

Squeezing Collective Atomic Spins with an Optical Resonator

by

Ian Daniel Leroux

Bachelor of Applied Science
University of Toronto, 2005

Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Physics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2011

© Massachusetts Institute of Technology 2011. All rights reserved.

Author
Department of Physics
February 28, 2011

Certified by
Vladan Vuletić
Lester Wolfe Associate Professor of Physics
Thesis Supervisor

Accepted by
Krishna Rajagopal
Associate Department Head for Education

Squeezing Collective Atomic Spins with an Optical Resonator

by

Ian Daniel Leroux

Submitted to the Department of Physics
on February 28, 2011, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Physics

Abstract

This thesis describes two methods of overcoming the standard quantum limit of signal-to-noise ratio in atomic precision measurements. In both methods, the interaction between an ultracold atomic ensemble and an optical resonator serves to entangle the atoms and deform the uncertainty distribution of the collective hyperfine spin so that it is narrower in some coordinate than would be possible if the atoms were uncorrelated. The first method uses the dispersive shift of the optical resonator's frequency by the atomic index of refraction to perform a quantum non-demolition measurement of the collective spin, projecting it into a squeezed state conditioned on the measurement outcome. The second method exploits the collective coupling of the atoms to the light field in the resonator to generate an effective interaction that entangles the atoms deterministically. Both methods are demonstrated experimentally, achieving metrologically relevant squeezing of 1.5(5) dB and 4.6(6) dB respectively, and simple analytical models, including the effects of scattering into free space, show that much greater squeezing is realistically achievable. To demonstrate the potential usefulness of such squeezing, a proof-of-principle atomic clock whose Allan variance decreases 2.8(3) three times faster than the standard quantum limit is also presented, together with a discussion of the conditions under which squeezing improves its performance.

Thesis Supervisor: Vladan Vuletić

Title: Lester Wolfe Associate Professor of Physics

Acknowledgments

The work described here was carried out in tight collaboration with Monika Schleier-Smith and Prof. Vladan Vuletić, two of the most enthusiastic and gifted scientists it has ever been my privilege to know. Vladan went far beyond the call of advisory duty, spending hours in the lab most days of the week, teaching electronics, control theory, atomic physics and scientific politics, hashing out quantum measurement theory at the whiteboard or tracking down one more noise source in the laser system, and never taking umbrage as I did my best to contradict everything he said. Monika joined the group at the same time I did and we have been a team ever since. Her patience, insight, and curiosity stand as a perpetual lesson to me. The hours we spent building, debugging, measuring, and thinking together were the best scientific education I can imagine.

The Vuletić group was a remarkably helpful and collaborative environment in which to spend my graduate years. Igor Teper and Yu-ju Lin built the apparatus with which this work was performed, taught us how to use it, and then graciously stood aside as we renovated the result of their labours and made it our own. The single-photon generators, ion trappers, and hollow-core fiber stuffers in the surrounding labs were ever ready to set aside their own problems to help solve mine, whether by sharing equipment, technical experience, a different scientific perspective, or a good story. My thanks in particular to Haruka Tanji, Jon Simon, Marko Cetina and Andrew Grier who were with me throughout the journey, and to Saikat Ghosh for his shining example of steady good nature and wide-ranging scientific curiosity. I would also like to thank Dr. Guy Lamouche and Prof. Joseph H. Thywissen for setting me on the path that led me here and for their continued support and mentorship.

Finally, I am indebted to old friends, to Brian and Katherine Sniffen and their band of merry intellectuals, to Sofia Magkiriadou and especially to my family, André Leroux, Ilea Tant, and Yves Leroux, for supplying me with thoughtful listening ears, a sense of perspective, and a broadening education.

Contents

| | | |
|----------|--|-----------|
| 1 | Introduction | 11 |
| 2 | Spins and Squeezing | 15 |
| 2.1 | A Single Atom: Spin-1/2 | 15 |
| 2.2 | Atomic Ensembles: Spin-N/2 | 17 |
| 2.3 | Spin Squeezing and Figures of Merit | 19 |
| 3 | Spin Squeezing in an Optical Resonator: Theory | 23 |
| 3.1 | Basic Model | 23 |
| 3.2 | Conditional Squeezing by Cavity Assisted Measurement | 26 |
| 3.2.1 | Transmission-based spin measurement | 27 |
| 3.2.2 | Spin readout by optical phase shift | 33 |
| 3.3 | Deterministic Squeezing by Cavity Feedback | 36 |
| 3.3.1 | Ideal Unitary Cavity Feedback | 39 |
| 3.3.2 | Realistic Cavity Feedback Squeezing | 44 |
| 4 | Apparatus | 51 |
| 4.1 | Pertinent properties of ^{87}Rb | 51 |
| 4.2 | Microwave and RF Drive | 54 |
| 4.3 | Lasers | 56 |
| 4.3.1 | Reference Laser | 58 |
| 4.3.2 | Repumper Laser | 59 |
| 4.3.3 | MOT Lasers | 61 |

| | | |
|----------|---|------------|
| 4.3.4 | Trap Laser | 62 |
| 4.3.5 | Probe Laser | 63 |
| 4.3.6 | Beam Switching | 66 |
| 4.4 | Optical Resonator | 70 |
| 4.5 | Photodetector | 76 |
| 4.6 | Dipole Trap | 78 |
| 5 | Experimental Procedures and Apparatus Characterization | 81 |
| 5.1 | Standard Experimental Sequence | 81 |
| 5.2 | Atom Cloud Characterization: Position and Temperature | 83 |
| 5.3 | Optical Pumping | 86 |
| 5.4 | Microwave Rotations | 88 |
| 5.5 | Atomic-State Readout | 90 |
| 5.6 | Measurements of Atomic Coherence | 98 |
| 5.7 | Magnetic-Field Calibration | 101 |
| 6 | Squeezing by Quantum Non-Demolition Measurement | 103 |
| 7 | Squeezing by Cavity Feedback | 109 |
| 8 | Squeezing in an Atomic Clock | 117 |
| 8.1 | Lifetime of Squeezed States | 119 |
| 8.2 | Allan Deviation Measurement | 121 |
| 9 | Conclusion | 125 |

List of Figures

| | | |
|------|---|----|
| 2-1 | Spin-1/2 Bloch Sphere | 16 |
| 2-2 | Projection noise | 18 |
| 3-1 | Basic model | 24 |
| 3-2 | Transmission-based measurement | 27 |
| 3-3 | Measurement squeezing performance | 32 |
| 3-4 | Reflected phase measurement | 33 |
| 3-5 | Comparing transmission- and phase-based measurement squeezing . . | 35 |
| 3-6 | Shearing the CSS | 38 |
| 3-7 | Degrees of freedom of the ideal cavity feedback squeezing model . . . | 39 |
| 3-8 | Continuum modes outside a symmetric optical cavity | 45 |
| 3-9 | Theoretical cavity feedback squeezing | 49 |
| 4-1 | Experiment chamber | 52 |
| 4-2 | Relevant levels of the ^{87}Rb D ₂ line | 53 |
| 4-3 | Microwave generation chain | 55 |
| 4-4 | Horn-side optical pumping beam | 57 |
| 4-5 | Camera-side optical pumping beam | 58 |
| 4-6 | Reference Laser | 59 |
| 4-7 | Repumper laser | 60 |
| 4-8 | MOT laser battery | 60 |
| 4-9 | Trap laser | 62 |
| 4-10 | Probe Laser | 63 |
| 4-11 | Probe Laser Sidebands | 64 |

| | | |
|------|---|-----|
| 4-12 | Cavity Input Optics | 65 |
| 4-13 | Laser switchyard | 67 |
| 4-14 | Shutter Driver Circuits | 68 |
| 4-15 | Optical resonator mounted on atom chip | 70 |
| 4-16 | Cavity output filtering optics | 76 |
| 5-1 | Radial Compression by the Dipole Trap | 84 |
| 5-2 | Longitudinal Position of the Atoms | 84 |
| 5-3 | Temperature Measurement | 85 |
| 5-4 | Typical Rabi Spectrum | 88 |
| 5-5 | Decay of Rabi Nutations | 89 |
| 5-6 | Typical Transmission Trace | 91 |
| 5-7 | Readout Noise vs. Photon Number | 97 |
| 5-8 | Ramsey decay | 98 |
| 5-9 | Fixed-Time Ramsey Signal | 99 |
| 5-10 | Contrast Measurement by Rabi Nutation | 100 |
| 5-11 | Typical Magnetic Field Calibration Curve | 102 |
| 6-1 | Projection Noise With and Without Measurement Suppression | 104 |
| 6-2 | Measurement Squeezing | 107 |
| 7-1 | Transverse Variance vs. Angle | 111 |
| 7-2 | Modeling Observed Cavity Feedback Squeezing | 113 |
| 7-3 | Metrological Squeezing by Cavity Feedback | 114 |
| 8-1 | Lifetime of a squeezed spin state | 120 |
| 8-2 | Allan deviation of a squeezed clock | 123 |

Chapter 1

Introduction

The tale of atomic physics is one of a quest for ever more detailed control of simple quantum systems. Starting with the ability to sort atoms by their magnetic quantum number in a Stern-Gerlach experiment [1], atomic physicists have learned to modify atoms' internal state with resonant electromagnetic pulses [2] and purify it with optical pumping [3], and then to control the external or motional state of the atom with laser cooling [4] techniques, so that the mean values of all atomic degrees of freedom are now under experimental control. A logical next step in this quest for control is to learn to manipulate higher moments of the atomic variables than the mean, starting with the variance.

Among the first to seriously consider manipulating the variances of internal atomic degrees of freedom at the quantum level were Kitagawa and Ueda who, in a seminal paper [5], laid out the essential properties of squeezed spin states. These are states of an ensemble of atoms for which the variance of some internal observable is smaller than the natural scale set by the independent quantum fluctuations of the atoms comprising the ensemble. Around the same time, Wineland *et al.* pointed out that such states had important practical applications for precision measurement, since such independent quantum fluctuations limit the performance of even technically perfect atomic precision measurements [6, 7]. Indeed, since 1999 there have been atomic clocks good enough that their short-term statistical uncertainty is bounded by this so-called standard quantum limit (SQL) [8]. To overcome this limit, the variance

of the atomic observable detected in the measurement must be reduced further than is possible if the atoms are statistically independent; in other words it is necessary to introduce correlations between the outcomes of measurements on different atoms or, in quantum terms, to entangle the atoms. Research into spin squeezing has thus been driven by dual motivations: from a fundamental perspective squeezing offers a doorway into the still nebulous world of many-particle entanglement that can be opened using only the experimentally-accessible concept of measurement variance; while from a practical standpoint squeezing holds out the hope of improved performance for atomic precision measurements, which are of such fundamental importance that they underlie the definitions of all SI units but the kilogram.

Unfortunately it was not immediately obvious how squeezed states might be produced in practice, and early experimental progress was slow. Kitagawa and Ueda postulated a “twisting” Hamiltonian which would deform a fiducial initial state’s uncertainty region, reducing its variance [5], but at the time no method for implementing this Hamiltonian in an experimentally convenient system was known [7]. The first demonstration of spin squeezing used absorption of squeezed light in an atomic vapor to transfer some of the light’s squeezing to the atomic ensemble [9], but this method was soon abandoned because delivering highly squeezed light to the atoms is difficult: massless photons are easily lost in transmission, and losses are detrimental to squeezing. The exquisite control of few-particle systems achieved in ion traps has allowed the direct synthesis of near-perfect squeezed states of a few ions by the use of entangling gates [10], though it is not clear how such manipulations might be extended to squeeze large ensembles where the potential benefits of correlating atoms are the greatest. More recently, the control of dense atomic ensembles has improved to the point that the twisting Hamiltonian can be implemented using the internal-state-dependent collisional interactions among atoms in a Bose-Einstein Condensate [11].

The work presented in this thesis¹ attacks the problem of spin squeezing using the tools of cavity quantum electrodynamics, by placing the atoms within an opti-

¹Significant portions of this work were reported in four previous publications, references [12–15].

cal resonator which isolates particular modes of the electromagnetic spectrum and strengthens their interaction with the atoms. These tools allow spin squeezing to be carried out in an ideal context for precision measurements. The atomic ensembles to be squeezed are dilute, with no direct collisional interactions or associated density-dependent energy shifts. In fact the only interactions between atoms in the ensemble are those introduced by the resonator light field itself, which can be switched on and off at the experimenter’s discretion. The techniques require only that the atomic states to be manipulated have different optical resonance frequencies, so that they may be applied to well-protected internal degrees of freedom that couple only very weakly to external perturbing fields. This work, for instance, uses the standard hyperfine “clock” transition for its atomic degree of freedom, the same one used in state-of-the-art fountain clocks [16]. Finally, the number of atoms in the ensemble can be large: the proof-of-principle demonstrations in this thesis involve around $10^4 \dots 10^5$ atoms, and more could be added subject only to the finite volume of the optical resonator.

Two methods of optical-resonator-assisted spin squeezing are demonstrated here. The first, termed measurement-based squeezing, is a new implementation of a proposal first put forward by Kuzmich *et al.* which produces squeezed states conditionally by precisely measuring the variable of interest and making a note of the measurement result [17]. The mean and variance of the observed quantity are then given by the result and the uncertainty of the measurement. Though the measurement result cannot be predicted and will in general fluctuate from realization to realization of the experiment, the measurement uncertainty can in principle be made small. In practice, scattering of photons into free space introduces noise processes that limit the achievable measurement uncertainty; the optical resonator introduced in this work mitigates these effects by strengthening the controlled atom-light interaction relative to these undesirable scattering processes. The second squeezing technique to be described, christened cavity feedback squeezing, is a novel implementation of the twisting Hamiltonian of Kitagawa and Ueda using an effective interaction between the atoms mediated by the light field circulating in the resonator. The method is deterministic, in that it reliably produces a specific squeezed state given a certain ex-

perimental procedure, without the need for monitoring by a measurement apparatus. As of this writing, the 4.6(6) dB of directly observable spin squeezing produced by this method is, to the best of the author's knowledge, the strongest such squeezing yet achieved.

Finally, as a first study of the practical applicability of squeezed states to precision measurement, atomic ensembles squeezed by cavity feedback squeezing are used to operate an atomic clock. The clock is used to study the vulnerability of the squeezed state to various noise processes in the lab, demonstrating certain parameter regimes where they are capable of improving metrological performance. Measurements of the clock's Allan deviation show that it is, along with the one reported in reference [18], the first atomic clock to operate beyond the SQL, by a factor of 2.8(3).

The remainder of the thesis is structured as follows. Chapter 2 gives an abstract presentation of the pseudospin variables used here in discussing the atoms' internal degrees of freedom and defines precisely what is meant by spin squeezing. Chapter 3 presents a theoretical analysis for an idealized atom of both the measurement-based and cavity-feedback-based squeezing schemes, including the fundamental limits to their performance due to scattering of light into free space. Chapter 4 describes the experimental apparatus used in this work, and in particular the optical resonator which is its essential tool, and chapter 5 narrates the basic experimental techniques for preparing and characterizing atomic ensembles and manipulating their internal degrees of freedom. Chapters 6, 7 and 8 present, analyze and discuss the experimental data obtained in tests of the two squeezing methods and the atomic clock studies, respectively. Finally, chapter 9 concludes with some suggested directions for further studies building on the work described here.

Chapter 2

Spins and Squeezing

This chapter briefly reviews the quantum mechanics of two-level atoms, first with a single atom, then with an ensemble of such atoms. The first two sections establish useful notation, introduce the class of highly-symmetric states most commonly used in the theoretical description of precision metrology experiments, and identify the fundamental quantum fluctuations that afflict these states. The final section considers the squeezed states where such fluctuations are suppressed, and in particular the problem of correctly defining a figure of merit or criterion for squeezing. The ideas presented here are mostly to be found in Kitagawa and Ueda's seminal paper on squeezed spin states [5] and the analysis by Wineland *et al.* of their use in metrology [6, 7].

2.1 A Single Atom: Spin-1/2

The simplest non-trivial quantum system is a two-level atom or qubit. The atom has two states, which we shall label $|\uparrow\rangle$ and $|\downarrow\rangle$. A completely general pure quantum state of this system can be written as $|\theta, \phi\rangle = \cos(\theta/2)e^{-i\phi/2}|\uparrow\rangle + e^{i\phi/2}\sin(\theta/2)|\downarrow\rangle$. If we




Figure 2-1: Bloch sphere for a single pseudospin-1/2. The \hat{z} component of the Bloch vector is one-half the population difference between the states $|\uparrow\rangle$ and $|\downarrow\rangle$, while its phase angle is the quantum phase between them.

introduce the pseudospin component operators

$$\begin{aligned}
 s_x &= (|\uparrow\rangle\langle\downarrow| + |\downarrow\rangle\langle\uparrow|)/2 \\
 s_y &= (|\uparrow\rangle\langle\downarrow| - |\downarrow\rangle\langle\uparrow|)/2i \\
 s_z &= (|\uparrow\rangle\langle\uparrow| - |\downarrow\rangle\langle\downarrow|)/2,
 \end{aligned}
 \tag{2.1}$$

which obey the usual angular momentum commutation relations, then we find

$$\langle s_x \rangle = \frac{1}{2} \cos \phi \sin \theta \tag{2.2}$$

$$\langle s_y \rangle = \frac{1}{2} \sin \phi \sin \theta \tag{2.3}$$

$$\langle s_z \rangle = \frac{1}{2} \cos \theta, \tag{2.4}$$

as expected for a vector of length 1/2 whose orientation is given by the polar angle θ and azimuthal angle ϕ . This equivalence between two-level systems in quantum mechanics and coordinates on a sphere runs very deep [19], and enables the use of geometrical reasoning to understand their evolution.

Note that once the orientation of the spin vector is given, the quantum state is fully specified. This has two important consequences, the first being that the state

of a single two-level atom can only evolve through rotations of its pseudospin vector. It follows that the signal in any measurement performed on a single two-level atom can be interpreted as a deflection of the pseudospin. The second consequence of the completeness of the mean spin as a specification of the two-level atom's state is that all variances of spin operators are determined as soon as the expectation value of the spin vector $\langle \mathbf{s} \rangle$ is given. This leads to a straightforward fundamental limit on the precision of any measurement performed on a single two-level system: any attempt to measure small deflections of the spin vector are limited by the unmodifiable variance $\langle \Delta s_{\perp}^2 \rangle = 1/4$ of all spin components perpendicular to the mean spin $\langle \mathbf{s} \rangle$. This is the origin of the saying that “one cannot squeeze spin-1/2”.

2.2 Atomic Ensembles: Spin-N/2

If, instead of a single two-level atom, our system contains an ensemble of N such atoms, we can define collective pseudospin operators by adding up the individual spin-1/2 vectors of each atom in the ensemble: $S_x = \sum_{i=1}^N s_{x,i}$ and similarly for S_y and S_z . The resulting collective Bloch vector \mathbf{S} is no longer a complete description of the state of the ensemble: the Hilbert space for N qubits has 2^N dimensions, whereas specifying the expected Bloch vector $\langle \mathbf{S} \rangle$ only fixes three real parameters. However, the vast majority of the states in this huge 2^N -dimensional Hilbert space have very short Bloch vectors $|\langle \mathbf{S} \rangle| \ll S_0 = N/2$ [20], giving them poor sensitivity to rotations. The states of interest are those where the individual $\langle \mathbf{s}_i \rangle$ point in nearly the same direction, states where $|\langle \mathbf{S} \rangle|$ approaches its maximum value of S_0 and which are symmetric under permutations of the atoms in the ensemble [20].

Consider, in particular, the coherent spin state [21] along $\hat{\mathbf{x}}$

$$|\text{CSS}\hat{\mathbf{x}}\rangle = \left(\frac{|\uparrow\rangle + |\downarrow\rangle}{\sqrt{2}} \right)^{\otimes N}, \quad (2.5)$$

which plays the role of a canonical initial state in this thesis. It can be prepared by initializing all atoms in the ensemble to the $|\downarrow\rangle$ level and then rotating them into an




Figure 2-2: Projection noise: small deflections of a macroscopic spin S (top row) are detected by measuring a transverse component S_{\perp} . The measurement of the transverse component projects each constituent spin-1/2 to $\pm 1/2$ with equal probability (middle row), so that their sum yields a fluctuating total transverse component ΔS_{\perp} (bottom row) and a corresponding uncertainty $\Delta\phi$ on the deflection of the spin.

equal superposition of $|\uparrow\rangle$ and $|\downarrow\rangle$ with a $\pi/2$ pulse. It is manifestly a product state, and has the maximum spin length $|\langle \mathbf{S} \rangle| = \langle S_x \rangle = S_0$. Its transverse spin variances can be derived formally from the commutation relations for angular momentum and the Heisenberg uncertainty relation, or by noting that when the transverse spin component is measured, each individual atom is projected with equal probability onto a transverse component of either $+1/2$ or $-1/2$ because of the quantization of spin. The total transverse component, the difference between the number of spins which are projected each way, is then binomially distributed with variance $\langle \Delta S_y^2 \rangle = \langle \Delta S_z^2 \rangle = S_0/2$. This fixes the radius of the uncertainty region at the tip of the Bloch vector, which is circular since the same argument can be made for any transverse spin component.

If the spins all rotate through some small angle $\delta\phi$, so will the collective Bloch vector, and the mean squared error in measuring such small rotations is $\langle \Delta\phi^2 \rangle = \langle \Delta S_y^2 \rangle / \langle S_x \rangle^2 = (2S_0)^{-1} = N^{-1}$. This is the best rotation sensitivity achievable with independent atoms, as can be seen by noting that it is simply the best possible single-atom rotation sensitivity (from the previous section) improved by the usual N^{-1} factor

for averaging N independent measurements, one on each atom. This sensitivity limit is referred to as the standard quantum limit (SQL), and can only be overcome by entangling the atoms so that the outcomes of measurements on the individual atoms are no longer independent.

Note, in passing, that for large ensembles with $N \gg 1$, the standard deviation $\sqrt{S_0/2}$ in the transverse components of the coherent state is much larger than the discrete step size (of 1 unit of angular momentum) imposed by the quantization of spin. Thus, for states not too far from coherent states of large ensembles, the transverse spin components can be treated as effectively continuous variables and can, indeed, be rescaled to form a canonically conjugate pair:

$$[S_y, S_z] = iS_x, \quad \text{where } \langle S_x \rangle \approx S_0 \gg 1 \quad (2.6)$$

$$[\frac{S_y}{\sqrt{S_0}}, \frac{S_z}{\sqrt{S_0}}] = [X, P] = i. \quad (2.7)$$

This so-called Holstein-Primakoff approximation [22] is frequently used in the spin squeezing literature to connect results for the Bloch sphere of a collective spin vector to those for the usual phase plane [23], and thus to draw on the accumulated wisdom of the light squeezing community.

2.3 Spin Squeezing and Figures of Merit

While the fluctuations in any single atom's transverse spin components cannot be reduced, they can be correlated with those of other atoms so as to partially cancel each other. Thus, while $\langle \Delta \phi^2 \rangle = (2S_0)^{-1}$ is the best rotation sensitivity achievable for N independent atoms, it can be improved upon by entangling the atoms. Wineland *et al.* [6] introduced a squeezing criterion

$$\zeta = \frac{\langle \Delta S_{\perp}^2 \rangle}{|\langle \mathbf{S} \rangle|^2} \cdot 2S_0 \quad (2.8)$$

$$= \frac{2\langle \Delta S_z^2 \rangle}{S_0} \cdot \frac{1}{C^2} \quad (2.9)$$

which compares the sensitivity with which the given state may be used to detect rotations in the direction of S_\perp to the best possible rotation sensitivity for unentangled atoms, as achieved by the coherent spin state. The second line expresses this criterion directly in terms of experimentally accessible quantities for a state whose mean spin is oriented along the $\hat{\mathbf{x}}$ axis: the fluctuations of the population difference $2S_z$ between $|\uparrow\rangle$ and $|\downarrow\rangle$, and the contrast or fringe visibility $C = \langle S_x \rangle / S_0$ of a Rabi or Ramsey oscillation starting from this state. This gives a clear interpretation of the squeezing criterion as an improvement in signal to noise ratio: the first factor measures the suppression of measurement noise relative to the unentangled $|\text{CSS}\hat{\mathbf{x}}\rangle$, the second factor quantifies the loss of signal relative to that same state. $\zeta < 1$ implies that the signal to noise ratio of a rotation measurement is better than it could possibly be in the absence of quantum correlations, and therefore implies entanglement. References [12, 14, 15] use a slight modification of the Wineland squeezing criterion, multiplying it by an additional factor of the initial contrast C_0 . Ideally, $C_0 = 1$ and this makes no difference, but in practice background decoherence processes can reduce the signal contrast even before the squeezing procedure is applied, in which case equation 2.9 overestimates the loss of signal due to the squeezing. While this modified criterion preserves the practical meaning of ζ^{-1} as an increase in signal-to-noise ratio relative to the unsqueezed state with projection noise and the initially available signal contrast C_0 , there exist in principle pathological scenarios where $C_0\zeta < 1$ but the final state is not entangled [24]. This thesis therefore adopts the strict Wineland criterion of equation 2.9 where $\zeta < 1$ guarantees both entanglement and improved rotation sensitivity.

There are two other inequivalent definitions of squeezing common in the literature. One definition compares the fluctuations of the transverse spin components to the minimum for unentangled atoms given the same mean spin vector length

$$\zeta_e = \frac{2\langle \Delta S_\perp^2 \rangle}{|\langle \mathbf{S} \rangle|} \quad (2.10)$$

and calls a state squeezed if $\zeta_e < 1$. This definition is interesting from a fundamental

perspective, since $\zeta_e < 1$ is sufficient to prove that some pair of atoms in the ensemble must be entangled [25]. It is not, however, the relevant figure of merit for metrology, because a state can be squeezed by this definition and still yield worse sensitivity than the unentangled $|\text{CSS}\hat{\mathbf{x}}\rangle$. The essential point is that a state with $|\langle\mathbf{S}\rangle| < S_0$ yields a smaller signal (i.e. contrast $C < 1$) than $|\text{CSS}\hat{\mathbf{x}}\rangle$, and the noise must be suppressed sufficiently to overcome this handicap before one can see a real improvement in signal-to-noise ratio.

The other alternative figure of merit is the quotient obtained by dividing the observed transverse spin noise to the expected projection noise for the ensemble

$$q = \frac{2\langle\Delta S_{\perp}^2\rangle}{S_0}. \quad (2.11)$$

This is a convenient parametrization of the spin noise suppression, but the condition $q < 1$, sometimes referred to as “number squeezing” when the transverse spin component is S_z (i.e. number difference between $|\uparrow\rangle$ and $|\downarrow\rangle$), does not imply entanglement. It is easily satisfied by a product state, for instance one with one atom each in $|\uparrow\rangle$ and $|\downarrow\rangle$ —with no uncertainty along $\hat{\mathbf{z}}$ —and with the remaining atoms aligned along $\hat{\mathbf{x}}$ so that $\hat{\mathbf{z}}$ is transverse to the mean spin. Nor does it imply an improvement in measurement performance: q quantifies the suppression of noise, but does not account for the concomitant loss of signal contrast. That some loss of signal is inevitable is easily seen by noting that $|\text{CSS}\hat{\mathbf{x}}\rangle$ is the unique state for which $\langle S_x \rangle = S_0$ or, more generally, that the coherent spin states are the only states for which $|\langle\mathbf{S}\rangle| = S_0$. Any other state, including the “number squeezed” states, must have a shorter mean spin. In the absence of information about the mean spin length, properly quantified by the parameters ζ and ζ_e introduced above, $q < 1$ does not imply any interesting non-classical properties of the state for which it holds, and it is unfortunate that this misleading definition of squeezing persists in the literature.

Finally, it is worth emphasizing that the squeezed states are but a tiny subset of all the possible entangled states of N two-level atoms. Squeezing is not even a necessary condition for metrologically useful entanglement. The GHZ states $|\uparrow\rangle^{\otimes N} + e^{iN\phi}|\downarrow\rangle^{\otimes N}$,

used in some of the first demonstrations of entanglement-enhanced atomic phase estimation with trapped ions [10], have a mean spin length of zero and none of their three spin components has fluctuations below the projection noise for independent atoms. They yield enhanced measurement sensitivity not by reducing noise, but rather by an amplified signal: the phase between the all-up and all-down states evolves at N times the single-atom precession rate, and a suitable measurement, of parity rather than of a Cartesian spin component, can exploit this N -fold increase. GHZ states, however, suffer from several practical disadvantages for metrological applications. Though GHZ states of a few ions can be prepared by the application of pairwise entangling gates, it is not clear how to prepare them in large ensembles of atoms where there is the most to gain from overcoming the SQL. They are also very sensitive to decoherence: a single photon scattered into free space can completely destroy a GHZ state, whereas we shall see in chapter 8 that squeezed states are remarkably robust. Finally, the use of GHZ states requires a readout of the parity of the ensemble, which is much more difficult to measure for large ensembles than the simple population difference between $|\uparrow\rangle$ and $|\downarrow\rangle$ used as the readout signal in standard or squeezed quantum measurements.

Chapter 3

Spin Squeezing in an Optical Resonator: Theory

This chapter presents the theory of atom-cavity interactions useful for spin squeezing in an idealized setting. While later chapters will introduce corrections and complications necessary to describe specific experimental implementations, the goal for now is to understand the basic principles, scaling laws, and fundamental limitations that govern the use of the collective atom-cavity coupling for squeezing. Section 3.1 presents the simplified model of three-level atoms and optical resonators used in this analysis. Section 3.2 discusses the use of the optical cavity to measure a collective pseudospin component and thereby project the atoms into a squeezed state, with the squeezing conditioned on the measurement outcome. Section 3.3 analyzes the effective interaction between atoms introduced by the cavity, which serves to produce a squeezed state deterministically.

3.1 Basic Model

The model considered here consists of an ensemble of N three-level atoms, each comprising a pair of long-lived levels $|\uparrow\rangle$ and $|\downarrow\rangle$ with energies $\pm\hbar\omega_a/2$ respectively and an excited state $|e\rangle$ with energy $\hbar\omega_c$. The two optical transitions $|\uparrow\rangle \leftrightarrow |e\rangle$ and $|\downarrow\rangle \leftrightarrow |e\rangle$ are taken to have equal oscillator strength; that is, $|e\rangle$ decays to each of the




Figure 3-1: Basic model: ideal three-level atoms (left) uniformly coupled to an optical resonator (right)

two long-lived states with equal rates $\Gamma/2$, so that its linewidth is Γ .

The atoms are uniformly coupled, with the same Rabi frequency $2g$ for each optical transition, to an optical cavity of linewidth κ and resonant frequency ω_c . It is convenient, though not essential, to choose the cavity resonance frequency halfway between the two atomic transition frequencies, so that the detuning has the same magnitude $\Delta = \omega_a/2$ for both transitions. The detuning is taken to be large enough to keep the atom-cavity interactions in the linear, dispersive regime: $\Delta \gg \kappa, \Gamma, \sqrt{N}g$. The intracavity photon number $\langle c^\dagger c \rangle$ is assumed to be low enough to avoid saturating the atomic transitions: $\langle c^\dagger c \rangle (g/\Delta)^2 \ll 1$.

Under these conditions the far off-resonant coupling between photons and atoms in the cavity produces three effects, all of which can be calculated by textbook methods [26]. First, there is a dispersive shift of the cavity resonance frequency due to the index of refraction of the atomic gas

$$\delta\omega_c = 2S_z \frac{g^2}{\Delta}, \quad (3.1)$$

which is just the usual off-resonant level shift g^2/Δ multiplied by the difference $2S_z$

between the population of atoms in $|\uparrow\rangle$ which shift the cavity to the blue and the population of atoms in $|\downarrow\rangle$ which shift the cavity to the red. Generally there is an additional uninteresting cavity shift proportional to the total atom number N , but it vanishes for the symmetric choice of cavity detuning used here. This simplifies the theory and is also valuable in experiments since it minimizes the effect of fluctuations in the number of atoms loaded into the cavity. Corresponding to this atomic shift of the cavity resonance frequency there is an AC Stark shift of the atomic levels by the light in the resonator

$$\delta\omega_a = 2\langle c^\dagger c \rangle \frac{g^2}{\Delta}, \quad (3.2)$$

where the factor of 2 arises because $\delta\omega_a$ is the energy difference between two levels that shift in opposite directions. Finally, the atoms scatter photons into free space at an average rate per atom

$$\Gamma_{\text{sc}} = \langle c^\dagger c \rangle \frac{g^2}{\Delta^2} \Gamma. \quad (3.3)$$

It will be convenient to rewrite these quantities in terms of the cavity cooperativity $\eta = 4g^2/\kappa\Gamma$, which compares the rate Γ for an excited atom to decay by spontaneous emission into free space to the rate $4g^2/\kappa$ for decay via the cavity mode when the latter is resonant with one of the atomic transitions.

$$\frac{\delta\omega_c}{\kappa S_z} = \eta \left(\frac{\Gamma}{2\Delta} \right) \quad (3.4)$$

$$\frac{\delta\omega_a}{\langle c^\dagger c \rangle \kappa} = \eta \left(\frac{\Gamma}{2\Delta} \right) \quad (3.5)$$

$$\frac{\Gamma_{\text{sc}}}{\langle c^\dagger c \rangle \kappa} = \eta \left(\frac{\Gamma}{2\Delta} \right)^2 \quad (3.6)$$

The first equation gives the cavity shift (expressed as a fraction of a linewidth) per atom transferred from $|\downarrow\rangle$ to $|\uparrow\rangle$, while the latter two give the atomic energy shift and scattering rate relative to the rate at which the photons would leave the cavity in the absence of an external drive.

3.2 Conditional Squeezing by Cavity Assisted Measurement

Since the defining quality of a squeezed spin state is its well-known transverse spin component, a straightforward approach to squeezing is to measure such a transverse component. Ideally, the transverse spin component after the measurement agrees with the measurement result; if the measurement uncertainty is less than the projection noise of the spins, and if the coherence of the system remains high enough that the metrological squeezing criterion ζ decreases, then this final state is squeezed. Unfortunately, the ideal quantum measurement that leaves the system in a state which agrees with the measurement result is difficult enough to realize in practice that it merits its own adjective: Quantum Non-Demolition (QND). Kuzmich, Bigelow and Mandel were the first to point out that the dispersive interaction between an atomic ensemble and a propagating light field could be used to perform a QND measurement of a spin component [17]. Efforts to implement squeezing by such optical measurement have been one of the most active areas in spin squeezing research for the past decade [12, 27–31]. In all such schemes there is a competition between the coherent, dispersive interaction of the atoms with the light field which is used for measurement and the incoherent scattering into free space which disrupts the squeezing. The measurement schemes discussed in this section can be seen as variations on the theme of reference [17] where an optical resonator enhances the atomic sample’s interaction with the probing light field mode relative to the free space modes.

Since the experimentally-accessible spin component is S_z (i.e. the population difference between $|\uparrow\rangle$ and $|\downarrow\rangle$), the mean spin of the ensemble should be oriented in the equatorial plane such that S_z is a transverse component. For definiteness, the mean spin will be taken to be along $\hat{\mathbf{x}}$, and the initial state of the system will be the coherent spin state $|\text{CSS}\hat{\mathbf{x}}\rangle$ which saturates the SQL. This state is readily obtainable in experiments by optically pumping all atoms into the ground state $|\downarrow\rangle$ and applying a microwave $\pi/2$ pulse to all atoms at once. Two schemes to produce a squeezed state from this starting point are considered here. One, used in the experiments




Figure 3-2: Transmission-based measurement: for a two-sided cavity driven at one half-linewidth detuning, the average transmitted number $\langle p_t \rangle$ depends on the shift of the resonant frequency $\delta\omega_c$ due to the atomic index of refraction.

described in subsequent chapters, is based on a simple transmission measurement through the cavity. The other, which has yet to be implemented, uses an optical phase measurement on a probe laser at the cavity resonance frequency and offers better performance. In each case, the objective is to measure the transverse spin component S_z with an uncertainty much less than the projection noise of the coherent state $|\text{CSS}\hat{\mathbf{x}}\rangle$ while preserving the coherence, as quantified by the Ramsey fringe contrast or mean spin length.

3.2.1 Transmission-based spin measurement

Consider a symmetric optical resonator with separate input and output modes, driven on the input side by a laser pulse containing on average $\langle p_i \rangle$ photons and whose duration is long compared to the cavity lifetime κ^{-1} so that transient effects from the cavity dynamics may be neglected. The photon number transmitted through the cavity p_t is a Lorentzian function of the detuning δ of the drive laser relative to the cavity resonance frequency:

$$\frac{\langle p_t \rangle}{\langle p_i \rangle} = \mathcal{L} = \frac{1}{1 + (2\delta/\kappa)^2}. \quad (3.7)$$

A measurement of this transmitted photon number limited by photon shot noise (Poisson fluctuations) will have a variance $\langle \Delta p_t^2 \rangle = \langle p_t \rangle$, allowing the detection of small shifts $\delta\omega_c$ of the cavity resonance frequency with a mean square fractional error

$$\frac{\langle \Delta \delta\omega_c^2 \rangle}{\kappa^2} = \langle \Delta p_t^2 \rangle \left(\kappa \frac{\partial \langle p_t \rangle}{\partial \delta} \right)^{-2} = \langle p_t \rangle \left(\kappa \frac{\langle p_t \rangle}{\mathcal{L}} \frac{d\mathcal{L}}{d\delta} \right)^{-2} \quad (3.8)$$

$$= \frac{1}{\langle p_t \rangle} \frac{1}{(2\xi)^2} \quad (3.9)$$

where

$$\xi = \frac{\kappa}{2\mathcal{L}} \left| \frac{d\mathcal{L}}{d\delta} \right| = \frac{2(2\delta/\kappa)}{1 + (2\delta/\kappa)^2} \quad (3.10)$$

is the logarithmic derivative of the Lorentzian transmission profile. The maximum value of ξ is 1, attained for a detuning of one half linewidth $\delta = \kappa/2$. As this chapter is concerned with fundamental limits for an idealized model, $\xi = 1$ will be assumed henceforth.

Since the cavity shift $\delta\omega_c$ depends on the population difference between the spin states $|\uparrow\rangle$ and $|\downarrow\rangle$, measuring the cavity shift yields a measurement of the transverse spin component S_z whose variance is

$$\langle \Delta S_z^2 \rangle = \frac{\langle \Delta \delta\omega_c^2 \rangle}{\kappa^2} \left(\frac{\kappa S_z}{\delta\omega_c} \right)^2 = \frac{1}{4\langle p_t \rangle} \frac{1}{\eta^2} \left(\frac{2\Delta}{\Gamma} \right)^2. \quad (3.11)$$

Given a sufficiently large collected photon number $\langle p_t \rangle$ this measurement variance will be less than the projection noise $S_0/2$, potentially enabling spin squeezing.

The price paid for this knowledge of S_z is that while photons are being collected at the cavity output at a rate $\kappa\langle c^\dagger c \rangle/2$, they are being scattered into free space at a rate $N\Gamma_{\text{sc}}$. Defining r as the average number per atom of Rayleigh (spin-preserving) or Raman (spin-changing) scattering events over the duration of the probe pulse, the total probability for each atom to scatter a photon into free space is

$$2r = \langle p_t \rangle \frac{\Gamma_{\text{sc}}}{\kappa\langle c^\dagger c \rangle/2} = 2\langle p_t \rangle \eta \left(\frac{\Gamma}{2\Delta} \right)^2. \quad (3.12)$$

The spin measurement uncertainty can then be expressed as

$$\langle \Delta S_z^2 \rangle = \frac{1}{4\eta r} = \frac{S_0}{2} \frac{1}{2S_0\eta r}. \quad (3.13)$$

which depends only on the photon scattering probability and the cavity cooperativity. In particular, it is independent of the detuning Δ from atomic resonance.

It may seem surprising that increasing Δ does not reduce the photon scattering probability $2r$ for a given measurement performance: the signal $\delta\omega_c$ scales with Δ^{-1} , so that to maintain constant signal size one should scale the photon number as Δ , while the scattering rate scales with Δ^{-2} . However, what matters to measurement performance is not the signal itself but the signal-to-noise ratio, and the photon shot noise is not a fixed noise floor but grows with photon number. Thus, to maintain a constant signal-to-noise ratio the transmitted photon number must grow as Δ^2 , exactly canceling the reduction in scattering rate.

The scattering of photons has two undesirable effects, loss of coherence which shortens the Bloch vector and added noise which counteracts the variance reduction from the measurement. First, consider the coherence loss. Photons scattered into free space could, in principle, be imaged using a microscope to identify which atoms in the ensemble they came from. If they reveal the spin state of the atom that scattered them, these photons spoil the symmetry of the state; an atom projected into $|\uparrow\rangle$ or $|\downarrow\rangle$ by the observation of a scattered photon contributes nothing to the collective spin vector along $\hat{\mathbf{x}}$, which therefore gets shorter. All Raman-scattered photons reveal the state of the atom that scattered them (it is encoded in the photon frequency), while Rayleigh-scattered photons reveal the atomic state to the extent that the scattering amplitudes are different for the two spin states [32]. Even if the scattering rates are the same for each spin state, as they are in the model we are considering, they can differ by a sign, so that in the worst case every single scattered photon effectively removes an atom from the symmetric ensemble and shortens the Bloch vector by $1/2$. On its own this effect limits length of the Bloch vector to $S/S_0 = e^{-2r}$.

Raman scattering events, by flipping the spin of atoms in the ensemble, change

the state populations during the measurement. These changes lead to a decay of the autocorrelation of S_z for different times

$$\langle S_z(t_1)S_z(t_2) \rangle = \frac{S_0}{2} e^{-2\rho_{t_1,t_2}} \quad (3.14)$$

where ρ_{t_1,t_2} is the fraction of atoms which Raman-scatter a photon between t_1 and t_2 . This autocorrelation can be used to compute the correlation

$$\langle S_z \bar{S}_z \rangle = \frac{S_0}{2} \frac{1 - e^{-2r}}{2r} \xrightarrow{r \ll 1} \frac{S_0}{2} (1 - r) \quad (3.15)$$

between the time average \bar{S}_z recorded by the measurement and the actual value of S_z once the measurement is complete, as well as the variance of the time-averaged measurement result

$$\langle \bar{S}_z^2 \rangle = \frac{S_0}{2} \frac{e^{-2r} - 1 + 2r}{2r^2} \xrightarrow{r \ll 1} \frac{S_0}{2} \left(1 - \frac{2}{3}r \right). \quad (3.16)$$

The resulting additional measurement error is

$$\langle (S_z - \bar{S}_z)^2 \rangle = \langle S_z^2 \rangle - 2\langle S_z \bar{S}_z \rangle + \langle \bar{S}_z^2 \rangle \xrightarrow{r \ll 1} \frac{S_0}{2} \frac{4r}{3}. \quad (3.17)$$

Optimizing the measurement uncertainty therefore involves a tradeoff between reducing fractional photon shot noise by increasing the scattered photon number and reducing the spin noise due to Raman scattering.

There is one additional fundamental effect which can limit squeezing performance: the curvature of the Bloch sphere. As the measurement narrows the range of possible S_z values, its backaction inevitably perturbs the azimuthal angle ϕ , decreasing the length of the mean spin which is proportional to $\langle \cos \phi \rangle$. In the extreme case where ϕ is completely uncertain the mean spin vanishes even if the coherence between the atoms is preserved: the individual atomic Bloch vectors all point in the same direction, but that direction is completely random. The tools to calculate this phase broadening properly will be developed in section 3.3, but for now the following heuristic argument

may suffice: for a cavity driven at half-slope, the photon shot noise on the intracavity intensity has the same fractional variance as the shot noise on the transmitted light, hence the phase variance can be computed as the product of the shot noise on p_t with the average phase shift per transmitted photon:

$$\langle \Delta \phi^2 \rangle = \langle \Delta p_t^2 \rangle \left(\frac{\delta \omega_a}{\kappa \langle c^\dagger c \rangle / 2} \right)^2 = 4 \langle p_t \rangle \eta^2 \left(\frac{\Gamma}{2\Delta} \right)^2 = 4r\eta. \quad (3.18)$$

Approximating the phase distribution by a Gaussian of this variance, the shortening of the mean spin vector is $\langle \cos(\phi) \rangle = e^{-2r\eta}$.

Given these constraints, the squeezing is determined by comparing the uncertainty on S_z to the contrast remaining after the free space scattering and Bloch sphere curvature have taken their toll:

$$\zeta = \frac{2\langle \Delta S_z^2 \rangle}{S_0} \frac{1}{C^2} = \frac{1}{1 + \left(\frac{1}{2S_0\eta r} + \frac{4r}{3} \right)^{-1}} \frac{1}{(e^{-2r(1+\eta)})^2} \quad (3.19)$$

The leading 1 in the denominator accounts for the *a priori* knowledge that the system started in the state $|\text{CSS}\hat{\mathbf{x}}\rangle$ with $\langle \Delta S_z^2 \rangle = S_0/2$, the next two terms describe the measurement variance due to photon shot noise and the additional uncertainty due to Raman scattering, and the second fraction includes the contrast loss due to scattering into free space and curvature.

If the collective cooperativity $N\eta = 2S_0\eta$ is small, then no squeezing is possible: the loss of contrast from scattering into free space shortens the Bloch vector faster than the measurement can reduce $\langle \Delta S_z^2 \rangle$. However, in the regime of strong collective cooperativity $N\eta \gg 1$, the squeezing parameter is approximately given by

$$\zeta \approx \frac{1}{2S_0\eta r} + \frac{4r}{3} \geq 2\sqrt{\frac{2}{3S_0\eta}} \quad (3.20)$$

where the best squeezing, saturating the inequality on the right, occurs for a scattered photon number $2r = \sqrt{3/2S_0\eta} \ll 1$ for which the contrast loss can be neglected, and where the knowledge of the initial state can be ignored since the measurement provides much better resolution. Note that the resonant optical depth can be written as




Figure 3-3: Theoretical predictions of squeezing by transmission-based measurement. The curves show the metrological squeezing parameter ζ as a function of the measurement strength ηr for $S_0 = 10^4$ ($2 \cdot 10^4$ atoms) and cavity cooperativities $\eta = 0.01$ (dashed curve), $\eta = 0.1$ (chain-dotted curve) and $\eta = 1$ (solid curve). In all cases the squeezing initially improves as $(2S_0\eta r)^{-1}$ as the measurement improves knowledge of the transverse S_z component, before reaching an η -dependent limit set by random spin flips due to Raman scattering during the measurement. $\eta = 0.1$ is representative of the experiments discussed later in this thesis.

$2N\eta = 4S_0\eta$. Thus, when Raman scattering is the dominant decoherence mechanism, the achievable improvement in measurement variance from this scheme scales as the square root of the optical depth. This scaling is the prime motivation for using an optical resonator for spin squeezing: while a free-space laser beam making a single pass through an atomic cloud might see an optical depth of tens or hundreds, optical depths of 10^4 or more are easily achievable in an optical resonator where the beam has to make thousands of round trips through the atomic cloud before escaping. In the limit of extremely strong coupling $\eta \gg S_0$ the measurement backaction wraps the state around the curved Bloch sphere before the decoherence from Raman scattering becomes important, and the achievable squeezing comes within a constant factor of the Heisenberg bound: $\zeta \geq 2e/S$.




Figure 3-4: The ideal measurement of the atom-induced cavity shift uses a one-sided (strongly overcoupled) resonator and measures the phase of the reflected light. Here the phase measurement is schematically illustrated as a Michelson interferometer, though a practical implementation would probably use a Pound-Drever Hall scheme. The dispersive output signal of such an interferometer allows the cavity shift to be measured by a probe laser on resonance (red line on the plot), maximizing the coupling of the atoms to the cavity.

3.2.2 Spin readout by optical phase shift

The side-of-slope measurement described in the preceding subsection, while straightforward to implement, sacrifices performance in two respects: it neglects the information available in the beam reflected from the cavity, and by driving the atoms at a frequency different from the cavity resonance frequency it decreases the coupling of the atoms to the resonator relative to their coupling to free space. Both effects increase the number of photons that must be scattered into free space before attaining a given measurement resolution. A better scheme is to use a one-sided cavity, driven at its resonance frequency, and detect the cavity shift by a Pound-Drever Hall measurement of the reflected field.

The analysis of this case is very similar to that of the side-of-slope case, so only the differences will be summarized here. The phase of the reflected field $2 \arctan(2\delta/\kappa)$ can be detected with a variance $1/4\langle p_i \rangle$ by an ideal homodyne measurement which saturates the number-phase uncertainty relation for the $\langle p_i \rangle$ -photon pulse reflected

from the cavity. For a measurement on cavity resonance this gives a resolution of the cavity shift

$$\frac{\langle \Delta \delta \omega_c^2 \rangle}{\kappa^2} = \frac{1}{64 \langle p_i \rangle} \quad (3.21)$$

and thus a measurement of the spin S_z with variance

$$\langle \Delta S_z^2 \rangle = \frac{1}{64 \langle p_i \rangle \eta^2} \left(\frac{2\Delta}{\Gamma} \right)^2. \quad (3.22)$$

For a single-sided cavity driven on resonance, the incident or reflected photon rate is $\kappa \langle c^\dagger c \rangle / 4$, so that the total scattered photon number is

$$2r = \langle p_i \rangle \frac{\Gamma_{\text{sc}}}{\langle c^\dagger c \rangle \kappa / 4} = 4 \langle p_i \rangle \eta \left(\frac{\Gamma}{2\Delta} \right)^2 \quad (3.23)$$

and the measurement variance, expressed in terms of scattered photons, becomes

$$\langle \Delta S_z^2 \rangle = \frac{1}{32\eta r} = \frac{S_0}{2} \frac{1}{16S_0\eta r} \quad (3.24)$$

The effects of scattering into free space, calculated in terms of r , are the same as in the previous section, while the phase broadening due to photon shot noise is now

$$\langle \Delta \phi^2 \rangle = \langle \Delta p_i^2 \rangle \left(\frac{\delta \omega_a}{\kappa \langle c^\dagger c \rangle / 4} \right)^2 = 16 \langle p_i \rangle \eta^2 \left(\frac{\Gamma}{2\Delta} \right)^2 = 8\eta r \quad (3.25)$$

leading to a contrast reduction $\langle \cos(\phi) \rangle = e^{-4\eta r}$. Note that the product

$$\langle \Delta S_z^2 \rangle \langle \Delta \phi^2 \rangle = \frac{1}{4} \quad (3.26)$$

saturates an uncertainty relation:

$$\langle \Delta S_z^2 \rangle \langle \Delta \phi^2 \rangle \approx \langle \Delta S_z^2 \rangle \frac{\langle \Delta S_y^2 \rangle}{\langle S_x \rangle^2} \geq \frac{|\langle [S_z, S_y] \rangle|^2}{4 \langle S_x \rangle^2} = \frac{1}{4} \quad (3.27)$$

where the mean spin is aligned with the $\hat{\mathbf{x}}$ axis. In this sense the phase measurement on cavity resonance is ideal: it extracts as much information about the spin as its




Figure 3-5: Relative squeezing performance of ideal transmission-based and reflected-phase measurements. The chain-dotted line is for transmission-based measurement squeezing with $S_0 = 10^4$ and $\eta = 0.1$, as in figure 3-3, while the solid curve shows the squeezing performance of a phase measurement on a one-sided cavity with the same parameters. The latter scheme's improved measurement resolution allows greater squeezing before encountering the limit due to Raman scattering.

backaction will allow.

The achievable squeezing is now

$$\zeta = \frac{2\langle\Delta S_z^2\rangle}{S_0} \cdot \frac{1}{C^2} = \frac{1}{1 + \left(\frac{1}{16S_0\eta r} + \frac{4r}{3}\right)^{-1}} \frac{1}{(e^{-2(1+2\eta)r})^2}. \quad (3.28)$$

In the regime of strong collective cooperativity $2S_0\eta \gg 1$ this improves over the performance of the transmission-based scheme of subsection 3.2.1 by a factor of $2\sqrt{2}$ in the variance

$$\zeta \approx \frac{1}{16S_0\eta r} + \frac{4r}{3} \geq \frac{1}{\sqrt{3S_0\eta}} \quad (3.29)$$

and is achieved at a lower scattered photon number $2r = \sqrt{3/16S_0\eta}$. The improvement in squeezing performance is illustrated in figure 3-5.

To achieve even better performance in measurement-based squeezing, it is possi-

ble to make use of richer atomic level schemes than the simple three-level structure assumed here. The Polzik group have demonstrated a measurement squeezing scheme using separate lasers to probe the $|\uparrow\rangle$ and $|\downarrow\rangle$ states on closed transitions [33], thus eliminating the Raman scattering that is the dominant limitation to the squeezing schemes discussed here. Instead of being set by a tradeoff between the measurement and spin noise from Raman scattering, the maximum squeezing is then determined by the competition between the ever-decreasing measured $\langle\Delta S_z^2\rangle$ and the shrinking contrast, and scales with the collective cooperativity $2S_0\eta$ instead of its square root. This scheme has so far only been demonstrated in free-space squeezing where the collective cooperativity is limited [31]; implementing it in a system with a resonator to enhance the optical depth might allow squeezing to within a small constant factor of the Heisenberg limit.

3.3 Deterministic Squeezing by Cavity Feedback

In the schemes discussed in the preceding section the squeezing is conditioned on the result of the measurement, in the sense that in a given realization of the experiment the state produced can have any transverse spin component within the projection noise $\langle\Delta S_z^2\rangle = \frac{S_0}{2}$; only given the measurement outcome that was recorded for that realization is the state well-known. This has a practical drawback: any imperfections of the measurement apparatus, such as finite quantum efficiency or inhomogeneous backaction, directly limit the achievable squeezing. A procedure that deterministically generates a specific quantum state can be used to generate squeezed states reliably even without a measurement apparatus. Of course, precise state readout is still needed to verify the production of such squeezed states or make use of them in a measurement, but relaxing the requirement that the readout be non-destructive substantially eases the measurement problem. For instance, a measurement of S_z used for squeezing must preserve the phase of the spins to avoid shortening the Bloch vector, while a measurement of S_z used only for readout can inflict arbitrary dephasing on the spins provided only that it correctly counts the number of up and down spins.

Deterministic spin squeezing entails constructing a system whose Hamiltonian evolution converts some easily-prepared initial state to a squeezed state. Since the squeezed state is entangled, the Hamiltonian in question must necessarily include interactions between the spins so that some information ends up shared between them. These interactions must be well-controlled in order to preserve the system coherence. While suitable Hamiltonians—referred to as one- and two-axis twisting—were known from the origins of the spin squeezing literature [5], physical realizations of these Hamiltonians have been elusive, which partly explains the interest in squeezing by QND measurement. The complete control of the Hilbert space of a few qubits in trapped-ion systems has allowed the desired squeezing operations to be synthesized by suitable sequences of entangling gates [10], though applying this approach to large ionic ensembles remains an open experimental challenge. There have also been efforts to obtain squeezing using the direct collisional interactions between atoms in a Bose-Einstein condensate [11, 34–38]. Such techniques may well have a valuable role to play in the future of atom interferometry [36, 38], and are important test cases for precise control of quantum-degenerate gases. On the other hand, they intrinsically require the use of dense, collisionally-interacting atomic clouds; precisely the opposite of the dilute, non-interacting conditions currently preferred for minimizing systematic errors in precision measurements [39]. Techniques for precisely controlling density-dependent shifts, perhaps exploiting the suppression of *s*-wave collisions due to Pauli blocking in ultracold fermions [40], must be developed before squeezing by collisional interaction can become a tool for precision measurement.

This section presents an approach to deterministic spin squeezing in which the necessary interactions for squeezing are generated in a dilute sample of atoms by the light field of an optical resonator, which acts as a shared bus over which the atoms exchange state information. The use of a light-mediated interaction allows the atoms to be placed far apart and the interactions switched off at will (by switching off a drive laser), thus avoiding interaction shifts that might perturb a subsequent measurement.

Qualitatively, the principle of such cavity feedback squeezing is as follows. The population difference between spin states shifts the cavity resonance frequency. If




Figure 3-6: Geometrical interpretation of cavity feedback squeezing: an S_z -dependent phase shift (dashed arrows) shears the uncertainty region of the $|\text{CSS}\hat{x}\rangle$ into an ellipse with a narrower minor axis. The effect is rather as though a circle were drawn on the side of a stack of paper and the top of the stack were pushed to the side, sliding the sheets relative to each other and shearing the circular image.

the resonator is driven by a laser detuned from cavity resonance by half a linewidth, then the intracavity intensity, being detuning-dependent, is sensitive to the resonator frequency shift and thus to the spin component S_z . The intracavity intensity, in turn, imparts an intensity-dependent AC Stark shift to the atomic levels, so that the individual atomic pseudospins precess at a rate which depends on the collective spin component S_z . Microscopically, the net result is an effective interaction where each pseudospin's evolution depends, through the cavity-shift-dependent AC Stark shift, on the spin state of all other pseudospins in the ensemble. That this interaction can lead to squeezing can be understood geometrically by noting that the S_z -dependent precession rate of the Bloch sphere causes a shearing of the initially circular uncertainty region of the state $|\text{CSS}\hat{x}\rangle$, stretching it out into an ellipse whose minor axis is shorter than the radius of the initial circle (figure 3-6).

The remainder of this section serves to formalize this intuitive picture of cavity feedback and give it a quantitative footing. Subsection 3.3.1 analyzes the light-driven interaction for a lossless resonator coupled to a single external continuum (figure 3-7) and driven by a photon Fock state, showing that it can reproduce the one-axis twisting Hamiltonian of Kitagawa and Ueda [5]. Subsection 3.3.2 considers the use of such one-axis twisting for squeezing, including the effects of scattering into free space and noise on the intracavity photon number.




Figure 3-7: The ideal quantum model of cavity feedback squeezing considered in this section includes a single quantized cavity mode (creation operator c^\dagger) interacting with a single continuum of modes outside the cavity (spectrum of creation operators b_ω^\dagger) and with an ensemble of symmetrically-coupled atoms (total pseudospin vector operator \mathbf{S}).

3.3.1 Ideal Unitary Cavity Feedback

The model analyzed in this section is the one presented in section 3.1, but now explicitly including the continuum of modes outside the cavity's input mirror, described using the input-output formalism of Collett and Gardiner [41, 42]. For now, only a single such continuum is included (i.e. the resonator cavity is taken to be one-sided, as in figure 3-7), and scattering into free space is neglected in order to keep the dynamics unitary.

Under these assumptions the system Hamiltonian becomes

$$\frac{H}{\hbar} = \omega S_z + \omega_c c^\dagger c + \int d\omega \omega b_\omega^\dagger b_\omega + \Omega c^\dagger c S_z + i\sqrt{\frac{\kappa}{2\pi}} \int d\omega (b_\omega^\dagger c - c^\dagger b_\omega). \quad (3.30)$$

The first three terms of the Hamiltonian describe the energy of the bare atoms, the intracavity field, and the continuum of modes outside the cavity. The fourth term provides the coupling between the atoms and the cavity field, where $\Omega = 2g^2/\Delta = \eta\Gamma\kappa/2\Delta$ can be interpreted either as the differential AC Stark shift on the atomic levels due to a single intracavity photon or as the cavity frequency shift due to a single spin flip, as can be seen by grouping it with either of the first two terms. The final term describes the coupling between the cavity field and the external modes through the input mirror.

Both the population difference operator S_z and the total photon number operator $c^\dagger c + \int d\omega \omega b_\omega^\dagger b_\omega$ commute with this Hamiltonian, so their eigenvalues can be taken as good quantum numbers. This immediately leads to a first set of energy eigenstates of the form $|m\rangle \otimes |0\rangle$, product states of the atoms in a state of definite population

difference (Dicke state) with a vacuum of all field modes. For any given atomic state $|m\rangle$, the space of states with a single photon in the system, consisting of a discrete level (the photon in the cavity) coupled to a continuum (the photon in one of the external field modes), can be diagonalized by a procedure due to Fano [26, 43]. The result is conveniently expressed by defining an eigenphoton creation operator

$$a_\omega^\dagger = \frac{1}{\sqrt{(\omega - \omega_c - \Omega S_z)^2 + \kappa^2/4}} \times \\ \times \left[-i\sqrt{\frac{\kappa}{2\pi}}c^\dagger + \frac{\kappa}{2\pi}\mathcal{P} \int d\omega' \frac{b_{\omega'}^\dagger}{\omega - \omega'} + (\omega - \omega_c - \Omega S_z)b_\omega^\dagger \right] \quad (3.31)$$

with \mathcal{P} designating the Cauchy principal value. a_ω^\dagger satisfies the usual commutation relations $[a_\omega, a_{\omega'}^\dagger] = \delta(\omega - \omega')$ for a creation operator. It creates a photon in a superposition of intracavity and external field modes with a definite energy. The Hamiltonian can now be rewritten as

$$\frac{H}{\hbar} = \omega_a S_z + \int d\omega \omega a_\omega^\dagger a_\omega \quad (3.32)$$

which is just the free energy of the ensemble of two-level atoms and a single continuum of field eigenstates. It is important to note that, while c^\dagger and b_ω^\dagger act only on the light field and commute with all spin operators, the eigenstates of the field created by a_ω^\dagger must depend on the atomic state, since the field and atoms interact, so that a_ω^\dagger does not commute with S_x and S_y . It does, however, commute, with S_z , which will prove useful.

This form of the Hamiltonian can be used to find the evolution of the atoms under the action of an incident light pulse. Consider the product state of an arbitrary atomic state $|\psi_a\rangle$ with a Fock state of n photons whose field spectrum $B(\omega)$ describes an incident pulse outside the cavity:

$$|\psi(-t_0)\rangle = |\psi_a\rangle \otimes \frac{1}{\sqrt{n!}} \left(\int d\omega e^{i\omega t_0} B(\omega) b_\omega^\dagger \right)^n |0\rangle. \quad (3.33)$$

Introducing a new atom-dependent operator

$$\Theta_\omega = \arctan\left(\frac{\kappa/2}{\omega - \omega_c - \Omega S_z}\right) - \frac{\pi}{4} \quad (3.34)$$

the system state can be rewritten in terms of the field eigenstates as

$$|\psi(-t_0)\rangle = \frac{1}{\sqrt{n!}} \left(\sqrt{-i} \int d\omega e^{i\omega t_0 - i\Theta_\omega} B(\omega) a_\omega^\dagger \right)^n |\psi_a\rangle \otimes |0\rangle \quad (3.35)$$

where the field operators a_ω^\dagger must be allowed to act on both the atomic and the photonic parts of the state. The initial time $-t_0$ is taken to be far in the past, $t_0 \gg \max|dB/d\omega|^{2/3}$, before the photons arrive at the cavity. This form of the initial state, taken together with equation 3.32, makes it easy to compute the time evolution of the system until a time $+t_0$ far after the photons have left the cavity. The usual spin precession gives a factor $e^{-i\omega_a S_z 2t_0}$ and each photon operator evolves by a phase $e^{-i\omega 2t_0}$:

$$|\psi(+t_0)\rangle = \frac{e^{-i\omega_a S_z 2t_0}}{\sqrt{n!}} \left(\sqrt{-i} \int d\omega e^{-i\omega t_0 - i\Theta_\omega} B(\omega) a_\omega^\dagger \right)^n |\psi_a\rangle \otimes |0\rangle. \quad (3.36)$$

Reverting to separate field and atomic operators

$$|\psi(+t_0)\rangle = \frac{e^{-i\omega_a S_z 2t_0}}{\sqrt{n!}} \left(-i \int d\omega e^{-i\omega t_0 - 2i\Theta_\omega} B(\omega) b_\omega^\dagger \right)^n |\psi_a\rangle \otimes |0\rangle, \quad (3.37)$$

shows that this final state, in general, is an entangled state of the (exiting) light field and the atoms since the integral over the various frequency modes b_ω^\dagger depends on the atomic operator Θ_ω . However, provided the spectrum $B(\omega)$ has a bandwidth much less than the cavity linewidth κ , the operator Θ_ω is effectively constant over the range of frequencies that contributes to the integral, and the final state factors into

$$|\psi(+t_0)\rangle = e^{-2i\omega_a S_z t_0 - 2ni\Theta_{\omega_p}} |\psi_a\rangle \otimes \frac{1}{\sqrt{n!}} \left(-i \int d\omega e^{-i\omega t_0} B(\omega) b_\omega^\dagger \right)^n |0\rangle, \quad (3.38)$$

where ω_p is the center frequency of the incident pulse spectrum $B(\omega)$. Thus, as far as the atoms are concerned, the net effect of the incident light pulse has been

a unitary transformation $U_n = e^{-2ni\Theta_{\omega_p}}$. Choosing the incident pulse frequency one half-linewidth detuned from the bare cavity resonance $\omega_p = \omega_c + \kappa/2$ and provided that the atom number fluctuations never shift the cavity by a large fraction of a linewidth $\langle \Delta S_z^2 \rangle \ll \kappa^2/\Omega^2$, the transformation effected by the light pulse becomes approximately

$$U_n = e^{-ni(\phi_0 S_z + \frac{1}{2}\phi_0^2 S_z^2 + \mathcal{O}(\phi_0^3))}, \quad (3.39)$$

an overall rotation of the Bloch sphere combined with a shearing or one-axis twisting [5]. The rotation angle per incident photon $\phi_0 = 2\Omega/\kappa$ is closely related to the cavity frequency shift per atom, as seen in section 3.2, which is why it appears twice in the prefactor of the shearing term: once for the atomic shift of the cavity resonance and once for the ensuing backaction of the light field on the atoms.

It is both theoretically and experimentally convenient to consider the action of a pair of incident light pulses with photon numbers n_1 and n_2 , between which the atomic spin is inverted by a π pulse which sends $S_z \rightarrow -S_z$. Neglecting terms of third and higher order in ϕ_0 , the overall transformation produced by this pulse sequence is

$$U_{n_1, n_2} = e^{-i(\alpha S_z + \frac{1}{2}\mu S_z^2)} \quad (3.40)$$

where the rotation angle $\alpha = (n_2 - n_1)\phi_0$ and the shearing factor $\mu = (n_2 + n_1)\phi_0^2$ can now be independently specified. Note that the coordinate rotation $e^{-i\pi S_x}$ associated with the π pulse itself has been absorbed into the definition of the initial atomic state. The coherent input state $|\text{CSS}\hat{\mathbf{x}}\rangle$ of interest for this thesis is an eigenstate of the π pulse, so that this only changes a physically meaningless global phase. For equal photon numbers in the two pulses, the transformation is then $U_{n,n} = e^{-i\mu S_z^2}$, precisely the one-axis twisting operation considered in Kitagawa and Ueda's original spin squeezing proposal [5].

The transformation U_{n_1, n_2} does not modify the S_z operator, with which it com-

mates, but it does change the spin-raising operator

$$S_{+(+t_0)} = S_x + iS_y = S_{+(-t_0)} e^{i(\alpha + \mu(S_z + \frac{1}{2}))} \quad (3.41)$$

$$S_{+(+t_0)}^2 = S_x^2 - S_y^2 + i(S_x S_y + S_y S_x) = S_{+(-t_0)}^2 e^{i(2\alpha + \mu(2S_z + 2))}. \quad (3.42)$$

From these expressions one obtains the expectation values for the first and second moments of the spin operators in the transformed $|\text{CSS}\hat{\mathbf{x}}\rangle$:

$$\langle S_x \rangle = S_0 \cos(\alpha) \cos^{2S_0-1} \left(\frac{\mu}{2} \right) \quad (3.43)$$

$$\langle S_y \rangle = S_0 \sin(\alpha) \cos^{2S_0-1} \left(\frac{\mu}{2} \right) \quad (3.44)$$

$$\langle S_z \rangle = 0 \quad (3.45)$$

$$\langle S_x^2 \rangle = \frac{S_0}{2} \left[1 + \left(S_0 - \frac{1}{2} \right) \left(1 + \cos(2\alpha) \cos^{2S_0-2}(\mu) \right) \right] \quad (3.46)$$

$$\langle S_y^2 \rangle = \frac{S_0}{2} \left[1 + \left(S_0 - \frac{1}{2} \right) \left(1 - \cos(2\alpha) \cos^{2S_0-2}(\mu) \right) \right] \quad (3.47)$$

$$\langle S_z^2 \rangle = \frac{S_0}{2} \quad (3.48)$$

$$\langle S_x S_y + S_y S_x \rangle = \frac{S_0}{2} (2S_0 - 1) \sin(2\alpha) \cos^{2S_0-2}(\mu) \quad (3.49)$$

$$\langle S_y S_z + S_z S_y \rangle = S_0 (2S_0 - 1) \cos(\alpha) \cos^{2S_0-2} \left(\frac{\mu}{2} \right) \sin \left(\frac{\mu}{2} \right) \quad (3.50)$$

$$\langle S_z S_x + S_x S_z \rangle = S_0 (2S_0 - 1) \sin(\alpha) \cos^{2S_0-2} \left(\frac{\mu}{2} \right) \sin \left(\frac{\mu}{2} \right). \quad (3.51)$$

These are quantum expectation values for the pure state that results from shearing the initial state $|\text{CSS}\hat{\mathbf{x}}\rangle$ with known photon numbers n_1 and n_2 and without scattering into free space. Note that the variance of the spin component S_z has not changed, while that of the perpendicular component S_y has increased; there is, however, an oblique angle between the two where the variance is reduced, as suggested in figure 3-6 and encoded mathematically in the correlation term $\langle S_y S_z + S_z S_y \rangle$. The squeezing obtained in this idealized scenario is analyzed in some detail in reference [5]. Here,

consideration of the squeezing will be deferred until certain realistic corrections have been brought to this model.

3.3.2 Realistic Cavity Feedback Squeezing

Adapting the preceding results to describe the outcome of a real experiment will require accounting for three effects: the coupling of the symmetric Fabry-Pérot resonator used in experiments to two continua (one behind each mirror) instead of the single external continuum considered so far; uncertainty in the photon numbers n_1 and n_2 used to drive the resonator for squeezing; and the effect of scattering into free space. Each of these will be considered in turn. Three additional assumptions will help lighten the notation in this section: the mean spin will be taken along \hat{x} after the shearing ($\langle \sin(\alpha) \rangle = 0$), the total spin $S_0 \gg 1$ will be taken to be large so that $\cos^{2S_0-2}(x) \approx \cos^{2S_0-1}(x) \approx \cos^{2S_0}(x) \approx e^{-S_0 x^2}$, and the shearing will be confined to the moderate regime where $\mu \ll S_0^{-1/2}$. The first condition is naturally satisfied in the experiments to be described, where $\langle n_1 \rangle = \langle n_2 \rangle$, and can always be satisfied after a suitable coordinate rotation. The second condition is also comfortably fulfilled in experiments with $S_0 \sim 10^4$. The third condition is required for any appreciable squeezing, since for $S_0 \mu^2 > 1$ the mean spin length is exponentially suppressed as the sheared uncertainty distribution is appreciably affected by the curvature of the Bloch sphere. In the limit of large μ the direction of the collective Bloch vector becomes completely uncertain and the mean spin length vanishes.

If the one-sided cavity of figure 3-7 is replaced by the symmetric resonator of figure 3-2, then there are two degenerate external field modes at each frequency, one to the left of the cavity (input) and one to the right (output). However, the cavity field only couples to a single linear combination of these modes; the orthogonal combination does not couple to the cavity at all. This can be understood classically by noting that if the cavity is driven from both sides with appropriately chosen phases, the two drives can interfere destructively within the cavity and the field intensity vanishes there. This corresponds to sending photons into the uncoupled mode. With the opposite choice of phase, the drives interfere constructively within




Figure 3-8: In a symmetric Fabry-Pérot resonator, one superposition of external fields to the left and right (with creation operator b_ω^\dagger) interferes constructively in the cavity, while the orthogonal superposition (d_ω^\dagger) interferes destructively in the cavity and therefore does not couple to the cavity field (c^\dagger). The symmetric resonator is therefore equivalent to a one-sided cavity whose single outside continuum (b_ω^\dagger) is mixed on a beam splitter with a separate continuum unconnected to the cavity (d_ω^\dagger).

the cavity, corresponding to photons sent into the coupled mode. Since the coupled and uncoupled modes for a symmetric cavity have equal field amplitudes to the left and to the right of the cavity, photons sent in only from the left (input) side of the cavity are in an equal superposition of coupled and uncoupled modes. In this sense the real symmetric resonator used in experiments is equivalent to the ideal one-sided resonator considered heretofore placed behind a 50/50 beam splitter (figure 3-8).

On average, half the photons incident from one side on a symmetric resonator do not couple to the cavity at all, while the other half impart a phase shift as calculated above. This is why the heuristic argument of section 3.2.1 for estimating the atomic phase broadening during a transmission measurement worked: the transmitted photon number at half-slope is also half the incident photon number, so the average number of transmitted photons $\langle p_t \rangle$ and the average phase per transmitted photon are the same as the average photon number interacting with the atoms $\langle n_1 + n_2 \rangle$ and the phase shift per interacting photon ϕ_0 . By the same token, the shearing parameter $Q = S_0 p_t \phi_0^2$ used in Refs. [13, 14] is equal to $S_0 \mu$. This equivalence only holds at one-half-linewidth detuning. A photon in the coupled mode always imparts

a deterministic detuning-dependent phase shift $2\mathcal{L}\phi_0$ to the atoms; ϕ_0 at half-slope, $2\phi_0$ on resonance since the intracavity amplitude of the coupled mode is higher there. Meanwhile, a transmitted photon may or may not have been in the uncoupled mode so that the associated phase shift on the atoms is unknown, with an expectation value ϕ_0 that is independent of detuning because the average intracavity intensity associated with a given transmitted photon rate is fixed. The Lorentzian detuning dependence instead appears in the number of transmitted photons $\langle p_t \rangle \propto \mathcal{L}$. Since the experiments are conducted at the half-slope point, it will often be convenient to work in terms of the transmitted photon number, which is directly measurable, instead of the cavity-coupled photon number, which is not, but bear in mind the distinction between the two.

This leads to the next complication to be added to the model: fluctuations of the photon number incident on the cavity mode. The experiments to be described are performed at a large detuning, with a correspondingly small single-photon phase shift $\phi_0 \sim 10^{-4}$, and for large photon numbers $n_{1,2} \sim 10^4 \dots 10^6$ whose fluctuations are independent and close to photon shot noise. The distribution of the rotation angle α is therefore approximately Gaussian with a variance $\langle \Delta\alpha^2 \rangle = \gamma \langle n_1 + n_2 \rangle \phi_0^2 = \gamma\mu$ where γ is the ratio of the actual photon number variance to photon shot noise. The variance of the shearing μ is suppressed by two additional powers of ϕ_0 , and its effect is negligible compared to that of the uncertainty on rotation angle. The average over uncertain photon numbers thus yields the following spin moments relevant to squeezing, within the approximations outlined above:

$$\langle S_x \rangle = S_0 e^{-\frac{1}{2}v} \quad (3.52)$$

$$\langle S_y \rangle = \langle S_z \rangle = 0 \quad (3.53)$$

$$\langle S_y^2 \rangle = \frac{S_0^2}{2} + \frac{S_0^2}{2} (1 - e^{-2v}) \quad (3.54)$$

$$\langle S_z^2 \rangle = \frac{S_0^2}{2} \quad (3.55)$$

$$\langle S_y S_z + S_z S_y \rangle = 2 \frac{S_0}{2} \mu S_0 e^{-\frac{1}{2}v} \quad (3.56)$$

where the exponent $v = \gamma\mu + \langle S_z^2 \rangle \mu^2$ is the characteristic phase variance due to both the uncertainty on rotation angle from the photon number fluctuations and the shearing of the atomic S_z distribution. The shortening of the mean spin length $\langle S_x \rangle$ has the same form as in section 3.2. The variance of S_y is initially just the projection noise $\frac{S_0}{2}$, grows at first as the angular variance scaled by the length of the Bloch vector $S_0^2 v$, and eventually saturates at $S_0^2/2$ when the uncertainty distribution is completely wrapped around the Bloch sphere. The correlation term is just the product of the variance of S_z with the factor which converts fluctuations in S_z to fluctuations in S_y : μ radians of phase shift per unit change in S_z , converted into S_y fluctuations by a factor of the spin length.

Finally comes the effect of scattering into free space. As in section 3.2, the length of the Bloch vector is reduced by a factor of $S/S_0 = e^{-2r}$. The spin-flips due to Raman scattering average out some of the atom-induced cavity frequency fluctuations, so that the phase variance is reduced to $v' = \gamma\mu + \langle \bar{S}_z^2 \rangle \mu^2 \approx \gamma\mu + \frac{S_0}{2}(1 - 2r/3)\mu^2$. Raman scattering also reduces the correlation term between S_y and S_z , since the shift imparted to S_y depends on the time-average \bar{S}_z , not the final value S_z at the end of the squeezing operation. Combining these effects yields a final set of spin moments

$$\langle S_x \rangle = S e^{-\frac{1}{2}v'} \quad (3.57)$$

$$\langle S_y \rangle = \langle S_z \rangle = 0 \quad (3.58)$$

$$\langle S_y^2 \rangle = \frac{S_0}{2} + \frac{S^2}{2} (1 - e^{-2v'}) \quad (3.59)$$

$$\langle S_z^2 \rangle = \frac{S_0}{2} \quad (3.60)$$

$$\langle S_y S_z + S_z S_y \rangle = 2 \langle S_z \bar{S}_z \rangle \mu S e^{-\frac{1}{2}v'}. \quad (3.61)$$

To find the squeezing, we first compute the variance of an arbitrary transverse spin component $S_\perp = S_z \cos \beta - S_y \sin \beta$, which would be observed in an experiment by rotating the spins through an angle β about the \hat{x} axis using a microwave pulse

and measuring the population difference between $|\uparrow\rangle$ and $|\downarrow\rangle$:

$$\langle S_{\perp}^2 \rangle = \frac{1}{2} \left(\langle S_y^2 \rangle + \langle S_z^2 \rangle - \cos(2(\beta - \beta_0)) \sqrt{(\langle S_y^2 \rangle - \langle S_z^2 \rangle)^2 + \langle S_y S_z + S_z S_y \rangle^2} \right) \quad (3.62)$$

where the angle that minimizes the variance is given by

$$\beta_0 = \frac{1}{2} \arctan \left(\frac{\langle S_y S_z + S_z S_y \rangle}{\langle S_y^2 \rangle - \langle S_z^2 \rangle} \right). \quad (3.63)$$

For the parameter regime $S_0^{-1} \ll \mu \ll S_0^{-1/2}$ where the shearing is substantial but does not yet wrap around the Bloch sphere, and for small scattering probability $2r \ll 1$, the variance of the best-defined transverse component (for $\beta = \beta_0$) can be closely approximated by

$$\langle S_{\perp}^2 \rangle|_{\beta=\beta_0} \approx \frac{\langle S_y^2 \rangle \langle S_z^2 \rangle - \frac{1}{4} \langle S_y S_z + S_z S_y \rangle^2}{\langle S_y^2 \rangle + \langle S_z^2 \rangle} \quad (3.64)$$

$$\approx \frac{S_0}{2} \left(\frac{1}{(S_0 \mu)^2} + \frac{2\gamma}{S_0 \mu} + \frac{(S_0 \mu^2)^2}{24} + \frac{4r}{3} \right) \quad (3.65)$$

The first term shows the suppression of the fixed initial projection noise by the shearing. The second term shows the suppression of the phase fluctuations due to intra-cavity light noise (ideally photon shot noise); since the cumulative phase uncertainty due to the light noise is growing even as the shearing proceeds, the second term decreases more slowly than the first. Usually the second term dominates and the variance initially decreases as μ^{-1} ; however, if the photon shot noise could be suppressed ($\gamma = 0$), for instance by sending photon Fock states into a one-sided cavity as assumed in section 3.3.1, then the variance would initially decrease as μ^{-2} . This initial variance decrease is stopped by one of two limits, either the deformation of the uncertainty region due to the curvature of the Bloch sphere captured by the third term (the “swirliness” of reference [5]), or by the random spin flips due to Raman scattering in the fourth term. Which of these terms dominates depends on the single-atom cooperativity since the shearing is related to the scattering probability by $\mu = 4\eta r$. For a weakly coupled resonator ($\eta \ll 1$) the scattering term dominates




Figure 3-9: Curves of metrological squeezing ζ as a function of shearing parameter μ for $S = 10^4$. The solid line is for an ideal system with no scattering into free space ($\eta \rightarrow \infty$) and no intensity noise of the incident pulse ($\gamma = 0$). The dotted curve adds scattering into free space ($\eta = 0.1$) to the model, the dashed curve adds photon shot noise on the incident pulse ($\gamma = 1$), and the chain-dotted curve adds both.

while for a strongly coupled resonator ($\eta \gg 1$) the curvature of the Bloch sphere limits the variance reduction before the scattering can become significant.

For the experiments to be described here, with no suppression of the photon shot noise $\gamma \gtrsim 1$ and with a weak single-atom cooperativity $\eta \sim 10^{-1}$, the second term due to photon shot noise and the fourth term due to Raman scattering dominate the transverse uncertainty and the achievable squeezing takes on a familiar form

$$\zeta \approx \frac{\gamma}{2S_0\eta r} + \frac{4r}{3} \geq 2\sqrt{\frac{2\gamma}{3S_0\eta}}. \quad (3.66)$$

This is the same bound as for the transmission-based measurement of section 3.2, but limited by the intracavity light noise (parametrized by γ) rather than the measurement performance. Since generating photon-shot-noise-limited light is a much simpler task than performing photon-shot-noise-limited detection, this is a significant advantage in practice.

Chapter 4

Apparatus

This chapter describes the experimental setup used to test the theory of squeezing in cavities developed in chapter 3. All this work is carried out using ^{87}Rb , whose pertinent properties are summarized first. The vacuum chamber, magnetic field coils and overall layout of the experiment are unchanged from the work reported in two previous theses [44, 45] (figure 4-1). The bulk of this chapter is devoted to those parts of the apparatus that were upgraded for the experiments to be described here: the microwave and laser systems used to manipulate the atoms, improved characterization of the optical resonator, photodetectors capable of reading out the hyperfine pseudospin of the atoms with a precision well beyond the SQL, and confining potentials used to trap the atoms in the resonator mode during the squeezing while maintaining the coherence of the atomic ensemble.

4.1 Pertinent properties of ^{87}Rb

In the experiments to be described here the ideal three-level atoms of chapter 3 are replaced by ^{87}Rb atoms, with the $|F = 2, m_F = 0\rangle$ and $|F = 1, m_F = 0\rangle$ hyperfine levels of the $5^2\text{S}_{1/2}$ electronic ground state serving as the pseudospin states $|\uparrow\rangle$ and $|\downarrow\rangle$ respectively, and the $5^2\text{P}_{3/2}$ excited-state manifold of the D_2 line serving as $|e\rangle$. The relevant atomic levels are summarized schematically in figure 4-2. For a compendious reference on the physical and optical properties of this atomic system, see Steck [46].




Figure 4-1: Experiment chamber

The $|F = 1, m_F = 0\rangle \leftrightarrow |F = 2, m_F = 0\rangle$ transition is one of two so-called clock transitions in the $5^2S_{1/2}$ ground state whose frequency is insensitive to magnetic field up to linear order. It has a minimum frequency of 6 834 682 610.904 32(2) Hz for an unperturbed atom [16], with a quadratic Zeeman shift coefficient of 575.14 Hz/G² [47]. This is the same transition as is used in ^{87}Rb fountain clocks [16] and is available because the atoms are held in an optical dipole trap which confines all the Zeeman sublevels. For operation in a magnetic trap, where the $|F = 1, m_F = 0\rangle$ state is unconfined, the other clock transition between $|F = 1, m_F = -1\rangle$ and $|F = 2, m_F = 1\rangle$ would be used instead, with a minimum frequency of 6 834 678 113.59(2) Hz [48] at the “magic” field $B_0 = 3.228917(3)$ G [49] and a quadratic Zeeman shift coefficient of 431.35957(9) Hz/G². As will be discussed in section 4.6, this choice would reduce sensitivity to external fields and correspondingly increase coherence times. However, driving the two-photon transition $|F = 1, m_F = -1\rangle \leftrightarrow |F = 2, m_F = 1\rangle$ is more difficult than driving the single-photon $|F = 1, m_F = 0\rangle \leftrightarrow |F = 2, m_F = 0\rangle$ transition used in the proof-of-principle experiments reported here.

The D₂ optical transition, with a linewidth $\Gamma = 2\pi \cdot 6.065$ MHz [46], is used for all




Figure 4-2: Pertinent levels and transitions of the ^{87}Rb D₂ line. The $5^2\text{S}_{1/2}$ ground state has hyperfine levels $|F = 1, 2\rangle$, while the $5^2\text{P}_{3/2}$ excited state has hyperfine levels $|F' = 0, 1, 2, 3\rangle$. The hyperfine splitting in the excited state has been exaggerated by a factor of five for visibility. The dotted arrow shows the 780.245 nm $|F = 2, m_F = 2\rangle \leftrightarrow |F' = 3, m_F' = 3\rangle$ cycling transition used for initial trapping and cooling of the atomic sample. The chain-dotted arrows show the 780.232 nm $|F = 1, m_F = \pm 1\rangle \leftrightarrow |F' = 0, m_F' = 0\rangle$ transitions used for optical pumping into the state $| \downarrow \rangle = |F = 1, m_F = 0\rangle$, and the solid arrows show the far-off-resonant probe light used to squeeze and probe the $\{| \uparrow \rangle, | \downarrow \rangle\}$ superposition.

optical measurements and manipulation of the atomic internal state. For squeezing and probing, the atoms are driven with light far-detuned compared to the hyperfine splitting in the excited state $5^2P_{3/2}$. The squared matrix element g^2 with which this light couples to the $m = 0$ pseudospin states $|\uparrow\rangle$ and $|\downarrow\rangle$, summed over excited hyperfine levels, is then $2/3$ of what it would be for a cycling dipole-allowed transition such as $|F = 2, m_F = 2\rangle \leftrightarrow |F' = 3, m'_F = 3\rangle$ [46]; this $2/3$ factor reduces the coupling η of the atoms to the cavity.

4.2 Microwave and RF Drive

Hyperfine transitions between $|\uparrow\rangle$ and $|\downarrow\rangle$ are driven with resonant microwave radiation. These transitions correspond to rotations of the collective pseudospin, and since the characteristic angular scale of the spin states to be manipulated is around a milliradian (sub-projection-noise angular resolution for several tens of thousands of spins), these rotations must be well-controlled and reproducible to avoid introducing technical noise on the orientation of the rotated states. The two microwave synthesis chains used for this purpose each consist of a stable local oscillator (LO) running near the ^{87}Rb hyperfine transition frequency and a frequency-tunable and phase-agile RF source generating an intermediate frequency (IF). These two signals are then combined in a single-sideband mixer to produce a tunable phase-agile microwave signal at the hyperfine transition frequency, which is amplified and broadcast into the vacuum chamber.

For the measurement squeezing experiments of chapter 6, the LO was a Giga-Tron 600 microwave synthesizer running at 6.82 GHz. The long-term frequency stability of the LO was assured by referencing the synthesizer to a rubidium frequency standard (SRS FS725) with a specified accuracy of $\pm 5 \cdot 10^{-11}$ and Allan variance of $1 \cdot 10^{-11}$ at 10 s (the approximate experimental cycle time). The 14 MHz IF was supplied by a standard laboratory function generator (SRS DS345) without any special stabilization measures, since the fractional stability requirements for the IF are much less stringent. The experiment-control software can automatically set the function generator




Figure 4-3: Microwave generation chain schematic for the measurement squeezing (upper) and cavity-feedback squeezing (lower) experiments

frequency via GPIB between experimental runs, but to allow synchronized phase hops and pulses in the midst of an experimental sequence the IF passed through an analog phase shifter and RF switch (Minicircuits SPH-16+ and ZYSWA-2-50DR) before being combined with the LO in a single-sideband mixer (MITEQ SSM 0408LC2CDA). The combined signal was amplified up to ~ 1 W (by a Narda DB02-0097 and HD Communications HD23686 in series) and broadcast into the vacuum chamber by a 10 dBi horn antenna (HD Communications HD18502, figure 4-4). Since the vacuum chamber contains large metallic surfaces which might act as reflectors, a circulator (JQL JIC6500T7200S1) was added between the last amplifier and the horn to protect the amplifier output stage from reflected power.

For the cavity feedback squeezing experiments of chapters 7 and 8, the microwave synthesizer serving as the LO was replaced by a dedicated 6.8 GHz low-noise phase-locked quartz oscillator (Wenzel Associates 500-20132), still referenced to the rubidium frequency standard for long term stability. This change required an increase of the IF frequency (to 34 MHz) beyond the range of the DS345 function generator, which was therefore replaced by a direct digital synthesis chip (DDS, Analog Devices

AD9959) referenced to a 100 MHz crystal oscillator (Wenzel Associates 501-08579). The DDS can perform synchronized phase and frequency hops under real-time control from the experiment software. The analog phase shifter thus became superfluous and was removed. This proved to be a significant improvement, as the analog phase shifter’s insertion loss varies by nearly a dB depending on the phase setting, making independent control of the rotation axis and rotation angle unnecessarily difficult.

The AD9959 evaluation boards used as DDS signal synthesizers for generating the microwave IF and certain laser sidebands (section 4.3 below) are computer-controlled via USB using an undocumented protocol, which is not a problem as long as the vendor-supplied evaluation software is used to drive the board. To allow the experiment-control software to control the board directly, integrating it with the rest of the apparatus, it was necessary to reverse-engineer the control protocol (using the **SniffUSB** packet sniffer). The resulting protocol description and LabVIEW drivers are freely available from <http://cua.mit.edu/AD9959-USB-drivers/> under the terms of the MIT/X11 open-source license.

Though not necessary for any of the experiments described here, it was occasionally convenient for diagnostic purposes to drive RF transitions between adjacent Zeeman levels. When necessary, this was accomplished by connecting an RF function generator (SRS DS345) to the on-chip wire Q4 (figure 4-15), coupling the RF signal in through a pair of capacitors so that the normal Q4 current controller could still set the DC current.

4.3 Lasers

The light used to manipulate and probe the atomic state is produced by a bank of seven lasers, whose output is switched using shutters and acousto-optic modulators, split and combined to provide the various mixtures of wavelengths needed in different regions of the experimental apparatus, then coupled into polarization-maintaining single-mode fibers that transmit the light to one of eight beams that enter the experimental chamber:




Figure 4-4: Horn-side optical pumping beam

1 horizontal MOT beam along the \hat{x} axis.

2 diagonal MOT beams in the $\hat{y}\hat{z}$ plane.

1 circularly-polarized optical pumping beam launched from below the chamber along the \hat{z} axis, capable of addressing both the MOT in the center of the chamber and the final trapped cloud on the chip.

1 horizontally-polarized optical pumping beam from the East or horn side in the $\hat{x}\hat{y}$ plane, at a 70° angle from the cavity axis, addressing atoms in the cavity (figure 4-4).

1 vertically-polarized optical pumping beam from the West or camera side counter-propagating with the preceding one (figure 4-5).

1 imaging beam along the \hat{x} axis, intersecting the cavity mode at right angles and collected by an asphere inside the chamber for lateral absorption imaging of the atomic cloud. This beam is usually blocked by the microwave horn used for driving hyperfine transitions, but is occasionally valuable for establishing the location of the atomic cloud.

The optical resonator mode along the \hat{y} axis.




Figure 4-5: Camera-side optical pumping beam

4.3.1 Reference Laser

The absolute optical frequency reference used to establish all laser frequencies in this work is an external cavity diode laser (ECDL) with a diffraction grating in Littrow configuration locked to the $|F = 3\rangle \rightarrow |F' = 4\rangle$ transition of the D₂ line in ⁸⁵Rb using Doppler-free polarization spectroscopy according to the scheme of Wieman and Hänsch [50] (figure 4-6). The laser lock point is 1130 MHz to the blue of the $|F = 2\rangle \rightarrow |F' = 3\rangle$ transition in ⁸⁷Rb. Reference laser light is used only for frequency locking of other lasers and is not transmitted to the vacuum chamber.

The Wieman-Hänsch polarization spectroscopy scheme relies on Faraday rotation of a probe laser in the atomic vapor to find the atomic transition frequency. Through the end of the experiments described in chapter 6, this Faraday rotation was measured by observing the fraction of light transmitted by a polarizing filter onto a single silicon photodiode. As part of an overhaul of the reference laser setup to allow some of the reference light to be shared with another lab, the readout photodiode was changed to a New Focus 2307 balanced photodetector, using both output ports of a polarizing beam splitter to act as a polarization analyzer. By suppressing the error signal's sensitivity to total laser power (which causes common-mode fluctuations between the two output ports of the analyzer), this setup eliminated the problem of slow drifts of




Figure 4-6: Reference laser and vapor-cell locking setup

the reference laser set point.

4.3.2 Repumper Laser

The repumper laser, a distributed feedback diode laser (DFB, Eagleyard EYP-DFB-0780-00080-1500-SOT02-0000), is used for all operations addressing atoms in $|F = 1\rangle$ (figure 4-7). Its ~ 5.3 GHz beat note with the reference laser is detected by a fast MSM photodiode (Hamamatsu G4176) and combined with a 6.8 GHz local oscillator in a mixer whose output is divided by 4096 to obtain a signal at ~ 350 kHz that serves as input to a conventional PI feedback lock (using a PLL to detect the frequency). The frequency set point can be dynamically adjusted during the experimental sequence to reach any of the $|F' = 0 \dots 2\rangle$ excited states accessible from $|F = 1\rangle$. The repump laser light can be sent to the experiment through both the diagonal MOT beams in the $\hat{y}\hat{z}$ plane, and through all three of the optical pumping beams: from below, from the East (horn) side and from the West (camera) side.




Figure 4-7: Repumper laser




Figure 4-8: MOT laser battery

4.3.3 MOT Lasers

The MOT master laser is used for cooling, optical pumping and imaging of atoms in $|F = 2\rangle$. It is very similar to the repumper laser, being a DFB diode of the same type, with its ~ 1.2 GHz beat note against the reference laser detected on a fast photodiode (another G4176) and divided by 4096 to obtain a signal at ~ 300 kHz that serves as input to a PI feedback lock. The frequency set point can again be adjusted during the experimental sequence to reach any of the $|F' = 1 \dots 3\rangle$ excited states accessible from $|F = 2\rangle$. The MOT master can be sent into the experiment chamber through all three optical pumping beams and from the imaging beam when the latter is in operation.

For initial cooling and trapping of the atomic sample, more power is needed than can be provided by a single 80 mW diode. This is supplied by a pair of 120 mW MOT slave lasers, injection-locked to the MOT master, devoted solely to the MOT beams (figure 4-8). The so-called YZ slave is used to drive both diagonal MOT beams while the X slave drives the horizontal one. For the measurement squeezing experiment (chapter 6) the YZ slave used an inexpensive CD-player diode (Sharp GH0781JA2C), which was replaced for subsequent experiments (chapters 7, 8) by an uncoated SLI-CW-9mm-780-0.15S-R diode; both operated reliably. The X slave used an AR-coated diode (Sacher SAL-780-100) whose free-running wavelength at room temperature was 790 nm. Contracting the diode sufficiently to allow reliable injection locking at the MOT wavelength of 780.245 nm required cooling the diode to below 10 °C, well below the typical dew point in the lab. The laser casing was therefore half-filled with silica dessicant and completely sealed with Torr-Seal adhesive, to prevent ambient humidity from condensing on the laser diode. This gave satisfactory performance for a period of four years, encompassing all the experiments to be described here. However, the temperature differential between the dessicant—most of which was in thermal contact with the laser case, i.e. at room temperature—and the cold laser diode mount meant that any residual humidity tended to condense on the mount and surrounding wires, which eventually corroded. The X slave has now been replaced by an inexpensive




Figure 4-9: Trap laser

CD-player diode (Sharp GH0781JA2C).

4.3.4 Trap Laser

The optical dipole trap in the resonator is generated by a third DFB diode (Eagleyard EYP-DFB-0852-00150-1500-TOC03-0000) operating at 852 nm. As the resonator cavity has a finesse at this wavelength of 4.3×10^4 , a few milliwatts is enough to generate tens of watts of intracavity circulating power for trapping; a high-power trapping laser is therefore not needed. However, as frequency fluctuations of the laser relative to cavity resonance are converted to intensity fluctuations that heat the atomic ensemble [51], it is essential that the laser frequency be stabilized to a small fraction of the cavity linewidth. Two techniques are employed for this purpose. The first is optical feedback from a glass plate mounted 18 cm in front of the laser diode to form an extended low-finesse cavity. AR-coated on one face to reduce loss from secondary reflections, the glass plate is mounted on a cylindrical piezo to allow the external cavity length to be tuned together with the Bragg cavity built into the laser diode chip (figure 4-9).

The second technique used is a Pound-Drever-Hall lock [52] with very large bandwidth. Before being coupled into the cavity, 30 MHz sidebands are applied to the trap laser beam using a fiber-coupled electro-optic modulator (EOM, EOSPACE PM-0K5-20-PFA-PFA-780). The amplitude modulation obtained when these sidebands reflect from the cavity is detected by a fast MSM photodiode (Hamamatsu G4176), amplified




Figure 4-10: Probe Laser

and mixed with the 30 MHz modulation frequency, thus producing the Pound-Drever-Hall error signal which is used to lock the laser carrier to a TEM_{00} cavity resonance. In addition to the usual PI controller, implemented with operational amplifiers and feeding back to the laser diode current and the external cavity's piezo, a special fast path connects the error signal directly to the laser diode current input. This fast path, which increases the lock's unity gain bandwidth from a few tens of kHz to over 1 MHz, is described in more detail in the thesis of Monika Schleier-Smith [53].

4.3.5 Probe Laser

The probe laser (figure 4-10), usually referred to as the cavity lock laser (CLL) in the notes documenting these experiments, is an ECDL in Littrow configuration like the reference laser, but with an additional mirror mounted to the diffraction grating in order to keep the output beam alignment independent of grating angle adjustments [54]. Its frequency is locked to a cavity resonance using a Pound-Drever-Hall lock with a similar fast feedback path to the one used for the trap laser, ensuring a similarly high lock bandwidth of 2 MHz [53]. The need to minimize undesired perturbations from the near-resonant probe light led, however, to the use of a more complex




Figure 4-11: Probe laser sideband schemes (not to scale). The two grey bars indicate the optical transitions from the two hyperfine ground states. Dashed lines indicate light used for locking the laser to the resonator. Pale resonances are not addressed by any light. In both schemes the TEM_{00} resonance between the two atomic transitions is used for probing. In scheme (a) the same carrier is used for probing and locking, with a 113 MHz sideband to lock to the TEM_{10} odd mode and symmetric 14.08 GHz sidebands addressing the opposite slopes of two TEM_{00} Lorentzian resonances, so that frequency shaking of the laser has no effect on transmission to first order. In scheme (b) separate carriers are used for probing and locking, with only the upper 36.6 GHz sideband resonant with a TEM_{00} mode for probing.

locking scheme than the straightforward carrier-on- TEM_{00} -resonance setup used for the trap. Two different schemes were used, one for the measurement-based squeezing of chapter 6 and one for the cavity feedback squeezing of chapters 7 and 8. For both schemes, the goal is the same: except for short periods when probe light is deliberately introduced into the cavity's TEM_{00} mode for measurement or manipulation purposes, there should be as little probe light as possible in the resonator, it should be detuned as far as possible from atomic resonance, and it should be in a TEM_{10} mode with a node at the cavity axis to minimize its interaction with the atoms.

The first scheme, illustrated in figure 4-11a, places the probe laser carrier exactly halfway between two TEM_{00} resonances, 2.5 free spectral ranges or 14.08 GHz to the red of the resonance used for probing the atoms. Because the cavity is not exactly confocal, the carrier is 113 MHz away from the nearest atomic resonance. The beam from the probe laser passes through a fiber-coupled EOM (EOSPACE PM-0K5-20-PFA-PFA-780), which generates the requisite 113 MHz sidebands for the Pound-




Figure 4-12: Cavity input optics, showing the three incident beams used for trapping the atoms, locking the probe laser to the cavity, and probing the atoms

Drever-Hall lock, before being coupled into the cavity. The incident beam is displaced from the cavity's optical axis so that it couples both to the TEM_{00} mode used for probing and to the TEM_{10} mode used for locking. Light reflected in the TEM_{10} mode is collected on an avalanche photodiode and demodulated to provide the lock error signal. Adjustments of the laser frequency relative to the resonator are performed by changing the lock sideband frequency. The fiber-coupled EOM also generates the 14.08 GHz sidebands used for probing. Since the carrier is exactly halfway between two TEM_{00} modes, both these sidebands are simultaneously resonant with resonator modes. The upper or probing sideband is much closer to atomic resonance than the lower or compensation sideband, so it is the one principally affected by the atom-cavity interaction. The compensation sideband is useful, however, to reduce the effect of technical noise: when the probe sideband is tuned slightly to the blue of one cavity resonance, the compensation sideband is slightly to the red of another so that fluctuations of the laser frequency which decrease the probe transmission increase the compensation sideband transmission, yielding a first-order cancellation of the technical transmission noise.

The second scheme, shown in figure 4-11b, moves the probe laser carrier 6.5 free spectral ranges or 36.6 GHz to the red of the resonance used for probing. Since SMA

connector losses are substantial at this frequency, the probe sideband is actually generated as a second-order sideband by driving the EOM with an 18.3 GHz source. The probing and locking light now arrive at the cavity in separate fibers (figure 4-12), with separate fiber-coupled EOMs, so that the probing light can be correctly aligned to the optical axis for maximum mechanically stable coupling to the TEM_{00} mode while the locking light is offset for maximum coupling to the TEM_{10} odd mode. To avoid interference effects between the lock light and any residual stray probe light which reaches the lock photodiode, the two beams are offset in frequency by 80 MHz by having the probe beam bypass the AOM whose first-order output is used for locking. The probe carrier is no longer centered between TEM_{00} modes; the typical Pound-Drever-Hall sideband frequency is now around 50 MHz which, together with the 80 MHz AOM frequency offset, puts the carrier 130 MHz away from the TEM_{10} mode used for locking or 17 MHz away from the symmetry point. The lock performance is sufficiently good that the compensation sideband suppresses less technical noise than the excess photon shot noise that it adds, so that it is better to remove the compensation sideband from the scheme by detuning it from cavity resonance so that it interacts with neither the atoms nor the transmission photodetector.

4.3.6 Beam Switching

Laser beams destined for the experiment pass through acousto-optic modulators (AOMs, Isomet 1201-C). The zeroth-order, unshifted output is discarded or used for monitoring purposes, and the diffracted output, deflected by a few degrees and typically containing 80% of the output when the AOM is driven at full power, is mode-matched into a polarization-maintaining single-mode fiber and sent to the experiment. The same AOM drive frequency (80 MHz) and diffraction order (first-order on the blue side) are used for all AOMs to reduce the risk of confusion and to simplify frequency accounting: the frequency difference between two lasers is the same before and after their respective AOMs. Switching the RF drive power to the AOM allows rapid modulation of the laser power in any beam.

While AOMs are useful for rapid ($\sim 1 \mu\text{s}$) switching of beam power, they do not




Figure 4-13: Laser switchyard. Labels on the left designate the beams from individual lasers. Labels on the right identify optical fibers destined for the experiment. Shutters are circled.

allow a beam to be entirely switched off; for that purpose we use solenoid-operated shutters (manufactured by TK-CMD, see figure 4-13). These shutters are held open by a coil spring by default, but can be closed by applying a short current pulse (nominally 210 mA for 10 ms) to overcome the inductance of the solenoid and the inertia of the shutter blade and held closed against the restoring force of the spring by a much smaller current (nominally 50 mA). Every laser beam, except for the reference and probe lasers which are never switched off, passes through one such shutter so that it may be blocked when stray light would be unacceptable.

Until the summer of 2007, these shutters were driven by a simple RC circuit which used a capacitor in series with the solenoid to supply the additional current for the initial pulse (figure 4-14, upper left). This design suffered from certain drawbacks:

- In order to allow the use of a single transistor with TTL input voltages to control the shutter current, the shutter coil was connected on the power supply side of the driver circuit. Both wires going to the shutter coil were thus at 5 V even when the shutter was unpowered, and an accidental connection between one of those wires and ground (e.g. the optical table) could send several hundred mA




Figure 4-14: Old (upper left) and new (lower right) shutter driver circuits

through the coil for an extended period, overheating and potentially destroying it.

- After switching off power to the coil, the capacitor had to be allowed to discharge before it could sink current for the next shutter-closing pulse. This made it impossible to cycle shutters quickly: if the shutter were not left open for at least 20 ms it would usually fail to close again.
- The DC hold current supplied by the circuit was 150 mA, much more than necessary, which led to the shutters running hot enough to burn stray fingers and limited their lifespan.
- The initial current pulse, though it satisfied the nominal requirements for closing the shutter (300 mA initial current with a time constant of 20 ms) was not large enough reliably to close the shutter in practice, given that the shutters are typically cycled several times every few seconds all day long and must close every time. Slight changes in the friction on the shutter blade would sporadically cause a shutter to stick open, which could lead to intermittent failures anywhere in the experimental cycle. Options to increase the initial pulse energy were

limited: the capacitor, at 2.2 mF, was already unreasonably large and making it larger would have lengthened the recovery time between cycles. Increasing the supply voltage would also have increased the hold current and aggravated the overheating problem.

An improved driver (figure 4-14, lower right), tested in July 2007 and installed for all shutters by December of that year, corrects these deficiencies:

- The transistor receiving the TTL input is now separate from the power transistor controlling current to the coil, which is now on the ground side of the controller. Shorts of the coil wire to the optical table are now harmless.
- An additional transistor discharges the timing capacitor as soon as the solenoid current is switched off so that it is immediately ready for the next shutter-closing pulse. The only limit on the cycling speed is the mechanical delay while the return spring opens the shutter blade.
- The DC hold current is reduced to 90 mA, still more than necessary for holding the shutter closed but not enough to cause noticeable overheating of the solenoid.
- Using the power transistor to amplify the current through the timing capacitor and having a separate reset circuit allows a free choice of the initial pulse energy. This is now sharp enough (600 mA for 10 ms) that the shutters all reliably close within 5 ms.

Since the installation of the new driver circuits, shutters have not failed to close. There have been a few instances of shutters jamming in the closed position, their return springs unable to reopen them, usually because the shutter was deformed by overtightening its mounting screws. Such failures are rare, generally not intermittent (since the spring, unlike the solenoid, does not get a new pulse every experimental cycle) and fairly easy to debug.




Figure 4-15: Optical resonator mounted on atom chip

4.4 Optical Resonator

The optical resonator which plays a central role in this work is a symmetric Fabry-Pérot cavity consisting of two dielectric mirrors of 2.5 cm radius of curvature, mounted below and to either side of the atom-chip support structure (figure 4-15), which is itself suspended from the top of the vacuum chamber. The input (South-side) mirror is glued to a fixed mount, while the output (North-side) mirror is glued to a cylindrical piezo-electrical transducer (from Channel Industries) that allows the length of the cavity to be adjusted through approximately three free spectral ranges (figure 4-15).

As long as the probe laser is locked to the resonator, the probe laser frequency is a measure of the resonator length. Therefore, in order to stabilize the resonator length, the beat-note frequency between probe and reference lasers is used as an error signal for a slow feedback loop that controls the output mirror piezo. Put differently, the high-bandwidth probe lock maintains the relative frequency of the laser and cavity, and the low-bandwidth cavity length lock ensures that the absolute frequencies of both resonator mode and laser give the desired detuning from atomic resonance. In the first probe laser scheme, used for the work of chapter 6, the 12.1 GHz beat note between probe laser carrier and reference laser is directly observed on a fast MSM photodiode and converted to a voltage by a delay-line interferometer, using a 40 dB

pre-amplifier (JCA812-600) to overcome the roll-off in the photodiode response at this frequency. In the second scheme, used for chapters 7 and 8, the 34.5 GHz beat-note frequency between probe laser carrier and reference laser is too high to be directly observable with a photodiode. Instead, light from both the probe and reference lasers is sent through a fiber-coupled amplitude modulator (EOSPACE AZ-0K5-20-PFA-PFA-780) driven by the same 18.3 GHz synthesizer as is used to generate the probe sidebands. The interference between the AM sidebands imparted to the two laser carriers introduces a modulation at the difference frequency between the separation of the probe and laser carriers and the 36.6 GHz probe sideband frequency, which is exactly the frequency difference between the probe sideband that interacts with the atoms and the reference laser. This difference frequency is detected by a fast photodiode and delay-line interferometer, as before. Note that while the probe laser frequency is subject to high-frequency noise and must be locked to a small fraction of the \sim MHz cavity linewidth, thus requiring a high-performance feedback loop, the cavity length is subject only to acoustic noise and need only be stable compared to the \sim GHz detuning from atomic resonance, so that a simple PI loop with a bandwidth of \sim 200 Hz suffices for the cavity length lock.

The mirror coatings were applied by REO Inc. for a design wavelength of 852 nm and target transmission of 30 ppm with losses below 5 ppm [55]. At 852 nm the measured transmission of a sample mirror from the same coating run was 27 ppm. To measure the mirror transmission, a weakly focused 852 nm laser beam was sent through two apertures and on to a silicon photodiode. The mirror was mounted between the two apertures and aligned by overlapping the incident and retroreflected light beams. The mirror alignment is important both to maximize the efficiency of the dielectric coating on the mirror and to minimize the amount of light reflected into the back of the first aperture and re-scattered onto the photodiode. The radius of the first aperture was adjusted to minimize the light transmitted to the photodiode: too large and stray light off the mirror axis would be allowed to reach the photodiode, too small and the incident laser beam would diffract outwards, allowing stray light to go around the mirror and again causing an excess reading on the photodiode. The

measurement was carried out in darkness since it was found that draping the setup in black cloth only served to reflect off-axis stray light back towards the photodiode. Once this mirror transmission reading was collected, the mirror was removed and the power transmitted through the same aperture openings was recorded to normalize the result to a fractional transmission.

The cavity linewidth was measured either by direct transmission measurements or by ringdown. For the direct transmission measurements, a large-amplitude EOM sideband of a laser locked to the cavity was scanned over the cavity resonance by adjusting the modulation frequency. Since the transmission profile obtained is the convolution of the cavity and laser lineshapes, a Lorentzian fit will tend to overestimate the linewidth of the cavity unless the laser is locked to within a small fraction of a cavity linewidth. For the ringdown measurements, the laser amplitude incident on a cavity resonance was given a sharp-edged square-wave modulation, either by locking the trap laser carrier to cavity resonance and intermittently modulating sidebands to steal power from the carrier, or by locking the probe carrier away from cavity resonance and intermittently adding sidebands on cavity resonance. Either way, the response of the cavity transmission to the step changes in the input power is the inverse Fourier transform of the Lorentzian cavity lineshape, an exponential decay with the resonator lifetime as its time constant. This method will underestimate the linewidth (overestimate the lifetime) unless both the optical power switching and the photodetector response are much faster than the cavity lifetime. For the 852 nm trapping wavelength, the linewidths obtained were $\kappa_t = 2\pi \cdot 131(3)$ kHz for the transmission spectrum by scanning a sideband and $\kappa_t = 2\pi \cdot 135(2)$ kHz from the ringdown measurement. For the 780 nm probing wavelength, the transmission spectrum gave $\kappa = 2\pi \cdot 1.012(3)$ MHz and the ringdown measurement gave $\kappa = 2\pi \cdot 1.01(3)$ MHz. The parenthesized values are statistical fit uncertainties. The good agreement between two measurement techniques with opposite biases is evidence that the cavity linewidth/lifetime has been correctly identified, and incidentally that the feedback loops are locking the laser to within much less than a cavity linewidth.

The resonator free spectral range of 5632.0(2) MHz was measured by recording

the sideband modulation frequency required to tune a pair of symmetric laser sidebands to simultaneous resonance with two longitudinal modes of the resonator, 5 free spectral ranges apart. From the free spectral range we obtain a cavity length $L = 26.615(9)$ mm. The uncertainty comes primarily from thermal drift: in normal operation the cavity length is locked, but only up to a half-integer multiple of the optical wavelength; as the experiment warms up and the mirror mounting structure expands, the length to which the cavity is locked is stepped from one mode to the next to stay within the range of the longitudinal adjustment piezo.

With the symmetric sidebands each on a TEM_{00} resonance (and thus with the laser carrier exactly halfway between such resonances), the Pound-Drever-Hall sideband frequency for the probe laser (i.e. the detuning of the carrier from the nearby TEM_{10} mode) is 113(1) MHz, with the error coming principally from the freedom of adjustment of the Pound-Drever-Hall error offset and from possible deformations of the error signal. This yields a resonator g -parameter [56] of -0.0629(6), assuming that the two cavity mirrors have the same curvature. The radius of curvature is then 25.040(16) mm, in good agreement with the nominal value of 25 mm.

From the curvature, cavity length and wavelength we can compute the TEM_{00} mode waist at cavity center according to [56]

$$w_0^2 = \frac{L\lambda}{\pi} \sqrt{\frac{1+g}{4(1-g)}} \quad (4.1)$$

For the 780.245 nm probe wavelength, we obtain a waist $w_0 = 55.71(2)$ μm at cavity center. For the 852 nm trap laser, which sees the same mirror geometry but has a longer wavelength and somewhat more diffraction, the waist at cavity center is 58.2(1) μm , with a larger uncertainty due to the uncertainty of the absolute frequency of the trap laser. From measurements described in section 5.2 the atom cloud is known to be displaced by some 2.6(5) mm from the cavity center. Using the usual Gaussian

beam diffraction formula

$$w(z) = w_0 \sqrt{1 + \left(\frac{z}{z_R}\right)^2}, \quad (4.2)$$

$$z_R = \frac{\pi w_0^2}{\lambda} \quad (4.3)$$

the mode waists at the location of the atoms are $w_p = 56.9(4) \mu\text{m}$ and $w_t = 59.4(5) \mu\text{m}$ for the probe and trap laser respectively.

The cooperativity for a closed dipole transition of a (two-level) atom on the cavity axis at an antinode of the standing-wave intracavity probe field is then given by

$$\eta_0 = \frac{24\mathcal{F}/\pi}{(kw)^2} \approx 0.203 \quad (4.4)$$

where \mathcal{F} is the resonator finesse and $k = 2\pi/780.24 \text{ nm}$ is the probe wavenumber. The use of an open transition with a smaller matrix element reduces this by a factor of $2/3$ (see section 4.1). Moreover, for atoms displaced from the cavity axis or from the antinode of the standing wave, the cooperativity is further reduced by a position-dependent factor proportional to the local intensity of the probe light mode.

Since the cooperativity is no longer the same for atoms at different locations, the ideal model of chapter 3, which assumes identical atoms, must break down in the limit where individual atomic pseudospins are resolvable. However, provided that the number of atoms and the variance in any given collective pseudospin component remain large, the collective pseudospin can be treated as an effectively continuous variable [23], obtained from a suitably weighted sum over the atoms. In this limit the composition of the collective pseudospin in terms of microscopic atomic pseudospins no longer matters and the behaviour of chapter 3 is recovered. The leading-order departure from this idealization is due to dephasing from the position-dependent AC Stark shift of the probe light, which can be addressed as discussed in section 5.5 and reference [33].

What effective cooperativity η is associated with this weighted collective pseudospin? First of all, η should reproduce the observed shift of the cavity resonance

when all atoms are optically pumped into $|\uparrow\rangle$ or $|\downarrow\rangle$, since that is the experimentally observable manifestation of collective pseudospin length (effective atom number). Recalling from equation 3.5 that the cavity shift is proportional to ηS_z , this requires that

$$S_0\eta = \int d\mathbf{r} \rho_r \frac{1}{2} \eta_r = \frac{N_t}{2} \langle \eta_r \rangle \quad (4.5)$$

where ρ_r is the number density of the cloud, $\frac{1}{2}$ is the S_z component of each optically pumped atom, η_r is the position-dependent cooperativity, N_t is the total number of atoms in the cloud and $\langle \eta_r \rangle$ is their average cooperativity. Additionally, the fluctuations of the cavity shift due to the projection noise of uncorrelated atoms should also be preserved, requiring

$$\frac{S_0}{2} \eta^2 = \int d\mathbf{r} \rho_r \frac{1}{4} \eta_r^2 = \frac{N_t}{4} \langle \eta_r^2 \rangle \quad (4.6)$$

where $1/4$ is the spin variance contributed by each atom in the cloud. These two requirements preserve the meaning of ζ as a change in the noise-to-signal ratio relative to the best possible ratio achievable without entanglement and with a fully polarized ensemble of independent atoms. They impose the definition

$$\eta = \frac{\langle \eta_r^2 \rangle}{\langle \eta_r \rangle} \quad (4.7)$$

for the effective cooperativity.

As will be discussed in sections 4.6 and 5.2, the atomic distribution is effectively uniform along the cavity axis and Gaussian in the two transverse directions, so that

$$\langle \eta_r \rangle = \frac{2\eta_0}{3} \int_0^{2\pi} \frac{d(ky)}{2\pi} \cos^2(ky) \int_0^\infty dr r \frac{e^{-\frac{r^2}{2\sigma^2}}}{\sigma^2} e^{-\frac{2r^2}{w^2}} = \frac{2\eta_0}{3} \frac{1/2}{1 + 4\frac{\sigma^2}{w^2}} \quad (4.8)$$

$$\langle \eta_r^2 \rangle = \left(\frac{2\eta_0}{3} \right)^2 \int_0^{2\pi} \frac{d(ky)}{2\pi} \cos^4(ky) \int_0^\infty dr r \frac{e^{-\frac{r^2}{2\sigma^2}}}{\sigma^2} e^{-\frac{4r^2}{w^2}} = \left(\frac{2\eta_0}{3} \right)^2 \frac{3/8}{1 + 8\frac{\sigma^2}{w^2}} \quad (4.9)$$

$$\eta = \frac{2\eta_0}{3} \left(\frac{3}{4} \right) \frac{1 + 4\frac{\sigma^2}{w^2}}{1 + 8\frac{\sigma^2}{w^2}} \quad (4.10)$$




Figure 4-16: Cavity output filtering optics

where the $2/3$ comes from the matrix element of the open transition used in these experiments (section 4.1), y is the axial coordinate, r is the distance from the cavity axis and σ is the RMS cloud radius. For a cloud radius of $7\text{ }\mu\text{m}$ (section 5.2) this yields an effective cooperativity $\eta = 0.096(3)$.

4.5 Photodetector

Nearly all measurements performed on the atoms in these experiments rely on measurements of the transmission fraction of the cavity, as in section 3.2.1, and so the performance of the photodetector used to collect the transmitted light is critical. Performance here has three aspects: collection of the transmitted light with high quantum efficiency, rejection of stray light, and bandwidth sufficient to allow rapid measurements (compared to the coherence time). The photodetector itself, a silicon avalanche photodiode (Hamamatsu S3884) followed by a specially tuned amplifier with a bandwidth of $\sim 100\text{ kHz}$, is described in reference [53]. The focus here will be on the stray light rejection optics (figure 4-16).

The beam transmitted by the cavity contains three frequency components: trap

light, light from the Pound-Drever-Hall sideband used to lock the probe laser to the resonator, and the transmitted probe sideband itself. Of these, only the weakest, the probe sideband, is of immediate interest. The \sim mW trap light is much the brightest of the three components, but as it is at 852 nm it can be easily filtered out: an interference filter (CVI HRF-780 laser line filter) suppresses it to undetectable levels while transmitting the 780 nm probe laser with over 95% efficiency. The locking sideband, with a power 1...100 times greater than the probe sideband depending on the experiment, also fits into the several-nm-wide passband of the filter and must be disposed of separately. It can be distinguished from the probe sideband of interest on two grounds: it is transmitted on the TEM_{10} odd mode of the resonator while the probe sideband leaves in a fundamental TEM_{00} mode, and it is generally chosen to have the orthogonal polarization. The light exiting the cavity is therefore transformed by a quarter-waveplate and half-waveplate so that the probe light is horizontally polarized, passes through a polarizing beam splitter which deflects most of the locking sideband (as well as a portion of the trap laser light used for power monitoring), then passes through the laser line filter used to block the trap light, before finally being coupled into a polarization-maintaining single-mode fiber for delivery to the photodetector. Since the TEM_{00} mode transmitted by the cavity is a clean Gaussian mode, it can be coupled into the fiber with over 80% efficiency while the orthogonal TEM_{10} mode (as well as ambient stray photons from overhead lighting or other lasers) are rejected by the fiber which acts as a spatial filter. Depending on beam alignment, which is unfortunately subject to thermo-mechanical drifts, polarization and spatial mode-matching each give two to three orders of magnitude of suppression of the locking sideband, reducing it to negligible levels.

The overall gain of the photodetection path was periodically verified by comparing the difference in signal voltage at the detector output when the probe sidebands were tuned on and off resonance to the change in photocurrent measured with an unamplified large-area photodiode at the cavity output port of the vacuum chamber. The observed gain of 10(2) nW/V is consistent with an estimated optical efficiency of 60%, photodiode quantum efficiency of 80%, avalanche gain of 13, and amplifier

transimpedance gain of 25 M Ω .

4.6 Dipole Trap

The atoms progress through four different trapping potentials in an experimental cycle (c.f. section 5.1). The first three are used only for sample preparation: a six-beam magneto-optical trap in the centre of the chamber used to collect the atomic ensemble, a macroscopic magnetic quadrupole trap used to transfer it to the vicinity of the resonator and a microscopic quadrupole trap based on the Q4 chip wire (figure 4-15) used to compress it radially into an elongated cloud similar in diameter to the cavity mode. These are described in references [44, 45]. For the actual experiments the atoms are confined in the optical dipole trap produced by the standing wave built up by the 852 nm trap laser inside the resonator mode.

The use of an intracavity dipole trap of this type has several advantages. The first is automatic alignment: coupling the trap laser to the same TEM₀₀ transverse mode as the probe laser guarantees that the minimum of the trapping potential is on the resonator axis. The second is state-independent confinement: since all magnetic sub-levels are trapped in a dipole potential, the atoms can be optically pumped between states without fear of losing them to magnetically untrapped states. Furthermore, the easiest “clock” (linear-Zeeman-free) transition to drive in the ⁸⁷Rb hyperfine system is $|F = 1, m_F = 0\rangle \leftrightarrow |F = 2, m_F = 0\rangle$, which connects magnetically-untrapped states. Finally, since the distance from the resonator mode to the chip (250 μ m) is significantly greater than the mode waist, it is easier to provide tight radial confinement with an optical potential whose transverse scale is set by the resonator mode than with a magnetic potential whose gradient scales inversely with the distance to the wires which generate it. Figure 5-1 shows the substantial radial compression when loading into the dipole trap, which brings the atoms close to the resonator axis where they see the most coupling to the probe light.

Because the trap laser forms a standing wave inside the resonator, the trapping potential actually consists of a large number of “pancakes”, one for each antinode

of the standing wave. The trapping potential in each pancake, calculated from the transmitted trap laser power and the measured transmission of the cavity mirror (section 4.4), has a depth of $U_0 = h \cdot 20 \text{ MHz}$, an axial trapping frequency $\omega_t = 2\pi \cdot 500 \text{ kHz}$ and a radial trapping frequency $\omega_r = 2\pi \cdot 1.5 \text{ kHz}$. The trap standing wave is incommensurate with the probe standing wave, the spatial beat period between them (i.e. the distance between points where the probe antinode overlaps with the minimum of a trapping potential) being $4.6 \mu\text{m}$. Since this is much less than the $\sim 1 \text{ mm}$ length of the atom cloud in the trap (section 5.2), the atoms occupy trapping sites with all possible relative phases between trap and probe, and the distribution of atoms over the axial standing wave of the probe is effectively uniform. Also, since the axial trapping frequency is large compared to all experimental timescales, the axial distribution of atoms is frozen: there are atoms trapped at both nodes and antinodes of the probe standing wave, but an atom at an antinode remains there for the duration of the experiment.

The intensity of the trap laser light in the cavity mode is sensitive to frequency fluctuations of the laser relative to the cavity resonance frequency. The resulting intensity fluctuations cause exponential heating of the atoms with a characteristic e-folding time $4(\omega_t^2 S(2\omega_t))^{-1}$, where $S(2\omega_t)$ is the one-sided spectral power density of the fractional intensity fluctuations at twice the trap frequency [51]. Because of the high axial trap frequency ω_t , this corresponds to a very short time unless the laser is locked to a very small fraction of the cavity linewidth; fortunately, the combination of optical feedback and high-bandwidth Pound-Drever-Hall locking described above is enough to suppress the intensity noise density at twice the trap frequency to $S(1 \text{ MHz}) \lesssim 10^{-12} \text{ Hz}^{-1}$ so that the observed lifetime of atoms in the dipole trap is around 7 s, longer than any experimentally relevant time scale.

The trapping potentials experienced by the two hyperfine levels differ fractionally by $2.2 \cdot 10^{-4}$ due to the 6.8 GHz difference in detunings from their respective optical transitions. Given the thermal distribution of the atoms (with $k_B T \sim h \cdot 1 \text{ MHz}$) in the trapping potential, this leads to an inhomogeneous broadening of the hyperfine transition frequency by tens of hertz, significantly limiting the coherence time of

the pseudospin. To overcome this limitation, the trap laser is circularly polarized, thus introducing a vector Stark shift which acts as an effective magnetic field along the trap beam. In combination with a (real) magnetic offset field along the cavity axis, this laser-induced effective field produces an intensity-dependent quadratic Zeeman shift which can be used to cancel out the inhomogeneous scalar Stark shift. This trick, along with its experimental implementation in our apparatus, trade-offs and optimization of parameters, is discussed in reference [53]. Unfortunately, this technique for canceling the effect of the trap Stark shift requires a bias field B_y of several gauss. Since the minimum of the quadratic Zeeman shift for the $|F = 1, m_F = 0\rangle \leftrightarrow |F = 2, m_F = 0\rangle$ transition is at zero field, this introduces a linear sensitivity of the clock transition frequency to magnetic field fluctuations at the level of several kHz/G as the price paid for insensitivity to trap intensity variations. While this compromise is adequate for the proof-of-principle demonstrations of this work, where the atomic phase is never interrogated for longer than a few ms, true precision experiments with long interrogation times will probably have to give up the convenience of the dipole trap and use magnetically trapped atoms and the $|F = 1, m_F = -1\rangle \leftrightarrow |F = 2, m_F = 1\rangle$ two-photon clock transition at its magic field of 3.23 G. With no AC Stark shifts from the trap and no linear sensitivity to magnetic fields, such a system has been shown to allow phase coherence times of several seconds [49].

Polarization-gradient cooling requires degenerate Zeeman sublevels and thus cannot be performed in a circularly polarized dipole trap. To retain the ability to perform such cooling in the dipole trap at the beginning of the experimental sequence while using the vector-scalar Stark shift cancellation described above for squeezing experiments later in the sequence, the trap laser polarization is dynamically adjustable using a liquid crystal variable retarder (Meadowlark Optics). Unfortunately, the retarder is very slow, requiring around 75 ms to switch and several hundred ms before the polarization drift settles to an undetectable level.

Chapter 5

Experimental Procedures and Apparatus Characterization

Building on the description of the experimental apparatus in the previous chapter, this chapter describes a number of basic experimental procedures used either to characterize the operation of the apparatus or as basic steps in the more complex experiments described in subsequent chapters. Since all subsequent experiments depend on the presence of a cloud of cold ^{87}Rb atoms in the resonator, the first two sections explain how this cloud is prepared and characterized. The next section describes the basic tools for manipulation of the hyperfine pseudospin: optical pumping to initialize the Bloch vector in a known direction and resonant microwave pulses to rotate it. The last three sections describe standard measurement protocols: readout of S_z and S_0 based on the cavity shift, measurement of the length of the Bloch vector and characterization of the decoherence processes that shorten it, and *in situ* magnetic field calibration.

5.1 Standard Experimental Sequence

^{87}Rb vapor emitted by an electrically heated getter is first collected in a standard six-beam magneto-optical trap (MOT). The MOT loading time ranges from 3s for diagnostic experiments, where small atom number and rapid cycle time are desirable,

to 6 s for squeezing experiments requiring maximum optical depth. In the last 30 ms of the loading time the magnetic field minimum is adjusted from the center of the intersection region of the MOT beams to the upper edge, nearest the chip. The repumper intensity is turned down, allowing the atoms to spend some time in the $|F = 1\rangle$ dark states, and the MOT laser detuning is increased, reducing outwards radiation pressure from scattered photons and compressing the atomic cloud (“dark MOT” [57]). To transfer the atoms from this compressed MOT to the macroscopic magnetic quadrupole trap for transport to the chip, the MOT beams are switched off and a $\sim 10\,\mu\text{s}$ current pulse to a solenoid in the vacuum chamber displaces the magnetic field minimum faster than the atoms can follow, placing them in a uniform vertical bias field. The vertical circularly polarized optical pumping beam, aligned with the quantization axis thus defined, drives the atoms into the magnetically trappable $|F = 2, m_F = 2\rangle$ state. The magnetic field minimum is then returned to the position of the atoms, which are now confined in a magnetic quadrupole trap.

Over the next 650 ms the position of the quadrupole trap is ramped up through the remaining centimeter which separates it from the chip. As the cloud approaches the chip, a U-shaped on-chip wire (Q4) adds strong magnetic field gradients in the transverse (\hat{x} and \hat{z}) directions, compressing the cloud into a cigar shape (c.f. section 5.2) parallel to the cavity axis. The dipole trap laser intensity is then turned down to minimize the energy gained by the atoms as they fall into the AC Stark potential, the magnetic trap is aligned with the cavity axis, the trap laser is ramped back up in 20 ms to confine the atoms on the cavity axis and the magnetic trapping fields are ramped off over the next 300 ms.

The atomic cloud heats as it is compressed into the dipole trap, which is much tighter and deeper than the magnetic trap. The atoms are recoiled to below 50 μK by the “grey molasses” variant of polarization-gradient cooling [58], using 5 ms of illumination by the counterpropagating horizontal beams at an angle of 70° to the cavity axis with orthogonal linear polarizations and the circularly polarized beam from below, orthogonal to the cavity axis. Each beam contains cooling light tuned to the blue of the $|F = 2\rangle \leftrightarrow |F' = 2\rangle$ transition and repumper light tuned to the blue

of the $|F = 1\rangle \leftrightarrow |F' = 1\rangle$ transition. This cooling stage requires that the different Zeeman levels in each hyperfine state be degenerate, except for the light shifts from the cooling light, and thus it must take place in zero magnetic field (see section 5.7 for magnetic field calibration procedures) and with linear trap polarization to avoid vector Stark shifts.

For maximum coherence lifetime it is desirable to carry out experiments on the atoms in the polarization-compensated trap mentioned in section 4.6, so the trap laser polarization is switched to circular with the variable retarder and a magnetic bias field along the cavity axis (\hat{y}) is ramped on. The sample is now ready for one of the experiments described in the following sections. Once that is completed, all settings are ramped back to their initial values to begin the loading of the next MOT. The whole experimental cycle repeats with a period of between 5 s and 9 s, depending principally on the MOT loading time.

5.2 Atom Cloud Characterization: Position and Temperature

Detailed spatial information on the ^{87}Rb cloud can be obtained from absorption imaging on the $|F = 2\rangle \leftrightarrow |F' = 3\rangle$ transition, collecting the light from the imaging beam using an aspheric lens mounted to the chip support structure inside the chamber. Typical images collected in this fashion are shown in figure 5-1, showing the atom cloud in the magnetic trap just before and in the dipole trap just after the trap laser is ramped on. Two bridge wires, visible as trapezoids projecting downwards from the chip surface, provide useful position references in these images: their height and separation are both $500\text{ }\mu\text{m}$. Note that the atom cloud is only $250\text{ }\mu\text{m}$ from the chip surface: the resonator mode actually passes between the bridge wire and the chip surface [44, 45].

Figure 5-2 shows the atomic cloud in the final position used for the squeezing experiments. Unfortunately, the cloud goes past the edge of the field of view of




Figure 5-1: Absorption images of the atomic cloud in the magnetic (upper) and optical dipole (lower) traps, showing the strong radial compression provided by the latter. The upper edge of the images corresponds to the chip surface, the circular arcs in the lower corners to the edge of the field of view of the imaging lens, and the two vertical structures are bridge wires on the chip $500\text{ }\mu\text{m}$ high and spaced $500\text{ }\mu\text{m}$ apart. The cloud image in the dipole trap appears dimmed due to inhomogeneous broadening of the optical transition by the steep dipole potential.




Figure 5-2: Absorption image of the atom cloud in the position used for squeezing experiments




Figure 5-3: Cloud temperature measurement by radial expansion. The solid black line shows the fit for $k_{\text{B}}\bar{T} = h \cdot 0.83(1)$ MHz, or an absolute radial temperature of $T = 43(1)$ μK .

the imaging lens. However, by assuming it to be of similar length to the clouds in figure 5-1, and using the bridge wires as position markers, it is possible to infer that the cloud center is displaced by 2.6(5) mm from the center of the cavity towards the cavity input mirror, which information is useful in calculating the optical mode waist size at the location of the atoms (section 4.4).

The radial size of the cloud is determined from the temperature of the atoms' radial motion, which can be evaluated in a single experimental cycle by abruptly switching off the dipole potential and observing the decrease of the cavity shift as the atoms freely expand away from the cavity axis (figure 5-3). Assuming a thermal kinetic energy distribution, the RMS radius of the cloud during this expansion is given by $\sigma(t)^2 = (1 + \omega_r^2 t^2) k_{\text{B}} T / m \omega_r^2$ where the radial trap frequency is given by $\omega_r = 4U_0 / m w_t^2$. It follows that the average atom-cavity coupling, and hence the atom-induced cavity shift, fall as

$$\delta\omega_c(t) = \frac{\delta\omega_c(\sigma=0)}{1 + 4\frac{\sigma^2}{w_p^2}} = \frac{\delta\omega_c(t=0)}{1 + k_{\text{B}}\bar{T} \frac{4t^2}{m w_p^2}} \quad (5.1)$$

where the reduced temperature \overline{T} is related to the thermodynamic temperature T by

$$\frac{\overline{T}}{T} = \frac{1}{1 + \frac{w_i^2}{w_p^2} \frac{k_B T}{U_0}} \xrightarrow{k_B T \ll U_0} 1. \quad (5.2)$$

From the observed temperature, which is typically below 50 μK for properly optimized polarization gradient cooling parameters, the cloud radius can be inferred to be 7(1) μK , as used in calculating the effective cavity cooperativity in section 4.4.

Note that this technique only measures the temperature of the radial degrees of freedom. Since the polarization-gradient cooling beams are all nearly perpendicular to the cavity axis, they do not efficiently cool the atoms' axial degrees of freedom, which are therefore probably somewhat hotter. Fortunately, determining this temperature precisely is not so important as the incommensurate probe and trap lattices ensure that the axial position of the atoms relative to the probe standing wave is uniformly distributed regardless of the atomic temperature (section 4.6).

5.3 Optical Pumping

Optical pumping between $|F = 1\rangle$ and $|F = 2\rangle$ can be performed with very high efficiency since the excited-state linewidth Γ is much smaller than the ground-state hyperfine splitting. Within the $|F = 1\rangle$ manifold, optical pumping into the initial state $|\downarrow\rangle = |F = 1, m_F = 0\rangle$ can be performed by driving the $|F = 1\rangle \rightarrow |F' = 0\rangle$ transition with the camera-side optical pumping beam, which contains only σ_+ - and σ_- -polarized light when the magnetic bias field and quantization axis are chosen along the cavity axis. The state $|F = 1, m_F = 0\rangle$ is dark to such illumination since there are no $|F' = 0, m'_F = \pm 1\rangle$ states to which to couple, but off-resonant scattering from the $|F' = 1, m'_F = \pm 1\rangle$ states limits the efficiency of this optical pumping to around 90%.

Microwave pulses can spectroscopically address individual Zeeman sublevels and further purify the state. For the measurement squeezing experiments of chapter 6 the unpumped atoms were disposed of by first transferring the pumped atoms from

$|F = 1, m_F = 0\rangle$ to $|F = 2, m_F = 0\rangle$ by a microwave pulse, optically pumping all remaining atoms into $|F = 2\rangle$, returning the pumped atoms from $|F = 2, m_F = 0\rangle$ to $|F = 1, m_F = 0\rangle$ by a second microwave pulse and finally blowing away the residual atoms left behind in $|F = 2\rangle$ with resonant light on the $|F = 2\rangle \leftrightarrow |F' = 3\rangle$ cycling transition. While this method successfully produces atomic samples with 99% of the remaining atoms in the desired $|F = 1, m_F = 0\rangle$ initial state, it leads to the loss of around 12% of the atomic sample.

If the same atomic sample is to be reused multiple times this loss is undesirable, as it changes the atom number in the cloud after each initialization, so that the results obtained after different initializations are not readily comparable. This became particularly important in the cavity feedback squeezing experiments of chapters 7 and 8, where each atomic sample was initialized as many as 10 times before being discarded, in order to accelerate data taking. For these experiments, the 10% of atoms in $|F = 1, m_F = \pm 1\rangle$ after the initial optical pumping were allowed to remain there and their number was simply subtracted from the measured total atom number in order to obtain the number of “participating” atoms initially in $|F = 1, m_F = 0\rangle$. Ignoring the additional unpumped atoms in this way is legitimate to the extent that they contribute neither signal nor noise to the measurement of the collective pseudospin S_z , as described in section 5.5, so that they are truly invisible spectators to the squeezing experiments. The spectator atom number (typically 10% of the total) was determined by comparing the population transfer obtained using optical pumping, addressing all atoms, to that obtained by driving the $|F = 1, m_F = 0\rangle \leftrightarrow |F = 2, m_F = 0\rangle$ clock transition with microwaves, addressing only the participating atoms. By allowing the spectator atoms to remain in the trap while awaiting the next initialization, the atom number can be kept stable to within 10% as the atomic state is reinitialized 10 times in a row.




Figure 5-4: Typical Rabi spectrum, scanning the frequency of a $60\text{ }\mu\text{s}$ microwave pulse across the $|F = 1, m_F = 0\rangle \leftrightarrow |F = 2, m_F = 0\rangle$ clock transition. The 85% contrast on resonance is limited by a combination of imperfect pumping into the $|F = 1, m_F = 0\rangle$ initial state and inhomogeneity of the Rabi frequency.

5.4 Microwave Rotations

While optical pumping is used for destructive initialization of the atomic state, coherent rotations of the collective Bloch vector are performed using the microwave system described in section 4.2, which can drive the $|F = 1, m_F = 0\rangle \leftrightarrow |F = 2, m_F = 0\rangle$ clock transition with a Rabi frequency up to 12 kHz. For experimental convenience the microwave power is typically adjusted to give a $\pi/2$ pulse time of $30\text{ }\mu\text{s}$ (8.3 kHz Rabi frequency). The microwave frequency is tuned to the hyperfine resonance frequency so that the axis of microwave-driven rotations lies in the equatorial plane of the Bloch sphere. Since the hyperfine transition frequency is affected by the trap potential, which may drift over time, this frequency adjustment is regularly checked by taking a Rabi spectrum. The spectrum is obtained by optically pumping the atoms into $|F = 1, m_F = 0\rangle$, driving the atoms with $60\text{ }\mu\text{s}$ of microwave radiation and measuring the number of atoms transferred to $|F = 2\rangle$ (see section 5.5). This number is normalized to the total atomic population, measured by optically pumping the atoms between $|F = 1\rangle$ and $|F = 2\rangle$, to obtain a fractional signal that is insensitive to




Figure 5-5: The lifetime of the Rabi nutations is primarily limited by the inhomogeneous strength of the microwave drive: when the microwave power is decreased by a factor of 10 (dashed fit, open circles), the nutation rate and the decay rate decrease by the same factor.

fluctuations in the number of atoms loaded on each cycle (figure 5-4).

As the atoms are located within a microwave wavelength of several metal structures (chip wires, the copper chip mount, steel mirror mounts, etc.) it is to be expected that boundary conditions will alter the magnetic field produced by the microwave horn, introducing some inhomogeneity in the microwave Rabi frequency of different atoms. This effect is clearly visible in the decay of the Rabi oscillations shown in figure 5-5: when the microwave power is decreased, the Rabi frequency and the decay rate are reduced in equal measure, so that the oscillations decay in a fixed number of cycles, not a fixed absolute time. This indicates that the dominant decoherence mechanism limiting the fidelity of the microwave-driven rotations is the inhomogeneity of the rotation rate of the various atomic spins in the ensemble.

To overcome this limitation, π rotations of the collective spin are driven using a composite sequence of three microwave π pulses with phases chosen to suppress errors from differing rotation rates (SCROFULOUS sequence [59]). The price paid for this suppression is an increased vulnerability to errors in the microwave frequency,

which is most noticeable when the effective axis of the composite rotation is parallel to the state's Bloch vector. It was therefore found necessary in practice to choose the effective rotation axis to be along $\hat{\mathbf{y}}$, rather than along the initial Bloch vector along $\hat{\mathbf{x}}$, so that drifts in the microwave transition frequency would not lead to broadening of the state beyond the projection noise limit as it was flipped by the microwaves. With the phase chosen correctly, the π pulse contrast, checked by driving multiple composite π pulses in a row and comparing the population transfer by successive pulses, is better than 98%.

Note that to avoid inhomogeneous broadening of the microwave transition frequency by the probe laser AC Stark shift, all probe laser sidebands resonant with the cavity are switched off and the laser frequency locking loop is disabled during the microwave pulses. The probe laser lock bandwidth is sufficient that re-enabling the locking sideband and feedback loop 20 μs before a subsequent measurement gives adequate time for the laser to stabilize its frequency relative to the cavity resonance.

5.5 Atomic-State Readout

Aside from the cloud position measurements described in section 5.2, all experiments described in this thesis rely on detection of the atom-induced frequency shift of the optical cavity resonance, following the side-of-slope transmission measurement squeezing analyzed in section 3.2.1. This section discusses some of the practical aspects of this measurement.

By way of introduction, figure 5-6 shows a typical experimental trace, which can be interpreted as follows. The probe sideband is tuned to the blue slope of the cavity resonance, so that its transmission increases when the resonance frequency increases and vice versa (see also figure 3-2). The sideband is switched on as the trace begins, with the atoms optically pumped into $|F = 2\rangle$. After 500 μs the atoms are pumped into $|F = 1\rangle$, shifting the cavity resonance to the red and reducing the probe transmission. The size of this reduction indicates the total atom number $2S_0$. The probe light is then switched off briefly while a 60 μs microwave π pulse




Figure 5-6: A typical trace of the cavity transmission, showing the raw voltage output of the avalanche photodetector recorded as a function of time

transfers atoms back to $|F = 2\rangle$, so that a subsequent $50\text{ }\mu\text{s}$ probe light pulse shows an increased transmission once again. The difference between this level and the fully pumped level seen in the first $500\text{ }\mu\text{s}$ indicates the finite contrast of the simple microwave π pulse. After these preliminaries, the state $|\text{CSS}\hat{x}\rangle$ is prepared by optical pumping into $|\downarrow\rangle = |F = 1, m_F = 0\rangle$ (section 5.3) and a pair of $50\text{ }\mu\text{s}$ probe pulses, $330\text{ }\mu\text{s}$ apart, is used to measure its S_z component. This measurement is repeated for verification. In the example trace shown, a fresh $|\text{CSS}\hat{x}\rangle$ state is then prepared and measured twice in the same way. There follows a second pair of $500\text{ }\mu\text{s}$ readings with the atoms optically pumped into $|F = 2\rangle$ and $|F = 1\rangle$, to measure the change in total atom number during the squeezing experiments, then the trapping potential is dropped to release the atoms and a radial expansion measurement is performed to check the temperature (section 5.2). Finally, a trio of pulses on the blue slope, the peak and the red slope of the transmission resonance serve to check that the probe sideband is correctly tuned relative to the empty transmission resonance and to check the laser power.

Whenever the probe sideband is on, the detuning of the probe sideband from cavity resonance can be determined from the transmitted fraction of the light simply

by inverting the Lorentzian transmission function:

$$\frac{\delta}{\kappa} = \frac{1}{2} \sqrt{\frac{\langle p_i \rangle}{\langle p_t \rangle} - 1}. \quad (5.3)$$

The transmitted fraction $\langle p_t \rangle / \langle p_i \rangle$ can be determined straightforwardly given the transmission signal, the known photodiode output voltage when the sideband is off (measured around the 9 ms mark in figure 5-6) and the output voltage when the probe sideband is known to be on cavity resonance. This last quantity is obtained from the final triplet of pulses. Ideally the middle pulse is on resonance and gives the peak transmission directly. Even if it is not, the peak transmission can be derived knowing that the three pulses are spaced in frequency by 0.5κ : labeling the transmitted power of the red-slope, peak and blue-slope pulses as j_r , j_p and j_b respectively, the approximation for the peak transmission

$$\frac{j_r + j_p + j_b}{2} + \frac{(j_b - j_r)^2}{2(j_r + j_p + j_b)} \quad (5.4)$$

is insensitive to detuning errors up to sixth order. The pulse triplet also serves to monitor the frequency of the probe sideband: when the probe is correctly tuned to the half-slope of the unperturbed cavity, the transmission pulses are symmetric. Errors in laser detuning relative to the cavity will imbalance j_r and j_b by moving one of the pulses up the slope of the cavity resonance and the other down.

In the simplified model of chapter 3, the cavity resonance frequency was chosen halfway between the frequencies of the optical transitions from $|\uparrow\rangle$ and from $|\downarrow\rangle$. While not necessary in principle, this choice has the practical benefit that the resonance frequency shift depends only on the population difference $2S_z$ between the states $|\uparrow\rangle$ and $|\downarrow\rangle$, and has no explicit dependence on the total atomic population $2S_0$, so that fluctuations in the number of atoms loaded do not add noise to the cavity shift signal so long as the Bloch vector lies near the equator. To reproduce this feature in the experiment, the effective detuning of the probed resonance from the transitions $|F = 1\rangle \leftrightarrow |F' = 0, 1, 2\rangle$ and $|F = 2\rangle \leftrightarrow |F' = 1, 2, 3\rangle$, averaged over the different

hyperfine levels in the excited state and weighted by the coupling matrix elements for the various transitions, must be chosen equal and opposite. For the linearly polarized probe light used in these experiments this symmetric frequency is 3182 MHz to the blue of the $|F = 2\rangle \leftrightarrow |F' = 3\rangle$ cycling transition, yielding a weighted detuning of ± 3286 MHz for the optical transitions from the two states $|F = 1\rangle$ and $|F = 2\rangle$ and a resonance frequency shift of $\kappa \cdot 4.4(2) \cdot 10^{-5}$ for every atom of population difference between the two hyperfine states. This is the cavity detuning used in chapters 7 and 8.

Note that, due to the opposite detuning for the two hyperfine states the Rayleigh scattering amplitudes are opposite for atoms in $|\uparrow\rangle$ and $|\downarrow\rangle$, so that the worst-case situation considered in section 3.2.1 applies: every Rayleigh-scattered photon projects the scattering atom into either $|\uparrow\rangle$ or $|\downarrow\rangle$, removing it from the coherent subensemble that contributes to the clock signal. However, this effect is not noticeable in practice because technical sources of decoherence shorten the Bloch vector much more than does scattering into free space.

The use of a compensation sideband to suppress technical noise (section 4.3) introduces two additional complications that apply only to the measurement squeezing experiments of chapter 6. First, the transmitted fraction $\langle p_t \rangle / \langle p_i \rangle$ must be calculated after subtracting the power of the compensation sideband which also reaches the transmission photodiode. Second, since the frequency of the cavity resonance near the compensation sideband is somewhat modified by the presence of atoms in the resonator, it is necessary to adjust the detuning of the probe from atomic resonance to 3570 MHz blue of the $|F = 2\rangle \leftrightarrow |F' = 3\rangle$ transition, so that the S_0 -dependent shifts of the two resonances will balance to give an S_0 -independent total transmission signal. Neither of these complications arise in the work of chapters 7 and 8, where the probe sideband is the only light to reach the transmission photodiode. The small spatial overlap between the atomic cloud on the cavity axis and the TEM₁₀ resonance to which the laser is locked suppresses the atomic shift of the locking resonance by 4 orders of magnitude to an entirely negligible level. No additional correction for the shift of the lock resonance is necessary.

At the detunings employed here the optical transition matrix elements are the same to within less than $< 3\%$ for all Zeeman sublevels, so that the resonator shift is essentially only sensitive to the relative population of the two hyperfine levels, without regard to whether the atoms are actually in the $m = 0$ clock state pair or not. This allows the total atom number to be measured using only efficient hyperfine optical pumping to transfer the atoms between $|F = 2\rangle$ and $|F = 1\rangle$, as seen in the first 1 ms of the trace in figure 5-6.

For the measurements of S_z used in the squeezing experiments, a more sophisticated procedure is called for in order to detect only those atoms in the $m_F = 0$ clock manifold while minimizing sensitivity to technical noise and unnecessary measurement-induced decoherence. The S_z measurement consists of two 50 μs pulses of probe light separated by a composite π pulse which inverts S_z . The change in transmission between the two pulses reveals the cavity frequency shift associated with this inversion, which changes the population difference between the hyperfine levels by $4S_z$. Note that while the cavity shift depends only on the population of the hyperfine states without regard to the magnetic sublevels, this differential signal is only sensitive to the atoms in $m_F = 0$. Spectator atoms in $m_F \neq 0$ are not addressed by the π pulse, and so do not lead to a change in cavity resonance frequency between the two probe pulses. They merely produce a common-mode shift that is ignored by the differential-mode measurement. Other sources of technical noise, such as slow laser power fluctuations, are similarly suppressed to the extent that the errors they produce are common to both pulses in the measurement.

More important than its ability to suppress technical noise is, however, the two-pulse measurement scheme's role in suppressing atomic decoherence. Since atoms at different locations along the cavity axis see different probe light intensities, they suffer different phase shifts due to the probe laser AC Stark shift. This inhomogeneous dephasing dramatically shortens the collective Bloch vector, prohibiting squeezing if left uncorrected. The two-pulse measurement with an intervening π rotation forms a spin-echo sequence: provided the atoms suffer approximately the same phase shift in each probe pulse, these phase shifts cancel, largely undoing the inhomogeneous

dephasing. The remaining dephasing is due to the difference in intensities between the two probe pulses, which difference is proportional to S_z and is small for the states near the equator of the Bloch sphere to be considered here.

The echo effect only occurs if the coupling of the atoms to the cavity is the same for the two pulses. Axially this condition is well satisfied: the atoms are tightly confined by the standing wave lattice of the dipole trap, with a trap oscillation frequency that is large compared to the probe pulse duration so that any thermal motion will be averaged over within each probe pulse. Radially, however, the atoms are free to move on timescales comparable to the measurement time. The 330 μs spacing of the probe pulses is therefore chosen to match one-half of the radial oscillation period, so that when the second probe pulse arrives the atoms will be the same distance away from the cavity axis as for the first probe pulse, with the same coupling to the probe field.

The use of an echo measurement which distinguishes between participating atoms in $m_F = 0$ and spectator atoms in $m_F \neq 0$ complicates the analysis of the readout noise due to Raman scattering in section 3.2.1 by introducing several qualitatively different scattering processes beyond the simple spin-preserving and spin-flipping processes considered in the idealized model. Aside from Rayleigh scattering, whose effect is the same as in the idealized model, there are now three different classes of Raman scattering events: those that take an atom out of the $m_F = 0$ clock manifold without changing its hyperfine state, occurring with probability r_o per atom; those which leave the atom in $m_F = 0$ so that it continues to count as a participating atom but flip its hyperfine state, occurring with probability r_f per atom; and those which change both m_F and F , occurring with probability r_{of} per atom in the ensemble. To understand their effect it is important to distinguish between $2S_{zi}$, the population difference between $|F = 2, m_F = 0\rangle$ and $|F = 1, m_F = 0\rangle$ at the start of the measurement; $2S_{zf}$, the same population difference at the end of the measurement and neglecting those atoms which have left the $m_F = 0$ clock manifold due to scattering; and $2\bar{S}_z$, the average population difference between $|F = 2\rangle$ and $|F = 1\rangle$ measured via the cavity shift. Working to first order in these scattering probabilities one finds that the

variances of the initial and final S_z are given by

$$\langle S_{zi}^2 \rangle = \frac{S_0}{2} \quad (5.5)$$

$$\langle S_{zf}^2 \rangle = \frac{S_0}{2} (1 - r_o - r_{of}), \quad (5.6)$$

which is just the projection noise of the atoms in $m_F = 0$ at the beginning and end of the measurement respectively. The variance of the measurement and its correlations with the initial and final S_z values are given by

$$\langle \bar{S}_z^2 \rangle = \frac{S_0}{2} \left(1 - \frac{2r_f}{3} - \frac{r_o}{2} - \frac{2r_{of}}{3} \right) \quad (5.7)$$

$$\langle S_{zi} \bar{S}_z \rangle = \frac{S_0}{2} \left(1 - r_f - \frac{r_o}{2} - \frac{r_{of}}{2} \right) \quad (5.8)$$

$$\langle S_{zf} \bar{S}_z \rangle = \frac{S_0}{2} (1 - r_f - r_o - r_{of}), \quad (5.9)$$

leading to identical scattering-induced measurement errors on S_{zi} and S_{zf} :

$$\langle (S_{zi} - \bar{S}_z)^2 \rangle = \langle (S_{zf} - \bar{S}_z)^2 \rangle = \frac{S_0}{2} \left(\frac{4r_f}{3} + \frac{r_o}{2} + \frac{r_{of}}{3} \right). \quad (5.10)$$

At the detuning used here and given the overlap of the atomic cloud with the TEM_{00} cavity mode used for probing, the scattering probabilities are related to the transmitted photon number $\langle p_t \rangle$ by

$$r_f = 2.5 \cdot 10^{-8} \langle p_t \rangle \quad (5.11)$$

$$r_o = 1.5 \cdot 10^{-8} \langle p_t \rangle \quad (5.12)$$

$$r_{of} = 1.3 \cdot 10^{-8} \langle p_t \rangle \quad (5.13)$$

so that the scattering-induced measurement error is expected to be

$$\langle (S_{zi} - \bar{S}_z)^2 \rangle = \langle (S_{zf} - \bar{S}_z)^2 \rangle = \frac{S_0}{2} \frac{\langle p_t \rangle}{2.2 \cdot 10^7}. \quad (5.14)$$

To evaluate the performance of the readout, two successive identical measurements




Figure 5-7: Readout noise as a function of the photon number $\langle p_t \rangle$ used in the measurement. The open circles and dotted line correspond to the case of no atoms in the cavity and show the intrinsic detector noise. The filled circles and dashed line correspond to the case of $2S_0 = 3.1 \cdot 10^4$ atoms in the cavity, so that noise from Raman scattering also contributes. The solid line indicates the photon shot noise limit.

of the same atomic state are performed. The sequence of state preparation and double measurement is repeated 50...100 times. Ideally, S_z should be the same for both measurements and their results should agree. The variance of the difference between the successive measurements corresponds to the sum of the noise variance of the two measurements; provided the measurements are similar and their noise is uncorrelated, the observed variance can simply be divided by two to obtain the single-measurement precision. Typical data for the readout noise as a function of the photon number $\langle p_t \rangle$ employed in the measurement are presented in figure 5-7.

Consider first the noise measured without atoms in the cavity, so that the true S_z is identically zero (open circles). The fit (dotted line) shows that the three dominant noise contributions are a background electronic noise with a fixed variance of $3.2(6) \cdot 10^5 \text{ photons}^2$, which degrades the measurement for small readout photon numbers, photon shot noise at five times the ideal level (solid line) due to finite quantum efficiency and avalanche noise in the photodetector and fractional light noise—due to




Figure 5-8: Ramsey decay

the laser frequency shaking relative to the cavity resonance or to power fluctuations—that limits the measurement resolution to $\langle \Delta S_z^2 \rangle = 100(10)$ at high readout power. Fixing these parameters and adding the expected noise from Raman scattering for $2S_0 = 3.1 \cdot 10^4$ atoms (dashed line) gives good agreement with the measurement noise observed with that many atoms in the resonator (filled circles). For a more detailed analysis of the measurement performance and of the technical factors that limit it, see reference [53]. For the squeezing experiments with $2S_0 \approx 3 \cdot 10^4$, the optimum readout photon number is around $\langle p_t \rangle \approx 10^6$ and yields a measurement variance $\langle \Delta S_z^2 \rangle \approx 10^3$. Note, however, that for smaller atom numbers the noise from Raman scattering does not intrude so quickly and resolutions down to $\langle \Delta S_z^2 \rangle \approx 10^2$ (20-atom resolution) are achievable.

5.6 Measurements of Atomic Coherence

As discussed in chapter 2, spin squeezing depends not only on the suppression of atomic projection noise but also on the preservation of atomic coherence, so that the signal-to-noise ratio is improved. It is therefore crucial to measure the length of the




Figure 5-9: Ramsey signal at fixed time, obtained by sweeping the phase of the second Ramsey $\pi/2$ pulse through 2π . In this case the curve shows an available signal contrast of 90%.

mean collective Bloch vector, which corresponds to the signal amplitude obtainable in a Ramsey-type atomic clock. Figure 5-8 shows the decaying sinusoidal signal from such a clock as a function of the Ramsey precession time T_R , showing a coherence lifetime of 10(1) ms, limited by inhomogeneous shifts of the transition frequency by the trap potential. Most of the squeezing experiments to be described take place in under 1 ms, so that this background decoherence rate is slow enough to be unimportant. The contrast is instead limited by decoherence mechanisms introduced by the squeezing procedure itself, primarily inhomogeneous dephasing of the atomic spins by the AC Stark potential of the probe light and of the light used to lock the probe laser to the resonator.

To measure the contrast at a given moment in time—such as immediately after the preparation of a putative squeezed state—it is convenient to record the Ramsey signal not as a function of time but rather as a function of the variable phase of the second Ramsey pulse. Sweeping this phase through 2π without altering the timing of the sequence provides a single-cycle sinusoidal fringe (figure 5-9) whose amplitude indicates the available signal contrast, provided only that the Bloch vector of the




Figure 5-10: Measuring the contrast of a state with large collective phase fluctuations by Rabi nutation. In a Rabi experiment the phase fluctuations appear as fluctuations in the amplitude of the sinusoidal nutation signal. Each open circle is the result of a single Rabi experiment. The dashed line indicates the envelope of the data (with amplitude $\langle |\mathbf{S}| \rangle / S_0 = 0.48$ in this case), the filled circles and solid lines the average signal and corresponding sinusoidal fit respectively (yielding a metrologically relevant contrast $|\langle \mathbf{S} \rangle| / S_0 = 0.28$ in this example). Note that some of the data appear to have an inverted amplitude, indicating phase fluctuations greater than $\pi/2$.

state to be characterized lies somewhere in the equatorial plane ($S_z \approx 0$) before the final Ramsey pulse. This condition is generically satisfied at the end of the squeezing procedures to be discussed here, which work with states in the vicinity of the \hat{x} pole of the Bloch sphere. This fixed-time Ramsey technique was used to verify the signal contrast for the measurement squeezing experiments of chapter 6.

Rather than sweeping the phase of the final Ramsey pulse, it is also possible to measure the signal contrast by sweeping its length and recording the resulting sinusoidal Rabi nutation signal. Because the microwave Rabi frequency is much higher than typical decoherence times in the experiment, this measurement is still an effectively instantaneous measurement of the available signal contrast at the moment of the final microwave pulse.

To see the full contrast, this method requires that the phase of the microwaves used

to drive the rotations be perpendicular to the mean Bloch vector; for any other phase choice the amplitude of the resulting oscillations in S_z will be reduced. This allows a straightforward illustration of the distinction between the length of the mean spin $|\langle \mathbf{S} \rangle|$, which dictates the signal amplitude relevant to metrology and to calculating the squeezing parameter ζ , and the radius of the Bloch sphere itself $\langle |\mathbf{S}| \rangle$, which is important for modeling the behavior of the system (see chapter 7). The contrast measurement shown in figure 5-10 was taken under conditions with large fluctuations of the phase of the collective Bloch vector, so that the amplitude of the sinusoidal Rabi oscillations varied dramatically from shot to shot. The envelope of the curve, indicated by the dashed line, corresponds to the maximum amplitude and thus to the radius of the Bloch sphere $\langle |\mathbf{S}| \rangle / S_0$, used in chapter 7 when verifying the model of chapter 3, while the metrologically relevant signal contrast is only the mean amplitude $|\langle \mathbf{S} \rangle| / S_0$ indicated by the solid line. This was the contrast-verification technique used in the cavity feedback squeezing experiments of chapters 7 and 8.

5.7 Magnetic-Field Calibration

Control of magnetic fields is an important, though often implicit, requirement throughout this work. The same coils and current controllers used in references [44, 45] were used for all the experiments described in this thesis, but the sensitive atomic-state detection techniques described in section 5.5 allow much improved calibrations of the effect of those coils by using the atomic cloud as an *in situ* magnetometer.

With atoms optically pumped into $|F = 1\rangle$, a microwave drive tuned to one of the magnetic-field sensitive transitions (such as $|F = 1, m_F = 1\rangle \leftrightarrow |F = 2, m_F = 2\rangle$) produces a population transfer to $|F = 2\rangle$ observable as a shift of the cavity resonance frequency. Repeating this experiment and measuring the frequency of maximum population transfer as a function of the current in one of the magnetic bias coils gives




Figure 5-11: Typical plot of hyperfine transition frequency as a function of bias field used for magnetic field calibration. The upper curve is the 2.10 MHz / G $|F = 1, m_F = 1\rangle \leftrightarrow |F = 2, m_F = 2\rangle$ transition and the lower curve shows part of the trace for the 0.70 MHz / G transitions $|F = 1, m_F = 1\rangle \leftrightarrow |F = 2, m_F = 0\rangle$ or $|F = 1, m_F = 0\rangle \leftrightarrow |F = 2, m_F = 1\rangle$.

a hyperbolic curve of the form

$$\nu = \nu_0 \pm g_f \mu_B (m_{F=2} + m_{F=1}) \sqrt{B_0^2 + \left(\frac{dB}{dI} (I - I_0) \right)^2} \quad (5.15)$$

where ν_0 is the transition frequency at zero field (the hyperfine transition frequency modified by the scalar Stark shift from the trap), $g_f \mu_B \approx 0.70$ MHz/G is the linear Zeeman shift between adjacent substates, B_0 is the field orthogonal to the axis of the coil under study, dB/dI is the conversion from coil current to magnetic field at the atoms' location and I_0 is the current for which the field along the coil axis vanishes. One such sweep for each of three perpendicular coils yields sufficient information to null the magnetic field at the location of the atomic cloud—by setting $I = I_0$ for each axis—or to give it any convenient value, since the conversion from current to local field dB/dI is known.

Chapter 6

Squeezing by Quantum

Non-Demolition Measurement

This chapter presents an experimental demonstration of the measurement-based squeezing scheme discussed in section 3.2. To recapitulate: provided that S_z is a conserved quantity, or in other words that the populations of the two clock states are stable, two successive measurements of S_z on a given atomic ensemble must give the same result to within their measurement uncertainty. After the first (squeezing) measurement, the result of an ideal second (readout) measurement of S_z can be predicted with an uncertainty no greater than that of the squeezing measurement. The quantum state of the system after the squeezing measurement, conditioned on its result, can therefore have a spin variance $\langle \Delta S_z^2 \rangle$ no greater than the squeezing measurement's uncertainty. If this uncertainty is smaller than the projection noise for an unentangled ensemble, and if the coherence (signal contrast) of the system after the squeezing measurement remains sufficiently large, the state produced by this measurement and conditioned on its outcome can be squeezed [17].

The experimental sequence used to demonstrate this squeezing mechanism is essentially the same as that described in section 5.5 for characterizing the effectiveness of the atomic state readout. An atomic sample is loaded into the dipole trap as described in section 5.1, then optically pumped into the $|F = 1, m_F = 0\rangle$ state and purified as discussed in section 5.3, producing the trivial coherent spin state with all




Figure 6-1: Projection noise as a function of atom number. The empty triangles show the variance in S_z for a single measurement of a state prepared near $|\text{CSS}\hat{\mathbf{x}}\rangle$, while the open circles show the measurement variance obtained by comparing pairs of measurements on independently prepared atomic states. Both lie somewhat above the projection noise limit for unentangled atoms indicated by the dashed line, due to technical noise and imperfect state preparation. The filled circles show the measurement variance for two successive measurements of the same atomic state, as discussed in section 5.5. The readout photon number for all these data was $\langle p_t \rangle \approx 5 \cdot 10^5$.

atoms spin-down. A microwave $\pi/2$ pulse rotates the Bloch vector into the equatorial plane to produce a state near to $|\text{CSS}\hat{\mathbf{x}}\rangle$. A two-pulse measurement, as in section 5.5, serves to suppress conditionally the transverse spin variance $\langle \Delta S_z^2 \rangle$ of the state, and a second two-pulse measurement verifies this suppression. Optical pumping into $|F = 1, m_F = 0\rangle$ and a $\pi/2$ pulse then reinitialize the atomic ensemble for an S_z measurement on an independently prepared state. Finally, the atoms are optically pumped back and forth from $|F = 2\rangle$ to $|F = 1\rangle$ to check the total atom number and then discarded. In order to be able to calculate variances of the various measurements, the whole cycle is repeated 100 times for each set of experimental parameters, i.e. for each point in the plots.

To begin, consider the unconditional spin variance $\langle \Delta S_z^2 \rangle$ of the initial state. The open triangles in figure 6-1 show the variance of the first measurement of S_z ,

immediately after the microwave pulse brings the Bloch vector into the equatorial plane. As expected, this variance grows approximately linearly with the length of the Bloch vector S_0 , and lies slightly above the projection noise limit $S_0/2$ indicated by the dashed line. The excess variance is due in part to the noise of the measurement itself (note the non-zero variance even for $S_0 = 0$, with no atoms in the resonator), and partly to slow drifts in the microwave power of 0.4% over the course of the data-taking.

Another way of quantifying the unconditional spin variance, one that is closer to the method used later in establishing the conditional spin noise suppression, is to compare the outcomes of S_z measurements made on independently prepared coherent states. Since the projection noise seen in the two measurements is then uncorrelated, it appears as a fluctuation in the difference between them. Taking the variance of the difference between the measured S_z for the first and second coherent state preparations in the sequence described above and dividing it by two to obtain the variance of each preparation individually yields the open circles in figure 6-1, again slightly above the ideal projection noise.

The ability to prepare states with the same S_z fluctuations as the coherent spin state $|\text{CSS}\hat{\mathbf{x}}\rangle$ is not essential. The spin squeezing mechanisms to be demonstrated suppress spin fluctuations regardless of their origin and can be used to overcome technical noise in the state preparation just as well as quantum noise. Moreover, when calculating the squeezing achieved in the experiments, it is the ideal projection noise limit, calculated from the independently measured atom number $2S_0$, that must be used as a benchmark, and not merely the observed noise of the best initial states that can be prepared experimentally. Nonetheless, it is reassuring to see that preparation procedures supposed to produce states near $|\text{CSS}\hat{\mathbf{x}}\rangle$ do nearly reproduce that state's expected $\langle \Delta S_z^2 \rangle$, as this corroborates the correct calibration of the atom number and projection noise limit and indicates that the experiment is operating in a regime where the fundamental quantum projection noise is large compared to technical effects.

The first requirement for conditional squeezing by measurement is that the measurement uncertainty be substantially less than the fundamental projection noise of

the coherent state. The filled circles in figure 6-1 demonstrate this feature. They show the variance in the discrepancy between the squeezing and readout measurements of S_z performed on the same state, divided by two to obtain the variance of a single measurement. Because the quantum projection noise of the initial state is common to both measurements, it does not appear in the difference between them. As a result, the variance is only weakly dependent on S_0 , and for a large atomic ensemble it lies well below the projection noise limit. In this regime the squeezing measurement provides a prediction of the outcome of the readout with much less uncertainty than would be possible if the readout were seeing the quantum fluctuations in S_z of a coherent state of unentangled atoms.

In order to show that the quantum state produced by the squeezing measurement has suppressed fluctuations, one must be careful to account for possible discrepancies between the actual S_z value of the atomic state and the well-correlated results of the squeezing and readout measurements. Scattering into free space and microwave pulse errors, for example, can introduce correlated errors that would not appear in the difference between the two measurements but would nonetheless reduce the squeezing measurement's value as a predictor of the atomic state. References [12] and [53] contain a detailed analysis of the normalized variance $q = 2\langle\Delta S_z^2\rangle/S_0$ of the state produced by the measurement, bounding the effects of the technical imperfections alluded to above, the possibility that the two measurements might not be identical so that the uncertainty ascribed to the squeezing measurement might not be just half the uncertainty on the difference between them, and including the additional information available from the known initial state before the squeezing (recall the leading 1 in the denominator of equation 3.19).

Figure 6-2 shows the resulting normalized transverse spin variance q of the state after the squeezing measurement, plotted as a function of the transmitted photon number $\langle p_t \rangle$ used in the measurements (open squares). Aside from this normalization, the data and model (dotted line) are similar to those in the measurement noise plot of figure 5-7, but with a finite variance even in the limit of weak measurement (small $\langle p_t \rangle$) thanks to the known initial state. As in the measurement noise analysis of section 5.5,




Figure 6-2: Spin noise suppression q (open squares, dotted line), signal contrast C (filled circles, solid line) and metrological squeezing parameter ζ (filled triangles, chaindotted line) as a function of readout photon number $\langle p_t \rangle$ for an atom number $2S_0 = 3.3(2) \cdot 10^4$. The dashed line at 1 indicates the level that all these parameters would have for the ideal unentangled state $|\text{CSS}\hat{x}\rangle$.

the minimum measurement variance occurs for a photon number of $\langle p_t \rangle \approx 10^6$. For this measurement strength the noise is suppressed by a factor of $q = 0.12(2)$. This 9(1) dB of noise suppression does not, however, imply any squeezing, because the poor contrast at this photon number costs as much in signal as the measurement has earned in noise.

The solid line and open circles show the contrast C of a Ramsey measurement (section 5.6) made immediately after the first (squeezing) measurement of S_z , indicating the fraction of the ideal signal amplitude that would be available to a clock that used the squeezed state as its input. The tradeoff between reduced noise and shrinking signal gives the Wineland metrological squeezing parameter ζ (filled triangles and chaindotted line), which is optimized for a transmitted photon number $\langle p_t \rangle = 3 \cdot 10^5$. There the remaining signal contrast is 54% and the measurement suppresses the projection noise by a factor $q = 0.20(3)$, yielding a metrological squeezing parameter of $\zeta = 0.7(1)$ or a signal-to-noise ratio improvement of 1.5(5) dB. This is substantially

worse than the 14 dB of squeezing that might have been expected based on the model of section 3.2 and the collective cooperativity $S_0\eta = 1.6 \cdot 10^3$. This is partly because of the imperfect performance of the transmission measurement, with a variance 5 times larger than the photon shot noise limit (section 5.5), and partly due to poor signal contrast: of the 7.0 dB of noise suppression afforded by the measurement at this photon number, 5.4 dB are needed simply to compensate for the reduction in signal strength relative to the ideal coherent state. Some of this contrast loss is due to inhomogeneous phase shifts from the probe light; a side effect of the squeezing measurement that could be avoided in principle if the atoms were uniformly coupled to the probe field.

Much of the contrast loss, however, has nothing to do with the squeezing process itself: even without any probe light, the contrast of the state produced is 68% because of inhomogeneous dephasing by the locking light. This is a significant technical handicap for squeezing: with such a reduced signal the projection noise variance must be suppressed by more than a factor of two just to reach the same sensitivity as the unentangled coherent state $|\text{CSS}\hat{x}\rangle$. Reducing this handicap was the prime motivation for increasing the lock sideband's detuning for the cavity feedback squeezing experiments (section 4.3): reducing the purely technical decoherence due to the locking light brings the signal-to-noise ratio of the initial state closer to the SQL and makes it easier for the squeezing to produce beyond-SQL sensitivity. Increasing the initial contrast from 68% to the 80% obtained in later experiments would, without any other technical improvements, have allowed for an additional 1.4 dB of metrological squeezing in this experiment.

Chapter 7

Squeezing by Cavity Feedback

While conditional squeezing by measurement can overcome the SQL, as demonstrated in the previous chapter, it still leaves room for improvement. In particular, it is desirable to have a deterministic squeezing scheme, relying on predictable dynamics to generate a specific squeezed state rather than on measurement projection to collapse the system randomly onto one of a family of squeezed states. Such a deterministic squeezing procedure is conceptually tidier than a conditional scheme, since it can be understood solely in terms of the state and evolution of the system itself, without having to keep track of the information recorded by the measurement apparatus. It is also practically advantageous since the squeezing performance in a deterministic scheme is independent of the quantum efficiency of the detection apparatus, which limits the performance of measurement-based schemes where only detected photons contribute to the squeezing.

This chapter discusses the experimental demonstration of cavity feedback squeezing, the deterministic scheme analyzed in section 3.3. There, the squeezing mechanism was presented as an effective S_z^2 interaction mediated by the cavity light field which deforms the uncertainty region on the Bloch sphere of an initial coherent state, accompanied by decoherence due to classical uncertainty on the evolution rate when the number of photons interacting with the resonator is unknown. An alternative picture for the squeezing dynamics is to consider a squeezing measurement of the kind considered in chapter 6 but, instead of using the information in the exiting light field

to project the atomic ensemble into a squeezed state, to use the information directly stored into each atom's state by the light field in the form of an S_z -dependent phase shift. In this picture the cavity light field acts as an information bus between the atoms, allowing each to adjust its state as a function of that of its peers. The decoherence that increases the area of the uncertainty region arises, in this picture, from the information leaked to the observer—and thereby lost to the atoms—by the exiting light field.

Consequently, it should not come as a surprise that the experimental sequence used to demonstrate cavity feedback squeezing is, aside from a few technical improvements, close to that used for measurement-based squeezing. After loading an atomic sample into the dipole trap and cooling it as described in section 5.1, the atomic ensemble's initial size is checked by optically pumping it from $|F = 2\rangle$ to $|F = 1\rangle$. The ensemble is then prepared in a state near $|\text{CSS}\hat{\mathbf{x}}\rangle$ by optical pumping into $|\downarrow\rangle = |F = 1, m_F = 0\rangle$ followed by a $\pi/2$ microwave pulse (sections 5.3 and 5.4). The same pair of probe pulses separated by a composite π pulse is used for the cavity feedback squeezing as for measurement, and for the same reasons: to reduce inhomogeneous dephasing of the atoms and because the net phase shift imparted to the atoms, being proportional to the intensity difference between the two pulses caused by the intervening inversion of S_z , is first-order insensitive to the effects of the spectator atoms in $m_F \neq 0$ and to fluctuations in overall laser power. The Bloch vector is then rotated by an angle β about its mean direction along $\hat{\mathbf{x}}$ by a short microwave pulse so that a subsequent measurement of S_z can be used to measure any of its transverse spin components. Since there is no need for this final readout measurement to preserve the atomic coherence, it is performed with around 10^6 transmitted photons to obtain optimum resolution (section 5.5). The whole procedure of state initialization, squeezing and readout is repeated 10 times in a row before the total atom number is checked once more and the ensemble is discarded. Repeating this experimental cycle 10 times therefore gives 100 measurements of the transverse spin component for a given squeezed state preparation procedure, from which the variance of the transverse spin fluctuations can be extracted.




Figure 7-1: Normalized transverse spin variance as a function of the angle by which the state is rotated between the cavity feedback squeezing and the S_z measurement. Data shown are for shearing parameters $\mu = 0$ (circles), $\mu = 7.1 \cdot 10^{-5}$ (squares), $\mu = 5.2 \cdot 10^{-4}$ (triangles), $\mu = 1.8 \cdot 10^{-3}$ (bowties). The curves are sinusoidal fits, distorted by the log-log scale used to emphasize the small-angle behavior.

Representative variance data is plotted in figure 7-1 as a function of β for $S_0 = 1.6 \cdot 10^4$ and transmitted photon numbers of $\langle p_t \rangle = 0, 2.4 \cdot 10^3, 1.8 \cdot 10^4, 6.2 \cdot 10^4$, corresponding to shearing parameters $\mu = 0, 7.1 \cdot 10^{-5}, 5.2 \cdot 10^{-4}, 1.8 \cdot 10^{-3}$. The dependence of variance on rotation angle is generically sinusoidal for any distribution of transverse spin components, and from a sinusoidal fit to the observed variance the minimum and maximum (normalized) variances q_{\min} and q_{\max} and the angle β_0 that minimizes the observed variance can all be extracted. It is convenient, and consistent with the model presented in section 3.3, to think of the uncertainty region as an ellipse whose orientation, semimajor and semiminor axis lengths are given respectively by β , $\sqrt{q_{\max}}$ and $\sqrt{q_{\min}}$ in units of the RMS width of the coherent spin state. As the shearing parameter is increased, the minimum variance dips lower and appears at a smaller β , while the maximum variance increases; geometrically the uncertainty ellipse gets narrower and longer, and its long axis approaches the horizontal as it is sheared out, in qualitative agreement with the intuitive geometrical picture of figure 3-6.

Section 3.3 derives quantitative predictions for these parameters in an idealized setting. Testing them against the actual experiment requires the following adjustments. The effects of scattering into free space are again complicated by the presence of additional spectator Zeeman levels into which the atoms can scatter; in calculating the variance and covariance of equations 3.59 to 3.61, one must therefore use the variances and correlations of \bar{S}_z and S_{zf} given in section 5.5 to account properly for the connection between the average cavity shift during the squeezing and the actual collective S_z value of the atoms in the $m_F = 0$ clock states at the end of the squeezing. The length of the Bloch vector being limited by inhomogeneous dephasing rather than by scattering into free space, the ratio S/S_0 is measured experimentally from the envelope of a Rabi oscillation measurement (figure 5-10) and inserted into the theory by hand. The observed variances include a contribution from the finite resolution of the readout measurement (section 5.5), which must be added to the model predictions for q_{\min} and q_{\max} . Finally, the fluctuations in intracavity light power contain a technical contribution from fractional light noise, which must be determined experimentally by measuring the fluctuations in transmitted photon number with no atoms present in the cavity. This is just the fractional noise contribution obtained when characterizing the atomic readout in section 5.5. There it was expressed as an effective S_z variance $\langle \Delta S_z^2 \rangle = 100(10)$. Converting this into a fractional intensity noise by multiplying by ϕ_0^2 and normalizing to photon shot noise by multiplying by the number of photons interacting with the resonator mode yields $\gamma = 1 + 100\mu$, with the leading 1 representing the photon shot noise itself.

The resulting parameter-free predictions of maximum variance q_{\max} , minimum variance q_{\min} and angle for minimum variance β_0 are plotted as a function of shearing parameter μ as dotted curves in figure 7-2, displaying fairly good agreement with the values extracted from the sinusoidal fits over a range of two orders of magnitude in the strength of the shearing interaction. Note, in particular, the effects of the curvature of the Bloch sphere visible as an increase in q_{\min} for large shearing parameter. For comparison, the solid curves show the predictions of the model when all technical effects are removed by setting $S = S_0$ and $\gamma = 1$ and assuming a perfect readout




Figure 7-2: Modeling the observed behavior of maximum and minimum normalized transverse spin variances q_{\max} and q_{\min} and the rotation angle for minimum variance β_0 as a function of shearing parameter μ . The dotted lines show parameter-free predictions including known technical effects, while the solid lines are model predictions for a technically perfect system. Data points show the observed variances and angles extracted from sinusoidal fits like those of figure 7-1, with statistical error bars.




Figure 7-3: Noise suppression (open squares and dotted curve), contrast (open circles and solid curve) and Wineland metrological squeezing (filled triangles and chain-dotted curve) obtained by cavity feedback squeezing as a function of photon number. Dashed line indicates the SQL.

which adds no variance to q . These show that the minimum variance q_{\min} is strongly affected by the readout noise, which accounts for more than half of the observed variance in the range $\mu = 10^{-3} \dots 10^{-2}$, while the antisqueezing q_{\max} and ellipse orientation β_0 are relatively insensitive to technical effects.

How effective is cavity feedback squeezing as a tool to improve signal-to-noise ratio? Taking the minimum transverse variance q_{\min} achieved by rotating the state through an angle β_0 before the final readout of S_z and comparing it to the average signal contrast obtained in a Rabi nutation experiment immediately after the squeezing light is applied yields a Wineland metrological squeezing parameter of $\zeta = 0.34(5)$, or a signal-to-noise ratio improvement of 4.6(6) dB, for the best shearing parameter $\mu = 1.2 \cdot 10^{-3}$ where $q_{\min} = 0.21(3)$ and the signal contrast is 78%. Though the performance of this scheme is theoretically equivalent to that of the measurement-based technique described in the previous chapter, in practice it fares substantially better, providing an extra 3 dB of demonstrated signal-to-noise ratio improvement. Partly this is due to intervening technical improvements; for instance, the contrast

is 80% before squeezing light is applied in this work as against less than 68% in the measurement squeezing experiments, thanks to the modified laser locking scheme. Much of the improvement is, however, attributable to the deterministic nature of the squeezing process: not needing to compensate for the photons lost in the detection path, it is possible to suppress the transverse spin noise with weaker probe pulses (compare the photon number axes in figures 7-3 and 6-2), thus avoiding dephasing-induced contrast loss. Another way of understanding this improvement is to note that the detection operates at a noise level 5 times higher than the photon shot noise (section 5.5), while $\gamma = 1.12$ at the optimum squeezing point so that the atoms see only 12% more intensity fluctuations than photon shot noise.

While there are still 2 dB of squeezing to be gained by avoiding the dephasing from the locking light that reduces the signal contrast to 80% before any squeezing light is applied, the main technical obstacle to further squeezing by this method is currently readout performance, as indicated above. That is, based on the agreement with the model shown in figure 7-2, it is probably the case that the squeezing procedure is preparing states that are more strongly squeezed than the apparatus is capable of detecting. As technical improvements to the readout performance are certainly possible (some are suggested in chapter 9), this bodes well for the progress of cavity feedback squeezing as a method of producing states with reduced quantum fluctuations.

Chapter 8

Squeezing in an Atomic Clock

Having demonstrated the possibility of producing squeezed spin states on a clock transition, it is now time to consider their applicability. Can squeezed states actually improve the performance of a real metrological experiment and, if so, under what conditions? For concreteness, this question will be addressed in the context of a Ramsey-type atomic clock, using a squeezed input state to reduce the quantum uncertainty on the clock's initial phase. Such reduced phase uncertainty can be helpful whenever the initial phase uncertainty is the dominant source of fluctuations on the clock output signal. However, Huelga *et al.* pointed out, in reference [60], that initial phase uncertainty should never be the fundamental performance limit for an atomic clock. For any given initial phase uncertainty, the clock's stability is optimized by extending the Ramsey interrogation time—dividing the phase uncertainty by a larger and larger accumulated phase—until environmental decoherence contributes as much to the noise of the clock output signal as the initial phase fluctuations, at which point the initial state's phase uncertainty is no longer the dominant uncertainty. They further showed that entangled states and squeezed states, because they are generically more sensitive to environmental decoherence than the unentangled $|\text{CSS}\hat{x}\rangle$, require shorter Ramsey precession times to reach this limit, and that this reduced interrogation time cancels the benefit from the reduced phase fluctuations of the initial state. The result is that entanglement in general, and squeezing in particular, even if it improves the phase sensitivity of the clock's initial state, yields little or no improvement

in actual clock performance.

This assumes, however, that the Ramsey precession time can be extended freely until the initial phase uncertainty ceases to be the dominant limit to clock stability. In practice, other limits to the precession time are common. For example:

- The interrogation time of a fountain clock is limited by the ballistic flight time of the atoms in a reasonably sized apparatus. Without the option to extend interrogation time and with systematic efforts to suppress technical noise sources, fountain clocks have been operating at or near the SQL, limited by the quantum phase uncertainty of the initial state, for more than a decade¹ [8, 16].
- In optical clocks using trapped atoms, the interrogation time is limited by the phase coherence time of the laser which acts as the local oscillator: the laser phase uncertainty must not reach π , to avoid losing track of the number of optical cycles elapsed during the Ramsey precession. The case of clocks limited by local oscillator dephasing was analyzed in references [61, 62], where it was shown that squeezing could produce substantial improvements in clock stability depending on the spectral properties of the local oscillator noise.
- More generally, there may be external requirements on the clock readout rate: if, for technological reasons, an atomic clock signal is required every second—perhaps as part of a servo loop to stabilize a flywheel oscillator—then a squeezed input state may be useful to improve the phase stability at one second. The fact that a clock operating without squeezing could in principle have reached the same stability with a longer interrogation time is then irrelevant.

The next two sections address the question of squeezed state usefulness experimentally. Section 8.1 studies the fragility of a squeezed state, observing its gradual degradation under the action of environmental perturbations to the atomic phase

¹Sadly, fountain clocks are in other respects a poor candidate system for the use of spin squeezing because most of the atoms launched at the start of the Ramsey sequence are lost due to the lateral expansion of the cloud. Even if the whole cloud were squeezed before being launched into the fountain, once the lost atoms were traced over, the remaining subensemble used for the final readout would still display projection noise.

and demonstrating that the observed behavior can be understood in terms of simple uncertainty addition. Section 8.2 examines the performance of a fully operational atomic clock using a squeezed input state, demonstrating by a measurement of its Allan deviation spectrum that it can provide a stability well beyond the SQL. Both sections discuss work that was previously published in reference [15].

8.1 Lifetime of Squeezed States

To measure the lifetime of squeezed states, squeezed ensembles of $2S_0 = 3 \cdot 10^4$ atoms were repeatedly prepared by cavity feedback squeezing, yielding initial states with $\zeta^{-1} \approx 4$ dB. The squeezed states were then rotated with a microwave pulse so that the direction of maximum angular sensitivity was oriented either in the equatorial plane (phase squeezing) or along S_z (number squeezing). After a variable hold time T_R , the surviving squeezing was measured by verifying both the fluctuations of the initially squeezed quadrature and the surviving signal contrast.

For the phase-squeezed initial state, this corresponds to evaluating the noise in a Ramsey-type atomic clock: an ensemble of atoms is prepared with a known (in fact squeezed) initial phase, the phase is allowed to precess for a time T_R , and the final phase is read out by rotating the final state with a microwave $\pi/2$ pulse and measuring the resulting population difference. Consequently, any technical noise which affects the stability of an atomic clock also degrades the observed squeezing by increasing the fluctuations seen in the squeezed quadrature after the hold time T_R . In the present case, the dominant noise term is the fluctuating magnetic field, which perturbs the atomic transition frequency. These fluctuations are slow compared to $T_R \sim \text{ms}$, but fast compared to the experimental cycle time of 9 s, so that they lead to shot-to-shot frequency fluctuations and a phase variance which scales quadratically with interrogation time: $\langle \Delta\phi^2 \rangle = \langle \Delta\omega_a^2 \rangle T_R^2$. The result is shown in the left panel of figure 8-1. The initial 4 dB of squeezing persists for around 200 μs , then the phase uncertainty increases due to the magnetic field fluctuations until it is no longer below the SQL after 500 μs . A fit to the model $\zeta(T_R) = \zeta(0) + 2S_0 \langle \Delta\omega_a^2 \rangle T_R^2$ (figure 8-1,




Figure 8-1: Lifetime of squeezed spin states under different conditions. When the phase quadrature is squeezed (left panel, filled triangles) the state is vulnerable to phase noise in the form of shot-to-shot fluctuations in the atomic precession frequency that set a noise floor (solid line). The same noise affects the phase of an initially unsqueezed state (open triangles), though the effect is only noticeable at later times due to the poorer initial phase sensitivity in the absence of squeezing. The dotted and chain-dotted lines are fits to the noise model described in the text. For states insensitive to frequency fluctuations, either because the population difference was squeezed rather than the phase (lower right panel, filled triangles), or because a spin echo was used to cancel the effect of fluctuating frequency shifts (filled bowties), the squeezed state can be much longer-lived, lasting until the shortening of the Bloch vector (loss of contrast, upper right panel) degrades the angular sensitivity of the state back to the SQL (dashed line).

chaindotted line) yields a frequency variance $\langle \Delta\omega_a^2 \rangle = (2\pi \cdot 1.3 \text{ Hz})^2$ (figure 8-1, solid line). Of course, the same frequency fluctuations afflict initially unsqueezed states as well (figure 8-1, open triangles and dotted line in the left panel). The effect takes a little longer to appear (500 μs instead of 200 μs) simply because the unsqueezed state has poorer phase sensitivity to begin with.

The squeezed state's lifetime can be extended by making it insensitive to the magnetic field fluctuations discussed in the previous paragraph. One way of doing this is to insert a microwave π pulse halfway through the precession time T_R . The result is a spin echo sequence in which the effect of frequency fluctuations slow compared to T_R cancels out. As shown in the solid bowties in the lower right panel of figure 8-1, the phase can then remain squeezed for several milliseconds. Even longer squeezed lifetimes, up to 5 ms, can be obtained by orienting the squeezed quadrature along S_z , so that phase uncertainty does not affect it at all (filled triangles). Since S_z -changing events (atom loss and Raman scattering) are rare as long as the probe light is off (which it is during the hold time T_R), the uncertainty on the squeezed S_z quadrature essentially does not increase with time; instead it is loss of signal contrast (figure 8-1, upper right panel) due to dephasing between the atoms which eventually degrades the angular sensitivity of the state back to the SQL.

8.2 Allan Deviation Measurement

The standard figure of merit for quantifying the stability of clocks is the Allan deviation [8]

$$\sigma(\tau) = \sqrt{\frac{\sum_{k=1}^N (y_k - y_{k-1})^2}{2N}} \quad (8.1)$$

where $y_k - y_{k-1}$ is the change in clock frequency between two successive averaging bins of duration τ . This frequency change is normalized to the clock transition frequency to obtain a dimensionless fractional fluctuation. Conceptually, $\sigma(\tau)$ measures the jitter in the clock's output frequency at the time scale τ . For an atomic clock operating

at the SQL, on a transition of frequency ω_a with a Ramsey interrogation time T_R and an experiment cycle time T_c (i.e. with a duty factor T_R/T_c), the clock output fluctuations due solely to quantum projection noise are uncorrelated from shot to shot and the fractional Allan deviation is simply [47]

$$\sigma(\tau) = \frac{1}{\omega_a T_R} \sqrt{\frac{T_c}{2S_0 \tau}}. \quad (8.2)$$

This is just the phase uncertainty of $(2S_0)^{-1/2}$ normalized to the total accumulated phase $\omega_a T_R$ in the Ramsey precession time to give the single-shot fractional stability, then divided by the square root of the number of cycles in a time bin $\sqrt{\tau/T_c}$. A measurement of the Allan deviation of a clock using squeezed input states to overcome this limit is a concrete demonstration of the potential benefits of spin squeezing to metrological performance.

To measure the Allan deviation, the clock's output must be compared to another equally stable clock over a period $N\tau$, with $N \gg 1$ to ensure good statistics. Of course, measuring fluctuations in the clock frequency can only be done by comparison to another clock of equal or better stability; as will be seen, the absolute stability of the clocks to be demonstrated is fairly poor, so that the rubidium frequency standard used as a reference for the microwave synthesizers (section 4.2) is for present intents and purposes an exact clock. The phase angle measured at the end of a Ramsey sequence is essentially a comparison of the atomic clock frequency with this microwave frequency standard. To obtain the Allan deviation, one can therefore run a Ramsey sequence repeatedly over many hours, convert the observed S_z output values into fractional frequency offsets between atomic clock and microwave frequency according to

$$y = \frac{S_z}{CS_0} \frac{1}{\omega_a T_R} \quad (8.3)$$

and then average those readings over time bins of duration τ , computing the fluctuations between successive bins as in equation 8.1.




Figure 8-2: Allan deviation with (solid triangles) and without (open triangles) squeezing. Error bars are statistical. The dashed line indicates the SQL at $1.85 \cdot 10^{-9} \text{ s}^{1/2}/\sqrt{\tau}$, while the chaindotted line is at $1.1 \cdot 10^{-9} \text{ s}^{1/2}/\sqrt{\tau}$.

The open triangles of figure 8-2 show the resulting measurement of Allan deviation for a conventional Ramsey sequence using an unsqueezed input state (approximately $|\text{CSS}\hat{x}\rangle$) of $2S_0 = 3.54 \cdot 10^4$ atoms, a 9 s experiment cycle time, 200 μs Ramsey interrogation time and with a signal contrast that was not measurably less than 1. As expected, the fluctuations average down as $\tau^{-1/2}$, though they are somewhat larger than the $1.85 \cdot 10^{-9} \text{ s}^{1/2}/\sqrt{\tau}$ expected for a clock operating at the SQL with these parameters (figure 8-2, dashed line), due to technical imperfections in the preparation of the initial state.

For a clock operating with a cavity-feedback-squeezed input state as described in the previous chapter, with a contrast reduced to 81% and otherwise identical parameters to the unsqueezed clock above, the Allan deviation starts at $1.1 \cdot 10^{-9} \text{ s}^{1/2}/\sqrt{\tau}$ (figure 8-2, filled triangles and chaindotted line). In variance or integration time, this is an improvement by a factor of 2.8(3) relative to the SQL, which persists for averaging times up to a minute. Thereafter, having reached a fractional stability of 10^{-10} or an absolute frequency sensitivity of 0.7 Hz, the clock becomes sensitive to slow drifts of the magnetic field at the level of a few hundred microgauss. Whether or not the input state is squeezed, improving the fractional stability beyond this level requires

reducing the drifts in the frequency of the transition used for the Ramsey sequence, either by stabilizing the magnetic field with shielding or feedback or by reducing the clock transition’s sensitivity to magnetic fields as discussed in section 4.6.

While the squeezed clock’s short-term stability improvement relative to the SQL—by nearly a factor of 3—is striking, its absolute performance is not. This is not due to poor phase sensitivity: the few-milliradian phase stability of the clocks demonstrated here is comparable to that of current fountain clocks [16]. Rather, it is a result of the 200 μs Ramsey interrogation time used in these experiments, which is much shorter than the $\sim 1\text{ s}$ interrogation time used in typical primary frequency standards. If environmental perturbations on the transition frequency could be controlled at the level of 100 μHz , then the squeezed phase precision could be preserved for an interrogation time of a second, enough to allow a short term stability $\sigma(\tau) \approx 2 \cdot 10^{-13} \text{ s}^{1/2}/\sqrt{\tau}$ comparable to that of current realizations of the SI second [63].

Chapter 9

Conclusion

This thesis presents two methods of exploiting the collective interaction of a dilute atomic ensemble with an optical resonator to produce squeezed states with suppressed quantum fluctuations of an atomic pseudospin variable. The first method is a practical implementation of conditional squeezing by quantum non-demolition measurement, as proposed in reference [17]. Using the optical resonator to reduce the harmful effects of scattering into free space, dispersive measurement is shown to be capable of projecting the atomic ensemble onto a state with 1.5(5) dB of metrological squeezing, according to Wineland's strict criterion of signal-to-noise ratio improvement. The second method employs a new scheme using the light field in the resonator to mediate an interaction between the atoms, producing a deterministic shearing deformation of the Bloch sphere that stretches the circular uncertainty region of a coherent spin state into an ellipse with a shrinking minor axis. 4.6(6) dB of metrological spin squeezing have so far been demonstrated by this approach.

The potential of these methods to improve practical precision measurements has been studied by operating an atomic clock with a squeezed input state. The lifetime of the squeezed state has been examined in the presence of real laboratory noise sources, illustrating how simple uncertainty analysis can indicate parameter regimes where the initial states' fluctuations dominate the clock output noise and where squeezing is potentially helpful. Under certain conditions the clock with the squeezed input state has an Allan deviation which averages down 2.8(3) times faster than an ideal

clock operating at the standard quantum limit for an unentangled input state.

Tripling the measurement rate is already a practically significant potential performance improvement, but substantial further gains in squeezing are possible with the methods presented here. As argued in chapter 7, the demonstrated squeezing is currently limited by the precision of state readout, not state preparation. Improved state detection, either by the on-resonance phase measurement proposed in section 3.2.2 or by the use of cycling transitions to avoid scattering-induced noise as in reference [33], could provide at least an additional 3 dB of squeezing by removing this limitation. Avoiding loss of contrast by reducing the inhomogeneous dephasing between atomic pseudospins would also help improve the achievable signal-to-noise ratio. This could be achieved by using matched lattices for the trapping and probing light fields so that the trapped atoms are uniformly coupled to the relevant resonator mode [64], or by using a two-color probing scheme as in reference [33] to cancel out the average inhomogeneous phase shift. Finally, even the current 14 dB Raman-scattering limit can be improved, either by increasing the collective cooperativity with additional atoms or a higher-finesse resonator, by employing probe pulses with known photon number to reduce the phase broadening from photon shot noise in cavity feedback squeezing (section 3.3), or by using cycling transitions to reduce scattering noise in measurement-induced squeezing [33].

The work presented here can be developed in other directions than quantitative improvements to the metrological squeezing parameter. The effective S_z^2 interaction used in cavity feedback squeezing is sufficient in principle, when combined with rotations of the Bloch sphere by microwave pulses, to generate arbitrary transformations of the Bloch vector's $2S_0$ -dimensional Hilbert space [28]. While decoherence currently limits this control to very simple operations such as the shearing demonstrated here, the theoretical analysis of section 3.3 shows that the information leakage out of the resonator via the exiting probe light can in principle be stopped using squeezing light pulses of known photon number, to allow interactions very close to the ideal S_z^2 and correspondingly more precise and complex operations. Even with current decoherence rates, the non-linearities built into the Lorentzian resonance of the optical cavity

and the spherical geometry of the space of states for the Bloch vector may allow the production of states with a more interesting shape than the elliptical Gaussians considered here. Detecting such shapes, by measuring higher moments of the transverse spin distribution, would be a significant first step on the path from simple manipulation of the uncertainties of semi-classical Bloch vectors to the exercise of complete control over a Hilbert space of mesoscopic size.

Bibliography

- [1] Walther Gerlach and Otto Stern. Das magnetische moment des silberatoms. *Z. Phys.*, 9:353–355, 1922. URL <http://dx.doi.org/10.1007/BF01326984>.
- [2] I. I. Rabi, J. R. Zacharias, S. Millman, and P. Kusch. A new method of measuring nuclear magnetic moment. *Phys. Rev.*, 53(4):318, February 1938.
- [3] Alfred Kastler. Quelques suggestions concernant la production optique et la détection optique d’une inégalité de population des niveaux de quantification spatiale des atomes. Application à l’expérience de Stern et Gerlach et à la résonance magnétique. *J. Physique*, 11(6):255, June 1950.
- [4] D. J. Wineland, R. E. Drullinger, and F. L. Walls. Radiation-pressure cooling of bound resonant absorbers. *Phys. Rev. Lett.*, 40(25):1639–1642, June 1978.
- [5] Masahiro Kitagawa and Masahito Ueda. Squeezed spin states. *Phys. Rev. A*, 47(6):5138–5143, June 1993.
- [6] D. J. Wineland, J. J. Bollinger, W. M. Itano, F. L. Moore, and D. J. Heinzen. Spin squeezing and reduced quantum noise in spectroscopy. *Phys. Rev. A*, 46(11):R6797–R6800, December 1992.
- [7] D. J. Wineland, J. J. Bollinger, W. M. Itano, and D. J. Heinzen. Squeezed atomic states and projection noise in spectroscopy. *Phys. Rev. A*, 50(1):67–88, July 1994.
- [8] G. Santarelli, Ph. Laurent, P. Lemonde, A. Clairon, A. G. Mann, S. Chang, A. N. Luiten, and C. Salomon. Quantum projection noise in an atomic fountain: A high stability cesium frequency standard. *Phys. Rev. Lett.*, 82(23):4619–4622, June 1999.
- [9] J. Hald, J. L. Sørensen, C. Schori, and E. S. Polzik. Spin squeezed atoms: A macroscopic entangled ensemble created by light. *Phys. Rev. Lett.*, 83(7):1319–1322, August 1999.
- [10] V. Meyer, M. A. Rowe, D. Kielpinski, C. A. Sackett, W. M. Itano, C. Monroe, and D. J. Wineland. Experimental demonstration of entanglement-enhanced rotation angle estimation using trapped ions. *Phys. Rev. Lett.*, 86(26):5870–5873, June 2001.

- [11] J. Estève, C. Gross, A. Weller, S. Giovanazzi, and M. K. Oberthaler. Squeezing and entanglement in a Bose-Einstein condensate. *Nature*, 455:1216–1219, October 2008.
- [12] 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(7):073604, February 2010.
- [13] 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(2):021804, February 2010.
- [14] Ian D. Leroux, Monika H. Schleier-Smith, and Vladan Vuletić. Implementation of cavity squeezing of a collective atomic spin. *Phys. Rev. Lett.*, 104(7):073602, February 2010.
- [15] Ian D. Leroux, Monika H. Schleier-Smith, and Vladan Vuletić. Orientation-dependent entanglement lifetime in a squeezed atomic clock. *Phys. Rev. Lett.*, 104(25):250801, June 2010.
- [16] R. Wynands and S. Weyers. Atomic fountain clocks. *Metrologia*, 42:S64–S79, 2005.
- [17] A. Kuzmich, N. P. Bigelow, and L. Mandel. Atomic quantum non-demolition measurements and squeezing. *Europhysics Lett.*, 42(5):481, 1998.
- [18] Anne Louchet-Chauvet, Jürgen Appel, Jelmer J. Renema, Daniel Oblak, Niels Kjærgaard, and Eugene S. Polzik. Entanglement-assisted atomic clock beyond the projection noise limit. *New J. Phys.*, 12(6):065032, June 2010.
- [19] Richard P. Feynman, Frank L. Vernon, and Robert W. Hellwarth. Geometrical representation of the Schrödinger equation for solving maser problems. *J. Appl. Phys.*, 28(1):49, January 1957.
- [20] R. H. Dicke. Coherence in spontaneous radiation processes. *Phys. Rev.*, 93(1):99, January 1954.
- [21] F. T. Arecchi, Eric Courtens, Robert Gilmore, and Harry Thomas. Atomic coherent states in quantum optics. *Phys. Rev. A*, 6(6):2211–2237, December 1972.
- [22] T. Holstein and H. Primakoff. Field dependence of the intrinsic domain magnetization of a ferromagnet. *Phys. Rev.*, 58(12):1098–1113, December 1940.
- [23] Lars Bojer Madsen and Klaus Mølmer. Spin squeezing and precision probing with light and samples of atoms in the Gaussian description. *Phys. Rev. A*, 70(5):052324, November 2004.
- [24] Lieven M. K. Vandersypen, June 2010. Private communication.

- [25] Anders S. Sørensen and Klaus Mølmer. Entanglement and extreme spin squeezing. *Phys. Rev. Lett.*, 86(20):4431–4434, May 2001.
- [26] Claude Cohen-Tannoudji, Jacques Dupont-Roc, and Gilbert Grynberg. *Atom-Photon Interactions: Basic Processes and Applications*. Wiley-VCH, Weinheim, 2004.
- [27] A. Kuzmich, L. Mandel, and N. P. Bigelow. Generation of spin squeezing via continuous quantum nondemolition measurement. *Phys. Rev. Lett.*, 85(8):1594–1597, August 2000.
- [28] Souma Chaudhury, Seth Merkel, Tobias Herr, Andrew Silberfarb, Ivan H. Deutsch, and Poul S. Jessen. Quantum control of the hyperfine spin of a Cs atom ensemble. *Phys. Rev. Lett.*, 99(16):163002, October 2007.
- [29] Igor Teper, Geert Vrijsen, Jongmin Lee, and Mark A. Kasevich. Backaction noise produced via cavity-aided nondemolition measurement of an atomic clock state. *Phys. Rev. A*, 78(5):051803, November 2008.
- [30] 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(3):033601, January 2009.
- [31] 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, July 2009.
- [32] H. Uys, M. J. Biercuk, A. P. VanDevender, C. Ospelkaus, D. Meiser, R. Ozeri, and J. J. Bollinger. Decoherence due to elastic Rayleigh scattering. *Phys. Rev. Lett.*, 105(20):200401, November 2010.
- [33] M. Saffman, D. Oblak, J. Appel, and E. S. Polzik. Spin squeezing of atomic ensembles by multicolor quantum nondemolition measurements. *Phys. Rev. A*, 79(2):023831, February 2009.
- [34] C. Orzel, A. K. Tuchman, M. L. Fenselau, M. Yasuda, and M. A. Kasevich. Squeezed states in a Bose-Einstein condensate. *Science*, 291(5512):2386–2389, March 2001.
- [35] Fabrice Gerbier, Simon Fölling, Artur Widera, Olaf Mandel, and Immanuel Bloch. Probing number squeezing of ultracold atoms across the superfluid-Mott-insulator transition. *Phys. Rev. Lett.*, 96(9):090401, March 2006.
- [36] G.-B. Jo, Y. Shin, S. Will, T. A. Pasquini, M. Saba, W. Ketterle, D. E. Pritchard, M. Vengalattore, and M. Prentiss. Long phase coherence time and number squeezing of two Bose-Einstein condensates on an atom chip. *Phys. Rev. Lett.*, 98(3):030407, January 2007.

- [37] J. Sebby-Strabley, B. L. Brown, M. Anderlini, P. J. Lee, W. D. Phillips, J. V. Porto, and P. R. Johnson. Preparing and probing atomic number states with an atom interferometer. *Phys. Rev. Lett.*, 98(20):200405, May 2007.
- [38] Wei Li, Ari K. Tuchman, Hui-Chun Chien, and Mark A. Kasevich. Extended coherence time with atom-number squeezed states. *Phys. Rev. Lett.*, 98(4):040402, January 2007.
- [39] P. Rosenbusch. Magnetically trapped atoms for compact atomic clocks. *Appl. Phys. B*, 95(2):227–235, May 2009.
- [40] Martin W. Zwierlein, Zoran Hadzibabic, Subhadeep Gupta, and Wolfgang Ketterle. Spectroscopic insensitivity to cold collisions in a two-state mixture of fermions. *Phys. Rev. Lett.*, 91(25):250404, December 2003.
- [41] M. J. Collett and C. W. Gardiner. Squeezing of intracavity and traveling-wave light fields produced in parametric amplification. *Phys. Rev. A*, 30(3):1386–1391, September 1984.
- [42] C. W. Gardiner and M. J. Collett. Input and output in damped quantum systems: Quantum stochastic differential equations and the master equation. *Phys. Rev. A*, 31(6):3761–3774, June 1985.
- [43] U. Fano. Effects of configuration interaction on intensities and phase shifts. *Phys. Rev.*, 124(6):1866–1878, December 1961.
- [44] Igor Teper. *Ultracold atoms in microfabricated magnetic traps near surfaces and inside optical resonators*. PhD thesis, Stanford University, 2006.
- [45] Yu-Ju Lin. *Ultracold atoms magnetically trapped inside an optical resonator on a microfabricated chip*. PhD thesis, Stanford University, 2006.
- [46] Daniel A. Steck. Rubidium 87 D line data, August 2009. URL <http://steck.us/alkalidata/rubidium87numbers.pdf>.
- [47] J. Vanier and C. Audoin. *The Quantum Physics of Atomic Frequency Standards*. Adam Hilger, Philadelphia, 1989.
- [48] H. J. Lewandowski, D. M. Harber, D. L. Whitaker, and E. A. Cornell. Observation of anomalous spin-state segregation in a trapped ultracold vapor. *Phys. Rev. Lett.*, 88(7):070403, January 2002.
- [49] D. M. Harber, H. J. Lewandowski, J. M. McGuirk, and E. A. Cornell. Effect of cold collisions on spin coherence and resonance shifts in a magnetically trapped ultracold gas. *Phys. Rev. A*, 66(5):053616, November 2002.
- [50] C. Wieman and T. W. Hänsch. Doppler-free laser polarization spectroscopy. *Phys. Rev. Lett.*, 36(20):1170–1173, May 1976.

- [51] T. A. Savard, K. M. O'Hara, and J. E. Thomas. Laser-noise-induced heating in far-off-resonance optical traps. *Phys. Rev. A*, 56(2):R1095–R1098, August 1997.
- [52] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward. Laser phase and frequency stabilization using an optical resonator. *Appl. Phys. B*, 31:97–105, 1983.
- [53] Monika H. Schleier-Smith. *Thesis*. PhD thesis, Massachusetts Institute of Technology, March 2011.
- [54] C. J. Hawthorn, K. P. Weber, and R. E. Scholten. Littrow configuration tunable external cavity diode laser with fixed direction output beam. *Rev. Sci. Instr.*, 72(12):4477–4479, December 2001.
- [55] Donna Berns, May 2007. Private communication.
- [56] Anthony E. Siegman. *Lasers*. University Science Books, Sausalito, CA, 1986.
- [57] Wolfgang Ketterle, Kendall B. Davis, Michael A. Joffe, Alex Martin, and David E. Pritchard. High densities of cold atoms in a dark spontaneous-force optical trap. *Phys. Rev. Lett.*, 70(15):2253–2256, April 1993.
- [58] D. Boiron, C. Triché, D. R. Meacher, P. Verkerk, and G. Grynberg. Three-dimensional cooling of cesium atoms in four-beam gray optical molasses. *Phys. Rev. A*, 52(5):R3425–R3428, November 1995.
- [59] Holly K. Cummins, Gavin Llewellyn, and Jonathan A. Jones. Tackling systematic errors in quantum logic gates with composite rotations. *Phys. Rev. A*, 67(4):042308, April 2003.
- [60] S. F. Huelga, C. Macchiavello, T. Pellizzari, A. K. Ekert, M. B. Plenio, and J. I. Cirac. Improvement of frequency standards with quantum entanglement. *Phys. Rev. Lett.*, 79(20):3865–3868, November 1997.
- [61] D. J. Wineland, C. Monroe, W. M. Itano, D. Leibfried, B. E. King, and D. M. Meekhof. Experimental issues in coherent quantum-state manipulation of trapped atomic ions. *J. Res. NIST*, 103(3):259, 1998.
- [62] A. André, A. S. Sørensen, and M. D. Lukin. Stability of atomic clocks based on entangled atoms. *Phys. Rev. Lett.*, 92(23):230801, June 2004.
- [63] V. Gerginov, N. Nemitz, S. Weyers, R. Schröder, D. Griebisch, and R. Wynands. Uncertainty evaluation of the caesium fountain clock PTB-CSF2. *Metrologia*, 47(1):65, 2010.
- [64] Igor Teper, 2009. Private Communication.