. Because cold atoms with narrow linewidth optical transitions

form a very narrow linewidth and stable gain medium, this could be a major advantage.

Today's state-of-the-art stable lasers [53, 54, 52], which are used in state-of-the-art clocks [29, 27, 55, 56, 26] consist of a fairly standard laser whose frequency is stabilized [57] to the resonance of a highly engineered optical cavity, which consists of a spacer to which two mirrors are stuck. The frequency of the laser is thus determined by the length of the optical cavity. Over the last several decades, heroic efforts have been made to design cavities with very stable lengths [58]. Currently, the stability of optical cavities is limited by thermal brownian motion in the spacer, mirror substrates, and reflective coatings [59, 60]. To mitigate these effects, researchers are currently developing cavities with crystalline spacers, substrates and reflective coatings that operate at cryogenic temperatures [54, 52]. This is very hard, I am told. Reaching the regime where these thermal fluctuations limit the performance of the laser also requires elaborate vibration isolation techniques, which might make it difficult to use these lasers outside of a carefully controlled laboratory environment.

A superradiant laser would offer an alternate, and perhaps complementary approach to creating a stable laser. Because the frequency of a superradiant laser is determined by the atomic transition frequency rather than the cavity resonance frequency, a superradiant laser should be largely insensitive to these effects that continue to limit current stable lasers.

From another perspective, because the light emitted from a superradiant laser comes directly from a narrow-linewidth atomic transition, one can view these devices as active optical frequency standards. A superradiant laser would basically be the optical equivalent of a maser, which is an active microwave frequency standard [61]. The fact that the superradiant laser operates in the optical domain instead of the microwave domain is a major advantage.

The advantage of operating a frequency reference at high frequency is quite general. For a given phase resolution and measurement time, the uncertainty in one's ability to measure frequency, expressed in absolute terms, is fixed. However, because a reference tone can be transferred to different frequency bands while maintaining its original fractional frequency error [62], the key figure of merit for a frequency reference is its fractional uncertainty. This gives high-frequency

standards a huge advantage over low frequency standards with similar absolute uncertainties.

Fixed phase resolution is a reasonable assumption to make here, as most frequency measurements are fundamentally limited by shot noise associated with measuring a finite number of quantum particles (e.g. atoms or photons). This shot noise generally limits one’s ability to resolve a phase to $1/\sqrt{N}$ radians, where N is the number of particles measured. Apart from phase resolution, many other technical limitations such as shifts from magnetic fields contribute frequency noise whose magnitude is independent of the frequency at which the frequency standard operates. These errors are also fractionally smaller for optical frequency references than for their microwave counterparts.

The power of working at optical frequencies can be seen in the great improvement in the performance of atomic clocks when the jump was made from the microwave to optical domain. Because of their higher operating frequency, the fractional frequency stability and accuracy of these clocks quickly overtook their microwave counterparts, and are now several orders of magnitude more favorable. Because a superradiant laser would operate in the optical domain, one might reasonably expect to see similar improvements over microwave masers to those that optical clocks achieved relative to microwave clocks.

In this thesis, I will present a series of milestones on the path towards a technologically useful superradiant laser. Before I joined the group, previous members of the Thompson lab — Justin Bohnet, Zilong Chen, Joshua Weiner, and Kevin Cox — performed the first demonstrations and characterizations of a quasi steady-state superradiant laser [63, 64, 65, 66]. This experiment used Raman transitions to create an effective long-lived optically excited state in an ensemble of cold rubidium atoms. While not useful as a frequency reference (the frequency noise on the “dressing” laser used to create the Raman system appears on the laser output), this system was an ideal test-bed for studying many aspects of steady-state superradiance. Unlike rubidium, strontium has several narrow-linewidth optically excited states, suitable for the operation of a true narrow linewidth optical superradiant laser.

As a first step towards a superradiant laser, we demonstrated pulsed and quasi-steady state

lasing on the $\gamma = 2\pi \times 7.5$ kHz linewidth $^1\text{S}_0$ to $^3\text{P}_1$ transition in ^{88}Sr [24]. The partially forbidden nature of the 7.5 kHz transition puts the system in an interesting regime. The cavity linewidth ($\kappa = 2\pi \times 160$ kHz) is much broader than the atomic transition linewidth, so excitations leave the cavity primarily through the cavity mirrors rather than through spontaneous emission. The frequency of coherent, collective interactions between the atoms and cavity is typically of order $\Omega = 2\pi \times 1$ MHz, much larger than either decay rate. In pulsed operation, a laser operating with $\Omega \gg \kappa, \gamma$ exhibits interesting ringing behavior as photons are emitted into the optical cavity and reabsorbed.

We also demonstrated that it is possible to repump this laser, returning atoms that had emitted a photon to the excited state to achieve quasi steady-state operation. In this steady-state regime, the collectively enhanced emission rate is limited by the repumping rate. The repump process also broadens the linewidth of the lasing transition, which we can use to continuously tune the laser from the conventional “good cavity” regime of laser physics where the laser frequency is set by the cavity to the “bad cavity,” or superradiant regime where the laser frequency is set by the atoms.

The next step was a demonstration of lasing on the 1 mHz linewidth clock transition in ^{87}Sr [25]. In contrast to the laser based on the 7.5 kHz transition, this laser operates deep in the superradiant regime, where cavity decay far exceeds both atomic dissipation and the rate of coherent interactions. We demonstrated collectively enhanced emission in which atoms emitted photons into the cavity a factor of up to 10,000 times faster than would independently radiating atoms.

After demonstrating and characterizing superradiance from the perspective of the radiated power, we turned to a characterization of the spectral properties of the emitted pulses of light by comparison with a very stable laser from Jun Ye’s lab down the hall. We characterized key elements of the performance of our superradiance as a frequency reference, including the short-time stability, absolute accuracy, and sensitivity to perturbations such as cavity frequency shifts and magnetic fields. We have found the short-term stability to be within an order of magnitude of

the current state-of-the-art optical lattice clocks (making ours the most precise active frequency reference demonstrated to date), and the absolute frequency to be within roughly 1 Hz of its predicted value. We have demonstrated a method to effectively eliminate sensitivity to magnetic field fluctuations, and a reduction in sensitivity to changes in cavity frequency by more than five orders of magnitude. These observations are very encouraging for the prospect of creating even more precise superradiant frequency references.

1.1.4 Cavity-Mediated Spin-Spin Interactions

When we tune our cavity on resonance with a narrow linewidth optical transition, the primary effect is an increase in the rate of decay from the excited state (superradiance). If we instead detune the cavity from atomic resonance, interesting Hamiltonian dynamics emerge, which we can describe as cavity-mediated interactions between the effective spins formed by the upper and lower states of the lasing transition. In our case, these interactions manifest as both an energy shift to the atomic transition that depends on both atomic inversion and cavity frequency, and as a self-energy term that leads to a so-called many-body gap protection. The energy shift leads to the buildup of atom-atom correlations via an effect known as “One-Axis-Twisting” (OAT). OAT has been demonstrated in driven systems as a way to create entangled spin-squeezed states [17, 67, 68, 69]. In our system, the atoms themselves create the drive that leads to twisting. This idea was independently proposed in ref. [70]. The second effect — many-body gap protection — leads to a suppression of decoherence in the presence of interactions mediated by the optical cavity. A similar effect, mediated by collisions between atoms, has been observed and utilized in optical magnetometers [71], and cold-atom systems [72, 73, 74].

Experimentally, the superradiant light exiting the cavity gives us a precise realtime nondestructive probe of the atomic evolution [65]. By observing frequency shifts of the superradiant light, we can probe the energy shifts of the atomic transition. We use this to characterize the OAT dynamics at the mean-field level. We can also simultaneously populate two sets of atomic states, and observe simultaneous lasing on transitions between pairs of these states. This allows us to

controllably introduce a frequency splitting between the two transitions, which in turn allowed us to observe signatures of many-body gap protection. Finally, we demonstrate a technique in which we introduce controllable reversible dephasing to perform a direct spectroscopic measurement of the many-body gap.

This work has important implications for understanding potential perturbations to the superradiant laser, for generating metrologically useful entangled states, and also as a demonstration of a potentially appealing new type of system for studying many-body physics phenomena.

1.1.5 A Novel Cooling Mechanism

The following two projects do not rely on coupling atoms to an optical cavity, but do rely on properties enabled by narrow-linewidth optical transitions. The first concerns several related new mechanisms for laser cooling and trapping.

The development of laser cooling has revolutionized atomic physics. Typically, laser cooling has relied on broad-linewidth optical transitions. Because atoms decay very quickly from the excited state of these transitions, we usually view the atoms as primarily living in the ground state, and occupying the excited state only fleetingly before emitting a photon. Similar techniques to those used on broad transitions have been extended experimentally to narrow linewidth transitions [75, 76, 77, 78, 79, 80], but the conceptual framework used to understand the cooling is very similar to the broad linewidth case.

The use of narrow linewidth optical transitions opens up fundamentally new possibilities for laser cooling beyond those previously explored. Because the excited state of such transitions is long-lived, we can now consider the use of mechanisms that rely on populating the excited state for significant periods of time. In this work, we demonstrate two related but distinct mechanisms that rely on transferring significant population to the excited state of a narrow linewidth optical transition.

The first mechanism, which we have named SWAP (sawtooth wave adiabatic passage) cooling, relies on subsequent Landau-Zener transitions to and from an excited state driven by counterprop-

agating laser beams whose frequencies are swept over atomic resonance in a sawtooth waveform. This method enables the generation of large forces even on very weak transitions, and the possibility of cooling with reduced amounts of spontaneous emission. This may make it an appealing mechanism for the cooling of molecules, for which spontaneous emission causes major challenges. SWAP cooling generates a velocity dependent force that leads to cooling and can increase phase-space density in an atomic ensemble.

We also demonstrate a related mechanism that provides a confining force as well as cooling. This mechanism utilizes a similar sawtooth ramp to SWAP cooling, but relies on shelving atoms in one of several long-lived excited states to generate a differential scattering force that depends on the local magnetic field. This allows for the formation of a magneto-optical trap (MOT). From a practical perspective, this second mechanism has proved to be very useful in our experiment, and forms the basis of our cooling in everyday operation. It’s use has enabled us to increase our atom number, improve the robustness of our experiment, and greatly simplify our cooling and trapping protocol.

1.1.6 Enhanced Gravitational Wave Detection with Atoms

The final project that I will describe here involves a new class of proposals for gravitational wave detectors based on cold atoms with narrow linewidth optical transitions [81, 82, 83, 84]. These proposals rely on ensembles of atoms contained within two distant satellites. Laser pulses are passed back and forth between the two satellites, interacting with the atoms on each end. A passing gravitational wave modifies the phase of the light as it interacts with the far ensemble, and this phase shift is encoded in atomic observables. The use of narrow linewidth optical transitions enables protocols that appear to be impossible with more conventional approaches using broad linewidth transitions or Raman transitions between ground states.

These new proposals fall into two categories: those described as “clocks,” where the atoms are tightly confined by an optical potential, and those described as “atom interferometers” (AI), in which the atoms are not confined and as a result experience momentum kicks from absorbing

or emitting photons. My goal was to understand the fundamental differences between these two types of proposals. In particular, the AI proposals rely on a technique known as “large momentum transfer” (LMT) to enhance sensitivity, while clocks had no such analog.

In this work, we provided a framework in which the fundamental mechanism for sensitivity of the two types of proposal can be understood. We reached the conclusion that the two types of proposals were fundamentally the same, and that the photon recoils and resulting spatial separation of the atoms’ wavefunctions in the AI proposals have no fundamental role in the gravitational wave detection. We further showed that a technique analagous to LMT could be applied to a clock, though no momentum is actually transferred to the atoms. Finally, we proposed an enhancement mechanism in which useful evolution time can be achieved with the atomic population entirely in ground states, which could be separated by acoustic rather than optical frequencies. These enhancement mechanisms do not make sense if we think of the detectors that utilize confined atoms simply as clocks whose sole purpose is to measure frequency. Rather, our perspective highlights that we should think of the atoms as phase memories that allow the phases of light pulses arriving at different times to be compared.

1.2 Outline of Thesis

The rest of this thesis is organized as follows: In Chapter 2 I will outline language and theory needed to understand systems where atoms are coupled to cavities. I will particularly focus on those elements that are different when using narrow-linewidth optical transitions. In Chapter 3, I will describe the design, construction and basic operation of the experimental apparatus. In chapters 4-10, I cover a series of experiments performed using the new apparatus and narrow-linewidth optical transitions. Chapter 11 describes a new proposal for improved detection of gravitational waves using cold atoms with narrow-linewidth optical transitions. Finally, in Chapter 12 I provide an outlook to future extensions of this work and rough outlines for possible experimental implementations of spin squeezing in an optical lattice clock, and for a steady-state superradiant laser.

Chapter 2

An Introduction to Atoms in Cavities and Narrow-Linewidth Transitions

Much of the work presented in this thesis centers around coupling large ensembles of atoms to optical cavities via narrow-linewidth optical transitions. This chapter provides a basic introduction to the physics of atoms interacting with cavities. I will introduce some of the language and parameters used to define atom-cavity systems, along with simple scaling arguments where applicable. This will enable a discussion of different regimes in which one might operate, and the advantages and challenges of each. I will particularly focus on these tradeoffs in the context of coupling atoms to cavities via narrow-linewidth dipole-forbidden optical transitions. Finally, I will present a simple and intuitive approach to describing and visualizing the atom-cavity system, taking advantage of the Bloch sphere formalism and an analogy to a classical pendulum.

2.1 Narrow-Linewidth Transitions in Strontium

The primary distinguishing feature of strontium, as compared to the more commonly used alkali atoms, is the presence of two valence electrons. This property leads to rich internal structure. Importantly, the spin of the two valence electrons can either be aligned in a symmetric manner (spin $S = 1$), known as a triplet configuration, or in an anti-symmetric or singlet configuration ($S = 0$). Any transition between a singlet and a triplet state is in principle forbidden, as it would involve flipping the spin of one electron relative to the other.

In reality, the states we write down are not actually triplet or singlet states. Rather, when we write down a state like 3P_1 , where the 3 stands for triplet and which decays to the nominally

singlet 1S_0 state at a rate of $\gamma = 2\pi \times 7.5$ kHz, we are actually referring to a state that is mostly triplet-like, but has some singlet character mixed in. In the case of strontium (and other alkaline earth atoms), the finite lifetime of 3P_1 is due to the presence of spin-orbit coupling, which mixes in some 1P_1 character into the state that we call 3P_1 . In ^{87}Sr , which has nuclear spin $I = 9/2$, $I \cdot S$ coupling mixes some 3P_1 character into 3P_0 , which allows it to decay very weakly, resulting in a transition linewidth of roughly 1 mHz.

Great detail and calculations of these effects can be found in [85, 86]. For the rest of this work however, the key takeaway is that atomic strontium has two particularly useful forbidden transitions, with very weak dipole matrix elements. The 7.5 kHz linewidth transition is present in all isotopes of strontium, while the 1 mHz linewidth transition is present only in the fermionic isotope, ^{87}Sr , as this is the only isotope with $I \neq 0$.

2.2 Atoms Interacting with a Cavity

Two of the simplest systems in quantum mechanics are the two-level system and the harmonic oscillator. The interaction between an atom and a cavity provides a clean experimental realization of the coupling of these two canonical systems, with the atom serving as the two-level system (with ground and excited states $|g\rangle$ and $|e\rangle$) and the cavity providing the harmonic oscillator. Despite (or more likely because of) the simplicity of this system, the field of cavity quantum electrodynamics (cavity QED) has yielded a wealth of fundamental understanding of quantum mechanics, as well as great potential for technologically useful systems.

The typical goal of cavity QED is to realize a situation where an atom (or ensemble of atoms) interacts primarily with a single mode of electromagnetic radiation. The cavity supports discrete resonant modes, each with a well-defined frequency and spatial profile. Each of these modes is analogous to a harmonic oscillator, with a ladder of evenly spaced energy levels representing integer numbers of photons occupying the mode. When an atom is located within one of these modes whose resonance frequency is near that of the atomic transition, then the interaction between the atom and this mode may dominate over coupling to other modes of the cavity or the environment.




Figure 2.1: A cavity QED system. Atoms exchange excitations with the cavity mode at a rate $\Omega = 2g\sqrt{N}$. Atoms can decay into free-space via single-particle spontaneous emission at a rate γ , leading to loss of coherence. Photons in the cavity leak out through the mirrors at rate κ . This channel of dissipation is easy to detect, and provides only collective information about the state of the atomic ensemble.

Starting with the case of a single atom in the cavity, the atom-cavity system can be described by the so-called Jaynes-Cummings Hamiltonian:

$$\hat{H} = \hbar\omega_a\hat{\sigma}^+\hat{\sigma}^- + \hbar\omega\hat{c}^\dagger\hat{c} + \hbar g(\hat{\sigma}^-\hat{c}^\dagger + \hat{\sigma}^+\hat{c})$$

The first term corresponds to the energy of the atom — ω_a is the atomic transition frequency and $\hat{\sigma}^+ = |e\rangle\langle g|$ ($\hat{\sigma}^- = |g\rangle\langle e|$) is the atomic raising (lowering) operator. The second term represents the energy of the cavity field — ω is the cavity resonance frequency, and \hat{c}^\dagger (\hat{c}) are the photon creation (annihilation) operators. Finally, the third term represents the interaction between the atom and the cavity. The atom can go from $|e\rangle$ to $|g\rangle$ by emitting a photon, or from $|g\rangle$ to $|e\rangle$ by absorbing one. The frequency at which this happens is given by g , the single particle coupling.

We can rewrite this Hamiltonian in a rotating frame at the atomic transition frequency ω_a (basically just subtracting out the atom's excited state energy, and assuming that the transition from $|e\rangle$ to $|g\rangle$ is accompanied by the addition of a photon to the cavity) as:

$$\hat{H} = \hbar\delta_c\hat{c}^\dagger\hat{c} + \hbar g(\hat{\sigma}^-\hat{c}^\dagger + \hat{\sigma}^+\hat{c})$$

where $\delta_c = \omega - \omega_a$ is the cavity detuning .

In addition to the coherent exchange of excitations between the atom and cavity, we must consider the coupling of the system to the environment. The atom can decay by emitting a photon into a mode other than the cavity mode at a rate γ , the single-atom spontaneous decay rate from $|e\rangle$ to $|g\rangle$ in free space. Photons in the cavity decay at a rate κ , the cavity linewidth. Note that κ represents a power decay rate, not an amplitude decay rate. This is illustrated in Fig. 2.1. Additional sources of broadening such as Doppler shifts of the atomic transition cause dephasing of the collective excitations stored in the atoms. These additional sources of broadening must be carefully controlled so as not to dominate the dynamics.

The single atom coupling g is a very important parameter in cavity QED, so it is worth understanding how it scales. Following the definitions in Dan Steck's quantum optics notes [87] (which I will do as much as possible in this thesis), g is given by:

$$\hbar g = -\hat{\epsilon} \cdot \vec{d} \sqrt{\frac{\omega}{2\epsilon_0 \hbar V}}$$

$\hat{\epsilon}$ represents the polarization of the cavity field, \vec{d} is the dipole matrix element of the atomic transition, and V is the volume of the cavity mode (the cross-sectional area at the location of the atoms times the cavity length). Physically, $\hbar g$ represents the energy $\vec{d} \cdot \vec{E}$ associated with the atomic dipole interacting with the field associated with a single photon in the cavity. A summary of the scalings of basic and derived parameters is given in Table 2.1.

From an experimental design perspective, for a fixed frequency ω , g depends only on atomic properties (\vec{d}), and the volume of the cavity mode. A short cavity with a tightly focused mode will thus have a large g , while a long cavity with a spatially broad mode will have a small value of g . Note that the quality of the mirror coatings do not effect the value of g because we have assumed that there is a single photon in the cavity.

If N atoms are placed inside the cavity mode and excited by the cavity field, the Hamiltonian becomes

$$\hat{H} = \hbar\delta_c\hat{c}^\dagger\hat{c} + \hbar g(\hat{J}_-\hat{c}^\dagger + \hat{J}_+\hat{c})$$

where $\hat{J}_\pm = \frac{1}{2}\sum_{i=1}^N\hat{\sigma}_i^\pm$ are the collective raising and lowering operators for the atoms. If we focus on the case where the number of atomic excitations is much less than the total number of atoms, we can approximate this Hamiltonian [88, 89] as

$$\hat{H} = \hbar\delta_c\hat{c}^\dagger\hat{c} + \hbar g\sqrt{N}(\hat{a}\hat{c}^\dagger + \hat{a}^\dagger\hat{c}) \quad (2.1)$$

where we replace \hat{J}_\pm with creation and annihilation operators $\hat{a} = J_+/\sqrt{N}$ and $\hat{a}^\dagger = J_-/\sqrt{N}$.

The takeaway from this exercise is that when there are many spins inside the cavity, the rate at which a photon is exchanged between the cavity and the atoms is enhanced from $2g$ to $2g\sqrt{N}$, which I will follow convention and label Ω , the collective vacuum Rabi frequency. While this derivation assumed a low-excitation limit, the basic scaling is true in general so long as dissipation rates are sufficiently low.

Quantity/Label	Formula/Definition	Key Scaling
Atomic Transition Frequency ω_a		
Cavity Field Frequency ω	$\omega \simeq \omega_a$	
Cavity Length L		
Loss	power loss in round trip	
Cavity Finesse F	$2\pi/\text{Loss}$	
Wavelength λ	$2\pi c/\omega$	
Mode Waist w	$1/e$ radius of field	
Cavity Mode Volume V	$\pi w^2 L$	
Dipole Matrix Element d		
Atomic Linewidth γ	$\frac{\omega_a^3 d^2}{3\pi\epsilon_0 \hbar c^3}$	$d^2 \omega_a^3$
Cavity Linewidth κ	$c/2LF$	Loss/L
Coupling g	$d\sqrt{\frac{\omega}{2\epsilon_0 \hbar^3 V}}$ (rms value of averaged over standing wave)	$d\sqrt{\omega/V}$
Cooperativity Parameter C	$4g^2/\kappa\gamma$ (averaged over standing wave)	$\lambda^2/(w^2 \text{Loss})$

Table 2.1: Key parameters and scalings for cavity QED.

Generally, the goal of cavity QED is to create a situation where interactions between the atoms and the cavity dominate in some way over other effects. These situations are broadly referred to as the strong coupling regime. There are several useful definitions of the strong coupling regime, which are each relevant in different situations and for different tasks. The simplest criteria is that the rate of coherent interactions — $2g$ in the single-atom case and $2g\sqrt{N}$ in the multi-atom case — should exceed the dissipative rates κ and γ (as well as other sources of inhomogeneous broadening). These criteria define the resolved vacuum Rabi splitting regime.

A second definition of the strong coupling regime involves the cooperativity parameter of cavity QED, $C = 4g^2/\kappa\gamma$. In the limit that $\kappa \gg \gamma$, the cooperativity parameter quantifies the relative probability that an atom placed in the excited state in the cavity mode will decay by emitting a photon into the cavity mode, versus by emitting a photon into free space. For many applications of cavity QED, $C \gg 1$ is a key figure of merit. If many atoms are used, C becomes collectively enhanced to a value of NC , which can greatly exceed unity.

Historically, several approaches have been used to attain strong coupling. I will briefly review

these, to provide context for the new regime explored in this thesis.

The first of these approaches, exemplified by the work of the Haroche group, involves coupling single Rydberg atoms to a superconducting microwave cavity. In these experiments, Rubidium atoms prepared in circular Rydberg states traverse the mode of a microwave resonator coated in superconducting niobium. Because the cavity is superconducting, losses are very low (small κ). Further, because transitions between circular Rydberg states have very large dipole moments, they couple very strongly to the microwave field (large g). Finally, because the transitions occur at microwave frequencies, the available phase-space for spontaneous emission is relatively small, which leads to a long lifetime (small γ). These qualities, combined with efficient state-selective detection techniques, made this an ideal system for early explorations of the strong coupling regime, leading to a wealth of pioneering work [90, 49, 30] and a Nobel prize.

An alternate approach, initially pursued by the groups of Jeff Kimble and Michael Feld, among others, is to couple single Alkali atoms to optical resonators by dipole-allowed optical transitions. Certain physical scalings lead to persistent challenges with this approach, but the allure of coupling atoms to optical fields (whose thermal occupation at room temperature is low [91]) has justified great efforts in this domain and some impressive successes [92, 31].

In order to overcome the rapid spontaneous emission rate associated with optical transitions (which scales as the cube of frequency), one must confine the cavity mode to very small volumes to increase the local electric field seen by the atoms, and thus the rate of coherent interactions. Even if one assumes that the mode volume scales with the wavelength of the radiation as $1/\omega^3$ (which practically gets really hard to do), the ratio g/γ scales unfavorably from the perspective of optical transitions as $1/\omega$. Counterintuitively, it is the relatively small dipole matrix element of these optical transitions that makes them viable at all for cavity QED applications — g scales with \vec{d} , while γ scales with \vec{d}^2 , so achieving $g/\gamma \gg 1$ is actually easier with a weak matrix element \vec{d} , though this does make it harder to achieve $g \gg \kappa$. This is a hint that transitions with even smaller dipole matrix elements may be a favorable way to go.

In order to maximize g (to achieve $g \gg \gamma$), these systems end up employing very short

optical cavities (sub-millimeter lengths) with very tightly focused waists in order to attain a small mode volume V . Unfortunately, while $g \propto 1/\sqrt{L}$, $\kappa \propto 1/L$. This means that the short cavity length L makes it challenging to achieve the condition $g \gg \kappa$, which requires very high mirror reflectivity. Breakthroughs in dielectric mirror coatings enabled the realization of cavities with finesse F exceeding 10^5 , and the attainment of single-atom strong collective coupling.

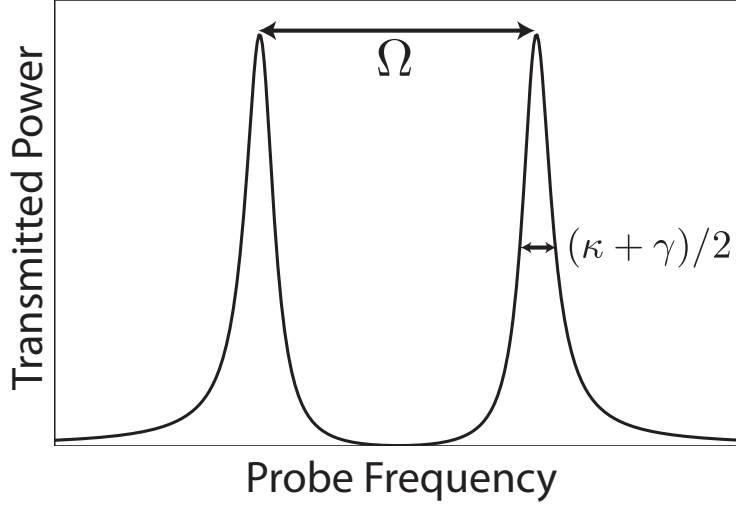


Figure 2.2: One of the key manifestations of the strong coupling regime defined by $\Omega \gg \kappa, \gamma$ is the vacuum Rabi splitting. As a weak probe tone coupled to the cavity is swept over resonance with the cavity (here, $\delta_c = 0$), a single transmission peak splits into two. The splitting between these peaks is set by Ω , and the width of each peak is a weighted average of atomic and cavity decay rates.

The key signature of the single-atom resolved vacuum Rabi splitting regime ($g \gg \kappa, \gamma$) is the single atom vacuum Rabi splitting (pictured in Fig. 2.2), an effect where a single transmission peak of the optical cavity splits into two peaks separated by $2g$ (when $\delta_c = 0$), corresponding to hybrid atom-cavity modes [93]. The equivalent feature for the collective regime is the collective vacuum Rabi splitting, in which the peaks are separated by $\Omega = 2g\sqrt{N}$. The eigenfrequencies of our Hamiltonian for $\delta_c = 0$ are $\pm\Omega/2$. These determine the center frequency of the two peaks. The eigenmodes that correspond to these eigenfrequencies are half atom-like and half photon-like in character, and thus decay at a rate given by the average of the atom and cavity decays $\kappa' = (\kappa + \gamma)/2$.

This decay rate determines the width of each of the two peaks.

If the two peaks are resolved, it means that their separation, which is set by the rate of coherent interactions, exceeds their width, which is set by decay rates from the system. A very nice review of the origin and characteristics of the vacuum Rabi splitting can be found in ref. [88].

For many applications, it is not necessary (or desirable) to operate with only a single atom. For these applications, collective enhancement may be used to ease technical requirements. The strong collective coupling regime has been applied to the generation of entanglement through effective unitary dynamics [17] and measurement [14, 15, 33], the creation of single photon sources [94] and studies of optomechanical effects with cold atoms [95, 96], among others.

2.3 Cavity QED on Dipole-Forbidden Transitions

Much of my thesis work centers around extending the collective coupling approach to dipole-forbidden optical transitions. This has been previously explored in the context of cavity-enhanced non-linear spectroscopy in the group of Jan Thomsen, where an ensemble of roughly milliKelvin temperature strontium atoms are coupled to a moderate-finesse optical cavity via a dipole-forbidden 7.5 kHz linewidth optical transition to obtain cavity-enhanced spectroscopic features useful for laser stabilization [97, 98]. In these experiments, the relatively high temperature of the atoms led to inhomogeneous Doppler broadening that exceeded the rate of coherent interactions between atoms and cavity and prevented the observation of some of the canonical features of cavity QED, such as the vacuum Rabi splitting.

In the work I present in this thesis, we both further cool the atoms to temperatures of a few microKelvin and tightly confine the atoms to less than a wavelength using an optical standing wave lattice. This enables the observation of canonical cavity QED features on the 7.5 kHz linewidth strontium transition, as well as the exploration of a far more extreme regime of cavity QED using the even more forbidden 1 mHz linewidth optical clock transition.

In this section, I will discuss the differences between using narrow-linewidth optical transitions and the more conventional regimes explored previously. I will focus on what gets harder or easier

with narrow transitions, and what advantages they may provide over broad-linewidth transitions.

Let us consider the first criteria for strong coupling presented earlier, the resolved vacuum Rabi splitting: $\Omega \gg \gamma, \kappa$. A forbidden transition has a long-lived excited state because the dipole matrix element d associated with the transition is small. This leads to a reduction in g , and therefore $\Omega = 2g\sqrt{N}$ compared to a dipole-allowed transition. Fortunately for the goal of achieving this definition of the strong coupling regime, the atomic linewidth $\gamma \propto d^2$, so $\Omega \gg \gamma$ is actually easier to achieve for a dipole-forbidden transition.

However, if we account for sources of atomic broadening besides photon emission into free space, things get much harder when using narrow linewidth transitions. Effects like Doppler shifts ($\Delta\omega_D = 2\pi v/\lambda$, where v is the atomic velocity and λ is the wavelength of the radiation) do not scale with the strength of the atomic transition. Ensuring that these effects do not overwhelm coherent interactions, which are now much slower, requires far more care and lower temperatures than for dipole-allowed transitions. Specifically, assuming a fixed wavelength, in order to fix the same ratio of $g/\Delta\omega_D$ between a narrow linewidth transition and a broad linewidth transition, one must scale the temperature by the ratio of the transition linewidths.

Further, achieving the condition that $\Omega \gg \kappa$ is also more challenging. For the 7.5 kHz linewidth strontium transition, this is still relatively easy to achieve with modern mirrors. However, accessing this regime on the 1 mHz linewidth transition would be experimentally challenging.

For our other definition of strong coupling, $NC \gg 1$, the scalings work out somewhat differently. Because $g \propto d$ and $\gamma \propto d^2$, the dipole matrix element actually cancels from C . Thus, for this definition of the strong coupling regime, the fact that we use dipole-forbidden transitions is fundamentally unimportant. However, it is sometimes important that the rate $NC\gamma$, the rate of collectively enhanced emission per atom in an overdamped regime, exceeds the rate of inhomogeneous dephasing. This leads to a more stringent requirement that the temperature must scale as the square of the transition linewidth in order to maintain the same ratio of $NC\gamma/\Delta\omega_D$, which would be very hard to achieve using unconfined atoms and narrow linewidth transitions.

There is one more consideration in a cavity QED system where the transition linewidth is

critical: the ratio of γ to κ . The regime where $\kappa \gg \gamma$ is often referred to as the bad-cavity regime, while $\gamma \gg \kappa$ defines the good-cavity regime. In the bad-cavity regime, excitations are more likely to leave the atom-cavity system as photons emitted from the cavity mirrors, while in the good-cavity regime they are more likely to leave as spontaneously emitted photons. In many applications of cavity QED, photons transmitted through the cavity mirrors are desirable, as they provide only collective information about the atomic ensemble and are easily collected. Photons scattered into free-space are typically detrimental, as they lead to dephasing of the atomic ensemble and are difficult to utilize.

This makes the bad-cavity regime desirable for many applications, and therefore makes the use of dipole-forbidden transitions very advantageous. In cavity QED experiments that involve dipole-allowed optical transitions (which typically have linewidths of a few MHz), one must work hard to access the bad-cavity regime. In order to achieve a high value of C , one would like to use a high-finesse cavity. The down-side is that this implies a long photon storage time and corresponding narrow linewidth, unless one makes the cavity very short and difficult to work with. In contrast, when working with a narrow-linewidth optical transition, one can work in the bad-cavity regime with a cavity of convenient length (allowing ample optical access to atoms trapped within the cavity and a relatively small free spectral range) and reasonably high finesse.

In certain applications one can attain some of the advantages of being in the bad-cavity regime by detuning the optical cavity resonance from atomic resonance or using Raman transitions. While these techniques have been applied with great success in a variety of systems [63, 17, 14, 15, 99], their use does lead to certain limitations not present with narrow-linewidth transitions. For example, by working in the bad-cavity regime with the cavity on resonance with a narrow-linewidth transition, we were able to demonstrate a noise-insensitive probing scheme that would not be possible with a detuned cavity [20].




Figure 2.3: (a) The state of a two-level system with states $|e\rangle$ and $|g\rangle$ can be represented by a Bloch vector that lives on a Bloch sphere. The figure shows a general state $|\psi\rangle = \sin(\theta/2)|g\rangle + \cos(\theta/2)e^{i\phi}|e\rangle$. The vertical projection of the Bloch vector tells us the population inversion of the state (the relative likelihood of finding the atom in the ground or excited state), while the azimuthal angle ϕ tells us the relative phase between the $|e\rangle$ and $|g\rangle$ components. (b) An ensemble of many two-level systems can be represented on a collective Bloch sphere by simply adding up all of the individual Bloch vectors associated with the two-level systems to form a collective Bloch vector \vec{J} . The inversion of the collective Bloch vector is labelled \vec{J}_z , and the atomic coherence (projection onto the equatorial plane) is labelled \vec{J}^+ .

2.4 Cavity QED with Bloch Spheres and Pendula

This section introduces a very powerful and intuitive way of thinking about cavity QED using Bloch spheres. This model is totally classical, yet captures almost all important aspects of cavity QED for the purposes of the work in this thesis. Those that it does not can often be added in by hand as necessary. I will also introduce an even more clearly classical picture, where the behaviour of the atoms and cavity are mapped onto that of a pendulum experiencing gravity.

The goal is to provide a visual formalism for describing the dynamics of an atom interacting with a cavity. Detailed mathematical formalism can be found, for example, in Daniel Steck's quantum optics notes [87]. We find that this simple visual picture provides not only great explanatory power, but rapid predictive power that is key to running the experiment.

I will start with a description of how we represent the different parts of the physical system. A single atom will be treated as a two-level system, which can be represented by a vector on a

Bloch sphere of radius $1/2$. For pure quantum states, the tip of the vector lies on the surface of the sphere. If the atom is in the excited state $|e\rangle$, the vector points to the north pole of the sphere. If the atom is in the ground state $|g\rangle$, the vector points to the south pole. An atom in the state $|\psi\rangle = \sin(\theta/2)|g\rangle + \cos(\theta/2)e^{i\phi}|e\rangle$ is depicted in Fig. 2.3a.




Figure 2.4: We can describe the dynamic of an ensemble of atoms interacting with the cavity field using a Bloch sphere picture. The cavity field is represented by an axis \check{C} about which the Bloch vector rotates at a rate proportional to $g|\check{C}|$. As pictured, the rotation is leading to collective decay of the atoms from $|e\rangle$ to $|g\rangle$.

For systems with more than one atom, we simply add the Bloch vectors of the individual atoms together to form a collective Bloch vector \vec{J} , which lives on a collective Bloch sphere. In order to account for the fact that the atoms are located at different points along the cavity axis, and therefore experience a different optical phase from the cavity field, we define the phase of each atom's Bloch vector relative to the phase of the optical field of the resonant mode. So long as the atoms do not move around, and coherent rotations are only performed using the cavity mode (as is the case for pretty much everything in this thesis), we can safely ignore this detail. We describe \vec{J} in terms of its vertical projection J_z , which is $-N/2$ at the south pole, 0 at the equator and $N/2$ at the north pole, and its projection onto the equatorial plane $J^+ = J_x + iJ_y$, which is a complex number whose real and imaginary parts are $J_x = J\sin\theta\cos\phi$ and $J_y = J\sin\theta\sin\phi$ respectively. These quantities are labelled in Fig. 2.3b.

We describe the cavity field C (not the cooperativity parameter introduced earlier, but this

is a somewhat common notation) as a vector that originates from the center of the Bloch sphere. For a field whose frequency is on resonance with the atomic transition (which is the relevant case for the primary applications of this model in this thesis), this vector lies in the equatorial plane of the Bloch sphere. The Bloch vector rotates about this cavity field vector at a frequency $2g|C|$, with C expressed in units of the amplitude associated with a single photon. This is shown in Fig. 2.4.

In general, C can be any arbitrary cavity field, with any general source. Here though, I will specialize to the case where the cavity field is radiated by the atoms themselves (hence the assumption that it is on resonance with the atomic transition). We wish to describe the effects of this self-radiated field back on the atoms. Following [64], we may write down differential equations (known as optical Bloch equations) that describe the interactions with the cavity field C , as well as dissipation. We assume that both the atomic coherence and cavity field oscillate at the same as of yet unknown frequency ω_γ , and define $\check{J}^+ = J^+ e^{i\omega_\gamma t}$ and $\check{C} = C e^{i\omega_\gamma t}$. I will keep the $\check{}$ s around for the remainder of this chapter, but will treat them as implied after that. We may then write equations of motion for \check{J}^+ and \check{C} :

$$\dot{\check{J}}^+ = -[\gamma_\perp + i(\omega_a - \omega_\gamma)]\check{J}^+ + i2g\check{C}J_z \quad (2.2)$$

$$\dot{J}_z = -\gamma(J_z + N/2)/2 + ig(\check{J}^+\check{C}^* - \check{J}_-^*\check{C}) \quad (2.3)$$

$$\dot{\check{C}} = -[\kappa/2 + i(\omega_c - \omega_\gamma)]\check{C} - ig\check{J}^+ \quad (2.4)$$

Stepping through the terms in these equations, γ_\perp defines the decay rate of J^+ due to excited state decay or other processes. In the absence dephasing processes other than excited state decay, $\gamma_\perp = \gamma/2$. Note that these equations only capture processes that would lead to an exponential decay of J^+ , such as spontaneous emission or collisions that lead to large, instantaneous phase shifts. Other processes that lead to a Gaussian decay of J_\perp such as Doppler decoherence require a more sophisticated treatment. ω_a is the frequency of the bare atomic transition, and $2g$ is the single-atom vacuum Rabi frequency. γ represents the excited state decay rate, κ is the power decay rate of the cavity, and ω_c is the bare cavity resonance frequency.

We can also understand the nature of the collective enhancement of the vacuum Rabi frequency in this picture. Each atom's Bloch vector rotates at a rate proportional to $|\check{C}|$. If we start with N atoms near the excited state, they will emit photons into the cavity, causing the Bloch vector to swing down to the south pole. At this point, $|\check{C}| \propto \sqrt{N}$, as each atom has emitted one photon into the cavity. The rate at which the Bloch vector is rotating is thus enhanced by a factor of \sqrt{N} , which is one manifestation of the collective enhancement. In general, whenever the atoms are interacting with a self-generated field, the rate of the interactions is enhanced because each atom sees the field radiated by all of the other atoms.

We can gain further intuition though a simple mapping: For a cavity tuned to resonance with the atomic transition and no damping of the atomic coherence (though cavity damping is still captured), the equations of motion for the Bloch vector are equivalent to those for a pendulum under gravity. This picture can be very useful for describing cavity QED phenomena.




Figure 2.5: With the substitutions given here and in the text, the dynamics of an ensemble of atoms interacting with the mode of a resonant cavity can be mapped onto the motion of a damped pendulum under gravity. The angular velocity of the pendulum takes the role of the cavity field while the decay of photons from the cavity is equivalent to viscous damping of the motion of the pendulum. The length of the pendulum scales as $(g^2J)^{-1}$, which leads to faster dynamics for many atoms or large coupling g .

Because we are neglecting damping of the atomic coherence (setting $\gamma = \gamma_{\perp} = 0$), the length J of the Bloch vector is fixed. Further, because we assume that the cavity is on resonance

($\omega_a = \omega_c = \omega_\gamma$), we can assume that \check{J}^+ is real. We can then describe the position of the Bloch vector by its polar angle θ alone. Rewriting the above equations of motion in terms of θ , we get:

$$\dot{\theta} = i2g\check{C} \quad (2.5)$$

$$\ddot{\theta} = 2g^2J\sin\theta - \kappa\dot{\theta}/2 \quad (2.6)$$

We can compare these to the equation of motion for the angular velocity $\dot{\theta}$ of a pendulum:

$$\ddot{\theta} = g_{\text{grav}}\sin\theta/L - \beta\dot{\theta} \quad (2.7)$$

Here, g_{grav} is the gravitational acceleration, L is the length of the pendulum, and β is the velocity damping rate, due for example to the pendulum being immersed in a viscous fluid.

The cavity field \check{C} thus corresponds to the angular velocity of the pendulum (up to a factor of $2g$). The strength of gravity and length of the pendulum are constrained by $g_{\text{grav}}/L = 2g^2J$, and damping of the cavity field corresponds to viscous drag on the pendulum. Strong coupling (high g) or large atom number (large J) correspond to a short pendulum in a strong gravitational field.

I will return to this pendulum analogy several times throughout this thesis, but for now will give one example of its utility. An obvious quantity of interest for a pendulum is its resonance frequency. Does this correspond to anything meaningful in cavity QED? The resonance frequency of a pendulum is $\sqrt{g_{\text{grav}}/L}$, which maps to $g\sqrt{2J} = g\sqrt{N}$. If we excite the pendulum, it will swing back and forth with a frequency $g\sqrt{N}$. The angular velocity (and thus the cavity field) of the pendulum also oscillates at this frequency. If we look at the field inside the cavity, it will thus be amplitude modulated as $C = \cos(g\sqrt{N}t)e^{i\omega_c t}$. The Fourier spectrum of this field consists of two peaks, centered around the cavity resonance ω_c , and split by $2g\sqrt{N}$. This spectrum corresponds to the famous collective vacuum Rabi splitting of cavity QED.

This pendulum mapping is also very useful for providing intuition into different regimes of superradiance, and I will return to it in later sections of this thesis.

Chapter 3

Experimental Apparatus

In this chapter, I describe the basic experimental hardware and processes that allows us to operate our experiment. I will focus on those elements that are general to most or all experiments that we perform, and will leave atomic manipulation and detection techniques specific to particular experiments to their corresponding chapters.

From a broad perspective, the key requirement of the apparatus is to achieve a situation where the rate of collective interactions between the atoms and cavity, mediated by the mHz linewidth transition, exceeds all atomic decoherence rates in the system. This is the requirement to obtain superradiance, and dictated most design choices in the experiment.

For superradiance experiments operating in the deep bad-cavity regime (where the cavity is much more lossy than the atoms), the rate of collective interactions is given by $NC\gamma$, where N is the atom number, C is the cavity cooperativity parameter, and γ is the decay rate of the transition. We want this rate to be large, which for a given atomic transition requires many atoms and a high-cooperativity cavity. In practice, much effort has gone into optimizing the number of atoms in the cavity.

Achieving a low enough decoherence rate involves removing any factors that would cause a scrambling of the relative phase between the superposition states of the different atoms. Such factors include collisions with background gas, which is eliminated through the use of an ultra-high vacuum environment, Doppler shifts, which are eliminated by cooling and tightly confining the atoms, and interactions with electromagnetic fields, which are managed by carefully controlling

the optical and magnetic fields experienced by the atoms. These requirements provide the key constraints for the experimental design and implementation.

3.1 Experimental Overview

As is typical of cold-atom experiments, this experiment runs in a cyclical manner. In each cycle, a new ensemble of atoms is prepared, manipulated, measured, then discarded. This is repeated with a cycle time of roughly one second. The details of the manipulation and measurement portions of the sequence depend on the exact experiment being performed, but the initial preparation step is for the most part common.




Figure 3.1: A simplified atom loading sequence. Atoms begin as a roughly 500 C collimated beam exiting the oven. Some of the atoms are slowed to around 50 m/s by a Zeeman slower, then redirected by two 2D MOTs into the main experimental chamber. They are then further cooled and spatially compressed by a 3D MOT, which we typically load for 500 ms. The Zeeman slower and both MOT stages also use 461 nm (blue) light. The atoms are then transferred from the blue MOT to a narrow-line red MOT over roughly 50 ms. The magnetic field gradient (represented by the slope of the black line) is reduced by a factor of roughly ten during this step. An asymmetric sawtooth frequency sweep applied to the red MOT beams greatly improves the efficiency of this transfer (See Chapter 10). Atoms are further cooled, and loaded into the optical lattice with a typical temperature of around 10 μ K.

The process of preparing an ensemble of confined cold atoms involves several steps that are summarized in Fig. 3.1. The atoms originate as a collimated beam emanating from a hot oven. They are then slowed and cooled from this initial high temperature using dissipative laser cooling

techniques. This portion of the cycle uses the broad-linewidth “blue” transition at 461 nm, and brings the atoms down to a few milliKelvin. At this point, further cooling is performed using the “red” transition at 689 nm, which has a much narrower linewidth and is compatible with cooling to much lower temperatures. For this step, we utilize a new form of cooling and trapping, which we refer to as Sawtooth-Wave Adiabatic Passage (SWAP) cooling. This technique has significant advantages over conventional laser cooling techniques, and is described in detail in Chapter 10.

Once the atoms are appropriately cooled using SWAP cooling, they are loaded into the conservative potential of the optical lattice. This optical lattice is supported by the same cavity that the atoms interact with via narrow-linewidth transitions, ensuring spatial overlap between the atoms and cavity mode. The lattice both holds the atoms within the cavity mode for the duration of the experiment and tightly confines them along the cavity axis, eliminating Doppler shifts in that important direction. The wavelength of the lattice is set to the so-called “magic” wavelength, which minimizes perturbations of the milliHertz linewidth clock transition [100, 101].

After this point in the experimental sequence, the procedure differs depending on the specific experiment being performed. I will cover most of the specifics of these procedures in the corresponding experimental sections (Chapter 4-10). One important distinction between different experiments is the choice of isotope. We use one of two isotopes: ^{88}Sr (the easy isotope), and ^{87}Sr (the hard isotope). ^{88}Sr has roughly ten times the natural abundance of ^{87}Sr , making it much easier to achieve high atom numbers. The most significant difference between the isotopes however is the nuclear spin — ^{88}Sr has none, while ^{87}Sr has a nuclear spin $I = 9/2$. The presence of nuclear spin in ^{87}Sr causes the $^3\text{P}_0$ to $^1\text{S}_0$ transition to become weakly allowed [85]. The large nuclear spin of ^{87}Sr also leads to a great proliferation in the number of atomic states. $^3\text{P}_0$ and $^1\text{S}_0$ now each have ten Zeeman sublevels, and other states have additional hyperfine structure. While we are sometimes able to exploit this multitude of states, they generally lead to additional complications when working with ^{87}Sr .



Figure 3.2: Pictures of the experiment from July 2017, to give a sense of scale. Upper left: four-bore reference cavity, with associated optics. Upper right: main vacuum chamber end of experiment. Lower left: the experiment spans two optics tables. Vacuum system is in the foreground, with lasers and laser stabilization in background. Lower right: enclosure houses all red and infrared lasers.

3.2 Vacuum System

The atoms must be prepared, cooled, and manipulated within an ultra-high vacuum (UHV) environment in order to avoid loss and perturbations associated with background-gas collisions. Our vacuum system consists of three main sections: the atomic source (provided by AOSense, Inc.), the main science chamber, and the pump appendage. A model of the vacuum system is pictured in Fig. 3.3

The function of the atomic source is to provide a beam of pre-cooled strontium atoms to the main chamber, and will be described in more detail in Sections 3.3 and 3.5.2. From a vacuum perspective, the source does not present a significant load on the rest of the chamber. In fact,




Figure 3.3: The vacuum system consists of three main sections: the atom source, which contains an oven, Zeeman slower, and two 2D MOTs, the science chamber, which contains the optical cavity, and the pumping appendage. The system is pumped by a 75 liter/second ion pump and a titanium sublimation pump.

the pressure measured within the atom source chamber is lower than that of the main science chamber, even with the oven at full temperature. This is because strontium is highly reactive, and the coating that forms in the atomic source chamber serves as an excellent pump. The atom source is connected to the main chamber by an all-metal valve (VAT 48124-CE01-0001) that allows us to vent the main chamber without compromising the vacuum inside the atomic source chamber and oxidizing the strontium in the oven.

The main chamber is a Kimball physics 8 inch spherical octagon (MCF800-ExtOct-G2C8A16). This chamber houses the optical cavity and provides ample optical access. All windows have a broadband anti-reflective coating provided by TAKOS with a specification of less than 1 percent reflection for wavelengths between 461 and 1064 nm.

The pump appendage is connected to the bottom 8 inch conflat flange of the main chamber. The pump appendage contains an angled mirror to redirect the vertical MOT beams through the center of the cavity spacer. Pumping is performed by a 75 l/s ion pump (gamma vacuum TiTan

ion pump, 75S). We have also installed a titanium sublimation pump, but did not notice a major change in vacuum levels from running it. Typical pressures measured by the main ion pump are in the mid 10^{-10} Torr range, sufficient for our needs. This level of vacuum is probably limited by the relatively large amount of macor, epoxy, viton and Kapton coated wire associated with the cavity and its support structure, and by the relatively low bake temperature that these components require.

3.3 Atomic Source

We used an atomic source provided by AOSense, Inc. This source consists of an oven that provides a collimated beam of hot strontium atoms, a permanent-magnet Zeeman slower, two 2D MOT's with in-vacuum permanent magnets and optics, a heated sapphire window upon which the beam of strontium atoms impinges, and necessary vacuum pumps for UHV operation. During typical operation, the pressure inside the source chamber is 3×10^{-11} Torr, lower than the pressure in our main science chamber. Operating details are given in Section 3.5.2.

3.4 MOT Coils and Driver

Because of the relatively large size of the main science chamber and the broad linewidth of the 461 nm cooling transition (32 MHz), the MOT coils ended up being quite serious. The goal was to be able to create magnetic field gradients of up to 50 G/cm without the coils getting hot, as this could heat the chamber and the cavity spacer, causing drifts in beam alignment and in the cavity resonance frequency. The coils must also be capable of switching off within a few milliseconds.

I wound the MOT coils out of 3/16" hollow square wire potted with Araldite 2011 epoxy. The hollow core of the wire allows chilled water to flow through the coils. Each of the two MOT coils has an inner diameter of 8.25", an outer diameter of 10.5", and a length of 1.75". The centers of the coils are spaced by 6". These dimensions were chosen to maintain a relatively homogeneous field gradient (by staying near an anti-Helmholtz configuration), while fitting over the 8 inch flanges of our chamber. Each coil has 60 turns, comprised of three sets of 20 turns that are wired electrically

in series by jumper bars, but whose coolant passages are wired in parallel. Running the coolant through three parallel paths increases the total coolant flow by a factor of nine relative to a single path. In current operating conditions, the coils are wired together in series. We flow 80 Amps through the coils to generate a gradient in the vertical direction of 30 G/cm, and half that in the other two directions. Heating of the coils is negligible under these conditions.

Three sets of bias coils in near-Helmholtz geometry provide uniform bias fields in three orthogonal directions. Each coil that produces a horizontal field has 51 (yes, 51) loops of 14AWG copper wire. A closeup of the coils and support structures is shown in Fig. 3.4. The vertical bias coils (50 turns each) are wound around the MOT coils for structural support and cooling. The MOT coils are held in place by Garolite (cotton fabric impregnated with phenolic resin) supports. This material is chosen because it is nonmagnetic and non-conductive, which means it will not support eddy currents when the coils are switched on and off. The bias coils are wound on aluminum supports, which are connected in the corners by Delrin plastic blocks so they do not form complete conductive loops (also to eliminate eddy currents).




Figure 3.4: Closeup of MOT and bias coils. MOT coils (large circular coils) are wound from hollow square wire through which water flows. The vertical bias coils are wound onto the MOT coils for support and cooling. The two sets of horizontal bias coils are wound around square frames. Additional freely wound coil in the foreground creates an additional gradient along the cavity axis, but is not typically used.




Figure 3.5: Left: diagram of circuit used to control current through MOT coils. Right: actual circuit used to control current through MOT coils. All high-power circuit elements are mounted to a water-cooled plate. The MOSFET is the large white component in the top of the image. The shunt resistor is the large black component below it, and the diodes are below that.

The circuitry that controls and switches the current through the MOT coils also requires some degree of care in its design, and is shown in Fig. 3.5. The amount of current flowing through the coils is controlled by a very large MOSFET (microsemi APTM50UM09FAG, rated for 500 V, 497 A, 5000W dissipation), and sensed by a 0.5 m Ω shunt resistor (Ohmite TGHGCR0005FE, 1%, 60ppm/C) in series with the coils. A JILA coil driver with PID controller provides a control voltage to the gate of the MOSFET in order to drive the current through the shunt resistor to a desired level. The control electronics are mounted to a water-cooled plate. Copper bars and plates are used to connect the various components.

When the coils are fully energized, a substantial amount of energy is stored in the magnetic field. When the coils are turned off quickly, we must provide a path for the current to go, and safely dissipate this energy. Otherwise, the voltage across the MOSFET would rise until something fails dangerously. This is the job of the flyback circuit. Currently, the flyback circuit consists of a

1 Ohm resistor (TAP800K1R0E) and two diodes (UFL230FA60), all in series. Upon initial turnoff, the voltage across this resistor jumps up to $I_{max} * 1\text{Ohm}$, where I_{max} is the maximum current flowing through the MOT coils. The current through the coils then decays exponentially with a time constant of L/R until the voltage drop becomes dominated by the diodes, at which point it decreases linearly to zero. In their current configuration, the coils settle from their operating current to zero in 10 ms, and to 10 percent of operating current in 6 ms. If desired, this could be made faster by using a larger flyback resistor at the expense of creating a higher voltage spike in the flyback circuit. Quickly switching the coils also leads to mechanical shaking of the table as well as pickup on the cavity piezos, so we choose not to be more aggressive with the switching speeds.

As a word of caution, a previous version of the flyback circuit proved unable to handle the peak power dissipation that we needed, and after six months of operation failed in an open configuration. The current from the MOT coils then had nowhere to go and exploded the large switching transistor which was rated for 600 V and 300 A, presumably higher than a human. This highlights how dangerous these circuits can be if not properly designed, even when hooked up to a low-voltage power supply.

3.5 Laser Systems and Laser Cooling

In contrast to Alkali atoms, the more complicated level structure of Alkaline Earth atoms requires many more laser wavelengths to utilize. The laser systems used in our experiment are summarized in Table 3.1, along with their common functions. Details on the various systems follow.

An energy level diagram with relevant states and transitions is shown in Fig. 3.6. Dipole-forbidden transitions, which are of particular interest for this thesis, are shown as dashed lines. Spin singlet states are shown on the left, and spin triplet states are shown on the right. Transitions shown between states of the same total spin are dipole allowed, and have decay rates of at least several MHz, while transitions between singlet and triplet are dipole forbidden and decay much more slowly.

Laser systems			
Transition	Wavelength (nm)	Functions	Stabilization Method
1S_0 to 1P_1	461	Zeeman slower, 2D MOTs, 3D MOT, fluorescence	polarization spectroscopy in hollow-cathode lamp
1S_0 to 3P_1	689	narrow-line MOT, optical pumping, heterodyne frequency reference, cavity probe	reference cavity
1S_0 to 3P_0	698	seeding superradiance, heterodyne frequency reference	reference cavity, fiber from Ye lab
3P_0 to 3S_1	679	depopulate 3P_0	reference cavity or free-running
3P_1 to 3S_1	688	repump for steady-state lasing	reference cavity
3P_2 to 3S_1	707	depopulate 3P_2	FM spectroscopy in hollow-cathode lamp
NA	813	magic wavelength lattice	science and reference cavities

Table 3.1: Overview of laser systems.

3.5.1 461 nm Laser System

The current generation of the 461 nm laser system used for initial cooling is shown in Fig. 3.7. It provides power for the Zeeman slower, 2D MOTs and 3D MOT. A master laser (Toptica DL-pro) is frequency stabilized to a polarization-rotation spectroscopy signal detailed in [102] and modeled after the system described in [103].

It is desirable to have far more 461 nm power available than our single diode laser can supply. To achieve this, we use three injection locked “slave” diodes and a second master diode

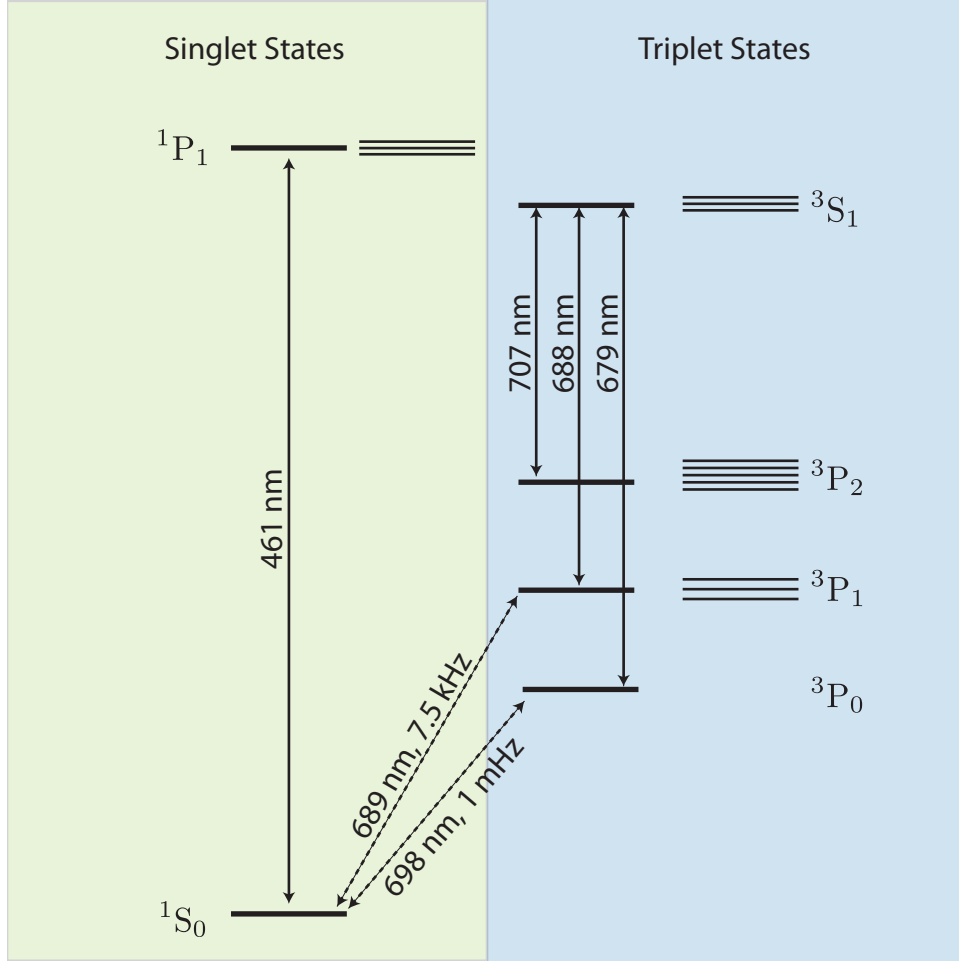


Figure 3.6: Relevant energy levels of strontium. Level structure of ^{88}Sr , which lacks nuclear spin and therefore hyperfine structure, are denoted with thick lines. When present, additional states resulting from the hyperfine structure of ^{87}Sr are shown with thin lines. Zeeman sublevels are neglected.

laser (another Toptica DL-pro). One of these diodes provides power for the Zeeman slower, the second to the 2D MOTs, and the third (seeded by the second master laser) provides power to the 3D MOT. There are two types of injection locked diodes in the current setup, which operate in a very different manner. The diodes used for the Zeeman slower and 2D MOTs do not have an antireflection (AR) coating on their output facet, while the diode used for the 3D MOTs does. All three are supplied (somewhat unwillingly) by the Nichia corporation (part numbers NDBA116T (AR coated) and NDB4216E (non AR coated)).

The diodes without the AR coating have a typical lasing threshold of roughly 30 mA. We




Figure 3.7: The 461 nm laser system consists of two master diode lasers and three slave diodes. The frequencies of the two diode lasers are stabilized to a polarization-rotation spectroscopy setup, and by a beatnote lock between the two lasers. The three slave diodes provide power for the Zeeman slower, 2D MOTs, and 3D MOTs respectively.

typically operate them around 140 mA. In the absence of seed light from the master laser, the diodes emit with an ill-defined frequency. When roughly 0.5 mW of seed light is applied the output power of the diode remains roughly the same, but the frequency of the emitted light matches that of the master laser if the current applied to the slave diode is within a certain range. These diodes output roughly 100 mW, giving them an effective gain of around 200.

The diode with the AR coating has a lasing threshold far above our operating current. In the absence of seed light, the diode emits negligible amounts of light. When seed light is applied, the diode acts to amplify it by a factor of around 5. A similar configuration is described in [103].

The non-AR coated diodes thus have a much higher effective gain than the AR coated diodes, making them a far more efficient way of amplifying optical power. They are also much cheaper (\$3300 instead of \$7600 at the time of purchase). The downside of the non-AR coated diodes is that they are much less stable to operate, and I would not recommend them for use at 461 nm.

The range of current, temperature and input power over which the diodes will injection-lock is quite narrow and also highly hysteretic. Drifts or transient changes in any of these parameters will cause the diode to either drop out of lock entirely or to operate in a partially injection-locked state, where only a portion of its power is at the correct frequency. Further, these properties vary greatly with individual diodes. We have one diode that will stay injection locked for days at a time. Other diodes, under seemingly identical operating conditions, drop at an unacceptably short timescale of a few hours or less. We have installed these diodes in Thorlabs laser diode mounts (TCLDM9), and also in a custom mount designed for thermal and mechanical stability. The behavior was largely unmodified with the new mount, and seems to be more a property of the diode than the mount. Other groups have reportedly had both better and worse success with these diodes. It is unclear whether this is due to variation between individual diodes or some element of the experimental design.

These frustrations lead us to replace the most critical of these diodes (the 3D MOT diode) with an AR coated version. Because of the lower gain of the AR coated diode, this required a second master laser devoted entirely to providing seed light for this diode. This second master laser is stabilized to the first master laser by a beat-note lock. This has the added advantage of allowing us to jump the frequency of the MOT beams quickly and precisely over at least 10's of MHz during the experimental sequence, which has proved useful for fluorescence imaging. Because the AR coated diode is operated below threshold, it cannot drop out of lock, and we can be confident that all emitted power is actually at the correct frequency.

3.5.2 Initial Slowing and Cooling




Figure 3.8: Closeup of AOsense source, with optics for Zeeman slower (far left) and 2D MOTs. The thin tube with bellows visible in the lower left corner feeds atoms into the main vacuum chamber.

Key Beam Parameters			
Role	Wavelength	Power	Beam Diameter
Zeeman Slower	461 nm	44 mW	8 mm
2D MOTs	461 nm	15 mW each	5x25 mm
Blue MOT	461 nm	13mW each	10 mm
Red MOT	689 nm	3 mW each	10 mm
$^3\text{P}_2$ Repump	707 nm	2 mW	5 mm
$^3\text{P}_0$ Repump	679 nm	2 mW	5 mm

Table 3.2: Cooling and trapping beam parameters.

The atoms begin at an initial temperature of roughly 500 C inside the AOSense oven, then exit the oven through a nozzle consisting of an array of narrow capillaries that collimate the atoms into a narrow beam. A portion of the atomic beam is slowed in the Zeeman slower to roughly 50 m/s, and is redirected into our main chamber by two 2D MOTs, whose axes are at small angles from the thermal beam. AOSense provides in-vacuum magnetics and optics for these. We provide light for the Zeeman slower and 2D MOTs from external optics, as shown in Fig. 3.8. Both Zeeman slower and 2D MOTs require circularly polarized light. We use powers of 44 mW in the Zeeman slower, and 15 mW in each of the 2D MOTs. All slowing light up to this point is near resonance with the $^1\text{S}_0$ to $^1\text{P}_1$ dipole allowed transition near 461 nm, which has a linewidth of roughly 32 MHz. Key beam parameters for lasers involved in slowing and cooling are summarized in Table 3.2.

At this point, the atoms have entered the main chamber, where we capture and cool them using a standard 3D MOT formed again on the 461 nm transition. The MOT beams here have a diameter of roughly 1 cm, limited by the optical access afforded by the cavity spacer (See Fig. 3.12 for geometry). The beams are retro-reflected after passing through the vacuum chamber, with the typical power in a single pass of around 13 mW. A 1X3 fiber splitter from Evanescence Optics is used to efficiently distribute the power.

3.5.3 Repump Lasers at 707 and 679 nm

The $^1\text{S}_0$ to $^1\text{P}_1$ transition is not perfectly cycling. There is a branching ratio of roughly 2×10^{-5} to the $^1\text{D}_2$ state [86], which eventually decays to $^3\text{P}_2$. To avoid this source of atom loss,

we apply a repump laser to the 3P_2 to 3S_1 transition near 707 nm. This pumps atoms into the 3P_1 and 3P_0 states. 3P_1 decays quickly (20 μ s) to the ground state, while 3P_0 is metastable. To avoid accumulation of population in 3P_0 , we apply a second repump laser to the 3P_0 to 3S_1 transition near 679 nm. Both lasers are homemade ECDLs using AR coated diodes from Sacher Laser.

The 707 nm repump laser is stabilized to an error signal derived from the same hollow cathode lamp as the 461 nm error signal. While this transition connects two excited states, there is sufficient population of 3P_2 in the hollow cathode lamp to stabilize a laser, presumably because of electron collisions. We have detailed the frequency stabilization technique in [102]. Unfortunately, the useful geometry of hollow cathode lamp for this application has become very difficult to find.

The 679 nm laser can either be left free-running, or can be locked via a frequency sideband derived from a fiber modulator to our stable reference cavity.

For ^{88}Sr , which has no nuclear spin and therefore no hyperfine structure, a single tone at each of 707 nm and 679 nm is sufficient to provide efficient depopulation of 3P_2 and 3P_0 . In ^{87}Sr , the presence of nuclear spin leads to hyperfine structure in 3P_2 , requiring multiple tones to efficiently depopulate this state. We generate these tones by sending the 707 nm repump light through a fiber modulator with high power-handling capabilities (AdvR WPM-K0707-P78P78ALO). We apply a tone generated by a DDS board (AD9914) to this modulator to create sidebands to address the various 3P_2 to 3S_1 transitions. One tone is applied at a time, with eight pre-programmed tones that we jump between at a rate of 100 kHz by toggling the DDS profile pins with an Arduino. We empirically find that tones at 225, 500, 550, 575, 800, 1150, and 1300 MHz lead to efficient depumping, with the carrier tone at roughly 40 MHz red relative to the 3P_2 to 3S_1 transition in ^{88}Sr because of various AOM offsets.

3.5.4 689 nm System and Narrow-Line Cooling

After initial slowing and cooling using the 461 nm transition, we use the narrow linewidth ($\gamma = 2\pi \times 7.5$ kHz) transition at 689 nm to provide further cooling and to load the atoms into our optical lattice. We use a new form of MOT that we have developed to do this, as described in




Figure 3.9: 689 nm laser system. An interference-filter ECDL from AOSense provides the master light for our 689 nm laser system. The laser is stabilized to our stable reference cavity by a Pound-Drever-Hall lock on a sideband that provides a tunable offset. Power for the MOT beams is increased using a tapered amplifier (TA). An AOM both switches the cooling light on and off, and allows us to sweep its frequency in a sawtooth ramp for efficient cooling. The 689 nm laser also provides optical pumping for state preparation, among other functions.

Chapter 10. This MOT uses standard polarizations and magnetic field configurations to operate, but the frequency of the MOT beams is swept in an asymmetric sawtooth waveform during the time from when the blue MOT is turned off until the atoms are loaded and cooled into the lattice. The frequency sweep goes to slightly above the transition between 1S_0 and 3P_1 $F = 11/2$, and covers a frequency range of 7 MHz every 50 μ s. Using this technique, we avoid the need for a stirring laser on the 1S_0 and 3P_1 $F = 9/2$ transition that is typically required.

The 689 nm laser system is shown in Fig. 3.9. An AOSense interference filter ECDL (IF ECDL) provides initial power to seed a tapered amplifier (Eagleyard EYP-TPA-0690, in JILA-made mount). The laser is stabilized to our reference cavity (see 3.5.6) with a variable offset provided by an EOM sideband. The power and frequency of the MOT beams are controlled by an AOM that shifts the light roughly 80 MHz to the red. Additional light is picked off both before and after the tapered amplifier to provide optical pumping, cavity probe tones, or for other experiment-dependent functions.

The red MOT beams are also distributed via a fiber splitter from Evanescence Optics. The red MOT beams pass through separate polarization optics from the blue beams, and are overlapped with and separated from the blue beams by dichroic beam-splitter plates before and after the atoms.

3.5.5 698 nm Laser System

The role of our laser system at 698 nm is two-fold. It is used to transfer population between 1S_0 and 3P_0 and to establish coherence on this transition, and also as a frequency reference against which to measure the frequency of superradiant emission. For population transfer, we use our own moderately stable laser (the local laser). For a frequency reference, we rely on very stable light from the Ye lab’s stable laser system (the Ye lab laser).

Our local laser is another IF ECDL from AOSense. It is stabilized to our stable reference cavity via a standard PDH lock, with a tunable fiber EOM providing a frequency offset (see 3.5.6). A beatnote between our local laser and the Ye lab laser indicates a short-term linewidth of 200 Hz (Gaussian FWHM) when we float the optical table on which the laser and cavity live. The relative linewidth is 2 kHz when the table is not floating. The 200 Hz FWHM represents an upper bound on the laser linewidth, as it includes noise from roughly five meters of unprotected optical fiber. A small portion of our 698 nm light is coupled to the science cavity via a second fiber-modulator. This allows us to apply agile frequency sweeps to a laser tone that addresses the atoms.

The light from the Ye lab comes from a laser stabilized to a 40 cm long ULE reference cavity [104], which goes by the name “Mike Martin.” We pick off a portion of this light to seed an injection-lock located in the Ye lab, and send the output to our lab through a 75 meter long fiber that runs through the ceiling. This injection lock (which uses a non-AR coated HL6738MG diode from Thorlabs) is very robust compared to the blue ones, and outputs 67 mW after an isolator. We use a typical seed power of 400 μ W, a home-built diode mount, and a JILA current and temperature controller.

To eliminate noise associated with vibrations in the fiber that connects the two labs (which contribute roughly 20 kHz of frequency deviation if left uncanceled, as inferred from the FWHM of

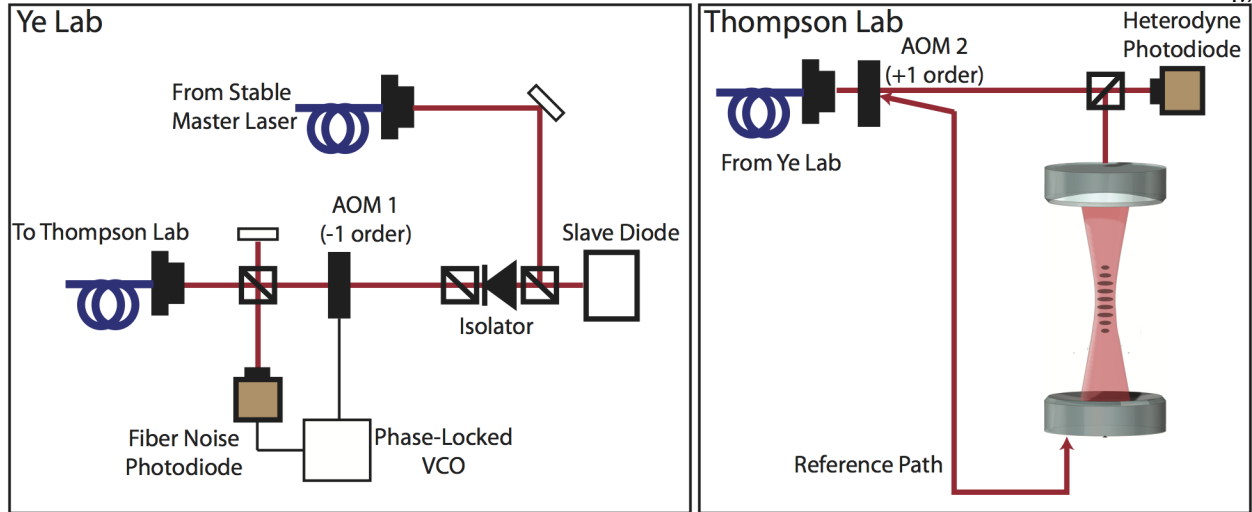


Figure 3.10: Fiber-noise cancellation system for clock reference light. This architecture involves two AOMs — one in the Ye lab (AOM1) and the other in the Thompson lab (AOM2). Light that is deflected by AOM2 (giving it a frequency offset f) is reflected off the back mirror of the science cavity, and sent back to the Ye lab. This light is overlapped with light that has not gone through the long fiber, and detected on a fast photodiode. The resulting beatnote (near $2f$ in frequency) is used to lock the frequency of a voltage controlled oscillator (VCO). The output of the VCO is divided by two and used to drive AOM1 at a frequency near f . Any noise present in the fiber leads to a perturbation of the frequency of the beatnote. Because we use opposite diffraction orders on the two AOMs, this frequency offset is cancelled by AOM1. Compared to simpler configurations that use a single AOM, this configuration has the advantage of being immune to unwanted reflections before AOM2.

an observed beatnote with our local laser), we retro-reflect a portion of the light back down the fiber and actively cancel the noise due to vibrations using the same architecture as demonstrated in [105]. In order to provide a stable reference for superradiant light emitted from one end of the cavity, we use the mirror on the opposite side of the optical cavity as a reference surface. This configuration is shown in Fig. 3.10, and provides cancellation of noise associated with both vibrations in the optical fiber and movement of the cavity. Doppler shifts from cavity motion contribute around 20 Hz of noise if left uncanceled.

There is a 2 m optical fiber that transfers reference light to the far side of the optical cavity. Acoustic noise in this fiber, as well as roughly 2 m of free-space path are imposed on the heterodyne beatnote. While it is possible that these parts of the path do contribute noise on our reference laser, we have found no significant increase in noise when we direct a fan over the optical table.

3.5.6 Reference Cavity




Figure 3.11: Transfer cavity from Stable Lasers Systems with four sets of mirrors sharing the same spacer.

All lasers that require precise frequency stabilization (those at 689 nm, 698 nm, 688 nm, and sometimes 679 nm), as well as the science cavity via the lattice laser, are stabilized via standard Pound-Drever-Hall (PDH) locks to a stable reference cavity from Stable Lasers Systems (SLS). SLS also provided a vacuum and isolation enclosure for the cavity. The reference cavity has high Finesse (roughly 25000) near both 700 nm and 813 nm. We measure a drift rate of 7 kHz/day, sufficiently low that we typically do not need to account for its drift during a typical dataset. The reference cavity spacer (a cylinder of 100 mm diameter and length, made from ultra low expansion glass) is temperature stabilized near its point of zero coefficient of thermal expansion (21C) by heaters surrounding the vacuum can. Unfortunately for us, the zero-crossing ended up being at room temperature for our spacer, meaning that we must bias a few degrees away from the optimal point to actively stabilize the temperature with heaters alone. Nonetheless, the measured drift rate is adequate for our needs. The spacer has four sets of mirrors mounted to it, which gives us four reference cavities to which we can stabilize lasers without overlapping beam-paths. We have found this to be advantageous to avoid issues of cross-talk from multiple lasers addressing the same

cavity.

The resonances of the reference cavities are at arbitrary frequencies, so we need a way to provide offsets to the laser locks of up to half a free spectral range of the reference cavity ($\text{FSR} = 3\text{GHz}$). To do this, we first send each laser through a fiber-coupled EOM (from EOSpace). We apply a tone to these EOMs to generate “offset sidebands” on the laser. We lock these offset sidebands to the reference cavity, and can tune the frequencies of the lasers by changing their frequencies. We set the modulation index to empirically maximize the size of the first order sidebands. After the fiber-coupled EOM, we send each laser through a free-space EOM. This free-space EOM is modulated near 20 MHz to provide PDH sidebands to lock the offset sideband to the cavity. The PDH sidebands on the different lasers are at slightly different frequencies (separated by roughly 1 MHz) in order to avoid coupling between the different locks.

3.6 The Optical Lattice

The optical lattice serves three key functions. First, it provides a way to trap many atoms within the cavity mode volume, and keep them there for the length of the experimental sequence. Additionally, the lattice localizes the atoms to a fraction of the optical wavelength along the axis of the cavity, eliminating the effects of Doppler shifts. Finally, the lattice laser provides a way to stabilize the length of the cavity. The lattice light forms a standing wave within the optical cavity, with regions of high intensity and low intensity separated by $813/2$ nm along the cavity axis. The atoms are attracted to regions of high intensity through optical dipole forces, which tightly confines them in the direction of the cavity axis. In the directions perpendicular to the cavity axis, the cavity mode has a Gaussian profile with a characteristic size of around $80\text{ }\mu\text{m}$. This leads to a much looser confinement of the atoms in the radial directions.

The energy shift associated with the optical dipole potential (also known as an AC Stark shift) arises from the lattice laser introducing off-resonant coupling between the state of interest (whichever state is populated, often the ground state) and other states. If there is a single transition involved, the shift of the lower-energy state is given by $\Delta E/\hbar = \Omega^2/4\Delta$, where Δ here is the

detuning of the laser frequency ω_L from the atomic transition frequency ω_a , $\Delta = \omega_L - \omega_a$, and Ω is the Rabi frequency associated with the light. If the laser is higher in frequency (blue detuned) than the atomic transition, the energy of the lower state is increased, while if the laser is lower in frequency (red detuned), the energy of the lower state is decreased, leading to an attractive potential.

Real atoms have more than one atomic transition, so to find the total energy shift associated with a state we must sum up the contributions from all relevant states. For our purposes, we want a potential that is attractive, and has the same magnitude for the 1S_0 and 3P_0 states. This ensures that the atomic transition frequency between these two states is unperturbed by the optical lattice. This is exactly the same requirement for strontium optical lattice clocks, and has thus been the subject of extensive research [101, 100]. It turns out that light near 813.4274(4) nm[106] fulfills these conditions, leading it to be referred to as a magic wavelength. Deviations from the magic wavelength lead to shifts of the clock transition of 0.07 Hz/(GHz μ K). For typical experiments, we

Typical Lattice Parameters	
transmitted lattice power (mW)	5.3
lattice depth (μ K)	100
longitudinal frequency (kHz)	170
radial frequency (Hz)	60
Lamb-Dicke parameter	0.17
atom temperature (μ K)	10
rms radial extent (μ m)	20

Table 3.3: Typical lattice parameters.

inject several mW of light into the optical cavity (which is described in the next section) to create a lattice with a depth of order 100 μ K. The lattice light is supplied by a home-built 813 nm diode laser, which uses an AR coated diode from Eagleyard.

Along the cavity axis, the atoms are tightly confined in the so-called Lamb-Dicke regime, where the effects of atomic motion are highly suppressed. This regime is of great importance to the operation of optical clocks, and so is discussed in great detail in early optical clock theses (for

example ref. [106]). Since this is pretty well established by now, I will provide a summary that highlights a few of the key implications rather than reproduce the full detail of the derivations here.

In the Lamb-Dicke regime, we can describe the atomic motion as a quantum harmonic oscillator along the direction of propagation of the light with which the atom is interacting. The Lamb-Dicke regime is quantified by the Lamb-Dicke parameter $\eta = kx_0 \ll 1$, where $x_0 = \sqrt{\hbar/2m\omega_m}$ is the extent of the ground-state atomic wavefunction along k and ω_m is the harmonic oscillator frequency in that direction. Equivalently, the Lamb Dicke parameter can be defined as $\eta^2 = \omega_R/\omega_m$, where ω_R is the recoil-frequency $\hbar k^2/2m$. The recoil frequency corresponds to the average change in kinetic energy of the atom when it absorbs or emits a photon. We can thus understand the Lamb-Dicke regime in two ways: the atom is localized to much less than the wavelength of the light, or the energy spacing of the harmonic oscillator is much greater than the recoil energy.

There are two limits in which the Lamb-Dicke regime is important. The first is when an atom absorbs or emits a photon quickly, for example through spontaneous emission or due to a strong drive. We are interested in the probability that the atom makes a jump from one harmonic oscillator level to another (assuming that the photon was emitted or absorbed in the direction of confinement, for simplicity). In this limit, we must calculate the matrix element associated with the transition from one harmonic oscillator state to another, coupled by the phase factor associated with the photon:

$$P(n \rightarrow n') \propto \left[\langle n' | e^{ikx} | n \rangle \right]^2 = \left[\langle n' | e^{i\eta(\hat{a} + \hat{a}^\dagger)} | n \rangle \right]^2$$

where \hat{a} , \hat{a}^\dagger are the raising and lowering operators for the atomic motion. In the Lamb-Dicke regime ($\eta \ll 1$), we can Taylor expand this to see that the probability of changing harmonic oscillator states $P(n \rightarrow n') \propto \eta^2$ for $n' = n \pm 1$, and is even smaller for larger changes in motional quantum number. Changes in the atom's axial quantum number are greatly suppressed when absorbing or emitting a photon in the Lamb-Dicke regime.

In fact, the effects of the Lamb-Dicke regime are even more pronounced when the Rabi frequency Ω_d associated with the light that drives the transition between internal states is much

lower than the harmonic oscillator frequency of the atom, as is typically the case for optical clocks. In this regime, not only does the photon lack the spatial variation to efficiently connect different harmonic oscillator states (the previous effect described), but it may be off resonance from all $n \rightarrow n'$ transitions but one, as each transition that is accompanied by a change in the harmonic oscillator level is separated by ω_m . This regime is referred to as the resolved sideband limit. For example, if the drive is tuned on resonance with the $n \rightarrow n$ transition, then the probability of exciting to $n' = n \pm 1$ is reduced by an additional ω_a^2/Ω_d^2 compared to the case of instantaneous transitions. We typically operate in this limit, meaning that transitions between different harmonic oscillator states in the axial direction are quite negligible. However, the finite value of the Lamb-Dicke parameter does lead to a slight reduction in the atom-cavity coupling g .

While we usually wish to operate in the Lamb-Dicke regime, for some applications, we wish to maintain radial confinement while allowing the atoms to move freely along the cavity axis. We can accomplish this by driving longitudinal modes of the cavity with opposite parity about the center of the cavity, such that the maxima of one mode are aligned to the minima of the other. In practice, we do this by creating frequency-modulation sidebands on the lattice laser using a fiber-coupled electro-optic modulator (EOM). If these frequencies are equal to the free spectral range (FSR) of the cavity, then we primarily drive one mode with the carrier tone, and two modes with opposite parity with the sidebands. By adjusting the sideband amplitude, we can flatten out the potential experienced by the atoms. To achieve our desired lattice depth we need more power than the EOM can handle, so we use a tapered amplifier (TA) after the EOM to increase the total available power.

We have measured atomic lifetimes in the lattice of up to 1.1 seconds when all atoms are in the ground state, though this value depends in a sensitive manner on the exact operating conditions of the lattice laser. Because of the relatively narrow linewidth of the optical cavity even at 813 nm, small frequency excursions of the lattice laser can lead to significant heating. Small tweaks to the laser current seem to dramatically alter the frequency noise properties of the laser. Because we are most sensitive to intensity fluctuations at even multiples of the axial trap frequency (typically several hundred kHz), suppressing these fluctuations with electronic feedback is difficult.

One annoying detail of the far-detuned optical lattice is that because the wavelength of the lattice differs from the wavelength associated with the forbidden transitions that we wish to couple to the cavity, the coupling between the atoms and the cavity is inhomogeneous. The atoms are trapped at antinodes of the optical lattice, which are spaced by $813/2 = 406.5$ nm, while the magnitude of the electric field associated with the cavity modes at 689 and 698 nm oscillates with a spacing of 344.5 and 349 nm, respectively. The location of the atoms and the cavity mode antinodes thus get out of phase with each other on a roughly $1\text{ }\mu\text{m}$ length-scale, and so over the roughly 1 mm extent of the atomic cloud can be treated as random. To account for this effect, we can treat the atoms as though they are randomly distributed along the cavity axis, and assign a modified coupling rate to atoms at different locations x : $g = g_0 \cos[kx]$, where k is the wavenumber of the transition of interest (though we typically average over a full cycle anyway so it does not really matter what k refers to), and $2g_0$ is the vacuum Rabi frequency for an atom located at an antinode of the cavity mode associated with the coupling. Depending on the experiment being performed, this inhomogeneous coupling will lead to either a simple renormalization of quantities of interest, or to qualitatively different dynamics. I will treat these cases as they arise.

3.7 The Optical Cavity

The heart of the experiment is the optical cavity. The cavity serves two main functions. First, it supports the magic wavelength optical lattice, and second it provides a mode with which the atoms interact via either of the narrow linewidth transitions at 689 nm and 698 nm. The cavity was designed to achieve a moderately high cavity cooperativity parameter ($C \simeq 0.4$) at 689 nm and 698 nm. This value was chosen with the goal of maximizing the collective cooperativity parameter NC within other technical constraints. The cavity has a lower (but still substantial) finesse at 813 nm to enable a deep lattice with moderate laser power. The cavity parameters are summarized in table 3.4.

We could have increased C (which would have been nice) by either using a smaller mode waist or higher finesse cavity. The smaller waist would be accompanied by a higher atomic density and




Figure 3.12: Optical cavity assembly and vibration isolation. The cavity consists of two mirrors separated by 4 cm. The axis of the cavity is horizontal. The mirrors are mounted to piezoelectric actuators (PZTs), which allow for control of the cavity length. To decouple the cavity from acoustic vibrations, the cavity spacer sits on viton balls on top of a support block, which also sits on viton balls.

collision rate, which we thought would be detrimental. Using a higher finesse would have decreased the cavity linewidth (for a fixed length). In doing so, we would have sacrificed the favorable ratio $\kappa/\gamma \gg 1$ on the 7.5 kHz linewidth transition, though likely benefitted from the higher product NC for work on the milliHertz linewidth transition.

We chose a relatively long cavity length of 4 cm. Several factors dictated this choice. First, we wanted to ensure ample optical access for MOT and other beams. We also chose to recess the cavity mirrors into the spacer to avoid a direct line-of-sight from the atomic source, as we were worried about coating the cavity mirrors with strontium. These constraints encouraged a relatively long cavity.

Additionally, we wanted to keep the free-spectral range (FSR, the frequency spacing between modes of the cavity, equal to $c/2L$) relatively small. For superradiance experiments using the milliHertz transition, a mode of the cavity near 698 nm must be made resonant with the clock transition. At the same time, we want to keep the lattice as close to the magic wavelength as




Figure 3.13: Mirror assembly. The cavity mirrors are mounted to piezo actuator tubes (PZTs) by an intermediate macor hat. This configuration allowed us to recess the mirrors into the cavity spacer in order to avoid coating them with strontium, and provided an easy way to align the mirrors. The two holes at either end of the cavity spacer in which the mirrors sit were drilled at the same time, ensuring that they share a common axis. All other alignments are referenced to those holes. All connections are made by very thin beads of Torrseal epoxy. Electrical connections to the inner and outer surfaces of the PZT tubes are made by custom clips (not shown) that attach through a slit in the macor hat (bottom, in image).

possible. In the worst-case scenario, the lattice laser is forced to be one half of a FSR from the magic wavelength. Reducing the FSR thus minimizes the expected frequency offset of the lattice relative to the magic wavelength.

To compensate for this magic wavelength offset, one could simultaneously drive cavity modes at $0, \pm 2$ FSRs. For the correct relative power between these modes, one could create an effective magic wavelength lattice. This leads to the practical constraint that the FSR should be less than half of the bandwidth of available phase-modulators, which would be used to generate the additional tones. This factor ultimately constrained our cavity length to its chosen value.

The cavity mirrors were supplied by Advanced Thin Films, with transmission specifications of 95 parts per million (ppm) at 689 nm, 106 ppm at 698 nm and 1339 ppm at 813 nm. The radius

Cavity Parameters			
	698 nm	689 nm	813 nm
C	0.4433(4)	0.4230(4)	NA
$2g_0/2\pi$ (Hz)	4.04 (spec)	11,000 (spec)	NA
L (cm)	4.03517(1)		
finesse F	25,570(20)	24,700 (100)	2346
$\kappa/2\pi$ (kHz)	145.3(1) (measured)	150.3(4) (measured)	1600 (spec)
mode waist $w(\mu\text{m})$	73.7	73.3	79.6
mirror transmission (ppm)	106 (spec)	95 (spec)	1339 (spec)
mirror losses (ppm)	30 (estimated)		
FSR (MHz)	3714.750(1)		

Table 3.4: Key parameters of optical cavity.

of curvature of the mirrors is 5 cm. We would expect to achieve a finesse of 33,000 at 689 nm, and 29,500 at 698 nm. We measured cavity linewidths of 150.3(4) kHz and 145.3(1) kHz at 689 nm and 698 nm respectively, implying finesses of 24,700(100) and 25,570(20). These measurements imply losses of around 30 ppm per mirror.

The length of the optical cavity can be tuned by roughly two FSRs using piezo actuators (PZTs) on each of the two mirrors. The PZTs are cut from tubes of PZT-5H material supplied by Piezomechanic GMBH, and are driven with a voltage of 0-120 V. The cavity mirrors nest inside of the PZT tubes and are recessed into the cavity spacer. This assembly is shown in Fig. 3.13, and described in detail in the figure caption.

To enable additional slow tuning of the cavity length, we have installed three automotive tail-light bulbs that warm the cavity through the large viewport. These are controlled by a tunable power supply, and can tune the cavity length by several FSR, enabling better alignment of the lattice to the magic wavelength. The settling time associated with this tuning is very slow (several hours), so we have not been using them recently.

The ability to tune and stabilize the resonance frequency of the science cavity is critical for our experiments. A convenient way of doing this is to use the optical lattice.

The lattice laser frequency is locked to the resonance frequency of the science cavity using a standard PDH lock. This lock is performed with fairly high bandwidth (roughly 100 kHz unity gain

frequency). Some power from the lattice laser is split off and sent through a fiber-coupled EOM to the stable reference cavity. As described earlier, this EOM generates offset sidebands at frequencies of up to a few GHz. Additionally, we use a free-space EOM to apply sidebands at roughly 10 MHz, which are used to generate PDH error signals on the carrier and each of the offset sidebands. We stabilize the frequency of the science cavity to the reference cavity by feeding back on the science cavity PZTs based on a PDH error signal generated from one of the offset sidebands. By tuning the frequency of the offset sideband, we can then control length of the science cavity in a well-defined manner.

The bandwidth with which the science cavity frequency is stabilized relative to the reference cavity is limited by the PZTs to which the cavity mirrors are mounted. Before installing the science cavity, I measured the resonance response of the cavity mirror/PZT system by forming a Michelson interferometer that included one of the cavity mirrors. The first mechanical resonance appeared at 10 kHz, which limits feedback bandwidth to the PZTs to roughly 200 Hz before oscillations become apparent at 10 kHz, if a simple PI controller is used. This dramatically limits the degree to which we can suppress harmonics of 60 Hz that appear on the output of our high voltage source (unless we decided to do something fancy [107]). These harmonics lead to cavity frequency excursions of order 5 kHz. These have not been a problem for us so far, but could be improved for future applications.

Chapter 4

Strong Collective Coupling on 7.5 kHz Linewidth Transition and Nondestructive Atom Counting

The first scientific project on the new system was the first demonstration of strong collective coupling between an optical cavity and atoms, via a forbidden optical transition [20]. Typically, strong transitions are used in cavity QED experiments to couple the atoms to the cavity field. In the optical domain, this means using dipole-allowed transitions, while in the microwave domain it means using Rydberg atoms with large dipole moments. Here, we instead demonstrate strong collective coupling using the dipole forbidden 1S_0 to 3P_1 transition in ^{88}Sr . This transition has a wavelength of 689 nm, and a decay linewidth of 7.5 kHz, roughly a factor of 1000 slower than a dipole allowed transition of the same wavelength.

Demonstrating strong collective coupling in our system is of particular interest, as strontium atoms (and other Alkaline Earth-like atoms with similar structure) have proved very useful for quantum metrology [108]. Example technologies include optical lattice clocks [55, 27, 29] and ultra-narrow lasers [109, 110, 111, 112], along with their associated broad range of potential applications [113, 114, 115, 116, 117, 118].

Strong coupling on dipole-forbidden transitions in a system like ours enables state-selective, non-destructive counting [119] of these useful species of atoms. Such counting has been used to generate highly spin-squeezed states that surpass the standard quantum limit on phase estimation [33, 37, 120, 121, 122, 123, 124]. If applied to optical lattice clocks, this could provide a rare opportunity to use entanglement to enhance the performance of a range of precision measurements.




Figure 4.1: (a) Experimental diagram. ^{88}Sr atoms are confined by 1D optical lattice in high-finesse cavity. π polarized probe light is coupled through cavity and is detected on a single-photon counting module. This probe interacts only with the 1S_0 to 3P_1 , $m_J = 0$ transition. A magnetic field B shifts the $m_J = \pm 1$ Zeeman sub-levels by ± 15 MHz. (b) Relevant level structure of ^{88}Sr , showing dipole forbidden transitions at 689 nm and 698 nm, and dipole allowed transition at 461 nm.

More simply, but quite importantly, non-destructive readout methods [125] can reduce the highly deleterious effects of local oscillator noise aliasing [126, 35] that limit current state of the art optical lattice clocks.

In this work, we utilize the vacuum Rabi splitting (a key signature of strong collective coupling) to make non-destructive differential atom number measurements at the equivalent of sub-projection noise sensitivity. We verify that the number of photon recoils imparted per atom to reach this sensitivity is much less than one (0.01). This is a key criteria for non-destructiveness, as it characterizes the degree to which atoms are heated by the measurement regardless of trap depth, as well as the measurement's applicability for the generation of spin-squeezed states[88].

Cavity-enhanced nonlinear spectroscopy [127] of the same transition has been performed in a MOT in [98, 97], but inhomogeneous doppler broadening prevented the observation of a collective vacuum Rabi splitting. Non-fluorescence based atom-counting in ^{87}Sr [125, 126] and ^{171}Yb [128] have been performed using the dipole-allowed transition 1S_0 to 1P_1 and no optical cavity. Atomic

projection noise level sensitivity was achieved, but with the scattering of ~ 100 photons per atom [125, 126] and with uncharacterized scattering [128]. While [125, 126] demonstrated that the atoms were retained in the lattice trap, the need to re-cool the atoms leads to additional stochastic losses [129]. Finally, at values of $\gtrsim 0.5$ recoils per atom one would be unable to generate entangled states for fundamental enhancements in measurement precision beyond the standard quantum limit [130, 88].

4.1 Demonstration of Collective Strong Coupling

As noted in Chapter 2, there are certain aspects of strong collective coupling that become hard when using a dipole-forbidden transition, and others that do not. For the work we present here, two conditions of the strong collective coupling regime must be met. The first is that the collective cavity cooperativity parameter $NC = \Omega^2/\kappa\gamma \gg 1$. As a reminder, $\Omega = 2g\sqrt{N}$ is the collective vacuum Rabi frequency, κ is the cavity linewidth and γ is the atomic transition linewidth. This condition is required in order to perform atom-counting measurements compatible with spin squeezing. Because $g \propto |\vec{d}|$ and $\gamma \propto |\vec{d}|^2$, where \vec{d} is the dipole matrix element of the transition which is much smaller for dipole-forbidden transitions than for dipole-allowed transitions, C is actually independent of the strength of the transition, and attaining $NC \gg 1$ is not harder when using a dipole-forbidden transition.

A second requirement for the observations that we present here, and a second definition of the strong collective coupling regime is that the collective vacuum Rabi frequency must exceed relevant decoherence rates in the system: $\Omega \gg \kappa, \gamma$. This requirement allows for the observation of a resolved vacuum Rabi splitting. Fundamentally, $\Omega \gg \gamma$ is actually easier to achieve on a weak transition, as $g \propto |\vec{d}|$ and $\gamma \propto |\vec{d}|^2$. Achieving $\Omega \gg \kappa$ is harder, and requires the use of a relatively high-finesse optical cavity. The collective vacuum Rabi frequency must exceed other non-fundamental sources of broadening as well, notably Doppler decoherence. This requirement is much harder to meet for a narrow linewidth transition, and has prevented the observation of this regime in previous works [98, 97].

To confirm that we had achieved this more stringent definition of the strong collective coupling regime, we show that we can observe a highly-resolved splitting of the normal modes of the coupled atom-cavity system, known as a collective vacuum Rabi splitting [131, 132]. This constitutes a direct verification of the hierarchy $\Omega \gg \kappa, \gamma$.

To do this, we begin by loading atoms into the lattice as described in the experimental apparatus section. Because we used ^{88}Sr for this experiment, which has a unique ground state, no further state preparation was required. We probe the atoms with a weak tone coupled to the cavity that is polarized along the direction of an applied magnetic field, as shown in Fig. 4.1a, and measure transmitted power on a single photon counting module (SPCM).

For this work, the probe laser frequency was stabilized to a transfer cavity whose length is actively stabilized by saturated absorption FM spectroscopy on the $^1\text{S}_0$ to $^3\text{P}_1$, $m_J = 0$ transition. A phase-modulation sideband of the 813 nm lattice laser was locked to the stabilized transfer cavity, providing a tunable offset frequency. The science cavity was frequency stabilized to the lattice laser.

At 689 nm, the cavity has a finesse of $F = 2.4(2) \times 10^4$. The atom-cavity coupling is enhanced by using up to $N = 1.25 \times 10^5$ atoms, such that $\Omega = \sqrt{N}2g = 2\pi \times 6.3$ MHz, where $g = 2\pi \times 9.2(1)$ kHz [133, 134]. The quoted N and g are the effective atom numbers and couplings used to account for inhomogeneous coupling of atoms to the standing wave probe mode, and for the finite extent of the atomic ensemble [121, 122]. The effective coupling is related to the maximum possible coupling g_0 by $g = \sqrt{3/4}g_0$. The effective atom number is related to the total number of atoms in the lattice N_{tot} by $N = 2N_{\text{tot}}/3$. Using these definitions consistently reproduces the size of the vacuum Rabi splitting and the magnitude of atomic projection noise as though all atoms were homogeneously coupled. A quick note on cavity linewidths: at the time that I performed this work, I had measured the cavity linewidth to be 160(16) kHz. A later, more precise measurement put the linewidth at 150.3(4) kHz. In this section, I will use the previous measurement, as that is what appears in the published version. It is important to note here that the key quantitative claims made in this work depend on g , which is independent of κ , and whose value we can calculate based on well-known geometric factors up to our knowledge of the transition linewidth γ .

Inhomogeneous doppler broadening is highly suppressed along the cavity axis by confining the atoms in the 1D optical lattice to much less than the 689 nm probe wavelength (Lamb-Dicke parameter of 0.16). We estimate that inhomogeneous transition broadening from the lattice is at most 44 kHz rms. This broadening results from the fact that 813 nm is not a magic wavelength for the 1S_0 to 3P_1 transition, which leads to polarization-dependent shifts [135]. Since the inhomogeneous broadening is much less than Ω , its impact on the dressed modes is negligible [136].

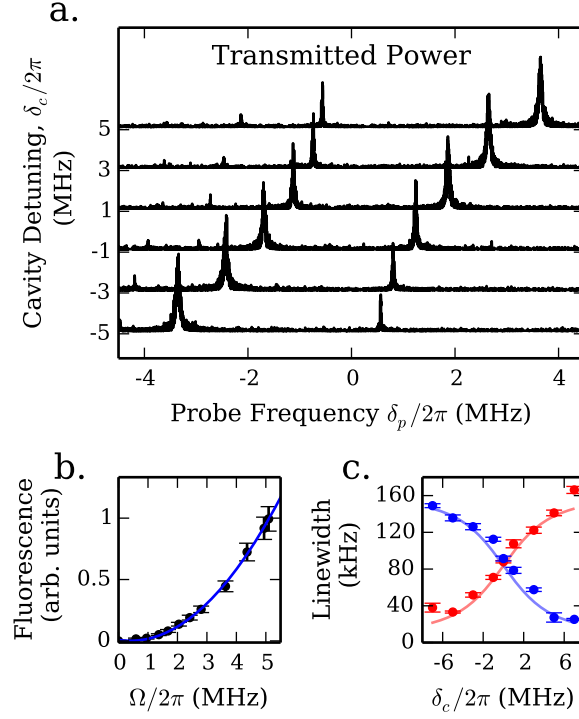


Figure 4.2: Observation of the collective strong coupling regime on the forbidden 1S_0 to 3P_1 transition. (a) We characterize the vacuum Rabi Splitting by recording transmitted power as probe frequency is swept. Traces correspond to different cavity detunings, δ_c . (b) Detected fluorescence versus vacuum Rabi splitting Ω . (c) Fitted linewidths κ'_\pm of the two resonances at ω_+ (red) and ω_- (blue) versus δ_c , with prediction based on known atomic linewidth γ , measured Ω and cavity linewidth κ .

Figure 4.2a shows a vacuum Rabi splitting obtained by tuning the empty cavity resonance ω_c near the atomic transition frequency ω_a , loading atoms into the lattice, and then recording the transmitted power as a cavity probe is linearly swept in frequency. The two transmission peaks correspond to the new normal modes of the dressed atom-cavity system at frequencies ω_+ and

ω_- . The different traces correspond to different detunings $\delta_c = \omega_c - \omega_a$, and exhibit the expected avoided-crossing behavior $\omega_{\pm} = (\delta_c \pm \sqrt{\delta_c^2 + \Omega^2})/2$ (see ref. [88] for a nice review of the features of a vacuum Rabi splitting). In Fig. 4.2b, we confirm the \sqrt{N} scaling of the size of the normal mode splitting with atom number by measuring the frequency splitting of the two modes at zero detuning, and comparing to the subsequently measured free-space fluorescence signal when 461 nm light is applied from the sides of the cavity.

The fitted widths κ'_{\pm} of the upper and lower transmission peaks are plotted in Fig. 4.2b, along with a theoretical prediction based on the known atomic decay rate γ and the measured cavity linewidth κ [131]. The width of each feature is a weighted average of the atomic and bare-cavity linewidths [88]. The more cavity-like the normal mode is, the closer its width is to κ , while a more atom-like mode has width closer to γ . The fact that the splitting between the normal modes is always much larger than the width of either mode indicates that we have indeed achieved the hierarchy $\Omega \gg \kappa, \gamma$.

4.2 Nondestructive Atom Counting

We next demonstrated that the strong coupling achieved here could in principle be used to achieve spin squeezing. By measuring the size of the vacuum Rabi splitting, we can infer the atom number N with great precision. If we can do so in a manner that does not reveal single-atom information (or equivalently without scattering photons into free space), then this can be used to create an entangled state. Similar techniques have been used in systems with dipole-allowed transitions to generate large amounts of spin squeezing [33, 15, 14].

We characterize the non-destructiveness of our differential population measurement by the average number of photon recoils per atom m_s imparted while achieving a measurement imprecision equal to the projection noise level of the atoms, m_s^{PN} [36, 137].

At the time that we performed these measurements, we did not have the ability to drive the atoms into a superposition state on the clock transition. Even now, we can only do so by taking advantage of a cavity resonance near the clock transition wavelength, which is incompatible with

the measurement scheme we demonstrate here. As a result, the atoms are not in a superposition of states, so there is no atomic projection noise. However, to give a point of comparison, we define the projection noise level as though our N atoms were actually $2N$ atoms, each in a superposition of the ground state and the excited 3P_0 clock state, represented in Fig. 4.1b. The rms quantum projection noise in the ground state population would then be $\Delta N_{PN} = \sqrt{N/2}$. We consider the rms noise $\Delta N = \Delta(N_f - N_p)$ in the difference of two consecutive measurements N_p and N_f of the population in the ground state 1S_0 . We define the spin noise reduction as $R = (\Delta N / \Delta N_{PN})^2$. Two figures of merit for the measurement are the minimum R attainable, and m_s^{PN} .

Two consecutive measurements of the vacuum Rabi splitting Ω , with measurement outcomes labeled Ω_p and Ω_f , are used to infer the populations of the ground state N_p and N_f . For $\delta_c = 0$, the rms fluctuations ΔN_{PN} lead to rms fluctuations in the measured vacuum Rabi splitting $\Delta \Omega_{PN} = g/(\sqrt{2}) = 2\pi \times 6.48(6)$ kHz, regardless of atom number N , provided $\Omega \gg \kappa, \gamma$ [88]. The spin noise reduction can then be written in terms of measured and known quantities as $R = (\Delta(\Omega_f - \Omega_p) / \Delta \Omega_{PN})^2$. The vacuum Rabi splitting frequency $\Omega = \omega_+ - \omega_-$ is determined by measuring the transmitted power when probe light is scanned over cavity resonance.

To gain common mode rejection of frequency noise between the cavity and probe laser, we use two frequency components or tones to simultaneously probe the two dressed cavity modes, as shown in Fig. 4.3a [122]. We create the two probe frequencies $\omega_{p\pm}$ by amplitude modulating a fixed probe laser whose center frequency is resonant with the empty cavity. As we ramp the modulation frequency ω_m over 1 MHz in 40 ms, the lower and higher frequency AM sidebands $\omega_{p\pm}$ simultaneously sweep across the resonances at ω_{\pm} .

When the probe tones are on the side of the resonance feature of the dressed modes, changes in the number of atoms in 1S_0 cause first order changes in the total transmitted power. In the small noise limit, relative frequency noise between the probe laser and the empty cavity causes no net change in the total transmitted power. To demonstrate the noise suppression of this two-tone technique, Fig. 4.3b shows the measured power spectral density of cavity-probe frequency noise, using a single sideband to probe the bare cavity resonance versus using the two tones $\omega_{p\pm}$

simultaneously. The single tone probing exhibits broad band noise and resonances, likely due to voltage noise on the science cavity's piezo driver, that are strongly suppressed using the two-tone technique.

To demonstrate the probing method with atoms, we perform sweeps as described above and fit the total transmitted power to a Lorentzian. We then extract the portion of the time data that corresponds to probing on side of resonance, defined as the portion between $.5$ and $.75 \times \kappa/2$. To convert changes in transmitted power into changes in frequency, we fit a linear slope S of transmitted power versus the known probe frequency. The time data is broken into bins of length τ , and the average transmitted power, P_i is computed within each time bin i . Differences are taken between adjacent bins, and the desired standard deviation $\Delta\Omega = \Delta P/S = \Delta(P_{i+1} - P_i)/S$ is computed over non-overlapping pairs of bins.




Figure 4.3: (a) To probe the Vacuum Rabi Splitting, two probe tones at $\omega_{p\pm}$ are simultaneously swept across the two normal mode resonances at ω_{\pm} while total transmitted power is recorded. (b) The noise power spectra relative to photon shot noise for cavity probe with a single tone (red trace) and two tones (black trace), demonstrate noise cancellation of two-tone probing.

Fig. 4.4a shows measurement noise relative to hypothetical projection noise R as a function of the chosen window length τ at $N = 9 \times 10^4$. Similar results were obtained for atom numbers between $N = 1.5 \times 10^4$ and 1.25×10^5 . At a given atom number, the incident probe power is held

fixed at a value that does not saturate the detector or the atomic ensemble. Therefore, increasing τ is equivalent to collecting more probe photons to average down photon shot noise. At large τ , R rises due to uncanceled low frequency probe noise. For window lengths between 50 and 175 μs , corresponding to detected photon numbers of 750 and 2650 per window, we achieve measurement noise $R < 1$. The lowest value measured is $R = 0.58(13)$, without any detector or shot noise subtractions applied.

The goal of being in the bad-cavity, strongly-coupled regime is to achieve a high ratio of detected photons to free-space scattered photons, which lead to atomic state collapse, heating, and potentially atom loss. We characterize these effects by measuring how passing photons through the cavity leads to changes in Ω , as shown in Fig. 4.4b. Changes in Ω allow us to infer changes in the product $g\sqrt{N}$. These changes could either be due to atom loss or heating in directions perpendicular to the cavity axis. We make our observations at much higher transmitted probe photon numbers M_t than those used for observing $R < 1$ by inserting a third sweep between two probe sweeps used to measure a change in Ω . From the fit to the change in Ω^2 vs M_t , we extrapolate to the typical number of transmitted photons $M_t \approx 5000$ in a single window that achieves $R \leq 1$. At $N = 5 \times 10^4$ atoms, we find 0.01 photons scattered per atom/measurement window. This corresponds to a temperature increase for two windows of 8 nK, or in the context of generating entanglement, a loss of atomic coherence of only 1% or 0.09 dB of signal loss due to the first measurement.

We theoretically expect one free-space scattered photon per 11(1) photons transmitted from a single end of the cavity, a value set simply by the ratio $\kappa/2\gamma$. The resulting change in Ω^2 predicted from radial recoil heating due to scattering is shown as a dashed line in Fig. 4.4b, and is consistent within the error of the fit to the data.

Because we operate with $\kappa \gg \gamma$ there is little fundamental loss for generating atomic squeezing or for non-destructive readout from probing with $\delta_c = 0$, and applying the noise-immune two-tone technique [88]. This would not be true for the current cavity geometry and finesse on a fully allowed dipole transition as the magnitudes of the atomic damping and cavity damping would be reversed $\kappa \ll \gamma$. This is a key advantage of working with a dipole-forbidden transition.

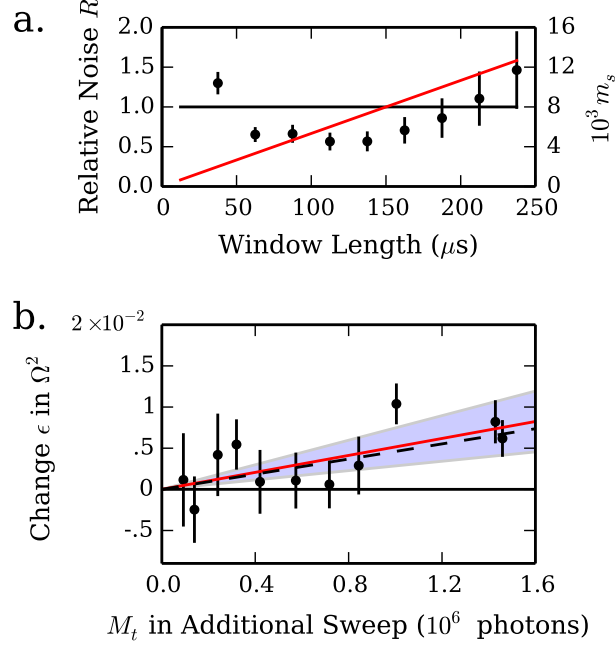


Figure 4.4: (a) Black points show measurement noise relative to hypothetical projection noise in adjacent windows R , versus length of a measurement window. The lowest two points represent $R = -2.4_{1.1}^{0.9}$ dB. Red line shows the the number of scattered photons per atom m_s in a single measurement window. (b) Fractional change in $\Omega^2 = 4g^2N$ due to additional sweep $\epsilon \equiv (\Omega_0^2 - \Omega_s^2)/\Omega_0^2$, versus M_t , the number of photons in the additional sweep. Ω_s and Ω_0 are the measured vacuum Rabi splittings with and without the additional scattering sweep. Solid red line is fit to data, with shaded region representing 1 sigma error on the fit. Dashed line represents a prediction based on the theoretically predicted free space scattering. The number of photons used in two typical 100 μs windows is represented by the thickness of the left axis.

For certain applications, including spin squeezing, another relevant metric for the non-destructiveness of a measurement is the degree to which it imparts inhomogeneous AC Stark or light shifts of the ground state 1S_0 . While we cannot measure this directly, we can estimate it by considering the AC Stark shifts associated with the two probe tones. For typical probing parameters, the time integrated phase shift contributed by a single probe tone during a measurement window would be of order π radians for maximally coupled atoms. With our two-tone probing scheme, the phase shifts cancel due to their opposite detunings from the excited state 3P_1 [88]. If we assume a power balance at the 10% level, we could expect a phase shift of order 0.1 radians for maximally coupled atoms, such that the loss of atomic contrast would be $< 0.5\%$ or 0.04 dB loss of

signal. This cancellation would potentially eliminate the need for a spin-echo sequence to maintain contrast in an experimental sequence. Cavity-optomechanical effects that can limit the observed spin noise reduction would be partially canceled as well [33].

From m_s^{PN} , we can estimate the fundamental limit of our probing scheme [88]. If technical sources of noise were addressed, one could achieve 22 dB of spin squeezing at 2×10^5 total atoms. This fundamental limit on spectroscopic enhancement could reach 30 dB by improving the net effective quantum efficiency for detecting the probe light from the current 3.5% to a value of 25%. For context, Kevin Graham and Josh obtained a quantum efficiency of 37% on the rubidium experiment [15].

4.3 Conclusions

While we were able to demonstrate an atom-counting method that is fundamentally compatible with spin squeezing, significant challenges remain to actually use it to create spin squeezing. The primary challenge is to create the necessary superposition on the optical clock transition. In addition to a very stable laser, this requires us to address the atoms in a Doppler-shift-free manner. We can currently do so only along the cavity axis, where the atoms are tightly confined. This requires us to shine the clock light through the cavity mirrors, which only works well when the cavity is on resonance with the clock transition at 698 nm. Because the atom-counting relies on placing a cavity mode very near resonance with the 1S_0 to 3P_1 transition at 689 nm, we would have to get very lucky for this to work out with a single cavity.

A way around this would be to use a two-dimensional lattice to provide a direction other than the cavity axis along which the atoms could be addressed by the clock laser in a Doppler free manner. This idea is discussed further in Chapter 12.

Chapter 5

Magnetically Induced Transparency

In this work, we explore a new cavity-enhanced spectroscopy technique for stabilizing a laser to narrow linewidth optical transitions. We call this technique “magnetically induced transparency” (MIT) in analogy to the much-studied phenomena of “electromagnetically induced transparency” (EIT) [38, 39, 40]. MIT allows us to create a narrow spectroscopic feature from an ensemble of atoms, whose frequency is set primarily by atomic properties, making it suitable for the stabilization of optical frequencies. EIT can also be used to generate a narrow-linewidth feature, but its frequency is set by that of an applied laser, making it unsuitable for optical frequency stabilization. It is the use of narrow-linewidth, dipole forbidden transitions that makes such applications possible with MIT. While this work focuses primarily on the utility of MIT for frequency metrology, it may also be suitable for the photon-storage type applications that have made EIT so famous [41, 42, 43, 44, 45]. I will provide a brief discussion of these applications at the end of this chapter.

As with much of the work in this thesis, this project is motivated by the quest to create lasers with ever-higher frequency stability, as well as novel kinds of absolute frequency references. Currently, the frequency stability of state-of-the-art lasers is limited by thermal fluctuations in the reference cavity mirror coatings, substrates, and spacer [59]. Further, the frequency of a laser is not fixed in absolute terms — the reference cavity frequency is arbitrary and can drift in an unbounded manner. One approach to getting around these limitations is to create systems whose frequency is referenced to narrow-linewidth optical atomic transitions. These techniques can take several forms: active frequency references (superradiant lasers [63, 24, 25, 51]), which will be discussed in detail

in later chapters, and passive frequency references.

Depending on one's desired tradeoffs between accuracy, complexity, and bandwidth, passive optical frequency references may take the form of a fully fledged optical clock, operated in a pulsed mode, or a continuously operated interferometer using an atomic beam [138]. These free-space beam solutions are appealing for their simplicity, but face challenges resulting from motional effects, limited optical depth, and the difficulty of collecting photons emitted in random directions. To mitigate these effects while keeping the system simpler than a state-of-the-art clock, cavity enhanced spectroscopy techniques have been proposed and developed for narrow-linewidth optical transitions [127, 98, 139].

In these techniques, an ensemble of cold atoms is placed inside an optical cavity, such that light incident on the cavity will bounce back and forth many times through the atoms, increasing the effective optical depth. Because the light is incident from both sides, certain errors resulting from atomic motion are suppressed. Further, because the cavity imposes a resonance condition on the system, the spectroscopic features that emerge are the result both of atomic absorption and of dispersive phase shifts caused by the atoms. MIT represents a new variant on these cavity-enhanced spectroscopy techniques.

5.1 Basic Mechanism

Very roughly (a detailed description follows), our technique works by placing atoms inside the mode of an optical cavity that is nearly resonant with a set of transitions between a unique ground state and several Zeeman sublevels of an excited state. A probe laser is shone through the cavity with a polarization that couples the ground state to two of the excited states that experience opposite frequency shifts in the presence of a magnetic field. If the cavity is placed on resonance with the zero-field frequency of the transitions (which are the same), very little light will get through the cavity, as the atoms efficiently absorb and scatter the light into other directions. If a magnetic field is applied, it opens up a window between the two transitions, allowing light to circulate in the cavity without being absorbed, and creating a transparent window between the two atomic

absorption features.

If this experiment were performed in free-space (no cavity), the width of this window would increase without bound with the magnetic field, making it a poor spectroscopic feature. If a large enough field were applied such that the probe were transmitted through the ensemble, the frequency width of the transmission window would be too wide to serve as a useful frequency reference. The key role of the cavity is to impose a strict requirement on the round-trip phase shift of the circulating light. For light with a frequency exactly between the two atomic transitions, the phase shifts contributed by the two transitions perfectly cancel, allowing the cavity resonance condition to be fulfilled and light to be transmitted. For probe frequencies not exactly centered between the two atomic transitions, the phase shift from the closer transition exceeds that of the farther transition and the two no longer cancel, breaking the cavity resonance condition. For appropriate system parameters, this leads the cavity resonance condition to be dominated by the atoms, rather than the cavity length. This can make the transmission feature narrower with atoms than without, and can make the frequency of the transmission feature largely insensitive to changes in the cavity length.

As in the previous chapter, we create a strongly coupled atom-cavity system by loading up to $N = 1.3 \times 10^6$ ^{88}Sr atoms (we got much better at this since the last chapter) into our usual 1D optical lattice supported by the high-finesse optical cavity. The peak trap depth is $100(10) \mu\text{K}$ and the atoms are laser-cooled to $10(1) \mu\text{K}$ (see Fig. 5.1a). We tune a TEM00 resonance of the cavity at frequency ω_c to be near resonance with the dipole-forbidden $\gamma = 2\pi \times 7.5 \text{ kHz}$ linewidth $^1\text{S}_0$ to $^3\text{P}_1$ transition at frequency ω_0 or wavelength 689 nm (see Fig. 5.1b). The cavity decay rate at this wavelength is $\kappa = 2\pi \times 150.3(4) \text{ kHz}$. In ^{88}Sr , the absence of nuclear spin means that the $^1\text{S}_0$ ground state is unique, while the $^3\text{P}_1$ excited state has three Zeeman sublevels (see Fig. 5.1b.)

In the limit of zero applied magnetic field, our system responds to an applied probe as though each atom were a simple two-level system. The ensemble can collectively absorb light from and then collectively reemit light into the cavity mode at the collective vacuum Rabi frequency $\Omega = \sqrt{N}2g$. Here $2g/2\pi = 15 \text{ kHz}$ is the rms value of the single-atom vacuum Rabi frequency, which accounts for

averaging over the standing-wave cavity mode. This exchange creates two new transmission modes (the usual vacuum Rabi splitting) that are shifted away from the empty cavity's transmission peak by $\pm\Omega/2$, as shown by the central red trace in Fig. 5.2a.

Because two orthogonal components of the probe light can couple to two distinct Zeeman sublevels, the coupled atom-cavity system has three normal modes of excitation, not two (see Fig. 5.1b and c.) In addition to the two modes at $\pm\Omega/2$ that lead to the transmission peaks, which we will refer to as the ‘bright’ modes, there is a third ‘dark’ mode whose frequency is equal to that of the atomic transition. The dark mode is composed of an equal superposition of the two atomic excited states and a photonic component that vanishes as the magnetic field approaches zero.

By applying a magnetic field B , we can mix photonic character into the dark mode, inducing transmission (also referred to as transparency) for probe light near ω_0 . The probe light is horizontally polarized and is perpendicular to the vertically oriented magnetic field. Each trace in Fig. 5.2a corresponds to a different applied magnetic field with strength parameterized by the induced Zeeman frequency splitting $\Delta/2\pi = (2.1 \text{ MHz/G})B$.

To describe the system, we extend the linearized input-output equations of [88] to include an additional atomic transition written in a rotating frame at the average atomic transition frequency ω_0 as:

$$\dot{a} = -\frac{1}{2}(\gamma + i\Delta)a - i\frac{1}{2\sqrt{2}}\Omega c \quad (5.1)$$

$$\dot{b} = -\frac{1}{2}(\gamma - i\Delta)b - i\frac{1}{2\sqrt{2}}\Omega c \quad (5.2)$$

$$\dot{c} = -\frac{1}{2}(\kappa + i2\delta_c)c - i\frac{1}{2\sqrt{2}}\Omega(a + b) + \sqrt{\kappa_1}c_i e^{i\delta_p t}. \quad (5.3)$$

Here, $\delta_c = \omega_c - \omega_0$ is the detuning of the cavity resonance frequency ω_c from atomic resonance, Ω is the observed collective vacuum Rabi splitting when $\Delta = 0$, γ is the decay rate of the excited atomic states, κ is the cavity power decay rate, and κ_1 is the coupling of the input cavity mirror that is driven by an externally incident probe field with complex amplitude c_i and at a probe frequency ω_p and detuning from atomic resonance $\delta_p = \omega_p - \omega_0$. The complex variables $a = \langle \hat{a} \rangle$, $b = \langle \hat{b} \rangle$,




Figure 5.1: (a) Simplified experimental diagram. The system is probed with horizontally polarized probe light. The light can be coherently absorbed by the atoms (brown ovals) and reemitted into the cavity at collective vacuum Rabi frequency Ω . The transmitted power is detected on a single photon counting module (SPCM). A magnetic field is applied along the vertical direction. (b) The atomic energy level diagram of the ground $1S_0$ and excited states $3P_1 |m_j\rangle$ with quantization axis \hat{q} parallel to the applied magnetic field \vec{B} . This is the viewpoint adopted in this letter. Here, the applied magnetic field creates a Zeeman splitting Δ between the excited states $|\pm 1\rangle$. Both of these transitions interact equally with the horizontally polarized light inside the cavity with collectively enhanced Rabi frequency $\Omega/\sqrt{2}$. The transmitted probe light is measured as the probe's detuning δ_p is swept. The $|0\rangle$ state is shown, but it does not interact with the horizontally polarized cavity-field.

$c = \langle \hat{c} \rangle$ are expectation values of bosonic lowering operators describing the cavity c , and collective excitations of the two atomic transitions a and b . The required Holstein-Primakoff approximation assumes weak excitation such that the number of atoms in the excited states $M_a = |a|^2$, $M_b = |b|^2 \ll N$ is a small fraction of the total atom number N . The average number of photons in the cavity is given by $M_c = |c|^2$, and the complex field transmitted through the cavity is $c_t = \sqrt{\kappa_2} c$ with κ_2 the coupling of the output mirror. The transmitted probe power relative to incident probe power is $P_T = |c_t/c_i|^2$ and the relative phase is $\psi = \arg(c_t/c_i)$.

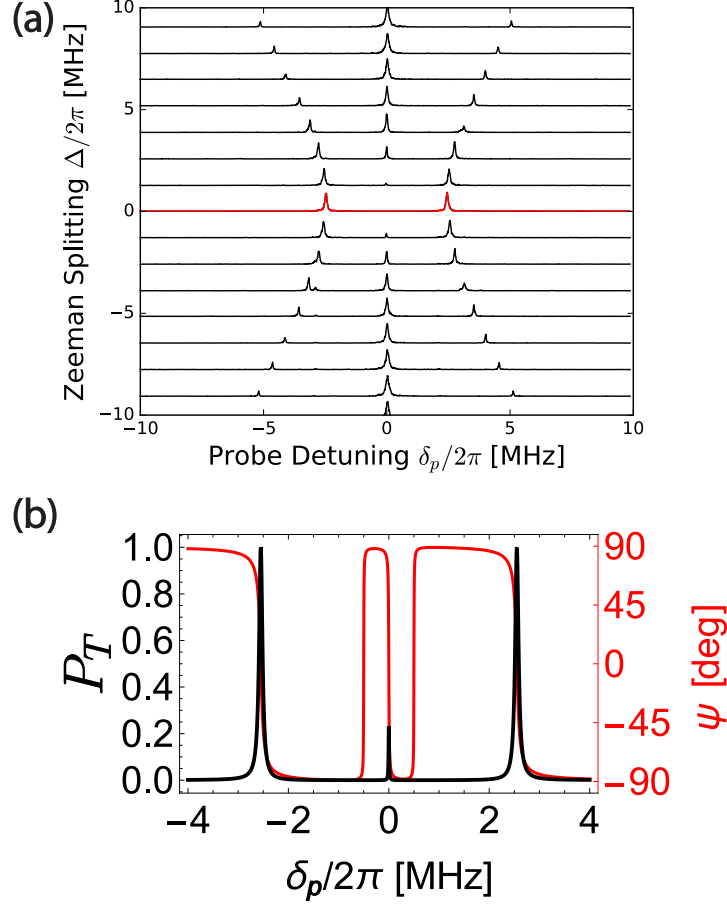


Figure 5.2: (a) The transmitted power through the cavity versus the probe detuning δ_p , with $\delta_c = 0$. Each trace was taken for different applied magnetic fields, creating different Zeeman splittings Δ labeled on the vertical. The central red trace is taken for $\Delta = 0$ and displays a collective vacuum Rabi splitting $\Omega/2\pi = 5$ MHz. When a magnetic field is applied perpendicular to the probe polarization, inducing a Zeeman splitting Δ , a new transmission feature appears in between the two original resonances of the vacuum Rabi splitting. (b) Linearized theory showing the power P_T and phase ψ of the transmitted light, plotted here for $\Omega/2\pi = 5$ MHz and $\Delta/2\pi = 1$ MHz.

Figure 5.2b shows the calculated steady-state transmitted power and phase for a single Zeeman splitting. The phase response changes rapidly near zero probe detuning, which results in a narrow MIT resonance compared to the broad vacuum Rabi splitting or bright modes for which the phase changes more slowly.

In order to describe the linewidth of the dark state resonance, we introduce a mixing angle θ defined by $\sin^2 \theta = \bar{\Delta}^2 / (\Omega^2 + \bar{\Delta}^2)$. Here, the effective detuning is $\bar{\Delta}^2 = \Delta^2 + \gamma^2$. The character of

the dark state excitation is given by the ratio of the probability that the excitation is photonic-like $P_c = M_c/(M_c + M_a + M_b) = \sin^2 \theta$ versus atomic-like $P_{ab} = (M_a + M_b)/(M_c + M_a + M_b) = \cos^2 \theta$. The dark state excitation can decay into free space at rate R_{ab} or by emission through the cavity mirrors R_c , with the ratio of the rates given simply by $R_{ab}/R_c = \gamma/(\kappa \tan^2 \theta) = NC(\gamma/\bar{\Delta})^2$, where the single particle cooperativity parameter is $C = 4g^2/\kappa\gamma$.

The linewidth of the dark state resonance can be written as:

$$\kappa' = (\gamma \cos^2 \theta + \kappa \sin^2 \theta)/b. \quad (5.4)$$

The term in parentheses is a weighted average of atom and cavity linewidths that reflects the character of the mode. The correction factor is $b = d \cos^2 \theta + \sin^2 \theta$, where $d = (\Delta^2 - \gamma^2)/\bar{\Delta}^2$. When $\Delta \gg \gamma$, both b and d approach unity. At small detunings $\Delta \sim \gamma$, the responses of the dark and bright modes to the applied drive become comparable, causing a modification of the correction factor. In the regime experimentally explored here ($b \approx 1$), κ' is simply the full width at half maximum linewidth of the power transmission feature. To define a linewidth valid in general, we define the linewidth via $\kappa' = 2(d\psi/d\delta_p)^{-1}|_{\delta_c=\delta_p=0}$. For $\Omega \gg \Delta \gg \gamma$, the mixing angle is small and the linewidth approaches the atomic linewidth $\kappa' \approx \gamma$, which can be much narrower than the cavity linewidth κ .

5.2 Experimental Characterizations of MIT

We measure the linewidth for the central dark resonance by linearly sweeping the probe laser's frequency over the cavity resonance and recording a time-trace of power transmitted on a single photon counting module. A Lorentzian is fit to the central feature to extract the full width at half maximum. This measurement is taken for a range of different Zeeman splittings and vacuum Rabi splittings by varying the applied dc magnetic field and atom number respectively. Figure 5.3a shows collected data plotted against the theoretical prediction at several Rabi frequencies. For very small Δ the feature becomes increasingly narrow, approaching the atomic transition linewidth

$\gamma = 2\pi \times 7.5$ kHz. For very large Δ the feature linewidth approaches the cavity decay rate κ .




Figure 5.3: (a) The measured linewidth of the central MIT transmission feature versus the induced Zeeman splitting between excited states. The traces are taken for three different collective vacuum Rabi frequencies $\Omega/2\pi = 4.6(5)$ (red), $10(1)$ (blue), and $16(1)$ (green) MHz, with values set by changing the total atom number N . The upper dashed line is the empty cavity's linewidth κ , and the lower dashed line is the atomic transition's linewidth γ . The minimum observed linewidth was 11 kHz. The shaded regions are no-free parameter predictions from the linearized model introduced in the text, indicating the ± 1 standard deviation uncertainty bands based on independent measurements of Ω . (b) The measured peak transmitted power of the central MIT transmission feature for the same collective Rabi frequencies. Here, the transmitted power is normalized to the peak transmitted power when the cavity is empty. Again the shaded regions indicate the ± 1 standard deviation uncertainty bands for the predictions.

Figure 5.3b shows the peak transmitted power at the MIT feature's resonance for the same data shown in Fig. 5.3a. The linearized theory predicts that the peak transmitted power is given by:

$$P_{max} = \frac{4\kappa_1\kappa_2}{\kappa^2} \frac{1}{\left(1 + \frac{\gamma}{\kappa \tan^2 \theta}\right)^2} \quad (5.5)$$

Note that the term in the denominator above is just the ratio of excitation decay rates R_{ab}/R_c . For large detunings $\Delta \gg \Omega, \gamma$, the peak transmission goes to that of an empty cavity $P_{max} \rightarrow P_{empty} =$

$$4\kappa_1\kappa_2/\kappa^2.$$

In the regime experimentally explored here, a change in the cavity resonance frequency ω_c by $\Delta\omega_c$ leads to a change in the dark state resonance frequency ω_D by a much smaller amount $\Delta\omega_D$. The pulling coefficient $P = \Delta\omega_D/\Delta\omega_c$ expresses this ratio. A small pulling coefficient $P \ll 1$ is desirable for a frequency reference as it will be less sensitive to thermal fluctuations and technical noise on the reference cavity. We can extract a pulling coefficient applicable to all parameter regimes from the linearized theory by considering how much the probe and cavity detunings would have to change to create equal changes in the quadrature amplitude of the transmitted field (such as one might measure using homodyne detection). This general pulling coefficient can be expressed as:

$$P = \frac{\sin^2 \theta}{b}. \quad (5.6)$$

In the typical regime of operation ($b \approx 1$), this is simply the cavity-like fraction of the dark excitation.

In Fig. 5.4, we show the measured pulling coefficient versus splitting Δ for several values of Ω , along with the predicted pulling coefficients from the linearized theory. The pulling coefficients were measured by sweeping the probe laser frequency across the dark resonance and fitting the center frequency ω_D with a Lorentzian fit model. This is then repeated while toggling ω_c between two values separated by 100 kHz, and the pulling coefficient is determined from the change in ω_D versus ω_c . Our lowest measured pulling coefficients are below $P = 0.05$.

In principle, Δ can be reduced further to reach a smaller pulling coefficient, at the expense of transmitted power. The theoretical pulling coefficient reaches a minimum value of $P = 1/(1 + N/(8M_c))$ at $\Delta = \sqrt{3}\gamma$, where $M_c = (\gamma/(2g))^2$ is the so-called critical photon number. The critical photon number is proportional to the cavity mode volume and atomic linewidth, but does not depend on the mirror reflectivity. As a result, small pulling coefficients are reached by working with small cavity volumes and very narrow linewidth transitions. For spectroscopic applications,

one would want to balance the desire for a low pulling coefficient against the need to collect transmitted photons without inducing heating in the atomic ensemble due to free-space scattering. The optimal parameter regime will depend on the specific requirements of the system.

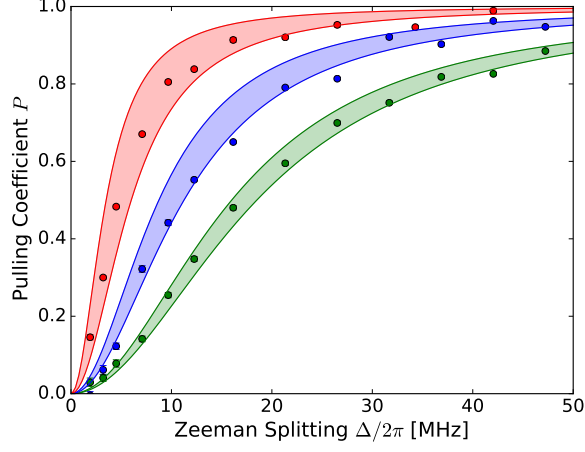


Figure 5.4: The pulling coefficient P versus Zeeman splitting Δ for several collective vacuum Rabi frequencies $\Omega/2\pi = 5(1)$ (red), $10(1)$ (blue), and $17(1)$ (green) MHz. The prediction from the linearized theory is shown with ± 1 standard deviation bands.

While the majority of this work has been done with the atoms trapped in the Lamb-Dicke regime (i.e. confined to much less than the wavelength of the probe light) with respect to the cavity axis, we have also performed scans of the cavity transmission spectrum in which the atoms were unconfined along the cavity axis. In this configuration, the rms Doppler shift along the cavity axis is roughly 45 kHz. Despite this inhomogeneous broadening, we observe a center feature linewidth of 18.5 kHz, which we believe is limited by technical noise on the cavity frequency that arises when we turn down the lattice depth to release the atoms. We expect the linewidth of the dark feature to be insensitive to inhomogeneous broadening so long as Δ is much larger than the inhomogeneous broadening [136]. This insensitivity to Doppler broadening may make such techniques suitable to continuously operating atomic beam experiments, where confining the atoms to the Lamb-Dicke regime would be challenging.

5.3 Mapping to EIT

In this section, we make explicit the connection between magnetically induced transparency (MIT) and a more traditional Raman electromagnetically induced transparency (EIT) system.

Previously, we adopted the perspective where the quantization axis \hat{q} is taken to be along the direction of the applied magnetic field \vec{B} . We will call this the ‘MIT picture’. In this picture, the horizontally polarized probe light can be decomposed into circularly polarized σ^\pm components that collectively couple to the $|m_j = \pm 1\rangle$ states respectively with Rabi frequency $\Omega/\sqrt{2}$. Since the quantization axis is along the magnetic field, these $|\pm 1\rangle$ states have a Zeeman splitting given by Δ . The $|0\rangle$ state does not interact with the horizontally polarized light. When no magnetic field is applied, the dark state mode is entirely atom-like, and no probe light is transmitted through the cavity. However, when the magnetic field is applied, it mixes a photon-like component into the dark state, resulting in transmitted light.

In this picture, we can think of the transmission as resulting from two simultaneous effects. First, the absorption profiles of the $|m_j = \pm 1\rangle$ states are pushed off resonance with the dark state. At the same time, the dark state is nearly pinned to the center of the two transition frequencies because of the opposite dispersive shifts from the two states $|m_j = \pm 1\rangle$.

However, we may also view the system from an alternative perspective where the quantization axis is perpendicular to both the cavity axis and the axis of the magnetic field (Fig. 5.5). We will call this the ‘EIT picture’. For clarity, the excited Zeeman states written in this new basis are denoted by $|m_{j\perp}\rangle$. The probe light is π polarized and collectively couples the ground state to the single excited state $|0_\perp\rangle$ with Rabi frequency Ω .

The three states $|m_{j\perp}\rangle$ are degenerate in this EIT picture, but are coupled to their neighboring states by the transverse magnetic field with a Rabi frequency $\Delta/\sqrt{2}$. Because of the symmetry of the couplings, the magnetic field couples $|0\rangle$ to the symmetric state $|S\rangle = (|+1_\perp\rangle + |-1_\perp\rangle)/\sqrt{2}$ with Rabi frequency Δ . This symmetric state $|S\rangle$ is not coupled to the cavity since this state physically corresponds to an optical dipole aligned along the cavity axis. The difference state

$|D\rangle = (|+1_\perp\rangle - |-1_\perp\rangle)/\sqrt{2}$ has no coupling to other states, but for completeness we note that the difference state would correspond to a vertical dipole moment that would couple to the vertically polarized cavity mode.

The ground state, $|0_\perp\rangle$, and $|S\rangle$ form a traditional three-level EIT system, where the control optical field has been replaced by a control dc magnetic field. In contrast to a traditional EIT system however, the state $|S\rangle$ has the same linewidth as $|0_\perp\rangle$. While there is a coherent cancellation of excitation to $|0_\perp\rangle$ as in EIT, the state $|S\rangle$ does become populated in our system, and may radiate into free space. This prevents the observation of spectral features narrower than $|0_\perp\rangle$, as are observed in EIT, and motivates the use of narrow optical transitions.




Figure 5.5: In the ‘EIT picture’, the quantization axis is perpendicular the magnetic field and parallel to the probe light’s polarization. (left) The cavity now couples the ground and excited Zeeman state $|m_{j\perp} = 0\rangle$. The magnetic field couples $|0_\perp\rangle$ to neighboring excited Zeeman states $|m_{j\perp} = \pm 1\rangle$. (center) Equivalently, the magnetic field couples to a symmetric combination of Zeeman states $|S\rangle$ while the antisymmetric combination $|D\rangle$ does not participate. (right) For comparison, a traditional Raman EIT system is shown for comparison with the equivalent coupling strengths labeled.

5.4 Limits on Photon Retrieval Efficiency

EIT systems have been used to store light. The hypothetical photon capture and retrieval efficiency is P_R , defined such that $P_R = 1$ corresponds to perfect photon storage or retrieval efficiency. Compared to more traditional EIT schemes in which the two ground states are long

lived, here the photon storage state $|S\rangle$ decays at the same rate γ as the intermediate state $|0_\perp\rangle$.

To estimate the theoretical maximum efficiency P_R , we adopt the EIT quantization axis. We assume that a photon is perfectly stored as a collective excitation in state $|S\rangle$ at zero magnetic field (i.e. $\Delta = 0$.) The goal is to convert the excitation into a photon in the cavity as quickly as possible, returning all atoms to the ground state.

To accomplish this conversion, at $t = 0$ one switches on a large magnetic field $\Delta \gg \Omega$ for a duration that creates a π -pulse. This swaps the collective excitation from the state $|S\rangle$ into the state $|0_\perp\rangle$. Physically, this corresponds to using a transverse magnetic field to rotate the axis of the collective optical dipole by 90° from being oriented along the cavity (i.e. cannot radiate into the cavity) to being oriented along the quantization axis (i.e. can radiate into the cavity.) The magnetic field is then returned to zero. We will assume that the magnetic field can be made arbitrarily large and can be switched on and off arbitrarily fast such that atomic decay at rate γ during this step can be neglected.

Having swapped the excitation into $|0_\perp\rangle$, the excitation will now automatically Rabi flop into a cavity photon and back at frequency Ω . We assume here that $\Omega \gg \kappa$, as is the case in our system. The time to accomplish this second π pulse, converting the excitation into a cavity photon, is just $t_\pi = \pi/\Omega$. At t_π one immediately switches the magnetic field back on to a large value $\Delta \gg \Omega$ coupling $|0_\perp\rangle$ to $|S\rangle$ such that the photon in the cavity is no longer re-absorbed by the atoms. The photon will then exponentially leak out of the cavity at power decay rate κ .

During this second π -pulse, the excitation spends on average 50% of the time as a collective atomic excitation that decays at rate γ . Accounting for the atomic decay, one finds a retrieval efficiency $P_R = 1 - \frac{1}{2}t_\pi\gamma = 1 - (\pi/2)\gamma/\Omega$, the result given in the main text. Numerical simulation of the photon retrieval process indicates that the above expression is accurate to 20% even when $\Omega = \kappa$. From time reversal symmetry, photon storage will have the same efficiency. For the current system parameters, the theoretical limit is $P_R \approx 99.8\%$. It is likely that at this level, other experimental imperfections (finite magnetic field switching speeds, finite magnitude maximum magnetic fields, and other) would limit storage and retrieval efficiencies to less than this fundamental limit.

We note that here the photon loss scales as $1 - P_R \propto \sqrt{m_c/N}$ where m_c is the critical photon number of cavity QED [134]. This is to be compared to more traditional EIT schemes in which one assumes $\kappa \gg \Omega$ and $\kappa\gamma/\Omega^2 \ll 1$, in which case $1 - P_R \propto n_c/N$ [140]. Here, n_c is the critical atom number $n_c = 1/C$ of cavity QED. The critical atom number does not depend on the strength of the atomic transition, while here, the critical photon number is improved (i.e. made smaller) by using narrower linewidth optical transitions since $m_c \propto \gamma$. Also note that the critical photon number does not depend on the cavity finesse, so long as the finesse is high enough such that $\Omega \gg \kappa$.

5.5 Conclusion

To summarize, we have demonstrated a technique to realize a narrow spectroscopic feature based on collective interaction between an ensemble of atoms and a high finesse optical cavity. The center frequency of the feature can be made highly insensitive to changes in cavity resonance frequency. In analogy to EIT, this technique may also be applicable to tasks relevant for information processing.

The observed linewidth of the feature approaches the natural linewidth of the 7.5 kHz optical transition and can be insensitive to inhomogeneous broadening of the atomic transition frequency. In the future, it might be possible to extend this technique to even narrower optical transitions for enhanced spectroscopic sensitivity in atoms such as calcium and magnesium.

Chapter 6

Lasing on the 7.5 kHz Linewidth Transition: the Superradiant Crossover Regime

Our first foray into superradiance on narrow-linewidth optical transitions was a study of superradiance on the 7.5 kHz linewidth $^3\text{P}_1$ to $^1\text{S}_0$ transition in ^{88}Sr . The hierarchy of decay rates $\kappa/\gamma \simeq 20 \gg 1$ places this laser in the bad-cavity regime with respect to decay rates — excitations in the system are much more likely to be lost as photons transmitted through the cavity mirrors than as photons emitted into free space by the atoms. However, for typical atom numbers used in this experiment, the collective vacuum Rabi frequency $\Omega = 2g\sqrt{N}$ was several MHz, much greater than either κ or γ . This can place the system in a regime where coherent dynamics dominate over dissipation — the resolved vacuum Rabi splitting regime.

Pulsed lasing in this regime ($\Omega \gg \kappa, \gamma$) exhibits interesting ringing behavior associated with these coherent dynamics as the atoms exchange photons with the cavity mode several times before they are lost to transmission through the mirrors.

We can also apply repumping lasers to achieve quasi steady state operation. This modifies the hierarchy of rates, and thus the qualitative dynamics of the system. This occurs because the repumping process (which involves spontaneous emission) both causes damping of the atomic coherence and takes the atomic system out of a fully symmetric state (reduces the length of the Bloch vector J), which reduces the rate of coherent interactions. A more general expression for the rate of coherent interactions between the atoms and cavities is $\Omega' = 2g\sqrt{2J}$, which can be much less than Ω . By introducing variable rates of repumping, we can continuously tune between the resolved

vacuum Rabi splitting regime $\Omega' \gg \kappa, \gamma$ to the more traditional bad-cavity regime $\kappa \gg \Omega', \gamma$.

For context, the bad-cavity regime of laser physics has been explored in the optical domain in gas lasers [141], with homogeneous and inhomogeneous transition linewidths of hundreds of MHz, and in a 4-level system in Cesium with a Doppler-broadened gain bandwidth of 9 MHz [142]. In the microwave domain, masers operate deep in the bad-cavity regime [143]. Rare-earth doped solid state lasers utilize transitions with long-lived excited states, but with inhomogeneous broadening of hundreds of GHz [144]. Raman dressing has been used to create long-lived virtual states with low inhomogeneous broadening, which have been used to demonstrate key properties of optical lasing very deep into the bad-cavity regime [145, 63, 65, 66], and deep into the good-cavity regime [146]. This work is unique in that we explore lasing on a true optical transition in a regime where both the homogeneous and inhomogeneous linewidths of the gain medium can be made small compared to the cavity linewidth, and in which the rate of coherent interactions between atoms and cavity can be tuned to be either large or small compared to the cavity decay rate.




Figure 6.1: Energy level and experimental diagrams. The lasing transition is the 7.5 kHz wide 3P_1 , $m_J = 0$ to 1S_0 transition in ^{88}Sr . A magnetic field B is oriented perpendicular to the cavity axis, and pump lasers at 688 and 689 nm are applied with directions and polarizations as shown. Collectively enhanced emission into the TEM00 mode of the cavity is collected in an optical fiber and sent to one of various detectors.

For this work, we used up to $N \sim 100\text{k}$ ^{88}Sr atoms. The atoms were cooled to $9\text{ }\mu\text{K}$ and confined as usual in our 813 nm optical lattice (Fig. 6.1). The lasing occurs on the $|e\rangle \equiv |^3\text{P}_1, m_J = 0\rangle$ to $|g\rangle \equiv |^1\text{S}_0\rangle$ transition at 689 nm , which has a natural decay linewidth $\gamma = 2\pi \times 7.5\text{ kHz}$. A single atom trapped at an antinode of the lasing mode exchanges excitations with the TEM00 mode at a frequency $2g_0 = 2\pi \times 21.2\text{ kHz}$. Accounting for inhomogeneous coupling between the atoms and the cavity mode, this frequency is collectively enhanced to as much as $\Omega = g_0\sqrt{2N} \sim 2\pi \times (1\text{ to }5)\text{ MHz}$.

6.1 Pulsed Lasing in the Vacuum Rabi Splitting Limit




Figure 6.2: (a) Laser power emitted from the cavity during a collectively enhanced pulse, with $N = 36\text{k}$ atoms. The atoms are optically pumped into $|e\rangle$ at $t = 0$. (b) Numerical simulation of (a) with no free parameters. (c) The power emitted from the cavity during continuously repumped lasing with $N = 60\text{k}$ atoms. Each atom emits around 35 photons into the cavity before repump-induced heating causes laser operation to cease. For scale, 1 nW corresponds to 3500 intracavity photons. (d) Repumping scheme. Atoms are incoherently repumped through $^3\text{P}_1$, $m_J = -1$ to $^3\text{S}_1$, where they decay into the ^3P manifold. Atoms that fall into $^3\text{P}_0$ or $^3\text{P}_2$ are then repumped by additional lasers via $^3\text{S}_1$. For display purposes, $^3\text{P}_{0,2}$ and $^1\text{P}_1$ states are shown at arbitrary vertical positions.

First, we characterized pulsed lasing in the resolved vacuum Rabi splitting regime. When many atoms are placed in state $|e\rangle$, the collectively enhanced coupling to the cavity causes the atoms to quickly decay to $|g\rangle$ by emitting a pulse of light into the cavity. We detect this light on an avalanche photodiode, with a single trace shown in Fig. 6.2a. In this regime, the collectively enhanced emission rate exceeds the cavity linewidth, which in turn is much greater than the atomic decay rate; i.e., $\Omega' \gg \kappa \gg \gamma$. The light pulse is partially re-absorbed by the atoms, and re-emitted into the cavity several times before escaping the system via transmission through a cavity mirror, resulting in oscillations in output power from the cavity, as shown in Fig. 6.2a. The timescale of these oscillations is of order $2\pi/\Omega$, but because of inhomogeneous coupling to the cavity, the atoms do not remain in a fully symmetric state. This modifies the period of oscillations and results in incomplete reabsorption of the initially emitted pulse.

By measuring a vacuum Rabi splitting following the pulse, we infer an atom number of $N = 36,000$ for the trial shown in Fig. 6.2a, up to fluctuations between trials of around 20 percent. We simulate the pulse dynamics for this atom number by integrating a set of optical Bloch equations (see Chapter 2.4 and ref. [64]) that also account for the inhomogeneous coupling to the cavity mode. The results are shown in Fig. 6.2b, with good qualitative agreement to observed pulses.

We can get some intuition into the ringing behavior from the pendulum model presented in Chapter 2.4. Neglecting the fact that the atoms are inhomogeneously coupled to the cavity, the system is equivalent to an under-damped pendulum that starts out fully inverted. The pendulum swings downward, turning potential energy into kinetic energy (signifying atomic inversion being turned into excitation of the cavity field). Because the system is under-damped, the pendulum will oscillate about its resting position several times before settling.

6.2 Quasi Steady State Operation

We next apply continuous repumping from $|g\rangle$ back to $|e\rangle$ to operate the laser in a steady-state manner. The repumping process both maintains population inversion and causes decay of the transverse coherence of the atomic ensemble, which reduces the collectively enhanced emission

rate. When a steady state operating condition exists, the total rate of emission from $^3\text{P}_1$ matches the single-particle rate w at which atoms are optically pumped out of $|g\rangle$. The repumping process also homogeneously broadens the ground state to a width w . When operating with $w \ll \kappa$, we access the bad cavity or superradiant regime, where the cavity decay rate exceeds the collectively enhanced emission rate, other damping rates, and the bandwidth of the gain medium.

Fig. 6.2c, shows the laser output power for a representative run of the experiment, with repumping applied to an ensemble of $N \simeq 60\text{k}$ atoms initially in $|e\rangle$. We can sustain lasing for up to 1.5 ms, with higher power operation possible for durations of around $500 \mu\text{s}$, as shown. In such a trial, each atom emits up to 35 photons into the cavity before lasing ceases. This seems consistent with heating we would expect from photon recoils during the repump process.

In addition to operating in the bad-cavity regime defined by $\kappa \gg \gamma$, our laser operates in an unusual regime with respect to inhomogeneous broadening. In contrast to gas and solid state lasers that operate on similarly long-lived transitions, our laser has very low inhomogeneous broadening. We should compare the spread in frequencies associated with this broadening to the collectively enhanced emission rate, which in the steady state case is set by w [110, 112]. When inhomogeneous broadening is comparable to or less than the collectively enhanced emission rate, the majority of atoms in the ensemble can phase-synchronize and contribute to collectively enhanced emission.

Because we operate in the Lamb-Dicke regime, Doppler shifts are negligible in our system. However, the lattice contributes a polarization-dependent shift of the lasing transition [135]. The degree to which this shift is inhomogeneous depends on the temperature of the atoms, which increases during laser operation. At the beginning of laser operation, we estimate inhomogeneous broadening to be less than 30 kHz FWHM for our typical lattice polarization. In principle, this broadening can be tuned to zero by changing the polarization of the lattice. For our range of operating parameters, the majority of atoms have transition frequencies within w of the average transition frequency, allowing them to simultaneously contribute to lasing.

The repumping mechanism is shown in Fig. 2d, with beam directions and polarizations shown in Fig. 6.1a. A magnetic field is applied perpendicular to the cavity axis, splitting out the $m_J = \pm 1$

states by ± 7 MHz. The 689 nm repump light is applied near resonance with the 1S_0 to 3P_1 , $m_J = -1$ transition. Additional π -polarized 688 nm repump light resonant with the dipole-allowed 3P_1 to 3S_1 transition is applied from the same direction. Because of dipole selection rules, this 688 nm pump couples the 3P_1 $m_J = \pm 1$ states to 3S_1 , while the $m_J = 0$ state is unaffected. Single-particle spontaneous emission takes atoms from the 3S_1 states to the 3P states, and additional repump lasers at 707 nm and 679 nm are applied to depopulate the metastable 3P_2 and 3P_0 states. Because atoms reach 3P_1 by single-particle spontaneous decay, we do not expect any coherence between the pump lasers and the emitted light.

It is likely that lasing terminates when heating from free-space scattering caused by the repump process reduces the product NC both by causing atom loss and by reducing the coupling of remaining atoms to the cavity. For each lasing photon emitted into the cavity, we estimate roughly 13 photon recoils are imparted during the repump process. For 35 lasing photons per atom, this would contribute heating comparable to the lattice depth.

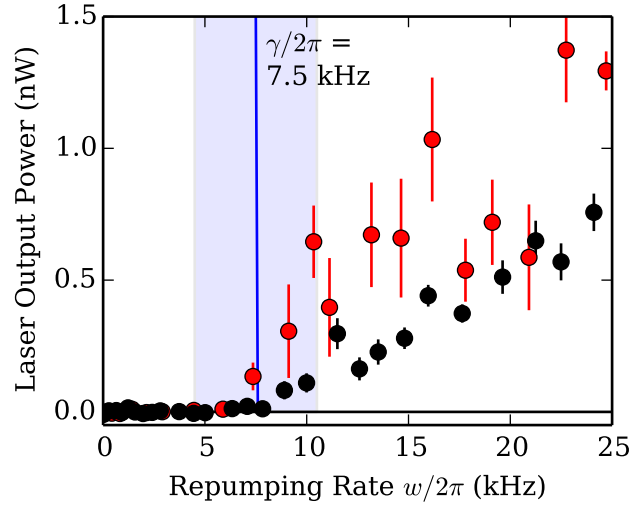


Figure 6.3: Threshold behavior of laser. Laser output power in quasi steady state is plotted versus repumping rate w for no added broadening $\gamma'_\perp = 0$, (red points) and added broadening, ($2\gamma'_\perp \simeq 2\pi \times 3$ MHz $\gg \gamma, \kappa$, black points). Atomic decay rate γ is displayed as vertical blue line, with shaded blue region representing uncertainty in calibration of w . For both conditions, the measured threshold repumping rate w_t is consistent with $w_t = \gamma$.




Figure 6.4: (a) Pulling coefficient P versus the total homogeneous broadening of the lasing transition, along with prediction (blue). The vertical line at 160 kHz represents the cavity decay rate κ . The inset displays the left five points on a linear scale. (b) Averaged heterodyne power spectral density (PSD) between output light and 689 nm pump laser, with recentering of each individual trial before averaging. Lorentzian (Gaussian) fits, shown as red (blue) lines indicate FWHM linewidths of 6.0(3) (4.7(3)) kHz.

6.3 Lasing Threshold Behavior

The output power in our laser exhibits a characteristic threshold behavior versus w , as shown in Fig. 6.3. In the regime that the total atomic transverse decoherence rate $\gamma_{\perp} = \gamma/2 + w/2 + \gamma'_{\perp}$ satisfies $\gamma_{\perp} \ll NC\gamma$, the threshold repumping rate for lasing is simply $w_t = \gamma$. This is expected to be true even in the presence of homogeneous or inhomogeneous broadening that exceeds the natural decay rate of the lasing transition [51]. In the above, γ'_{\perp} accounts for other mechanisms that lead to a decay of atomic coherence. Intuitively, this threshold requirement states that in order to establish population inversion, atoms must be removed from $|g\rangle$ at a rate faster than they can undergo single particle decay from $|e\rangle$ to $|g\rangle$.

We measure output power from the laser in a window between 20 and 60 μs after the beginning of laser operation. We tune w by changing the intensity of the 689 nm repump laser and measure $w_t = 2\pi \times 8(3)$ kHz, with the error dominated by the calibration of w . This is consistent with the predicted threshold $w_t = \gamma = 2\pi \times 7.5$ kHz.

By inducing Rayleigh scattering on the dipole-allowed $|g\rangle$ to 1P_1 transition at 461 nm using our MOT beams, we add additional homogeneous broadening $2\gamma'_\perp \simeq 2\pi \times 3 \text{ MHz} \gg \kappa, \gamma$. Despite this high scattering rate, we measure the same threshold value of $w_t = 2\pi \times 8(3) \text{ kHz}$, again consistent with the prediction. By rotating the lattice polarization, we can introduce inhomogeneous broadening that can be varied from less than 30 kHz to as much as 140 kHz FWHM. Between these two conditions, we measure w_t to differ by less than 3 kHz. This confirms that the dominant factor in determining threshold is indeed the bare atomic decay rate, not the effective transition linewidth.

6.4 Cavity Pulling, and Tuning Through the Superradiant Crossover Regime

An important technological appeal of bad-cavity lasers is their reduced sensitivity to fluctuations in cavity resonance frequency. We define the pulling coefficient $P = \Delta f_\ell / \Delta f_c$, where Δf_ℓ is the change in the lasing frequency created by a change Δf_c in the cavity resonance frequency. We expect the pulling coefficient to be given by $P = 2\gamma_\perp / (2\gamma_\perp + \kappa)$ when the laser is operating in steady-state [141, 63, 66]. A pulling coefficient $P \ll 1/2$ is a key signature that the laser is operating in the bad-cavity, or superradiant regime with coherence primarily stored in the atoms, not in the light field.

To measure the pulling coefficient, we overlap the laser light emitted from the cavity with a heterodyne beam from our 689 nm laser. We changed the cavity frequency between trials and computed power spectra from the heterodyne data. From Gaussian fits to these power spectra, we extract the peak frequency of the emitted light. We obtain the pulling coefficient from a linear fit to f_ℓ versus f_c , with f_c scanned by 1 MHz. We repeat this with different values of w . When we turn down the repump beams to just above threshold, $w = 2\pi \times 14 \text{ kHz}$, we measure a pulling coefficient of $P = 0.09(2)$. This indicates that we have reached the bad-cavity regime, where the spectral properties of the emitted light are dominated by the atomic gain medium. When we increase the repump rate such that the effective atomic linewidth becomes comparable to the cavity decay rate, we reach a crossover regime where coherence is shared between atoms and light field, and measure a pulling coefficient consistent with $P = 1/2$.

The right-most point in Fig. 6.4a was obtained by turning on the 461 nm MOT beams to further increase γ_{\perp} by Rayleigh scattering from $|g\rangle$ at a rate $2\gamma'_{\perp} \simeq 2\pi \times 3$ MHz. With this scattering rate, we measure a pulling coefficient of $P = 0.97(3)$. In this highly broadened regime, scattering from the 461 nm MOT beams removes coherence between 1S_0 and 3P_1 without directly affecting the population inversion. By removing the phase coherence from the atoms, we access the good-cavity regime, where the spectral properties of the emitted light are set by the optical cavity.

We bound the linewidth of the emitted light by computing an average power spectrum from heterodyne data taken over many trials (Fig. 6.4). From trial to trial, the center frequency of the power spectrum fluctuates due to low frequency noise on the heterodyne laser and optical cavity (at the time of this work, we did not have a very stable reference cavity, so our laser moved around quite a bit). We align each power spectrum by shifting its frequency axis by the center frequency obtained from a Gaussian fit, which reduces sensitivity to low-frequency noise. We then fit a Lorentzian and Gaussian to points from all recentered spectra simultaneously. From 150 μ s long subsets of the time data, we measure a Lorentzian (Gaussian) FWHM linewidth of 6.0(3) (4.7(3)) kHz. We expect that the fundamental quantum linewidth for these operating conditions would be around 1 Hz [141], far below our current ability to resolve. Our measured linewidth is Fourier-limited for shorter subsets, and limited by acoustic noise on the heterodyne laser for longer subsets. However, our measured linewidth is slightly narrower than the natural linewidth of the lasing transition (7.5 kHz), and far narrower than the linewidth imposed by repumping (~ 100 kHz), exhibiting the frequency narrowing characteristic of synchronization in a laser.

This work served as a first step in demonstrating that the results achieved in Raman systems could indeed be applied to a system with true narrow-linewidth optical transitions. It also demonstrated a basic mechanism for achieving quasi steady state operation in strontium. Beyond its current role as a scientific curiosity, a laser based on the 7.5 kHz linewidth strontium transition could perhaps be considered as a moderate performance but relatively simple frequency reference. If one wished to take this path, it would be advantageous to work with a broader linewidth optical cavity in order to further reduce the effects of cavity frequency noise. To achieve steady-state

operation, one would need to provide a constant flux of new atoms into the cavity mode in order to replace those lost to heating. Because of the relatively broad linewidth of the transition, it is conceivable that one could operate such a laser using as a gain medium a beam of pre-cooled strontium atoms that pass through the cavity mode without axial confinement. This would represent a substantial simplification relative to a laser operating on the milliHertz linewidth transitions.

Chapter 7

Pulsed Lasing on the milliHertz Linewidth Clock Transition

The core element of my thesis is the demonstration and characterization of superradiance from the milliHertz linewidth strontium clock transition. A superradiant laser based on this transition provides a promising path to a active optical frequency reference with performance far beyond the current state of the art [51].

The first major milestone towards this goal, described in this chapter and in [25], was to demonstrate the collective enhancement of radiation from the milliHertz linewidth transition by observing superradiant pulses. In this work, I was able to demonstrate a factor of up to 10,000 enhancement in the rate at which atoms radiate into the cavity, compared to the expected rate of emission for independently radiating atoms. The temporal form of the pulses, as well as key scalings are well-described by simple theory presented later in this section. Also in this section, I describe a method for seeding coherence in the atoms that has proved very useful in the study of our superradiant source. While a much more thorough characterization of the spectral properties of the superradiant pulses appears in the next chapter, I include here a preliminary characterization of the relative coherence between two simultaneously lasing milliHertz linewidth transitions.

7.1 Experimental System

For this work, we switched over from ^{88}Sr , for which the $^3\text{P}_0$ is effectively infinitely long lived for our purposes, to ^{87}Sr . In ^{87}Sr , the presence of nuclear spin causes $^3\text{P}_0$ to decay in about 160 seconds [85]. We used an ensemble of up to $N = 2.5 \times 10^5$ ^{87}Sr atoms cooled and tightly

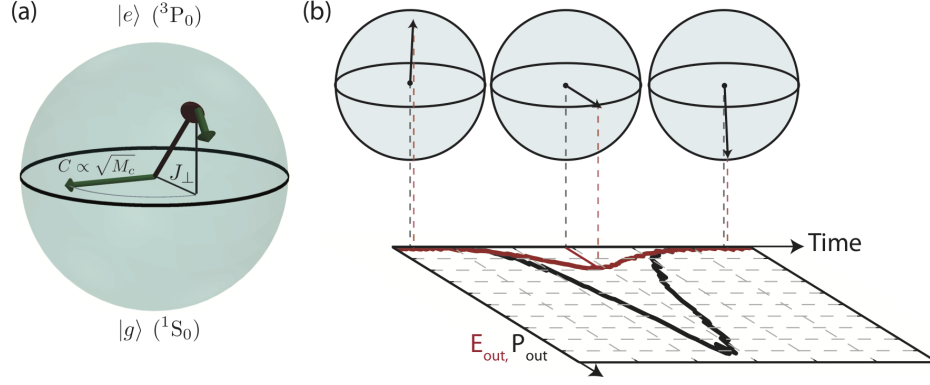


Figure 7.1: (a) The atoms can be represented by a collective Bloch vector on a Bloch sphere. The cavity field $C \propto \sqrt{M_c}$, where M_c is the number of photons in the cavity, is represented by a vector that lies in the equatorial plane perpendicular to the atomic coherence J_\perp (when the cavity is on resonance with the atomic transition). The Bloch vector rotates about the axis that represents the cavity field. (b) The Bloch vector behaves like a highly damped pendulum that starts inverted at the north pole of the Bloch sphere (all atoms in $|e\rangle$). Quantum fluctuations disturb the system from its unstable equilibrium position, causing the Bloch vector to swing down the Bloch sphere, emitting peak radiation at the equator and ultimately relaxing to the south pole (all atoms in $|g\rangle$) as inversion is lost. The radiated electric field (red trace) is proportional to the perpendicular projection of the Bloch vector, J_\perp , which at its peak is proportional to N . The radiated power (black trace) is proportional to the square of the radiated electric field, and at its peak is therefore proportional to N^2 . This is one way to understand the origin of the collective enhancement in emission rate. The black output power trace on the projection plane is actual data.

trapped within our usual magic wavelength lattice (near 813.4274 nm, which unlike in previous sections, is finally actually a magic wavelength). Our goal was to get the atoms to collectively emit light from the mHz linewidth transition.

To describe this process, I will apply our usual Bloch sphere description to this system: As a reminder, the state of the atomic ensemble can be represented by a collective Bloch vector, which is the vector sum of the Bloch vectors of the N individual atoms. An atom in the excited state, $|e\rangle (^3P_0)$, has a Bloch vector pointing up, while an atom in the ground state, $|g\rangle (^1S_0)$, has a Bloch vector pointing down. An atom in an equal superposition of $|e\rangle$ and $|g\rangle$ corresponds to a Bloch vector on the equator of its Bloch sphere, with phase ϕ determined by the phase of its superposition $|g\rangle + e^{i\phi}|e\rangle$. A collective Bloch vector on the equator of the Bloch sphere corresponds to each atom in a superposition of $|e\rangle$ and $|g\rangle$ with the appropriate phases to collectively radiate into the cavity

mode (See Fig. 11.1a).

As described in Chapter 2.4, the atoms radiate an electric field C into the cavity at a rate proportional to J_{\perp} , the magnitude of the projection of the Bloch vector onto the equatorial plane of the Bloch sphere. The collective enhancement of emission results from the fact that the power radiated is proportional to the square of the electric field. Because the electric field is proportional to the atom number N , the radiated power scales with N^2 [147].

For this work, we place the cavity on resonance with the atomic transition. In this configuration, the cavity field $C \propto M_c$, where M_c is the average number of photons in the cavity, is represented by a vector that in the equatorial plane of the Bloch sphere that is perpendicular to J_{\perp} . The electric field radiated into the cavity acts on the atoms by causing rotation of the Bloch vector about C . This is the mechanism of stimulated emission in the superradiant regime. For an ensemble whose Bloch vector lies in the northern hemisphere of the Bloch sphere, this leads to positive feedback for emission. The atoms will radiate into the cavity, and the radiated electric field will then cause the Bloch vector to tip further from the north pole, increasing J_{\perp} and thus causing the atoms to radiate more strongly. The result is a pulse of light that builds up gradually, reaches a peak in power as the Bloch vector passes the equator, and falls to zero as the atoms reach the ground state (Fig. 11.1b).

The dynamics of the Bloch vector in this regime (neglecting the effects of inhomogeneous coupling of the atoms to the cavity) are well-described by the motion of an initially inverted, highly damped pendulum. In contrast to the pulsed emission on the 7.5 kHz transition, the high level of cavity damping here relative to the rate of coherent interactions means that there is no ringing in the pulsed emission, which would be analogous to the pendulum swinging about its equilibrium position.

7.2 Observation of Superradiant Pulses

To observe superradiant pulses, we prepare atoms in $|e\rangle$, the nuclear $|F = 9/2, m_f = 9/2\rangle$ sub-level of $^3\text{P}_0$. We first optically pump the atoms to $|g\rangle$, the $|F = 9/2, m_f = 9/2\rangle$ sub-level of




Figure 7.2: Spontaneously generated superradiant pulses. Representative single time traces of photon output rate R for pulses at different atom number $N \approx 100 \times 10^3$ (green), 125×10^3 (blue), 150×10^3 (red), 200×10^3 (black). The equivalent average intracavity photon number is calculated on the right as $M_c = R/\kappa$.

1S_0 , then adiabatically transfer up to 75% of the atoms to $|e\rangle$ using a frequency-swept 698 nm transfer beam applied through the cavity.

This transfer beam is derived from our 698 nm diode laser, which is frequency-stabilized to our transfer cavity. At the time of this experiment we were using a home-built transfer cavity, which was not stable enough to allow for high-fidelity optical rotations without sweeping. Further, we are constrained to drive the optical rotations along the cavity axis, as this is the only direction in which Doppler shifts are sufficiently suppressed to efficiently drive the narrow-linewidth transition. This means that we are forced to drive the transition with a standing wave with spatially varying Rabi frequency. This would further limit our ability to do a high-fidelity optical rotation (although this actually works quite well for other modes of operation, as will be described in the following section). Our solution was to rely on Landau-Zener transitions (adiabatic transfer) from the ground to excited state. We accomplish this by coupling a sideband from the 698 nm laser to a TEM00 mode of the cavity. This mode is placed on resonance with the atomic clock transition. The laser is then swept in frequency over the atomic and cavity resonances. For this experiment, we typically swept the laser by 200 kHz in 20 ms. The Rabi frequency experienced by an atom at an antinode

of the standing wave was typically around 5 kHz. Because the transfer beam is linearly polarized along the magnetic field, the atoms maintain their spin polarization m_f when transferred to 3P_0 .

This adiabatic transfer process was sufficient to transfer most of the atoms to the desired state with minimal heating. Further, those that were left behind were presumably those that were located near nodes of the 698 nm cavity mode, and so would contribute little to the superradiant pulse anyway. To prepare the atoms with full inversion (no atoms in $|g\rangle$) and to ensure that the laser pulses are initiated by quantum noise rather than residual atomic coherence associated with the adiabatic transfer process, we next briefly apply lasers to the dipole-allowed 1S_0 to 1P_1 transition to heat any atoms remaining in the ground state out of the lattice.

When all atoms are initially prepared in $|e\rangle$, we observe collectively enhanced decay on the clock transition. To detect the superradiant laser pulses, the output of the cavity is coupled to a single-mode fiber and detected on a single-photon counting module (SPCM) whose TTL output is low-pass filtered to provide a signal proportional to the photon emission rate.

Both quantitative and qualitative features of the collectively enhanced emission are dramatically different from that of independently radiating atoms. Not only does the collective enhancement lead to an emission rate into the cavity mode of up to 10^4 times greater than that of independently emitting atoms, but the functional form of the decay versus time is distinctly non-exponential. Figure 7.2 shows the photon output rate R for four representative pulses recorded with different initial atom numbers. Because the rate of collectively enhanced emission per atom scales with N , for higher atom numbers the pulses appear sooner, have shorter duration, and have a higher peak power than for lower atom numbers.

We quantitatively characterize our superradiant pulses in terms of the scaling of key observables with atom number N in Fig. 7.3. Figure 7.3a shows the characteristic N^2 scaling of the peak output power R_{peak} versus atom number that one expects for superradiance. In the presence of decoherence or atom loss, the atom number must exceed a threshold N_t for superradiance to occur. N_t is set by the requirement that the collectively enhanced decay rate exceeds the atomic decoherence rate. Above this threshold, we predict $R_{peak} = \frac{1}{4\xi} N_x^2 C\gamma$, where $N_x = N - N_t$ is the




Figure 7.3: (a) Peak photon output rate, R_{peak} , versus initial total atom number. The black line is a quadratic fit to the output power. We observe a horizontal offset in the fit, indicating a threshold atom number N_t (black vertical line in b,c,d). The threshold results from decoherence and atom loss and is a signature of lasing that would not be present for single-atom emission. (b) FWHM pulse duration t_w (blue) and delay of peak power t_d (red) versus initial total atom number. The blue line is a fit to the predicted functional form for the pulse duration, with N_t determined from the fit to R_{peak} . The dashed red line is a fit to the pulse delay, assuming threshold is due to atomic homogeneous broadening without atom loss. The solid red line is a fit to the pulse delay, assuming the threshold is set only by atom loss from the lattice. (c) The ratio of emitted photons M_{tot} to the number of atoms in excess of threshold N_x , plotted versus atom number. The dashed line is the average ratio, showing that 48% of the atoms in excess of threshold participate in the superradiant pulse, largely independent of N .

total number of atoms in the lattice in excess of the threshold atom number. The inhomogeneous coupling of the atoms to the cavity mode is accounted for by the numerical factor $\xi \approx 2.95$ (Details on how we account for decoherence and inhomogeneous coupling follow). From a fit of this form, we extract a threshold of $N_t = 3.3 \times 10^4$ atoms. From the fitted N_t and known C and γ , the measured peak photon emission rate R_{peak} is 0.7(4) times the above predicted rate.

The time duration of the superradiant pulse provides a measure of the collectively enhanced decay rate. The measured FWHM time duration versus atom number N is shown in Figure 7.3b

(blue points). We predict that the FWHM duration t_w of the pulse is given by $t_w \approx 7.05/(N_x C \gamma)$, such that the enhanced decay rate scales linearly with the excess atom number N_x . We fit this functional form (blue line) to the data, with the threshold held fixed to $N_t = 3.3 \times 10^4$ atoms from above. From this fit we find that the measured FWHM is 1.4(7) times the predicted FWHM given the known C and γ .

The measured delay time t_d of the peak in output power versus atom number N is shown in Figure 7.3b (red points). In the presence of homogeneous broadening of the atomic transition, but with no atom loss, we expect the delay time to be given by $t_d \approx \frac{2(\ln N + \gamma_e)}{N_x C \gamma}$ (dashed red fit line with N_t fixed and C fitted), where $\gamma_e \approx 0.577$ is the Euler-Mascheroni constant. In the presence of atom loss from the lattice at rate γ_ℓ , we observe a delay time in numerical simulations of the form $t_d \approx \alpha \gamma_\ell + \frac{2(\ln N + \gamma_e)}{N C \gamma}$ where α is a constant, a functional form which seems to better describe the data (solid red fit line with N_t fixed and C and α fit.)

We define the number of atoms that participated in a superradiant pulse in terms of the integrated number of photons M_{tot} emitted from the cavity mode, i.e. one photon equals one participating atom. Figure 7.3c shows the number of emitted photons per atom in excess of threshold, M_{tot}/N_x . We observe that above threshold, $M_{tot}/N_x = 0.48(15)$, independent of atom number. Because the peak photon output rate R_{peak} and pulse duration t_w scale as N_x^2 and $1/N_x$, respectively, we do not expect this quantity to depend on atom number. Inhomogeneous coupling to the cavity mode would predict $M_{tot}/N_x = 0.7$ partially accounting for the observed participation. Atomic collisions, which lead to an atom-number dependent contribution to N_t may account for the additional reduction of participation. We may contrast this level of participation to the case where atoms emit independently: if there were no stimulation, only 0.1% of atoms would emit a photon into the cavity mode during our measurement time.

The preceding measurements all rely on a calibration of the number of atoms in the optical lattice. As is often the case, this is somewhat tricky. After allowing a fixed time T (typically 300 ms) to record the superradiant pulses, we perform atom counting using a resonant 461 nm fluorescence beam and a CCD camera. We calibrate the fluorescence signal by measuring a vacuum

Rabi splitting on the 1S_0 to 3P_1 , $F = 9/2$ transition. To measure total atom number, we apply pump lasers to the 3P_0 and 3P_2 to 3S_1 transitions to drive all atoms to the ground state. We then infer the number of atoms at the beginning of the 300 ms superradiance time window from a measured background loss rate from the lattice of $\gamma_0 = 2.0(4)/\text{s}$ and a measured additional atom-number dependent collisional loss rate of $\gamma_N = 10(5) \times 10^{-6}/(\text{atom s})$ for excited state atoms. To account for the state-dependence of the collisions, we make a lowest order correction for the decrease in excited state population after the superradiant pulse, which leads to a 5% correction at 175,000 atoms, and a $<1\%$ correction below 100,000 atoms.

7.3 Seeding Atomic Coherence

Instead of preparing the atomic ensemble in the excited state with no initial coherence as before, we can terminate the adiabatic transfer process early and prepare the atoms in a superposition of $|g\rangle$ and $|e\rangle$. Because the transfer is performed with light that is resonant with the cavity mode at 698 nm, the phase factors written in to the atoms are exactly right to allow those atoms to radiate in phase into the cavity mode. This seeds the collectively enhanced emission. Unlike in the non-seeded pulses of figure 7.2, we detect an immediate output of light from the cavity. Figure 7.4a shows a typical output trace resulting from terminating the adiabatic transfer with the Bloch vector just above the equator of the Bloch sphere.

Seeding the coherence in this manner also eliminates the threshold behavior observed in Figure 7.3. Figure 7.4b shows the peak output power R_{peak} versus N for pulses seeded with the Bloch vector just above the equator of the Bloch sphere. This data is well described by a quadratic fit with no offset, i.e. $N_t = 0$.

In contrast to the non-seeded case described earlier, the positive feedback mechanism is not required for seeded pulses. If prepared with the Bloch vector at the equator, the atoms will immediately radiate with a power proportional to N^2 . Sources of dephasing will lead the pulse to die out sooner than it otherwise would, but will not effect the initial radiated power.

We view this technique for establishing collectively enhanced emission with no threshold




Figure 7.4: Seeded superradiant pulses. (a) By terminating adiabatic transfer with atoms in a superposition of $|e\rangle$ and $|g\rangle$, we seed coherence in the atomic ensemble. Here, the atomic Bloch vector is rotated to just above the equator (i.e. a small amount of initial inversion.) Seeding leads to the immediate onset of superradiant emission, in contrast to the non-seeded pulses shown in Fig. 7.2 for which quantum noise seeds the coherence. (b) Peak output power for seeded pulses exhibits N^2 scaling. In contrast to results of Figure 2, seeded pulses exhibit a peak photon output rate consistent with no threshold: $N_t = 0$.

or delay time to be a key tool for the development of superradiant sources. In this work, being able to use these signatures of collectively enhanced emission to incrementally tune the system to meet threshold was essential. In subsequent work, we typically operate with seeded pulses, as they are more robust, less sensitive to fluctuations in atom number and cavity detuning, and offer the additional control of being able to target a desired population inversion. More fundamentally, seeding allows collectively enhanced emission to be achieved in systems that are incapable of meeting threshold, but that may still be of metrological value. If superradiant pulses are used to stabilize

the frequency of another laser, seeding could be used to reduce dead-time that contributes to Dick noise aliasing [35].

7.4 Simultaneous Lasing on Multiple Transitions




Figure 7.5: Lasing on multiple transitions at the same time leads to beating in output power (a) Atoms prepared in the $m_f = 9/2$ and $m_f = 7/2$ Zeeman sub-levels radiate simultaneously into the cavity. (b) Interference between the electric fields radiated from the two transitions leads to a modulation of the output power. An average of 20 time-traces is plotted, showing that the modulation has the same phase between trials of the experiment. (c) We compute a Fourier transform of an averaged time trace, showing a peak at the frequency splitting of adjacent Zeeman sub-levels (≈ 200 Hz). The peak at ≈ 400 Hz indicates that a smaller number of atoms have been prepared in $m_f = 5/2$. A Lorentzian fit to the average power spectrum (red line) returns a FWHM linewidth of 11 Hz, due to the finite length of the pulse. The inset of (b) shows the center frequency of the Lorentzian fit versus applied magnetic field. The blue line shows the predicted splitting between adjacent Zeeman sub-levels.

For all preceding data, we applied optical pumping to populate primarily the 1S_0 , $m_f = 9/2$ sub-level before adiabatic transfer, resulting in a single relevant lasing transition. We can deliberately reduce the efficiency of the optical pumping to populate both the $m_f = 9/2$ and $m_f = 7/2$ ground states, and then adiabatically transfer the atoms into superpositions of ground and excited states with different m_f projections, represented in Figure 7.5a. This creates two separate sub-ensembles of atoms that interact with the same cavity mode but have slightly different

transition frequencies. We observe a modulation in the output power at the magnetic-field induced frequency difference between the two transitions. This modulation is the result of interference between the fields radiated by the atoms lasing on the two transitions. Because more atoms are prepared in the $m_f = 9/2$ sublevel, the total field radiated never goes through zero and the contrast of the modulation is not full. Later on (as described in the next chapters), we became more sophisticated with this kind of technique, but this represents our first attempt at spectral characterization of the superradiant pulse.

Figure 7.5b shows the average of 20 time traces recorded under these conditions, illustrating that the phase of the modulation is the same between trials, a result of seeding coherence into the two transitions. To verify that the observed amplitude modulation is the result of beating between adjacent Zeeman transitions, we compute a Fourier transform of the emitted power (Figure 7.5c) and fit the peak in the power spectrum that corresponds to the output power modulation. The frequency of this peak is plotted against our applied magnetic field in the inset of Figure 7.5b. The slope and offset of beat frequency are consistent with the expected Zeeman splitting between the $m_f = 9/2$ and $m_f = 7/2$ transitions. The smaller peak near 400 Hz indicates that a smaller number of atoms have been left in the $m_f = 5/2$ state.

A Lorentzian fit to the peak in the average power spectrum returns a FWHM linewidth of 11 Hz, primarily reflecting the finite length of the pulse. Because many photons are detected in a trial of the experiment, we can fit the center of the Lorentzian peak with deviation much smaller than its width.

Treating this as a differential frequency measurement of two lasers, we compute a fractional Allan deviation of 2.6×10^{-15} at ≈ 2 seconds, the repetition rate of our experiment. Because many sources of frequency errors that are common-mode to the two lasing transitions are not captured by this measurement, this number does not indicate the ultimate performance of the system as a frequency reference. It does, however, reflect a bound on its quantum-limited instability. In the next chapter, we perform a more complete characterization of the spectral properties of our superradiant pulses by comparison to a very stable reference laser.

7.5 Simulations of Superradiance

To provide theoretical predictions for superradiance in the presence of inhomogeneous coupling to the cavity mode and atomic decoherence and loss, we integrate a set of simplified optical Bloch equations:

$$\dot{J}_z = -C\gamma J_\perp^2 - \gamma_\ell J_z, \quad (7.1)$$

$$\dot{J}_\perp = (C\gamma J_z - \gamma_\perp - \gamma_\ell)J_\perp. \quad (7.2)$$

Here, J_z and J_\perp are the atomic inversion and the magnitude of the coherence, γ_ℓ accounts for atom loss, and γ_\perp accounts for other homogeneous atomic decoherence. In reality, both γ_ℓ and γ_\perp depend on N , and on the instantaneous distribution of population in ground and excited states. For simplicity, and to arrive at the analytic expressions given below, we assume γ_ℓ and γ_\perp to be constants. In arriving at these equations, we have assumed that the cavity mode is on resonance with the atomic transition, and that the cavity mode occupation quickly equilibrates to the instantaneous value of J_\perp . As discussed in ref. [47], the dynamics of superradiance consist of an initial time period, during which quantum fluctuations are important in providing initial coherence, and a later time period during which the classically radiated field greatly exceeds quantum fluctuations and the dynamics proceed classically, governed by the above equations. In the later time period, these equations can easily be solved for the case of homogeneous coupling and $\gamma_\perp, \gamma_\ell = 0$, with solutions of the form:

$$J_\perp(t) = \frac{N}{2} \text{sech} \left(\frac{NC\gamma(t - t_d)}{2} \right),$$

$$J_z(t) = -\frac{N}{2} \tanh \left(\frac{NC\gamma(t - t_d)}{2} \right).$$

To simulate the quantum fluctuations in the cavity field that lead to the onset of superradiance, we include a drive to J_\perp that represents random vacuum fluctuations that are replaced at the cavity amplitude decay rate $\kappa/2$. This allows us to simulate the initial conditions that lead to the time delay t_d . We verify that in the limit of large atom number our simulation reproduces the

expected peak power, as well as the pulse duration t_w and time delay t_d of the peak output power derived in ref. [47].

7.5.1 Effects of Atomic Decoherence and Loss

Atomic decoherence, or a decay of J_\perp at a rate γ_\perp , has the effect of setting a threshold atom number N_t below which a superradiant pulse will not occur without seeding of coherence. This threshold can be derived from Equation 7.2 by setting $J_z = N/2$ to determine the minimum value of N for which fluctuations in J_\perp will grow. For atoms homogeneously coupled with cooperativity C , $N_t = 2\gamma_\perp/(C\gamma)$. For atom numbers above this threshold value, N atoms radiating in the presence of decoherence produce a pulse identical to the pulse that the number of excess atoms, $N_x = N - N_t$, would produce in the absence of decoherence. The solutions for J_\perp and J_z for homogeneous coupling in the presence of decoherence γ_\perp are

$$J_\perp(t) = \frac{N_x}{2} \text{sech} \left(\frac{N_x C \gamma (t - t_d)}{2} \right),$$

$$J_z(t) = -\frac{N_x}{2} \tanh \left(\frac{N_x C \gamma (t - t_d)}{2} \right) + N_t/2.$$

To treat the effects of atom loss, we rely on numerical simulation. In the presence of atom loss at a rate γ_ℓ , we observe a threshold-like behavior at an atom number of $N_{t\ell} \approx 13.6\gamma_\ell/(C\gamma)$, assuming homogeneous coupling. For atom numbers well above threshold, peak output power and pulse duration are well-described by assuming an ensemble of $N - N_{t\ell}$ atoms radiating in the absence of loss. Near threshold, however, the peak power radiated is higher and the pulse duration is shorter than that predicted by this analogy. The delay time of the pulse exhibits a more striking difference: at fixed atom number, the delay time of the pulse peak increases linearly with loss rate up to the threshold loss rate.

From our measured atom loss rate, assuming uniform inhomogeneous coupling (defined below), we expect a contribution to the threshold atom number of $3.0(6) \times 10^4$. The measured value is $3.3(8) \times 10^4$, which may indicate that the threshold behavior is due primarily to atom loss.

The rates of atom loss and decoherence due to atomic collisions scale in proportion to the number of atoms, N . This leads to a threshold atom number whose value depends on N . This raises the threshold atom number compared to its value in the absence of N -dependent decoherence and leads to a constant fractional reduction in N_x . For our measured atom-number dependent loss rate, we expect the threshold atom number to increase by 0.13(7) for each atom added. This effectively decreases N_x by 13% compared to its value if the same threshold were the result of atom number independent processes only. Collisions that do not lead to atom loss, but whose rate would also scale with N could lead to additional decrease in N_x .

7.5.2 Effects of Inhomogeneous Coupling

In our system, the atoms are trapped at anti-nodes of the 813 nm lattice, which are not aligned with the anti-nodes of the lasing cavity mode at 698 nm. The coupling of an individual atom to the lasing mode thus depends on its location along the cavity axis.

We numerically simulate the effects of this inhomogeneous coupling by dividing the atoms into many classes, each with its own location x_i and value of coupling to the lasing mode $g_i = g \cos(kx_i)$, where k is the k -vector associated with the 698 nm cavity mode. We may reasonably assume that before atoms are transferred to 3P_0 , they are effectively uniformly distributed with respect to the phase of the standing-wave 698 nm cavity mode. For simplicity, we analyze the case where the atoms maintain this uniform distribution after being transferred to 3P_0 , meaning that the values of x_i are sampled from a uniform distribution. In reality, the adiabatic transfer is probably more effective for atoms located at the anti-nodes of the 698 nm cavity mode, as these experience a higher Rabi frequency from the transfer beam. The true distribution of couplings for different transfer parameters will thus lie somewhere between the uniform distribution and the homogeneously coupled case.

Below, we summarize the results of the numerical simulations for key parameters under these two conditions: homogeneous coupling and uniformly inhomogeneous coupling. Homogeneous coupling refers to the case where all N atoms are at anti-nodes of the lasing mode and are coupled

with peak coupling g . Uniform inhomogeneous coupling refers to the case where N atoms are distributed uniformly along the cavity mode. In all expressions, C stands for the peak cooperativity

$$C = \frac{4g^2}{\kappa\gamma}.$$

	Homogeneous Coupling	Uniform Inhomogeneous Coupling
peak output rate R_{peak}	$\frac{1}{4}(N - N_t)^2 C\gamma$	$\approx \frac{1}{11.8}(N - N_t)^2 C\gamma$
delay time t_d	$\approx \frac{(\ln N + \gamma_e)}{(N - N_t)C\gamma}$	$\approx \frac{2(\ln N + \gamma_e)}{(N - N_t)C\gamma}$
duration t_w	$\approx 3.5/(NC\gamma)$	$\approx 7.05/(NC\gamma)$
threshold atom number N_t	$2\gamma_{\perp}/C\gamma$	$4\gamma_{\perp}/C\gamma$
participation M_{tot}/N	$(N - N_t)/N$	$\approx 0.7(N - N_t)/N$

Table 7.1: Summary of key parameters and effects of inhomogeneous coupling for pulsed superradiance.

The effect of uniform inhomogeneous coupling on t_d , t_w , and N_t is to reduce C by a factor of 2, corresponding to taking the spatial average of C over the ensemble of atoms. Its effects on the peak photon output rate R_{peak} and the fraction of atoms emitted into the cavity mode M_{tot}/N , however, are more complex. Differing Rabi frequencies within the ensemble cause the atoms to dephase with respect to the polar angle θ . This shortens the length of the collective Bloch vector at the time of peak emission and strands poorly coupled atoms in $|e\rangle$ when the pulse terminates, leaving 30% of atoms in the excited state.

The finite temperature of the atoms within the lattice leads to a reduction of C through several mechanisms. The finite temperature of the atoms leads to a non-zero radial extent of the ensemble, which causes the atoms to sample regions with lower coupling to the cavity mode. At 10 μ K, this leads to a 10% reduction in the spatially averaged C . The finite Lamb-Dicke parameter $\eta = 0.16$ also leads to a reduction C for the motional carrier transition observed here. We estimate

this effect to lead to a 7.5% reduction in C . To compare measured values of R_{peak} and t_w to predictions in the main text, we take into account these two temperature-related effects, as well as the 13% reduction in N_x from the measured atom number dependent loss rate.

Chapter 8

Frequency Stability Characterization of Pulsed Lasing

8.1 Introduction

After demonstrating that we could achieve pulsed superradiant emission from the mHz linewidth transition and characterizing the pulses in the amplitude domain, our next step was to characterize the spectral properties of the pulses and their sensitivity to various perturbations. To do this, we needed a frequency reference to compare against. Fortunately for us, our lab is down the hall from the Ye lab, which has developed one of the most stable laser systems in the world for their optical lattice clocks. Because their clocks also use the strontium clock transition, their laser is also at exactly the frequency we need to compare to our superradiant laser. Ye lab students and postdocs John Robinson, Ed Marti, Aki Goban and Ross Hutson were particularly instrumental in enabling these comparisons.

Through these comparisons, we were able to characterize the short-term frequency stability of our laser light, and found the fractional instability to scale with a coefficient below $10^{-15}/\sqrt{\text{Hz}}$ (a standard deviation of around 300 mHz between subsequent laser pulses). This makes our superradiant laser the most precise active atomic frequency standard realized to date. We also performed a preliminary characterization of the frequency accuracy of our source by comparing to an optical lattice clock operating in the Ye lab. We have found agreement between the clock and our superradiant light at the Hz (fractionally 10^{-15}) level, on par with the stability floor of a modern hydrogen maser (e.g. Microsemi MHM 2010) achieved for averaging times of order one day.

In order to understand the sensitivity of our laser's frequency to environmental perturbations,

we have primarily focused on effects that are unique to our superradiant laser system (as opposed to optical lattice clocks in general). The most important of these effects is cavity pulling — a shift in the output frequency of the laser associated with the optical cavity’s detuning from atomic resonance. We have found this pulling to exhibit an interesting dependence on atomic inversion, which is explored in more detail in Chapter 9. For typical operating conditions, we have found that for a given shift in the cavity frequency, our laser frequency shifts by five orders of magnitude less. This is an encouraging confirmation of an important promise of superradiant lasers: they have the potential to be many orders of magnitude less sensitive than conventional lasers to both thermal and technical cavity length noise sources that limit the performance of today’s most stable lasers. We also demonstrate a mode of operation that provides real-time insensitivity to shifts due to magnetic fields, which may also be of great use for future superradiant lasers.

8.2 Basic Techniques

The basic idea of this experiment is to generate seeded pulses of superradiant light, and compare their frequency to that of the stable reference laser. We do this by using a heterodyne beat-note, a standard technique that involves overlapping the light from our weak superradiant source with the much stronger reference light on a photodiode. In addition to a DC signal corresponding to the total optical power incident on the photodiode, the detector produces an AC signal at the difference frequency between the two lasers. The frequency of this AC signal is our measured quantity of interest.

We can describe the superradiant light incident on the photodiode as a time-varying electric field $E_\ell = E_{0\ell}\cos[\omega_\ell t]$, where ω_ℓ is the frequency of the superradiant laser light. The reference light has an associated electric field of $E_r = E_{0r}\cos[\omega_r t]$, where we assume $E_{0r} \gg E_{0\ell}$. The photodetector is sensitive to optical power, which is proportional to the square of the total electric field $P \propto (E_r + E_\ell)^2 \simeq (E_{0r}\cos[\omega_r t])^2 + 2E_{0r}E_{0\ell}\cos[\omega_r t]\cos[\omega_\ell t] = (E_{0r}\cos[\omega_r t])^2 + E_{0r}E_{0\ell}(\cos[(\omega_r + \omega_\ell)t] + \cos[(\omega_r - \omega_\ell)t])$. The detector is too slow to respond to signals at frequencies near ω_r or ω_ℓ , and so provides a time-average over these signals. The first term in the final expression for P then

corresponds to the DC voltage from the photodetector.

The term $\propto E_{0r}E_{0\ell}\cos[(\omega_r - \omega_\ell)t]$ is our signal of interest, and is within the bandwidth of the photodetector because the reference laser is near in frequency to our superradiant light. Note that the large reference field E_{0r} boosts the size of the signal. By using a sufficiently powerful reference beam, we can make the heterodyne beat large compared to the electronic noise of the detector, even though our superradiant laser is very weak. In our case, the signal frequency $\omega_r - \omega_\ell$ is near 60 MHz, a convenient frequency band for low-noise detection. We further mix the signal down using an IQ demodulator to obtain a signal at a few 10 kHz, which we directly sample into our data acquisition system. Details of the reference laser system can be found in Chapter 3, and in Mike Martin's thesis [148].

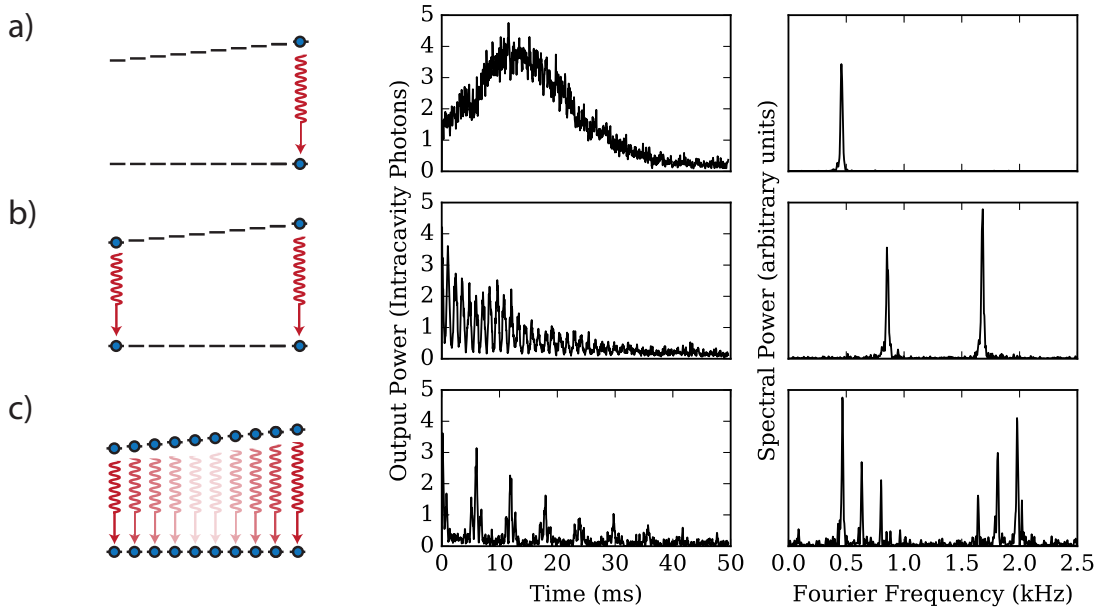


Figure 8.1: Superradiant pulses in the time and frequency domains for different configurations of spin states. (a) We can populate a single nuclear spin state to observe lasing on a single transition. In the time domain (center column) this leads to a simple superradiant pulse. In the frequency domain (right column), we see a single peak. (b) When two nuclear spin states are populated, we observe a beating of the two transitions in the time domain, and two peaks in the frequency domain. (c) When we populate all ten nuclear spin states, we observe an interesting interference pattern in the time domain from simultaneous lasing on the ten transitions. In the frequency domain, we observe peaks associated with the different transitions, with heights determined by the Clebsch Gordan coefficients of the different transitions.

Fig. 8.1 shows typical signals in the time and frequency domain for a variety of experimental configurations. The first row shows a configuration familiar from the previous chapter, with all atoms prepared with a single nuclear spin projection m_f . On a given run of the experiment, the superradiant light comes out in a pulse lasting of order 100 ms or shorter, shown in the power domain in the center panels. We perform a Fourier transform of the heterodyne signal to view the pulse in the frequency domain. For the single transition of the first row, we see a single spectral feature, with width of order 10 Hz, corresponding to the finite length of the pulse. What we are most interested in here is the center frequency of this feature, which we can obtain by fitting a Lorentzian to the peak in the Fourier domain. The choice of fit function is somewhat arbitrary, but we find that a Lorentzian has a reasonable form at higher offset frequencies than does a Gaussian, and was chosen for simplicity.

Instead of populating a single nuclear spin state, we may populate both stretched (maximal $|m_f|$) states simultaneously, and observe lasing on both transitions at the same time. This is shown in the second row of Fig. 8.1. In the power domain, we see a beating at the frequency of the magnetic-field induced splitting between the two transitions. In the frequency domain, we see two distinct peaks corresponding to light emitted from the two transitions. We can independently fit the center frequencies of these two peaks. Because the sensitivity to magnetic fields is equal and opposite for the two transitions, we can use this configuration to obtain a magnetic-field insensitive signal (characterized in more detail in section 8.5).

Finally, we can populate all 10 nuclear spin states, and observe a fun interference effect between the fields radiated by all 10 transitions (bottom row). In the amplitude domain, this manifests as a series of short bursts of power as ten transitions come into phase with one-another, separated by dark periods when their radiated fields cancel. In the frequency domain, we see peaks associated with each transition, with heights that reflect the relative strengths of the Clebsch Gordan coefficients of the different transitions. This configuration may prove useful for schemes that combine dark evolution periods with bright periods during which phase may be read out.

In each of these panels, the data shown represents a single run of the experiment, giving an

indication of the high signal to noise obtained from a single pulse. In contrast to a passive clock, where one would have to run the experiment at least once for every point on the frequency axis, we get all of this information from a single pulse. Further, in a passive clock one must find the transition by setting the clock laser frequency to within the width of the feature to be probed. Here, so long as the reference laser is within the detection bandwidth of the heterodyne system (which can easily be GHz), the signal can be measured with uncompromised precision. This could be a substantial advantage for active frequency references in terms of bandwidth and dynamic range.

8.3 Stability Characterization

We characterize the frequency stability of our superradiant light in the two-stretched-state configuration. As a figure of merit, we compute the Allan deviation for different time windows τ . The fractional Allan deviation defined as $\sigma(\tau) = \frac{1}{f_0} \sqrt{\langle (f_{n+1} - f_n)^2 \rangle / 2}$, where f_n is the average frequency in the n th window of length τ from the beginning of the measurement record and f_0 is the absolute frequency of the superradiant light, is a standard metric for quantifying the performance of a frequency reference [149]. The average is performed over all adjacent pairs — each f_n is a part of two pairs.

While each pulse has a width in the Fourier domain of roughly 10 Hz, we can extract the center frequency with much greater precision due to the high signal to noise. This comes from the fact that we collect many photons per pulse.

The Allan deviation of the frequencies of our superradiant pulses is shown in Fig. 8.2, after subtracting a linear drift associated with the reference laser. We compute the Allan deviation for two quantities: the average frequency of the two peaks corresponding to the two transitions (black points), and half the difference between the two peaks (red points). The black points correspond to the stability of the system if used as a frequency reference. We observe an initial behavior of $\sigma = 1.04(4) \times 10^{-15} / \sqrt{\tau}$ for the data shown, which quantifies the short-term stability of our laser. Our most favorable measurement of this value was $\sigma = 6.7(1) \times 10^{-16} / \sqrt{\tau}$, obtained for a separate shorter set of data. This corresponds to a standard deviation of the frequencies of individual

trials of 300 mHz. The recently-obtained record for this quantity is $6 \times 10^{-17}/\sqrt{\tau}$, achieved by simultaneously interrogating two optical lattice clocks based on ytterbium at NIST [26]. Our value is encouragingly close, given the relative maturity of superradiant lasers and optical lattice clocks.




Figure 8.2: Stability of superradiant pulses expressed as a fractional Allan deviation $\sigma(\tau)$ (see text for definition). Two stretched ($m_f = \pm 9/2$) transitions lase simultaneously. The black points correspond to the average frequency between the two transitions, while the red points correspond to half the difference. Lines represent fits of the small- τ points to $\sigma(\tau) \propto 1/\sqrt{\tau}$, the expected scaling for uncorrelated noise between measurements.

The red points (corresponding to half the measured difference in frequency between the two transitions) provide a useful characterization of the fundamental limit on the precision of our measurement. The difference frequency is insensitive to many forms of perturbation that plague the average frequency, for example to noise on the reference laser, cavity pulling effects, and shifts associated with atomic density. The difference frequency is sensitive to magnetic field noise, but this appears to be small on short time-scales. The red points, which exhibit a short-term stability of $4.5(2) \times 10^{-16}/\sqrt{\tau}$ for the data shown provide an estimate for the quantum-limited stability of our laser.

The quantum limit on our ability to estimate the frequency of the laser is given by $\Delta f \simeq 1/(T\sqrt{M})$, where T is the duration of the pulse and M is the number of photons detected. This is the limit imposed by photon shot noise, which can in principle be improved by improving detection efficiency. A more fundamental quantum limit for this system is $\Delta f \simeq 1/(T\sqrt{N})$ where N is the number of atoms, which is the limit imposed by atom shot noise. We confirm that we are photon shot noise limited by placing an OD 0.6 ND filter after the output of the cavity (reference power is unmodified). We measure a reduction in detected power to 0.25 of its original value, and an increase in the short-term instability (first points in Allan deviation plots) of the difference points of by a factor of 1.9(2), consistent with the expected photon shot noise scaling. For the average frequencies, the short-term instability increases by a factor of 1.4(2). This confirms that the average frequency stability is partially limited by effects other than photon shot noise, though of course we already knew this because the sum frequencies are noisier than the difference frequencies.

A major contributor to the added noise on the average frequencies versus difference frequencies is noise on the reference laser. This appears worse than one might think in our measurements, because the duration of our superradiant pulses is short relative to the experimental cycle time (an effect known as Dick noise aliasing [35]). At short averaging times τ , this should contribute a noise contribution of around $5 \times 10^{-16}/\sqrt{\tau}$. For our most favorable dataset, with fractional stability $\sigma = 6.7(1) \times 10^{-16}/\sqrt{\tau}$, this contribution explains the difference between sum and difference frequency standard deviations. For our longer dataset shown in Fig.8.2, there appear to be additional contributions, perhaps due to imperfect performance of the fiber-noise cancellation system or to fluctuating atom number.

For longer averaging windows (large τ), we observe that the measured Allan deviations flatten out. The reference laser has a noise floor set by thermal perturbations to the length of the reference cavity at around 1×10^{-16} , which should set the floor for the black points. We measure a floor that is slightly higher than this, most likely due to long-term atom number drifts. The Allan deviation for the difference frequencies (red points) should be insensitive to noise on the reference cavity. These points also hit a floor at $\sigma(\tau) \simeq 2 \times 10^{-16}$, perhaps due to slowly varying magnetic fields in

the lab.

Contribution to Detection Efficiency	
cavity losses	0.7
other cavity mirror	0.5
heterodyne overlap beamsplitter	0.82
other optics	0.95
fiber-coupling to detector	0.7
detector quantum efficiency	0.84
detector electronic noise	0.9
heterodyne detection	0.5

Table 8.1: Summary of quantum efficiencies in detection chain.

To determine how close we are to being atom shot-noise limited, we must estimate the total efficiency of our detection, η . Contributions to η for typical operating conditions are summarized in Table 8.1. From this, we estimate a total quantum efficiency of $\eta = 0.06$. If we were to work for a heroic total efficiency of $\eta = 0.37$, as demonstrated in [15] (which would involve asymmetric cavity mirrors and homodyne detection), we would expect to improve our quantum-limited precision by a factor of 2.5. This, however, is not a priority at the moment.

8.4 Cavity Pulling

A detailed description of the mechanism behind cavity pulling is provided in the next chapter. For our purposes here, the important consequence is that for small perturbations of the cavity detuning Δ_c from atomic resonance, the frequency of the output of the laser ω_ℓ is “pulled” by an




Figure 8.3: Sensitivity of laser output frequency ω_ℓ to changes in the cavity detuning Δ_c for different initial atomic inversions. Red points correspond to an inversion empirically tuned to minimize sensitivity to cavity detuning. Black and blue points correspond to slightly higher and lower inversions. The slope of the red points is 2×10^{-6} .

amount $\Delta\omega_\ell = P\Delta_c$. We refer to P as the pulling coefficient. The value of P depends linearly on atomic inversion, with a sign that switches between states prepared above and below the equator of the Bloch sphere. Because inversion changes during a superradiant pulse, when we fit the center frequency of the Fourier peak, we are sensitive to a value of P that is averaged over inversion. By preparing states for which the average inversion during the superradiant pulse is near zero, we can make P very small.

Unlike in the previous chapter, where we used adiabatic transfer to prepare the superradiant state, here we prepare the superradiant states by performing a coherent rotation using light resonant with the clock transition and coupled through the optical cavity. By varying the duration of the coherent rotation, we can prepare states with differing initial inversion. In Fig. 8.3 we show the sensitivity of the frequency of ω_ℓ to changes in cavity detuning Δ for different initial inversions. The red points correspond to a case where the initial inversion is empirically tuned to give a low pulling coefficient (measured $P = 2(3) \times 10^{-6}$). Compared to the red points, the black and blue points

correspond to preparation pulse areas 13% higher and 7% lower respectively, and have pulling coefficients roughly an order of magnitude higher ($P = -2.5(5) \times 10^{-5}$ and $P = 3.0(3) \times 10^{-5}$ for the black and blue points, respectively). While these values do differ, for any of these configurations the output frequency of the laser is highly insensitive to the cavity frequency. This is a key promise of superradiant lasers.

8.5 Magnetic Field Insensitivity




Figure 8.4: A magnetic-field insensitive configuration. The frequencies of the $m_f = 9/2$ and $m_f = -9/2$ transitions (labelled f_{upper} and f_{lower} here) respond in equal and opposite manners to an applied magnetic field, making the average frequency of the two peaks highly insensitive to magnetic fields.

By operating with population in both stretched states, and taking the average frequency of the upper and lower frequency peaks, we can obtain a high degree of insensitivity to magnetic field fluctuations, which would become relevant at long timescales in our system. Fig. 8.4 shows the center frequency of the upper frequency peak plotted against that of the lower frequency peak as a varying magnetic field is applied on different trials. A linear fit to f_{Upper} versus f_{Lower} yields a value of $-1.000(5)$. The average frequency between the two peaks is thus highly insensitive to magnetic

field fluctuations. At the bias field of 1 G used here, the second order Zeeman shift is 0.233 Hz/G. By measuring the splitting between the two frequency components at the 5 Hz level (which would be easy), one could subtract out this second order correction at the mHz level (provided sufficiently precise knowledge of the second order Zeeman shift coefficient). Note that in contrast to standard optical lattice clocks, where the two stretched transitions are alternately probed on subsequent runs of the experiment, this method provides real-time insensitivity to fluctuating magnetic fields. This provides insensitivity to magnetic fields that are fluctuating on timescales short compared to the cycle time of the experiment, which can actually become quite long in clocks, especially as the duration of the spectroscopic sequence is extended.

8.6 Absolute Accuracy and Other Perturbations

We performed an initial measurement of the absolute accuracy of our superradiant light by making frequency measurements at the same time as measurements were being made on an optical lattice clock down the hall (Thanks go to Aki Goban, Ed Marti and Ross Hutson for making this possible). The goal of this project was not so much to make a claim about the final accuracy of our system, as the current iteration of the experiment was not designed to have great accuracy. Rather, the goal was to determine if there are any unexpectedly large shifts to our laser frequency that would have to be controlled at an unreasonable level in order to achieve a desired stability. With that in mind, we sought to characterize key perturbations at roughly the 1 Hz level, acknowledging that further improvements could be made in the future.

The clock (described in [150]) had systematic errors calibrated at the 0.3 Hz level at the time of our measurement. These errors were dominated by uncertainty in the wavelength of the lattice, as it was referenced only to a wavemeter with 100 MHz accuracy at the time of our measurement.

For this comparison, we take into account known sources of error for both the clock and the superradiant laser, then compare the resulting corrected frequency. This comparison was not performed in a blind manner, but as this is an immature experiment and we are not measuring a quantity of universal importance, I think that this is acceptable for our rough comparison.

Table 8.2 summarizes shifts and uncertainties in our comparison, each of which is described in more detail below.

8.6.1 Lattice Shifts

By far the largest perturbation, if we do not account for it properly, is the shift due to the optical lattice. Because we are constrained to operate with the cavity on resonance with the clock transition at 698 nm and the lattice must be on resonance with another mode near the magic wavelength at 813.4272, we do not have complete freedom in setting the lattice wavelength — we must work with the closest mode available, which could be up to half a FSR away. A deviation from the magic wavelength Δ_{latt} (in frequency), in a trap of depth U contributes a shift to the clock transition frequency of $-0.0155 \text{ Hz}/(\text{GHz } E_{\text{rec}}) \times U\Delta_{\text{latt}}$ [151], where $E_{\text{rec}} = \hbar k^2/2m$ is the photon recoil energy, which is used in the clock community as the unit for lattice depth. We use a trap depth of up to around $1000 E_{\text{rec}}$ (around $185 \mu\text{K}$) deep (which is very deep compared to most clocks).

We can fit away this shift by measuring the frequency of our laser at several wavelengths and power levels. The wavelength of the lattice can be changed either by simply changing the frequency of the lattice laser by 1 FSR (3.7 GHz) and locking to a different longitudinal mode of the cavity, or by changing the length of the cavity such that the clock transition is on resonance with a different longitudinal mode of the cavity. This shifts the frequency of the mode to which the lattice locks by $\text{FSR} \times 698/813$. The cavity PZTs have enough tuning range to place two different longitudinal modes on resonance with the clock transition, and the lattice can lock to either of two modes for each of these, one above and one below resonance. This gives us access to four lattice detunings. In principle, we could detune the lattice laser further from resonance to access more detunings, but we do not find this to be necessary.

For each of these lattice detunings, we measure the frequency of our laser at three different lattice depths. We then perform a fit of the form: $f = a(f_{\text{latt}} - f_{\text{magic}}) \times (P - P_0) + f_0$ to these twelve points. Here, f_{latt} and P are the measured lattice frequencies (using a wavemeter) and

Corrections and Uncertainties		
Effect	Value or Shift (Hz)	Uncertainty (Hz)
Measured Offset	0.1	1
Cavity Drift, Magic Wavelength	0	0.7
Collisions	1	1
Hyperpolarizability	0.3	0.3
Clock Lattice	0	0.3
2nd order Zeeman	-0.04	0.01
DC Stark Shift	0	+0, -0.1
Differential Blackbody Radiation	0	0.06

Table 8.2: Summary of shifts and uncertainties in frequency comparison.

transmitted power for the various points, and a , f_{magic} , P_0 , and f_0 are fit parameters. The value a represents the sensitivity of the clock transition to the magic wavelength. We find a value of $-0.0147(3)$ Hz/(GHz E_{rec}), close to the previously measured value.

The results of this fit are shown in Fig. 8.5. The vertical dashed line represents the fit value of f_{magic} . Our measured value of f_0 is 0.1 ± 1 Hz (prior to the application of corrections described below). To check the robustness of this fit, we have tried fixing the lattice power offset $P_0 = 0$, and including a shift that depends quadratically on the measured PZT voltage for each point. These did not shift the measured frequency by more than the reported error on the fit.




Figure 8.5: Frequency difference between our laser and Ye lab optical lattice clock frequency at different lattice laser detunings and power levels. Different colors of points correspond to different lattice depths. A simultaneous fit (see text) to all data returns a nominal frequency offset consistent with zero, as represented by the horizontal dashed line, with the uncertainty of the fit represented by the blue band. The vertical dashed line represents the fitted magic wavelength (relative to the expected value based on our wavemeter) with fit uncertainty given by the green band.

8.6.2 Atomic Collisions

The next largest source of frequency shifts that we characterized were those associated with collisions between strontium atoms resulting from the relatively high atomic density in our lattice. To get a sense of the scale of these shifts for our typical operating conditions, we spread out the atoms in order to modify the atomic density without modifying the atom number (which affects the pulling coefficient, and could confuse the measurement). To spread out the atoms, we applied sidebands to our lattice laser at ± 1 FSR for 30 ms after initial lattice loading (see Chapter 3). Cavity modes that differ by 1 FSR have opposite symmetry about the center of the cavity — the nodes of one mode correspond to the antinodes of the other. By setting the proper power ratio between the lattice carrier and sidebands, we can flatten out the potential from the optical lattice in a region near the center of the cavity. The atoms are still confined radially, but can expand axially. Once the atoms have expanded to about four times their original extent, we turn off the sidebands to restore the lattice, and then apply cooling for 20 ms. In doing so, we cause less than a 10% change in atom number, and observe no measurable change in temperature between when the atoms have expanded and when they have not. The density in the configuration in which the atoms were not spread out was roughly $5 \times 10^{12}/\text{cm}^3$, though we did not make a precise calibration of this number as our goal was to get a rough estimate of the size of collisional shifts for our typical operating parameters.

By toggling between either spreading the atoms or not spreading the atoms, we can modify the density of the atoms by a factor of roughly 4. We observe shifts to the frequency of the emitted light associated with this toggling. Previous groups have observed a strong inversion dependence to the collisional frequency shift [148], so we measured these density shifts versus the initial state preparation angle (Fig. 8.6). Importantly, our data was taken with two spin projections populated, as in our typical operating condition. In a typical optical lattice clock, a single spin state is used, which prevent so-called s-wave collisions from occurring in fermionic atoms, as these collisions only occur between symmetric motional states of atoms. In our case, s-wave collisions are allowed and

may be dominant for us. We do not observe a substantial inversion dependence in the density shift of our laser.

During the absolute frequency measurement, we used a substantially lower atom number compared to the density measurements. This was so that we could maintain the same atom number when turning down lattice power. We thus apply a conservative shift and error of 1 ± 1 Hz for our comparison. If a more precise absolute frequency measurement were desired, it would probably be preferable to work in the low-density configuration for which these shifts are a factor of four smaller.




Figure 8.6: Measured shift to superradiant output frequency associated with density-dependent collisions versus fraction of atoms prepared in 3P_0 at beginning of superradiant pulse. Data taken with roughly 5×10^5 atoms in lattice at beginning of superradiance.

8.6.3 Other Shifts

Other effects contribute fairly negligible contributions to our measured absolute frequency and uncertainty at the level of our measurements. These include hyperpolarizability, 2nd order

Zeeman shifts, DC Stark shifts and Blackbody Radiation shifts. These are summarized below.

To estimate the magnitude of the hyperpolarizability shift — a shift that scales with the square of the lattice intensity — we use the coefficient $0.3(3)\mu\text{Hz}/E_{\text{rec}}^2$ from Travis Nicholson's thesis [152]. For the deepest lattice depths that we use, this would give a shift of $0.3(3)$ Hz.

To calculate the second order Zeeman shift, which is a symmetric shift to the $m_f = \pm 9/2$ transitions that depends on the square of the magnetic field, we simply need to know the value of our field, which we extract from the spacing between the upper and lower transition peaks. We use the coefficient -0.233 Hz/G^2 , also from Travis Nicholson's thesis, to infer a shift of $-0.04(1)$ Hz.

DC Stark shifts describe shifts to the transition frequency due to static electric fields. It has been reported that dielectric surfaces under vacuum can accumulate surprising amounts of charge, and lead to electric fields that can cause significant perturbation to the frequency of atomic clocks [153]. This effect could be especially concerning for us, due to the ceramic cavity spacer and mirrors that are in close proximity to the atoms. Further, the piezo tubes that control the position of the cavity mirrors are driven with up to 120 V (although the geometry of the piezos should lead to a high degree of cancellation of the associated electric fields at the location of the atoms).

The magnitude of the DC stark shift scales with the square of the electric field at the location of the atoms. This leads to a convenient way of measuring it's value by applying an additional field with external electrodes. The shift $\Delta f_{\text{stark}} \propto |\vec{E}|^2 = |\vec{E}_{bg} + \vec{E}_{\text{applied}}|^2$ where \vec{E}_{bg} is the preexisting background field and \vec{E}_{applied} is the applied field. As the strength of the applied field is scanned, one observes a parabolic frequency shift of the clock transition (Fig. 8.7). By comparing the frequency measured at the maximum of the parabola to frequency measured at $\vec{E}_{\text{applied}} = 0$, one can measure the shift due to the component of \vec{E}_{bg} that lies along $\vec{E}_{\text{applied}} = 0$.

To generate an applied field \vec{E}_{applied} , we simply placed an electrode (a screw embedded in a Teflon block) on the large upper window of our vacuum chamber. We applied a voltage to this electrode relative to the vacuum chamber (which is grounded), and varied this voltage over roughly ± 500 V. A sample scan is shown in Fig. 8.7. We repeated this measurement with different electrode positions on the large top window in order to vary the direction of \vec{E}_{applied} , and observed frequency




Figure 8.7: Calibration of DC Stark shifts. Measured frequency is plotted against voltage applied to electrode on large window of vacuum chamber. Frequency shifts due to background fields in the direction of the applied field from the electrode manifest as an offset between the minimum of the parabola and its value at zero electrode voltage. Shifts here are consistent with zero.

offsets always well below 0.1 Hz. We thus assign an uncertainty associated with DC Stark shifts $+0, -0.1$ Hz (since the sign of the shift is well-defined).

Finally, we consider the shifts to the clock transition frequency associated with blackbody radiation, which have merited great effort in the clock community (See refs [29, 28, 55, 154, 155], for example). Because we are performing a comparison against another strontium system, we do not care about the absolute value of these shifts, but rather the difference in these shifts between our system and the clock. The slope of the shift versus temperature is roughly $0.03 \text{ Hz}/\mu\text{K}$. We measure the temperature of the two vacuum chambers to be well within 1 K of each other. However, because the lattice clock has recessed windows that are close to the atoms and to the MOT coils, and which are not accessible for temperature probes, we assign a more conservative temperature uncertainty of 2 K, corresponding to a differential error associated with blackbody shifts of 0.06 Hz.

Accounting for these corrections and uncertainties, we arrive at a measured frequency difference between our superradiant light and the value measured in the Ye lab optical lattice clock of 1.2 ± 1.6 Hz.

Chapter 9

Cavity-Mediated Spin-Spin Interactions

In the previous chapters, all measurements were performed with the optical cavity tuned right on resonance with the atomic transition. In that regime, the collective interactions lead primarily to an enhanced decay rate of the atoms into the cavity, which we can naturally describe as a dissipative process.

In this chapter, I present a detailed exploration of the effects of detuning the cavity from atomic resonance. This work was performed in collaboration with Ana Maria Rey's theory group. As one might expect, detuning the cavity from atomic resonance reduces the rate of collectively enhanced decay from the cavity, representing a decrease in the strength of dissipative interactions. More interestingly, interactions mediated by the detuned cavity also lead to inversion-dependent frequency shifts of the atomic transition and radiated light as well as spin-locking effects, which can be viewed as Hamiltonian or unitary dynamics.

The inversion-dependent frequency shifts lead to an effect known as One-axis-twisting (OAT), which can generate useful spin squeezing [156], Schrodinger cat states [157], quantum phase magnification [158], and new measures of entanglement [159]. The spin-locking effects can protect collective dynamics against single-particle sources of dephasing that limit the atomic coherence times needed for high precision measurements and for preserving interesting quantum correlations [160].

This work has implications for two fairly distinct types of scientific pursuits: the development of precision metrology tools, and quantum simulation. For the development of a narrow-linewidth

superradiant laser, it is clearly very important to understand how changes in the frequency of the optical cavity perturb the frequency of the emitted light. We refer to this effect as cavity-pulling, and the results presented here can be viewed as a detailed exploration of cavity pulling effects in a superradiant laser. Further, there are good reasons why one might want to operate an optical lattice clock inside of an optical cavity, such as enabling non-destructive atom counting. This work provides tools for understanding the nature and magnitude of frequency shifts to the clock that may be introduced in this configuration.

From the perspective of quantum simulation, this work provides a new way of generating interactions between atoms. Previous implementations of atomic coherent interactions result from direct atomic collisions [161] direct electric and magnetic dipole interactions [162, 163, 164, 165, 166, 167], phonon-mediated couplings in trapped ions [168, 169, 170], and photon-mediated coupling in a driven optical cavity [67]. We add to this list spin-exchange interactions between ultra-long-lived optical dipoles mediated by photons in an undriven optical cavity, which are decoupled from atomic motion, and are fully tunable by controlling the cavity parameters.

9.1 Quantum Description




Figure 9.1: (a) The cavity field mediates interactions between the effective spins corresponding to the optical $1S_0(g)$ - $3P_0(e)$ clock transition. This leads to both a collectively enhanced decay at rate Γ which leads to light leaking out of the cavity, and to interesting unitary dynamics that lead to shifts in the frequency of the emitted light ω_l . (b) The resonance of the optical cavity is detuned from the atomic transition by a frequency Δ .

The interactions present in our system can be described as a collective spin-exchange interaction acting on the effective spins formed by the $^1\text{S}_0(g)$ - $^3\text{P}_0(e)$ clock transition — the ensemble of atoms emits a photon into the cavity, which is then reabsorbed (Fig. 9.1). This process flips one spin from down to up and another from up to down (although because the interaction is collective, we do not know which ones).

The Hamiltonian for the system (which describes only the unitary dynamics) can be written in terms of collective spin operators $\hat{J}_{x,y,z}$, with \hat{J}_z corresponding to atomic inversion, acting on the e - g clock states:

$$\hat{H}_{\text{eff}} = \hbar\chi\hat{J}^+\hat{J}^- = \hbar\chi[\hat{J}^2 - \hat{J}_z^2 + \hat{J}_z] \quad (9.1)$$

Here $\hat{J}^\pm = \hat{J}_x \pm i\hat{J}_y$ and $\chi = 4g^2\Delta/(4\Delta^2 + \kappa^2)$. In the second form, the term $\propto \hat{J}_z$ induces a small single-particle rotation and can be safely ignored in the large N limit where our experiment operates. For now, I will ignore the fact that the atoms couple with different strengths to the cavity mode. This is important for a quantitative understanding of the dynamics, but for now is not especially important for a qualitative description.

The two other terms each contribute distinct and interesting spin dynamics (see Fig. 9.2). The $\chi\hat{J}_z^2$ term leads to an effect known as one-axis twisting (OAT). OAT represents a frequency shift of the atomic transition whose value depends linearly on the atomic inversion J_z . This leads to a shearing effect of the Bloch sphere, as though it were a rubber ball that you grabbed with a hand on each of the two poles and twisted. If applied to a coherent spin state, represented by the dashed circle in Fig. 9.2a, OAT will shear the noise distribution, making it narrower along one direction but wider along the other, generating useful squeezing and entanglement[156, 67, 69].

The final $\chi\hat{J}^2$ term leads to an effect known as many-body gap protection or spin-locking [160]. By creating an energy splitting between states of different J , this term suppresses transitions between states of different J that are caused by static (or slow) frequency shifts to subsets of the atoms. For example, if one prepares a maximally symmetric state with $|J| = N/2$, this term will




Figure 9.2: Illustration of the effects of the two terms in $H = \hbar\chi[\hat{J}^2 - \hat{J}_z^2]$ (a) The $\chi\hat{J}_z^2$ term causes a twisting of the Bloch sphere, referred to as one-axis-twisting (OAT). A quantum noise distribution (represented by the dashed circle) will be deformed by this twisting, leading to a potentially useful entangled state. (b) The $\chi\hat{J}^2$ term leads to a so-called many-body gap protection effect, where dephasing of the atoms due to inhomogeneous noise sources is suppressed.




Figure 9.3: A Bloch sphere representation of the atomic ensemble radiating into a detuned cavity. The phase of the cavity field C (ϕ_C) is locked to the phase of the atomic coherence J^+ (ϕ) up to a fixed offset $\phi_r = \arctan(2\Delta/\kappa) + \pi/2$. The Bloch vector precesses about the cavity field axis, which is the source of both collectively enhanced decay and of shifts in the output frequency ω_ℓ .

suppress transitions to states of lower J , which could be very useful in metrological applications where maintaining atomic coherence is important.

Spin-locking effects generated by exchange interactions [171, 172] have been observed in a variety of physical systems including NMR experiments in spin polarized hydrogen [173], ^3He - ^4He mixture [174] and in trapped cold atoms [72, 73, 74]. In all of these systems, the exchange interactions emerge from direct collisions. In contrast, these effects emerge in our system from the exchange of optical photons in a cavity.

9.2 Classical Description of OAT

I will now present a classical picture for these dynamics using the Bloch sphere. To describe the system, we need to keep track of expectation values corresponding to the atomic inversion J_z and the atomic coherence $J_+ = J_x + iJ_y$, a complex number which represents the projection of the Bloch vector onto the equatorial plane of the Bloch sphere, and the cavity field C (another complex number that we represent on the equatorial plane of the Bloch sphere). As discussed in Chapter 2,

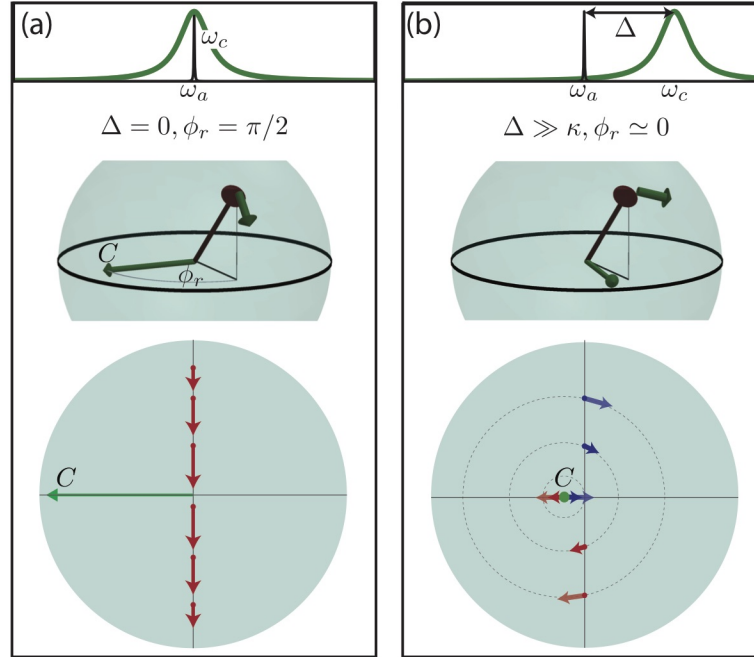


Figure 9.4: (a) When the cavity is on resonance with the atomic transition ($\Delta = 0$) the cavity field C is perpendicular to J^+ . This causes a rotation of the Bloch vector downwards, and is the physical origin of superradiant emission. (b) When the cavity is tuned far off resonance from the atomic transition ($\Delta \gg \kappa$), the cavity field C is nearly aligned with J^+ . Now, the rotation of the Bloch vector about C is primarily sideways. Because C follows the Bloch vector as it rotates, this leads to a frequency shift of the Bloch vector precession, and of the emitted light.

the Bloch vector rotates about the vector that represents the cavity field. When the cavity is on resonance with the atomic transition, the cavity field vector is always perpendicular to the atomic coherence. This causes the Bloch vector to rotate downwards, and is the origin of the collectively enhanced emission (Fig. 9.4). The collective emission leads to a decay of the collective excitation at a rate $\Gamma = 4g^2\kappa/(4\Delta^2 + \kappa^2)$.

The situation can be generalized to the case where the cavity is tuned off resonance with the atomic transition by a frequency Δ . In this case, the cavity mode represents a harmonic oscillator that the atoms are driving off resonance. Like any harmonic oscillator, the response of the cavity to an off-resonant drive will be shifted in phase relative to its resonant response — it will exhibit a

phase lag if the cavity is driven above resonance and a phase advance if it is driven below resonance. On the Bloch sphere, this manifests as a phase shift of the rotation axis C . The field's azimuthal angle ϕ_C follows the azimuthal angle of the Bloch vector's projection onto the x-y plane ϕ . The relative azimuthal angle is $\phi_r = \phi - \phi_C$, and is set by the cavity detuning $\phi_r = \arctan(2\Delta/\kappa) + \pi/2$.

When the cavity is tuned far from resonance and the field becomes nearly parallel with the atomic coherence ($\phi_r = 0$ or π), depending on the sign of the detuning. The effect of the field-generated rotation on the Bloch vector is no longer a rotation downward, but sideways — the cavity field drives a rotation of the Bloch vector that changes its azimuthal angle, but not its polar angle (See Fig. 9.4b). This manifests as a modification of the rate at which the Bloch vector accumulates azimuthal phase. We interpret this additional precession of the Bloch vector's azimuthal angle as an inversion-dependent frequency shift $\omega_{OAT} = -2\chi J_z$ of the atomic coherence $J^+ \sim J^+(0)e^{i\omega_{OAT}t}$. Because the cavity mode decays quickly, the cavity field vector will readjust to the new phase of the Bloch vector, maintaining the same relative phase ϕ_r . These effects lead to OAT — the rate at which the Bloch vector accumulates phase, and the sign of the phase shift, depend on whether the Bloch vector is above or below the equator, and by how much.

In general, the dynamics will have contributions from both of these effects, with classical equations of motion for the Bloch vector given by:

$$\frac{dJ^+}{dt} = -i(2\chi + i\Gamma)J^+(t)J_z(t) \quad (9.2)$$

$$\frac{dJ_z}{dt} = -\Gamma J^+(t)J^-(t) \quad (9.3)$$

In the absence of collective decay ($\Gamma = 0$), and assuming homogeneous coupling between the atoms and the cavity, the atomic system evolves according to:

$$\frac{dJ^+}{dt} = -2i\chi J_z(0)J^+. \quad (9.4)$$

Because J_z is constant, this simply represents an azimuthal precession of the Bloch vector at a rate $2\chi J_z(0)$, and is the origin of OAT. The inhomogeneous coupling of the atoms to the cavity makes things a bit more complicated. I will go through how we treat this here, following derivations by Robert Lewis-Swan.

We prepare the atoms by performing a coherent rotation from the bottom of the Bloch sphere using resonant light coupled through the cavity. Atoms contribute to superradiant emission to a degree that depends on their location along the cavity axis both because their coupling to the radiating mode varies and because the state preparation angle depends on the local intensity of the laser that performs the initial coherent rotation. We wish to generalize Eq. 9.4 to take these effects into account. Fortunately, this method of state preparation (in which the inversion and coupling of a given atom are correlated) retains the form of Eq. 9.4 with simple redefinitions of g and $J_z(0)$.

The atom-light coupling for a given atom is characterized as $g_j = g \cos(k_0 j)$, where j indexes the location of the atom, k_0 is the wave-number of the cavity mode at 698 nm, and $2g$ is the single-photon Rabi frequency of a maximally coupled atom. Because each atom radiates into the cavity by an amount proportional to its coupling to the cavity, we define new collective operators that account for this weighting: $\hat{J}^+ \equiv \frac{1}{g} \sum_j g_j \hat{j}_j^+$, where \hat{j}^+ represents the coherence of an individual atom at location j . It is \hat{J}^+ that we can infer from the light leaking out of the cavity. The mean-field equation of motion for \tilde{J}^+ is given by:

$$\frac{d\tilde{J}^+}{dt} = -i\tilde{J}^+ \sum_j g_j^2 j_j^z(t) \quad (9.5)$$

where $j_j^z(t) = \langle \hat{\sigma}_j^z(t) \rangle / 2$, the inversion of the atom at location j . We assume here that $j_j^z(t) = j_j^z(0)$, valid for times short compared to the characteristic evolution time χN , again assuming $\Gamma = 0$. With these assumptions, we can associate $\Delta\omega_\ell = \sum_j g_j^2 j_j^z(t)$ with the frequency shift of the radiated light.

We now plug in the known forms for g_j and $j_j^z(0)$: $g_j = g \cos(k_0 j)$ and $j_j^z(0) = \cos[\theta_0 \cos(k_0 j)]$, where θ_0 is the rotation of a maximally coupled atom ($g_j = g$) and the subscript j now simply provides an index for the location of the atoms that we will average over. We obtain:

$$\Delta\omega_\ell = 4 \frac{g^2 \Delta}{4\Delta^2 + \kappa^2} N \left[\frac{J_1(\theta)}{\theta} - J_2(\theta) \right]. \quad (9.6)$$

Note that the form of this expression follows the homogeneous case, up to a re-definition of g : $g' = g/\sqrt{2}$ and $J_z : J'_z = N(\frac{J_1(\theta_0)}{\theta_0} - J_2(\theta_0))$. I will call these quantities the “effective coupling” and the “effective inversion.” As long as we use these quantities instead of the original, unweighted

ones, we recover the simple linear dependence of $\Delta\omega_\ell$ on (effective) inversion. Note that for small θ_0 , $J'_z \simeq (N/2)\cos(\theta_0)$, and the primary effect of the inhomogeneous coupling is a factor of two reduction in $\Delta\omega_\ell$.

9.3 Experimental Characterization of OAT Term




Figure 9.5: Shifts in frequency of emitted light ω_ℓ versus cavity detuning Δ , showing expected dispersive behavior. Blue line is a fit with cavity linewidth held to its independently measured value.

Experimentally, we explore the variation of ω_ℓ with cavity detuning Δ (Fig. 9.5) at a given J_z (and J'_z). On each trial, we prepare the atoms in the same state near the bottom of the Bloch sphere (for which the shifts are large), and scan the detuning of the cavity mode while recording changes in ω_ℓ . Note that in this case because θ_0 is small and we are interested in the dependence of ω_ℓ on Δ , rather than its overall magnitude, we can ignore effects of inhomogeneity for this measurement. We observe the expected dispersive behavior of the frequency shift versus detuning, as can be seen

by the fitted dispersive curve with the cavity linewidth held fixed to its independently measured value.




Figure 9.6: Main: Pulling coefficient P (black points) and atomic inversion J_z (blue points) versus state prep pulse duration. Blue line is a fit to the blue points with the expected functional form of inversion versus state preparation angle, allowing for a fixed fraction of atoms that remain in the ground state. Black line is predicted value for pulling coefficient given the fit parameters of the blue line, up to an overall scale factor. Inset: the same data can be viewed by plotting P against the value of $J'_z(0)$ inferred from the blue fit line, showing the expected linear scaling.

Next, we demonstrate the linear scaling of the frequency shift with effective atomic population inversion J'_z . To do so, it is useful to define the pulling coefficient $P = \frac{d\omega_\ell}{d\Delta}$, the sensitivity of emission frequency to cavity detuning. Measuring P provides a way of measuring frequency shifts of the emitted light that are due to shifts in the cavity frequency while being insensitive to technical drifts in the stable reference laser. In our system, P is always much less than one, a characteristic feature of superradiant lasers. We expect P to depend linearly on effective atomic inversion.

We vary θ_0 by changing the duration of the state preparation pulse and measure the resulting actual population inversion J_z (Fig. 9.6) using fluorescence detection. Because of the inhomogeneous coupling to the standing wave cavity mode, we expect $J_z/N = \frac{1}{2}J_0(\theta_0)$. We find that in practice

somewhat fewer atoms are transferred to the excited state than we would expect. We account for these by assuming that a certain fraction remain in the ground state, and that the coupling of these atoms is representative of the coupling of the atoms in general. The imperfect transfer may be due to imperfect spin polarization, or to finite laser linewidth. We extract the value of θ_0 at a given preparation pulse duration from a fit to J_z . We allow for an offset in the fit to account for the finite transfer fraction

On separate trials, we measure P by toggling the cavity detuning Δ by $\pm 2\pi \times 30$ kHz, and measuring the resulting change in the emission frequency during the first 8 ms of the superradiance (black points in Fig. 9.6). Vertical error bars reflect statistical uncertainty of each point, while horizontal error bars reflect uncertainty in the initial preparation angle due to laser power and frequency drifts. We expect P to be proportional to J'_z , so we fit the pulling coefficient data to the functional form of J'_z . We hold fixed the values of θ_0 and the fraction of atoms left in the ground state due to imperfect state preparation to the values returned from the fit to J_z , and allow only an overall scale factor to be fit to the measured values of P . This is shown as the black line in Fig. 9.6. There seem to be some additional fluctuations in the data points beyond the pure statistical error bars, which are of unknown origin. One contribution could be due to small changes in J'_z during the 8 ms window that we analyze. Still, the black line clearly captures the main structure of the data.

We may view the same data by plotting P against the value of J'_z extracted from the fluorescence measurement Fig. 9.6, inset. This makes clear the linear dependence between P and J'_z . The ratio of the measured slope to our independently predicted value is $0.7(3)$. The error on this measurement is due to uncertainty in our atom number calibration and in the actual linewidth of the mHz linewidth transition (which determines g).

9.4 The χJ^2 Term

From a classical perspective, we can understand the $\chi \hat{J}^2$ term as contributing a rotation of the collective Bloch vector about itself — there is an energy difference between states aligned with




Figure 9.7: Illustration of the effects of the χJ^2 term. We consider two sub-ensembles, represented by the green and orange Bloch vectors. (a) When the two sub-ensembles have the same transition frequency, the χJ^2 interactions simply lead to a precession about the total Bloch vector J (blue arrow). (b) When the transition frequencies of the two sub-ensembles are split by δ , the two Bloch vectors rotate in opposite directions along the equator when the χJ^2 interactions are absent. (c) When both δ and χ are finite, and $\chi N \gg \delta$, the two sub-ensembles follow small circles and never develop a large opening angle. This is the origin of the many-body gap protection.

the collective Bloch vector, and states anti-aligned with it. For a fully symmetric ensemble of atoms (meaning that each atom has the same Bloch vector) this interaction does not lead to an observable effect on the Bloch sphere. If however the atoms in the ensemble are distinguishable, this term has important consequences.

The simplest way to see this is to consider two distinguishable sets of atoms inside the cavity (for example occupying two different nuclear spin projections). One of these sets contributes to a collective Bloch vector \vec{J}_1 , while the other contributes to a collective Bloch vector \vec{J}_2 . The total Bloch vector is $\vec{J} = \vec{J}_1 + \vec{J}_2$. If \vec{J}_1 and \vec{J}_2 are prepared with some opening angle between them, they will precess in circles about their sum \vec{J} , as shown in Fig. 9.7a.

Next, consider a situation where the two sub-ensembles each have a slightly different frequency, split by $\pm\delta/2$ and start out aligned on the equator of the Bloch sphere. In the absence of interactions, the two Bloch vectors would precess along the equator in opposite directions, leading to a reduction in the length of \vec{J} , followed by revivals.

In the presence of strong χJ^2 interactions however, the rotation about \vec{J} redirects the precession of the Bloch vectors so that they each follow a small loop, one above the equator and the other below, before returning to their original locations. Because the opening angle remains small, the magnitude of J remains fairly constant, providing suppression of the dephasing effects caused

by the frequency splitting δ .

9.5 Experimental Characterization of Effects from the χJ^2 Term

While it is conceptually nice to understand the gap protection in terms of Bloch vectors near the equator, we found it experimentally easier to observe the effects of the χJ^2 term for states near the south pole of the Bloch sphere. We perform two distinct types of experiment that both show experimental manifestations of the effect of the χJ^2 term of the Hamiltonian.

The first experiment mimics the model that I presented in the previous section, with two distinct spin sub-ensembles. To achieve this in practice, we prepare half of the atoms in $m_F = -9/2$ (described by a Bloch vector \vec{J}_1) and the other half in $m_F = 9/2$ (described by a Bloch vector \vec{J}_2). The two ensembles experience a differential Zeeman energy shift $\pm\hbar\delta/2$ proportional to an applied magnetic field.

It is useful to consider the behavior of both the sum and difference of the two Bloch vectors. We define operators $\hat{J} = \hat{J}_1 + \hat{J}_2$ and $\hat{j} = \hat{J}_1 - \hat{J}_2$ that represent the sum and difference respectively. The Hamiltonian can then be written as $\hat{H} = \chi\hat{J}^+\hat{J}^- + \delta\hat{j}_z$ (neglecting Γ , as it does not effect the azimuthal phase of the Bloch vectors).

The mean-field equations of motion (neglecting Γ for simplicity) are given by:

$$\frac{dJ^+}{dt} = -i2\chi J^+(t)J_z(0) - i\delta j^+(t) \quad (9.7)$$

$$\frac{dj^+}{dt} = -i2\chi J^+(t)j_z(t) - i\delta J^+(t) \quad (9.8)$$

$$\frac{d^2 j_z}{dt^2} = -4[\delta^2 + J_z^2(0) + J_\perp^2(0)]j_z + 6\delta\chi j_z^2 + 2\delta\chi|J^+(0)|^2 \quad (9.9)$$

The detuning δ converts the amplitude of J^+ into amplitude in j^+ and back. The presence of interactions ($\chi \neq 0$) suppresses this interconversion through an effect that we may describe as an energy gap.

In general, the variation of j_z with time makes the interpretation of these equations somewhat complicated. However, when the total Bloch vector is prepared near a pole of the Bloch sphere,

energy conservation requires that $|j_z| \ll |J_z|$. If we neglect the term proportional to j_z , then it is straightforward to analyze the normal modes of the system. When $\delta = 0$, we find that J^+ and j^+ represent eigenstates — a system prepared with full coherence ($j^+(0) = 0$) will maintain full coherence, while a system prepared with no coherence ($J^+(0) = 0$) will always have no coherence. On the other hand, if $\delta \gg \chi N$, then we know intuitively that the degree of coherence will oscillate with time as the Bloch vectors of the two ensembles align and cancel, and the eigenstates will be equal combinations of j^+ and J^+ .

In general, we can write the eigenstates as $B_+ = \cos(\alpha/2)J^+ + \sin(\alpha/2)j^+$ and $B_- = \sin(\alpha/2)J^+ - \cos(\alpha/2)j^+$ where the mode mixing angle α is $\tan \alpha = 2\delta/\chi N$. The normal mode frequencies associated with these eigenstates are $\omega_{\pm} = \left(\chi N \pm \sqrt{(\chi N)^2 + 4\delta^2} \right) / 2$.

The frequency splitting at zero detuning $\omega_+ - \omega_- = \chi N$ is exactly the frequency gap that suppresses decoherence from inhomogeneous noise sources. For values of $\delta \ll 2\chi N$, the normal mode characters remain nearly either fully coherent or fully incoherent. The energy gap between the two modes suppresses interconversion between coherent and incoherent states.

In our experiment, we detect the field radiated into the cavity, which is proportional to J^+ . In contrast, the j^+ atomic coherence does not emit light. We thus think of the J^+ and j^+ as bright and dark modes. For small δ , we thus expect that the mode B_+ will be bright, while the mode B_- will be dim, while for large δ the two modes should be equally bright. Further, the two modes should undergo an avoided-crossing type behavior.

We experimentally observe these effects by recording the power spectra from simultaneous lasing on the $m_f = \pm 9/2$ transitions for atoms prepared near the bottom of the Bloch sphere. We scan the value of the inhomogeneous frequency shift δ between trials. The experimental results are displayed in the spectrograms of Fig. 9.8b with the cavity detuned 150 kHz above resonance (upper panels), and 150 kHz below resonance (lower panels). The center panels show spectrograms from a simulation of the mean-field equations using our estimated experimental parameters, which qualitatively reproduce the features of the experimental data, especially an imbalance of the power in the two radiating branches that switches sign with the cavity detuning. The far-right panels

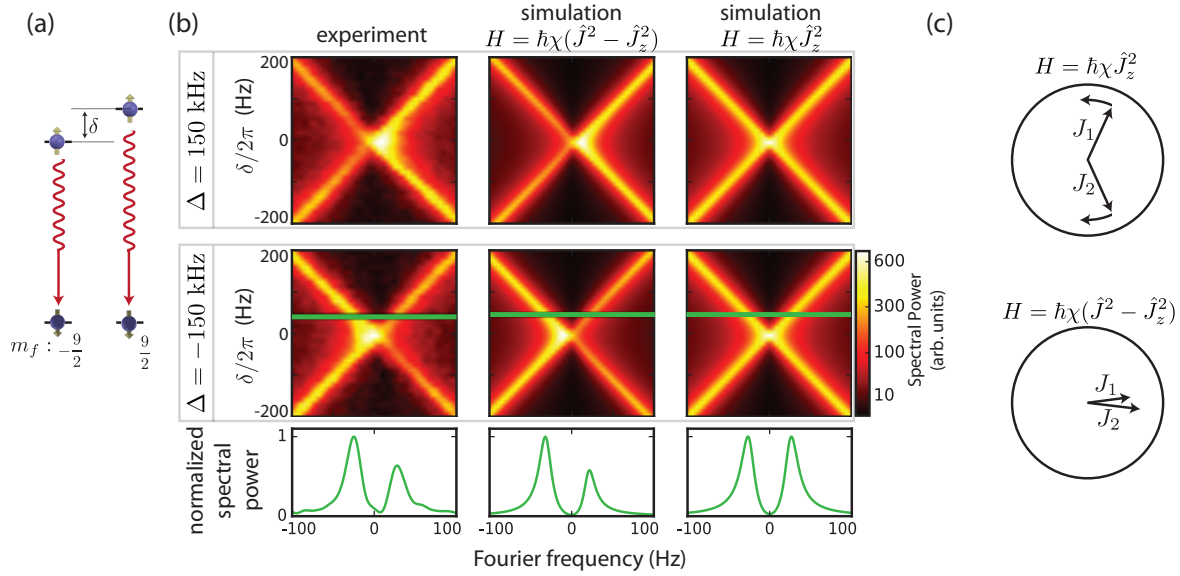


Figure 9.8: (a) We introduce controlled inhomogeneity by simultaneously populating the $m_F = \pm 9/2$ states and applying a magnetic field to split the Zeeman sublevels by a frequency δ . (b) Output power spectra for lasing on the two transitions, with atoms prepared near the south pole of the Bloch sphere. In the presence of a large detuning δ , the two branches correspond to light emitted independently on the $m_F = \pm 9/2$ transitions. When the cavity mode is detuned from atomic resonance by $\Delta/2\pi = 150$ kHz $\simeq \kappa/2\pi$ (top row), we experimentally observe an asymmetry in the two branches (left panel). This asymmetry is reversed for opposite cavity detuning (middle row). A simulation with the full Hamiltonian $\hat{H} = \hbar\chi[\hat{J}^2 - \hat{J}_z^2]$ qualitatively reproduces these features (center column), while a simulation that includes only $\hat{H} = \hbar\chi\hat{J}_z^2$ does not (right column). This indicates the important role of the term $\chi\hat{J}^2$, which is associated with many-body gap protection. Power spectra at fixed δ (bottom row) highlight this comparison, taken at a detuning δ indicated by the horizontal green lines. (c) Cartoon projections of the Bloch vectors associated with the two spin ensembles onto the xy plane for pure OAT Hamiltonian (upper) and full Hamiltonian (lower) in reference frame of total Bloch vector. For pure OAT, the two ensembles precess independently, leading to a large opening angle. For full dynamics, the two Bloch vectors synchronize, preventing a large opening angle from developing.

show the results of a simulation with only the OAT contribution to the Hamiltonian $\hat{H} = \hbar\chi\hat{J}_z^2$. These do not reproduce the asymmetry observed in the data, highlighting the role of the $\chi\hat{J}^2$ term in \hat{H} .

The second feature that we expect to see is a repulsion of the normal modes near $\delta = 0$, corresponding to the many-body energy gap. Instead of using two separate spin states to introduce controlled dephasing of the atomic ensemble, our next experiment relies on inhomogeneous light shifts to controllably destroy and restore atomic coherence. We apply a laser near resonance with

the 1S_0 to 3P_1 transition, whose intensity varies over the axial extent of the atoms in the lattice (Fig. 9.9a). In practice, we achieve this by simply misaligning the center of the roughly 3 mm diameter laser spot from the center of the atoms in the direction along the cavity axis. When we apply this laser with a detuning of typically 14 MHz and a duration of typically 1 ms, we see that the inhomogeneous dephasing destroys superradiance. We can then reverse the sign of the detuning and apply a second pulse of equal duration to reverse the dephasing and restore superradiance (Fig. 9.9b).

This procedure allows us to controllably swap the atoms between the bright state (non-dephased) and the dark state (dephased). While we are not working with the two spin-ensembles that we used to derive the forms of the bright and dark states, the concept is general and can be meaningfully applied here anyway. We are interested in measuring the frequency differences between these two states. However, because the dark state does not radiate light for us to detect, we need to be a bit more sophisticated to measure its frequency.

We prepare a bright state near the south pole of the Bloch sphere using our usual method. We then measure the optical phase (ψ_1) of the emitted light relative to the reference laser. We next dephase the atoms using the inhomogeneous laser, populating a dark state. After a variable amount of time T_{hold} , we rephase the atoms, restoring population to the bright state and make a second phase measurement of the emitted light (ψ_2). We may then interpret the slope of the phase difference $\psi_2 - \psi_1$ versus T_{hold} as the frequency of the dark state. If we follow exactly the same procedure, but do not apply the dephasing and rephasing beams, we instead interpret this same slope of the phase difference as the frequency of the bright mode.

In practice, we actually measure how the phase difference changes when we toggle the cavity between two detunings $\Delta = \pm 30$ kHz: $\Delta\psi = (\psi_2 - \psi_1)|_{30 \text{ kHz}} - (\psi_2 - \psi_1)|_{-30 \text{ kHz}}$, which is related to the pulling coefficient (Fig. 9.10). For the atom number and inversion used here, we find that when we leave the atoms in the bright state by not applying the dephasing or rephasing pulses, the slope of $\Delta\psi$ versus T_{hold} represents a frequency shift of the bright state due to cavity-mediated interactions of ± 4.5 Hz. When the dephasing and rephasing steps are introduced, the slope of




Figure 9.9: (a) We introduce controlled, reversible dephasing by illuminating the atoms with light near resonance with the 1S_0 to 3P_1 transition. The beam is offset from center with the atomic cloud, so the atoms experience a spatially varying intensity, and therefore a spatially varying light shift. (b) Applying the beam for roughly 1 ms scrambles the phases of the atoms, reducing atomic coherence and superradiant emission. Applying the laser again, this time with the opposite detuning from the 1S_0 to 3P_1 transition unscrambles the atomic phases, restoring atomic coherence and superradiance.

$\Delta\psi$ versus T_{hold} is consistent with zero, indicating no shift of the dark state due to the cavity.

The difference in frequency between the bright and dark states is the direct manifestation of the many-body gap for a system near the south pole of the Bloch sphere.




Figure 9.10: We measure the change in the difference between optical phase measurements before dephasing (ψ_1) and after rephasing (ψ_2), as the cavity is toggled between opposite detunings $\Delta/2\pi = \pm 30$ kHz: $\Delta\psi = (\psi_2 - \psi_1)|_{30 \text{ kHz}} - (\psi_2 - \psi_1)|_{-30 \text{ kHz}}$. Red (black) points correspond to trials where the atoms were dephased (not dephased) during T_{hold} , corresponding to the bright (dark) normal modes. $\Delta\psi$ shows no dependence on T_{hold} for the dephased case, and a phase shift that scales linearly with T_{hold} when the atoms are not dephased, confirming a frequency shift between the bright and dark portions of the atomic coherence.

To make this connection explicit, we may consider the effect of dephasing and rephasing in a Dicke basis for states prepared near $J_z = -N/2$ under the effect of a Hamiltonian $\hat{H} = \hbar\chi\hat{J}^2$. We neglect the effects of the $\chi\hat{J}_z^2$ term for now, as this commutes with both \hat{J}^2 and \hat{J}_z so we can account for its effects at the end. States in this basis $|J, m_j\rangle$ are eigenstates of \hat{J}^2 and \hat{J}_z , and are labelled by the magnitude of their spin J and their eigenvalue m_j . A maximally symmetric coherent spin state $|\theta, \phi\rangle$ with $J = N/2$ (and corresponding Bloch vector with polar angle θ and azimuthal angle ϕ) is represented by a superposition of Dicke states as

$$|\theta, \phi\rangle = \sum_{m_j=-N/2}^{N/2} c(m_j, \theta) e^{-i\phi m_j} |N/2, m_j\rangle \quad (9.10)$$

where the coefficients $c(m_j, \theta)$ are real. The constant relative phase between adjacent values of m_j sets the azimuthal angle ϕ of the coherent spin state.




Figure 9.11: The connection between our observed phase shift and an energy gap can be understood explicitly by considering the effect of dephasing in a Dicke basis for states prepared near $J_z = -N/2$, subject to a pure χJ^2 Hamiltonian. Atoms begin in a superposition of maximally symmetric states with $J = N/2$ (far left states). Dephasing (red arrows) transfers the atomic ensemble to states of minimal coherence for a given J_z . Because the difference in J between the states of maximal and minimal coherence for a given J_z is $J_z + N/2$ for states with $J_z \simeq -N/2$, the energy difference between these states is approximately $\chi(J_z + N/2)$. When coherence is restored in the rephasing step (blue arrows), the different J_z states will have accumulated a phase proportional to this energy shift. Because a linear phase shift between adjacent J_z states corresponds to an azimuthal rotation of the Bloch vector, the measured phase shifts of (b) provide a direct measurement of the energy splitting between manifolds of different J , which is the origin of many-body gap protection.

The relevant Dicke states for the three lowest values of m_j are shown in Fig. 9.11, along with the energy gaps created by a $\hbar\chi\hat{J}^2$ Hamiltonian. When ideal dephasing is applied, Dicke states with constant m_j are coupled such that the maximally symmetric Dicke state $|N/2, m_j\rangle$ is mapped to the state $||m_j|, m_j\rangle$, the Dicke state with the smallest possible value of $J = |m_j|$ compatible with the initial spin projection m_j of each Dicke state.

For Dicke states near the bottom of the Bloch sphere, the energy gap between $||m_j|, m_j\rangle$ and $|N/2, m_j\rangle$ scales linearly with m_j as $E(m_j) \approx \hbar\chi N(m_j + N/2)$. After an ideal rephasing step, the complex amplitudes that evolved phase during T_{hold} will have been mapped back from $||m_j|, m_j\rangle$ to the original Dicke states $|N/2, m_j\rangle$. Because of the time spent in the minimally symmetric states, each Dicke state's amplitude will have acquired a phase shift $\propto \chi N T_{hold} m_j$, equivalent to a change of the coherent state's azimuthal phase ϕ by $\chi N T_{hold}$. Therefore, we can understand that measuring the difference in phase accumulated by the Bloch vector during T_{hold} with and without dephasing/rephasing is a direct spectroscopic measure of the many-body energy gap. The additional OAT dynamics changes which configuration appears to have a measurable phase shift,

but importantly, the two configurations retain a relative phase shift caused by the many-body energy gap.

Finally, we can consider what happens if we only partially dephase the atoms (Fig. 9.12). This corresponds to populating a mix of the bright and dark states. Will the frequency of the emitted light be determined by the bright mode frequency, or the dark mode, or some combination of the two? We vary the power in the dephasing beam, and infer the fractional remaining coherence from the amplitude of the emitted field (which is proportional to the coherence J_+). We then measure the pulling coefficient P associated with the partially dephased state by measuring the frequency shift of the emitted light when the cavity detuning is toggled between $\Delta = \pm 2\pi \times 30$ kHz. We find no dependence of the pulling coefficient on the degree of coherence, as the coherence is modified by a factor of roughly three — the frequency of the emitted light is determined by the bright mode even when the system is partially dark.

This observation is consistent with the form of the Hamiltonian $\hat{H}_{\text{eff}} = \hbar\chi\hat{J}^+\hat{J}^- \equiv \hbar\chi[\hat{J}^2 - \hat{J}_z^2 + \hat{J}_z]$. The J^2 term does not contribute here, as J is unmodified during the relevant part of the experiment. All we see is the effect of the J_z^2 term, which explicitly does not depend on the degree of atomic coherence.




Figure 9.12: We may vary the power in the dephasing beams to only partially dephase the atoms. We can infer the residual coherence from the amplitude of the superradiant emission. We observe that the pulling coefficient is unaffected by this process, maintaining the same value as the fully coherent case (to within our experimental uncertainty).

9.6 Conclusion

Apart from being potentially useful in the context of quantum simulation, the interactions that we have explored here will likely have immediate impacts in the field of precision optical metrology. Understanding the inversion dependence of the pulling coefficient is key for the design and specification of superradiant lasers. Further, it is important to understand the nature of cavity pulling effects if one wishes to use an optical cavity in a clock — for example to provide non-destructive atom counting. Further the same dynamics that we observe here could in principle be used to create spin-squeezed states (as independently proposed in ref.[175]).

Creating useful spin squeezing requires operating in a far-detuned regime $\Delta \gg \kappa$ to ensure that the squeezing dynamics occur before superradiant decay destroys coherence. In our current system, background dephasing at a rate $\gamma_{\perp} \simeq 2\pi \times 4$ Hz, due at least partly to atomic collisions, becomes dominant over collective decay in limiting the pulse duration at a detuning comparable to the cavity linewidth. This prevents us from winning by going into the far-detuned regime, as χ decreases while γ_{\perp} remains constant. In addition to challenges associated with observing any spin squeezing that did occur, generating useful entanglement would require substantial experimental upgrades. Besides addressing the issue of background dephasing γ_{\perp} , it would be useful to operate with a higher cooperativity cavity in order to speed up the dynamics relative to any remaining dephasing mechanisms, and to operate with a somewhat stronger transition such as the Yb clock transition considered in Ref. [175].

Chapter 10

Sawtooth-Wave Adiabatic Passage (SWAP) Cooling

Most implementations of laser cooling of atoms and molecules has utilized broad-linewidth optical transitions. Because the excited state of such transitions is short-lived, it is typical to treat the atoms as generally occupying one or more ground states. The excited state population is typically treated only as a proxy for spontaneous emission, or as a means of modifying the energy and character of the ground state. The same cooling techniques used on broad-linewidth transitions have been extended to narrow-linewidth dipole-forbidden optical transitions [176], including the 7.5 kHz linewidth transition in strontium [75, 76, 77, 78, 79, 80]. Even in these experiments however, the atoms are generally thought of as primarily occupying the ground state, and the cooling mechanisms are fundamentally very similar to their broad-linewidth counterparts.

This chapter concerns a fundamentally new mechanism for cooling that is enabled by the long lifetime of the excited state of a dipole-forbidden transition. We call this mechanism Sawtooth-Wave Adiabatic Passage (SWAP) cooling. SWAP cooling explicitly relies on transferring atomic population to the excited state, and storing it there for significant periods of time. In addition, this cooling mechanism relies heavily on the unitary dynamics of stimulated emission and absorption. Spontaneous emission plays an important, though greatly reduced role in the cooling process relative to many previous cooling methods. This mechanism thus brings up interesting questions relating to entropy redistribution, and may be applicable to the cooling of particles for which spontaneous emission causes severe problems, such as molecules.

10.1 How it Works

10.1.1 Adiabatic Passage

As the name suggests, SWAP cooling relies on a mechanism called adiabatic passage, also known as a Landau-Zener crossing. A detailed explanation can be found in many places, for example in Chapter 5 of ref. [87]. The basic idea of is that a drive (in our case a laser) that couples two atomic states is swept in frequency over the transition frequency between the two states. If the strength of the drive is sufficiently high compared to the rate of the frequency sweep, then the atomic population may be transferred from one state to the other with very high efficiency.

Several intuitive pictures can be used to illustrate this mechanism. The first picture involves the concept of dressed states and an avoided crossing, and is illustrated in Fig. 10.1. The Hamiltonian, written in the reference frame of the drive is given by:




Figure 10.1: Illustration of Landau-Zener crossing in a dressed state picture. Black lines represent energy eigenstates in the presence of drive coupling Ω . An atom that begins in $|g\rangle$ when the drive has a large negative detuning δ_d will be transferred to $|e\rangle$ as the drive sweeps to large positive detuning (red arrows).

$$H = \begin{bmatrix} -\delta_d & \Omega \\ \Omega & 0 \end{bmatrix} \quad (10.1)$$

$\delta_d = \omega_d - \omega_a$, where ω_d and ω_a are the frequencies of the drive and atomic transitions, respectively, and Ω is the Rabi frequency associated with the drive. This Hamiltonian acts on the state vector:

$$\begin{bmatrix} g, n \\ e, n - 1 \end{bmatrix} \quad (10.2)$$

where g, e represent the ground and excited states of the atom and $n, n - 1$ indicate the numbers of photons in the drive field associated with these states. The eigenvalues of this Hamiltonian are shown in Fig. 10.1, along with the character of the eigenstates at large detuning $|\delta_d|$. An atom in the ground state with a large negative detuning corresponds to the left side of the lower mode on the plot. As the drive frequency is swept, we move to the right on the plot. So long as the gap between the two states remains sufficiently large, the atom will stay in the lower-energy state, ultimately being transferred to the bottom right branch, which represents an atom in the excited state with large positive drive detuning.

The condition for the gap being big enough is that $\Omega^2 \gg \alpha$, where α is the rate of the frequency sweep, $\alpha = d\delta_d/dt$. If this condition is not fulfilled, the chance of ending up on the wrong branch is given by $e^{-\pi\Omega^2/2|\alpha|}$.

The second picture involves the Bloch sphere. The atoms start in the ground state, and the drive starts out with a large negative detuning. This drive is represented by a vector that points near the south pole of the Bloch sphere (at an angle $\theta = \tan^{-1}(\Omega/\delta_d)$) about which the Bloch vector

rotates at the effective Rabi frequency $\Omega' = \sqrt{\Omega^2 + \delta_d^2}$. As the detuning of the drive is swept over resonance, the drive vector sweeps up the Bloch sphere. If the rate of precession of the Bloch vector about the drive vector is high enough, then the Bloch vector will follow the drive vector up the Bloch sphere, making tiny circles about it. When the drive ends at a large positive detuning, both the drive vector and the Bloch vector will point near the north pole, indicating nearly complete transfer of the atomic population.

10.1.2 Basic Cooling Mechanism in One Dimension




Figure 10.2: Illustration of basic cooling mechanism. (a) Atoms moving in one dimension with velocity v are illuminated by counter-propagating laser beams. (b) Atoms have a single relevant atomic transition at frequency ω_a . In our case, this transition is dipole-forbidden, so $|e\rangle$ has a relatively long lifetime. (c) Representation of frequencies in lab frame. Because of Doppler shifts resulting from atomic motion, the co-propagating (counter-propagating) laser beams are resonant with $|g\rangle$ to $|e\rangle$ transition when the laser frequency $\omega = \omega_a - kv$ ($\omega_a + kv$). Because ω increases with time, the counter-propagating laser sweeps over the transition before the co-propagating laser. If the atom starts in $|g\rangle$, the counter-propagating laser transfers the atom from $|g\rangle$ to $|e\rangle$ as it sweeps over resonance, and the co-propagating laser transfers the atom back to $|g\rangle$. This results in the transfer of two photon recoils of momenta to the atom in the correct direction to cool.

To understand our cooling mechanism in its simplest form, we may consider a two-level

atom with a long-lived optically excited state $|e\rangle$ and ground state $|g\rangle$ moving in one dimension with velocity v , as shown in Fig. 10.2a. Two counterpropagating laser beams with wave-number k and frequency ω_d are linearly ramped in frequency from below to above the atomic transition frequency ω_a , with full sweep range Δ_s . This ramp is repeated to form a sawtooth pattern in time, as in Fig. 10.2b. Each laser interacts with the two-level atom with a Rabi frequency Ω , which is tuned to be larger than the spontaneous decay rate γ from $|e\rangle$ to $|g\rangle$. The sweep range is adjusted such that $\Delta_s > \Omega, 4kv$. We control the frequency sweep rate α to ensure that it fulfills the Landau-Zener condition $\alpha \ll \Omega^2$ for adiabatic transfer of atomic populations between ground and excited states. Lastly, the jump in laser frequency at the end of each ramp is considered to be instantaneous, and therefore perfectly diabatic.

In the reference frame of the atom, both cooling beams start below resonance with the atomic transition. Doppler shifts due to the atomic motion cause the beam counter-propagating to the atom's velocity to appear kv higher in frequency and the co-propagating beam to appear kv lower in frequency. As the beams sweep upward in frequency, the counter-propagating beam sweeps over resonance first, and adiabatically transfers the atom from $|g\rangle$ to $|e\rangle$. Because of the long lifetime of the excited state, the atom remains in $|e\rangle$ until the co-propagating beam sweeps over resonance and adiabatically transfers it back to the ground state. In this process, the atom has absorbed one photon from the beam propagating against its motion, and emitted one photon into the beam propagating along its motion, resulting in a net momentum transfer of $2\hbar k$ against its motion. The laser frequency is then diabatically jumped back to its start frequency such that the atom remains in $|g\rangle$, and then the process is repeated. Because the direction of the net momentum kick is dictated by the time-ordering of the adiabatic transfers, which is in turn determined by the atom's velocity, these momentum kicks will always be in the correct direction to reduce the velocity of the atom regardless of the direction of its initial velocity. This leads to a compression up of the atomic velocity towards zero.

10.1.3 Where Does the Entropy Go?

So far, the dynamics that I have introduced are completely unitary, and thus cannot in fact lead to true cooling as defined by an increase in the phase-space density of the atoms. The missing ingredient in the preceeding section is that I assumed the atom always starts each sweep in the ground state. If the transition being used has an infinite excited-state lifetime, this assumption would be problematic. As the velocity of the atoms decreases, the time-ordering of the two adiabatic transfers becomes ambiguous, and the atom may end up in the excited state at the end of the sweep. If this happens, the atom will then receive momentum kicks in the direction of its motion rather than against its motion, and will begin to heat.

The presence of even small amounts of spontaneous emission can be sufficient to allow steady-state cooling to occur, by ensuring that the atoms preferentially begin each sweep in the ground state. So long as the total time during a sweep before the first adiabatic transfer and after the second adiabatic transfer exceeds the time between the two adiabatic transfers, excited-state decay will ensure that the atom is more likely to start each sweep in the ground state than in the excited state. This breaks the time-reversability of the otherwise unitary dynamics, and allows this cooling mechanism to remove entropy from the atoms.

Intuitively, how should we think about where this entropy goes? Imagine an idealized scenario where an atom starts with an arbitrary initial velocity. It is then slowed by repeated sweeps, with each removing $2\hbar k$ of momentum, until it reaches some low velocity. At that point, it will begin to scatter photons. For the sake of this example, let us then assume that the atom decays into some other state with high probability during the first scattering event, and stays there forever, decoupled from the cooling process. Is this single photon enough to remove the initial entropy of the atom? To answer that question, we need to ask whether the photon carries enough information to reconstruct the initial state of the atom. In this case, it does (up to a possible ambiguity in the sign of the initial velocity). If the atom were initially moving with high velocity, it would take more sweeps to slow it to the point where it scatters a photon than if it were moving slowly. The initial

momentum of the atom is thus encoded in the time-delay before the first emitted photon.

In a realistic system, things are obviously more complicated. There is not some well-defined cutoff at which the atom begins to scatter, and the atom will likely emit at least several photons during the cooling process. This idealized version illustrates a limiting case, and should hopefully make the idea that large amounts of entropy can be carried away by a single photon seem acceptable, when combined with carefully engineered unitary dynamics. Similar ideas have been pursued in the context of optical pumping in large spin systems [177] and cooling with small numbers of spontaneous emission events [178].

10.1.4 A Dressed-State Picture

We can also understand our cooling mechanism in a dressed-state picture, as illustrated in Fig. 10.3.

We imagine that the left and right propagating lasers represent counter-propagating modes of a ring cavity, and label the number of photons in each mode m for the left going mode, and n for the right going mode (the ring cavity is just a conceptual tool to justify a conserved number of photons.) We may then label the state of the combined atom-light system as $|m, n, a\rangle$, where $a = e, g$ labels the state of the atom. In this simplified model, we will neglect spontaneous emission, so the total number of photons $n + m$ is conserved when the atom is in $|g\rangle$, and is reduced by one when the atom is in $|e\rangle$. The atom can transfer photons from one mode to the other via absorption and stimulated emission. Each of these processes is accompanied by the transfer of momentum $\pm\hbar k$ to the atom. By counting changes in the relative population of the two modes, we implicitly track the change in momentum of the atom.

We plot the energies of the combined atom-field eigenstates (in the atom's rest frame) versus the detuning of the light modes from atomic resonance (as measured in the lab frame). At each laser detuning, we reference the energy of the combined atom-light states to the atom frame energy of the initial photon number. With this subtraction, the uncoupled states (for which $\Omega = 0$) $|m, n, g\rangle$ appear as horizontal lines of constant energy, while the states $|m, n, e\rangle$ appear as slanted

lines because one of the photons has been turned into an atomic excitation, whose energy does not change with laser frequency. These uncoupled states are represented by black lines in Fig. 10.3. For this analysis, we assume that the atom is moving to the right, so the left-going photons (labeled m) have energies increased by $\hbar kv$ relative to the lab frame, while right-going photons have energies decreased by $\hbar kv$. This means that states with different photon number are shifted vertically from each other on our plot.

When we turn on coupling between the atoms and light at frequency Ω , the eigenvalues resemble those shown in color in Fig. 10.3. Avoided crossings with energy splitting $\hbar\Omega$ occur when two states for which either m or n (but not both) change by exactly one cross.

We consider the system to start in $|M, N, g\rangle$ with a large negative detuning of the lasers from atomic resonance (point 1 in Fig. 10.3). As the detuning of the cooling beams is swept to the blue, the system adiabatically follows the dressed eigenstate, until the atom encounters a crossing with $|M - 1, N, e\rangle$. Now, the atom will adiabatically transfer to $|M - 1, N, e\rangle$, and then to $|M - 1, N + 1, g\rangle$, ultimately ending up at point 2. When the detuning of the lasers diabatically jumps back to the red, the system stays in $|M - 1, N + 1, g\rangle$ but is now at point 3. In one cycle of the sweep, we see that a photon has been transferred from the left-going mode to the right-going mode, implying a transfer of $2\hbar k$ of momentum to the atom.

If Ω is increased, or kv is decreased, avoided crossings where more than one photon is transferred become significant. At crossings such as the one between $|M, N, g\rangle$ and $|M - 2, N + 1, e\rangle$, the energy splitting is $\hbar\Omega^3/16(kv)^2$. In the regime where $kv \gg \Omega$, these avoided crossings are very small and for a range of laser frequency sweep rates, the atom will diabatically traverse these crossings (for example remaining in $|M, N, g\rangle$). When this inequality no longer holds, as will be the case once the atom is sufficiently cold, then the atom may adiabatically follow the transition from $|M, N, g\rangle$ to $|M - 2, N + 1, e\rangle$. When this occurs, the atom may then follow a path that leads to the exchange of more than one photon between the two light modes, and as a result experience multiple photon recoils worth of momentum change. This is the physical mechanism for the onset of multiphoton processes, which we observe in both experiment and simulation at low velocity or

high Rabi frequency.




Figure 10.3: Eigenstate diagram of coupled atom-light modes versus detuning of the lasers from atomic resonance. Uncoupled eigenstates are labeled as $|m, n, a\rangle$ (defined in text). Atoms start at point (1), then are adiabatically transferred to (2) by the laser sweep. When the laser detuning jumps, they non-adiabatically project to (3), having given up $2\hbar k$ of momentum by transferring a photon from the left-going to right-going modes. Higher-order crossings, which lead to the exchange of multiple photons become significant when $kv \sim \Omega$

10.1.5 Regimes of Applicability

There are several conditions that need to be fulfilled for this cooling mechanism to work. First, as mentioned before, there is the requirement that the frequency ramp must be slow enough to be adiabatic. This means that we need $\alpha \ll \Omega^2$. In addition, in order to be in the regime in which $2\hbar k$ of momentum is transferred per sweep, which is also likely the regime in which spontaneous emission is highly suppressed, we need the ramp to be fast enough to cross the smaller level crossings in Fig. 10.3 diabatically. This means that $\alpha \gg (\Omega^3/16(kv)^2)^2$. In practice, violation of this second condition does not prevent the cooling mechanism from working. What we see instead is that much

larger amounts of momentum are transferred in a single sweep. This appears to lead to larger forces and faster cooling than would be possible if only $2\hbar k$ of momentum were transferred per sweep, but has uncertain implications for the final temperature of the atoms and for the degree to which spontaneous emission is suppressed.

Fulfilling both of the above requirements at the same time requires $4kv \gg \Omega$, which will always be violated at some point if the atoms cool far enough. Intuitively, we can understand this requirement as follows: the characteristic frequency range over which adiabatic transfer occurs is set by Ω . The frequency difference between the co-propagating and counter-propagating laser frequencies is $2kv$. So, when $kv \lesssim \Omega$, the time-ordering of the adiabatic transfers becomes ambiguous. Interestingly, it appears that rather than the effects from the two beams cancelling out, it turns out that in this regime the typical momentum transfer is much larger than in the regime with well-resolved adiabatic transitions. Very roughly, in the dressed state picture of Fig. 10.3 we can see that the lines generally go down and to the right, so if the atom makes all available transitions it will end up quite a ways down on the graph, indicating a large change of momentum.

The next condition is much simpler: we don't want the atom to decay between absorbing a photon from the counter-propagating beam and emitting one into the co-propagating beam. This simply means that $1/\gamma \gg 2kv/\alpha$, where γ is the excited state decay time. This condition is much easier to achieve with a dipole-forbidden transition than a dipole-allowed transition.

Finally, there is an important though slightly subtle requirement on the total sweep range Δ_s . Clearly, we must sweep far enough to address the two Doppler-shifted transitions, which would imply $\Delta_s > 2kv$. However, the requirement is actually more strict than that. Cooling relies on the atom preferentially beginning each sweep in the ground state than the excited state. This means that the time between adiabatic transfers in the middle of the sweep must exceed the total time before and after the adiabatic transfers at the end of the sweep, requiring $\Delta_s > 4kv$. The net force experienced by the atoms goes to zero when $\Delta_s = 2kv$, then switches sign for $\Delta_s > 2kv$.




Figure 10.4: Simulated trajectories for atoms with different initial momenta. Produced by a simulation based on a Monte Carlo wave function trajectory method written by Murray Holland and John Bartolotta.

10.1.6 Simulated Trajectories

Colleagues John Bartolotta and Murray Holland have developed a simulation of SWAP cooling based on a Monte Carlo wave function trajectory method. Everything but the external electric field associated with the cooling lasers is treated quantum mechanically. The atomic wave function contains ground and excited internal states in addition to a discrete family of external momentum states, which are equally spaced in units of $\hbar k$ up to a maximum cutoff. Unitary dynamics are dictated by the Hamiltonian

$$\hat{H} = \frac{\hat{p}^2}{2m} + \frac{1}{2}\hbar\omega_a\hat{\sigma}^z + \hbar\Omega\cos(k\hat{z})\left(\hat{\sigma}^+e^{i\eta(t)} + \text{h.c.}\right), \quad (10.3)$$

where the instantaneous phase of the applied coherent field $\eta(t)$ is the time integrated instantaneous frequency (a phase), where the frequency is ramped in the sawtooth pattern previously described. Because this is an open quantum system, a decay operator proportional to the spontaneous emission rate is also included to fully simulate the master equation.

Fig. 10.4 shows simulated trajectories for atoms that begin at a spread of different initial momenta. Atoms are prepared in a pure quantum state specified by their internal ground state

and fixed momentum. The average momentum over 50 trajectories is calculated at each time step. At high momentum, each sweep lowers the momentum of the atom by roughly $2\hbar k$. At lower momentum, we find that the change in momentum per sweep can greatly exceed $2\hbar k$. This is a manifestation of the multi-photon exchanges that occur when the atomic velocity is relatively low compared to the Rabi frequency Ω .

10.2 Experimental Demonstration

We experimentally explore this cooling mechanism in one dimension using a cloud of freely falling ^{88}Sr . We first cool the atoms to typically around $10\text{ }\mu\text{K}$ using a version of SWAP cooling that creates a restoring force, discussed later in this section. We then apply laser beams to the atoms from both sides. Each beam passes through a separate AOM, which is used to modulate the frequencies of the beams either independently or identically. After applying the beams for some period of time, we turn off all light and allow the cloud of atoms to freely fall and expand for a variable amount of time. We then image the cloud using fluorescence from the 461 nm transition and a camera that views the cloud in a direction perpendicular to that of the cooling beams. We can extract the average velocity and temperature of the atoms by looking at the displacement of the cloud center and the cloud width as a function of time. We use the states $|g\rangle \equiv |^1\text{S}_0, m_j = 0\rangle$ and $|e\rangle \equiv |^3\text{P}_1, m_j = 0\rangle$, where m_j labels the magnetic zeeman sub-level, and apply a magnetic field of several Gauss along the direction of the cooling beam polarization to split out the $m_j = \pm 1$ sublevels by more than our typical sweep range, so the cooling beams nominally interact with a single excited state.

10.2.1 Forces in 1D

First, we wish to characterize the average forces on the atoms as a function of their velocity. The easiest way for us to accomplish this experimentally is to apply a frequency offset Δ between the two cooling beams, as shown in Fig. 10.5a. This creates a moving reference frame F moving at velocity $v_F = \Delta/2k$ in which the two beams appear to have equal instantaneous frequency. From




Figure 10.5: Cooling into a moving reference frame. (a) We apply a frequency offset in the lab reference frame between the two beams, which creates a moving reference frame in which the beams have equal frequency. (b) If laser frequency is swept upwards, (green points) the atoms settle into the moving reference frame, which has a velocity indicated by the dashed black line. If frequency is swept downwards (orange points), the atoms accelerate in the other direction, without cooling. If the frequency sweep is symmetric (purple points), atoms experience a small acceleration due to radiation-pressure imbalance.

the perspective of F , atoms at rest in the lab frame appear to have a velocity $v = -v_F$.

We apply the frequency-swept beams for a variable amount of time T_a , then measure the resulting velocity of the atoms. This is shown in Fig. 10.5b for several sweep configurations. The green points represent the cooling configuration described above, with the laser frequency adiabatically swept from low to high and diabatically jumped from high to low. The two beams are offset by a detuning $\Delta/2\pi = 400$ kHz, and are swept by $\Delta_s/2\pi = 8$ MHz every $33 \mu\text{s}$. We observe that the atoms undergo initial acceleration until they reach equilibrium with the velocity of the moving frame $v_F = 0.14$ m/s, as one would expect for cooling into this moving frame

To rule out the interpretation that the atoms are merely being dragged by the moving standing wave formed by the detuned lasers, we reverse the direction of the sawtooth sweep without changing the relative detuning Δ of the two beams. The standing wave still moves in the same direction

as before, but now the atoms accelerate in the other direction. While the upwards sweep causes saturation of the atomic velocity at v_F , no such saturation is apparent for a downwards sweep. This confirms that the downwards sweep configuration does not lead to cooling.

Finally, when we apply a symmetric triangle-ramped frequency sweep to the lasers, the atoms undergo a much smaller acceleration, likely due to residual radiation pressure from an intensity imbalance between the two beams.




Figure 10.6: Acceleration versus polarization angle between two beams. Acceleration is enhanced when the polarizations of the two beams are aligned. The blue line represents the maximum acceleration for a saturated two-level atom without stimulated emission. The dashed blue line represents the measured acceleration with a single beam. Black points represent acceleration with both beams applied. The dashed black line represents the acceleration that two photon recoils per sweep of the laser frequency would produce.

Because our mechanism relies on stimulated emission, much larger accelerations can be achieved than would be possible with Doppler cooling on such a narrow transition. To quantify this acceleration, we apply the beams for a time much shorter than it takes the atoms to reach equilibrium velocity, and measure the resulting change in velocity. If we apply only one of the two cooling beams, the atoms experience an acceleration consistent with $(\gamma/2)(\hbar k/m) = 155 \text{ m/s}^2$, the expected value for a maximally saturated atom.

We see a far more dramatic effect when we apply both beams at the same time (Fig. 10.6). When the polarizations of the two beams are aligned, we observe a maximum acceleration of around 600 m/s^2 for a sweep period of $20 \text{ } \mu\text{s}$, a factor of almost 4 above both the observed acceleration from the leftward beam alone and the maximum expected acceleration for a two-level atom without

stimulated emission. This measured acceleration is within ten percent of the value we would expect if each sweep of the laser frequency led to two photon recoils of momentum transfer, though this agreement may partly be due to a cancellation between the effects of imperfect adiabatic transfers and multi-photon processes. When the polarizations are made orthogonal, the acceleration returns to near the single-beam value, as only one of the beams interacts with $|e\rangle$.

In contrast to Doppler cooling, the presence of the laser co-propagating with atomic motion enhances slowing. Interestingly, because this laser interacts with inverted atoms, the associated radiation serves to pull the atom toward toward the laser source.




Figure 10.7: Measured acceleration versus the sweep center frequency. As long as the lasers sweep over both atomic resonance (blue region), the acceleration is roughly constant. The sweep range here was 4 MHz.

One advantage of this cooling mechanism is that while it operates on a narrow linewidth transition, its performance is largely insensitive to drifts in the cooling laser frequency. To demonstrate this insensitivity, we apply a shift to the average frequency of the two laser beams and measure the resulting acceleration when both beams are applied. We find that as long as the lasers sweep over resonance, the measured acceleration is largely unmodified. In our case, this allows a large acceleration to be achieved over a range of roughly 4 MHz, even though the natural linewidth of the transition is 7.5 kHz.

Interesting features emerge when we attempt to measure the atomic acceleration as a function of their apparent velocity in the moving reference frame F . For the previous force characteriza-




Figure 10.8: Force versus velocity of moving reference frame, as set by the relative detuning Δ between counterpropagating cooling beams. Red points correspond to an upwards frequency sweep, while black points correspond to a downwards frequency sweep. Blue line represents half of the sum of the red and black points, while the green line represents half of the difference. We attribute the contributions to the acceleration represented by the blue line to transient effects not directly associated with our cooling mechanism.

tions, we were able to fix the apparent velocity of the atoms to be large enough to be well-described by our simple picture of subsequent adiabatic transfers, and to allow time for several sweeps before the atoms reach equilibrium with the moving reference frame. When we lower the velocity of the moving frame, we reach a regime where the number of sweeps required for the atoms to reach equilibrium becomes very small, forcing us to work with very short acceleration times T_a . This makes the results very sensitive to transient dynamics. Further, for fixed Ω , we reach a regime where the requirement $kv \gg \Omega$ no longer holds, and we observe strong effects associated with multi-photon processes.

The results of this measurement are shown in Fig. 10.8. While the center region is contaminated by these transient artifacts, for larger velocity we do observe a roughly dispersive acceleration profile typical of the kind of velocity-dependent force that leads to cooling. The sign of the force switches sign the direction of the sweep is reversed from increasing to decreasing.




Figure 10.9: Experimental study of an atomic ensemble using one-dimensional cooling by adiabatic transfer. (a) An ensemble of atoms, precooled to around 600 μK , is illuminated by frequency-swept counterpropagating beams for a variable amount of time. The one-dimensional temperature is decreased to a steady-state value of 45 μK by the cooling lasers. (b) The phase space density in one dimension (grey trace) and three dimensions (black trace) is increased during the cooling process.

10.2.2 Cooling and Phase-Space Compression in 1D

To characterize how these forces lead to cooling in one dimension, we eliminate the detuning between the two laser beams. We prepare an atomic ensemble with a temperature of roughly 600 μK , which we achieve by briefly turning on the 461 nm MOT beams after initial cooling. We then apply the sawtooth-ramped cooling beams for a variable amount of time before measuring the temperature of the atoms using our usual method. The frequency of the two beams is swept upwards by $\Delta_s = 2\pi \times 6.6 \text{ MHz}$ every 50 μs . In the direction of beam propagation, we find that for these parameters the atoms are cooled from their initial temperature of roughly 600 μK to 45 μK in of order 300 μs or 6 sweeps. It should take roughly 25 photon recoils to slow an atom from the rms velocity associated with a temperature of 600 μK to that of 45 μK , which should take 12.5 sweeps if an average momentum of $2\hbar k$ were transferred per sweep. The fact that we observe faster

cooling than this prediction is most likely due to the presence of multi-photon processes transferring larger amounts of momentum per sweep.

Importantly, we directly observe an increase in phase-space density during the cooling process, and not simply velocity reduction. We measure atom loss during cooling to be negligible and the relative increase in phase space density is $\rho/\rho_o = \Delta x_o \Delta v_o / (\Delta x \Delta v)$, where Δx and Δv (Δx_o and Δv_o) are the measured cloud size and velocity spread after (before) cooling. In Fig. 10.9b, this quantity is shown as the grey line. The black line in Fig. 10.9b accounts for measured heating in the orthogonal directions (which for simplicity we have assumed to be equal in the two directions). Note that because the atoms are not confined and the cooling is only applied along one dimension, this increase of phase space density is much smaller than one would obtain with the same decrease in temperature for three-dimensional cooling in a loose harmonic trap. In that scenario, the spatial extent Δx in each direction would decrease with the square root of the temperature, as would the velocity. The same temperature increase in a three-dimensional trap would lead to an increase in phase-space density of more than a factor of 2000.

We study the scaling of temperature with Rabi frequency by varying the power in the cooling beams. An analytic or intuitive prediction of the final temperature has proved difficult, especially in the presence of multi-photon processes. We compare our experimentally measured temperature to John and Murray's simulation in Fig. 10.10.

Both experiment and simulation show a minimum temperature as a function of Rabi frequency. At low Rabi frequency, the Landau-Zener condition $\alpha \ll \Omega^2$ breaks down, leading to inefficient adiabatic transfers. For the parameters used in (Fig. 10.10), $\alpha = \Omega^2$ when $\Omega = 130$ kHz, roughly the point at which temperature is minimized in both the simulation and experiment.

At larger Ω , both the experimentally measured and simulated temperatures rise with Ω , though the predicted and observed values appear to diverge at large Ω .

Several experimental factors could lead to this divergence. Because an acousto-optic modulator (AOM) is used to sweep the frequency of the cooling lasers, there are significant power variations during the sweep. Including this effect in our simulation results in lower temperatures

at high Rabi frequency, and this effect can explain much of the discrepancy (as indicated by the red band in Fig. 10.10). Furthermore, imperfect laser polarization and finite bias field may lead to cooling on the transitions to 3P_1 , $m_j = \pm 1$, especially at high Rabi frequency.

The exact scaling of the equilibrium temperature is as yet an open question. Because of the limited experimental tunability of important parameters such as atomic mass and photon momentum, we will rely on simulation to provide more details on temperature scalings. This is currently an ongoing project.




Figure 10.10: Temperature versus Rabi frequency for a sweep range of 7 MHz and repetition rate of 15 kHz. Red line represents the results of our simulations, and black points are experimental data. The red band indicates uncertainty of prediction due to statistical uncertainty in simulation results, experimental calibration of Ω and variations in Rabi frequency during the sweep. The vertical line at 130 kHz represents the approximate point at which we expect adiabaticity to break down: $\Omega^2 = \alpha$.

10.2.3 Practical Advantages

SWAP cooling has several key advantages over more common Doppler cooling techniques. For Doppler cooling, the minimum achievable temperature, known as the Doppler temperature T_D is given by ($T_D = \hbar\gamma/2k_B$), where k_B is Boltzmann's constant. The maximum force, set by the rate at which atoms can be returned to the ground state so they can scatter more photons is

set by $F_{max} = \hbar k \gamma / 2$. Both are thus set by immutable properties of a given atomic transition. In order to achieve a low temperature, one would like to use a narrow-linewidth transition. This leads to challenging technical requirements on laser stability, and requires substantial pre-cooling to compensate for the small achievable forces. For even narrower transitions, this force can be so small that atoms cannot be supported against gravity using standard Doppler cooling techniques [176].

SWAP cooling provides tunability in that quantities like the peak force and final temperature depend on parameters that are under the control of the experimentalist, such as the sweep repetition rate and laser power. This allows large forces to be generated using transitions that are also compatible with reaching low temperatures. Further, the requirements for long-term laser stability are greatly reduced compared to Doppler cooling, as the laser must now be stable relative to the entire sweep range, rather than the atomic transition linewidth. In a similar manner, SWAP cooling also provides a great deal of insensitivity to shifts of the atomic transition frequency, for example due to Doppler and Zeeman shifts when capturing from a relatively warm sample, or due to light-shifts associated with deep optical traps. So long as these shifts are small relative to the range of the sweep Δ_s , they do not have a major impact on the cooling performance.

10.3 Creating a Restoring Force

The mechanism that I have presented so far generates a viscous force, but does not lead to spatial confinement. A related mechanism allows us to generate both viscous drag and a spatially varying force, allowing us to form a new type of magneto-optical trap (MOT). I will call this a SWAP MOT. From a practical perspective, this mechanism has proved very useful to us, and is used in the everyday operation of our experiment. Advantages include an increased atom number, a simplified laser system and greater robustness against experimental drifts.




Figure 10.11: Mechanism for realizing a confining force with swept lasers. Counter-propagating beams with opposite circular polarizations illuminate the atoms in the presence of a magnetic field gradient, as in a standard MOT (left). The atoms are located near the zero of the magnetic field. The MOT lasers are swept upwards in frequency in a sawtooth manner, as before. In the simplest system, each laser connects the ground state to a unique excited state, which shift in opposite directions in an applied magnetic field. The lasers sweep over the transition to the state that is shifted to lower energy first, causing acceleration in the direction of the beam that addresses that state, and depleting the population in the ground state (right). Because the excited state is relatively long lived, the atom is less likely to be in the ground state when the other laser sweeps over its resonance. This leads to an imbalance in the number of photons absorbed from each beam, and causes a magnetic field dependent force that can be used to spatially confine the atoms near a magnetic field zero.

10.3.1 How it Works

The experimental configuration for a SWAP MOT is very similar to that of a standard MOT. Three pairs of counter-propagating beams, each pair containing beams with opposite circular polarizations, intersect at the location of a magnetic field zero. Photon recoil forces, which are modified by the presence of a magnetic field gradient, lead to a restoring force towards the field zero. The only important difference between our SWAP MOT and a conventional MOT using a narrow-linewidth transition is that our laser frequencies are swept in a sawtooth manner all the way across atomic resonance.

While in practice we use this in a three-dimensional configuration, I will focus on a single dimension here for simplicity. Also, while this mechanism is appropriate for species with more complex internal structure, I will focus on the simplest applicable level structure, which happens to represent ^{88}Sr . The atom has a single ground state, which connects to a trio of optically

excited states by dipole-forbidden transitions labeled by different angular momentum projections $m_j = -1, 0, 1$, as shown in Fig. 10.11. In the presence of a magnetic field, the states experience Zeeman shifts $\propto Bm_j$. The two lasers are polarized with opposite circular polarizations — in our case the laser beam impinging on the atom from the right has $\sigma+$ polarization and couples the ground state to the $m_j = 1$ excited state. The laser beam incident from the left has $\sigma-$ polarization, and couples the ground state to the $m_j = -1$ excited state. The two laser beams are swept in frequency in a sawtooth fashion as before, with the sawtooth ramp centered on the frequency of the $m_j = 0$ state (though the exact centering is not critical).

We will assume that the atom starts in the ground state at the beginning of a sweep. As the two lasers sweep upwards in frequency, the laser that addresses the transition to the state that was shifted downward by the magnetic field sweeps over resonance first. This transfers the population to that excited state, where it will stay for a while because the transition is long-lived. The other laser will cross resonance with its own transition shortly after. At this point, the atoms are still stored in the other excited state, and thus cannot absorb a photon from the second beam. We assume for the SWAP MOT that the remainder of the sweep is long enough for the atom to decay back down to the ground state before the next sweep begins. This process leads to an imbalance in the number of photons that the atom absorbs from the two beams, and in turn a force on the atom. The direction of the imbalance is set by the sign of the magnetic field, so will switch directions on opposite signs of a magnetic field zero, as desired. While the direction of the magnetic field and sweep direction can both be reversed and still lead to a restoring force, only for an upward sweep do we also get a damping force that enables a stable MOT. For atoms near the magnetic field zero, where the Zeeman shifts are small, the Doppler shift associated with atomic motion can break the tie to determine which beam the atom absorbs a photon from.

In contrast to SWAP cooling, the SWAP MOT relies heavily on spontaneous emission to reset the atom to its ground state after every absorbed photon. I think one could combine symmetrically swept laser frequencies with a rotating magnetic field to transfer the atoms back down though, but this is a project for another time.

10.3.2 Demonstration in One Dimension




Figure 10.12: Atomic acceleration in the configuration described in Fig. 10.11. The direction of the force switches sign depending on the sign of the magnetic field applied along the direction of the counter-propagating beams, and depending on the direction of the sweep (blue and red points). This is the necessary criteria for generating a restoring force toward a zero in a spatially dependent magnetic field. A symmetric frequency sweep generates a smaller force, which does not switch sign with the magnetic field (black points), likely due to an intensity imbalance between the two beams.

Using time-of-flight images as before, we demonstrate an atomic acceleration whose direction depends on the sign of an applied magnetic field (Fig. 10.12, blue and red points). This is the necessary condition to generate a restoring force when combined with an appropriate spatially varying magnetic field. Whereas before we set the polarization of the two lasers to be in the direction of the magnetic field, we now set the two beams to have opposite circular polarizations with respect to a direction along which we apply a variable magnetic field. The direction of the acceleration also depends on the direction of the asymmetric frequency ramp (Fig. 10.12, blue and red points). When we instead apply a symmetric frequency ramp (black points), we see a smaller acceleration whose sign does not depend on the direction of the magnetic field, presumably due to radiation-pressure imbalance between the two beams.

10.3.3 Application to a Real Experiment

We utilize the SWAP MOT in the everyday operation of our experiment. As discussed in Chapter 3, our experimental sequence consists of an initial phase where atoms are accumulated and




Figure 10.13: Fraction of atoms transferred from the broad-line MOT at 461 nm to the narrow-line MOT at 689 nm versus center frequency of the sweep for ^{87}Sr and ^{88}Sr for different sweep profiles. Red points represent an upwards-swept asymmetric sawtooth ramp, which we refer to as the SWAP MOT configuration. Black points represent the more commonly used symmetric triangle ramp. Blue points represent a downwards-swept asymmetric sawtooth ramp, which captures almost no atoms.

pre-cooled in a broad-line MOT formed using the dipole-allowed $^1\text{S}_0$ to $^1\text{P}_1$ transition at 461 nm (blue MOT). This stage prepares the atoms with a temperature of roughly 1 mK. At this point, the atoms are transferred to a narrow-line MOT formed using the 7.5 kHz linewidth $^1\text{S}_0$ to $^3\text{P}_1$ transition (red MOT).

In Fig. 10.13 we explore the efficiency with which we can recapture the atoms from the blue MOT into the red MOT. For this step, it is common practice to frequency-modulate the 689 nm laser to cover the range of Doppler shifts experienced by the atoms coming from the relatively warm blue MOT. Typically, people use a symmetric triangle ramp for this modulation. This configuration is represented as the black points in Fig. 10.13.

If instead we apply an upward-swept sawtooth ramp profile, we capture substantially more atoms, and over a broader range of laser frequencies. The improvements are more pronounced in our system for ^{87}Sr than for ^{88}Sr . This technique both allows us to capture more atoms than the standard approach, and makes our experiment very robust to drifts in the frequency of our 689 nm

laser. With our 689 nm laser locked to our reference cavity, we can go for multiple weeks without tuning its offset frequency with only minor reductions capture and cooling performance.

After recapturing the atoms in the SWAP MOT, the next step in our experiment is to cool and load the atoms into the optical lattice. We typically operate with a deep optical lattice, sometimes up to 200 μK . Because the 813 nm lattice is not a magic wavelength for the $^1\text{S}_0$ to $^3\text{P}_1$ transition, it causes substantial shifts (up to several hundred kHz) to the cooling transition. This would cause challenges for conventional Doppler cooling due to the narrow linewidth of the transition. Because SWAP cooling and the SWAP MOT are largely insensitive to differences between the atomic transition frequency and average laser frequency, we observe no decrease in the efficiency with which we can load the lattice when we work at high depths.

The whole transfer process from the blue MOT to the lattice can be performed with a single set of sweep parameters and magnetic field gradient. We typically sweep the MOT beams by roughly 7 MHz every 50 μs , in the presence of a roughly 2 G/cm magnetic field gradient in the vertical direction. For ^{87}Sr , the laser addresses $^1\text{S}_0$ to $^3\text{P}_1$, $F = 11/2$ transition. We typically transfer up to 50% of the atoms from the red MOT to the lattice. For typical operating conditions, with a lattice depth of 100 μK , we typically achieve temperatures of roughly 10 μK , sufficient for our needs.

In other experiments that use standard Doppler cooling to form a narrow-line MOT, a separate set of beams called “stirring beams” are used that address the $^1\text{S}_0$ to $^3\text{P}_1$, $F = 9/2$ transition in addition to the MOT beams that address the $^1\text{S}_0$ to $^3\text{P}_1$, $F = 11/2$ transition. Because the magnetic moment of the $F = 9/2$ states is smaller than the $F = 11/2$, the stirring beams provide more robust Doppler cooling than the the MOT beams. The stirring beams also prevent population from becoming trapped in unconfined Zeeman sublevels of the ground state [179].

The SWAP MOT does not require stirring beams. Because the laser sweeps over all m_f transitions, even if they are shifted by a large amount, the relatively large magnetic shifts of the $F = 11/2$ states (compared to the $F = 9/2$ excited states) do not present a problem, and there are no unconfined ground states. This substantially simplifies the laser systems required to perform

narrow-line cooling on ^{87}Sr , removing one of the main experimental barriers to working with this isotope.

Overall, the use of the SWAP MOT has greatly simplified our experimental sequence and hardware, and has lead to far more robust operation. While it is likely that lower ultimate temperatures could be obtained with standard Doppler cooling (though we do not yet understand the true limits of SWAP cooling), the temperatures that we obtain with our SWAP MOT are sufficient for our application.

10.4 Extensions and Applications

We can envision a host of extensions and application for SWAP cooling and related techniques. Perhaps the most exciting would be as a means to aid in the cooling of molecules that have narrow-linewidth transitions. Because molecules have many vibrational and rotational state to which they can decay when a photon is spontaneously emitted from an excited state, Doppler cooling techniques have proven very difficult for molecules. The reduced reliance on spontaneous emission for SWAP cooling could remedy this issue.

A promising application could be to the cooling of a molecule like YO, which has a narrow-linewidth optical transition with a linewidth of roughly 160 kHz [180]. In a typical experiment [181], the molecules emerge from a cryogenic buffer gas source with a longitudinal velocity of order 70 m/s. They must then be slowed substantially in order to be captured in a MOT. Because of the large forces and low scattering possible with SWAP cooling, it could be used either for longitudinal slowing of the molecules, or for a transverse cooling stage prior to longitudinal cooling, that would reduce the divergence of the molecular beam during the longitudinal slowing stage.

As another extension, SWAP cooling techniques can be applied using Raman transitions. This allows their application to species that do not have narrow-linewidth optical transitions, such as rubidium, or a broad class of molecules. Further, because spontaneous emission can be added in a controlled and tunable manner, because the effective excited state does not naturally decay, we can imagine using it as a tuning parameter to improve the performance of the cooling. On the

rubidium side of the Thompson lab, Graham Greve and Baochen Wu are currently studying SWAP cooling using Raman transitions in rubidium, and have demonstrated initial cooling results.

From a theory perspective, much work still remains to be done in terms of understanding the final temperature achievable with SWAP cooling, and its dependence on the various experimental parameters.

Chapter 11

A new perspective on gravitational wave detection with cold atoms in space

Recently, several proposals have been published that involve using cold atoms with narrow linewidth optical transitions to detect gravitational waves (GWs) [81, 82, 84, 182]. I found these proposals very appealing as a way to apply the kinds of technologies we are working on towards a very exciting problem in physics.

The proposals fall into two apparently (at first) disconnected categories: atom interferometers (AI) [82, 84, 182], and optical lattice clocks [81]. I became very interested in understanding the fundamental similarities and differences between these two kinds of proposals, which lead to the work presented in this chapter. Several specific questions arise from this quest:

- The AI proposals involve spatial separation of the atom's spatial wave-function, while the clocks do not. Is this important?
- AI proposals include an enhancement from so-called 'Large Momentum Transfer' (LMT) sequences. Is there a way to mimic these in a clock?

The short answers to these questions are 'No,' and 'Yes,' but the rest of this chapter will outline a perspective that should hopefully make this clear and provide further intuition into what is possible with atomic detectors.

In addition to co-authors James and Julia, conversations with Shimon Kolkowitz, Jason Hogan and Peter Graham were key in developing these ideas.

In both types of proposals, an ensemble of atoms with a long-lived optically excited state is

prepared in a drag-free reference frame in or near each of two satellites separated by a large distance. A passing gravitational wave leads to a periodic stretching of space between the two satellites that create phase shifts on light transmitted between the satellites. The laser light interacts with the atoms in each satellite, imprinting the gravitational wave-imposed optical phase shift on the distantly located atoms. The technical laser phase noise is imprinted on the atoms both satellites. As a result, the gravitational wave signal can be retained but the technical laser phase noise canceled by looking at the difference in the phases imprinted by the laser pulses at the two satellites.

In the clock community, the gravitational wave detection is described as a frequency measurement of Doppler shifts that result from the stretching of space between the satellites. In the AI community, the stretching of space is described as causing a phase shift between the laser phase and atomic coherence. A key element of AI proposals is enhanced sensitivity to gravity waves through the use of Large Momentum Transfer (LMT) pulses in which the atoms acquire many photon recoil momentum kicks. In contrast, in optical lattice clocks the atoms are tightly confined in an optical lattice and photon recoil is suppressed. This difference in particular, as well as the language used to describe the devices, would seem to point toward the two sensors somehow being fundamentally different.

In what follows, we provide a comparison of these two sensors to show that the fundamental mechanism for sensitivity is in fact the same between the clock and AI type detectors. In both types of sensors, an internal quantum superposition state of the atoms has imprinted on it a relative phase set by the local phase of the laser pulses that couple a ground and a long-lived optically excited atomic state. The atoms primarily act as a highly coherent phase memory that keeps track of these phase imprints, and allows them to be read out through atomic observables.

By viewing clock-like detectors as phase memories, rather than simply as clocks whose sole capability is to measure frequency, we show that they support the implementation of LMT-like protocols with enhanced sensitivity even when negligible momentum is transferred to the atom. We further show that the relevant phase information can actually be stored in stable ground states that evolve phases at acoustic rather than optical frequencies. This enables useful coherent evolution

times beyond the lifetime of the optically excited state.




Figure 11.1: Optical lattice clock and atom interferometer detectors for gravitational waves. Each type of proposal relies on two ensembles of atoms, one in each satellite. Lasers transmitted between the satellites encode phase shifts due to changing optical path lengths caused by a gravitational wave. The key difference between clock and atom interferometer proposals is that in clocks, atomic recoils due to the momentum of absorbed/emitted photons are suppressed by tightly confining the atoms, while in atom interferometers the atoms are free to recoil.

11.1 Optical Path Length Changes due to Gravitational Waves

For both types of proposals, the important effect of a gravitational wave is a modification of the optical path length L between the two satellites. We will describe the effects of a gravitational wave at angular frequency ω_g by replacing the gravitational wave with a phase modulator that fills the space between the two satellites (e.g. an electro-optic modulator or EOM) sinusoidally driven at frequency ω_g . The drive applied to the modulator leads to a phase shift $\alpha(t)$ on laser light that is launched from one satellite and detected at the other satellite. Physically, $\alpha(t)$ is the relative phase that one would measure at the location of the far atom between a laser pulse that passes through the driven modulator (i.e. mimicing the gravitational wave) and a second laser pulse that runs parallel to the first pulse but passes through an identical but undriven modulator (i.e. mimicing

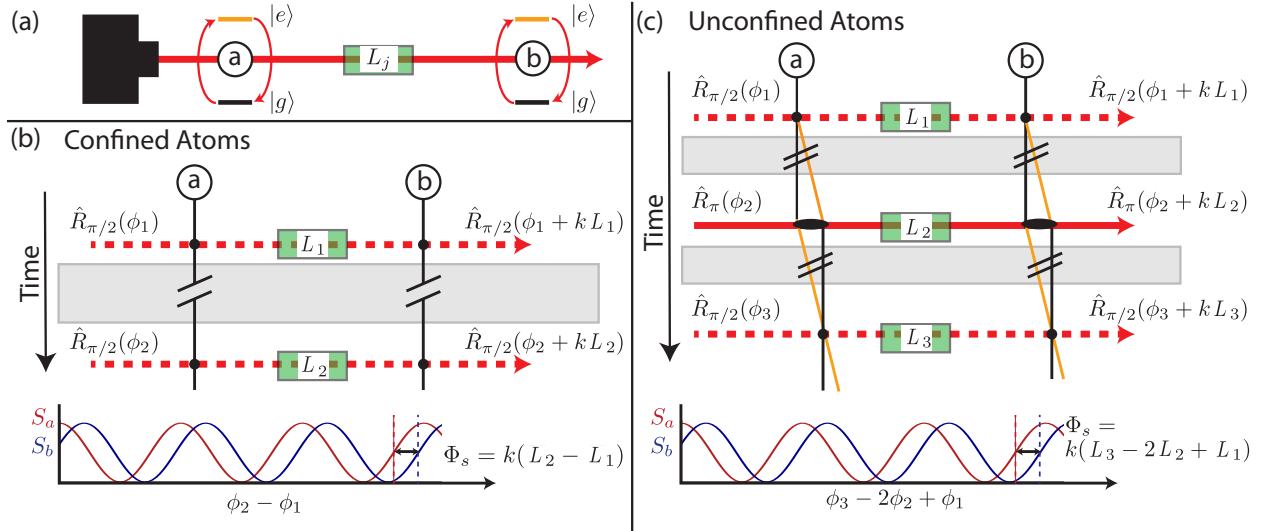


Figure 11.2: A toy model for gravitational wave detection. (a) Two atoms (a and b) are addressed by a single laser resonant with a transition between the ground state $|g\rangle$ and a long-lived excited state $|e\rangle$. A phase modulator mimics the effect of a gravitational wave by modifying the optical path length between the two atoms. (b) For tightly confined atoms (clocks), a simple Ramsey sequence can be used to detect changes in optical path length. A first $\pi/2$ pulse imprints the laser phase on the two atoms when the path length is shifted by L_1 , and a second $\pi/2$ pulse acts when the path length is shifted by L_2 . Random variations in the laser phase (ϕ_1 and ϕ_2) are common to the two atoms, while phase shifts resulting from changes in path length are not. The signal from the phase modulator manifests as phase shift $\Phi_s = k(L_2 - L_1)$ between the excited state probability oscillations of the two atoms. (c) For unconfined atoms (atom interferometers), transitions between $|g\rangle$ and $|e\rangle$ are accompanied by momentum kicks, which necessitate an additional laser π pulse in the middle of the sequence. The phase shift between excited state oscillations of the two atoms is now $\Phi_s = k(L_3 - 2L_2 + L_1)$.

no gravitational wave.) The light is launched from the first satellite at time t , but arrives at the second satellite where it is detected at time $t + T_d$, where T_d is the nominal delay or transit time between satellites.

The total phase shift $\alpha(t_j)$ for a pulse launched at time t_j is calculated by integrating along the optical path. We can define a net effective path length for the pulse $L + L_j = L + \alpha(t_j)/k$ where the small change in optical path length is

$$L_j = \frac{hc}{2\omega_g} (\sin(\omega_g(t_j + T_d)) - \sin(\omega_g t_j)) . \quad (11.1)$$

We have approximated that the gravitational wave's strain is very small $h \ll 1$. In this expression, c is the speed of light in the undriven modulator, $k = 2\pi/\lambda$ is the wavenumber of the laser light, and λ is the laser wavelength.

11.2 Detecting Changes in Optical Path Length

The goal is to understand how AI and clock-like sensors can be used to precisely estimate changes in the optical path length L_j . In our model, shown in Fig. 11.2, we consider two cold atoms labeled a and b that are separated by distance L . The atoms may either be cooled in free space (atom interferometers) or tightly confined (clocks). The atoms have a long-lived optical transition such as exist in alkaline-earth and similar atoms such Sr and Yb. A laser located near atom a can launch pulses of light that interact with both atoms with equal intensity.

The laser pulses interact with the atoms by applying so-called $\pi/2$ and π pulses between ground $|g\rangle$ and excited $|e\rangle$ states. We will assume that the coupling or Rabi frequency is much larger than the atomic decay rate from $|e\rangle$ and any relevant Doppler shifts due to atom motion. We will also assume for simplicity that the laser frequency is exactly equal to the atomic transition frequency. These assumptions rule out some of the capabilities of atom interferometers (such as addressing atoms that have experienced different recoils independently), but retains the basic mechanism for sensitivity. The effect of the laser pulses on the atoms written in a rotating frame

at the atomic transition frequency can then be expressed using the operators:

$$\hat{R}_{\pi/2}(\phi) = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & -e^{-i(\phi+kx)} \\ e^{i(\phi+kx)} & 1 \end{bmatrix} \quad (11.2)$$

$$\hat{R}_{\pi}(\phi) = \begin{bmatrix} 0 & -e^{-i(\phi+kx)} \\ e^{i(\phi+kx)} & 0 \end{bmatrix} \quad (11.3)$$

These operators act on the basis states :

$$|g\rangle \otimes \Psi_g(x) = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \otimes \Psi_g(x) \quad (11.4)$$

$$|e\rangle \otimes \Psi_e(x) = \begin{bmatrix} 0 \\ 1 \end{bmatrix} \otimes \Psi_e(x) \quad (11.5)$$

that are a product of an internal state label $|e\rangle$ or $|g\rangle$, and an external state wavefunction $\Psi_{e,g}(x)$.

Here x represents the distance from a fixed plane where the laser's phase is defined as ϕ . The effect of these interactions is both to transfer amplitude between the internal states of the atoms and

also to imprint the laser's local phase upon the transferred portion of the atom's wave function. The fact that $\Psi_g(x)$ may differ from $\Psi_e(x)$ indicates the possibility of entanglement between the internal and external degrees of freedom of the atom.

For simplicity, we account for laser frequency noise by allowing the laser phases to vary between pulses, but taking the actual laser frequency to be fixed at the atomic transition frequency such that k is constant. One can relax this assumption by accounting for an additional phase that results from the lasers detuning from atomic resonance and cancels the effect of associated changes in the wave-vector k . This leads to the laser frequency noise insensitivity demonstrated in previous proposals [81, 82], but doing this careful accounting does not alter the main results here.

First consider the effect of these rotations for a clock-like sensor in which the atoms are confined to much less than the laser wavelength (known as the Lamb Dicke regime.) In this limit, we can think of the $e^{i(\phi+kx)}$ term as imprinting a spatially constant phase onto the atom, whose value is determined by the location of the atom along the laser's path. For example, a π pulse applied to an atom tightly confined in a trap centered at position $x = A$ and in the ground state $|\Psi_0\rangle = |g\rangle \otimes \Psi_0(x)$ transfers the atom to the state $|\Psi\rangle \approx |e\rangle e^{-i(\phi+kA)} \otimes \Psi_0(x)$. The external state wave-function is to very good approximation unmodified by the pulse, but its internal wavefunction can be seen to have acquired a net phase $\phi + kA$.

In contrast, in AI sensors the atoms are not confined and a change in the internal state is accompanied by a change in the external state. In particular, one cannot neglect the variation of the optical phase factor over the spatial extent of the atomic wavefunction. For concreteness, a π pulse applied to an unconfined atom centered at $x = A$ and in the ground state $|\Psi_0\rangle = |g\rangle \otimes \Psi_0(x)$ transfers the atom to the state $|\Psi\rangle = |e\rangle e^{-i\phi} \otimes e^{-ikx} \Psi_0(x)$. As written, the internal portion of the wavefunction appears identical to that of the confined case, but now the external wavefunction has a spatially varying phase corresponding to one photon's worth of momentum recoil.

A change in the optical path length L_j between the atom and the laser producing the π -pulse manifests in the confined case as $|\Psi\rangle = e^{-i(\phi+k(L_j+A))} |e\rangle \otimes \Psi_0(x)$ and in the unconfined case as $|\Psi\rangle = e^{-i(\phi+kL_j)} |e\rangle \otimes \Psi_0(x) e^{-ikx}$. The sensitivity to changes in path lengths from gravitational

waves is due to the imprinting of an additional phase kL_j , which is the same whether the atoms are confined or not. In the case of AI sensors one must add additional pulses to become insensitive to the terms associated with the photon recoil.

11.3 Clock-like Detectors

A clock-like gravitational wave detector with confined atoms could be used to detect the gravitational wave phase shift using a basic Ramsey sequence [81], pictured in Fig. 11.2b. A change in the optical path length L_j leads to a modification of the laser phase experienced by atom b of kL_j for the pulses of light that are launched by the laser near atom a . The phase experienced by atom a is unmodified by the change in optical path length. In this sequence, the role of atom a is then simply to record any variation of the phase of the laser itself due to technical sources of noise so that this laser phase noise can be subtracted out from the final measurement.

Stepping through the Ramsey measurement, atoms in both locations are initially prepared in $|g\rangle$. At time t_1 atom a interacts with a pulse of light created by laser a which drives the first $\pi/2$ pulse with phase $\phi = \phi_1$ (we keep track of this phase only to demonstrate insensitivity to its value.) The rotation applied to atom a is $R_a = R_{\pi/2}(\phi_1)$. When the same pulse arrives at atom b it creates a rotation $R_b = R_{\pi/2}(\phi_1 + kL_1)$. At a later time t_2 , a second $\pi/2$ pulse is applied to atom a with a laser phase $\phi = \phi_2$: $R_a = R_{\pi/2}(\phi_2)$. The same pulse of light travels to atom b to drive a $\pi/2$ pulse $R_b = R_{\pi/2}(\phi_2 + kL_2)$.

The final signal that is detected to estimate the gravitational wave's signal is the difference in the probability for finding the atom in its excited and ground states. For atom a this observable can be parameterized as $S_a = P_{ea} - P_{ga} = \cos \Phi_{as}$ and for atom b as $S_b = P_{eb} - P_{gb} = \cos \Phi_{bs}$. For the above Ramsey sequence, the signal phases are given $\Phi_{as} = (\phi_2 - \phi_1)$ and $\Phi_{bs} = kL_2 - kL_1 + (\phi_2 - \phi_1)$. The gravitational wave signal is given by extracting the difference of the phase of these two signals [183]

$$\Phi_s = \Phi_{bs} - \Phi_{as} = k(L_2 - L_1) \quad (11.6)$$

The key result is that the technical laser phase noise is canceled by having been recorded on atom a , and only the gravitational wave's signal remains.

11.4 Atom Interferometer Detectors

In an atom interferometer, the atoms are unconfined and a slightly more complicated sequence is needed. The initial $\pi/2$ pulse entangles internal and external degrees of freedom of the atom by imparting a momentum kick to the portion of the atomic wavefunction transferred to $|e\rangle$. If this momentum kick is not reversed, the portions of the atomic wavefunction will not be spatially overlapped to interfere at the time of the second $\pi/2$ pulse. The simplest solution is to add a π pulse in the middle of the sequence (Fig. 11.2c). More complex sequences, with enhanced performance in different regimes are presented in [82, 84], though the core mechanism for sensitivity is the same as this simple version.

Stepping through the simplest AI sequence, the first $\pi/2$ pulse, π pulse and second $\pi/2$ pulse are launched from the laser at times t_1 , t_2 , and t_3 , respectively. When the three pulses arrive at atom b they will have experienced optical path length differences L_1 , L_2 , and L_3 .

The signal that is extracted is the same as above, and depends only on the optical path lengths as

$$\Phi_s = k(L_3 - 2L_2 + L_1) \quad (11.7)$$

where the random fluctuations in laser phase ϕ_3 , ϕ_2 , and ϕ_1 are again cancelled because they are common to both atoms. A similar analysis for the atom interferometer is presented in [?].

11.5 Enhanced sensitivity analogous to large-momentum-transfer

We now propose a mechanism by which the signal recorded by the clock-like sensor can be enhanced. In atom interferometry in general, and AI based GW detectors in particular [82], large momentum transfer (LMT) is a crucial tool used to enhance the size of measured signals. Our

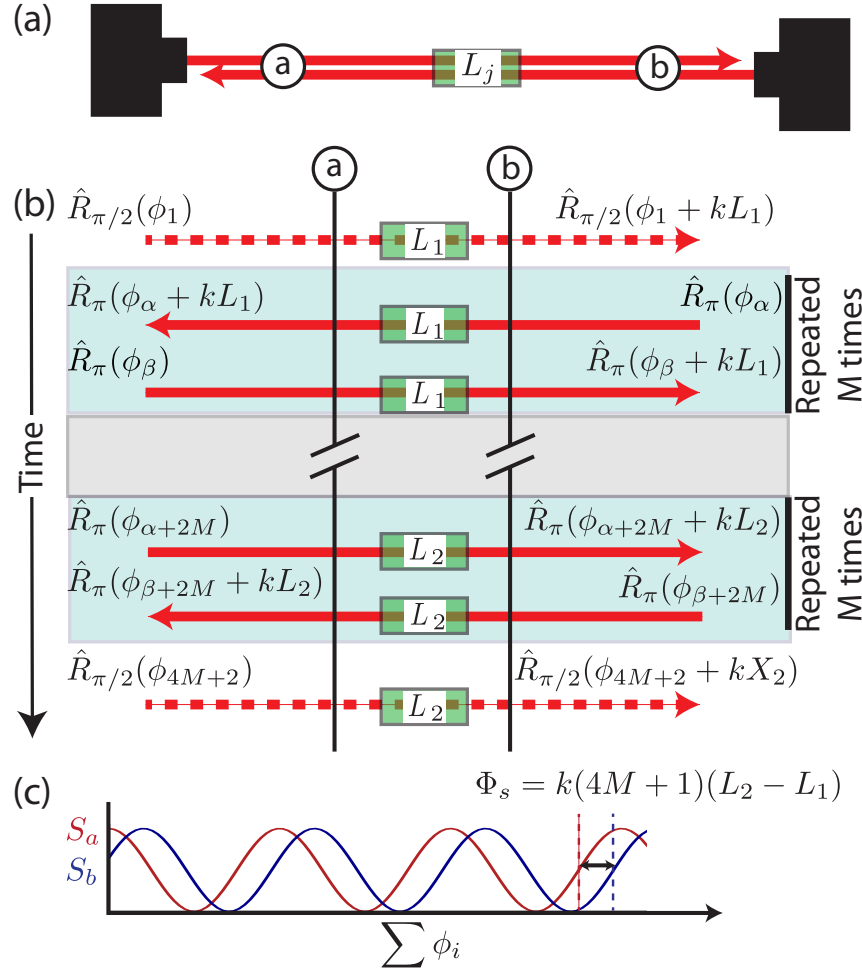


Figure 11.3: LMT-like enhancement with confined atoms. (a) A toy model to demonstrate enhancement mechanism. Atoms are addressed by counter-propagating lasers, both of which pass through the phase modulator that sits between the atoms. (b) Two blocks of M pairs of π pulses (blue boxes) are inserted between the $\pi/2$ pulses of the standard Ramsey sequence. Each pair of π pulses contains a pulse originating from each direction. The duration of each block is assumed here to be short relative to the GW period, while a long evolution time (grey box) between the two blocks may be comparable to the GW period. These additional pulses lead to a factor of $4M + 1$ enhancement in the signal phase Φ_s compared to simple Ramsey sequence.

enhancement mechanism is very similar to LMT, but is applied to confined atoms, so no momentum is actually transferred to the atoms. Instead, by allowing for multiple interactions between the lasers and atoms, a phase is repeatedly written in to the atomic coherence in a constructive manner to enhance the signal size Φ_s .

The key to this enhancement sequence is to apply the laser pulses from alternating sides (as

is proposed for LMT protocols for single-photon transitions with unconfined atoms [82, 182, 84]) using one laser near atom a and one laser near atom b , as illustrated in Fig. 11.3. Because the sign of the laser phase shift imprinted on the atoms is opposite in sign when the atoms are driven from $|e\rangle$ to $|g\rangle$ versus $|g\rangle$ to $|e\rangle$, if the same laser were used to try to imprint its phase multiple times, it would simply unwrite the phase that it had just written in. By interleaving π pulses from a second laser, which does not encode the same gravitational wave phase shift as the first, we can ensure that the phase shift corresponding to the GW is always written in with the same sign. The phases of the laser pulses are always referenced to a plane near the laser from which they originate so that there is negligible GW contribution to this phase.

For simplicity, we treat the case where a set of pulses is applied in quick succession at the beginning of the measurement sequence, with launch times roughly equal to t_1 and with optical path length changes all equal to L_1 . During a subsequent free evolution period T_e , much longer than the time over which the rotations were applied, the optical path length may change. At time $t_2 = t_1 + T_e$, a second set of rotations is quickly applied with launch times all roughly equal to t_2 such that all pulses experience an optical path length change equal to L_2 (a discussion of the effects of time delays is included in the appendix).

Consider the first group of pulses. From the perspective of one of the two atoms, the laser pulse that arrives from the distant laser encodes in its phase kL_1 the change in path length. This phase is then imprinted on the atomic superposition with a different sign between the portion that was transferred from $|e\rangle$ to $|g\rangle$ versus $|g\rangle$ to $|e\rangle$ such that the difference phase is $2kL_1$. The next laser pulse that is launched from the locally situated laser resets the two portions of the atomic superposition to their original internal states. It does imprint a phase shift of its own, but this phase does not encode information about the change in optical path length L_1 , and the laser's phase noise cancels in the final differential signal phase Φ_s . Because the two portions of the atomic superposition have been reset to their original internal states, a subsequent rotation from the distant laser will imprint a phase encoding the same path length difference kL_1 that adds constructively with the previously imprinted phase so that the total imprinted phase on the atom is now $2 \times 2kL_1$.

If we apply M pairs of π rotations after the first $\pi/2$ pulse and M pairs of π rotations before the last $\pi/2$ pulse (i.e. $4M$ total launched π pulses), with each pair containing a rotation originating from both the left and right sides, the differential phase shift between the two output channels is enhanced by a factor of $(4M + 1)$ relative to the simple Ramsey sequence presented above so that now

$$\Phi_s = (4M + 1)k(L_2 - L_1) \quad (11.8)$$

In the regime we consider here, even a single pair of LMT pulses (i.e. 4 total π -pulses), improves the estimate of $(L_2 - L_1)$ by a factor of 25 in variance. As a result, the same precision can be achieved with a reduction in required resources such as atom number or averaging time by a factor of 25. Alternatively, one could achieve the same sensitivity to gravitational waves with a 5 times shorter satellite separation or evolution time T_e for reduced technical complexity or enhanced bandwidth for detecting gravitational waves, respectively.

Previous proposals for detectors using confined [81] and unconfined atoms [84] include so-called dynamical decoupling (DD) sequences. A DD sequence would amount to applying a π pulse from the same laser every time t_j at which the magnitude of the optical path length change $|L_j|$ is maximal. Because the sign of L_j alternates between pulses, the resulting imprinted phase shifts add constructively. This is conceptually similar to our LMT-like enhancement mechanism, except that we can switch the sign of the phase shift by alternating which laser applies the pulse instead of waiting for the sign of the GW to switch.

Dynamical decoupling is useful when the evolution time T_e greatly exceeds the GW period T_g . In this regime the enhancement in signal scales as $T_e/T_g \gg 1$. This reflects the fact that the phase shifts from T_e/T_g cycles of the gravitational wave may be added constructively. There is no constraint on delay time T_d .

LMT-type sequences are useful when the period of the GW T_g and total evolution time T_e are long compared to the pulse transit time T_d . This is necessary in order to send multiple pulses

back and forth between the two satellites before the path length changes caused by the gravitational wave switches sign. The full signal response versus gravitational wave frequency is considered in the appendix. The main result is that one can build a long baseline experiment and set $T_e = T_d$ using a straightforward Ramsey sequence, or one can use LMT-type sequences to dramatically shorten the baseline such that T_e is the same but now $T_e \gg T_d$. In either case, the sensitivities are comparable, but LMT-like enhancement allows one to address technical constraints that might be relaxed by operating with shorter baselines.

Finally, as analyzed in [84], LMT can be combined with DD when both $T_e \gg T_g$ and $T_g \gg T_d$. Doing so simply amounts to varying from which satellite the π pulses are sent.

11.6 Evolution Times Beyond the Excited State Lifetime

We now present a further enhancement that can be realized by treating the atoms as a phase memory, rather than a clock. In a typical clock, the atoms are considered to be a two-state system. Because external degrees of freedom ideally remain unchanged, clocks lack additional quantum labels to specify other states. Any pulse that interacts with one ‘arm’ of the clock interferometer ($|g\rangle$) also interacts with the other ($|e\rangle$): the quantum mechanical amplitudes may be swapped between ground and excited states, but the portion of the wave-function in the ground state cannot be manipulated without also modifying the portion of the wave function in the excited state.

The same is not true of matter wave interferometers, if we relax the earlier assumption that the pulses were short enough to neglect Doppler shifts. When the atoms are unconfined, the momentum recoil that accompanies any change in internal state allows the two arms of the interferometer to be manipulated independently, as the laser frequency can be tuned to the Doppler-shifted resonance of one arm but not the other.

While in clock type proposals, a fraction of the atomic wave-function must always be in the excited state to achieve sensitivity, AI type detection sequences may include useful periods where both arms of the interferometer are in the ground state [82], allowing for evolution times greater than the excited state lifetime, particularly useful for enhancing the signal in the regime where the




Figure 11.4: Use of three-state system to allow atoms to stay in ground state for much of evolution time. (a) Our procedure utilizes two ground states, $|g_1\rangle$ and $|g_2\rangle$ and a single long-lived optically excited state $|e\rangle$. $|g_1\rangle$ and $|g_2\rangle$ can independently be coupled to $|e\rangle$ using laser pulses with different frequency and/or polarization (red and blue arrows). (b) After preparing an atomic superposition of $|g_1\rangle$ and $|e\rangle$ using a $\pi/2$ pulse from the left (atoms begin in $|g_1\rangle$), the portion of the atoms in $|e\rangle$ is transferred to $|g_2\rangle$ using a π pulse from the right. The atoms are then in a superposition of $|g_1\rangle$ and $|g_2\rangle$ with phase dictated by the lasers and the path length L_1 . The sequence of pulses is reversed after a long evolution time, during which time the atoms are not susceptible to excited state decay. The phase shift in the excited state probabilities of the two atoms $\Phi_s = 2k(L_2 - L_1)$ now encodes changes in the optical path length that occurred while the atoms were in the ground states.

period of the GW exceeds the lifetime of the excited state.

In a recoil-free GW detector, additional internal states of the atoms can provide additional quantum labels that allow for the independent manipulation of only one arm of the interferometer. Here, we consider an atomic species like ^{171}Yb with nuclear spin $I = 1/2$. There are two ground states with nuclear spin projections $m_I = \pm 1/2$ that we will label as $|g_1\rangle$ and $|g_2\rangle$. For our purposes, we care about a single excited state $|e\rangle$, as shown in Fig. 11.4a. A small magnetic field is applied to define a quantization axis. Transitions may be driven between either $|g_1\rangle$ and $|e\rangle$ or $|g_2\rangle$ and $|e\rangle$ by applying laser light with different polarizations or frequencies.

An experimental sequence for such a protocol is shown in Fig. 11.4b. A superposition is prepared by driving a $\pi/2$ pulse originating from the left between $|g_1\rangle$ and $|e\rangle$. A π pulse originating from the right then transfers the population of $|e\rangle$ to $|g_2\rangle$. The atoms remain in this configuration while the path length is allowed to change during an evolution time T_e which we take to be much

larger than the delay time T_d . The pulse sequence is then reversed and the change in path length is read out via a population measurement. The resulting signal is $\Phi_s = 2k(L_2 - L_1)$, and again with no contribution from laser phase noise. Memory storage in ground states could be combined with LMT-type enhancement by inserting π pulses on the $|g_1\rangle$ to $|e\rangle$ transition in between the π pulses on the $|g_2\rangle$ to $|e\rangle$ transition and the initial and final $\pi/2$ pulses to realize further enhancement.

In this sequence, the differential phase signal for a given change in path length is twice that of a simple Ramsey sequence, and the laser phase noise associated with all pulses is common mode for both atoms and cancels. In thinking about the signal size it is help to remember that compared to a π pulse applied to the $|g_1\rangle$ to $|e\rangle$ transition, the π pulse applied to the $|g_2\rangle$ to $|e\rangle$ transition contributes only half the relative phase, as it is applied to only one arm of the interferometer.

Because the atoms are in superpositions of ground states, there is no spontaneous emission during T_e . Instead, the relative phase that would normally be exist between $|g_1\rangle$ and $|e\rangle$, is now stored between $|g_1\rangle$ and $|g_2\rangle$ during the evolution period T_e . In addition to extending the evolution time beyond the spontaneous-decay limited coherence time of the atoms, this technique of shelving the atoms in ground states provides insensitivity to collisional effects that may limit the atomic density or coherence time achievable in practice. Inelastic collisions between excited state atoms, which contribute to atom loss in dense systems, would be eliminated by storing the atoms in the ground states [184, 185]. Because elastic collisional properties of alkaline-earth-like atoms are independent of nuclear spin state [114, 186], mean-field shifts due to elastic collisions are common to the two ground states, and will not lead to noise or dephasing. Other bias errors in clocks such as black-body, electrostatic, and lattice polarizability and hyperpolarizability shifts would all also be suppressed.

The key result is that it is not necessary for the phase memory to exist in a frame accumulating phase at an optical frequency. It is only necessary that the interactions with the laser pulses happen between optical transitions, but these interactions can in principle represent a small fraction of the total evolution time.

11.7 Conclusions

Gravitational waves create phase shifts on optical pulses, and the challenge is to detect these phase shifts as sensitively as possible. To detect these changes, at its most basic level, one would like to store an optical pulse of light sent at time t_1 with phase shift kL_1 until a second pulse of light arrives with a different phase shift kL_2 at time t_2 , then compare the phases of the two pulses. By providing a highly coherent phase memory with which the two pulses may interact, the atoms allow one to do just this.

Further, multiple pulses can be made to constructively imprint their phases onto the atoms by properly alternating their launch direction (LMT) or by synchronizing their launch times with the frequency of the gravitational wave one wishes to detect (DD). From a fundamental perspective, the ability to achieve LMT-type signal enhancement appears to have nothing to do with the transfer of momentum to the atoms, or that the two lasers originate from different directions, although for cancellation of technical sources of noise, the latter is required. In certain regimes, LMT allows the signal size to be increased such that the uncertainty in the estimate of the phase of interest $k(L_2 - L_1)$ can be greatly reduced to well below the atomic standard quantum limit $1/\sqrt{N}$ rad.

For gravitational wave detection, we do not need to measure the absolute frequency of a laser relative to an atomic transition frequency, as one would do in a clock. As we have shown, this allows the construction of a ground state shelving protocol for reduced sensitivities to perturbations and evolution times T_e greater than the optical transition lifetime. While in certain situations it is helpful to think of these systems as clocks, in other regimes such a picture is incomplete.

If we view a gravitational wave detector as a clock whose only function is to make frequency measurements, we would expect neither of the approaches presented in sections III and IV to work. Instead, these mechanisms suggest that instead of viewing the atom as a frequency reference, we should think of it as a phase memory. Often these two interpretations are equivalent: frequency is the derivative of phase, so one can infer frequency from two time-separated phase measurements. However, this view does not always give the complete picture.

The key point here is that we are not trying to measure the absolute frequency of a laser relative to an atomic transition, as one would do in a clock. Rather, we are using the atoms as an intermediary to measure the differential phase between various laser pulses (defined by their launch times and the direction from which they are applied).

To illustrate this point, let us imagine that we are trying to use an ensemble of atoms to measure the frequency difference between two lasers (lasers 1 and 2) in an allotted time T . There are two ways one could proceed — one could either use laser 1 to perform a Ramsey sequence of duration T on half of the atoms and laser 2 to perform a Ramsey sequence on the other half. This would represent a standard clock approach, and would independently yield the frequency of each laser, which could then be subtracted.

Alternatively, one could use the atoms to make a differential phase measurement between the two lasers at time $t = 0$, then again at time $t = T$. The difference in these differential phase measurements can then be used to extract the frequency difference between the two lasers. One can do much better with this approach, and not just because each laser is addressing all of the atoms instead of just half.

Instead of performing a simple Ramsey sequence that spans the time T , one could perform a short sequence at $t = 0$ and again at $t = T$ in which π pulses that alternately originate from each of the two lasers are applied between two $\pi/2$ pulses. These π pulses greatly enhance the phase imprinted on the atoms relative to the phase difference between the two lasers (ignoring issues of this phase exceeding 2π). Rather than independently measuring each laser's frequency over the time T , this protocol performs a direct measurement of the differential phase of the lasers at the beginning and end of the time T with enhanced precision, leading to a better estimate of the frequency difference between the two lasers than would be accomplished with independent clocks.

The point here is that by performing separate differential phase measurements of two lasers, we can do much better than we could if we tried to do independent frequency measurements of the lasers. This is exactly what LMT enhancement enables. Rather than making an absolute frequency measurement of a single laser as in a clock, LMT should be thought of as making

differential phase measurements between two lasers — the local laser that has not experienced a path length modification due to a GW, and the laser from the other satellite, which has. The use of two ensembles of atoms and pulses transferred back and forth between two satellites simply allows one to do this measurement in a manner that is insensitive to laser phase noise. From a fundamental perspective, the ability to achieve LMT-type enhancement has nothing to do with the fact that momentum may be transferred to the atoms, or that the two lasers originate from different directions (opposite sign on k). It also has nothing to do with the fact that we are performing a measurement that can be described in the language of a change in position. Fundamentally all that matters is that we are interested in making a differential phase measurement between two lasers. It is the photon shot noise on the lasers, and not the atom projection noise that sets the fundamental limit on such a measurement.

Ultimately, the decision to use confined or unconfined atoms will depend on a myriad of technical considerations that I have no business weighing in on. From a fundamental perspective, both methods have the same mechanism for sensitivity: atom interferometers and clocks both sense changes in phase that result from changes in optical path length between the two satellites. This interpretation should help distinguish to what degree differing expected sensitivities for future proposals are due to the choice of sensor architecture versus the specific parameters considered.

11.8 Accounting for Time Delays

In the main text, we considered an LMT pulse sequence in which the time delay T_d is much shorter than the evolution time T_e and gravitational wave period T_g . Rather than launching a series of LMT pulses at the beginning and at the end, one could in principle launch a continuous series of alternating π -pulses during the entire period of time between the $\pi/2$ pulses. For continuous LMT, we find that the signal size averaged over all phases of the gravitational wave is to good approximation given by:

$$\bar{\Phi}_s = 4\sqrt{2}h \frac{\omega_l |\sin(\omega_g T_d/2)|}{\omega_g^2 T_d} \sin^2(\omega_g T_e/4) \quad (11.9)$$

where ω_l is the laser angular frequency. When $T_d \gg T_g$, the sensitivity falls off rapidly as $1/\omega_g^2$. However, In the limit $T_d \ll T_g$ the scaling changes to

$$\bar{\Phi}_s \approx 2\sqrt{2}h \frac{\omega_l}{\omega_g} \sin^2(\omega_g T_e/4) \quad (11.10)$$

For comparison, if the spacing of the satellites is increased until $T_d = T_e$ then one cannot use LMT and the Ramsey sequence yields an averaged signal:

$$\bar{\Phi}_s = \sqrt{2}h \frac{\omega_l}{\omega_g} \sin^2(\omega_g T_e/2) \quad (11.11)$$

The oscillations in the signal versus ω_g differ by a factor of 2 and the envelope of the signal size is larger by a factor of 2 for LMT. Most importantly, approximately the same signal can be obtained using continuous LMT with of order T_e/T_d π pulses and an approximately a T_e/T_d times shorter satellite satellite spacing. It is likely that decreasing the spacing of the satellites will be advantageous for technical reasons, for instance, reduced requirements on laser power and pointing stability.

Chapter 12

Conclusion and future outlook

In this thesis, I have demonstrated several ways in which narrow-linewidth optical transitions can be used to create useful tools for precision metrology. These tools include a new method of laser cooling, a preliminary demonstration of nondestructive atom counting in a system suitable for an optical lattice clock, and the demonstration of the most precise active atomic frequency reference to date. While encouraging, this work really represents a step towards several long-term goals, which include the application of spin squeezing to an optical lattice clock and the development of a narrow linewidth, steady-state superradiant laser. In this chapter, I will give a perspective on the future of these challenges, and present what I think is an encouraging way to proceed towards each.

12.1 Spin squeezing in an optical lattice clock

The application of spin squeezing to an optical clock has been a long-standing motivation in the development of cavity QED systems for metrology. The ability to create useful spin squeezing using nondestructive measurements with an optical cavity is now well-demonstrated using states separated by microwave frequencies [14, 15]. In ref. [20], I demonstrated that the measurement portion of these techniques can be extended to a system with a narrow-linewidth optical transition suitable for use in an optical lattice clock.

The remaining challenge is to combine this measurement approach with the coherent rotations needed for clock operation. While this has been demonstrated in the microwave domain, these

rotations are far more difficult in the optical domain. The most fundamental challenge here is the difference in wavelengths between optical fields and microwave fields (micron vs centimeter scales). This means that while the Doppler shifts of the microwave transitions may be small compared to typical Rabi frequencies for coherent rotations, they are prohibitively large for the optical transitions (remember, Doppler shifts scale as v/λ , where v is velocity and λ is wavelength). For optical systems, this leads to the requirement that the atoms are tightly confined in the direction along which the clock laser (the one that addresses the narrow-linewidth clock transition) is applied.

In our current system, this presents a problem. The lattice is supported by a cavity that is highly reflective at the clock wavelength of 698 nm. Clock light would have to be applied along the cavity axis, through the cavity mirrors, as this is the only direction in which Doppler shifts are not present. This means that the clock light that interacts with the atoms forms a standing wave with inhomogeneous Rabi frequencies. This is not ideal for clock operation, as any rotation would be accompanied by a shortening of the collective Bloch vector.

In principle, the issue of inhomogeneous coupling could be overcome by preselecting atoms that couple homogeneously to the mode of the cavity at the clock transition wavelength. However, a second challenge is that the cavity should simultaneously have a mode on resonance with the clock transition (in order to couple the clock light to a well-defined mode), and a mode nearly resonant with the stronger transition used for atom counting (in our case at 689 nm). Achieving this would require precise predetermination of the absolute cavity length, which would be challenging.

Another solution that has been considered is to use cavity mirrors that are highly reflective at 689 nm and minimally reflective at 698 nm [187]. This requirement was found to be very difficult with current mirror coating technologies.

I currently think that the best solution to this problem is to confine the atoms in two dimensions, as shown in Fig. 12.1. This means that the atoms can be addressed in a Doppler-free manner by a clock laser incident from any direction within a plane. The two-dimensional confinement could be realized by intersecting the modes of two cavities, each of which supports an optical lattice near the magic wavelength. One of these cavities would have similar parameters at 689 nm to the one




Figure 12.1: A proposed architecture for spin squeezing in an optical lattice clock. Atoms are tightly confined in two dimensions by intersecting lattices supported by two crossed optical cavities. One of these cavities is used to perform atom counting measurements on the 7.5 kHz linewidth 3P_1 transition. Clock rotations are performed by a 698 nm laser applied in the plane defined by the two cavity modes, along which direction the atoms are tightly confined and Doppler-shift free.

used in this thesis, which I have already demonstrated to be sufficient for measurements compatible with spin squeezing, and would be used to perform atom counting.

12.2 Towards a steady-state superradiant laser

While the spectral properties of the pulsed superradiance explored here are encouraging, for a technologically useful device it would be desirable to have a laser that operates in steady state. This requires that inversion is constantly maintained as atoms emit photons. Two approaches may be employed: the atoms that have emitted a photon can be repumped back to the excited state as demonstrated in [63, 24], or fresh atoms already prepared in the excited state can be introduced as atoms that have decayed to the ground state are discarded. Even if the repumping approach is applied, replacement atoms will be required to compensate for atom loss due to heating and collisions with background gas.

The repumping approach would most likely involve exciting the atoms from 1S_0 to 3P_1 with a laser at 689 nm, then to 3S_1 with a laser at 688 nm. From there, the atoms will decay into all of the 3P states. Additional lasers at 707 nm and 679 nm would then be used to depopulate the

$^3\text{P}_2$ states and undesired $^3\text{P}_0$ sublevels. I am currently not optimistic about this approach for ^{87}Sr . The number of photon recoils imparted on the atoms to reach a particular state within $^3\text{P}_0$ is high (several 10's, depending on details), meaning that the atoms would heat out of the lattice within a few cycles. One could attempt to cool during the repump process, but these schemes have their own problems. Raman sideband cooling within $^3\text{P}_2$ would not address the loosely confined radial directions, and other sub-Doppler cooling methods require high intensities that may perturb the lasing transition frequency. We are considering cooling on the 50 kHz linewidth $^3\text{P}_2$ to $^3\text{D}_3$ transition at $2.9\ \mu\text{m}$, but this has yet to be demonstrated. Besides heating issues, the 679 nm laser used to optically pump within $^3\text{P}_0$ would likely lead to large perturbations to $^3\text{P}_0$ that would degrade the stability of the superradiant laser. Because of these issues, and the fact that a continuous stream of new atoms will be required anyway to compensate for losses, I think the first approach for ^{87}Sr should be to figure out how to continuously load fresh atoms into the lattice. If this works well enough, then repumping may not be required.

Steady-state operation will require a fundamentally different mode of experimental operation than used in our current apparatus. Like most cold atom experiments, the current apparatus operates in a time-sequenced manner. Many atomic manipulations occur in the same location, separated by time. In the continuous version, the operations that are currently time-sequenced will have to be performed in spatially separated regions. This leads to challenges, as it is much easier to create a magnetic field or laser beam whose properties vary abruptly in time than in space. Below, I will outline one possible architecture, shown in Fig. 12.2.

As in our current setup, atoms emerge from an oven and are slowed using a Zeeman slower. They are then confined in a 2D blue MOT, perhaps with additional beams present to plug the two ends of the MOT region. The magnetic field for this MOT has vertically oriented minimum. The MOT is operated without repump lasers, so the atoms eventually fall into $^3\text{P}_2$. Some of the sublevels within $^3\text{P}_2$ are magnetically trapped, and will fall under gravity, guided along the magnetic field zero.

At some point, the field gradient that provides confinement must fall off in order to allow for

superradiant laser operation. At this point, the atoms are pumped back to the ground states, and transversely cooled and perhaps compressed using the 7.5 kHz transition. The atoms continue to fall into a region where the optical cavity resides.

In order to load into the optical cavity, dissipation is required. Most dissipation schemes compatible with low enough temperatures for superradiant laser operation would perturb the lasing transition and be incompatible with simultaneous laser operation. To get around this, we divide the cavity mode into a loading region and a lasing region, probably separated by a carefully engineered divider. The cavity has a ring geometry, which allows us to create a moving lattice by driving the beams propagating in opposite directions with slightly different frequencies.

In the loading region, the atoms are cooled into the lattice using cooling light at 689 nm. There may be a magnetic field gradient present here to facilitate spatial compression. 679 nm light in the loading region provide strong dephasing of the clock transition and prevent the atoms from interacting with the superradiant light circulating the cavity. Once trapped and cooled within the lattice, the atoms are transferred to 3P_2 and transported to the lasing region, where 707 nm and 688 nm beams transfer them to 3P_0 . While travelling through the lasing region, each atom emits up to a single photon, before being removed from the participating ensemble by shelving in 3P_2 . The velocity of the moving lattice should be sufficiently high relative to the cavity linewidth that one direction of circulating light is shifted off resonance with the cavity, and the superradiant light propagates in a well-defined direction.

At the time of writing this thesis, we have begun to design and build a new experimental apparatus compatible with such a plan for steady-state operation. I am confident that getting this apparatus to work will be an exciting and unique challenge.




Figure 12.2: A proposed architecture for a continuously loaded superradiant laser. Atoms emitted from oven are slowed by a Zeeman slower, then captured in a 2D MOT formed on the 461 nm transition. Some atoms that fall into 3P_2 are magnetically trapped and fall guided along the magnetic field zero. They are then depumped to the ground state and cooled transversely with 689 nm light before falling into the mode of the optical cavity. Here, they are loaded and cooled into a moving optical lattice, which transports them to the lasing region, where each atom emits up to one photon.

Bibliography

- [1] William B Cairncross, Daniel N Gresh, Matt Grau, Kevin C Cossel, Tanya S Roussy, Yiqi Ni, Yan Zhou, Jun Ye, and Eric A Cornell. A precision measurement of the electron's electric dipole moment using trapped molecular ions. arXiv preprint arXiv:1704.07928, 2017.
- [2] Jacob Baron, Wesley C Campbell, David DeMille, John M Doyle, Gerald Gabrielse, Yulia V Gurevich, Paul W Hess, Nicholas R Hutzler, Emil Kirilov, Ivan Kozyryev, et al. Order of magnitude smaller limit on the electric dipole moment of the electron. Science, page 1248213, 2013.
- [3] D Hanneke, S Fogwell, and G Gabrielse. New measurement of the electron magnetic moment and the fine structure constant. Physical Review Letters, 100(12):120801, 2008.
- [4] Simon Rainville, James K Thompson, Edmund G Myers, John M Brown, Maynard S Dewey, Ernest G Kessler, Richard D Deslattes, Hans G Börner, Michael Jentschel, Paolo Mutti, et al. World year of physics: A direct test of $E=mc^2$. Nature, 438(7071):1096–1097, 2005.
- [5] Savas Dimopoulos, Peter W Graham, Jason M Hogan, and Mark A Kasevich. Testing general relativity with atom interferometry. Physical Review Letters, 98(11):111102, 2007.
- [6] Serge Haroche. Entanglement, decoherence and the quantum/classical boundary. Physics today, 51(7):36–42, 1998.
- [7] Alain Aspect, Jean Dalibard, and Gérard Roger. Experimental test of bell's inequalities using time-varying analyzers. Physical Review Letters, 49(25):1804, 1982.
- [8] Peter W Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. SIAM review, 41(2):303–332, 1999.
- [9] J Ignacio Cirac and Peter Zoller. Goals and opportunities in quantum simulation. Nature Physics, 8(4):264–266, 2012.
- [10] Harold J Metcalf and Peter Straten. Laser cooling and trapping of neutral atoms. Wiley Online Library, 2007.
- [11] Mike H Anderson, Jason R Ensher, Michael R Matthews, Carl E Wieman, Eric A Cornell, et al. Observation of bose-einstein condensation in a dilute atomic vapor. Science, 269(5221):198–201, 1995.

- [12] Kendall B Davis, M-O Mewes, Michael R Andrews, NJ Van Druten, DS Durfee, DM Kurn, and Wolfgang Ketterle. Bose-einstein condensation in a gas of sodium atoms. Physical Review Letters, 75(22):3969, 1995.
- [13] Adam M Kaufman, Brian J Lester, and Cindy A Regal. Cooling a single atom in an optical tweezer to its quantum ground state. Physical Review X, 2(4):041014, 2012.
- [14] Onur Hosten, Nils J Engelsen, Rajiv Krishnakumar, and Mark A Kasevich. Measurement noise 100 times lower than the quantum-projection limit using entangled atoms. Nature, 529(7587):505, 2016.
- [15] Kevin C Cox, Graham P Greve, Joshua M Weiner, and James K Thompson. Deterministic squeezed states with collective measurements and feedback. Physical Review Letters, 116(9):093602, 2016.
- [16] Waseem S Bakr, Jonathon I Gillen, Amy Peng, Simon Fölling, and Markus Greiner. A quantum gas microscope for detecting single atoms in a hubbard-regime optical lattice. Nature, 462(7269):74–77, 2009.
- [17] Ian D. Leroux, Monika H. Schleier-Smith, and Vladan Vuletić. Implementation of cavity squeezing of a collective atomic spin. Phys. Rev. Lett., 104:073602, Feb 2010.
- [18] Markus Greiner, Olaf Mandel, Tilman Esslinger, Theodor W Hänsch, and Immanuel Bloch. Quantum phase transition from a superfluid to a mott insulator in a gas of ultracold atoms. nature, 415(6867):39–44, 2002.
- [19] Kihwan Kim, M-S Chang, Simcha Korenblit, Rajibul Islam, Emily E Edwards, James K Freericks, G-D Lin, L-M Duan, and Christopher Monroe. Quantum simulation of frustrated ising spins with trapped ions. Nature, 465(7298):590–593, 2010.
- [20] Matthew A. Norcia and James K. Thompson. Strong coupling on a forbidden transition in strontium and nondestructive atom counting. Phys. Rev. A, 93:023804, Feb 2016.
- [21] Matthew N Winchester, Matthew A Norcia, Julia RK Cline, and James K Thompson. Magnetically induced optical transparency on a forbidden transition in strontium for cavity-enhanced spectroscopy. Physical Review Letters, 118(26):263601, 2017.
- [22] Matthew A Norcia, Julia RK Cline, John P Bartolotta, Murray J Holland, and James K Thompson. Narrow-line laser cooling by adiabatic transfer. arXiv preprint arXiv:1707.01944, 2017.
- [23] Matthew A Norcia, Julia RK Cline, and James K Thompson. A new perspective on atomic gravitational wave detectors. arXiv preprint arXiv:1707.04571, 2017.
- [24] Matthew A. Norcia and James K. Thompson. Cold-strontium laser in the superradiant crossover regime. Phys. Rev. X, 6:011025, Mar 2016.
- [25] Matthew A Norcia, Matthew N Winchester, Julia RK Cline, and James K Thompson. Superradiance on the millihertz linewidth strontium clock transition. Science advances, 2(10):e1601231, 2016.

- [26] M Schioppo, RC Brown, WF McGrew, N Hinkley, RJ Fasano, K Beloy, TH Yoon, Gianmaria Milani, D Nicolodi, JA Sherman, et al. Ultrastable optical clock with two cold-atom ensembles. Nature Photonics, 2016.
- [27] B. J. Bloom, T. L. Nicholson, J. R. Williams, S. L. Campbell, M. Bishof, X. Zhang, W. Zhang, S. L. Bromley, and J. Ye. An optical lattice clock with accuracy and stability at the 10⁻¹⁸ level. Nature, 506:71, February 2014.
- [28] TL Nicholson, SL Campbell, RB Hutson, GE Marti, BJ Bloom, RL McNally, W Zhang, MD Barrett, MS Safronova, GF Strouse, et al. Systematic evaluation of an atomic clock at 2 x 10⁻¹⁸ total uncertainty. Nature Communications, 6, 2015.
- [29] Ichiro Ushijima, Masao Takamoto, Manoj Das, Takuya Ohkubo, and Hidetoshi Katori. Cryogenic optical lattice clocks. Nat. Photon., 9:185, Feb 2015.
- [30] J. M. Raimond, M. Brune, and S. Haroche. Manipulating quantum entanglement with atoms and photons in a cavity. Rev. Mod. Phys., 73:565–582, Aug 2001.
- [31] A David Boozer, Andreea Boca, Russell Miller, Tracy E Northup, and H Jeffrey Kimble. Reversible state transfer between light and a single trapped atom. Physical Review Letters, 98(19):193601, 2007.
- [32] Andreas Wallraff, David I Schuster, Alexandre Blais, L Frunzio, et al. Strong coupling of a single photon to a superconducting qubit using circuit quantum electrodynamics. Nature, 431(7005):162, 2004.
- [33] J. G. Bohnet, K. C Cox, M. A. Norcia, J. M. Weiner, Z. Chen, and J. K. Thompson. Reduced spin measurement back-action for a phase sensitivity ten times beyond the standard quantum limit. Nat. Photon., 8:731–736, July 2014.
- [34] Kevin C Cox, Matthew A Norcia, Joshua M Weiner, Justin G Bohnet, and James K Thompson. Reducing collective quantum state rotation errors with reversible dephasing. Applied Physics Letters, 105(26):261102, 2014.
- [35] G Dick. Local oscillator induced instabilities in trapped ion frequency standards. Proc. of Precise Time and Time Interval, pages 133–147, 1987.
- [36] DJ Wineland, JJ Bollinger, WM Itano, and FL Moore. Spin squeezing and reduced quantum noise in spectroscopy. Phys. Rev. A, 46(11):R6797–R6800, 1992.
- [37] K. C. Cox, J. M. Weiner, G. P. Greve, and J. K. Thompson. Generating Entanglement between Atomic Spins with Low-Noise Probing of an Optical Cavity. ArXiv e-prints.
- [38] K-J Boller, A Imamoglu, and Stephen E Harris. Observation of electromagnetically induced transparency. Physical Review Letters, 66(20):2593, 1991.
- [39] Martin Mücke, Eden Figueroa, Joerg Bochmann, Carolin Hahn, Karim Murr, Stephan Ritter, Celso J Villas-Boas, and Gerhard Rempe. Electromagnetically induced transparency with single atoms in a cavity. Nature, 465(7299):755–758, 2010.
- [40] Gessler Hernandez, Jiepeng Zhang, and Yifu Zhu. Vacuum rabi splitting and intracavity dark state in a cavity-atom system. Physical Review A, 76(5):053814, 2007.

- [41] Chien Liu, Zachary Dutton, Cyrus H Behroozi, and Lene Vestergaard Hau. Observation of coherent optical information storage in an atomic medium using halted light pulses. Nature, 409(6819):490–493, 2001.
- [42] R Zhao, YO Dudin, SD Jenkins, CJ Campbell, DN Matsukevich, TAB Kennedy, and A Kuzmich. Long-lived quantum memory. Nature Physics, 5(2):100–104, 2009.
- [43] Thibault Peyronel, Ofer Firstenberg, Qi-Yu Liang, Sebastian Hofferberth, Alexey V Gorshkov, Thomas Pohl, Mikhail D Lukin, and Vladan Vuletić. Quantum nonlinear optics with single photons enabled by strongly interacting atoms. Nature, 488(7409):57–60, 2012.
- [44] YO Dudin and A Kuzmich. Strongly interacting rydberg excitations of a cold atomic gas. Science, 336(6083):887–889, 2012.
- [45] Jia Ningyuan, Alexandros Georgakopoulos, Albert Ryou, Nathan Schine, Ariel Sommer, and Jonathan Simon. Observation and characterization of cavity rydberg polaritons. Physical Review A, 93(4):041802, 2016.
- [46] R. H. Dicke. Coherence in spontaneous radiation processes. Phys. Rev., 93(1):99, Jan 1954.
- [47] Michel Gross and Serge Haroche. Superradiance: An essay on the theory of collective spontaneous emission. Physics reports, 93(5):301–396, 1982.
- [48] N Skribanowitz, IP Herman, JC MacGillivray, and MS Feld. Observation of dicke superradiance in optically pumped hf gas. Phys. Rev. Lett., 30(8):309, 1973.
- [49] Y Kaluzny, P Goy, M Gross, JM Raimond, and S Haroche. Observation of self-induced rabi oscillations in two-level atoms excited inside a resonant cavity: the ringing regime of superradiance. Physical Review Letters, 51(13):1175, 1983.
- [50] D. Meiser and M. J. Holland. Steady-state superradiance with alkaline-earth-metal atoms. Phys. Rev. A, 81(3):033847, Mar 2010.
- [51] D. Meiser, Jun Ye, D. R. Carlson, and M. J. Holland. Prospects for a millihertz-linewidth laser. Phys. Rev. Lett., 102(16):163601–163604, Apr 2009.
- [52] DG Matei, T Legero, S Häfner, C Grebing, R Weyrich, W Zhang, L Sonderhouse, JM Robinson, J Ye, F Riehle, et al. 1.5μ m lasers with sub 10 mhz linewidth. arXiv preprint arXiv:1702.04669, 2017.
- [53] T Kessler, C Hagemann, C Grebing, T Legero, U Sterr, F Riehle, MJ Martin, L Chen, and J Ye. A sub-40-mhz-linewidth laser based on a silicon single-crystal optical cavity. Nat. Photon., 6(10):687–692, 2012.
- [54] Christian Hagemann, Christian Grebing, Christian Lisdat, Stephan Falke, Thomas Legero, Uwe Sterr, Fritz Riehle, Michael J Martin, and Jun Ye. Ultrastable laser with average fractional frequency drift rate below $5 \times 10^{-19}/s$. Optics letters, 39(17):5102–5105, 2014.
- [55] N Hinkley, JA Sherman, NB Phillips, M Schioppo, ND Lemke, K Beloy, M Pizzocaro, CW Oates, and AD Ludlow. An atomic clock with 10^{-18} instability. Science, 341(6151):1215–1218, 2013.

- [56] CW Chou, DB Hume, JCJ Koelemeij, DJ Wineland, and T Rosenband. Frequency comparison of two high-accuracy Al^+ optical clocks. Phys. Rev. Lett., 104(7):070802, 2010.
- [57] RWP Drever, John L Hall, FV Kowalski, J. Hough, GM Ford, AJ Munley, and H Ward. Laser phase and frequency stabilization using an optical resonator. Applied Physics B, 31(2):97–105, 1983.
- [58] B. C. Young, F. C. Cruz, W. M. Itano, and J. C. Bergquist. Visible lasers with subhertz linewidths. Phys. Rev. Lett., 82(19):3799–3802, May 1999.
- [59] Kenji Numata, Amy Kemery, and Jordan Camp. Thermal-noise limit in the frequency stabilization of lasers with rigid cavities. Phys. Rev. Lett., 93(25):250602, Dec 2004.
- [60] Mark Notcutt, Long-Sheng Ma, Andrew D Ludlow, Seth M Foreman, Jun Ye, and John L Hall. Contribution of thermal noise to frequency stability of rigid optical cavity via hertz-linewidth lasers. Phys. Rev. A, 73(3):031804, 2006.
- [61] Harold Mark Goldenberg, D Kleppner, and NF Ramsey. Atomic hydrogen maser. Physical Review Letters, 5(8):361, 1960.
- [62] S. A. Diddams, Th. Udem, J. C. Bergquist, E. A. Curtis, R. E. Drullinger, L. Hollberg, W. M. Itano, W. D. Lee, C. W. Oates, K. R. Vogel, and D. J. Wineland. An optical clock based on a single trapped 199Hg^+ ion. Science, 293(5531):825–828, 2001.
- [63] Justin G Bohnet, Zilong Chen, Joshua M Weiner, Dominic Meiser, Murray J Holland, and James K Thompson. A steady-state superradiant laser with less than one intracavity photon. Nature, 484(7392):78–81, 2012.
- [64] Justin G Bohnet, Zilong Chen, Joshua M Weiner, Kevin C Cox, and James K Thompson. Linear-response theory for superradiant lasers. Physical Review A, 89(1):013806, 2014.
- [65] Justin G Bohnet, Zilong Chen, Joshua M Weiner, Kevin C Cox, and James K Thompson. Active and passive sensing of collective atomic coherence in a superradiant laser. Phys. Rev. A, 88(1):013826, 2013.
- [66] J. G. Bohnet, Z. Chen, J. M. Weiner, K. C. Cox, and J. K. Thompson. Relaxation oscillations, stability, and cavity feedback in a superradiant raman laser. Phys. Rev. Lett., 109:253602–253606, Dec 2012.
- [67] Ian D. Leroux, Monika H. Schleier-Smith, and Vladan Vuletić. Orientation-dependent entanglement lifetime in a squeezed atomic clock. Phys. Rev. Lett., 104:250801, Jun 2010.
- [68] Monika H. Schleier-Smith, Ian D. Leroux, and Vladan Vuletić. Squeezing the collective spin of a dilute atomic ensemble by cavity feedback. Phys. Rev. A, 81:021804, Feb 2010.
- [69] Justin G Bohnet, Brian C Sawyer, Joseph W Britton, Michael L Wall, Ana Maria Rey, Michael Foss-Feig, and John J Bollinger. Quantum spin dynamics and entanglement generation with hundreds of trapped ions. Science, 352(6291):1297–1301, 2016.
- [70] Jiazhong Hu, Wenlan Chen, Zachary Vendeiro, Alban Urvoy, Boris Braverman, and Vladan Vuletić. Vacuum spin squeezing. arXiv preprint arXiv:1703.02439, 2017.

- [71] JC Allred, RN Lyman, TW Kornack, and MV Romalis. High-sensitivity atomic magnetometer unaffected by spin-exchange relaxation. Physical Review Letters, 89(13):130801, 2002.
- [72] J. M. McGuirk, H. J. Lewandowski, D. M. Harber, T. Nikuni, J. E. Williams, and E. A. Cornell. Spatial resolution of spin waves in an ultracold gas. Phys. Rev. Lett., 89:090402, Aug 2002.
- [73] C. Deutsch, F. Ramirez-Martinez, C. Lacroûte, F. Reinhard, T. Schneider, J. N. Fuchs, F. Piéchon, F. Laloë, J. Reichel, and P. Rosenbusch. Spin self-rephasing and very long coherence times in a trapped atomic ensemble. Phys. Rev. Lett., 105:020401, 2010.
- [74] X. Du, L. Luo, B. Clancy, and J. E. Thomas. Observation of anomalous spin segregation in a trapped fermi gas. Phys. Rev. Lett., 101:150401, 2008.
- [75] Hidetoshi Katori, Tetsuya Ido, Yoshitomo Isoya, and Makoto Kuwata-Gonokami. Magneto-optical trapping and cooling of strontium atoms down to the photon recoil temperature. Phys. Rev. Lett., 82:1116–1119, Feb 1999.
- [76] Kurt R Vogel, Timothy P Dinneen, Alan Gallagher, and John L Hall. Narrow-line doppler cooling of strontium to the recoil limit. IEEE Transactions on Instrumentation and Measurement, 48(2):618–621, 1999.
- [77] Simon Stellmer, Benjamin Pasquiou, Rudolf Grimm, and Florian Schreck. Laser cooling to quantum degeneracy. Phys. Rev. Lett., 110:263003, Jun 2013.
- [78] P. M. Duarte, R. A. Hart, J. M. Hitchcock, T. A. Corcovilos, T.-L. Yang, A. Reed, and R. G. Hulet. All-optical production of a lithium quantum gas using narrow-line laser cooling. Phys. Rev. A, 84:061406, Dec 2011.
- [79] Tetsuya Ido, Yoshitomo Isoya, and Hidetoshi Katori. Optical-dipole trapping of sr atoms at a high phase-space density. Phys. Rev. A, 61:061403, May 2000.
- [80] Thomas H. Loftus, Tetsuya Ido, Andrew D. Ludlow, Martin M. Boyd, and Jun Ye. Narrow line cooling: Finite photon recoil dynamics. Phys. Rev. Lett., 93:073003, Aug 2004.
- [81] Shimon Kolkowitz, Igor Pikovski, Nicholas Langellier, Mikhail D Lukin, Ronald L Walsworth, and Jun Ye. Gravitational wave detection with optical lattice atomic clocks. Physical Review D, 94(12):124043, 2016.
- [82] Peter W Graham, Jason M Hogan, Mark A Kasevich, and Surjeet Rajendran. New method for gravitational wave detection with atomic sensors. Physical Review Letters, 110(17):171102, 2013.
- [83] Jason M Hogan and Mark A Kasevich. Atom interferometric gravitational wave detection using heterodyne laser links. arXiv preprint arXiv:1501.06797, 2015.
- [84] Peter W Graham, Jason M Hogan, Mark A Kasevich, and Surjeet Rajendran. A resonant mode for gravitational wave detectors based on atom interferometry. arXiv preprint arXiv:1606.01860, 2016.
- [85] Martin M Boyd, Tanya Zelevinsky, Andrew D Ludlow, Sebastian Blatt, Thomas Zanon-Willette, Seth M Foreman, and Jun Ye. Nuclear spin effects in optical lattice clocks. Physical Review A, 76(2):022510, 2007.

- [86] Martin M Boyd. High precision spectroscopy of strontium in an optical lattice: Towards a new standard for frequency and time. PhD thesis, University of Colorado at Boulder, 2007.
- [87] Daniel A Steck. Quantum and atom optics. Oregon Center for Optics and Department of Physics, University of Oregon, page 47, 2007.
- [88] Zilong Chen, Justin G. Bohnet, Joshua M. Weiner, Kevin C. Cox, and James K. Thompson. Cavity-aided nondemolition measurements for atom counting and spin squeezing. Phys. Rev. A, 89:043837, Apr 2014.
- [89] T Holstein and Hl Primakoff. Field dependence of the intrinsic domain magnetization of a ferromagnet. Physical Review, 58(12):1098, 1940.
- [90] JM Raimond, P Goy, M Gross, C Fabre, and S Haroche. Statistics of millimeter-wave photons emitted by a rydberg-atom maser: An experimental study of fluctuations in single-mode superradiance. Physical Review Letters, 49(26):1924, 1982.
- [91] H Jeff Kimble. The quantum internet. Nature, 453(7198):1023–1030, 2008.
- [92] Stephan Ritter, Christian Nölleke, Carolin Hahn, Andreas Reiserer, Andreas Neuzner, Manuel Uphoff, Martin Mücke, Eden Figueroa, Joerg Bochmann, and Gerhard Rempe. An elementary quantum network of single atoms in optical cavities. Nature, 484(7393):195–200, 2012.
- [93] MG Raizen, RJ Thompson, RJ Brecha, HJ Kimble, and HJ Carmichael. Normal-mode splitting and linewidth averaging for two-state atoms in an optical cavity. Physical Review Letters, 63(3):240, 1989.
- [94] James K Thompson, Jonathan Simon, Huanqian Loh, and Vladan Vuletić. A high-brightness source of narrowband, identical-photon pairs. Science, 313(5783):74–77, 2006.
- [95] Subhadeep Gupta, Kevin L Moore, Kater W Murch, and Dan M Stamper-Kurn. Cavity nonlinear optics at low photon numbers from collective atomic motion. Physical Review Letters, 99(21):213601, 2007.
- [96] Ferdinand Brennecke, Stephan Ritter, Tobias Donner, and Tilman Esslinger. Cavity optomechanics with a bose-einstein condensate. Science, 322(5899):235–238, 2008.
- [97] Bjarke TR Christensen, Martin R Henriksen, Stefan A Schäffer, Philip G Westergaard, Jun Ye, Murray Holland, and Jan W Thomsen. Non-linear spectroscopy of sr atoms in an optical cavity for laser stabilization. arXiv preprint arXiv:1507.05757, 2015.
- [98] Philip G. Westergaard, Bjarke T. R. Christensen, David Tieri, Rastin Matin, John Cooper, Murray Holland, Jun Ye, and Jan W. Thomsen. Observation of motion-dependent nonlinear dispersion with narrow-linewidth atoms in an optical cavity. Phys. Rev. Lett., 114:093002, Mar 2015.
- [99] Jonathan Simon, Haruka Tanji, James K Thompson, and Vladan Vuletić. Interfacing collective atomic excitations and single photons. Physical Review Letters, 98(18):183601, 2007.
- [100] Jun Ye, H. J. Kimble, and Hidetoshi Katori. Quantum state engineering and precision metrology using state-insensitive light traps. Science, 320(5884):1734–1738, 2008.

- [101] Hidetoshi Katori, Masao Takamoto, VG PalChikov, and VD Ovsiannikov. Ultrastable optical clock with neutral atoms in an engineered light shift trap. Physical Review Letters, 91(17):173005, 2003.
- [102] Matthew A Norcia and James K Thompson. Simple laser stabilization to the strontium 88sr transition at 707 nm. Review of Scientific Instruments, 87(2):023110, 2016.
- [103] Yosuke Shimada, Yuko Chida, Nozomi Ohtsubo, Takatoshi Aoki, Makoto Takeuchi, Takahiro Kuga, and Yoshio Torii. A simplified 461-nm laser system using blue laser diodes and a hollow cathode lamp for laser cooling of sr. Review of Scientific Instruments, 84(6):063101, 2013.
- [104] Michael Bishof, Xibo Zhang, Michael J Martin, and Jun Ye. Optical spectrum analyzer with quantum-limited noise floor. Physical Review Letters, 111(9):093604, 2013.
- [105] Long-Sheng Ma, Peter Jungner, Jun Ye, and John L Hall. Delivering the same optical frequency at two places: accurate cancellation of phase noise introduced by an optical fiber or other time-varying path. Optics letters, 19(21):1777–1779, 1994.
- [106] A. D. Ludlow, T. Zelevinsky, G. K. Campbell, S. Blatt, M. M. Boyd, M. H. G. de Miranda, M. J. Martin, J. W. Thomsen, S. M. Foreman, Jun Ye, M. Fortier, J. E. Stalnaker, S. A. Diddams, Y. Le Coq, Z. W. Barber, N. Poli, N. D. Lemke, K. M. Beck, and C. W. Oates. Sr lattice clock at 1×10^{-16} fractional uncertainty by remote optical evaluation with a ca clock. Science, 319(5871):1805–8, Mar 2008.
- [107] Albert Ryou and Jonathan Simon. Active cancellation of acoustical resonances with an fpga fir filter. Review of Scientific Instruments, 88(1):013101, 2017.
- [108] Hidetoshi Katori. Optical lattice clocks and quantum metrology. Nat. Photon., 5(4):203–210, March 2011.
- [109] Justin G. Bohnet, Zilong Chen, Joshua M. Weiner, Dominic Meiser, Murray J. Holland, and James K. Thompson. A steady-state superradiant laser with less than one intracavity photon. Nature, 484(6):78–81, 2012.
- [110] Joshua M Weiner, Kevin C Cox, Justin G Bohnet, and James K Thompson. Phase synchronization between two superradiant lasers. arXiv preprint arXiv:1503.06464, 2015.
- [111] Kevin C. Cox, Joshua M. Weiner, and James K. Thompson. Phase diagram for injection locking a superradiant laser. Phys. Rev. A, 90:053845, Nov 2014.
- [112] Minghui Xu, D. A. Tieri, E. C. Fine, James K. Thompson, and M. J. Holland. Synchronization of two ensembles of atoms. Phys. Rev. Lett., 113:154101, Oct 2014.
- [113] R Le Targat, L Lorini, Y Le Coq, M Zawada, J Guéna, M Abgrall, M Gurov, P Rosenbusch, DG Rovera, B Nagórny, et al. Experimental realization of an optical second with strontium lattice clocks. Nature Communications, 4, 2013.
- [114] X. Zhang, M. Bishof, S. L. Bromley, C. V. Kraus, M. S. Safronova, P. Zoller, A. M. Rey, and J. Ye. Spectroscopic observation of su(n)-symmetric interactions in sr orbital magnetism. Science, 345(6203):1467–1473, 2014.
- [115] Arne Bjerhammar. On a relativistic geodesy. J. Geod., 59(3):207–220, 1985.

- [116] A. Loeb and D. Maoz. Using Atomic Clocks to Detect Gravitational Waves. arXiv e-prints: 1501.00996 (2015).
- [117] T. Rosenband, D. B. Hume, P. O. Schmidt, C. W. Chou, A. Brusch, L. Lorini, W. H. Oskay, R. E. Drullinger, T. M. Fortier, J. E. Stalnaker, S. A. Diddams, W. C. Swann, N. R. Newbury, W. M. Itano, D. J. Wineland, and J. C. Bergquist. Frequency ratio of al^+ and hg^+ single-ion optical clocks; metrology at the 17th decimal place. Science, 319(5871):1808–1812, 2008.
- [118] S. Blatt, A. D. Ludlow, G. K. Campbell, J. W. Thomsen, T. Zelevinsky, M. M. Boyd, J. Ye, X. Baillard, M. Fouché, R. Le Targat, A. Brusch, P. Lemonde, M. Takamoto, F.-L. Hong, H. Katori, and V. V. Flambaum. New limits on coupling of fundamental constants to gravity using ^{87}Sr optical lattice clocks. Phys. Rev. Lett., 100(14):140801, Apr 2008.
- [119] P. Maioli, T. Meunier, S. Gleyzes, A. Auffeves, G. Nogues, M. Brune, J. M. Raimond, and S. Haroche. Nondestructive rydberg atom counting with mesoscopic fields in a cavity. Phys. Rev. Lett., 94:113601, Mar 2005.
- [120] J. Appel, P.J. Windpassinger, D. Oblak, U.B. Hoff, N. Kjærgaard, and E.S. Polzik. Mesoscopic atomic entanglement for precision measurements beyond the standard quantum limit. Proc. Natl. Acad. Sci., 106(27):10960–10965, 2009.
- [121] Monika H. Schleier-Smith, Ian D. Leroux, and Vladan Vuletić. States of an ensemble of two-level atoms with reduced quantum uncertainty. Phys. Rev. Lett., 104:073604, Feb 2010.
- [122] Zilong Chen, Justin G. Bohnet, Shannon R. Sankar, Jiayan Dai, and James K. Thompson. Conditional spin squeezing of a large ensemble via the vacuum Rabi splitting. Phys. Rev. Lett., 106:133601, Mar 2011.
- [123] R. J. Sewell, M. Koschorreck, M. Napolitano, B. Dubost, N. Behbood, and M. W. Mitchell. Magnetic sensitivity beyond the projection noise limit by spin squeezing. Phys. Rev. Lett., 109:253605, Dec 2012.
- [124] Georgios Vasilakis, Heng Shen, Kasper Jensen, Misha Balabas, Daniel Salart, Bing Chen, and Eugene Simon Polzik. Generation of a squeezed state of an oscillator by stroboscopic back-action-evading measurement. Nature Physics, 11(5):389–392, 2015.
- [125] Jérôme Lodewyck, Philip G. Westergaard, and Pierre Lemonde. Nondestructive measurement of the transition probability in a Sr optical lattice clock. Phys. Rev. A, 79:061401, Jun 2009.
- [126] Philip Westergaard, Jérôme Lodewyck, and Pierre Lemonde. Minimizing the dick effect in an optical lattice clock. IEEE Trans. Ultr. Ferr. Freq. Cont., 57(3):623–628, March 2010.
- [127] M. J. Martin, D. Meiser, J. W. Thomsen, Jun Ye, and M. J. Holland. Extreme nonlinear response of ultranarrow optical transitions in cavity qed for laser stabilization. Phys. Rev. A, 84:063813, Dec 2011.
- [128] T. Takano, M. Fuyama, R. Namiki, and Y. Takahashi. Spin squeezing of a cold atomic ensemble with the nuclear spin of one-half. Phys. Rev. Lett., 102:033601, Jan 2009.
- [129] C. W. Gardiner, J. Ye, H. C. Nagerl, and H. J. Kimble. Evaluation of heating effects on atoms trapped in an optical trap. Phys. Rev. A, 61:045801, Mar 2000.

- [130] A. Kuzmich, N. P. Bigelow, and L. Mandel. Atomic quantum non-demolition measurements and squeezing. Europhysics Letters, 42(5):481, 1998.
- [131] M. G. Raizen, R. J. Thompson, R. J. Brecha, H. J. Kimble, and H. J. Carmichael. Normal-mode splitting and linewidth averaging for two-state atoms in an optical cavity. Phys. Rev. Lett., 63:240–243, Jul 1989.
- [132] Yifu Zhu, Daniel J. Gauthier, S. E. Morin, Qilin Wu, H. J. Carmichael, and T. W. Mossberg. Vacuum Rabi splitting as a feature of linear-dispersion theory: Analysis and experimental observations. Phys. Rev. Lett., 64:2499–2502, May 1990.
- [133] E.T. Jaynes and F.W. Cummings. Comparison of quantum and semiclassical radiation theories with application to the beam maser. Proceedings of the IEEE, 51(1):89–109, 1963.
- [134] H J Kimble. Strong interactions of single atoms and photons in cavity qed. Physica Scripta, 1998(T76):127, 1998.
- [135] Tetsuya Ido and Hidetoshi Katori. Recoil-free spectroscopy of neutral sr atoms in the lamb-dicke regime. Phys. Rev. Lett., 91:053001, Jul 2003.
- [136] R. Houdré, R. P. Stanley, and M. Ilegems. Vacuum-field Rabi splitting in the presence of inhomogeneous broadening: Resolution of a homogeneous linewidth in an inhomogeneously broadened system. Phys. Rev. A, 53:2711–2715, Apr 1996.
- [137] WM Itano, JC Bergquist, JJ Bollinger, JM Gilligan, DJ Heinzen, FL Moore, MG Raizen, and DJ Wineland. Quantum projection noise: population fluctuations in two-level systems. Phys. Rev. A, 47(5):3554, 1993.
- [138] Richard W Fox, Jeffrey A Sherman, W Douglas, Judith B Olson, Andrew D Ludlow, and Christopher W Oates. A high stability optical frequency reference based on thermal calcium atoms. In Frequency Control Symposium (FCS), 2012 IEEE International, pages 1–3. IEEE, 2012.
- [139] Bjarke T. R. Christensen, Martin R. Henriksen, Stefan A. Schäffer, Philip G. Westergaard, David Tieri, Jun Ye, Murray J. Holland, and Jan W. Thomsen. Nonlinear spectroscopy of sr atoms in an optical cavity for laser stabilization. Phys. Rev. A, 92:053820, Nov 2015.
- [140] Alexey V. Gorshkov, Axel André, Mikhail D. Lukin, and Anders S. Sørensen. Photon storage in Λ -type optically dense atomic media. i. cavity model. Phys. Rev. A, 76:033804, Sep 2007.
- [141] S. J. M. Kuppens, M. P. van Exter, and J. P. Woerdman. Quantum-limited linewidth of a bad-cavity laser. Phys. Rev. Lett., 72:3815–3818, Jun 1994.
- [142] Zhichao Xu, Wei Zhuang, and Jingbiao Chen. Lasing and suppressed cavity-pulling effect of cesium active optical clock. arXiv preprint arXiv:1404.6021, 2014.
- [143] H. M. Goldenberg, D. Kleppner, and N. F. Ramsey. Atomic hydrogen maser. Phys. Rev. Lett., 5:361–362, Oct 1960.
- [144] A.E. Siegman. Lasers. University Science Books, 1986.
- [145] William Guerin, Franck Michaud, and Robin Kaiser. Mechanisms for lasing with cold atoms as the gain medium. Phys. Rev. Lett., 101(9):093002, Aug 2008.

- [146] Geert Vrijsen, Onur Hosten, Jongmin Lee, Simon Bernon, and Mark A. Kasevich. Raman lasing with a cold atom gain medium in a high-finesse optical cavity. Phys. Rev. Lett., 107(6):063904, Aug 2011.
- [147] M. Gross and S. Haroche. Superradiance: An essay on the theory of collective spontaneous emission. Physics Reports, 93(5):301 – 396, 1982.
- [148] Michael J Martin. Quantum metrology and many-body physics: pushing the frontier of the optical lattice clock. PhD thesis, University of Colorado at Boulder, 2013.
- [149] David W Allan. Statistics of atomic frequency standards. Proceedings of the IEEE, 54(2):221–230, 1966.
- [150] Sara L Campbell, RB Hutson, GE Marti, A Goban, N Darkwah Oppong, RL McNally, L Sonderhouse, JM Robinson, W Zhang, BJ Bloom, et al. A fermi-degenerate three-dimensional optical lattice clock. arXiv preprint arXiv:1702.01210, 2017.
- [151] C Shi, J-L Robyr, U Eismann, M Zawada, L Lorini, R Le Targat, and J Lodewyck. Polarizabilities of the sr 87 clock transition. Physical Review A, 92(1):012516, 2015.
- [152] Travis L Nicholson. A new record in atomic clock performance. 2015.
- [153] Jérôme Lodewyck, Michal Zawada, Luca Lorini, Mikhail Gurov, and Pierre Lemonde. Observation and cancellation of a perturbing dc stark shift in strontium optical lattice clocks. IEEE transactions on ultrasonics, ferroelectrics, and frequency control, 59(3), 2012.
- [154] Thomas Middelmann, Stephan Falke, Christian Lisdat, and Uwe Sterr. High accuracy correction of blackbody radiation shift in an optical lattice clock. Physical Review Letters, 109(26):263004, 2012.
- [155] MS Safronova, SG Porsev, UI Safronova, MG Kozlov, and Charles W Clark. Blackbody-radiation shift in the sr optical atomic clock. Physical Review A, 87(1):012509, 2013.
- [156] Masahiro Kitagawa and Masahito Ueda. Squeezed spin states. Phys. Rev. A, 47:5138–5143, Jun 1993.
- [157] Klaus Mølmer and Anders Sørensen. Multiparticle entanglement of hot trapped ions. Phys. Rev. Lett., 82:1835–1838, 1999.
- [158] O. Hosten, R. Krishnakumar, N. J. Engelsens, and M. A. Kasevich. Quantum phase magnification. Science, 352(6293):1552–1555, 2016.
- [159] M. Gärttner, J. G. Bohnet, A. Safavi-Naini, M. L. Wall, J. J. Bollinger, and A. M. Rey. Measuring out-of-time-order correlations and multiple quantum spectra in a trapped-ion quantum magnet. Nature Physics, advance online publication:doi:10.1038/nphys4119, 2017.
- [160] A. M. Rey, L. Jiang, M. Fleischhauer, E. Demler, and M. D. Lukin. Many-body protected entanglement generation in interacting spin systems. Physical Review A, 77, 2008.
- [161] I. Bloch, J. Dalibard, and W. Zwerger. Many-body physics with ultracold gases. Reviews of Modern Physics, 80:885, 2008.

- [162] M. Saffman, T. G. Walker, and K. Mølmer. Quantum information with rydberg atoms. Rev. Mod. Phys., 82:2313–2363, 2010.
- [163] H. Labuhn, D. Barredo, S. Ravets, S. de Lisle, T. Macr, T. Lahaye, and A. Browaeys. Tunable two-dimensional arrays of single rydberg atoms for realizing quantum ising models. Nature, 534:667, 2016.
- [164] J. Zeiher, J. T. Choi, A. Rubio-Abadal, T. Pohl, R. van Bijnen, I. Bloch, and C. Gross. Coherent many-body spin dynamics in a long-range interacting ising chain. arXiv:, 1705.08372, 2017.
- [165] T. Lahaye, C. Menotti, L. Santos, M. Lewenstein, and T. Pfau. The physics of dipolar bosonic quantum gases. Reports on Progress in Physics, 72:126401, 2009.
- [166] A. de Paz, A. Sharma, A. Chotia, E. Marechal, J. H. Huckans, P. Pedri, L. Santos, O. Gorceix, L. Vernac, and B. Laburthe-Tolra. Nonequilibrium quantum magnetism in a dipolar lattice gas. Physical Review Letters, 111:185305, 2013.
- [167] S. A. Moses, J. P. Covey, M. T. Miecnikowski, D. S. Jin, and J. Ye. New frontiers for quantum gases of polar molecules. Nature Physics, 13:13, 2017.
- [168] K. Kim, M. S. Chang, S. Korenblit, R. Islam, E. E. Edwards, J. K. Freericks, G. D. Lin, L. M. Duan, and C. Monroe. Quantum simulation of frustrated ising spins with trapped ions. Nature, 465:590, 2010.
- [169] J. W. Britton, B. C. Sawyer, A. C. Keith, C. C. J. Wang, J. K. Freericks, H. Uys, M. J. Biercuk, and J. J. Bollinger. Engineered two-dimensional ising interactions in a trapped-ion quantum simulator with hundreds of spins. Nature, 484:489, 2012.
- [170] J. T. Barreiro, M. Muller, P. Schindler, D. Nigg, T. Monz, M. Chwalla, M. Hennrich, C. F. Roos, P. Zoller, and R. Blatt. An open-system quantum simulator with trapped ions. Nature, 470:486, 2011.
- [171] C. Lhuillier and F. Laloe. Transport properties in a spin polarized gas, i. J. Phys. France, 43:197, 1982.
- [172] E P Bashkin. Spin waves and quantum collective phenomena in boltzmann gases. Soviet Physics Uspekhi, 29(3):238, 1986.
- [173] B. R. Johnson, J. S. Denker, N. Bigelow, L. P. Lévy, J. H. Freed, and D. M. Lee. Observation of nuclear spin waves in spin-polarized atomic hydrogen gas. Phys. Rev. Lett., 53:302–302, 1984.
- [174] W. J. Gully and W. J. Mullin. Observation of spin rotation effects in polarized ^3He - ^4He mixtures. Phys. Rev. Lett., 52:1810–1813, 1984.
- [175] J. Hu, W. Chen, Z. Vendeiro, A. Urvoy, B. Braverman, and V. Vuletic. Vacuum spin squeezing. arXiv:, 1703.02439, 2017.
- [176] T Binnewies, G Wilpers, U Sterr, F Riehle, J Helmcke, TE Mehlstäubler, EM Rasel, and W Ertmer. Doppler cooling and trapping on forbidden transitions. Physical Review Letters, 87(12):123002, 2001.

- [177] Simon M. Rochester, Konrad Szymański, Mark Raizen, Szymon Pustelny, Marcis Auzinsh, and Dmitry Budker. Efficient polarization of high-angular-momentum systems. Phys. Rev. A, 94:043416, Oct 2016.
- [178] Gabriel N Price, S Travis Bannerman, Kirsten Viering, Edvardas Narevicius, and Mark G Raizen. Single-photon atomic cooling. Physical Review Letters, 100(9):093004, 2008.
- [179] Takashi Mukaiyama, Hidetoshi Katori, Tetsuya Ido, Ying Li, and Makoto Kuwata-Gonokami. Recoil-limited laser cooling of ^{87}Sr atoms near the fermi temperature. Phys. Rev. Lett., 90:113002, Mar 2003.
- [180] Alejandra L Collopy, Matthew T Hummon, Mark Yeo, Bo Yan, and Jun Ye. Prospects for a narrow line mot in yo. New Journal of Physics, 17(5):055008, 2015.
- [181] Mark Yeo, Matthew T Hummon, Alejandra L Collopy, Bo Yan, Boerge Hemmerling, Eunmi Chae, John M Doyle, and Jun Ye. Rotational state microwave mixing for laser cooling of complex diatomic molecules. Physical Review Letters, 114(22):223003, 2015.
- [182] Jason M Hogan and Mark A Kasevich. Atom-interferometric gravitational-wave detection using heterodyne laser links. Physical Review A, 94(3):033632, 2016.
- [183] GT Foster, JB Fixler, JM McGuirk, and MA Kasevich. Method of phase extraction between coupled atom interferometers using ellipse-specific fitting. Optics letters, 27(11):951–953, 2002.
- [184] Ch. Lisdat, J. S. R. Vellore Winfred, T. Middelmann, F. Riehle, and U. Sterr. Collisional losses, decoherence, and frequency shifts in optical lattice clocks with bosons. Phys. Rev. Lett., 103:090801, Aug 2009.
- [185] Michael Bishof, Michael J Martin, Matthew D Swallows, Craig Benko, Yige Lin, Goulven Quéméner, Ana Maria Rey, and Jun Ye. Inelastic collisions and density-dependent excitation suppression in a ^{87}Sr optical lattice clock. Physical Review A, 84(5):052716, 2011.
- [186] Francesco Scazza, Christian Hofrichter, Moritz Höfer, PC De Groot, Immanuel Bloch, and Simon Fölling. Observation of two-orbital spin-exchange interactions with ultracold su (n)-symmetric fermions. Nature Physics, 10(10):779–784, 2014.
- [187] Michael N Bishof. Understanding atomic interactions in an optical lattice clock and using them to study many-body physics. PhD thesis, University of Colorado at Boulder, 2014.