First Steps toward Precision Measurements using Multicomponent Bose-Einstein Condensates of ^{87}Rb

by

James Michael Higbie

B.A. (Harvard University) 1998
M.A. (University of California, Berkeley) 2001

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

Doctor of Philosophy
in

Physics

in the

GRADUATE DIVISION

of the

UNIVERSITY OF CALIFORNIA, BERKELEY

Committee in charge:

Professor Dan M. Stamper-Kurn, Chair
Professor Dmitry Budker
Professor K. Birgitta Whaley

Fall 2005

The dissertation of James Michael Higbie is approved.

Chair Date

Date

Date

University of California, Berkeley

Fall 2005

First Steps toward Precision Measurements using Multicomponent Bose-Einstein Condensates of ^{87}Rb

Copyright © 2005

by

James Michael Higbie

Abstract

First Steps toward Precision Measurements using Multicomponent Bose-Einstein Condensates of ^{87}Rb

by

James Michael Higbie

Doctor of Philosophy in Physics

University of California, Berkeley

Professor Dan M. Stamper-Kurn, Chair

Experimental and theoretical progress toward the application of multicomponent Bose-Einstein condensates of ^{87}Rb to precision metrology are discussed. The building of a laser-cooling and trapping experimental apparatus is described. Experimental work on magnetization-sensitive imaging of spin-1 spinor condensates and their application to magnetometry is discussed, followed by a theoretical discussion of the periodically dressed condensate, with reference to enhancing the sensitivity of frequency and time measurements.

Professor Dan M. Stamper-Kurn
Dissertation Committee Chair

Contents

Contents i

List of Figures iv

List of Tables vi

Acknowledgements vii

1 Introduction 1
- 1.0.1 A Brief History of the Experiment . . . 3
- 1.0.2 Outline . . . 8

2 Apparatus 9
- 2.1 Oven . . . 10
  - 2.1.1 Characterization of Atomic Beam . . . 12
- 2.2 Zeeman Slower . . . 15
  - 2.2.1 Deceleration and Optimal Length . . . 19
  - 2.2.2 Coil Layout . . . 20
- 2.3 Diode Lasers and Tapered Amplifier . . . 26
- 2.4 Cold Plates . . . 32
- 2.5 Magneto-Optical Trap . . . 37
- 2.6 Magnetic Trap . . . 40
  - 2.6.1 Theory of Magnetic Trapping . . . 40
  - 2.6.2 Trap Design and Construction . . . 44
  - 2.6.3 Uncompressed and Compressed Operation . . . 49
  - 2.6.4 Switching Currents Off . . . 51

2.6.5 RF Noise Reduction 54
2.6.6 RF Evaporation 55
2.6.7 Water-Cooling and Interlock 59
2.7 Optical Trap 61
2.7.1 Optical Layout 63
2.7.2 Loading the Optical Trap 64
2.8 RF Coils and Spin Manipulation 66
2.9 Stern-Gerlach Analysis 69
2.10 Magnetic Field Control 69
2.11 Imaging Systems 72

3 Spinor Condensates **75**
3.1 Overview of Prior Work 75
3.2 Theoretical Description of Spin-1 Spinor Condensate 77
3.2.1 Mean-Field Solutions at Finite Magnetic Field 78
3.3 Limitations of Prior Imaging Techniques 79
3.4 In-situ Imaging 82
3.4.1 Phase-Contrast Technique 82
3.4.2 Magnetization Sensitivity 83
3.5 Imaging Larmor Precession 91
3.5.1 Image Processing 95
3.5.2 Zeeman Coherence in a BEC 96
3.5.3 Thermal Bose Gas 103

4 Prospects for Magnetometry **105**
4.1 Principle of the Measurement 105
4.2 Spatial Scale in Magnetometry 106
4.3 Resolution 109
4.3.1 Absence of Mean-Field Shift 113
4.3.2 Dynamic Range 115
4.4 Limitations 118
4.4.1 Optical Shot Noise 118

4.4.2 Optimum Probe Frequency for Measuring Magnetization . . . 119
4.4.3 Atomic Shot Noise 123
4.4.4 Self-Field 123
4.4.5 Effects of Quadratic Zeeman Shift 124
4.5 Test via A.C. Stark Shift 125
4.6 Possibility of Spin Squeezing 127

5 Theory of the Periodically Dressed Condensate 130
5.1 Single-Atom Theory 131
5.2 Bogoliubov Treatment of Many-Body Calculation 135
5.3 Superfluidity of Periodically Dressed Condensate 139
5.4 Periodically Dressed Condensates and the Double-Well Problem . . . 141
5.4.1 Raman Resonance and External Potential 142
5.5 Ground States of the Double-Well Problem 145
5.5.1 Feasibility of Schrödinger-Cat States 149
5.6 Application to Sub-Shot-Noise Interferometry 152
5.7 Stabilizing a Schrödinger-Cat Clock 155
5.8 Condensate-Number Fluctuations 158

Bibliography 159

List of Figures

2.1 Scheme of Apparatus 10
2.2 Diagram of Oven 11
2.3 Slowing Diagram 19
2.4 Optimal Length of Zeeman Slower 21
2.5 Zeeman Slower Coil Layout 22
2.6 Slower Profile 26
2.7 Optics Layout 27
2.8 External Cavity Diode Laser Diagram 28
2.9 Tapered Amplifier Mount 30
2.10 Diagram of TEC Stack 34
2.11 Zeeman Splitting 41
2.12 Magnetic Trap Layout 44
2.13 Axial Magnetic Trap Field Measurement 50
2.14 Transverse Gradient Measurement 50
2.15 Magnetic Trap Schematic 52
2.16 Ioffe-Pritchard Trap Compression 57
2.17 Interlock Circuit Schematic 58
2.18 Beam-Shaping for the Optical Trap 65
2.19 Rabi Level Diagram 67
2.20 Landau-Zener Sweep Error Rate 68
2.21 Imaging System Diagram 73

3.1 Phase-Contrast Imaging 81
3.2 D1 Transition Diagram 85

3.3 Schematic of Kinetics Imaging 94
3.4 Phase Gradient vs. Time 100
3.5 Sample Larmor Precession Image 101
3.6 Larmor Precession Decay 102
3.7 Decoherence of Thermal Cloud in Magnetic Gradient 103

4.1 Comparison to SQUID 108
4.2 Wound-Up Condensate Image 109
4.3 Phase Map of Wound-Up Condensate 110
4.4 Imaging Resolution Calibration 112
4.5 Signal/Noise Ratio vs. Detuning 121

5.1 Levels in Raman Transition 132
5.2 Dispersion Relation for Zero Rabi Frequency 134
5.3 Dispersion Relation For Finite Rabi Frequency 134
5.4 Quasiparticle Dispersion Relation 139
5.5 Squeezed State Histogram 147
5.6 Schrödinger Cat State 148
5.7 Pure and Partial Schrödinger Cat-State Energy Splitting 150
5.8 Clock Enhancement from Partial Cat State 151
5.9 Correlated/Uncorrelated Comparison 154
5.10 Fluctuations in Measured Phase vs. Feedback Step 156

List of Tables

2.1 Table of TEC Stack Performance Coefficients 37

Acknowledgements

Thank you to my parents, who made me the person I am today, among others. And thank you also to Wendy, whose kindness and good cheer have always made everything so much more pleasant. Thanks to my adviser, Dan Stamper-Kurn, whose many good ideas and quick grasp of the essential physics of any new problem he happened to be peeking or poking at have always impressed me. Thanks to Lorraine, whose levelheadedness and common sense helped us survive many more malfunctions than seemed strictly necessary at the time, and to whom I wish the best as she prepares to navigate the unfamiliar world of the Eastern Seaboard. Thanks to Shin, who brought to our labs a salutary thoroughness and precision which I am sure will continue to serve him well as he embarks on his exciting research career at the University of Tokyo. Thanks to Mukund, whose great work ethic and outstanding agreeableness have made him a pleasure to work with these last few months and give me confidence that the experiment will continue its march of progress, trampling in its path any such frivolities as sleep or solid nourishment. Thanks to Veronique, who did so much to get us started and whose elegant French handwriting, accidentally discovered on the odd lens or waveplate, never fails to elicit a nostalgic smile; to Ananth who, though he never quite managed to get our entire lab running on wind power, left us with a good deal more understanding of Bose condensates; to Sabrina, who is carrying on Shin's torch in matters of meticulousness, not least with respect to her plans to take a wrecking ball to the experiment as it now exists. Thanks also to the members of the other experiments–to Kevin, for all his help building lasers, tightening vacuum bolts, and helping us test our vibration isolation using his subwoofers; to Tom, whose sense of humor and electronics savvy have made our many conversations as entertaining as they were enlightening; to Kater, for letting us experience life on a Marin farm and for lending Kepco power supplies; and to Deep, for a number of very stimulating and informative conversations about physics and cricket. A big thank you to all the physics department staff, especially Anne Takizawa, Donna Sakima, Joseph Kant, Pete Thuesen, Dave Murai, Dave Nguyen, John Davis, Le Vern Garner, Pat Bonnefil, Dave Gallegos, and Ron Wyckoff, who have all been very friendly and helpful over the years. And, last and least (but still significantly), thanks to the writer of the UC Thesis LaTeX template.

Chapter 1

Introduction

One indication that a field of physics has reached full maturity is that it becomes well enough understood to be useful as a tool for understanding other physics. This sort of technological leapfrogging occurs constantly and on many levels in experimental work, but one of the most impressive examples is in atomic spectroscopy. Attempts to explain atomic spectra were at the center of the development of quantum mechanics. By the middle of the twentieth century, however, understanding of the structure of few-electron atoms and the technical ability to manipulate and probe them had progressed to the point that atoms were being used both to measure time at ever-greater levels of accuracy, reaching a symbolic milestone when, in 1967, the second was redefined in terms of the hyperfine frequency splitting of cesium, and to probe new physics, as in the evidence for quantum field theory provided by measurements of the Lamb shift [1]. The use of precise atomic spectroscopic methods to explore new physics has continued to flourish over the intervening decades and through the present, with work on the electron electric dipole moment [2–6], on possible variation of α or other physical constants [7–11], and on possible violation of Lorentz symmetry [12–16] or of the spin-statistics theorem [17].

As researchers learned to probe atoms more and more precisely, the realization was not long in coming [18, 19] that the motional states of atoms could be manipulated with laser beams and magnetic fields, and over the two decades between 1975 and 1995 the field of laser cooling grew from a theoretical suggestion into a crucial enabling technology in the creation of dilute-gas Bose-Einstein condensates (BECs)

[20, 21]. An important motivation for the development of laser cooling and Bose condensation was not only to spawn a new subfield in the study of quantum fluids, but also to provide a source of atoms with a very narrow velocity distribution which would allow cold atoms to repay their great debt to atomic spectroscopy. This motivation is apparent, for instance, in the often-repeated characterization of a Bose condensate as resembling an atom laser, which carries with it the implication that the BEC may, like the laser, prove to be a fruitful general-purpose tool in other areas of physics. It is certainly true that laser cooling has found important applications to spectroscopy in atomic fountain clocks [22, 23] and cold-atom precision measurements of the the atomic recoil frequency [24, 25]. Dilute-gas Bose condensates, however, while they have opened the way to many interesting and important experiments, have not yet proven of value to atomic spectroscopy or precision measurement, despite the fact that in their extremely small velocity distributions and relatively long available interaction times, BECs would appear in many respects to be an ideal source. This failure to date may be largely attributed to the small number of Bose-condensed atoms typically available per unit time, and to the systematic effects associated with ultracold collisions in a dense sample, which afflict even the comparatively rarefied atomic gases in an atomic fountain [26, 27].

The problem of insufficient "brightness" of BEC atom sources appears likely to be significantly ameliorated over coming years by technological improvements. Already, the cycle time of BEC experiments has been reduced from over a minute, in a "traditional" magnetic-trap-based apparatus, to a few seconds, by the use of all-optical techniques [28–30]. Moreover, techniques such as laser ablation and cryogenic pre-cooling [31] may substantially improve the initial conditions of BEC experiments, leading to higher condensate numbers. Various ways have been suggested to circumvent the problem of systematic frequency shifts from collisional interactions, such as preparing atoms in a carefully chosen mixture of states with opposite shifts [32], using the density dependence of cavity-pulling to cancel that of the collisional shift [33], employing fermions (whose S-wave cross section vanishes at low temperature) [34], and immobilizing atoms at the sites of an optical lattice [35, 36].

An alternative possibility for exploiting the desirable properties of Bose condensates for precision measurements without suffering from the effects of cold collisions is to exploit symmetry properties of the system. Such an approach has worked very well, for instance, in a small number of condensed matter systems such as the quantum Hall system, in which gauge invariance ensures the quantization of conductivity [37], or in superconducting quantum-interference devices (SQUIDs), in which the Meissner effect and single-valuedness of the wave function enforce flux quantization [38]. This approach is likewise analogous to the use of decoherence-free subspaces for quantum computation, in which the state-space is specifically chosen or engineered to be protected by symmetry from decoherence [39]. It is this sort of immunity from collisional perturbations that is invoked in chapter 4 with regard to using spinor condensates for magnetometry. In chapter 5 on the theory of the periodically dressed condensate, on the other hand, the problem of dealing with collisional interactions is turned on its head, and an attempt is made not to mitigate but to exploit these interactions to enhance the statistical sensitivity of a frequency or time measurement. This latter project is admittedly purely theoretical and of unknown practicality at present, but is offered as an illustration of the claim that with creativity and by exploiting the extra degrees of freedom offered by multicomponent condensates, it may still be possible to fulfil the promise of Bose condensation for improving precision measurements.

1.0.1 A Brief History of the Experiment

The work described in this thesis was much of it highly collaborative, so that several of the people involved may reasonably lay claim to many of the experimental accomplishments, while in other cases the primary responsibility or originating idea belonged clearly to an individual. It is not my intention in what follows to make any finely calibrated judgements about the degree to which credit should be apportioned, but merely to outline in broad terms what I consider to have been my contributions, both as part of an experimental team and as an individual researcher.

I joined the Stamper-Kurn group in January, 2001, which at the time consisted of Dan (the group's leader and eponym) and Lorraine Sadler, herself a member of one week's standing. While waiting for a laboratory to be made available to us, we busied ourselves with trying to organize, design, purchase, and build our way to a functioning Bose-Einstein-condensation machine. Given the size of the group, we were all involved in each other's work, but in the rough division of labor that emerged, I was responsible for designing and building the many electromagnets required for the experiment. This meant designing and building first the Zeeman slower, and then the magneto-optical/magnetic trap coils, as well as the power-switching, safety-interlock, and analog-control isolation electronics necessary to operate them. As we became more conscious of the expected fluxes and background pressures of rubidium in the vacuum system, we realized that it was desirable to have cold surfaces in the vacuum system to pump rubidium, and I accordingly built some Peltier-cooler stacks and cold plates for the purpose of controlling this background pressure. Before long, the Zeeman slower and the diode lasers that Dan, along with our new post-doc Veronique Savalli and graduate student Kevin Moore, had been building and locking to a rubidium frequency reference, came together, and Veronique quickly saw evidence of Zeeman-slowing in the atomic beam. Heartened by this success, we disassembled the skeleton vacuum chamber in which this first experiment occurred and replaced it with the full vacuum chamber, now incorporating the magnetic trap, including the interlock and the dedicated water-cooling system I had built.

This began a focused effort by Lorraine and me, with frequent help from Dan, to observe ultracold atoms in a magneto-optical trap (MOT). This stage of the experiment's progress took longer than expected, as a result both of some incorrect laser frequencies and of the initial lack of a systematic means of aligning the rather long Zeeman slower, but eventually the long-awaited MOT materialized, followed in short order by magnetically trapped atoms and the beginnings of radio-frequency evaporation. After eliminating some stray light and shuttering the atomic beam, we had magnetic-trap lifetimes in excess of 50 seconds, and believed ourselves in a favorable position to reach runaway evaporation in the near future. Unfortunately, it soon became apparent that the trap lifetime was strongly dependent on the magnitude of the bias field, limiting our ability to compress the atomic cloud and achieve the collision rates necessary for efficient evaporation. Finding the radio-frequency (RF) noise responsible for this number loss became a top priority, and involved eliminating DC-DC converters wherever possible, turning off ion gauges, and attempting to filter the output of the large switching magnet power supplies with capacitor banks and huge home-built inductors. At the same time, we worked at getting the best possible initial conditions for evaporation. This meant getting powerful MOT laser beams, and so I built a tapered amplifier assembly to increase the MOT power from 40 mW without fiber coupling to about 120 mW at the output of a single-mode optical fiber, with all the accompanying advantages of stability and isolation from upstream misalignment. We also worked at matching the size and temperature of the atoms from the MOT to the curvature of the magnetic trap (loosely called "mode-matching", in an optics analogy). Among my particular contributions at this period were the discovery that lowering the magnetic-field gradient could dramatically compress the cloud (an effect that was well known in the literature, but not in our lab), the implementation and optimization of a two-step evaporation procedure that prevented the atoms from overheating during compression, and the compensation of the stray magnetic field of the slower that was giving the MOT a transient kick during loading (although this latter turned out to be a transient problem.) The combination of a colder and larger initial atom cloud with reduced spin-flip rates from RF noise led to the production of our first Bose-Einstein condensates (BECs) in January, 2004.

Although it represented a large amount of time and work, the ability to produce BECs routinely was really just the starting point for the experiments we wished to perform. While progress setting up an optical trap for spinor condensates continued on the side, with help from Matt Pasienski and Craig Hetherington, Lorraine and I, now joined by Ananth Chikkatur, began some simple experiments on superradiant Raman scattering in a BEC, viewed from a nonlinear-optics standpoint, and ultimately aiming to observe an oscillation between superradiant optical pumping into an optically dark state and Larmor precession out of that dark state. These experiments, while interesting, were later judged a distraction from the core mission of the experiment, to study spinor condensates, and were postponed indefinitely. Progress towards putting condensates in an optical-dipole trap (ODT) went forward, and atoms were soon trapped, but heating rates in the optical trap were too high to maintain BEC. Following several efforts to shorten the ODT path length, improve the beam shape, and stabilize the optical power, I implemented a beam path entirely on one optical breadboard, which itself was clamped between pads of vibration-damping Sorbothane ®. After these improvements, we were able to transfer condensates from the magnetic trap to the ODT, but the resulting condensates were small and motionally excited. I had remained unconvinced by assertions that adiabatic transfer, as opposed to a sudden switch-off of the magnetic trap, was too difficult to be worth trying, and pushed forward with the attempt to maintain the location of the cloud as the magnetic-trap fields ramped down by adjusting the ramp trajectories. In the end, this approach worked very well for transferring thermal clouds, which could then be quickly evaporated into BECs in the tight, highly collisional optical trap.

With spinor condensates available to be studied, we needed only the imaging light, and so I put together a laser and lock circuitry at the D1 frequency of ^{87}Rb , using a mount built by our diploma student Stefan Schmidt, and we figured out how to align our phase dot and saw our first phase-contrast images soon thereafter. We quickly learned to manipulate the atomic spins with RF sweeps and pulses and set about learning to control the magnetic fields down to the 10 mG level, where we now planned to observe Larmor precession and see spontaneous ferromagnetism of the spinor condensate. After a few false positives, resulting from low-frequency RF noise, and some false negatives, probably due to small systematic differences between the calibration and observation procedures, we successfully observed Larmor precession and used our ability to image it to characterize the coherence-time of the thermal cloud and the condensate and to observe "magnon" excitations of the condensate excited by magnetic gradients [40]. After this exciting period, a small and intermittent vacuum leak and a failure of our CCD camera together cost us several months of progress. With a replacement camera suitable for absorption imaging but not for imaging Larmor precession, we undertook an interesting project outside of our main line of research, namely to use the enhanced absorption present in superradiant scattering from a BEC to image the condensate, obtaining information both about the coherence of the atomic sample and about the subtleties of spatially dependent superradiance. In order to enhance the signal further, I set up a chopper wheel in front of the camera, for use as a fast shutter, and implemented the necessary timing and logic to synchronize imaging with the chopper-wheel frequency. This development was superceded when our original camera, with its high-speed "kinetics" imaging, was finally back in order, allowing us to resolve the process of superradiance spatially and temporally. Although fascinating and qualitatively informative, this imaging technique turned out to be difficult to make quantitative. As Lorraine analyzed the superradiance data, I, with the help of our newest postdoc Mukund Vengalatore, de-mothballed the spinor experiment, quickly getting back spinor condensates in the ODT and the fine control of magnetic bias fields, before attempting the spontaneous magnetization experiment that had been cut short before. The experiment worked beautifully, and we were soon staring at two-dimensional images of condensates broken up into Larmor-precessing patches. Partly as a result of a MOT lens that had been moved for unrelated reasons, widening the ODT and making the condensates more optically resolvable, the quality of the images was high enough that Dan proposed analyzing them pixel by pixel to obtain phase and amplitude as a function of position in the condensate. This approach proved very productive, allowing detailed information about the spatial structure of spontaneous magnetization, and pointing toward topological defects, such as domain walls and spin vortices.

I worked on the analysis of the data as regarded spontaneous magnetization, while Lorraine moved forward in another front of the analysis, searching for vortices in these images. As the write-up of these results progressed, I then began work on the project, suggested by our earlier Larmor-precession work, of attempting to diagnose the spinor-condensate system as a precise, high-resolution magnetometer, using a fictitious optically-generated magnetic field. This project yielded preliminary results quite quickly, but measurements were limited by technical fluctuations in the magnetic-field inhomogeneities from shot to shot and within each shot. In order to be able to control the aspect ratio of the optical trap and give us a handle on spin instabilities, I designed and set up a cylindrical “zoom” lens configuration that would allow us to vary the aspect ratio over a wide range without making the beam astigmatic. Mukund and I succeeded in applying the spin-echo technique to eliminate some magnetic-field fluctuations, and then while I began writing my dissertation, Mukund refined the technique and, with Lorraine’s help, worked on several attempts to eliminate residual oscillation of the atomic sample in the optical trap. In a break from thesis-writing, I implemented one of the ideas for eliminating this motion, creating a small optical-lattice potential by retroreflecting a portion of the ODT, but the optical platforms in question were insufficiently stable, and the lattice served merely to heat the cloud severely. Further work on spinor-condensate magnetometry, using a blue-detuned laser to damp oscillations by means of a “speed-bump” potential, has been carried forward by Mukund in my absence.

Alongside these experimental endeavors, I was engaged also in a theoretical project, the so-called periodically dressed condensate, with the hope that it would be experimentally realizable. Although this has not happened during my time in graduate school, I have hope that the experiments may yet be attempted. This project began as an intuitive idea of Dan’s, on which I did some calculations as part of a final project for a condensed-matter physics class. The calculations, combined with further calculations by both Dan and me, were subsequently published [41]. I later did some related calculations, subsequently cast in a simpler form by Dan, exploring the possible use of the periodically dressed condensate as a model two-well system with the potential to create macroscopic entanglement [42].

1.0.2 Outline

The thesis reflects my graduate-school efforts in that considerable time and space are allotted to the experimental apparatus in chapter 2. Chapter 3 lays out the principal experimental innovation of this work, high-speed magnetization-sensitive phase-contrast imaging, and the observations of degenerate and nondegenerate spinor gases that it has enabled. In chapter 4, an exposition is made of attempts to use Larmor-precession imaging of spinor condensates for high-resolution magnetometry, along with principal advantages and limits to the technique. Finally, in chapter 5, the idea of the periodically dressed condensate, its predicted phenomenology, and a proposed application to metrology are discussed.

Chapter 2

Apparatus

This chapter describes the experimental apparatus that was constructed for producing, manipulating, and probing Bose-Einstein condensates. Many such machines have now been constructed, each with its own peculiar advantages and disadvantages, but with relatively little convergence on an optimal design. Particular attention will therefore be given to those aspects of the present design whose qualities, desirable, undesirable, or merely noteworthy, may modestly contribute to such convergence. Naturally, this chapter will also dwell at greater length on those portions of the experimental apparatus to which the present author has contributed most, as adumbrated in section 1.0.1.

The broad outlines of this apparatus are shown in figure 2.1. It consists of an atomic beam propagating from a rubidium oven to the main chamber (where condensates are ultimately formed), by way of a long deceleration region or Zeeman slower. Atoms exiting the Zeeman slower enter a magneto-optical trap (MOT), which in several steps cools and compresses the atoms until they are ready to be caught in a purely magnetic trap, suitable for forced radio-frequency (RF) evaporation close to the BEC transition. The final condensate formation is accomplished either in the magnetic trap or, for the bulk of the work described here, in an optical trap to which the atoms are transferred at a few times the critical temperature. Each of these experimental elements will be described below.

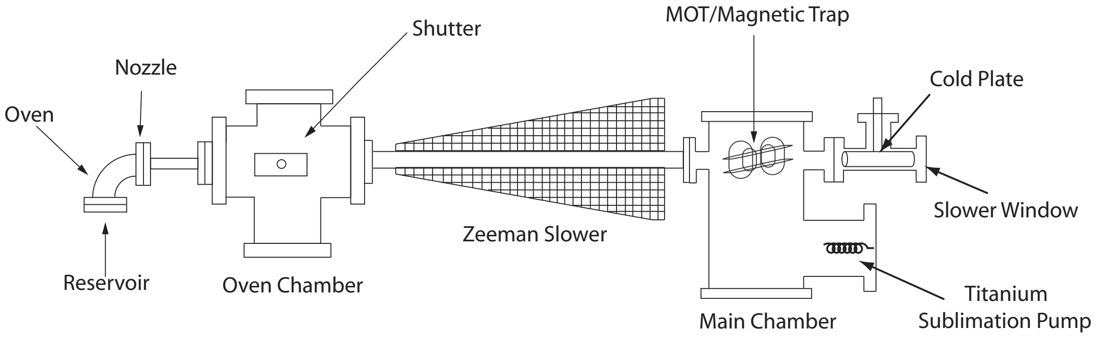

Figure 2.1. Over-all scheme of experimental apparatus (not to scale). An atomic beam propagates from left to right, originating at the oven. The beam may be shuttered in the oven chamber, or it may traverse the Zeeman-slower tube and attain the main chamber. Here a magneto-optical trap, magnetic trap, and optical trap are employed to create and study Bose condensates. Ultra-high vacuum in the main chamber is maintained by a titanium-sublimation pump, an ion pump (not pictured), and a cold plate.

2.1 Oven

The oven consists conceptually of a reservoir of liquid rubidium in equilibrium with its vapor and a small aperture from which the rubidium vapor can ballistically escape. In practice, the reservoir consists of a short section of 2.75 in.-flanged stainless steel ConFlat tubing in the form of an elbow, joined to the vacuum apparatus by another 3 in. long, $\frac{1}{2}$ in. diameter stainless steel vacuum nipple. The latter contains a piece of steel mesh, loosely-wound into a several-layer cylinder, which in turn is held in place by a smaller concentric steel tube. The end face of this innermost tube closest to the reservoir contains the aforementioned aperture or nozzle, whose diameter is 5 mm.

The flux of atoms from the oven is controlled by adjusting the temperature. This is a very effective control because of the fact that the vapor pressure of rubidium is a strong function of temperature, while the mean thermal velocity has a weak temperature dependence. Precise values of the vapor pressure may be found, for instance, in [43]. For many purposes, it is sufficient to remember a few points: at $-40°$C, the equilibrium pressure is around 10^{-11} torr; at 0°C, it is in the neighborhood of 10^{-8} torr; and at 100°C it approaches 10^{-4} torr. Very roughly, the vapor pressure increases by one order of magnitude every 20°C, though the change per degree is smaller at higher temperatures.

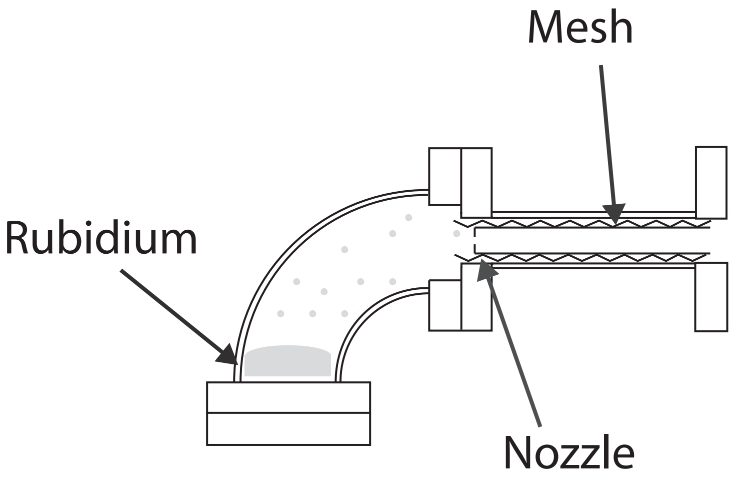

Figure 2.2. A diagram of the oven, showing the reservoir of liquid rubidium in equilibrium with an atomic vapor. Atoms traveling at the correct angle effuse through the nozzle and into the atomic beam, while those at larger angle strike the retaining tube or mesh and are wicked back to the hot (nozzle) end of the mesh, where they re-evaporate.

The Maxwell-Boltzmann phase-space distribution for a vapor of atoms at temperature $k_B T \equiv \tau$ and number density n is

$$\frac{\mathrm{d}N}{\mathrm{d}^3\mathbf{v}\mathrm{d}^3\mathbf{x}} = n\left(\frac{m}{2\pi\tau}\right)^{3/2}\exp\left(-\frac{m\mathbf{v^2}}{2\tau}\right). \tag{2.1}$$

If the direction of propagation of the atomic beam is chosen to be $\hat{\mathbf{z}}$, and the area of the oven aperture is A, then the number per time per velocity traversing the aperture is the well-known beam distribution

$$\frac{\mathrm{d}N}{\mathrm{d}^3\mathbf{v}\,\mathrm{d}t} = nAv_z\left(\frac{m}{2\pi\tau}\right)^{3/2}\exp\left(-\frac{m\mathbf{v^2}}{2\tau}\right). \tag{2.2}$$

A straightforward integration shows that the total number per time leaving the aperture is $nA\bar{v}/4$, where $\bar{v} \equiv \sqrt{8\tau/\pi m}$ is the mean magnitude of the velocity.

For typical parameters at 100°C of $n = 6 \times 10^{12}\,\mathrm{cm}^{-3}$, $\bar{v} = 3 \times 10^4\,\mathrm{cm/s}$, and an aperture area of $A = 0.2\,\mathrm{cm}^2$, this translates to a total effusion rate of $10^{16}\mathrm{s}^{-1}$. A reasonable charge of ten grams of rubidium, containing 7×10^{22} atoms, would thus last around 80 days if all rubidium emitted from the oven were permanently lost. Only emitted atoms possessing sufficiently small transverse velocity, however, will successfully be admitted into the atomic beam. The accepted fraction resulting from this solid-angle constraint is approximately the ratio of the 3/4-inch square ($\sim 4\,\mathrm{cm}^2$) open area at the far end of the beam (200 cm away from the oven) to the $2\pi \times (200\,\mathrm{cm})^2$ hemispherical surface area at that distance, numerically equal to 1.6×10^{-5} steradians.

As a result, the vast majority of atoms emitted from the oven are wasted, as far as condensate production is concerned. These considerations make it very desirable to incorporate some form of recirculation into the oven design, to prolong the period of operation between refills of the oven. This is the purpose of the steel mesh in the tube following the oven aperture. The portion of the steel mesh nearest the aperture is kept hotter than the remainder, with the intended effect that rubidium striking the mesh will be drawn toward the hot end, where it will evaporate, entraining other rubidium atoms into the depleted region and sustaining a wicking process which replenishes the reservoir.

In practice, the oven has operated approximately 60 to 70 hours per week over a period of as much as 10 months, or for a total of $\sim$ 3000 hrs. before requiring additional ampoules of rubidium. Compared to the non-recirculating estimate above, this suggests successful recirculation, although an alternative explanation would be that the emitted flux is lower than predicted. Given that different portions of the oven are maintained at different temperatures, it is necessary, absent a direct measurement, to make an assumption about the temperature characterizing the vapor. The most natural assumption is that the lowest temperature region at the bottom of the oven will collect the bulk of the liquid rubidium, and that the temperature of this reservoir will set the vapor pressure. If the other walls of the oven are substantially hotter, however, the kinetic temperature of the vapor may be higher, though a large temperature difference would be required to change the flux dramatically. Additional information about the oven emission may be obtained from measurements on the atomic beam, as described below.

2.1.1 Characterization of Atomic Beam

To characterize the performance of the oven, one would in principle like to know the differential flux as a function of velocity, which would suggest a Doppler-sensitive measurement. Near the output of the oven, however, where such measurements can be readily performed with reasonable signal-to-noise ratio, the current chamber offers only transverse (and therefore Doppler-insensitive) optical access to the atomic beam. This allows one to examine the absorption of a resonant laser beam transverse to the atomic beam, which can be related to the density and thereby to the flux, if the particular beam velocity distribution (equation (2.2)) is assumed.

The density of the beam on the far side of the oven aperture may be calculated from the relations

$$\frac{\mathrm{d}N}{\mathrm{d}V'\,\mathrm{d}v} = \frac{\mathrm{d}N}{r^2\mathrm{d}\Omega\, v\mathrm{d}t\mathrm{d}v} = \frac{v}{r^2}\frac{\mathrm{d}N}{\mathrm{d}^3\mathbf{v}\,\mathrm{d}t}, \tag{2.3}$$

where use has been made of the fact that the solid angle $\mathrm{d}\Omega$ of the velocity distribution at the aperture is the same as the solid angle that enters into the real-space volume element dV' far from the aperture. An integration with respect to velocity determines the density as a function of position:

$$n(r,\theta) = \frac{\mathrm{d}N}{\mathrm{d}V'} = \frac{n_{\mathrm{oven}}A}{4\pi r^2}\cos\theta. \tag{2.4}$$

The optical density of a resonant probe beam passing through this region is

$$O.D. = \int n(r,\theta)\sigma_{\mathrm{avg}}\mathrm{d}x, \tag{2.5}$$

where $n(r,\theta)$ is the total density irrespective of spin and $x = r\sin\theta$. Although the atomic beam at this point is relatively collimated, its transverse velocity spread corresponds to a Doppler shift of $\approx 20\,\mathrm{MHz}$, which is larger than the $6\,\mathrm{MHz}$ natural linewidth of rubidium, so that more properly the optical density is calculated as a convolution of the Doppler spread with the Lorentzian natural line shape of the atom, i.e., in terms of the Voigt profile:

$$\begin{aligned} O.D. &= \int \mathrm{d}x\mathrm{d}v \left.\frac{\mathrm{d}N}{\mathrm{d}V'\,\mathrm{d}v}\right|_{v,\theta} \sigma_{\mathrm{avg}}(v,\theta) \\ &= \frac{nA}{z}\left(\frac{m}{2\pi\tau}\right)^{3/2}\sigma_{\mathrm{avg}}\int \mathrm{d}v\,\mathrm{d}\theta\, v^2 e^{-mv^2/2\tau}\frac{1}{1+4k^2v^2\sin^2\theta/\gamma^2}. \end{aligned} \tag{2.6}$$

For $\gamma \ll 2kv$, the θ integral may be approximated by using the small-angle approximation $\sin\theta \approx \theta$ and taking the limits of integration to $\pm\infty$, since the effective collimation angle imposed by the resonance condition is narrower than the true collimation angle, while still small compared to 1 rad. With this approximation, the integral can be performed, with the result

$$O.D. = \frac{nA\sigma_{\mathrm{avg}}}{2\pi z}\frac{\gamma}{\bar{v}k}, \tag{2.7}$$

which is the answer that would be obtained by assuming a collimation angle of $\sin\theta_{\mathrm{col}} = \gamma/\bar{v}k$ and integrating the density profile (2.4). For the measurements described here, the linewidth of the laser is of order 1 MHz, so that broadening due to the laser linewidth is small.

The on-resonant cross section, averaged over the spin states of the atomic sample, is given by

$$\sigma_{\mathrm{avg}} = \sigma_0 \frac{1}{N_{\mathrm{g}}} \sum_g |c_{ge}|^2, \tag{2.8}$$

where c_{ge} are Clebsch-Gordon coefficients, $\sigma_0 = 6\pi k_0^{-2}$ is the resonant two-level-atom cross section, N_g is the total number of ground states, resonant or not, and the sum runs over transitions for which resonant light of the correct polarization is present. The temperature of the oven is large compared to the ground-state hyperfine splitting in ^{87}Rb (6.8 GHz $\sim$ 0.33 K), so that the Zeeman states in these levels are equally populated, resulting in a spin distribution which is isotropic, with no preferred spin axis. Thus the square amplitude summed over all states must be independent of the probe polarization. For the case where the excited manifold F' decays only to a single ground state manifold F, the average cross section is easily computed from considerations of symmetry, since

$$\sum_g |c_{ge}|^2 = \frac{1}{3} \sum_{\mathrm{pol.}} \sum_g |c_{ge}|^2 = \frac{N_e}{3}, \tag{2.9}$$

where N_e is the number of excited states in the F' manifold and where in the last step we have used the fact that all excited states decay at the same rate, so that the sum of the squares of the c_{eg}s from any excited state is unity. Thus for ^{87}Rb , probing on the $F = 2 \to F' = 3$ transition, $\sigma_{\mathrm{avg}} = \frac{7}{24}\sigma_0$. It is also useful to probe ^{85}Rb, since its greater natural abundance increases the signal size, and since none of the oven chemistry should distinguish among isotopes. For ^{85}Rb on the $F = 3 \to F' = 4$ transition $\sigma_{\mathrm{avg}} = \frac{1}{4}\sigma_0$.

We have performed measurements of the density in the atomic beam at the output of the oven. Because the atomic beam is already fairly collimated at this point, its transverse Doppler width is substantially smaller than that of the background vapor of rubidium in the oven chamber, as noted above. Moreover, it may be shuttered at will, allowing a ready distinction between the absorption due to the beam and that due to the background gas. To avoid saturation or optical-pumping effects, we perform the absorption measurements at low intensity, verifying that raising or lowering the power by a factor of two does not affect the measured absorption.

For experimental parameters $\theta_{\text{col}} = 0.5\,\text{cm}/10\,\text{cm}$, $z = 20\,\text{cm}$, $A = 0.2\,\text{cm}^2$, probing the $F = 2 \rightarrow F' = 3$ D2 transition of ^{87}Rb , and incorporating the fractional isotopic abundance of 27.8% for ^{87}Rb, this translates to

$$O.D. = \left(\frac{n_{\text{oven}}}{1.7 \times 10^{14}\,\text{cm}^{-3}} \right). \tag{2.10}$$

At the expected density for 100°C of $n_{\text{oven}} = 6 \times 10^{12}$ cm^{-3}, this implies about 3.5% absorption.

This value is in reasonable agreement with certain measured absorption values from a freshly loaded oven. At other times, however an absorption of 12.8% (August 26, 2004) or 13.3% (November 11, 2003) was recorded, both at an oven temperature of 96°C. This would appear to indicate that on these occasions, the vapor pressure in the oven was substantially higher or the collimation of the beam greater than has been assumed. The temperatures are measured by inserting thermocouples into the “sniffer” ports of the vacuum flanges. While inferior to a proper thermocouple well, these ports allow the thermocouple to penetrate to within a short distance of the inner surface of the oven, so that a large discrepancy in temperature is not expected. The measured values appear to droop over time, with a value of 1% or less indicating that the oven should be replenished. The observed droop, as contrasted to the ideal all-or-nothing behavior, is presumably an indication that the oven has a tendency to clog with rubidium, reducing its effective aperture, as has sometimes been observed when refilling it. On other refills, however, an exhausted oven was found with its output unobstructed, rendering generalization difficult.

2.2 Zeeman Slower

As noted above, the atoms emitted by the oven possess an average forward velocity of around 300 m/s, much larger than the $\sim$ 20 m/s capture velocity of the MOT. Consequently, only the very low end of the velocity distribution will be captured unless the atoms are decelerated before arrival at the MOT. In order to produce such a deceleration, we employ a Zeeman-tuned slower, which operates as follows. A circularly polarized laser beam of about 20 mW tuned 530 MHz to the red of the cycling $F = 2, m_F = -2$ to $F' = 3, m_F = -3$ transition propagates anti-parallel to the atomic beam; atoms with sufficiently large velocity are Doppler-shifted into resonance with the laser and scatter many laser photons. In order to maximize the initial number of atoms in the $F = 2$ level and to give any atoms that are inadvertently pumped to the $F = 1$ level during deceleration an exit pathway, a second laser detuned 257 MHz to the red of the $F = 1 \rightarrow F' = 2$ transition co-propagates with the slowing laser. This is referred to as the slower-repump laser. The reason for its specific detuning, which has been chosen by a empirical optimization, is unclear. For each atom-photon scattering event, the incoming photon has a small but nonzero momentum opposed to that of the atom, while the outgoing (or spontaneously emitted) photon may have an arbitrary momentum. As a result, the net effect of the scattering event is to reduce the forward momentum of the atom by $\hbar k$ and also to impart a random step in momentum of size $\hbar k$. After scattering many photons, the atoms will have decelerated, with the result that their Doppler shift no longer puts them in resonance with the slowing laser. The atoms may be kept near resonance as they decelerate if the Doppler shift is compensated by a Zeeman shift:

$$\omega_0 \approx \omega_L + kv + \mu B/\hbar, \tag{2.11}$$

where μ is the difference in magnetic moments between the excited and ground states, equal to $-\mu_B$ in the present case, ω_0 is the zero-field resonance frequency, and ω_L is the laser frequency. For uniform deceleration, $v^2 = v_0^2 - 2az$, then, the desired field profile is of the form

$$B = \frac{\hbar k v_0}{\mu}\sqrt{1 - 2az/v_0^2}. \tag{2.12}$$

This simple picture captures the essence of the functioning of the Zeeman slower. More precisely, one may calculate the average momentum imparted to an atom that scatters a σ_- photon. The kinematics of the atom-photon scattering event are easily analyzed in the center-of-momentum frame, where before scattering the atom carries momentum $\hbar k\hat{\mathbf{z}}$ and the photon an equal and opposite momentum $-\hbar k\hat{\mathbf{z}}$. After the collision, momentum conservation requires that the momenta still be equal and opposite, while energy conservation requires that the momenta have the same magnitude $\hbar k$ as before. The outgoing momenta can thus be written as $\pm k\,(\cos\theta\hat{\mathbf{z}} + \sin\theta\cos\phi\hat{\mathbf{x}} + \sin\theta\sin\phi\hat{\mathbf{y}})$. In the lab frame, the atom has incoming momentum $p\,\hat{\mathbf{z}}$, so that the non-relativistic velocity difference between the lab and center-of-mass frames is $(p - \hbar k)\hat{\mathbf{z}}/m$, and the outgoing atom momentum in the lab frame is

$$(\hbar k\cos\theta - \hbar k + p)\hat{\mathbf{z}} + \hbar k\sin\theta\cos\phi\hat{\mathbf{x}} + \hbar k\sin\theta\sin\phi\hat{\mathbf{y}},$$

so that the change in the atomic momentum is

$$\Delta p = \hbar k\left((\cos\theta - 1)\hat{\mathbf{z}} + \sin\theta\cos\phi\hat{\mathbf{x}} + \sin\theta\sin\phi\hat{\mathbf{y}}\right).$$

The angle θ was defined in the center-of-mass frame, while the spontaneous-emission pattern of the atom is most naturally analyzed in the rest frame of the atom. Because the momentum of the atom, of order 10^{-18} g-cm/s, is so much larger than that of the photon, which is of order 10^{-22} g-cm/s, however, the center of mass frame differs from the atom's rest frame only by a velocity of $\sim 1\,\mathrm{cm/s}$. For this velocity, the fractional change in the axial component of the photon wave vector is less than 10^{-10}, which implies that negligible error is made in treating the photon recoil angle as unchanged from one frame to the other. In the normal functioning of the Zeeman slower, both the optical excitation and subsequent decay occur on a σ_- transition. The two linear polarizations for a photon spontaneously emitted at polar angle θ and azimuthal angle ϕ can be written

$$\begin{aligned}\hat{\epsilon}_1 &= -\sin\phi\hat{\mathbf{x}} + \cos\phi\hat{\mathbf{y}}\\ \hat{\epsilon}_2 &= \cos\theta\cos\phi\hat{\mathbf{x}} + \cos\theta\sin\phi\hat{\mathbf{y}} - \sin\theta\hat{\mathbf{z}}.\end{aligned} \tag{2.13}$$

The allowed polarization for decay is $\hat{\epsilon}_- = (\hat{\mathbf{x}} - i\hat{\mathbf{y}})/\sqrt{2}$, which allows us to compute the overlap as a function of the decay angle

$$\begin{aligned}\hat{\epsilon}_-^* \cdot \hat{\epsilon}_1 &= -ie^{i\phi}\\ \hat{\epsilon}_-^* \cdot \hat{\epsilon}_2 &= \cos\theta e^{i\phi}.\end{aligned} \tag{2.14}$$

The sum of the squares of these amplitudes gives the emission pattern

$$P = \frac{3}{16\pi}(1 + \cos^2\theta), \tag{2.15}$$

where the normalization factor is chosen so that $\int d\Omega P = 1$. Performing an emission-weighted average over θ, then, yields

$$\langle \Delta p \rangle = \frac{3}{16\pi} \int d\Omega (1 + \cos^2\theta) \Delta p = -\hbar k \quad . \tag{2.16}$$

This shows that the simple argument given earlier is correct. On average exactly $\hbar k$ of momentum is removed from the decelerating atom. It is also of interest to compute the mean squared transverse momentum:

$$\langle \Delta p_\perp^2 \rangle = \frac{3}{16\pi} \int d\Omega (1 + \cos^2\theta) \hbar^2 k^2 \sin^2\theta = \frac{3}{5}\hbar^2 k^2, \tag{2.17}$$

from which it is readily observed that the r.m.s. transverse momentum acquired after scattering N photons is

$$\Delta p_\perp^{\text{r.m.s.}} = \sqrt{\frac{3N}{5}}\hbar k. \tag{2.18}$$

This is only slightly smaller than the value of $\sqrt{\frac{2N}{3}}\hbar k$ which one would calculate if the spontaneous emission were isotropic. This transverse heating is a matter of concern if it allows atoms to escape from the capturable solid angle of the beam. The transverse distance along one dimension diffused by an atom in its random momentum-space walk is

$$x = \sum_{j=1}^{N} \sum_{i=1}^{j} \epsilon_i \Delta v \Delta t, \tag{2.19}$$

where N is the number of photons scattered, $\Delta v = \hbar k/m\sqrt{3/10}$, and $\Delta t \approx 2/\Gamma$, while ϵ_i is a random variable whose average is 0 and whose r.m.s. value is 1. Thus

$$\langle x^2 \rangle = \frac{1}{2}(\Delta v \Delta t)^2 N^3, \tag{2.20}$$

assuming $N \gg 1$. The r.m.s. excursion of an atom from the beam center, using numbers of $N = 5 \times 10^4$, $\Delta v = 0.32\,\text{cm/s}$, and $\Delta t = 53\,\text{ns}$, is about 1.3 mm, which is rather insignificant. As a result, any loss due to transverse heating must result primarily from ballistic expansion of the slowed gas upon exiting the slower. From equation (2.18), the r.m.s. transverse velocity along one dimension is $\approx 70\,\text{cm/s}$, while the exit velocity of the slower is $\approx 20\,\text{m/s}$, so that the resulting loss in collimation is substantial, but presumably not detrimental since it remains smaller than the acceptance angle of the MOT as viewed from the exit of the slower ($\approx 1\,\text{cm}/10\,\text{cm}$).

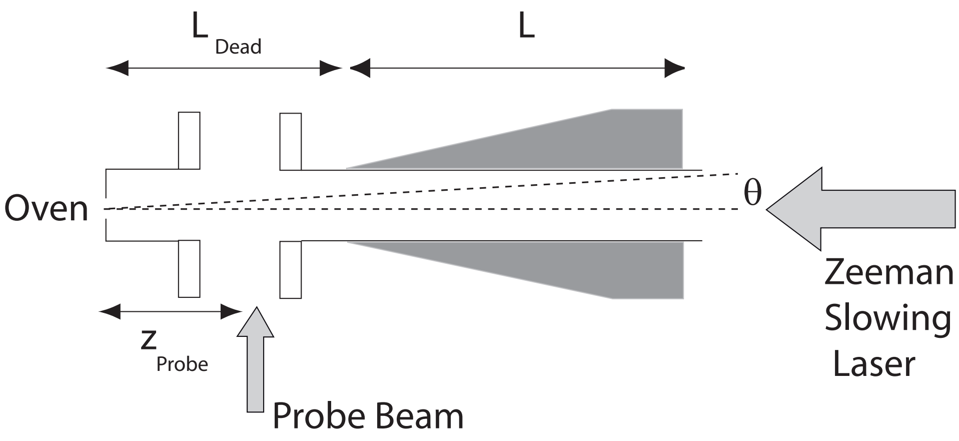

Figure 2.3. Zeeman slower diagram, showing angles and distances employed in the text. L_{dead} is the distance from the oven to the start of the slower, L is the length of the Zeeman slower, θ is the angle of an atomic trajectory with respect to the axis of the atomic beam or Zeeman slower. z_{probe} is the distance from the oven at which the atomic beam is characterized in section 2.1.1.

It is apparent, however, that lowering the exit velocity from the slower below about 10 cm/s would quickly begin to erode the capturable flux. This, in conjunction with the capture efficiency of the MOT as a function of velocity, is a likely explanation of the observed optimum flux around a final slower velocity of 20 m/s.

2.2.1 Deceleration and Optimal Length

A saturated two-level atom (which the $|F = 2, m_F = \pm 2\rangle \longrightarrow |F' = 3, m'_F = \pm 3\rangle$ transitions of ^{87}Rb well approximate) has on average half its population in the excited state, and consequently scatters photons at a rate $\gamma/2$, where $\gamma = 2\pi \times 6\,\text{MHz}$ is the natural linewidth. This implies a maximum deceleration (averaged over many photon scattering events) of

$$a_{\max} = \frac{\gamma}{2}\frac{\hbar k}{m}. \tag{2.21}$$

For a deceleration length L, this defines a maximum velocity $v_{\max}$ that can be brought to rest via $\frac{1}{2}mv_{\max}^2 = ma_{\max}L$, i.e.,

$$v_{\max} = \sqrt{2a_{\max}L}. \tag{2.22}$$

Given the beam distribution, equation (2.2), it is straightforward to calculate the number of slowable atoms per unit time

$$\begin{aligned}\frac{\mathrm{d}N_{\mathrm{slowed}}}{\mathrm{d}t} &= \int \mathrm{d}\Omega \int_0^{v_{\max}} v^2 \mathrm{d}v\, n_{\mathrm{oven}}\, A\, v \cos\theta \left(\frac{m}{2\pi\tau}\right)^{3/2} e^{-mv^2/2\tau} \\ &= \frac{n_{\mathrm{oven}} A \bar{v}}{4} \sin^2\theta \left(1 - e^{-mv_{\max}^2/2\tau} - \frac{mv_{\max}^2}{2\tau} e^{-mv_{\max}^2/2\tau}\right), \qquad (2.23)\end{aligned}$$

where θ is the maximum angle from the atomic-beam axis that an atom can take without being lost from the slower [44] and $\bar{v}$ is defined below equation (2.2). For a fixed final-aperture radius r, this angle is given approximately by $\theta = r/(L + L_{\mathrm{dead}})$. L_{dead} is the length of the "dead space", shown in figure 2.3 i.e., the region after the oven but before the Zeeman slower. Such a space is typically necessary for the maintenance of good vacuum in the main chamber: since the differential pressure across the Zeeman slower cannot be greatly larger than 1000, vacuum pumps are needed to keep the oven chamber at around 10^{-8} torr. The flux of slow atoms delivered to the experiment is shown in figure 2.4 for an oven temperature of 100°C (density $6 \times 10^{12}\,\mathrm{cm}^{-3}$) and $L_{\mathrm{dead}} = 55\,\mathrm{cm}$. If the assumption is made that the slower actually decelerates at the maximum deceleration, then it is found that the ideal length is quite short, approximately 60 cm. If, however, the more conservative assumption is made that deceleration occurs at a fraction (say half) of $a_{\max}$, then the ideal length is approximately 90 cm, and the length dependence becomes much softer. The actual length of the Zeeman slower employed for these experiments is 150 cm, considerably in excess of the ideal length. The longer length is helpful for maintaining good vacuum in the main chamber, although other means for attaining the same end (for instance, limiting the conductance with an aperture or tube) without sacrificing flux would no doubt be more efficacious. The slow flux indicated in figure 2.4 for a 150 cm slower is around 4×10^{10} slow atoms per second. Measured short-time loading rates for the MOT have been around 2.5×10^9/sec, indicating that between imperfect slowing efficiency and imperfect MOT capture efficiency, only 10% of the expected number of atoms are caught.

2.2.2 Coil Layout

In the design of the Zeeman slower, it is desirable (1) to minimize the number of independent current supplies required, (2) to ensure that the heat generated may be

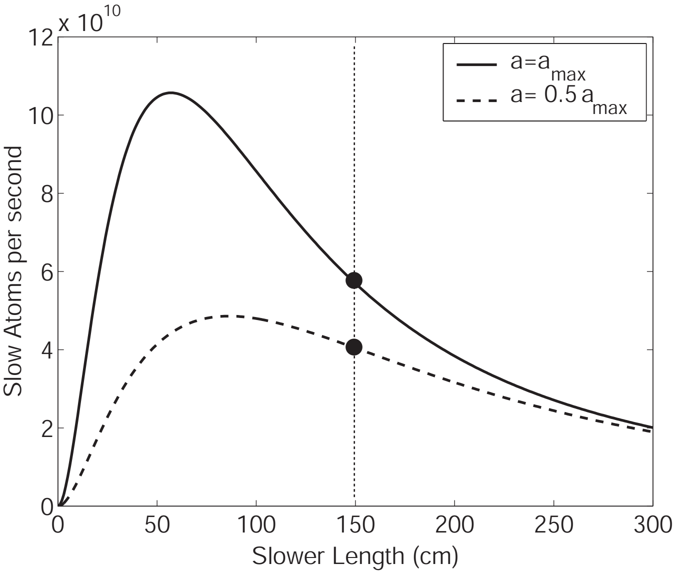

Figure 2.4. Slow flux vs. Zeeman slower length for $a = a_{\mathrm{max}}$ and $a = \frac{1}{2}a_{\mathrm{max}}$. The parameters for these plots are as follows: the dead space before the slower is $L_{\mathrm{dead}} = 55\,\mathrm{cm}$, the total oven density $n_{\mathrm{oven}} = 6 \times 10^{12}\,\mathrm{cm}^{-3}$, of which 27.8% is $^{87}\mathrm{Rb}$. The oven aperture is $0.2\,\mathrm{cm}^2$, and the mean velocity is $\bar{v} = 3 \times 10^4$. The radius of the clear aperture at the output of the slower is $r = 1\,\mathrm{cm}$. The dotted vertical line indicates the actual 150 cm length of the Zeeman slower that was built.

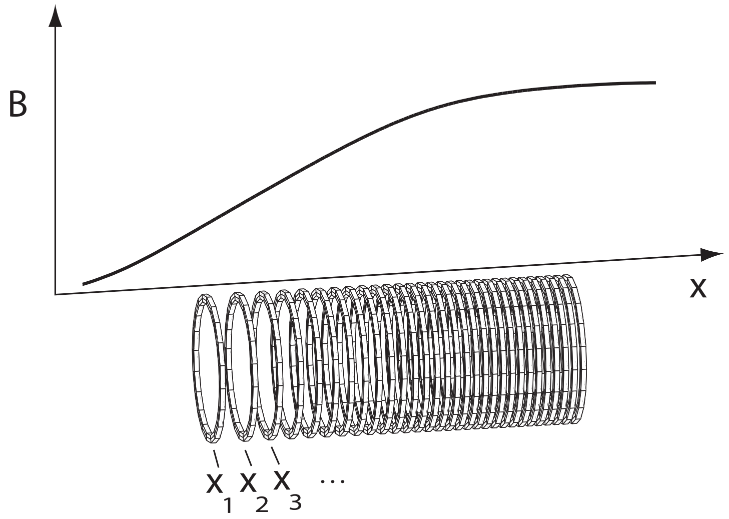

Figure 2.5. Illustration of a representative sample of the Zeeman slower coil, showing that a steady diminution in the turn spacing can produce a local gradient in the axial field. The turns, spaced according to equation (2.25), gradually become more tightly spaced As the turn density increases, so does the solenoid field, in an approximately linear fashion. The field along the axis is plotted in the background.

readily dissipated, and (3) to allow some margin of error in the maximal obtainable deceleration. The arrangement of coils to produce a field which will lead to near-uniform deceleration may be accomplished in several ways. The Zeeman slower wound for this experiment was designed by means of a simple approximation. The field in a solenoid of N/L turns per length carrying a current μ_0 is $B = \mu_0 IN/L$. Although derived for an infinite solenoid, this result depends only on the number of turns per length, which can be thought of as a locally-varying quantity producing a local field. This field is, however, smeared out over a distance corresponding to the radius of the coil. Consequently, by tailoring the local number of turns per length, one may produce nearly any slow variation in B. A linear gradient of the axial field, for instance, is obtained by approximating the discrete number of turns per length as a continuous variable and solving the differential equation (see figure 2.5)

$$\mu_0 I \frac{\mathrm{d}N}{\mathrm{d}x} = B'x, \tag{2.24}$$

which has the solution (reinterpreted in terms of the original discrete variables)

$$x_N = \sqrt{\frac{\mu_0 IN}{B'}}. \tag{2.25}$$

The approximation leading to this result is, of course, only valid provided the number of turns per length does not fall below $\sim 1/R$. Moreover, in any given layer of the winding, the spacing between turns cannot be less than the wire size. In the present instance, the wire used was 1/8" square cross-sectioned copper tubing. These constraints applied to $\mathrm{d}x = \frac{1}{2}\sqrt{\frac{\mu_0 I}{B' N}}$ may be summarized as $d_{\mathrm{wire}} < \mathrm{d}x < R$. For reasonable parameters of $I = 10\,\mathrm{A}$ and $B' = 5\,\mathrm{G/cm}$, this implies $0.15 < N < 7$. Since $x_N \approx 4\,\mathrm{cm}$, the total change in field over the length of the gradient is 20 G. Clearly, then to decelerate nearly to rest an atom initially at $3 \times 10^4\,\mathrm{cm/s}$, i.e., with a Doppler shift of 380 MHz, requires of order 10 such sections of gradient-producing coils in sequence to achieve a field of $380\,\mathrm{MHz} \longleftrightarrow 270\,\mathrm{G}$. When each layer reaches its maximum field, i.e., becomes close-packed, with wire spacing equal to wire size, it may continue the length of the Zeeman slower in the close-packed configuration, providing a uniform bias field on which subsequent gradient layers may build as well as offering a stable mechanical foundation for later windings. In practice, non-uniform precise spacings are difficult to achieve, and the calculated positions serve as guidelines; one tries, for instance, to make the number of turns per length vary smoothly from a minimum to a maximum value over a certain length.

The final design employs two currents, optimized empirically at 12 A and 44 A. The lower current is adequate for the slowly varying field along the largest portion of the slower, but to achieve a reasonable approximation of the near-vertical tangent of the ideal field, small windings at high current are desirable. In the actual experimental setup, a single power supply runs 12 A in series through both sets of coils, while a second supply runs 32 A through the latter set only. Figure 2.6 shows a comparison of measured fields to the ideal profile. The agreement is good, especially considering the fact that the Zeeman slower cools the atoms as well as slowing them, and that this velocity-narrowing force provides some immunity to deviations of the field profile from ideal.

The Zeeman slower coils were wound on a 1.5 in.-diameter, 60 in. long stainless-steel ConFlat vacuum nipple. A large lathe operating on very low speed was employed to facilitate winding and provide sufficient torque to pull the windings tight. For protection of the knife edges and easy gripping by the lathe, additional short half nipples were firmly bolted to either end. One of these was held in the lathe chuck, while the other was held in place but allowed to rotate by three lubricated symmetric radial supports on the tailstock of the lathe. The steel tube was covered with adhesive-coated polyimide tape to reduce the possibility of short circuits to the vacuum chamber. The wire used for the coils was square 1/8 in. tubing (1/32 in. wall thickness), insulated with polyimide. One end of the coil was typically anchored by means of a hose clamp, and the rest then slowly wound, making small adjustments to the number of turns per length. The newly wound coils were temporarily held in place by the use of several tight layers of polyimide tape around their circumference. Once a complete layer was wound, the assembly was removed from the lathe and Duralco^{TM}4525 high-temperature electrically resistant epoxy (manufactured by Cotronics Corporation) was liberally applied to the coil. When the epoxy had hardened, encasing the coil and holding it in place, a new layer could be wound on top. The extraction of leads to the separate coils, some from under as many as 10 other coils, requires special attention. The leads on the oven end of the Zeeman slower were easy to gain access to, but those at the far end of the slower all exit at the same axial position. An attempt was made to have these latter leads exit in a pinwheel pattern, but in practice the windings on the upper layers could not be continued flush with the lower layers without slipping off. This resulted in some degradation of the ideal profile, which had to be compensated by the high-current coils wound later. A next-generation design would probably incorporate a plate against which the coils could be wound flush, and which could optionally be removed once they were epoxied in place. Current was supplied to the coils via 14 A.W.G. wire that was soldered to the copper tubing. As a result of the high thermal conductivity of the tubing, it was difficult to obtain good solder joints without overheating the insulation on the tubing. A more satisfactory approach for a future Zeeman slower would be the use of barrel connectors. Alternatively, since the actual heat generated in the Zeeman slower is small enough to allow air cooling (although for maintenance of good vacuum, water cooling is employed on the bottom layer), a potentially attractive alternative is to use more turns of higher-gauge wire, allowing more precise control over the field profile via the number of turns per length. This would, of course, require higher-voltage, lower-current power supplies than on the present apparatus.

The Zeeman slower has been tested experimentally by probing at 60° with a near-resonant retroreflected laser beam and recording fluorescence as a function of the detuning of the probe. Because of the finite angle between the slower axis and the probe beam, this probing method was sensitive via the Doppler effect to the velocity distribution of the atoms. More particularly, this meant that the counterpropagating beam was not resonant when the forward beam was (and vice versa), but merely provided a convenient frequency marker for zero velocity, to wit, the symmetry point of the observed fluorescence-detuning curve. In order to improve the signal-to-noise ratio, the probe was chopped at a few hundred Hz, and the current from the photodetector was fed into the current input of a lock-in amplifier clocked at the chopping frequency. In the absence of slowing light or slower current, a broad peak in the detected signal versus detuning was observed. For the correct choice of current and detuning of the slower laser, however, a substantial portion of the peak was seen to move closer to the zero-velocity point. From the known sweep range of the probe frequency and the probe angle the frequency could be converted to velocity. The peak output velocity of the slower was around 20 m/s. Peaks at smaller velocities could be generated, but with diminishing signal height. The largest velocity that appeared to be slowed when comparison was made between the unslowed and slowed distributions was approximately 300 m/s, although this was difficult to determine precisely.

A more practical test of the Zeeman slower is its ability to enhance the loading rate of the MOT. In this regard, as noted in section 2.2.1, the present slower falls short of the theoretical prediction, delivering approximately 2.5×10^9 cold atoms per second, whereas a capturable flux roughly 10 times as large might have been expected. This difference may result merely from the combination of several small factors, such as the uniformity of the Zeeman-slower laser beam, its overlap with the slower repump beam, their common overlap with the atomic beam, and the transverse velocity spread of the slowed beam, among others. Another possibility is that the slower only slows efficiently down to a particular velocity, which could be comparable to or marginally higher than the capture velocity of the MOT, compromising the ability of each to perform optimally at any one final velocity. It is worth noting, however, that the number of atoms that can be held in a MOT scales weakly with the loading rate when the optical density of the MOT becomes substantial, and that the condensate numbers of $\sim 2 \times 10^6$ obtained in the present apparatus are on the

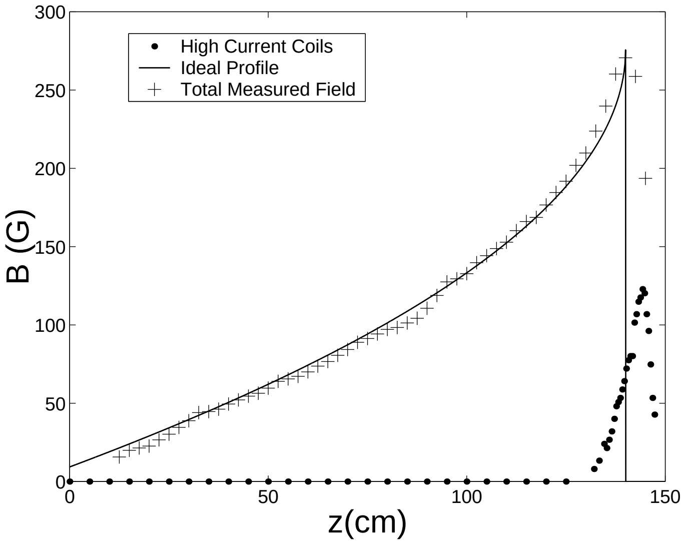

Figure 2.6. Measured slower fields compared to ideal slower field profile. Fields were measured by means of a Hall probe, mounted in a hole on the axis of plastic cylinder, and positioned by means of a long steel rod that had been epoxied to the cylinder. The solid curve shows the ideal field profile. Crosses show the total field due to both high- and low-current coils. Filled circles show the field from high-current coils only.

high end among contemporary ^{87}Rb BEC experiments, so that diminishing returns may be expected in efforts to increase the condensate size by increasing the MOT loading rate. Nevertheless, a higher loading rate is always a good thing, provided it does not increase the loss rate proportionally, and with sufficient optical access and laser power, the MOT beams can be made larger to accommodate a larger number of atoms. Thus, future efforts with a more compact Zeeman slower and less dead space between the oven and the beginning of the slower certainly remain worthwhile.

2.3 Diode Lasers and Tapered Amplifier

The laser system used for these experiments (see figure 2.7) consisted of three external-cavity diode lasers (ECDLs), one near the D2 $F = 2$ lines, one near the D2 $F = 1$ lines, and one near the D1 $F = 1$, as well as a free-running 825 nm diode laser, a slave laser injection locked to the D2 $F = 2$ master, and a tapered amplifier injected in turn by this slave.

The ECDLs which we have built and employed for this apparatus are of a stan-

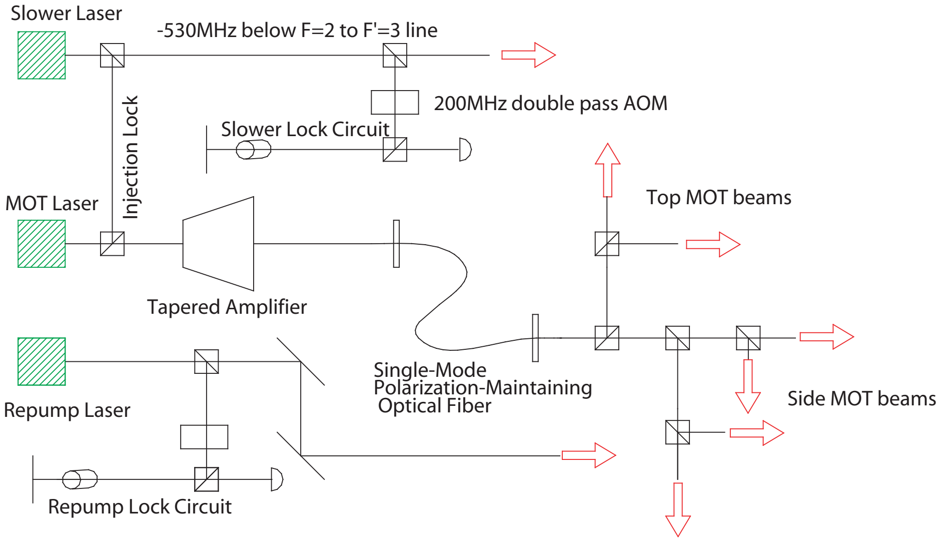
Slower Laser
-530MHz below F=2 to F'=3 line
200MHz double pass AOM
Slower Lock Circuit
Injection Lock
MOT Laser
Top MOT beams
Tapered Amplifier
Single-Mode
Polarization-Maintaining
Optical Fiber
Repump Laser
Side MOT beams
Repump Lock Circuit

Figure 2.7. Optics layout for laser cooling, consisting of two stabilized lasers, and an amplified injection-locked laser for the MOT. An imaging laser on the D1 line is unpictured. The master laser also serves as the Zeeman-slowing laser, while the other laser is resonant with the $F = 1$ hyperfine levels and provides repumping light to both Zeeman slower and MOT.

dard configuration, consisting of a commercially available (Thorlabs) diode housing with collimating lens and strain-relief cable (supplemented by a simple home-built protection diode network) screwed in place on the stationary portion of a commercial mirror mount. Most of the diodes used in the neighborhood of 780 nm have been the Sanyo DL-7140-201S, although the substantially less expensive Sharp GH0781JA2C diode has also been used with satisfactory results. To the movable plate of the mirror mount is affixed a small shelf which carries the grating, clamped in place. The grating employed is typically a UV grating with 1800 lines per inch, which for the correct orientation (polarization) of the diode should diffract around $10-20\%$ of the incident light. Between the fine-threaded screw of the mirror mount and the opposing point on the movable plate is placed a piezoelectric actuator (“piezo”). The actuator is glued to the plate on one end and on the other to a small metal disk which makes contact with and distributes the force from the end of the screw. The piezo is subject to a considerable buckling force from the spring-loaded mirror mount if it is not mounted precisely in line with the screw. Given the pivotal importance of the piezo for the grating stability of the laser, a next-generation ECDL design would probably

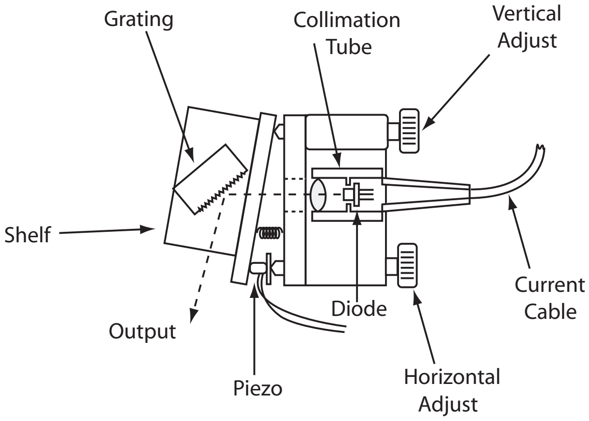

Figure 2.8. External cavity diode laser diagram, showing diode mounted in collimating tube, collimating lens, and grating. Horizontal and vertical adjustments of the grating allow coarse angular control over the diffracted spot from the grating, while in the horizontal direction a piezoelectric transducer provides fine control.

attempt to guard against shear or buckling of the piezo, perhaps by encasing it in a small hollow cylinder.

The entire mirror mount is then attached to a metal plate, which sits atop a thermoelectric cooler for temperature stabilization, and the whole assembly is housed in a metal box to shield against RF noise and against air currents. As a side note, the use of thermistors as temperature-sensing elements has generally been favored over that of AD590 sensors, because of their observed superior immunity to radio-frequency interference. Optimum temperature stabilization is achieved by placing the temperature sensor as close as possible to the cold surface of the TEC. This avoids heat propagation delays which tend to destabilize the feedback circuit. Of course, placing the sensor close to the laser diode is also desirable; together, these considerations make the use of a compact mount desirable. We have also successfully employed two-stage temperature stabilization, in which the entire mount is cooled (or heated) by a TEC, while fine tuning of the diode temperature is achieved by placing a resistive heater wire and a second sensor on the diode collimation tube itself.

The standard alignment procedure for these diodes is first to adjust the collimating lens to make the emitted beam as parallel as possible. Then, with the diode housing in its final location and heat-sink compound reapplied as necessary, the grating is adjusted to send the first-order diffraction spot back into the laser diode. It may be desirable, prior to this step, to adjust the temperature of the diode mount to bring the wavelength as close as possible to the desired wavelength, but depending on the difficulty of reading the wavelength this may be done later, at the expense of additional iterations. It is often useful to perform the initial adjustment of the grating coarsely by manually rotating the grating to be certain that the correct diffraction spot is being sent back. A finer adjustment is accomplished by clamping the grating in place and looking now at the forward-reflected beam with the diode current near its full operating value. When the diffraction spot is very closely aligned to the output of the diode, its reflection from the collimating lens or from the diode itself shows up as a secondary spot in the forward direction which can be adjusted using the screws on the mirror-mount that holds the grating. In the ideal case, this spot can be scanned from one side of the output beam to the other in both the vertical and horizontal directions. Optical feedback is observed when this spot appears to be drawn in to the larger output beam and to disappear. In practice, however, it sometimes occurs that the spot can be scanned only on one side of the beam. In these cases, one can verify that optical feedback is occurring by tuning the grating horizontally with the mirror-mount adjustment and seeing whether the wavelength tunes in a punctuated but nearly monotonic fashion with the grating angle. Another useful method of optimizing the optical feedback and the vertical tilt of the grating, particularly if the diffraction spot is close to being aligned, is to scan the current around the lasing threshold while monitoring the output power. A relatively sharp kink in the output power should be observed at threshold. When optical feedback is present, this kink should move to lower power and should exhibit extreme, flickery sensitivity to the grating alignment. Once optical feedback is achieved, the laser may typically be tuned by 7 nm or more with a combination of temperature and grating adjustments, although if the wavelength is pulled this far intermediate optimizations of optical feedback (e.g., of the vertical tilt of the grating) are likely to be necessary.

The laser is then stabilized by applying an electronic feedback signal to the current and to the piezo. This two-branch feedback is intended to reduce frequency noise below a cut-off frequency of several tens of kilohertz. The error signal for applying feedback is derived from a saturated absorption cell; both pump and probe are frequency-modulated by means of a double-passed acousto-optic modulator (AOM).

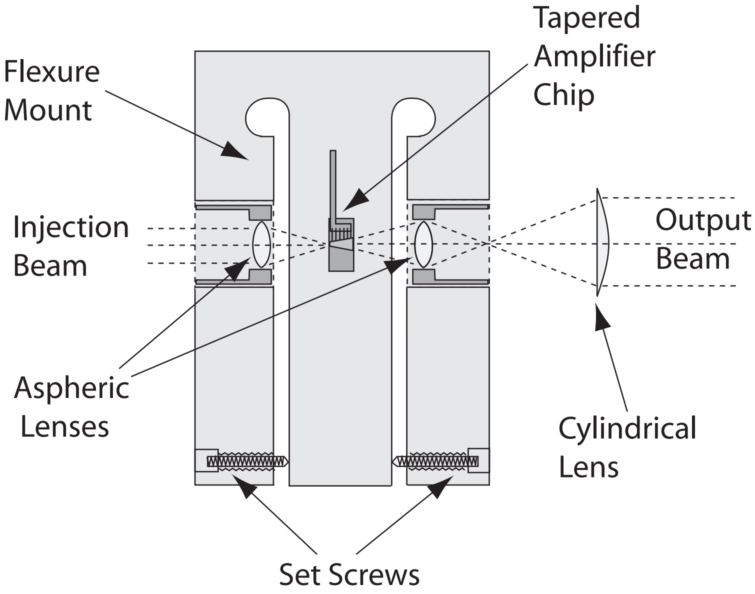

Figure 2.9. Tapered amplifier mount (not to scale), showing input lens (left) and output lens (right), whose axial positions are adjustable by set screws pushing against the elastic aluminum arms of the flexure mount. The output lens collimates the more diverging dimension of the astigmatic output beam. A cylindrical lens forms a Keplerian telescope with the output aspheric lens, collimating the other dimension.

In typical operation, the AOM induces many sidebands of approximately equal power, so that the final demodulated error signal contains contributions from many slightly different paths. Although this is expected to introduce additional noise, the method appears rather robust and stable, and is easy and inexpensive to implement.

The tapered amplifier consists of a home-built mount for a commercial amplifier chip (from the company Eagleyard). The chip is contained in a 2.75 mm C-mount package, and is specified at a maximum power of 500 mW. It is placed against a thin layer of indium foil and screwed into a metal mount so that the active region of the chip protrudes slightly above the top of the mount, allowing the injection beam to be coupled in. This coupling is accomplished by means of two aspheric microscope objective lenses, one on the input and one on the output. Each is mounted in a steel tube which slides securely into a small-tolerance bore traversing an aluminum block and is held in place by a set screw. These blocks are glued (using Torr Seal ®) onto two independent arms of a flexure mount with the aspheric lenses projecting inward toward the tapered amplifier chip. Additional setscrews in the arms of the flexure mounts allow fine positioning of the lenses along the optic axis. This entire assembly is mounted on top of a Peltier cooler, used to stabilize the temperature, and three Melcor spacers, which all reside on a large aluminum block. This final block, about 4 in. high, is clamped via standard baseplates to the optical table. Current is delivered by a Lightwave Electronics current driver across a protection board consisting of one reverse diode and several forward diodes in parallel (the number chosen to make the total diode drop exceed the operating voltage, $\sim 2.5\,\mathrm{V}$, of the tapered amplifier chip). An insulating feedthrough in the mount facilitates access to the cathode wire of the chip; its anode is in electrical contact with the C-mount body, so that a lead may be attached to the mount as convenient. The use of indium foil rather than the previously employed heat-sink compound, to maximize thermal contact, thus improves electrical contact as well, which otherwise would be guaranteed only by the mounting screw. It is not always easy, moreover, to ensure that heat-sink compound is applied in exactly the quantity and at exactly the points where it is desired in the tight confines of the tapered amplifier mount.

The tapered amplifier mount as implemented has functioned well for approximately three years, yielding $\approx 300\,\mathrm{mW}$ output power for $25\,\mathrm{mW}$ input at an operating current of $1100\,\mathrm{mA}$, although several chip replacements have been necessary. The thermo-mechanical stability of the mount when subjected to a change in heat load (e.g., when the amplifier current is turned on in the morning) is imperfect, requiring several small adjustments of the input coupling over the first half hour of operation to maintain good efficiency and occasional (every ~ 2 hours) such adjustments thereafter. This is presumed to result from gradual heating of the large aluminum base, which serves both as structural support and as heat sink for the Peltier cooler. A next-generation design might seek to stabilize the temperature of the base in addition to that of the amplifier chip, or perhaps more simply to separate the rôles of heat sink and support. For operation above room temperature, it is furthermore quite possible that no heat sink is needed.

When aligning the tapered amplifier after chip replacement, the input lens block is positioned first by means of a cylindrical lens holder or other clamp mounted on a triaxial translation stage. Running the amplifier at a $\sim 600\,\mathrm{mA}$, so that amplified spontaneous emission (ASE) is clearly visible in both the forward and backward directions with an infrared viewer, one attempts to position the lens so that the backward-emitted ASE is nearly collimated and centered on the objective. If the vertical position of the lens is not correct, it is necessary to mill, sand, or shim the block to the correct height at this stage. The backward ASE typically takes the form of three spots, the middle of which should be aligned and approximately mode-matched to the injection beam. The shaping of the injection beam is accomplished with spherical optics (and the aspheric objective lens), since the aspect ratio of the input to the tapered amplifier chip is small. It is important also that the injection beam be of the correct linear polarization to obtain amplification. Amplification should be apparent on an infrared up-conversion card even with only a few mW input power. After reasonable but not exhaustive optimization, which would in any event have to be repeated, the input lens block is epoxied in place, and the lens-holder clamp is removed when the epoxy has hardened. The same procedure is then followed for the output lens block, which is positioned to make the highly astigmatic amplified output beam (which differs in shape and collimation from the output ASE) approximately parallel in the vertical (i.e., rapidly diverging) direction. In the less diverging horizontal direction, this lens will serve to focus the beam, whose size will at some point beyond the focus then equal the vertical size. A cylindrical lens placed at this point, with its focal length chosen to collimate the horizontal direction, will then produce a non-astigmatic beam of near-unit aspect ratio. Fiber-coupling efficiencies as high as 70% have been obtained with beams shaped in this manner.

2.4 Cold Plates

Given the high vapor pressure of rubidium at room temperature and the relatively large flux of unslowable atoms from the oven (that is, atoms emitted at large angle to the beam and those of the other isotope), it is desirable for the maintenance of ultra-high vacuum to provide cold surfaces upon which excess rubidium may be accumulated. In the present chamber, one such surface has been provided in the oven chamber surrounding the atomic beam shutter and one in the main chamber, positioned to catch atoms re-emitted after striking the slower window (see figure 2.1). To maintain temperatures of -20° C to -40° C with minimal human intervention, a system of thermo-electric cooler (TEC) stacks was implemented. A TEC operates on the Peltier effect, according to which the flow of electric current through a conductor induces a corresponding heat flow. The ratio of the heat current to the electric current, denoted Π, has units of volts and depends on the material. By running a current through a wire consisting of alternating regions of large Π, in which current flows in one spatial direction, with regions of small Π, in which it flows in the opposite direction, a relatively modest current of a few Amperes can produce substantial heat flows of tens of Watts. This arrangement may be thought of as the operation of many small, individual Peltier devices electrically in series and thermally in parallel, and it is the basis of most commercial TECs.

The temperature difference obtainable with an arrangement of TECs is limited by the competition between the Peltier effect and the inevitable Ohmic heating. Indeed, the integral form of the heat continuity equation, written for a nearly isolated block of thermal conductor on the cold side (temperature T_0) of a TEC whose hot side is connected to an ideal heat sink at temperature T_1, is

$$K(T_1 - T_0) = \Pi I - \frac{1}{2} R I^2 - \dot{Q}, \tag{2.26}$$

where I is the TEC current, R its electrical resistance, K its thermal conductance, and $\dot{Q}$ the heat load to the block. In general, $\dot{Q}$ will depend on T_0, but if the heat load is small or independent of T_0, the temperature difference is easily seen to reach a maximum at

$$I_{\max} = \frac{\Pi}{R}, \tag{2.27}$$

corresponding to a temperature difference of

$$T_1 - T_0 = \frac{\Pi^2}{2KR} - \frac{\dot{Q}}{K}. \tag{2.28}$$

For reasonable TEC parameters of $\Pi \sim 6\,\mathrm{V}$, $R \sim 0.5\,\Omega$, and $K \sim 1W/{}^\circ\mathrm{C}$, this corresponds to a maximum temperature difference of approximately 40°C at zero heat load, or $T_0 \sim -20^\circ\mathrm{C}$ if the heat sink is at room temperature. To achieve larger temperature differences, and thus lower vapor pressures of rubidium, one might consider combining several TECs in series and/or in parallel. The general calculation of what temperatures may be obtained by such a combination is analogous to that for a single TEC, but has proven to be of some practical utility and will be reproduced here in outline for reference.

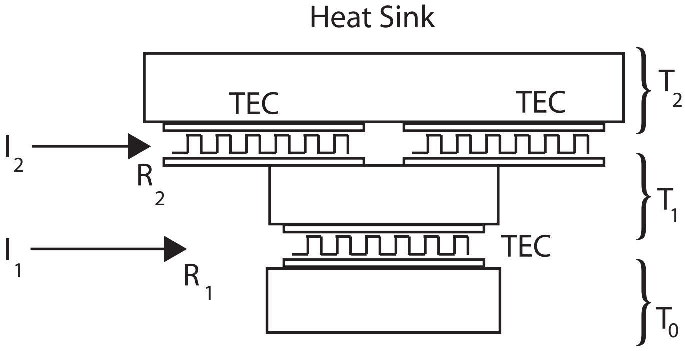

Figure 2.10. Diagram of TEC Stack. Peltier coolers of electrical resistance R_n and thermal conductance K_n carry currents I_n, maintaining a temperature difference of $T_n - T_{n-1}$ across their plates. Stacking TECs allows larger temperature differences, but upper layers must contend with a large heat load from lower layers.

Consider now a system of alternating layers of TECs and thermally conductive blocks, as shown in figure 2.10. The heat flow equation for the n-th block is

$$\Delta_n - \Delta_{n-1} = \Pi_n I_n - \Pi_{n-1} I_{n-1} - \frac{1}{2} R_n I_n^2 - \frac{1}{2} R_{n-1} I_{n-1}^2, n > 1 \tag{2.29}$$

$$\Delta_1 = \Pi_1 I_1 - \frac{1}{2} R_1 I_1^2 - \dot{Q}, \tag{2.30}$$

where $\Delta_n \equiv K_n(T_n - T_{n-1})$ and the quantities I_n,R_n, K_n and Π_n refer respectively to the electrical current, resistance, thermal conductance, and Peltier coefficient of the n-th TEC. This system of equations is easily solved for Δ_n as

$$\Delta_n = \Delta_1 + \sum_{i=2}^{n} (\Delta_i - \Delta_{i-1}) = \Pi_n I_n - \frac{1}{2} R_n I_n^2 - \sum_{i=1}^{n-1} R_i I_i^2 - \dot{Q}, \tag{2.31}$$

from which the total temperature difference is obtained

$$\begin{aligned} T_N - T_0 &= \sum_{n=1}^{N} \left(\frac{\Delta_n}{K_n} \right) \qquad (2.32) \\ &= \sum_{n=1}^{N} \frac{\Pi_n I_n}{K_n} - \sum_{n=1}^{N} \frac{1}{K_n} \left(-\frac{1}{2} I_n^2 R_n + \sum_{i=1}^{n} R_i I_i^2 \right) - \dot{Q} \sum_{n=1}^{N} \frac{1}{K_n}. \end{aligned}$$

Equating partial derivatives with respect to the N currents to zero yields

$$I_j^{\max} = \frac{\Pi_j}{R_j} \left(1 + 2 \sum_{n=j+1}^{N} \frac{K_j}{K_n} \right)^{-1}. \tag{2.33}$$

from which by straightforward algebra the maximum temperature difference is

$$T_N - T_0 = \frac{1}{2}\sum_{j=1}^{N}\frac{\Pi_j^2}{R_j K_j}\left(1 + 2\sum_{n=j+1}^{N}\frac{K_j}{K_n}\right)^{-1} - \dot{Q}\sum_{j=1}^{N}\frac{1}{K_j}. \tag{2.34}$$

For practical purposes, it is often convenient to cascade several of the same type of TEC, rather than to choose different components for each layer. In the special case that each layer n consists of a number w_n of identical TECs, each having specifications Π, K, and R as defined above, the maximum temperature difference may be simply parametrized as

$$T_n - T_0 = \alpha\frac{\Pi^2}{2KR} - \beta\frac{\dot{Q}}{K}, \tag{2.35}$$

where, noting that w_n TECs have resistance $w_n R$, Peltier coefficient $w_n\Pi$, and thermal conductance $w_n K$, we identify

$$\alpha = \sum_{n=1}^{N}\left(1 + 2\sum_{j=n+1}^{N}\frac{w_n}{w_j}\right)^{-1}, \quad \beta = \sum_{n=1}^{N}\frac{1}{w_n}. \tag{2.36}$$

Values of α and β for various TEC configurations are shown in table 2.1. To maximize the temperature drop, it is clearly advantageous to choose configurations for which α is large and β is small. For situations in which $\dot{Q}$ is small, maximizing α will be the primary concern, while minimizing β is of greatest importance when the heat to be dissipated is substantial. As noted earlier, the foregoing analysis requires modification when the heat load itself depends on the cold temperature, as would be the case for thermal leakage in a vacuum feedthrough, for instance. In this case, $\dot{Q} = K_{\text{leak}}(T_N - T_0)$, assuming that the heat load is due to leakage conductance to a point at the same (e.g., ambient) temperature as the heat sink. Referring to equation (2.32), it is apparent that the calculations for the optimal currents are unmodified, but that the total temperature difference is reduced by the factor $1 + K_{\text{leak}}\sum_{n=1}^{N}\frac{1}{K_n}$.

In practice, the configuration we have typically employed is the two-layer $w_1 = 1$, $w_2 = 2$ stack pictured in figure 2.10, for which $I_1^{\max} = I_{\max}/2$ and $I_2^{\max} = I_{\max}$ in terms of the single-TEC optimum of equation (2.27). A copper block matching the size of the single TEC is machined and endowed with a blind hole closely matched to the diameter of the copper thermal (or high-current) feedthrough. The feedthrough is usually slip-fitted to the block with heat-sink compound in the interstitial space, although the two pieces could be brazed together if provision were made to avoid overheating of the feedthrough. The TEC of the first layer is then placed on the other side of the copper block, followed by a second copper block approximately the size of two TECs side by side, followed by the two TECs themselves. Finally another larger water-cooled copper plate is placed on the hot side of the second layer of TECs. Heat-sink compound is applied between each thermal junction. The water cooling is effected by running tap water at a few tens of p.s.i. through copper tubes that have been brazed to the top copper plate in a meandering pattern. In order to hold the entire TEC assembly together without thermally short-circuiting the hot and cold surfaces, a system of thermally insulating clamps is sandwiched around the TEC stack. In one case, this consisted of Plexiglas plates joined by long steel bolts, while in another specially machined and threaded Delrin rods were used to provide compressional force between the plates of the TEC stack. To achieve the largest temperature differences, it is of course also important to insulate the entire cold area well. The TEC stacks on the present apparatus use flexible polyurethane foam sheets, cut and layered to enclose the stack as tightly as possible. Even with good insulation, however, the TEC assembly ices up over time, making it desirable to use TECs that have been sealed against moisture. Conduction by the ice is probably responsible for the gradual rise in the minimum temperature attainable from around -40°C for a newly optimized system to -20°C or -15°C for an assembly entirely encased in ice. To avoid such a decline in performance, a promising approach would be to enclose the assembly in a airtight plastic box with gasketed feedthroughs for the cooling water, electrical connections, and the thermal vacuum feedthrough. Most of these components are readily commercially available. The initial moisture content of the box could be extracted with desiccant packets. In conjunction with a titanium sublimation pump, however, a cold-plate temperature of -20°C has been adequate to achieve trapping lifetimes over 100 seconds, which is more than sufficient to allow creation of large BECs. The cold plate in the main chamber is tubular in form, surrounding the path of the atomic beam on the far side of the MOT region and designed to capture atoms (^{85}Rb and unslowed ^{87}Rb) which have struck the window by which the Zeeman-slowing laser enters. Although the total area of the cold plate is only $\sim 30\,\mathrm{cm}^2$, it occupies a large solid angle for atoms reflected from this window,

No./Layer				Coefficients	
w_1	w_2	w_3	w_4	α	β
1	0	0	0	1	1
2	0	0	0	1	0.5
1	1	0	0	1.33	2
1	2	0	0	1.5	1.5
2	1	0	0	1.2	1.5
1	1	1	0	1.5	3
1	3	0	0	1.6	1.33
2	2	0	0	1.33	1
1	1	2	0	1.75	2.5
1	1	1	1	1.68	4
4	0	0	0	1	0.25
2	3	0	0	1.43	0.833
2	4	0	0	1.5	0.75
1	2	3	0	1.8	1.83

Table 2.1. α and β for different numbers w_n of identical TECs in layer n.

so that at low temperature it should absorb with high probability a large fraction of the unslowed atomic beam.

2.5 Magneto-Optical Trap

The Magneto-Optical Trap (MOT) used in this work was of a standard configuration, consisting of six independent ~ 1 in.-diameter laser beams, each of around 15 mW, tuned 20MHz below the $F = 2$ to $F' = 3$ D2 transition of ^{87}Rb . These laser beams provide the trapping and cooling force, in conjunction with a quadrupolar gradient of 15 G/cm obtained by passing a current of 50 A through one curvature and one antibias coil of the magnetic trap (see section 2.6 and figure 2.15 below). The atoms are maintained in the $F = 2$ hyperfine level by a repumping laser resonant with the $F = 1 \rightarrow F' = 2$ transition. Approximately 8 mW total MOT repump laser power is combined with the four horizontal MOT beams, giving atoms off-resonantly scattered on the $F = 2 \rightarrow F' = 1, 2$ transitions into the $F = 1$ level a rapid optical-pumping pathway back to the $F = 2$ level, on which the cooling and trapping MOT laser operates.

In the neighborhood of 5×10^9 atoms are typically collected in a loading time of around 10 s. They are then compressed, roughly along the lines suggested by Lewandowski et al.[45], in a two-step process. In the first step, lasting 10 ms, the magnetic gradient is reduced to 7.5 G/cm and the MOT repump laser power from 8 mW to 130 μW. Further compression is achieved in a second stage, lasting 70 ms, in which the MOT repump power is reduced further to 55 μW and the MOT detuning increased to ~ -40 MHz at the same magnetic gradient. The shim fields provided by Helmholtz coils are adjusted at this stage to position the compressed MOT (CMOT) at the center of the magnetic trap. Then the magnetic gradient is switched off altogether and the MOT beams are reduced in power to ≈ 60 mW total and detuned to approximately -60 MHz and, while the repump power remains at 55 μW. This "optical molasses" step, lasting 7.5 ms, brings the temperature of the cloud to around $60 - 70\,\mu$K. Lower temperatures have been attained, but at the expense of smaller atom numbers or lower densities. Finally the repump beam is extinguished completely for 1 ms, followed 2ms later by the MOT. In this short period, nearly 100% of the atoms are optically pumped into the $F = 1$ hyperfine level. After 2 ms further delay, the magnetic trap is turned on in the uncompressed configuration, as described below.

Optimization of the MOT is accomplished in various ways. Initial alignment is performed by placing a centered iris in the MOT beams. On each of the horizontal axes, one of the counterpropagating pairs is chosen at random and walked into alignment with the vacuum viewports on either side, i.e., the beam is horizontally centered on the viewports and approximately 1/4 in. above center in the vertical direction. The opposing beam is then walked into alignment with the already aligned beam. The procedure for the vertical beams is similar, but more difficult because of the poor visibility and larger size of the viewports on which the beam should be centered. Various techniques have been employed to circumvent this difficulty. One method involves use of the small two-dimensional MOT that is formed with both vertical beams blocked. If fluorescence from this 2d MOT can be imaged, then the bottom-up MOT beam can be aligned to the location on the camera where the 2d MOT appears. Similarly, if some number of atoms can be captured in the magnetic trap and then imaged, this should provide an approximate idea of the position of the trap center, which is close to the position of the MOT zero. Still another way of aligning the bottom-up beam is to use its unbalanced radiation pressure to destroy the 2d MOT, with the bottom-down beam blocked. In any of these scenarios, the top-down MOT beam can then be walked to align it to the bottom-up beam. The longer path length between the bottom window and the trap ensures that the propagation direction of the bottom-up beam will be approximately vertical, even if its position on the bottom window is slightly displaced below the chamber. An alternative procedure is to place on the top and bottom viewports cardboard cutouts with holes in the center and around the circumference at the locations of the vacuum bolts. Then walking the two beams until they are aligned to the two center holes and to each other should result in correct alignment. The circular polarizations of the six beams are set so that the two beams along the MOT axis are of one helicity, while those on the other four axes are of the opposite helicity. In practice, we have observed functioning (albeit inferior) of the MOT with rather poor polarization, requiring almost perfectly incorrect polarization to extinguish trapping altogether, so that the complete absence of a MOT may usually be attributed to other causes.

In order to optimize an already functioning MOT, we have found it useful to monitor the fluorescence from the MOT. In particular, to optimize the number and density of the compressed MOT, an experimental cycle has been implemented, consisting of 2 s loading, followed by the usual CMOT stages. Then full-power near-resonant optical molasses (i.e., with no field gradient) is applied; fluorescence from the atoms is collected by a lens at an auxiliary viewport and imaged onto an iris immediately preceding a photodiode. The signal peaks rapidly when the optical molasses is turned on, and then diminishes on a timescale of hundreds of milliseconds. The initial height of the peak is taken as an indication of the number of atoms in the CMOT. It is this quantity that we attempt to maximize, while secondarily remaining cognizant of the rate of decay of the signal, which is a rough indication of temperature or of the degree to which the MOT beams are balanced. Typical signals have been 700 mV, measured through a 9/64 in.-diameter iris on a ThorLabs DET 110 photodiode, terminated by the 1 MΩ resistance of an oscilloscope, with light gathered by a ~ 0.14 numerical-aperture, near-unit-magnification imaging system. For the specified responsivity of the photodiode, 0.5mA/mW, this corresponds to about $1.4\,\mu$W optical power. For reference, saturated fluorescence from 10^9 atoms corresponds to about $23\,\mu$W within the 5 millisteradian solid angle of the imaging system, so that this value is quite reasonable, considering that the laser is detuned resonance and that the iris transmits only light from the central portion of the CMOT. The diffusion of atoms in this optical molasses is also monitored on a security camera, which detects fluorescence at 780 nm quite well. The atoms in a well aligned MOT should diffuse uniformly and slowly outward as the fluorescence fades, without swirls or rapid motion in any direction. The basic adjustments once the beams are in apparent alignment are the half-wave plates which control the relative intensities of the six beams. Slight tweaking of the Zeeman-slower laser alignment is frequently beneficial, as is tweaking of the MOT repump beam alignment. A more thorough optimization would include the gradient, the MOT beam power and detuning, the MOT repump power, and the Zeeman-slower currents during each of the stages (loading, CMOT, and molasses) when each is relevant. In the experience of the author, it is seldom necessary or desirable to adjust the alignment of any of the MOT beams to improve the fluorescence signal if they appear satisfactorily aligned to start with. Historically, when frequent unavailing MOT realignments have occurred in the laboratory, it has been an indication of another problem, such as oven failure.

2.6 Magnetic Trap

2.6.1 Theory of Magnetic Trapping

Magnetic traps, though now so routinely used in ultracold atomic physics as to be taken for granted, are nonetheless truly remarkable. They constitute what, described in the abstract, one might reasonably suppose to be a ludicrous idealization: a frictionless, temperature-less container which can defy gravity and hold atoms at a fraction of a microkelvin for upwards of a hundred seconds. The principles and technology of magnetic traps for neutral atoms have been extensively discussed [46–50].

The essential idea is that an atomic spin will adiabatically follow a time-varying magnetic field (or a spatially varying field through which the atom moves), provided the rate of change of the Larmor frequency $\omega_L \equiv \mu B/\hbar$ is small compared to its

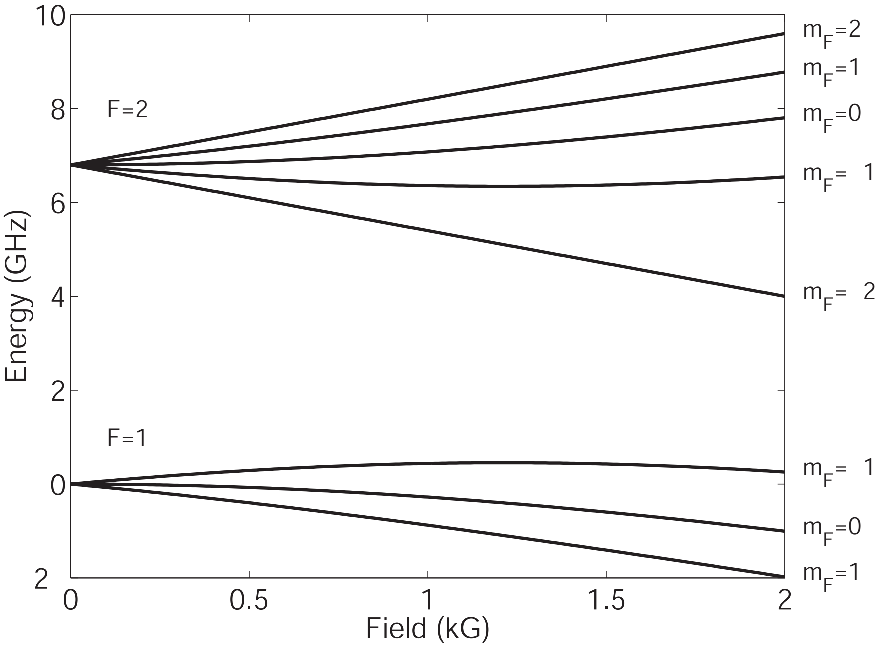

Figure 2.11. Zeeman splitting. States whose energies rise with increasing field (at low field) are magnetically trappable; these are $F = 1, m_F = -1; F = 2, m_F = 1; andF = 2, m_F = 2$. The spinor work of chapter 3 is performed in an optical trap, which allows simultaneous trapping of the three F=1 states.

square, i.e.,

$$\frac{\mathrm{d}\omega_L}{\mathrm{d}t} \sim \mathbf{v} \cdot \nabla \omega_L \ll \omega_L^2. \tag{2.37}$$

As a consequence, the adiabatic potential energy of an atom may be computed at each point by taking the quantization axis along the local field direction. The adiabatic approximation then consists in neglecting the "gauge" terms which appear when derivatives of the spatially-dependent change of basis are taken [51]. Thus the potential energy of an atom with spin F and spin projection m_F is

$$U(\mathbf{r}) = \mu_B g_F m_F |\mathbf{B}(\mathbf{r})|, \tag{2.38}$$

where g_F is the Landé factor. Adiabaticity, then, implies that the potential energy depends only on the magnitude of the field and the (unchanging) magnetic quantum number m_F. A magnetostatic trap is by definition a local minimum of the potential energy, which according to equation (2.38) must correspond to maximum or minimum of $|\mathbf{B}|$, accordingly as $g_F m_F$ is negative or positive.

A crucial further observation due to Wing [52] and generalized by Ketterle et al. [53] is that the magnitude of the magnetic field cannot have a maximum in free space, somewhat analogous to Earnshaw's theorem for the electrostatic potential. This may be seen quite simply from Maxwell's equations. Indeed, in free space $\mathbf{B}$ is divergenceless and curl-free, which implies that each of its vector components obeys Laplace's equation, $\nabla^2 B_i = 0$ for $i = \{x, y, z\}$. Thus each B_i is a harmonic function, and consequently its value at a point is equal to its average on a small sphere S surrounding the point, $\mathbf{B} = \langle \mathbf{B} \rangle_S$. As a result,

$$\langle \mathbf{B}^2 \rangle_S \geq \langle \mathbf{B} \rangle_S^2 = \mathbf{B}^2. \tag{2.39}$$

Since $\mathbf{B}^2$ at each point is less than its average at the surrounding points (at arbitrary radius), it cannot have a maximum. Fortunately, however, in contrast to Earnshaw's theorem, the inequality (2.39) permits $\mathbf{B}^2$ (and hence $|\mathbf{B}|$) to have a minimum. According to relation (2.38), then, only atoms in spin states for which $g_F m_F > 0$, (known as weak-field-seeking states), are magnetically trappable. For ^{87}Rb , the Landé factors are $g_2 = \frac{1}{2}$ and $g_1 = -\frac{1}{2}$, from which the only magnetically trappable states are $F = 1, m_F = -1$, $F = 2, m_F = 1$, and $F = 2, m_F = 2$, as shown in figure 2.11.

A distinction is commonly drawn between magnetostatic traps whose minimum field is zero, referred to as magnetic quadrupole traps, and those which possess a nonvanishing bias field. Of the latter, the simplest and most common type is known as the Ioffe-Pritchard trap, after its inventors [54, 55]. In the neighborhood of a field zero, even very slow atoms will fail to satisfy the adiabaticity condition of equation (2.37), and indeed for a quadrupole trap the tendency of cold atoms to cluster at the trap minimum more than offsets their lower velocity [56], resulting in an increasing rate of nonadiabatic spin flips with diminishing temperature. For this reason, quadrupole traps, although relatively simple to construct, are at present seldom used to confine ultracold atoms. A variant which has enjoyed much greater success is the Time-Orbiting Potential (TOP), in which a small, rapidly rotating bias field is superimposed on the quadrupole field, resulting in an average potential which varies harmonically about a nonzero minimum. The Ioffe-Pritchard trap configuration, however, circumvents the spin-flip problem almost completely without departing from the realm of magnetostatics, since its field is never zero in the trapping region. The form of the Ioffe-Pritchard field in the neighborhood of the trap bottom is given by

$$\begin{aligned}\mathbf{B}(x,y,z) = {} & \left(B'x - \tfrac{1}{2}B''zx\right)\hat{\mathbf{x}} - \left(B'y + \tfrac{1}{2}B''zy\right)\hat{\mathbf{y}} \\ & + \left(B_0 + \tfrac{1}{2}B''(z^2 - \tfrac{x^2+y^2}{2})\right)\hat{\mathbf{z}}, \end{aligned} \tag{2.40}$$

where $\hat{z}$ is the symmetry axis of the trap. It is easy to see that for $x = y = 0$, the field has a quadratic nonzero minimum at $z = 0$:

$$|\mathbf{B}(x=0, y=0, z)| = B_0 + \frac{1}{2}B''z^2. \tag{2.41}$$

The fact that this field configuration is confining in the transverse $\hat{x}$ and $\hat{y}$ directions may be appreciated by taking its square magnitude at $z = 0$

$$\begin{aligned}|B(x,y,z=0)| &= \sqrt{\left(B_0 - \frac{B''}{4}(x^2+y^2)\right)^2 + (B')^2(x^2+y^2)} \qquad (2.42) \\ &\approx B_0 + \frac{1}{2}B''_\rho(x^2+y^2),\end{aligned}$$

where we have retained terms of quadratic order and defined the radial curvature:

$$B''_\rho \equiv \frac{B'^2}{B_0} - \frac{1}{2}B''. \tag{2.43}$$

The ability to tailor the radial curvature by changing B_0 is of considerable practical importance (see figure 2.16 and discussion below). Although one might object that

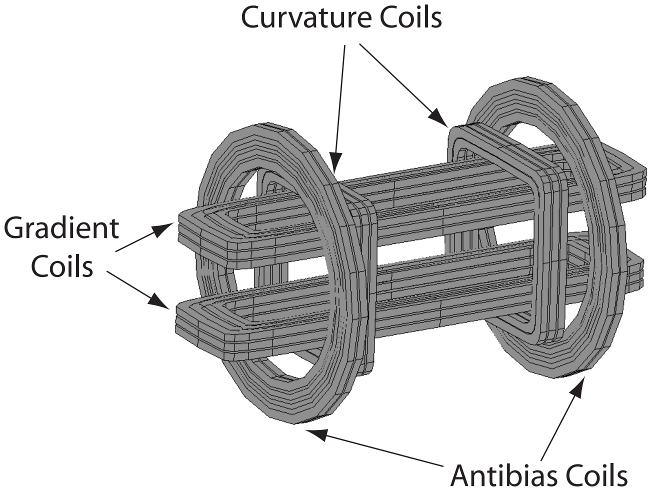

Figure 2.12. Spatial layout of magnetic-trap coils. Long rounded rectangular coils approximate Ioffe bars, yielding a transverse gradient. Square curvature coils provide bias field and axial confinement, while antibias coils produce opposing field to allow independent control of bias and curvature. Coils are wound from square copper tubing and mounted in vacuum by tie-wires to alumina plates (not shown) between each neighboring curvature/antibias coil pair.

the constructive approach outlined above produces a form (equation (2.40)) that is not unique, since any curl-free, divergenceless field may be added to it, in fact the conditions that $\hat{\mathbf{z}}$ be the symmetry axis and that $z = 0$ be the location of the minimum, which we may impose without loss of generality, dictate that any such additional fields will lead merely to a redefinition of the parameters in equation (2.40).

2.6.2 Trap Design and Construction

The inverse problem of determining coils which will produce the Ioffe-Pritchard fields of equation (2.40) is underdetermined, and has been solved in many ways, from the classic design of Ioffe [54, 57] to the more modern cloverleaf [50] and baseball [58] traps and the convertible QUIC trap [59]. The design employed in the present experiment is essentially of the classic Ioffe type. As shown in figure 2.12, it consists of two long (15 cm × 3 cm) nine-turn rounded rectangular coils at ±1.5 cm distance from the trap center. These coils produce a two-dimensional quadrupole field, i.e., a radial gradient. Surrounding these gradient coils are a pair of 5.5 cm-edge rounded-square six-turn curvature coils at approximately ±3.5 cm axial distance from the trap center and a larger pair of 9 cm-diameter eight-turn circular "antibias" coils at ±4.8 cm distance. These four axial coils are wired together so that the bias field at the center of the trap is in the same direction for the two curvature coils, but in the opposite direction for both of the two antibias coils. More specifically, the trap is designed so that the net bias field produced when a current is run through these four coils in series is close to zero, in which case magnetic field noise will be suppressed by the ratio of the actual bias field to the bias field due to either pair of coils separately. The gradient coils in this design are the innermost of the three pairs of coils and are so placed to maximize the transverse confinement in the trap. Indeed, the peak density of a classical gas of N atoms in a harmonic trap at temperature $\tau = k_B T$ is

$$n = N \left(\frac{m}{2\pi\tau} \right)^{3/2} \bar{\omega}^3, \tag{2.44}$$

where $\bar{\omega} = (\omega_x \omega_y \omega_z)^{1/3}$. Adiabatic compression of a harmonic trap preserves the ratio $\bar{\omega}/\tau$, so that the density increases as $\bar{\omega}^{3/2}$ in such a compression. The velocity increases as $\sqrt{\tau} \sim \bar{\omega}^{1/2}$, so that the net increase in the collision rate $\Gamma \equiv n\sigma_c v$ scales as $\bar{\omega}^2$. In terms of the trap parameters, this means that the collision rate scales, for large aspect ratio, as

$$\Gamma \sim \frac{B'^{4/3} B''^{1/3}}{{B_0}^{2/3}}. \tag{2.45}$$

The scaling of gradient with distance is I/R^2, while curvature scales as I/R^3, where it is assumed that the coils are positioned to optimize the gradient and curvature respectively. This optimum occurs at a distance on the order of the size R of the coil. Consequently, the collision rate scales as $R_{\text{grad.}}^{-8/3}$ and as $R_{\text{curv.}}^{-1}$, and it is relatively more beneficial, as far as the collision rate is concerned, to place the gradient coils closest to the atoms. This analysis neglects, of course, the possibility that tighter confinement may be required in the axial direction for a particular experiment. The relatively low ($\sim$ 6 Hz) axial frequency of the present trap makes experiments on one-dimensional condensates more feasible, but renders damping of "sloshing" motion in the trap more difficult.

It was decided to place the magnetic trap inside the vacuum system, in view of the fact that for a conventional steel vacuum system the minimum fractional solid angle occupied by blind vacuum hardware, as opposed to viewports, decreases as the size of the vacuum system increases. The magnitude of the attainable magnetic-field curvatures, however, decreases rapidly as the magnetic trap is made larger. Therefore, one might reason that the magnetic trap should be brought close to the atoms and the vacuum system made distant from them, for largest trap curvatures and best optical access. This logic, of course, is not watertight; it hardly applies, for instance, to a glass cell, nor does it take into account the practical difficulties of running large currents and high-pressure cooling water into a vacuum system operating at 10^{-11} torr. Nevertheless, this was the path chosen for the magnetic trap on the present apparatus. As a result, all materials as well as handling and cleaning procedures were of necessity compatible with ultra-high vacuum (UHV). Moreover, in order to minimize eddy currents during switch-off of the magnetic trap for the most rapid extinction of magnetic fields, considerable effort was made to avoid the use of conductive metallic materials in regions of high magnetic field.

The coils were wound from 1/8 in.-square oxygen-free high conductivity copper tubing (1/32 in. wall thickness) insulated with Kapton ®. To aid in the winding process, spools of the appropriate dimensions with removable end caps were machined from aluminum. Axial grooves were placed at regular intervals along the spool circumference, allowing thin Kapton-coated copper wires to be prepositioned under the tightly wound coils. For the rectangular coils, well deburred screw-tightened aluminum clamps were positioned along the edges to allow the springy copper tubing to be bent cleanly to the prescribed shape and held securely in place as winding progressed. The completed coils were than held together by the several tightly-twisted copper wires which had previously been laid underneath. Use of copper binding wire was chosen as an alternative to epoxy because of its vacuum compatibility. The now-monolithic coil was then gently pried and slid off the spool, with its end cap now removed, taking care not to tear the insulation. These coils were carefully wrapped in aluminum foil with only their leads exposed, and 1/4 in. VCR glands suitable for maintaining UHV were brazed onto the leads via round-to-square 1/4 in.-to-1/8 in. copper adapters, custom fabricated by the UC Berkeley Physics machine shop.

The separate coils were affixed to two alumina mounting plates, precisely machined by Andrew Mei of the Lawrence Berkeley National Laboratory machine shop, with each plate holding a curvature coil on one side and an antibias coil on the other side. A ~ 1 in.-square central hole was machined in the plates allowing optical access down the axis of the trap, with four rectangular cut-outs at the corners of this square hole accommodating the four Ioffe bars (i.e., the two sides of each of the two gradient coils). The alumina plates fit snugly over the gradient coils, while the curvature coils were pressed against the outer corners of the gradient coils, so that the entire assembly with the exception of the antibias coils was at this stage self-supporting. Additional thin Kapton-coated wires were then put in place binding each coil to the alumina mounting plates, and four 316 stainless steel 1/4" − 20 threaded rods were bolted into holes in the corners of the plates, joining them into a rigid structure. Finally, four alumina right-angle elbows pierced on each face by 1/4 in.-diameter holes and secured to each of the bottom corners of the alumina plates, allowed the magnetic-trap mount to be bolted to a ring-shaped stainless-steel mounting shelf that had been welded into the 6 in. tube of the main vacuum chamber. The leads from the six coils were brought down to three $4\frac{1}{2}$ in.-flanged ports at a lower level on the vacuum chamber, where the VCR glands mated (via a VCR gasket of steel or copper) with matching VCR glands brazed to the ends of 1/4 in. copper tube vacuum feedthroughs. Each feedthrough flange contained four independent insulated feedthroughs. The total resistance of connecting lugs, feedthroughs, and VCR fittings was measured at around 3 mΩ. In order to allow room for brazing the fittings VCR gaskets of steel had been employed in the earliest iteration of assembling the trap, but it later seemed that copper provided superior reliability, or at least greater immunity to imperfect attachment and tightening. In order to avoid blockage of optical access when the current feedthroughs were pushed into place, the leads were deliberately kept quite short, with the unfortunate side effect that very little working space was available to attach the VCR gaskets and tighten the VCR nuts. In later revision of the design, longer leads were left on the coils, and stainless-steel vacuum bellows were inserted between the vacuum chamber and the feedthrough flanges. With the bellows retracted, the VCR nuts could then be tightened with relative ease, and the bellows could subsequently be allowed to extend over the tightened connections. When the vacuum system was evacuated, the bellows naturally tend to retract, so that care must be taken to secure them at the desired length prior to pump-down.

The magnetic trap fields were measured at a current of 20 A before closing the vacuum system, with all coils in their final positions. The measurement was performed with a Hall probe, which had an approximate sensitivity of 0.5 G and required frequent re-zeroing. The current of 20 A was chosen because at the time the measurement was taken cooling water was not available. The results of a measurement of the axial fields are shown in figure 2.13. The theoretical curve (solid line) is based on the expressions

$$B_z = \frac{\mu_0 I N}{2} \frac{R^2}{\left(R^2 + z^2\right)^{3/2}} \tag{2.46}$$

for the axial field of a circular coil of radius R, N turns, carrying current I, and

$$B_z = \frac{\mu_0 I N}{\pi} \frac{2a^2}{a^2 + z^2} \frac{1}{\left(2a^2 + z^2\right)^{1/2}} \tag{2.47}$$

for the axial field of a square coil of edge $2a$, N turns, and current I. The positions of the coils and their sizes have been taken as free parameters in the theoretical curve. The curve fit yields values within a few millimeters of the design parameters, the largest deviation being in the radius of the antibias coils, which is required to be 6 mm smaller than the design value of 4.5 cm for good agreement. Such variation in the positions is not implausible, given the degree of stress to which the trap was subjected in mounting it within the vacuum chamber. The data suggest what was later confirmed by observation of trapped atoms, that the central minimum dips past zero, although the signal-to-noise ratio of figure 2.13 is insufficient to prove this. This slight zero-crossing is in fact desirable, since it implies that any nonzero bias field may be obtained by injecting a small current into the antibias coils.

A more rigorous characterization of the magnetic trap has been obtained, as suggested above, by using the trapped atoms themselves. Measurements of the minimum bias field, obtained by sweeping a RF field through resonance with the trapped atoms and observing trap loss, have yielded the approximate formula for the bias field

$$B_0 = \left(\frac{555\,\mathrm{kHz}}{A}\right) I_{\mathrm{Curv.}} - \left(\frac{1.7\,\mathrm{kHz}}{A}\right) I_{\mathrm{Main}} - \left(\frac{14\,\mathrm{kHz}}{A}\right) I_{\mathrm{Grad}}, \tag{2.48}$$

where $I_{\mathrm{Curv.}}$,I_{Main}, and $I_{\mathrm{Grad.}}$ are the curvature, main, and gradient currents respectively. As used here and elsewhere, the main current is that portion of the current which flows through both curvature and antibias coils, while the curvature current is the portion of the current in the curvature coils which is *not* in common with the antibias coils.

Oscillation frequencies of magnetically trapped atoms have also allowed the precise characterization of the magnetic trap. For measurements of the transverse frequency, pulsing on a transverse bias field to displace the transverse minimum of the trap is effective in initiating oscillations, while for measuring axial oscillations a brief kick from a magnetic gradient (e.g., the MOT gradient) is efficacious. Measurements of this sort, where the atom cloud is given a brief impulse and allowed to evolve for a variable period of time, have yielded frequencies of $143\,\mathrm{Hz} \times 5.7\,\mathrm{Hz}$ at a bias field of $2.11\,\mathrm{G}$, yielding an inferred transverse gradients of $225\,\mathrm{G/cm}$ and an inferred axial curvature of $38\,\mathrm{G/cm}^2$ at currents $I_{\mathrm{Grad.}} = 150\,\mathrm{A}$, $I_{\mathrm{Main}} = 400\,\mathrm{A}$, $I_{\mathrm{Curv}} \approx 0$. More precise frequency measurements have also been taken in a decompressed magnetic trap, i.e., at lower gradient current. The frequencies for this trap were determined to be $5.00 \pm 0.01\,\mathrm{Hz}$ and $48.6 \pm 0.35\,\mathrm{Hz}$ at currents of $I_{\mathrm{Grad.}} = 100\,\mathrm{A}$ and $I_{\mathrm{Main}} = 300\,\mathrm{A}$, and a bias field of $7\,\mathrm{G}$, implying an axial curvature of $29\,\mathrm{G/cm}^2$, a transverse curvature of $2920\,\mathrm{G/cm}^2$, and a transverse gradient of $144\,\mathrm{G/cm}$. The values of the axial curvature and transverse gradient obtained from these measurements are well fit by straight lines passing through the origin, whose slopes give the formula for the axial curvature and frequency and the radial gradient.

$$
\begin{aligned}
B_z'' &= \big(9.54 \pm .13\big)\,\mathrm{G/cm}^2 \left(\frac{I_{\mathrm{Main}}}{100\mathrm{A}}\right) \\
\nu_z &= 0.9\,\mathrm{Hz}\sqrt{\frac{B_z''}{\mathrm{G/cm}^2}} = 2.77\,\mathrm{Hz}\sqrt{\frac{I_{\mathrm{Main}}}{100\mathrm{A}}} \\
B_\rho' &= \big(145 \pm .8\big)\,\mathrm{G/cm} \left(\frac{I_{\mathrm{Grad.}}}{100\mathrm{A}}\right). \qquad (2.49)
\end{aligned}
$$

The radial frequency is determined from $\nu_\rho = 0.9\,\mathrm{Hz}\sqrt{B_\rho''/(\mathrm{G/\,cm}^2)}$, where B_ρ'' is given by (2.43).

2.6.3 Uncompressed and Compressed Operation

The cold, compressed cloud of atoms delivered by the MOT is in gravitational free-fall once the atoms are optically pumped to the $F = 1$ level and the lasers are extinguished. Thus the magnetic trap must be turned on quickly, before the atoms have had time to accelerate appreciably. If the trap curvatures or frequencies are chosen so that the atomic phase-space distribution matches the equilibrium distribution in

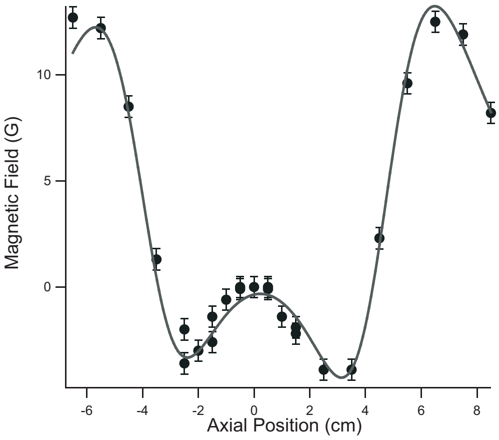

Figure 2.13. Axial magnetic-trap field measurement. The measurement was obtained by a Hall-probe while running a current of 20 A in series through the curvature and antibias coils. The magnitude (absolute value) displays a quadratic minimum of close to 0 G in the neighborhood of 0 cm.

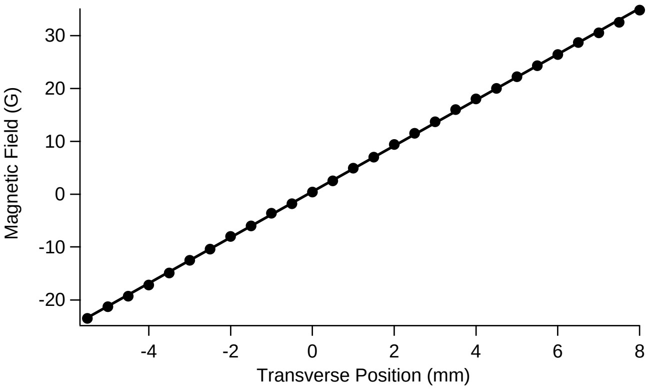

Figure 2.14. Transverse trap gradient measurement (solid circles) vs. transverse displacement, performed with a Hall probe in the unevacuated vacuum chamber at a current of 30 A. The solid line is a fit to the data, yielding a value of 144.3±0.3 G/cm.

the trap as nearly as possible, then heating of the atomic cloud in the transfer will be minimized. For a cloud at temperature T whose real-space distribution is Gaussian with $1/e$ radius r_0, this matching condition is

$$B'' = \frac{2k_B T}{\mu r_0^2}, \tag{2.50}$$

where $\mu = \frac{1}{2}\mu_B$. For $r_0 = 3\,\mathrm{mm}$ and $T = 75\,\mu\mathrm{K}$, this implies a relatively small trap curvature of $50\,\mathrm{G/cm}^2$ (equivalent to a trap frequency of 6.5 Hz) in all directions. On the other hand, for efficient evaporation one would like to have much higher curvatures to increase the density and rethermalization rate. Fortunately, the Ioffe-Pritchard trap provides a very convenient way to accommodate both these requirements. Referring to the expression for the transverse curvature in expression (2.43), it is clear that this curvature can be dramatically modified by changing the value of the bias field, as illustrated in figure 2.16. The near-isotropic trap configuration desired for loading the trap is provided by turning on a large bias field ($\sim 320\,\mathrm{G}$), while larger curvatures are achieved by ramping the bias field to fields of a few Gauss, at which curvatures of $15000\,\mathrm{G/cm}^2$ are readily obtained. In normal operation, the uncompressed magnetic trap is turned on with frequencies of approximately $(7.6 \times 10 \times 10)$ Hz, while the fully compressed trap has frequencies $5.7 \times 140 \times 140$ Hz at a bias field of approximately 2 G, suitable for efficient evaporation, as described in section 2.6.6.

2.6.4 Switching Currents Off

It is desirable to be able to take the trapping fields rapidly to zero for the purpose of free-expansion ("time-of-flight") imaging of atomic clouds. In order for the shut-off of the trap to be rapid with respect to the largest commonly employed magnetic trapping frequency of $\nu_\perp \approx 150\,\mathrm{Hz}$, one wishes the extinction-time of the trap current to be much less than a quarter period, or small compared to 1 ms. For a coil of $R \sim 3\,\mathrm{cm}$ radius and $N = 8$ turns with a wire radius of $r \sim 0.15\,\mathrm{cm}$, one estimates an inductance

$$L \approx \mu_0 R N^2 \log(R/r) \approx 7\,\mu\mathrm{H}. \tag{2.51}$$

Estimates for lead inductance give values of a few μH, so that a more conservative estimate would be $\sim 10\,\mu\mathrm{H}$.

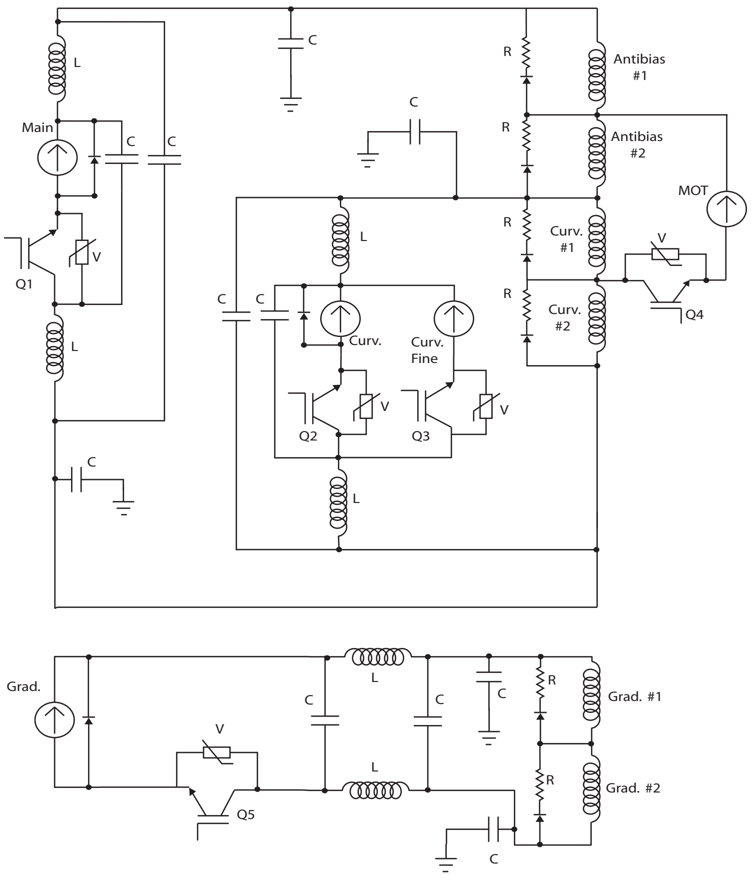
L
Main
C
C
Q1
V
L
C
C
C
R
R
Antibias
#1
Antibias
#2
MOT
R
R
Curv.
#1
Curv.
#2
V
Q4
L
C
C
Curv.
Curv.
Fine
V
V
Q2
Q3
L
Grad.
L
C
C
C
V
Q5
L
R
R
Grad. #1
Grad. #2
C

Figure 2.15. Magnetic trap circuit. A single ("Main") current-regulated supply runs current in series through curvature and antibias coils. Additional current may be added to the curvature coils only by the curvature and curvature-fine supplies, allowing coarse and precise control of the bias field, respectively. The MOT supply exploits the approximate anti-Helmholtz configuration of one curvature/ anti-bias pair to produce a magnetic gradient. The magnetic-trap gradient supply is on a separate circuit. Currents are switched by IGBTs (Q1-5), while shut-off voltage is limited by varistors (V). Each $L \approx 20\,\mu$H and $C \approx 2\,\mu$F belongs to an RF-noise filter. Secondary burn-off resistors R have values of around $5\,\Omega$.

In order to shut off the current as rapidly as possible, one wishes to sustain a large negative voltage across the trap coils. If one, for instance, wished to extinguish a 300 A current in 10 μs in a coil with $L \sim 10\,\mu$H, a reverse voltage of 300 V would be required. Such a large negative voltage is generated automatically by an inductor when its current is interrupted and is generally limited only by component breakdown. The integrated-gate bipolar transistors (IGBTs) used as switches, however, can only sustain from 500 V to 1000 V; therefore some means of limiting the peak inductive voltage must be implemented. The common method of limiting inductive spikes is to place a reversed diode in parallel with the inductor, which provides a path for the current when a switch is opened. A typical diode voltage drop of 0.6 V, however, is far less than the IGBT could sustain, and would place an unnecessary limitation on the switch-off speed of the circuit. A larger voltage drop may be obtained by placing a so-called "burn-off" resistor in series with the diode, with its resistance chosen to achieve the desired ramp-down voltage at the expected shut-off current. This solution has the advantage that it offers large initial ramp-down voltages, but as the current diminishes, the voltage available to ramp it down further diminishes in the familiar exponential L/R circuit decay. A solution which promises still faster ramp-down of the currents is to use a varistor in place of the diode/resistor combination. A varistor offers a nonlinear I-V curve, which allows it, much like a diode, to keep the voltage nearly constant above a certain threshold current. In contrast to the diode, however, this voltage may be hundreds of volts. Thus the varistor is capable of maintaining a large voltage for most of the ramp-down, speeding up the shut-off considerably. When the current drops below the varistor threshold, however, the varistor is no longer relevant, and a secondary burn-off path must be provided to deal with this case, such as the resistor-diode combination mentioned above. This two-stage ramp-down scheme is illustrated in figure 2.15. The resistances R are chosen to be a few Ohms, which is large enough to protect them from excessive current flow during the initial portion of the shut-off, but small enough that they do not generate a dangerously large voltage when the varistor ceases to conduct.

2.6.5 RF Noise Reduction

In addition to the need for rapid shut-off of the currents, a significant factor which has shaped the circuit represented in figure 2.15 is the need to filter and eliminate sources of RF noise. During initial attempts at evaporation, it was discovered that the lifetime of the magnetically trapped cloud depended strongly on the bias field at the minimum of the trap and on the temperature, going from 60 s at bias fields of several tens of Gauss to as little as 17 s at bias fields of $2 - 3$ G. Moreover "hot" clouds had substantially longer lifetimes than cold at the lower bias fields. There exist only a few causes for trap loss in a deep magnetic trap; collisions with room-temperature background gas molecules, near-resonant light, and radio-frequency noise constitute essentially an exhaustive list. Neither of the former are expected to have the same dependence on bias field and temperature, so that we soon concluded that RF noise in the range $1 - 10$ MHz was responsible. Attempts to probe directly the RF noise spectrum in the lab and more specifically on the magnetic trap coils by means of pickup coils or isolated oscilloscope probes connected to a spectrum analyzer yielded a number of initial culprits, mostly associated with DC-DC converters, which typically have switching frequencies in the 100 kHz range, but if poorly designed can have tens or hundreds of harmonics. Some of the culprits were relatively easy to replace, like the floating power supplies for the analog isolator circuits that fed control signals to the large magnetic-trap power supplies, or turn off, like the ion gauge we used for reading vacuum pressure. Others, like the magnetic-trap power supplies themselves, were integral to the experiment, while still others could not be identified at all, but appeared merely to be ambient or power-line noise. To mitigate the effects of these noise sources, we attempted to filter the outputs of the power supplies before they entered the vacuum system. This was accomplished by means of so-called π filters, i.e., with series inductors and parallel capacitors, as shown in figure 2.15. The inductors in question were required to carry currents of as much as 400 A, so that thick copper welding cable (4/0 gauge) was used for their construction. The resulting inductors, consisting of about 10 turns each around 15 cm diameter, had inductances of about 20 μH. Attempts to increase the inductance by the insertion of magnetically permeable cores were unsuccessful in reducing measured noise at the output in the few-megahertz frequency range of interest. At lower frequencies, however, these cores are expected to improve filter performance. Capacitances of a few μF were placed between the positive and negative current leads on either side of these inductors, which themselves were in series with both positive and negative leads, as indicated in the schematic. These filters appeared to reduce the noise associated with the trap power supplies, but a more dramatic effect was observed when the trap leads on the far side of the filters were capacitively coupled to ground (in this case, the optical table). The presence of these capacitors appeared well correlated with the atom lifetime, and their implementation shortly preceded the achievement of BEC in our laboratory. Further measures implemented later, based upon the results of RF noise measurements taken in a mock chamber on the same optical table, included twisting and shielding the 4/0 welding-cable current leads from the power supplies to the trap. The shield, made of overlapping sheets of aluminum foil taped in place, was grounded to the optical table.

2.6.6 RF Evaporation

Forced RF evaporation was originally proposed for magnetically trapped hydrogen by Hess [60] and has become a very widely used tool in ultracold atom experiments. In essence, a RF field of well-defined frequency is resonant with the Zeeman $m = -1 \rightarrow m = 0$ transition (for instance), at a particular magnetic field. As a result, it selectively flips the spins of atoms at those locations in space where the field has this magnitude, expelling them from the trap. If this field is chosen so that the expelled atoms have a higher-than-average energy, then the net effect is to lower the energy of the sample, which results, after rethermalization, in cooling it. Provided sufficient energy is removed per atom, the evaporation process will result in a nonzero atom number when the critical temperature for Bose condensation is reached. In a confining potential, colder atoms are also denser, resulting in more rapid rethermalization, so that a virtuous cycle may be initiated in which cooling leads to more rapid interatomic collisions and thereby to more efficient cooling. This desirable state of affairs is known as runaway evaporation.

In our experiment, the cloud is allowed to equilibrate after transfer from the MOT to the uncompressed trap described in section 2.6.3. This equilibration lasts for 2 sec, following which the cloud is partially compressed to a bias field of around 50 G. In the partially compressed trap, preliminary radiofrequency (RF) evaporation at fixed RF frequency for several seconds is used to prevent the temperature from rising excessively during compression. The frequency source for evaporation is an IFR model 2023A, which is computer-controlled via GPIB. The signal passes through an RF switch and an amplifier and is then capacitively coupled into one of the magnetic-trap gradient coils, which has the advantage of being close to the atoms and of being correctly orientated to produce a transverse driving field. A disadvantage to this scheme is that the gradient coil and its leads can broadcast the RF signal throughout the laboratory, destabilizing sensitive equipment. As no RF feedthrough was incorporated in the vacuum chamber, however, this has been the unavoidable if regrettable side-effect of evaporation throughout the history of the experiment.

After the preliminary evaporation, the trap is compressed to a bias field of approximately 2 G, corresponding to trap frequencies of $\sim (140 \times 140 \times 5)$ Hz, as described in section 2.6.3. During the compression, the RF frequency is ramped down so as roughly to track the dropping bias field. This ramp, from approximately 50 MHz to 10 MHz crosses a number of RF resonances in the combined system consisting of function generator, amplifier, coupling network, and broadcast coils. Certain fine tunings of this evaporation trajectory, other than those required for efficient evaporation, have proven useful in avoiding the destabilizing effects of RF on lasers and other laboratory equipment. Slowly ramping up the initial RF power, for instance, proved helpful in this regard, as have various short jumps in frequency and drops in power, e.g., around 47 MHz and 20 MHz. The remainder of the evaporation sequence, lasting approximately 20 s, occurs in this fully compressed trap. Condensates of approximately 2 million atoms can be produced, or alternately for the spinor experiments we terminate the evaporation at a temperature usually of $2\,\mu$K with remaining atom numbers of ≈ 30 million. Such a cold cloud is suitable for transfer to an optical trap, as described in section 2.7.

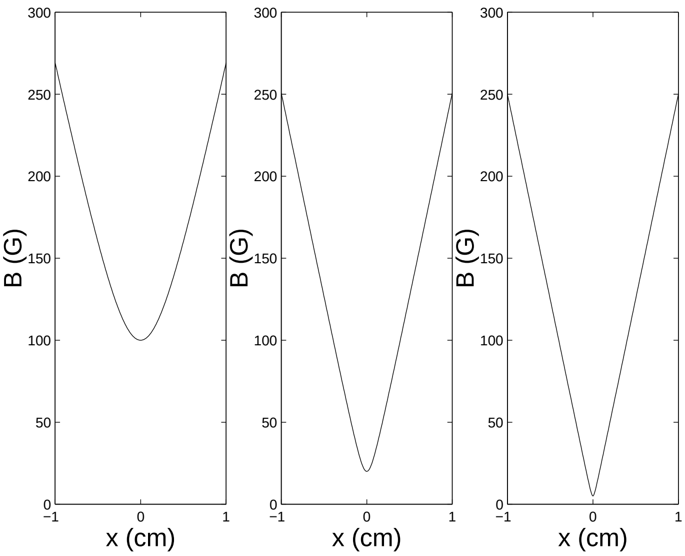

Figure 2.16. Illustration of compression in Ioffe-Pritchard trap. A radial gradient B' and an axial bias field B_0 add in quadrature, producing a hyperbolic field profile, which is close to linear for distances larger compared to B_0/B' and approximately parabolic for distances less than B_0/B'. At large B_0, the parabolic region is large, and the curvature is weak, while for small B_0 the curvature becomes very large. For hot clouds in a low-bias-field trap, the potential is a tight linear trap rather than a relatively weak quadratic trap, which aids in efficient evaporation.

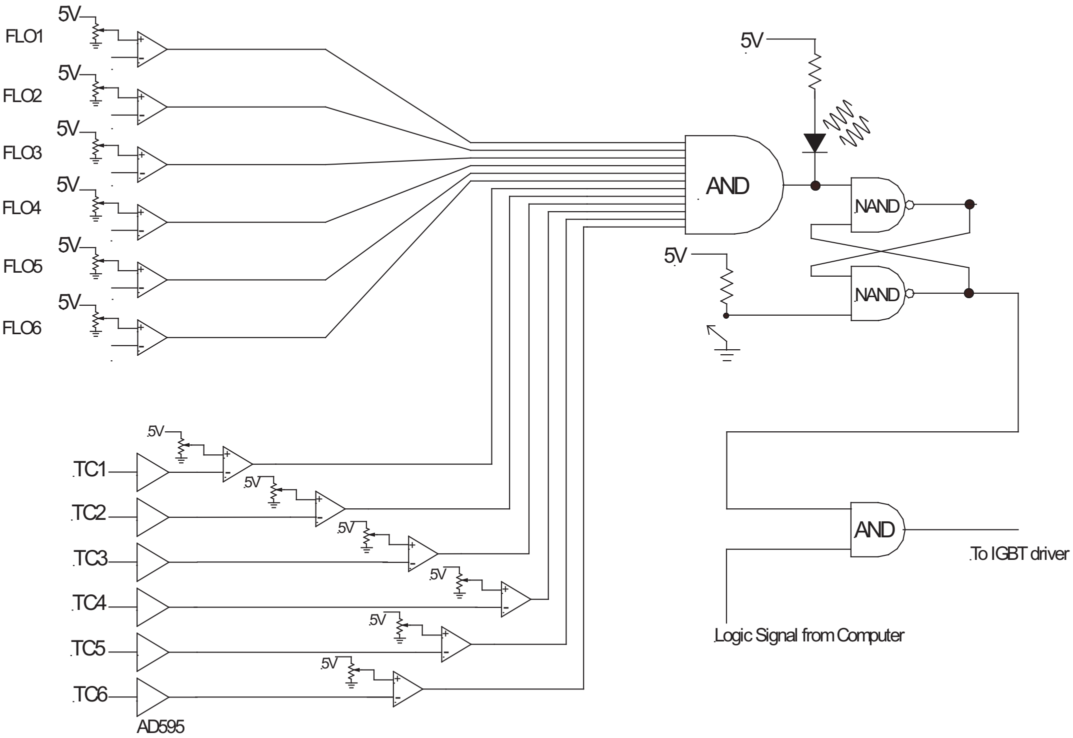
5V
FLO1
FLO2
FLO3
FLO4
FLO5
FLO6
AND
NAND
NAND
TC1
TC2
TC3
TC4
TC5
TC6
AD595
AND
To IGBT driver
Logic Signal from Computer

Figure 2.17. Interlock schematic. Flowmeters and thermocouples monitor water flow and temperature. Analog signals from each channel are compared to user-defined setpoints by hysteretic Schmitt trigger circuits with logic-level outputs. These outputs are ANDed and the resulting output controls a flip-flop circuit with manual reset, so that operation does not resume after a fault condition without human intervention. The final output gates the transistors that switch the high trapping currents as well as contactors that pass AC wall power to the trap power supplies. The contactors also require manual resetting after a fault.

2.6.7 Water-Cooling and Interlock

The magnetic trap at full power dissipates in excess of 4 kW power. For a single coil of conducting cross-section $0.068\,\mathrm{cm}^2$ and length 150 cm, using the room-temperature thermal conductivity of copper, 4.01 W/cm/°C, one calculates the unreasonable equilibrium temperature, limited by conduction to the ends, of 7×10^4 °C. Black-body radiation, assuming a surface area of $\sim 50\,\mathrm{cm}^2$ would limit the temperature to a more reasonable value of around 1500°C, but this is still in the neighborhood of and possibly exceeding copper's melting point of 1360°C. Even if melting of the coils could be forestalled, however, such extreme temperatures would result in highly undesirable thermal expansion of and stress on the trap, as well as destroying any insulation on the coils and probably making good vacuum impossible. It is clearly the case, then, that the magnetic-trap coils must be actively cooled, and indeed they were designed to be water-cooled by running high-pressure water down the center of the conducting copper tube.

The system for cooling the coils employs two pumps in series; the first, a Neslab recirculating chiller, provides a head of 60 p.s.i. and the second, a Bronze high-pressure turbine pump from McMaster-Carr, augments this value to $150 - 180$ p.s.i. The high-pressure water exiting these pumps is split into six parallel channels which feed the six coils of the magnetic trap, via 1/4 in. tubular copper vacuum feedthroughs. The water exiting these coils is forced to traverse six independent flowmeters (Lake Monitor style R) which transmit analog readings of the flow in each coil to the interlock logic box, described below. The flowmeters were chosen to have only metal parts in contact with the water, so as to preserve their integrity in the presence of hot water; in practice, this was unnecessary caution, since water temperatures exceeding 30°C are rarely attained. On each of the six pathways, thermocouples with a 1/16 in. stainless-steel sheath are immersed in the fluid, the watertight seal being provided by the compressed ferrules of a 1/16 in. Swagelok TM reducing union. The leads of these thermocouples are also fed into the interlock logic box via thermocouple extension cables, and only then are amplified to signal-level ($\sim 1\,\mathrm{V}$) voltages. A better design would place amplifiers much closer to the thermocouples to eliminate RF pickup on the long leads. After the thermocouples and flowmeters, the six pathways are recombined and return to the recirculator, where the water is passed through a heat exchanger (cooled by building water) before completing its closed loop by arriving at the reservoir of the recirculator.

Given the catastrophic consequences mentioned above that will occur in the absence of cooling water, it is important to ensure that high currents are never permitted to traverse the coils when adequate cooling is absent. This is ensured by monitoring flow rate and temperature of the cooling water as described above. In the dedicated interlock logic box, each flowmeter analog signal is compared to a corresponding threshold value set by a potentiometer on the front panel of the box. The comparison is made by Schmitt trigger with hysteresis, to provide rapid and chatter-free switching in the event of a fault condition. Thermocouple signals are amplified in the interlock box by six AD595 thermocouple amplifiers, and the resulting signals are similarly compared via Schmitt triggers to a single threshold, likewise set by a potentiometer. The outputs of these twelve Schmitt-trigger comparators are tied via pull-up resistors to 5 V, so that they possess TTL-compatible output levels, and these levels are ANDed together. The output of the AND feeds into a flip-flop consisting of two NAND gates. The other input is normally high, but can be manually pulled low by means of a reset button on the front panel. When the manual reset button is not being pressed and no fault condition (e.g., insufficient water flow or excessive temperature) is present, the flip-flop can exist in two states, with its output either high or low. Resetting it will force it into the output-high state, where it will remain until a fault condition occurs, after which it will remain in the output-low state. This feature is deliberate, designed to force a conscious decision from the user whether to continue operation when a fault has occurred, even if the fault condition no longer appears to be present. The output of the flip-flop gates the TTL signals which control the switches (IGBTs) for the high-current supplies, so that after a fault condition has occurred, no current should flow to the trap. As an added safety measure, the A.C. wall power going to the high-current magnetic-trap supplies passes through contactors (i.e., relays), which are similarly controlled by the interlock output, so that the power supplies are deprived of A.C. power when a fault condition has occurred. Like the interlock circuit, the contactor circuit has a “memory”, meaning that once it is triggered to shut off, it will not turn back on even when the logic signal from the interlock box ceases to indicate a fault, requiring a separate reset button to be pressed to resume normal operation.

2.7 Optical Trap

The optical-dipole trap (ODT) for cold atoms was proposed by Ashkin [19] and was first demonstrated by Chu et al. [61]. The present use of a shallow ODT for ultracold atoms previously evaporated in a magnetic trap was first implemented by [62].

The ODT operates via the A.C. Stark shift of the atomic ground state, which is to say that the oscillating optical field induces a dipole oscillating at the same frequency. Some intuition for this process may be gained from looking at the energy of a classical Lorentz atom, whose electron at sufficient detuning (i.e., neglecting radiative damping) obeys the equation of motion

$$m\ddot{x} = -m\omega_0^2 x - eE\cos\omega t, \tag{2.52}$$

with the solution

$$x = \frac{eE/m}{\omega^2 - \omega_0^2}\cos\omega t. \tag{2.53}$$

As a result, the average energy of the dipole interacting with the electric field is

$$\langle eEx\rangle \approx \frac{(eE)^2}{4m\omega_0}\frac{1}{\delta}, \tag{2.54}$$

where in the last step the detuning $\delta = \omega - \omega_0$ has been introduced and terms of order δ^{-2} neglected. Since the dipole is in phase with the electric field for $\delta < 0$, its energy is lower, while for $\delta > 0$ the dipole is π out of phase and therefore has higher energy.

Quantum mechanically, the A.C. Stark shift of an atomic ground state in the presence of a laser beam is

$$\Delta E_g = \sum_e \frac{\hbar|\Omega_{eg}|^2}{4\delta_e} = \frac{\hbar\gamma^2}{8}\frac{I}{I_{\text{sat}}}\sum_e \frac{|c_{eg}|^2}{\delta_e}, \tag{2.55}$$

where Ω_{eg} is the Rabi frequency, δ_e the detuning from resonance, I the laser intensity, I_{sat} the saturation intensity on the cycling transition, and c_{eg} the Clebsch-Gordan coefficient for the $|g\rangle$ to $|e\rangle$ transition.

For ^{87}Rb at detunings large compared to the 6.8 GHz ground-state hyperfine splitting, it is permissible to consider only fine structure, in which the D2 transition at 780 nm is a $J = \frac{1}{2} \longrightarrow J' = \frac{3}{2}$ transition and the D1 at 794.7 nm is a $J = \frac{1}{2} \longrightarrow J' = \frac{1}{2}$ transition. The Clebsch-Gordan coefficients for π-polarized light are $|c_{\frac{1}{2}\frac{1}{2}\rightarrow\frac{3}{2}\frac{1}{2}}|^2 = \frac{2}{3}$ and $|c_{\frac{1}{2}\frac{1}{2}\rightarrow\frac{1}{2}\frac{1}{2}}|^2 = \frac{1}{3}$, so that the ground state energy shift is

$$\Delta E_g = \frac{\hbar\gamma^2}{8}\frac{I}{I_{\text{sat}}}\left(\frac{2}{3\delta_{D2}} + \frac{1}{3\delta_{D1}}\right). \tag{2.56}$$

This shift is the same for $m_J = \pm\frac{1}{2}$, meaning that all spin-states in the ground-state manifold experience the same optical potential. If the detuning is chosen negative, then this potential is attractive, drawing atoms into regions of high intensity. For the wavelength of 825 nm employed in these experiments, and using the values $\gamma = 2\pi \times 6\,\text{MHz}$, $I_{\text{sat}} = 1.4\,\text{mW/cm}^2$, this translates to an energy shift of

$$\Delta E_g = -h \times 0.175\frac{\text{Hz}}{\text{mW/cm}^2}I = -k_B \times 8.3 \times 10^{-6}\left(\frac{\mu\text{K}}{\text{mW/cm}^2}\right)I. \tag{2.57}$$

In order to confine atomic samples with temperatures on the order of a microkelvin, it is clearly necessary to have a rather large spatial maximum in the optical intensity, i.e., a focus of the trapping laser beam. The description of a focused laser beam is elementary, but will be given briefly for reference in the general case of a possibly astigmatic focus. Such a focus is conveniently described in the paraxial slowly-varying-envelope approximation, where $\mathbf{E}(\mathbf{x},\mathbf{y},\mathbf{z}) = \mathcal{E}(x,y,z)\hat{\mathbf{x}}e^{ikz-i\omega t}$ and derivatives of $\mathcal{E}(x,y,z)$ with respect to z are neglected beyond first order. The particular linear polarization $\hat{\mathbf{x}}$ has been chosen for convenience. In this approximation $\mathcal{E}$ satisfies the paraxial wave equation

$$\left(\nabla_\perp^2 + 2ik\frac{\partial}{\partial z}\right)\mathcal{E} = 0. \tag{2.58}$$

This equation is separable, so that its solutions can be written as the product of two functions of the form

$$\mathcal{E}(y,z) = \frac{E_0 y_0}{\sqrt{y_0^2 + 2iz/k}}\exp\left(-\frac{y^2}{y_0^2 + 2iz/k}\right). \tag{2.59}$$

This solution smoothly interpolates between the form at the focus e^{-y^2/y_0^2} and the spherical wave at large distances $e^{ikR} \approx e^{ikz+\frac{iky^2}{2z}}$, of which the slowly-varying portion is $e^{\frac{iky^2}{2z}}$. If the focus in the x direction occurs a distance Δz away (i.e., the beam is astigmatic), then the total intensity distribution is

$$I = \frac{2P}{\pi x_0 y_0} \frac{\sqrt{1+\left(\frac{2\Delta z}{x_0^2 k}\right)^2}}{\left(1+\left(\frac{2(z-\Delta z)}{k x_0^2}\right)^2\right)\left(1+\left(\frac{2z}{k y_0^2}\right)^2\right)} \exp\left(\frac{-2x^2}{x_0^2+\frac{4(z-\Delta z)^2}{k^2 x_0^2}} + \frac{-2y^2}{y_0^2+\frac{4z^2}{k^2 y_0^2}}\right), \tag{2.60}$$

where P is the total power.

The general expression for the curvatures in the three directions around $z = 0$ (assuming a tighter focus in y than in x) is not difficult to derive. In the x and y directions, the curvatures are

$$\begin{aligned} \left.\frac{\partial^2 I}{\partial x^2}\right|_{z=0} &= -\frac{8P}{\pi x_0 y_0}\left(\frac{1}{x_0^2+\frac{4\Delta z^2}{k^2 x_0^2}}\right)\frac{1}{\sqrt{1+\left(\frac{2\Delta z}{k y_0^2}\right)^2}} \\ \left.\frac{\partial^2 I}{\partial y^2}\right|_{z=0} &= -\frac{8P}{\pi x_0 y_0}\left(\frac{1}{y_0^2}\right)\frac{1}{\sqrt{1+\left(\frac{2\Delta z}{k y_0^2}\right)^2}}. \end{aligned} \tag{2.61}$$

The full expression in the z direction is somewhat lengthy, but provided $y_0 \ll x_0$ is approximately given by

$$\left.\frac{\partial^2 I}{\partial z^2}\right|_{z=0} = -\frac{16P}{\pi x_0 y_0}\left(\frac{1}{k^2 y_0^4}\right)\frac{1}{\sqrt{1+\left(\frac{2\Delta z}{k y_0^2}\right)^2}}. \tag{2.62}$$

The trap frequencies are then determined by

$$\omega_i^2 = \frac{\hbar\gamma^2}{8 m I_{\mathrm{sat}} \delta_{\mathrm{eff}}} \left.\frac{\partial^2 I}{\partial x_i^2}\right|_{z=0}, \tag{2.63}$$

where

$$\delta_{\mathrm{eff}} = \frac{3\delta_{1/2}\delta_{3/2}}{2\delta_{1/2}+\delta_{3/2}}. \tag{2.64}$$

2.7.1 Optical Layout

The parameters of the optical trap, including beam waists, wavelength, and power, were chosen so that the trap in the vertical direction provided adequate force to support against gravity, while the transverse horizontal direction was loose enough to avoid excessive density, which would lead to high loss rates from three-body recombination. Additionally, it was desired to make the axial frequency reasonably tight ($\gtrsim 5\,\mathrm{Hz}$), to have a trap depth of at least several μK, and to avoid significant scattering from spontaneous emission. Availability of diode lasers of adequate power at the wavelength was an important practical consideration. These criteria were satisfied for the chosen wavelength of 825 nm and powers of a few mW, with beam waists of $12\,\mu$m in the tight vertical direction and $25-80\,\mu$m in the horizontal direction. Optical trap light from a free-running diode laser was fiber-coupled, and all the optics beyond the output of the fiber were placed on a single $\sim$ 1 ft. × 2 ft. optical breadboard, clamped above and below with the rubbery vibration-damping material Sorbothane. This appeared to decrease vibrational heating in the trap quite dramatically. The light exiting the fiber was shaped by means of cylindrical optics in order to obtain the desired aspect ratio. In the most recent design, a system of lenses (a composite "zoom" lens) was put in place that allowed quick changes of aspect ratio in order to facilitate studies of the rôle of dimensionality in spinor dynamics. This system of lenses contained a cylindrical telescope to change the vertical size of the beams. This telescope consisted of a 75 mm lens on a rotatable mount and a 300 mm lens mounted on translation stage. In the horizontal direction, a separate telescope of three cylindrical lenses was constructed, consisting of a rotatable 100 mm lens, a translatable 25 mm lens in the middle, and finally a 50 mm translatable lens. By moving the middle lens closer to the first lens of the telescope, the magnification of the telescope could be made larger, while moving it toward the final lens of the telescope would reduce the telescope. For each position of the middle lens, a corresponding position existed for the final 50 mm lens which would collimate the beam. A final spherical lens immediately before the vacuum viewport focused the beam at the location of the ODT. By imaging the ODT beam onto a camera focused on the location of the trap, it was possible to fine tune the size, focus, and astigmatism of the beam in real time.

2.7.2 Loading the Optical Trap

The transfer procedure from the magnetic to the optical trap is approximately adiabatic, contrary to the experience reported in [62] that such a transfer is impractical because of irregularities in the ramp-down of the magnetic-trap power supplies.

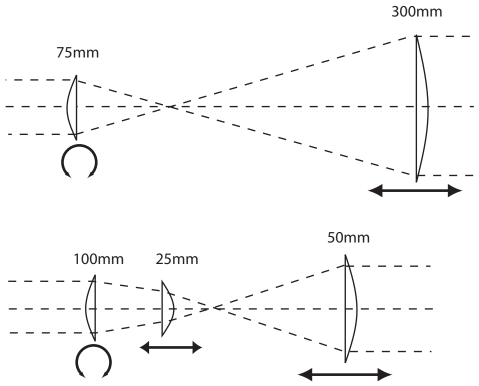

Figure 2.18. Beam-shaping the ODT. A fixed cylindrical telescope in the vertical dimension (top) creates a large beam, suitable for the tight ($\sim 12\,\mu$m) focus in the direction of gravity. A variable telescope in the horizontal direction allows the aspect ratio of the trap to be varied at will. Conceptually, the 25 mm lens moves on a translation stage (linear arrow) to change the beam size, while the 50 mm follows to keep the beam non-astigmatic. Transverse aspect ratios from 3 : 1 to 10 : 1 were readily created. Cylindrical lenses on mounts (circular arrows) allow compensation for cross-dimensional aberrations.

Standard conditions prior to transfer are around 3×10^7 atoms at $1.5 - 2\,\mu$K in a 150Hz × 150Hz × 4Hz magnetic trap. The ODT power is ramped to maximum in a few tens of milliseconds, as the magnetic trap beings to ramp down over 150 ms. By making fine (one-time) adjustments in the ramping rates of the main and gradient supplies (see figure 2.15), it has been possible to avoid displacing the atomic cloud excessively with transient gradients. After transfer, the optical trap power is exponentially ramped down with a $1/e$ time constant of ~ 200 ms and held at a final value of ~ 2.5 mW for the tightest traps employed, or closer to 5 mW for higher-aspect-ratio traps. This lowering of the trap depth causes rapid cooling at the high collision rates (~ 1 kHz) of the optical trap, leading to efficient condensate formation. At this trap depth, the balance between evaporation and heating allows condensates to be held for approximately a second in a tight trap, or several seconds in a large-volume trap.

Initial alignment of the optical trap is performed by imaging the magnetically trapped condensate along the trap axis; the ODT beam is then aligned to and focused at the observed condensate position. This is generally sufficient to capture some countable number of atoms. Further optimization is performed by making small adjustments to the position of the optical trap while monitoring the number held in the trap after 100 ms, when any untrapped atoms have fallen away.

2.8 RF Coils and Spin Manipulation

For the experiments on spinor condensates, various means of manipulating the atomic spins are necessary. In general, we exploit the fact that the magnetic trap purifies the ensemble to $m_F = -1$, since only that state is trapped, so that if a sufficient guide field is present during the transfer to the ODT, the initial state of the atoms in the ODT is also $m_F = -1$. At low ($\lesssim 1\,\mathrm{G}$) bias field, the application of resonant RF pulses merely rotates the spin state. This is useful for the experiments of chapter 3 on induced Larmor precession, which require such a rotation in the form of a $\pi/2$ pulse. In order to prepare other states, such as $m_F = 0$, that are not related to $m_F = -1$ by simple rotation, it is useful to exploit the quadratic Zeeman shift. Including the quadratic Zeeman shift, the dressed-state Hamiltonian for the three near-degenerate states

$$\begin{aligned} &|m_F = -1; (N-1)\omega\rangle \\ &|m_F = 0; N\omega\rangle \\ &|m_F = +1; (N+1)\omega\rangle \end{aligned} \tag{2.65}$$

is

$$H = \hbar \begin{pmatrix} -\delta & \Omega/\sqrt{2} & 0 \\ \Omega/\sqrt{2} & -q & \Omega/\sqrt{2} \\ 0 & \Omega/\sqrt{2} & \delta \end{pmatrix}, \tag{2.66}$$

where $\delta = \omega - \frac{1}{2}\mu_B|B|/\hbar$ and $q \approx 70\,\mathrm{Hz/G^2}$ is (minus) the quadratic Zeeman shift.

As shown in figure 2.19, the energy spectrum of this Hamiltonian exhibits the usual avoided crossings, at $\delta \approx \pm q$ and at $\delta = 0$, of magnitude $\hbar\Omega\sqrt{2}$ and $\hbar\Omega^2/q$ respectively. For $q \gg \Omega$, this means that nearly all the atoms can be placed in $m_F = 0$. Indeed, for a uniform sweep rate $\dot{\delta}$, the fraction in $m_F = 0$ after a sweep across resonance from negative to positive δ is given by the Landau-Zener formula,

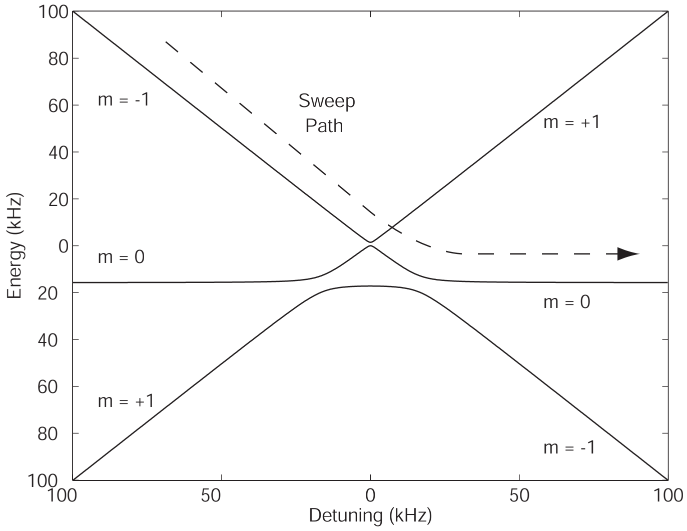

Figure 2.19. Dressed-state level diagram for $m_F = 0$ state preparation, with Rabi frequency $\Omega = 2\pi \times 5\,\mathrm{kHz}$. A sweep from low RF frequency to high RF frequency, with atoms initially in the $m_F = -1$ state, can nonadiabatically follow the $m_F = -1$ state at the small avoided level crossing with $m_F = +1$, but adiabatically follow the level curve into $m_F = 0$ at the second, larger, avoided level crossing. A sweep in the opposite direction at the same sweep rate would adiabatically transfer atoms from $m_F = -1$ to $m_F = +1$.

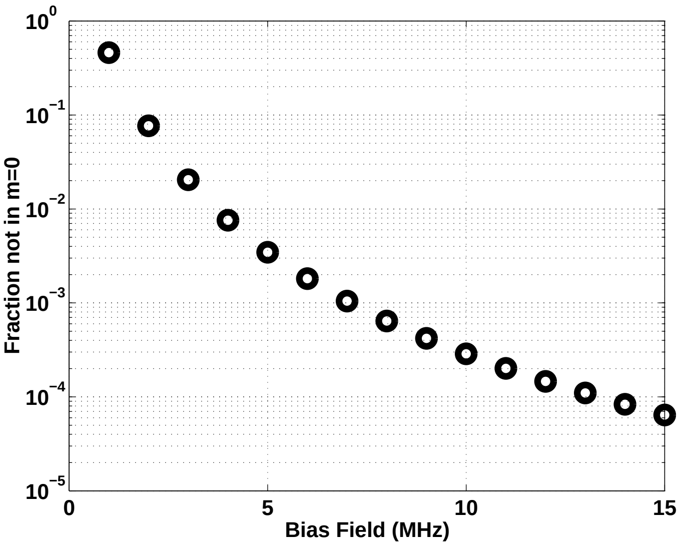

Figure 2.20. Calculated fraction of atoms not successfully swept into $m_F = 0$. The fraction has been minimized numerically with respect to Rabi frequency Ω and sweep rate $\dot{\delta}$ for each bias field, and the resulting minimum is plotted versus bias field.

applied to each avoided crossing separately (i.e., assuming $\Omega \ll q$):

$$f_0 = \exp\left(-\frac{\pi\Omega^4}{2q^2\dot{\delta}}\right)\left(1 - \exp\left(-\frac{\pi\Omega^2}{\dot{\delta}}\right)\right). \tag{2.67}$$

For a given bias field, i.e., fixed q, this probability has a maximum as a function of Ω and $\dot{\delta}$. The difference between this maximum value of f_0 and unity is plotted logarithmically in figure 2.20. One could, of course, sweep δ at a nonuniform rate, so that the $m = -1 \to m = +1$ avoided crossing is maximally nonadiabatic, while the $m = -1 \to m = 0$ crossing is adiabatic. If q is large enough, moreover, one could begin the sweep already on the far side of the first avoided crossing, thereby avoiding it altogether. Thus this result does not purport to represent a fundamental limit on the purity of $m = 0$ preparation, but merely to show the probable imperfections resulting from the most convenient experimental implementation of the state preparation.

In practice, Rabi pulses for $\gtrsim 100\,\text{kHz}$ are generated by an SRS 30 MHz function generator on "burst" mode, whose output is multiplexed with the evaporation-RF source and fed to a common amplifier. For lower-frequency pulses, the combined efficiency of the amplifier and of the coupling onto the gradient coil is too low, so that a different coil must be used. For this purpose, we have wound a $\sim$ 20-turn 15 cm-diameter coil which rests on the top imaging window of the vacuum chamber.

This coil can be driven directly by an unamplified SRS function generator in the tens-of-kilohertz range.

2.9 Stern-Gerlach Analysis

In order to ascertain the populations in the various Zeeman sublevels, we pulse on a magnetic gradient of order 5G/cm for 2 to 5ms and allow the differential impulse imparted to the respective populations to be translated into spatial separation by waiting for another 30ms before taking an absorption image. To avoid excessive interatomic collisions during separation, the gradient is usually pulsed on after the cloud has expanded for 2 ms, thereby reducing its density. A bias field of several hundred milligauss is present when the gradient is pulsed on, to avoid spin flips from transient switching effects. The gradient is of the spherical quadrupole form,

$$\mathbf{B}_{\text{grad}}(z) = B' z\hat{\mathbf{z}} - \frac{1}{2}B'\big(x\hat{\mathbf{x}} + y\hat{\mathbf{y}}\big), \tag{2.68}$$

so that depending on the orientation of the bias field, separation may be made to occur along any of the Cartesian directions, although the separation will be half as large in the two transverse directions.

2.10 Magnetic Field Control

For the spinor-condensate work described in chapters 3 and 4, it was essential to have precise control over the magnitude and direction of the magnetic field, as well as of magnetic gradients. Magnetic bias fields were controlled by means of three large coil pairs situated on the outside of the vacuum system. One, the “up-down” pair, was wrapped on and secured to the body of the vacuum chamber itself. The remaining pairs, labelled according to a geographical approximation “east-west” and “north-south,” provided fields along the two horizontal axes. As a result of imperfect planning, these coils were added to the system after optical breadboards for the MOT had been put into place, which, along with various other constraints from vacuum hardware, prevented the implementation of true Helmholtz pairs along orthogonal axes. The coils are sufficiently large, however, (between 10 cm and 20 cm) and close enough to the Helmholtz configuration that the gradients they produce are small. They are, moreover, linearly independent, so that an arbitrary field may be produced in a unique way, if not quite so simply as by orthogonal coils. As a technical point, the power supplies driving these coils were unipolar, so that large fields could only be produced in one spatial octant. Exploiting the background field (principally the earth's magnetic field), however, which had significant components along each direction, fields of either sign along any direction could be produced. The measured background field, including a calculated vertical field of approximately 200 mG from one always-on coil, was 580 mG at an angle of 147° to vertical and 68° to the magnetic-trap axis. This is approximately of the correct magnitude for the earth's field, which is reported by the National Oceanic and Atmospheric Administration to be 494 mG with an inclination of 61° in Berkeley, CA. Detailed numerical agreement is lacking, however, and is made worse by taking into account the vertical 200 mG field; this may be attributed either to the presence of residual ferromagnetism in neighboring vacuum hardware or to small leakage currents through the magnetic trap coils.

Whatever the source of the background field, it was necessary to eliminate it to the few-milligauss level. This was accomplished by transferring atoms into the optical trap and iterating a process of measuring the magnetic field and attempting to reduce it by changing the applied bias field. For initial measurements, the magnetic field was determined by sweeping a RF field and measuring Stern-Gerlach-separated populations in absorption imaging. By initially making large sweeps at high power and gradually reducing the sweep range and power, it was possible to localize the RF resonance to a few kHz, although care needed to be taken to ensure that the fundamental and not a harmonic of the RF source was resonant. In this manner, it was possible, although somewhat laborious, to obtain curves of magnetic field versus current in each of the three pairs of bias coils. These curves were of the expected hyperbolic shape, linear at large currents but with a minimum value determined by the transverse field in each case. By fitting these curves to determine the center current and the large-current slope, it was possible to obtain preliminary estimates for the dependence of field on the applied currents, although the nonorthogonality of the horizontal coils meant that these minima were shifted with respect to the zero-field values. In order to measure the angles among the field components of the different coils and between these and the trap axis, the magnitude of each separate component was measured, followed by measurements of the resultant when pairs of components were applied simultaneously. A field along the trap axis was applied by running current through the curvature coils. From three such measurements, the angle between any two magnetic fields could then be determined by the law of cosines. Once these angles were known, then orthogonal linear combinations of east-west and north-south fields could be formed, and by measuring the magnitude of the total field as these were varied, hyperbolic curves could again be measured, but now centered on the true zero-field currents.

This procedure worked well, but required many cycles of the experiment to measure the field at each current. Moreover, measurements at many different currents were needed to construct the field-current curve. In addition, at low bias fields it was difficult to obtain reliable RF resonance measurements, as a result of the limited current output of our RF source. A more efficient method for obtaining finer zeroing of the bias field was then implemented, exploiting the kinetics-imaging capability of our CCD camera (see chapter 3). The magnetic field was slowly swept from a large downward-pointing value to a large upward-pointing value while phase-contrast images were recorded at regular known intervals. The spins of the atoms, initially in $m_F = -1$, were assumed to follow the field adiabatically. As a result of the stronger interaction of the circular-polarized probe σ_- light with the spin-down atoms, the image went from dim when the field was pointing down to bright when the field was pointing up. By plotting the brightness of the image versus the applied vertical field at the time the image was taken, both the field zero in the vertical direction (the half-maximum point) and the magnitude of the transverse field (the width of the transition region) could be obtained in a single shot. By varying the transverse field components to make the transition from dim to bright images as sharp as possible, the magnitude of the transverse field could be efficiently nulled out. This method produced minimum fields of a few milligauss, at which level fluctuations in the bias field made further reduction impracticable.

Two separate pairs of coils have been put in place to null out magnetic gradients, one along the trap axis and one at an angle of 60° to the axis. These pairs are likewise geometrically constrained, so that they are not of the anti-Helmholtz configuration. Together, they allow gradients of the form $\frac{\mathrm{d}B_z}{\mathrm{d}z}$ and $\frac{\mathrm{d}B_x}{\mathrm{d}z}$ to be nulled out. The measurement of magnetic gradients is described in chapters 3 and 4.

2.11 Imaging Systems

Several imaging systems were employed for this work. For initial characterization and optimization of the MOT, an imaging system of magnification $m \approx \frac{1}{2}$ along an otherwise unused diagonal side viewport was employed. For time-of-flight (i.e., free expansion) imaging of BECs, a higher magnification system ($m \approx 2.5$) imaging from above was initially employed. For in-situ imaging, higher magnification yet was required, and a high-quality two-stage $m = 12$ system was put in place. On a daily basis, moving the expensive and delicate camera from one imaging station to another is undesirable, and so a way of adding lenses to the $m = 12$ system to convert it to a $m = 1.8$ imaging system with the same image plane was devised. This quick conversion scheme has significantly lowered the barrier to switching between time-of-flight and in-situ imaging and encouraged quick diagnosis of problems in the imaging system best suited to the particular task at hand. An additional high-quality $m = 5$ imaging system was implemented along the condensate axis, primarily for the purpose of aligning the optical trap and eliminating aberrations in its focus.

A diagram of the $m = 12$ imaging system is shown in figure 2.21. An primary imaging system of $m = 6$ forms an image at the focus of the 750 mm lens. Both the 145 mm achromatic lens and the 750 mm plano-convex lens are employed at infinite conjugate ratio for optimum performance. A secondary system, consisting of two more lenses at infinite conjugate ratio, namely another 145 mm achromatic lens and a 300 mm lens, re-magnifies the primary image by a factor of 2 for a total magnification of 12 at the camera. The two-stage imaging system allows a mask to be placed at the intermediate image, so that only a small portion of the camera may be exposed at a particular time. This is important for so-called kinetics imaging, as described in chapter 3. In the absence of refraction by the sample, probe light is focused at the approximate composite focal length of 120 mm beyond the 750 mm lens; a $\pi/2$ phase dimple (drawn for convenience as a dot) is placed at this focus for phase-contrast imaging. To switch rapidly to $m = 1.8$ imaging, typically used for free-expansion

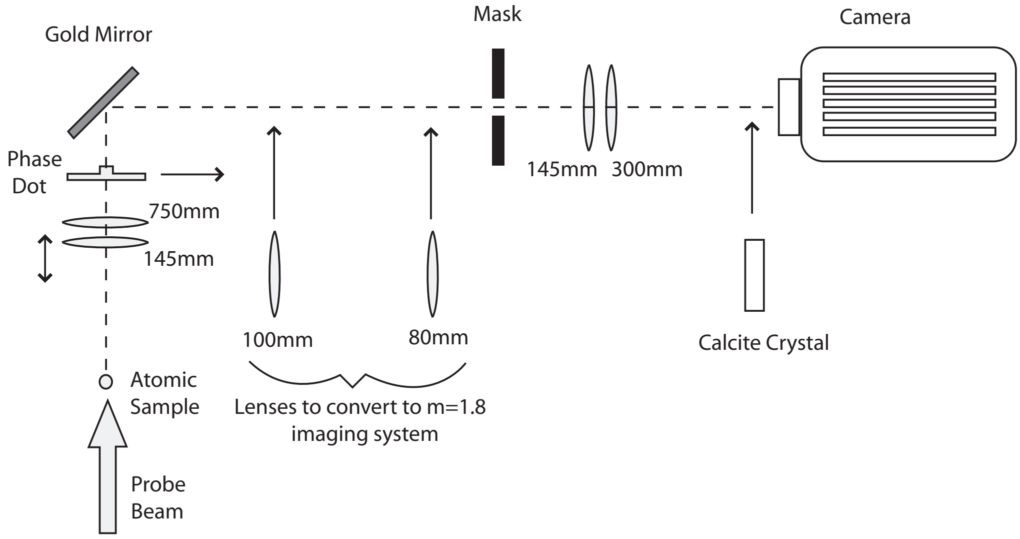
Gold Mirror
Mask
Camera
Phase Dot
750mm
145mm
145mm
300mm
100mm
80mm
Calcite Crystal
Atomic Sample
Lenses to convert to m=1.8 imaging system
Probe Beam

Figure 2.21. Diagram of imaging system. Front lens pair (145 mm and 750 mm) form ×6 magnified image at location of mask. Secondary pair (145 mm and 300 mm) re-magnify ×2, for total magnification ×12. Front 145 mm lens can be translated to keep free-falling atoms in focus. An optional phase dot (in reality, a dimple) ≈ 120 mm from the front lens pair is used for phase-contrast imaging, as described in chapter 3. Two more lenses (100 mm and 80 mm) can be inserted to create lower-quality ×1.8-magnification system. A calcite crystal may be inserted to allow separate simultaneous imaging of both linear polarizations.

or time-of-flight absorption imaging, a 100 mm lens is inserted ≈ 340 mm after the 750 mm lens, forming an image of the atomic sample that occurs near the focal point of the 100 mm lens. The net magnification of this new image is approximately 1.3. In order not to have to move the camera, a second 80 mm lens is placed 130 mm from the image plane of the $m = 6$ imaging system, so as to re-image onto this same image plane, with a further demagnification of about 0.65. Then, as before, the secondary (now tertiary) $m = 2$ imaging system re-images onto the camera, for a net magnification of $m = 1.8$. Although the $m = 1.8$ imaging system has too many lenses of too short focal lengths to be of very high resolution, its ease of use has frequently outweighed its somewhat inferior quality. The $m = 12$ imaging system, on the other hand, has been very carefully aligned, with considerable pains taken to ensure centering and correct tilt of lenses. Magnification has been measured by first focusing on a small atomic sample, then intercepting the imaging path with a mirror and inserting a USAF test pattern at approximately the same distance as the atoms, and finally adjusting the position of the test pattern along the imaging axis to bring it into focus. This procedure should ensure that the test pattern is indeed at the same distance as the atomic sample. Comparing the known size of the test-pattern figures to the size on the camera (using the specified pixel size of 13 μm) allows accurate calibration of the magnification. More recently, we have been able more fully to characterize the resolution of the $m = 12$ imaging system by imaging oscillatory magnetization patterns in spinor condensate samples; this is described in detail in chapter 4.

Chapter 3

Spinor Condensates

This chapter will discuss work on spinor Bose condensates and a new technique for probing their magnetization; this work was presented in the following publication:

- *J. M. Higbie, L. E. Sadler, S. Inouye, A. P. Chikkatur, S. R. Leslie, K. L. Moore, V. Savalli, and D. M. Stamper-Kurn. "Direct Nondestructive Imaging of Magnetization in a Spin-1 Bose-Einstein Gas" Phys. Rev. Lett.* **95***, 050401 (2005).*

3.1 Overview of Prior Work

BECs with a spin degree of freedom are naturally more complex objects than single-component or scalar BECs. They allow the exploration of many interesting questions, such as the nature of and stability of their topological excitations, the structure of their ground states, and the character of macroscopic coherence among multiple condensates.

The first multicomponent quantum fluid to be discovered and studied was ^{3}He. Interest in ascertaining whether ^{3}He could undergo a superfluid transition analogous to that for ^{4}He had existed even before 1948, when the former was first successfully liquefied by Sydoriak et al. [63]. Further interest was aroused when Bardeen et al. [64] proposed their famous pairing theory of superconductivity, stimulating speculation that ^{3}He could undergo a similar pairing transition to a superfluid state.

The era of speculation, however, was ended by the discovery of the A and B phases of ^{3}He by Osheroff et al. [65]. The difficulty of ascertaining the true nature of the samples under study is illustrated by the fact that Osheroff et al. initially ascribed the kinks in their pressurization curve to solidification of the sample, rather than to a superfluid transition. This misapprehension was quickly rectified, and within a remarkably short period of time the complex structure of ^{3}He was elucidated [66].

Although fascinating and intricate, the structure of ^{3}He is sufficiently complicated that it is difficult to consider it as the generic prototype of quantum fluids with spin. In order to increase the number of such systems available for study, and to gain some more general understanding of which features are general and which system-specific, it is appealing to explore other multicomponent quantum fluids than the one nature (with the aid of dilution refrigerators and tritium production) has provided. [1]

The creation of dilute-gas Bose-Einstein condensates in 1995 opened many new possibilities for the study of different varieties of quantum fluid. The first multicomponent condensates were produced shortly after the production of the first alkali-atom BECs by Myatt et al. [67] and consisted of two magnetically trappable Zeeman states of ^{87}Rb in distinct hyperfine levels. The JILA group followed up this initial success with an impressive series of experiments, in which the coherence and dynamics of this two-component system were studied.

The following year, the Ketterle group at MIT succeeded in confining a BEC of sodium in a far off-resonant optical trap [62]. As discussed in section 2.7, the AC Stark shift at sufficiently large detuning is independent of the spin projection of an atom, so that a trap formed by a focused red-detuned laser beam will confine all spin states equally, allowing spin-changing collisions to contribute to dynamics, and not merely to loss, as in a magnetic trap. The resulting condensates, in which the full rotationally-closed hyperfine multiplet was allowed to participate in dynamics, were christened ‘spinor’ condensates by the Ketterle group, using ‘spinor’ in its nonrelativistic-quantum-mechanics sense, i.e., a several-component wave function representing spin. This work was extended to studies of the phase diagram of spinor condensates of sodium, including the miscibility of spin components [68] and of tunneling dynamics starting from metastable domain configurations [69, 70].

---

[1] A possible exception is the proposed superfluid quark-gluon system, which, however, is very challenging to observe in the laboratory.

A new generation of spinor-condensate experiments has followed closely on the heels of these pioneering experiments. Notable work in recent years has included the observation in spin-1 ^{87}Rb condensates of spin relaxation dynamics providing evidence for the ferromagnetic character of the rubidium ground state [71] as well as coherent spin oscillations [72]. Also extending the frontier of multicomponent quantum fluids have been the experiments by the Sengstock group and the Hirano group on $F = 2$ spinor condensates of ^{87}Rb [73, 74].

3.2 Theoretical Description of Spin-1 Spinor Condensate

The Hamiltonian of a system of noninteracting spin-1 Bosons is

$$H_{\text{free}} = \sum_m \int \mathrm{d}^3\mathbf{x}\Psi_m^\dagger(\mathbf{x},t)\left(-\frac{\hbar^2}{2m}\nabla^2 + V(\mathbf{x})\right)\Psi_m(\mathbf{x},t), \tag{3.1}$$

where $\Psi_m(\mathbf{x},t)$ is a three-component position-space bosonic field operator, satisfying the equal-time commutation relation $[\Psi_m(\mathbf{x},t), \Psi_{m'}^\dagger(\mathbf{x}',t)] = \delta^3(\mathbf{x}-\mathbf{x}')\delta_{m,m'}$. Here, the spin index of Ψ has been written explicitly, but for what follows it will be convenient to suppress this index and to consider Ψ as a three-component column vector with $\Psi^\dagger$ a three-component row vector. To the free Hamiltonian must be added an interaction term describing the S-wave collisions of ultracold atoms. The interaction Hamiltonian is required to be a scalar by rotational symmetry, and should be quartic in the field operators Ψ and $\Psi^\dagger$ (specifically containing two creation and two annihilation operators), as it describes two incoming and two outgoing particles in the scattering process. There are precisely three independent quartic terms, corresponding to the three possible values of the total incoming spin, 0, 1, and 2. Moreover, the total spin-1 composite state of two individual spin-1 bosons is antisymmetric, which is incompatible with Bose symmetry when the incoming particles are in the same motional quantum state, as is certainly the case for atoms in a BEC. As a result, the collision Hamiltonian may be parametrized by two independent parameters, a_0 and a_2, the total-spin-0 and total-spin-2 scattering lengths. A very convenient way of writing this interaction Hamiltonian due to Ho [75] is in terms of the single-atom

3×3 spin matrices $\mathbf{F}$:

$$H_{\text{int}} = \frac{c_0}{2}\int \mathrm{d}^3\mathbf{x}\left(\Psi^\dagger(\mathbf{x})\cdot\Psi(\mathbf{x})\right)^2 + \frac{c_2}{2}\int \mathrm{d}^3\mathbf{x}\left(\Psi^\dagger(\mathbf{x})\cdot\mathbf{F}\cdot\Psi(\mathbf{x})\right)^2, \tag{3.2}$$

where, using the fact the $\mathbf{F}_1 \cdot \mathbf{F}_2 = \frac{1}{2}(\mathbf{F}_{\text{tot}}^2 - \mathbf{F}_1^2 - \mathbf{F}_2^2)$, the coefficients $c_{0,2}$ can be related to the scattering lengths $a_{0,2}$ as

$$\begin{aligned} c_0 &= \frac{4\pi\hbar^2}{m}\left(\frac{a_0 + 2a_2}{3}\right) \\ c_2 &= \frac{4\pi\hbar^2}{m}\left(\frac{a_2 - a_0}{3}\right). \end{aligned} \tag{3.3}$$

For ^{87}Rb, the values of a_0 and a_2 are inferred to be $a_0 = 101.8a_{\text{Bohr}}$ and $a_2 = 100.4a_{\text{Bohr}}$, implying that c_2 is negative. As a result, states with large net spin, and hence large expectation value of $\mathbf{F}$, are lower in energy. Thus, spinor condensates of ^{87}Rb are called ferromagnetic, since their mean-field ground state has all spins aligned in a single direction. In this context, it is of interest to note that when spin-conservation is taken into account, the many-body ground state can be highly correlated [76]; these correlated states, however, are expected to be rather fragile [77], and have not been observed experimentally.

3.2.1 Mean-Field Solutions at Finite Magnetic Field

In a BEC, all atoms occupy the same single-particle state, to the extent that the quantum depletion can be ignored, which is generally an excellent approximation for dilute-gas Bose condensates. As a result, the Hartree solutions to the equations of motion, which approximate the many-body wave function as a product of identical single-particle wave functions, are usually very accurate. They are mean-field solutions, meaning that each atom responds to an average "potential" due to the other atoms, neglecting any possible correlations.

The Hamiltonian $H = H_{\text{free}} + H_{\text{int}}$ retains the same form in this approximation, but is reinterpreted as an energy functional,

$$E[\psi] = \int \mathrm{d}^3\psi^\dagger\left(-\frac{\hbar^2}{2m}\nabla^2 + V(\mathbf{x})\right)\psi + \frac{c_0}{2}(\psi^\dagger\psi)^2 + \frac{c_2}{2}(\psi^\dagger\mathbf{F}\psi)^2,$$

where ψ is a three-component c-number-valued function. Likewise, the mean field equations of motion for the condensate wave function are of the same form as the field equations of motion derived from the many-body Hamiltonian, except that the quantum field Ψ is replaced by the classical field ψ. The equation of motion so obtained is known as the Gross-Pitaevskii equation [78, 79] . For a spin-1 boson in the finite magnetic field $\mathbf{B}(\mathbf{r})$, neglecting the quadratic Zeeman shift, the Gross-Pitaevskii equation takes the form

$$i\hbar\frac{\partial}{\partial t}\psi_a = \Big[-\frac{\hbar^2}{2m}\nabla^2\delta_{ab} + V(\mathbf{x})\delta_{ab} + g_F m_F \mu_B \mathbf{B}(\mathbf{r})\cdot\mathbf{F}_{ab} \tag{3.4}$$

$$+ c_0\psi_c^\dagger\psi_c\delta_{ab} + c_2(\psi_c^\dagger\mathbf{F}_{cd}\psi_d)\cdot\mathbf{F}_{ab}\Big]\psi_b, \tag{3.5}$$

where repeated indices are summed.

3.3 Limitations of Prior Imaging Techniques

Although much can and has been learned about spinor condensates by destructively imaging their separate spin populations after Stern-Gerlach separation (see 2.9) and free expansion ("time-of-flight") absorption imaging, it is difficult or impossible by means of these techniques to bring to light some of the most important aspects of spinor condensates. Indeed, the one-body density matrix for atoms of spin 1 is of the form

$$\rho = \begin{pmatrix} \rho_{11} & \rho_{10} & \rho_{1-1} \\ \rho_{01} & \rho_{00} & \rho_{0-1} \\ \rho_{-11} & \rho_{-10} & \rho_{-1-1} \end{pmatrix}. \tag{3.6}$$

Adiabatic Stern-Gerlach separation allows one to extract the three on-diagonal components, i.e., the populations, but offers no information about the off-diagonal coherences, which constitute the bulk of the information contained in the density matrix (6 real parameters compared to 2 for the populations, assuming $\mathrm{Tr}[\rho] = 1$).

In terms of the magnetization of the condensate, a determination of the populations clearly determines the longitudinal magnetization (i.e., the magnetization along the magnetic field direction and quantization axis, taken to be $\hat{\mathbf{z}}$) by

$$\mathrm{Tr}[\rho\mathbf{F}_z] = \rho_{11} - \rho_{1-1}, \tag{3.7}$$

while the coherences determine the transverse magnetization

$$\mathrm{Tr}[\rho\mathbf{F}_+] = \sqrt{2}\left(\rho_{01} - \rho_{-10}\right). \tag{3.8}$$

Even for a pure state, which is of particular interest for the study of BEC, there are two independent phases which determine the spin state. A simple example is afforded by comparison of the two states (in the spherical basis)

$$\begin{pmatrix} \frac{1}{2} \\ \frac{1}{\sqrt{2}} \\ \frac{1}{2} \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} \frac{1}{2} \\ \frac{i}{\sqrt{2}} \\ \frac{1}{2} \end{pmatrix}. \tag{3.9}$$

The former represents a spin polarized with maximal spin projection +1 along a particular axis (i.e., a "ferromagnetic" state), in this case the $\hat{\mathbf{x}}$ axis. The second is an unpolarized state with zero spin projection in all directions; specifically, it is the $m = 0$ eigenstate of $\frac{1}{\sqrt{2}}(\mathbf{F}_y - \mathbf{F}_z)$. Such states, for which the vector equation $\langle \mathbf{F} \rangle = 0$ holds, are referred to as "polar" or unmagnetized states and are ground states of an antiferromagnetic condensate, i.e., one for which $c_2 > 0$, as is the case for sodium.

Although these states are very different in their spin character, a Stern-Gerlach analysis of a condensate in the two states would give identical results, since they differ only by a phase. Consequently, for the study of spin textures or other patterns of magnetization in the spinor condensate, a phase-sensitive imaging method is very desirable.

Of course, it is possible to map coherences into populations via Rabi pulses, as, for instance, in Ramsey spectroscopy, but such an approach is only possible if the relative phase differences are stable over the evolution period. Since typical evolution times for spinor Bose condensates are of the order of hundreds of milliseconds, the magnetic field stability required ($\sim 10\,\mu$G) is difficult to achieve without magnetic shielding of the experiment. Magnetically shielding a conventional BEC apparatus, containing a magnetic trap whose coils produce fields of hundreds of Gauss, is a significant engineering challenge. A BEC apparatus using small trapping coils close to the atoms and running small ($\sim$ 25A) currents has been successfully placed in a magnetically shielded environment by Esslinger et al. [59]. The development of all-optical BEC by [28] or transport to a dedicated "science chamber" as in [80] also lower the barrier to making a magnetically shielded BEC, though the requirements of optical access remain challenging.

It is worth noting, in this context, that the experiments performed by the JILA group on pseudospin-1/2 ^{87}Rb did not suffer from this difficulty, because the two

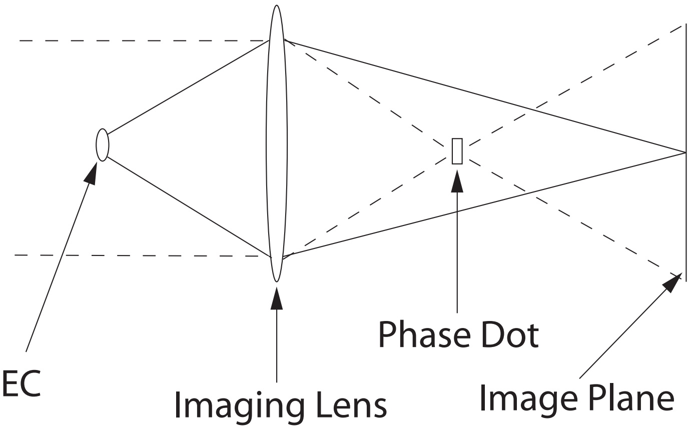

Figure 3.1. Diagram of phase-contrast imaging. Scattered light is imaged at the camera, where it interferes with unscattered light that has been phase-shifted by a phase dot at the Fourier plane.

spin states employed in those experiments possessed nearly equal magnetic moments. Consequently, magnetic field noise was common-mode as far as the relative phase between the two states was concerned, so that coherences could be measured as well as populations by converting the former into the latter.

A second limitation that has afflicted most prior experiments on multicomponent alkali-gas condensates results from the use of destructive absorption imaging. Destructive imaging is inefficient for studying dynamics, since a new sample must be prepared and imaged for each time step in the dynamics. In cases where the dynamics are not deterministic, moreover, as in the case of an instability nucleated by quantum or thermal noise, the comparison of images taken on separately prepared condensates is meaningless in any but a statistical sense. For such cases, if one wishes to follow the dynamics of an evolving sample, a nondestructive technique is necessary. An important precedent in the use of nondestructive probes of a multicomponent BEC is the experiment of [81], in which discrete nondestructive images and nearly continuous streak-camera recording of the phase of Rabi oscillation were obtained.

3.4 In-situ Imaging

Given the need for an imaging technique which is sensitive to the relative phase of the spin states in a spinor condensate (or equivalently, to the direction and magnitude of magnetization), it is natural to exploit the fact that the interaction strength of polarized light with an atomic medium depends both on the spin state or the magnetization of the atoms. A well-established method of obtaining nondestructive *in situ* images exists in the form of phase-contrast imaging (described below); by using polarized light appropriately detuned from an atomic resonance, then, it is possible to add magnetization-sensitivity to the list of the virtues of phase-contrast imaging.

3.4.1 Phase-Contrast Technique

The method of phase-contrast imaging was invented by Zernicke in the 1930s [82] and has been extensively applied to the imaging of refractive objects such as biological cells. In this method phase-shifted light is made to interfere with unshifted light, as is generally the case for measurements of optical phase. The elegance and practicality of the technique lie in the simplicity of the means by which the splitting and recombination occur. A nonuniform refractive object causes light to acquire an angular spread; consequently this spread may be used to differentiate the unscattered wave (that is, the wave front which would have existed with no object present) from the scattered wave (the wave which must be added to the unscattered wave to obtain the actual distribution of light). Indeed, if the unscattered light forms an intermediate focus as in figure 3.1 at a plane (referred to as the Fourier plane) prior to that of the image, then the scattered wave, which originates at the object and reconverges at the image plane, will possess a much larger spatial extent at the Fourier plane than the tightly focused unscattered wave. Thus, by placing a small extra thickness of dielectric at this focus, one may introduce a differential phase shift between the scattered and unscattered light beams. Moreover, since the two beams remain co-propagating, no special recombination is necessary and the requirements on the path-length stability of the two (non-distinct) arms of the interferometer are greatly reduced.

In order to provide a simple model for the subsequent calculation, a short version of the standard derivation of the phase-contrast signal will be presented here. If the incoming imaging light is a plane wave $E_0 e^{ikz}$, then after passing through a refractive medium, the electric field will be proportional to $E_0 e^{ikz+i\phi}$. This is trivially rewritten as the unscattered plus the scattered light

$$\begin{aligned} E_{\text{unscatt}} &= E_0 e^{ikz} \\ E_{\text{scatt}} &= E_0 e^{ikz}\left(e^{i\phi}-1\right). \end{aligned} \tag{3.10}$$

The effect of the phase dot is to phase shift the unscattered light by α, usually $\pi/2$, so that the final intensity distribution reconstituted at the image plane is

$$\begin{aligned} |E|^2 &= E_0^2 |e^{i\alpha}+e^{i\phi}-1|^2 \\ &= E_0^2\left(3+2\cos(\alpha-\phi)-2\cos\alpha-2\cos\phi\right), \end{aligned} \tag{3.11}$$

which for $\alpha=\pi/2$ and $\phi \ll 1$ reduces to

$$\frac{|E_{\text{with atoms}}|^2}{|E_{\text{ no atoms}}|^2} = 1+2\phi. \tag{3.12}$$

The linear dependence of the signal on the phase shift ϕ is a desirable feature for small phase shifts, since other schemes (e.g., dark-ground imaging) produce signals which are quadratic in ϕ. This may be thought of as the result of interfering a large signal (the unscattered light) with a small signal (the scattered light). For a dispersive object, these two waves do not interfere in amplitude at the image plane, but only in the phase quadrature. By deliberately introducing a phase shift between them, they may be made to interfere in amplitude.

3.4.2 Magnetization Sensitivity

To ascertain the magnetization sensitivity of phase-contrast imaging, one can calculate the polarization state of light having passed through an atomic medium in a particular spin state. Here, we assume that the state of the atoms is an uncorrelated mean-field state of the form $\sqrt{n(\mathbf{r})}\xi(\mathbf{r})$, where $\xi(\mathbf{r})$ is a unit-normalized three-component spinor and $n(\mathbf{r})$ is the density. More complicated collective spin states must be calculated in a more sophisticated formalism. This polarization state of the light will be a function of position, and is thus capable of being reformed into an image by means of the appropriate polarization optics. This calculation is straightforward, but will be recorded here for reference.

The connection between the microscopic spin state of the atoms and the classical concept of phase shifts and polarization rotation is provided by the dielectric constant, which becomes a tensor when different polarization states of light are considered. The dielectric tensor in component form is

$$\epsilon_{jk} = \delta_{jk} + \frac{6\pi n}{k_0^3} \sum_e \frac{c_j^{ge} c_k^{eg}}{\tilde{\delta}_e - i}, \tag{3.13}$$

where c_j^{ge} is the Clebsch-Gordan coefficient between ground and excited states $|g\rangle$ and $|e\rangle$, normalized to the cycling transition (i.e., $F = 2, m_F = 2 \rightarrow F' = 3, m'_F = 3$), k_0 is the resonant wave vector, and $\tilde{\delta}_e$ is the detuning from resonance with the excited state $|e\rangle$ in units of the half-linewidth $\frac{\gamma}{2}$. For the calculation of far off-resonant effects ($\tilde{\delta} \gg 1$), however, we neglect the i in the denominator of equation (3.13). For a typical condensate density of $n \sim 3 \times 10^{14}\,\mathrm{cm}^{-3}$ at the wave vector $k_0 = 8 \times 10^4\,\mathrm{cm}^{-1}$, the dimensionless prefactor $6\pi n/k_0^3$ is numerically close to 11. For typical detunings of $-230\,\mathrm{MHz}$ from the D1 $F = 1 \longrightarrow F' = 2$ transition (employed for the work described here–see figure 3.2), $\tilde{\delta} \sim 70$, while the maximum value of c_i^{eg} on this line is $\frac{1}{2}$, so that the dielectric tensor differs from the identity by no more than $\sim 4\%$. This fact allows a simplification of many subsequent expressions.

For brevity, we adopt the notation that $\epsilon_{jk} = \delta_{jk} + \mu_{jk}$. Furthermore, we consider only polarizations orthogonal to the direction of propagation of the light. This amounts to projecting the full 3×3 matrix ϵ_{jk} onto a two-dimensional subspace. The projection is valid provided effects such as the angular separation of ordinary and extraordinary rays need not be considered across the length of the medium. The precise choice of basis vectors is immaterial, and we will label the polarization basis vectors as $\hat{\mathbf{e}}_1 \leftrightarrow (1\,,\,0)^T$ and $\hat{\mathbf{e}}_2 \leftrightarrow (0\,,\,1)^T$, which may equally well represent an orthonormal pair of linear or circular polarizations.

The 2×2 matrix ϵ can be decomposed into Pauli matrices as

$$\epsilon = 1 + \bar{\mu} + A\hat{\mathbf{n}} \cdot \tilde{\sigma} \tag{3.14}$$

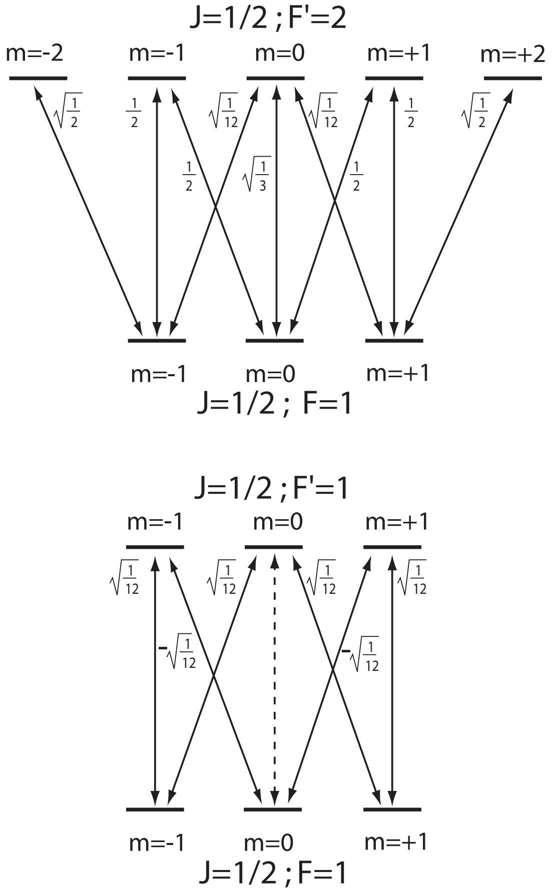
J=1/2 ; F'=2
m=-2
m=-1
m=0
m=+1
m=+2
$\sqrt{\frac{1}{2}}$
$\frac{1}{2}$
$\sqrt{\frac{1}{12}}$
$\sqrt{\frac{1}{12}}$
$\frac{1}{2}$
$\sqrt{\frac{1}{2}}$
$\frac{1}{2}$
$\sqrt{\frac{1}{3}}$
$\frac{1}{2}$
m=-1
m=0
m=+1
J=1/2 ; F=1
J=1/2 ; F'=1
m=-1
m=0
m=+1
$\sqrt{\frac{1}{12}}$
$\sqrt{\frac{1}{12}}$
$\sqrt{\frac{1}{12}}$
$\sqrt{\frac{1}{12}}$
$-\sqrt{\frac{1}{12}}$
$-\sqrt{\frac{1}{12}}$
m=-1
m=0
m=+1
J=1/2 ; F=1

Figure 3.2. Level diagram of optical transitions on the D1 $F = 1 \rightarrow F' = 2$ and $F = 1 \rightarrow F' = 1$ transitions. Clebsch-Gordon coefficients are listed next to the corresponding lines, indicating polarization dependence. Interaction strength of σ_+ light with an atom in $m = +1$ is six times stronger than with an atom in $m = -1$; this allows magnetization of an atomic cloud to be imaged, as described in the text.

where

$$\begin{aligned}
\bar{\mu} &\equiv \frac{\mu_{11}+\mu_{22}}{2} \\
A &\equiv \sqrt{\left(\frac{\mu_{11}-\mu_{22}}{2}\right)^2 + |\mu_{12}|^2} \\
\hat{\mathbf{n}} &\equiv A^{-1}\left\{\mathrm{Re}(\mu_{12})\,\hat{\mathbf{x}} + \mathrm{Im}(\mu_{12})\,\hat{\mathbf{y}} + \left(\frac{\mu_{11}-\mu_{22}}{2}\right)\hat{\mathbf{z}}\right\}. \qquad (3.15)
\end{aligned}$$

The eigenvalues of ϵ may easily be read off by comparison with the problem of a spin-$\frac{1}{2}$ particle in a magnetic field as $\lambda_\pm = 1 + \bar{\mu} \pm A$. The eigenvectors are obtained by means of the rotation operator

$$U \equiv \exp\left(-i\,\frac{\theta}{2}\,\hat{\mathbf{m}}\cdot\vec{\sigma}\right), \qquad (3.16)$$

where

$$\begin{aligned}
\hat{\mathbf{m}} &\equiv |\mu_{12}|^{-1}\Big(-\mathrm{Im}(\mu_{12})\,\hat{\mathbf{x}} + \mathrm{Re}(\mu_{12})\,\hat{\mathbf{y}}\Big) \\
\tan\theta &\equiv \frac{2|\mu_{12}|}{\mu_{11}-\mu_{22}}. \qquad (3.17)
\end{aligned}$$

The dielectric tensor is then

$$\epsilon = U\begin{pmatrix} \lambda_+ & 0 \\ 0 & \lambda_- \end{pmatrix}U^\dagger, \qquad (3.18)$$

from which it is clear that the index of refraction $n_{\mathrm{ref}} \equiv \sqrt{\epsilon}$ takes the form

$$n_{\mathrm{ref}} = U\begin{pmatrix} \sqrt{\lambda_+} & 0 \\ 0 & \sqrt{\lambda_-} \end{pmatrix}U^\dagger. \qquad (3.19)$$

The evolution operator corresponding to passage through a medium described by this index-of-refraction tensor is

$$U_{\mathrm{evol}} \equiv e^{ik_0 n_{\mathrm{ref}} z} = U\begin{pmatrix} e^{ik_0\sqrt{\lambda_+}} & 0 \\ 0 & e^{ik_0\sqrt{\lambda_-}} \end{pmatrix}U^\dagger. \qquad (3.20)$$

Using the more explicit form (3.16) of the operators U and $U^\dagger$, one can write U_{evol} as

$$\begin{pmatrix} e^{ik_0z\sqrt{\lambda_+}}\cos^2\frac{\theta}{2} + e^{ik_0z\sqrt{\lambda_-}}\sin^2\frac{\theta}{2} & \sin\frac{\theta}{2}\cos\frac{\theta}{2}\frac{\mu_{12}^*}{|\mu_{12}|}\left(e^{ik_0z\sqrt{\lambda_+}} - e^{ik_0z\sqrt{\lambda_-}}\right) \\ \sin\frac{\theta}{2}\cos\frac{\theta}{2}\frac{\mu_{12}}{|\mu_{12}|}\left(e^{ik_0z\sqrt{\lambda_+}} - e^{ik_0z\sqrt{\lambda_-}}\right) & e^{ik_0z\sqrt{\lambda_+}}\sin^2\frac{\theta}{2} + e^{ik_0\sqrt{\lambda_-}}\cos^2\frac{\theta}{2} \end{pmatrix}. \qquad (3.21)$$

If the initial polarization was $\chi_{\text{in}} \equiv (1\,,\,0)^\dagger$, which depending on the choice of basis could represent either linear or circular polarization, then after passing through the sample it would become

$$\begin{aligned} \chi_{\text{out}} &\equiv U_{\text{evol}} \cdot \chi_{\text{in}} \\ &= \begin{pmatrix} e^{ik_0 z\sqrt{\lambda_+}} \cos^2\frac{\theta}{2} + e^{ik_0 z\sqrt{\lambda_-}} \sin^2\frac{\theta}{2} \\ \sin\frac{\theta}{2}\cos\frac{\theta}{2}\frac{\mu_{12}}{|\mu_{12}|}\left(e^{ik_0 z\sqrt{\lambda_+}} - e^{ik_0 z\sqrt{\lambda_-}}\right) \end{pmatrix}. \end{aligned} \tag{3.22}$$

This polarization state can contain a considerable amount of information about the state of the atoms. It can be analyzed in a number of ways. First, let us compute the effect of phase-shifting the unscattered light, as in phase-contrast imaging. The light in the absence of atoms would have the polarization state

$$\chi_{\text{out}}^{\text{unscatt.}} \equiv \begin{pmatrix} e^{ik_0 z} \\ 0 \end{pmatrix}. \tag{3.23}$$

Thus, analogously to equation (3.10), we can write the scattered light as $\chi_{\text{out}} - \chi_{\text{out}}^{\text{unscatt.}}$, so that a relative phase shift of $\pi/2$ applied to the unscattered light will result in the state

$$\begin{aligned} \chi_{\text{out}}^{\text{shift}} &\equiv i\chi_{\text{out}}^{\text{unscatt.}} + \left(\chi_{\text{out}} - \chi_{\text{out}}^{\text{unscatt.}}\right) \\ &= e^{ik_0 z} \begin{pmatrix} i + e^{ik_0 z(\sqrt{\lambda_+}-1)} \cos^2\frac{\theta}{2} + e^{ik_0 z(\sqrt{\lambda_-}-1)} \sin^2\frac{\theta}{2} - 1 \\ \sin\frac{\theta}{2}\cos\frac{\theta}{2}\frac{\mu_{12}}{|\mu_{12}|}\left(e^{ik_0 z(\sqrt{\lambda_+}-1)} - e^{ik_0 z(\sqrt{\lambda_-}-1)}\right) \end{pmatrix}. \end{aligned} \tag{3.24}$$

As noted earlier, the detunings typically employed in the present work are sufficient to warrant use of the dilute approximation, according to which deviations of the dielectric tensor from the identity are small compared to unity. In the present instance, this is equivalent to the statement that $\sqrt{\lambda_\pm} - 1 \ll 1$. A related (but usually stronger) approximation is the thin-sample approximation, according to which $k_0 z(\sqrt{\lambda_\pm} - 1) \ll 1$. For the experiments described here, the condensate thickness along the imaging direction is approximately $z \approx 3\,\mu$m (the full width at half-maximum density), so that $k_0 z \approx 24$. For the strongest transitions, as estimated earlier, $\sqrt{\lambda_\pm} - 1 \approx 0.04$, so that the thin-sample approximation is on the verge of validity, $k_0 z(\sqrt{\lambda_\pm} - 1) \sim 1$, though the full expression (3.24) must be used to obtain the most accurate results. The simplicity of the results in the combined thin-sample and dilute approximations makes them useful for intuition and for rough calculations, even though they are not precisely applicable for the experiments under consideration.

In the dilute approximation, neglecting an over-all phase,

$$\chi_{\text{out}}^{\text{shift}} \approx \begin{pmatrix} i - 1 + e^{ik_0 z\bar{\mu}/2}\left(\cos\frac{k_0 zA}{2} + i\cos\theta\sin\frac{k_0 zA}{2}\right) \\ i\sin\theta\frac{\mu_{12}}{|\mu_{12}|}\left(e^{ik_0 z\bar{\mu}/2}\sin\frac{k_0 zA}{2}\right) \end{pmatrix}, \tag{3.25}$$

while in the dilute and thin-sample approximations,

$$\chi_{\text{out}}^{\text{shift}} \approx i\begin{pmatrix} 1 + \frac{1}{2}k_0 z\mu_{11} \\ \frac{1}{2}k_0 z\mu_{12} \end{pmatrix}. \tag{3.26}$$

Here use has been made of the relations $A\cos\theta = \frac{\mu_{11}-\mu_{22}}{2}$ and $A\sin\theta = |\mu_{12}|$.

The ultimate experimental signal is an intensity at the location of the camera (i.e., at the image plane), but polarization optics placed before the camera can modify the detected signal, allowing one to select which component or combination of components of state (3.22) to detect. We recall that a generic wave plate in the circular-polarization basis is described by the transformation matrix

$$\begin{pmatrix} \cos\frac{\psi}{2} & -i\sin\frac{\psi}{2}e^{-2i\phi} \\ -i\sin\frac{\psi}{2}e^{2i\phi} & \cos\frac{\psi}{2} \end{pmatrix}, \tag{3.27}$$

where, for instance, $\psi = \pi/2$ for a quarter-wave plate and $\psi = \pi$ for a half-wave plate. The angle ϕ determines the angle of the principle axes of the wave plate. Similarly, a linear polarizer, also in the circular basis, is represented by projection on the state

$$\frac{1}{\sqrt{2}}\begin{pmatrix} e^{-i\eta} \\ e^{i\eta} \end{pmatrix}. \tag{3.28}$$

Making use of these elements, we consider a number of special cases of experimental relevance, employing the approximation of equations (3.26) and the linearized phase-contrast signal analogous to equation (3.12).

Case 1: Phase dot present, σ_+ light in, no polarizer before camera.

This is the simplest case, and the results may be read off from equation (3.26). The signal recorded by the camera is, to linear order,

$$S = 1 + k_0 z\mu_{++}. \tag{3.29}$$

For imaging light detuned by $\tilde{\delta}\gamma/2$ from the D1 $F = 1 \to F' = 2$ transition and neglecting the weak $F = 1 \to F' = 1$ transition, the explicit form form μ_{++} is

$$\mu_{++} = \frac{6\pi n}{k_0^3 \tilde{\delta}} \left(\frac{1}{2} |\psi_{+1}|^2 + \frac{1}{4} |\psi_0|^2 + \frac{1}{12} |\psi_{-1}|^2 \right), \tag{3.30}$$

where ψ is the unit-normalized the condensate wave function quantized along the imaging direction $\hat{\mathbf{z}}$. This may also be written in terms of the spin operator F_z as

$$\mu_{++} = \frac{6\pi n}{k_0^3 \tilde{\delta}} \left(\frac{1}{4} + \frac{5}{24} \langle F_z \rangle + \frac{1}{24} \langle F_z^2 \rangle \right). \tag{3.31}$$

Case 2: Phase dot present, σ_+ light in, linear polarizer before camera.

The linearized phase-contrast signal in this case, for a polarizer at angle η, is

$$S = \frac{1}{2} \Big\{ (1 + k_0 z \big(\mu_{++} + \cos 2\eta \, \mathrm{Re}\mu_{+-} + \sin 2\eta \, \mathrm{Im}\mu_{+-} \big) \Big\}.. \tag{3.32}$$

For the transition of the previous case, we can write μ_{+-} explicitly as

$$\mu_{+-} = \frac{6\pi n}{k_0^3 \tilde{\delta}} \left(\frac{1}{12} \psi_{-1}^* \psi_{+1} \right) = \frac{6\pi n}{24 k_0^3 \tilde{\delta}} \langle F_-^2 \rangle, \tag{3.33}$$

where F_- is the spin lowering operator. As a concrete example, one can consider using this imaging method to observe the $m_F = 0$ state along $\hat{\mathbf{x}}$, $\psi = (\frac{1}{\sqrt{2}},\ 0,\ -\frac{1}{\sqrt{2}})^T$. This state is of particular interest, since vortex cores and domain walls in the ferromagnetic spinor condensate may be made up of this spin state when the magnetic field is along this $\hat{\mathbf{x}}$. In this case $\mu_{+-} = -6\pi n / 24 k_0^3 \tilde{\delta}$, so that the phase-contrast signal is

$$S = \frac{1}{2} \left\{ 1 + \frac{6\pi n z}{k_0^2 \tilde{\delta}} \left(\frac{7}{24} - \frac{1}{24} \cos 2\eta \right) \right\}.. \tag{3.34}$$

It is possible, by subtracting images where η differs by $\pi/2$ (i.e., images through orthogonal linear polarizers, such as the ordinary and extraordinary polarizations transmitted by a calcite block) to measure the $\Delta m = 2$ portion of this signal alone. If, on the other hand, the initial state had been $m_F = 1$ along $\hat{\mathbf{x}}$, or $\psi = (\frac{1}{2},\ \frac{1}{\sqrt{2}},\ \frac{1}{2})^T$, then the signal would have been

$$S = \frac{1}{2} \left\{ 1 + \frac{6\pi n z}{k_0^2 \tilde{\delta}} (\frac{13}{48} + \frac{1}{48} \cos 2\eta) \right\}.. \tag{3.35}$$

Comparison of equations (3.34) and (3.35) shows that it should be possible to distinguish $m = 0$ from $m = \pm 1$ regions by this method, particularly if images of both linear polarizations are recorded, but that the contrast remains relatively small, at around 2 : 1.

Case 3: No phase dot, σ_+ light in, linear polarizer before camera.

This case is similar to the preceding one, except that it takes as its point of departure the equation (3.22) rather than (3.26). The linearized signal is

$$S = \frac{1}{2}\left\{1 - k_0 z \operatorname{Im}\left(e^{-2i\eta}\mu_{+-}\right)\right\}. \tag{3.36}$$

If the aim is to see the $m_x = 0$ state, then this appears more favorable, since there is no large background scalar or density signal, but only a term depending on μ_{+-}. The signal (3.36) in the special cases $m_x = 0$ and $m_x = 1$ is

$$S = \frac{1}{2} + \frac{6\pi n z}{48 k_0^2 \tilde{\delta}} \begin{cases} -\sin 2\eta & \text{if } m_z = 0, \\ \frac{1}{2}\sin 2\eta & \text{if } m_z = 1. \end{cases}. \tag{3.37}$$

The contrast is thus not any larger, but the absence of a background is appealing.

Case 4: Phase dot present, $\hat{x}$ light in, $\lambda/4$ and linear polarizer before camera.

This is the same as case 1, since linear light can be decomposed into σ_+ and σ_-, with the exception that circular birefringence can cause the signal of one circular polarization to spill into the other.

The signal in this case for the two possible ports of the circular analyzer is

$$S_\pm = \frac{1}{2}(1 + k_0 z(\mu_{\pm\pm} + \operatorname{Re}\mu_{+-})),, \tag{3.38}$$

where the extra factor of $\frac{1}{2}$ is due to the fact that half of the linearly polarized incoming light is in σ_+ and half in σ_-, and where μ_{++} and μ_{+-} are as given above, while

$$\mu_{--} = \frac{6\pi n}{k_0^3 \tilde{\delta}} \left(\frac{1}{4} - \frac{5}{24}\langle F_z\rangle + \frac{1}{24}\langle F_z^2\rangle\right).. \tag{3.39}$$

Case 5: No phase dot, $\hat{x}$ light in, and linear polarizer at $\pi/4$ before camera.

The measured signal at the camera is easily computed in this case also. It is

$$\begin{aligned} S_\pm &= \frac{1}{4}(2 + k_0 z \operatorname{Re}(\mu_{++} - \mu_{--})) \\ &= \frac{1}{2} + \frac{6\pi n z}{k_0^2 \tilde{\delta}} \left(\frac{5}{24}\langle F_z\rangle\right). \end{aligned} \tag{3.40}$$

As in case 3, which also assumed no phase dot, the linearized signal shown here reveals the magnetization signal $\langle F_z \rangle$ without contamination from any $\Delta m = 2$ signal (e.g., $\langle F_z^2 \rangle$ or $\langle F_+^2 \rangle$). Unlike the simpler case 1, however, only half of the incident light contributes to the signal. If images of both polarizations are recorded, however, then this does not represent a waste of imaging light.

3.5 Imaging Larmor Precession

One of the primary contributions of the present work has been to apply the magnetization-sensitive imaging described above to time-resolved imaging of the Larmor precession of a Bose gas, allowing direct spatially and temporally resolved imaging of its transverse and longitudinal magnetization. An understanding of this imaging technique relies somewhat on the technical details of the camera, which I summarize briefly here. The camera used for these experiments, a Roper Scientific PI-MAX (EEV 515x512FT CCD57), is a 512 × 1024 pixel CCD cooled to −40°C. The CCD works in a standard fashion: during the exposure period, photons are converted to photoelectrons which are stored in an electrostatic "bucket" or potential well at each pixel. The number of photoelectrons in a pixel may be read out by removing the confining voltage and allowing the photoelectrons to flow into an integrating amplifier, whose output feeds a digital-to-analog converter. Because of manufacturing constraints, there is only one such read-out device, located at one corner of the chip (see figure 3.3, so that readout of the entire chip necessitates a two-dimensional bucket-brigade movement of the pixels via a phased change of voltages applied to the pins of the CCD chip. Pixels of the first column (i.e., that nearest the read-out corner) are shifted one by one along the column toward the read-out corner, the value of the pixel at the corner being extracted on each step. After this process has occurred 512 times, or once for each row, the second column is moved over to take the place of the first, the third to take the place of the second, etc., and a second iteration begins. Altogether there are 1024 iterations, since this is the number of columns. The time required to shift an entire column of pixels is approximately $1\,\mu$s, but the time required to read out a column at the lowest noise setting is approximately 5 ms. This time scale results from the fact that the integrating amplifier suppresses random noise at its input by a factor proportional to $1/\sqrt{T_\text{int}}$ where T_int is the integration time; longer integration time produces less noise, but the scaling with time is not very favorable. At the lowest-noise setting, corresponding to a 100 kHz digital-to-analog conversion rate, the total time to read out the contents of the chip is approximately 5 s, and for a typical image, derived from three separate raw images, this $\sim$ 15 s read-out time is a significant portion of the experimental cycle. Because of the large separation in timescales between the shift time and the read-out time, it is very desirable to be able to record information at a rate limited by the smaller shift time. This is accomplished by means the "kinetics" mode of the camera, which allows many a single image to be broken into many frames. Unlike the standard mode of the camera, in which a single trigger initiates read-out of the entire chip, in the kinetics mode each trigger causes the chip to shift all pixels a predetermined number of columns (the frame width) toward the read-out end. If the entire chip with the exception of a single frame is masked off, so that imaging light is blocked from reaching it, then the following rapid sequence is possible: an image is recorded on the unmasked portion by exposing it to a brief pulse of imaging light. At the end of the imaging pulse, an electronic trigger initiates shifting of the pixels of the camera, which for a 40-pixel frame is completed in about 40 μs, at which point the cycle begins again, with the prior images safely stored in the dark. The camera shutter is necessarily open throughout this process, since it requires several milliseconds to open and close.

The advantage of employing kinetics imaging becomes apparent if one considers the problem of trying to observe transverse magnetism in the condensate. The condensate is initially in the magnetically trappable $|m_F = -1\rangle$ state, its spin pointing along the field direction. We prepare a transversely magnetized state by applying a resonant quarter-cycle ($\pi/2$) Rabi pulse. The transversely magnetized state so prepared is no longer an eigenstate of the Zeeman Hamiltonian, but is rather a superposition of eigenstates of different energies, which accrue phase at different rates. The physical manifestation of these time-varying phases is precession of the spin at the Larmor frequency, $\omega_L = \frac{1}{2}\mu_B|B|/\hbar$. Because the experiments under consideration operate in a magnetically unshielded environment, in relatively close proximity to many dozens of pieces of electronic apparatus, the magnetic fields at the location of the condensate fluctuate by a few mG r.m.s, with a noise spectrum that extends to several tens of kHz. As a result, we have found it necessary to operate at a bias field of 55 mG or higher to avoid noise-induced spin flips, corresponding to a Larmor frequency of at least 38 kHz, or a period of 26 μs. Thus the imaging pulse must be short compared to 26 μs in order to avoid averaging the signal to zero. Indeed, if the imaging pulse lasts a duration T, the signal recorded would be (in the approximation that the signal is purely sinusoidal)

$$S = \frac{1}{T}\int_{t-T/2}^{t+T/2} \cos(\omega_L t')\mathrm{d}t' = \frac{\sin \omega_L T/2}{\omega_L T/2}\cos \omega_L t, \tag{3.41}$$

so that in order for the apparent reduction in amplitude to be less than 10%, one must have $T \lesssim 1.58/\omega_L \approx T_L/4$, where $T_L = 2\pi/\omega_L$.

Ideally, in order to resolve the Larmor-precession signal well, one would like to take four or more images over the course of a Larmor period. Given that an image of the condensate is approximately 20 μm wide, which corresponds to about 20 pixels and about 20 μs shift time on the camera, it is clear that this is not possible for magnetic fields of the size being considered. An alternative method is to sample at a frequency near an integer divisor of the Larmor frequency and to observe the low-frequency beat note (to alias the signal, in the language of audio processing). Indeed, if the sample frequency $\omega_S = 2\pi/T_S$ satisfies $\omega_L = m\omega_S + \Delta\omega$ for m an integer, then the signals recorded at times nT_S, for $n = 1, 2, 3...$ will be

$$\sin(\omega_L n T_S) = \sin(\Delta\omega n T_S),$$

so that indeed the apparent frequency is simply the difference frequency $\Delta\omega$. Thus, we can sample at 20 kHz with the bias field at 38 kHz, and the beat frequency of 2 kHz is sufficiently slow to allow upwards of 20 images to be recorded per cycle. In combination with the less precise, but absolute, determination of the bias field magnitude from the resonance frequency of the Rabi pulse (as observed with Stern-Gerlach “time-of-flight” absorption imaging), the frequency of Larmor precession signal allows us to measure in a single shot the instantaneous magnetic field to a fraction of a milligauss. A measurement of magnetic field which is yet much more precise can be envisaged by determining the *phase* of Larmor precession. This possibility will be treated in the following chapter.

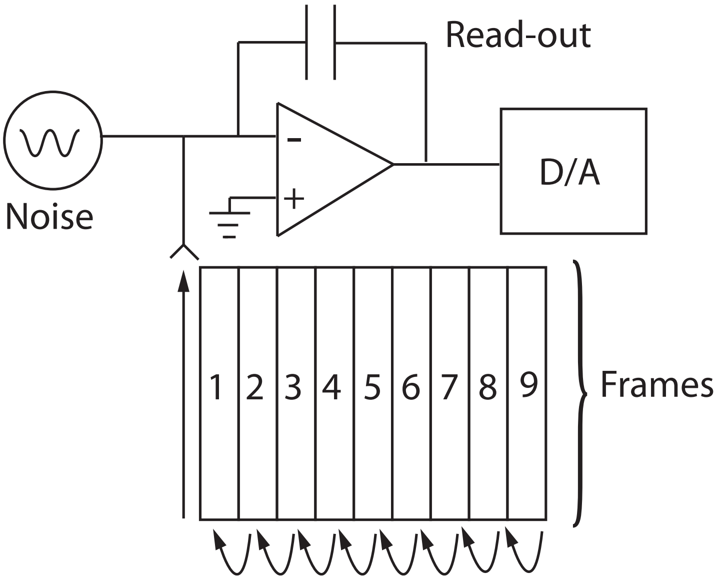

Figure 3.3. Schematic of kinetics imaging. The CCD pixel array is divided into frames of a certain width (usually around 30 pixels). Individual images are recorded on the last numbered frame. Imaging light does not affect the other frames, which are blocked by a mask. After the short pulse of imaging light, a trigger causes the camera to move each frame rapidly to the position of the frame to its left. Frames moving off the left edge of the chip are lost. When the camera has received as many triggers as there are frames, the entire contents of the chip are slowly read out through the amplifier and analog-to-digital converter shown.

3.5.1 Image Processing

The images obtained in kinetics mode consist, as described above, of a number N_{frames} of contiguous frames (usually $25-40$) of predetermined width. Since the condensate is nominally at rest and the temporal spacing of the images ($\sim 50\,\mu$s) is very short compared to motional timescales of at least a few milliseconds, it is straightforward to identify homologous pixels in each frame, i.e., those which represent the same subregion of the condensate. Conceptually, then, for each pixel-sized region of the condensate, one has a time sequence of N_{frames} signals sampled at $\sim 20\,$kHz. For a condensate that is Larmor precessing in the plane of the imaging laser beam, this signal is to good approximation sinusoidal, and may be Fourier transformed or fit to a sine wave to extract an amplitude, phase, offset, and frequency. For many of the experiments described in the present work, the magnetic inhomogeneity across the sample was less than $10\,\mu$G, corresponding to 7 Hz maximum difference in Larmor frequency, so that over the $\sim 1\,$ms image-collection time the accumulated phase difference is less than 50 mrad. In this situation, the Larmor frequency across the cloud can be treated as constant, while leaving the phase variable, even though ultimately the phase of Larmor precession is due to small differences in the Larmor frequency acting over the much longer evolution time. In other cases, magnetic gradients of up to 10 mG/cm have been applied, resulting in Larmor-frequency differences of as much as 210 Hz and phase accumulation over a 1 ms imaging time of 1.3 rad. For such images, it is more proper not to constrain the frequency across the cloud to be constant, although such a constraint could be applied in a subregion of the condensate if one had a priori knowledge about the length scale over which the magnetic field could be varying. The ability to constrain the frequency is of some practical utility, particularly for pixels with low signal-to-noise ratio; for such pixels, it may be possible to extract a phase at a known frequency but difficult to determine this frequency by searching for the maximum in the Fourier spectrum or by curve-fitting. Moreover, computation time is considerably reduced if the frequency does not need to be redetermined at each pixel. The former concern, and to some extent the latter as well, may be reasonably addressed by binning or performing a weighted average (e.g., a Gaussian blur) of neighboring pixels before Fourier transforming; this should result in no loss in information when the binning size is smaller than the imaging resolution.

The result of this temporal Fourier transform or curve fit (both of which methods are commonly employed), is a pixel-by-pixel map of the phase, amplitude, and, optionally, frequency of Larmor precession in the atomic cloud. The offset of the sine wave, i.e., the DC portion of the Fourier transform, is moreover a measurement of the density profile of the condensate. If only axial and not two-dimensional information is desired, the binning procedure may be applied across the full width of the condensate, or alternatively, the transverse profile at each axial point may be fit to an appropriate functional form and the peak value or area under the curve ascribed to that point. For the work described in this sections 3.5.2 and 3.5.3, the latter approach was taken, using the phenomenological functional form $A + B\,\mathrm{sinc}(\kappa(x - x_0))$ to take into account the imperfect imaging resolution in the transverse direction, while for more recent work a full two-dimensional phase and amplitude map have been preferred.

3.5.2 Zeeman Coherence in a BEC

The observation of Larmor precession in a spinor condensate proceeds along the following lines. A longitudinally magnetized condensate is produced in the optical trap by lowering the trap depth and evaporating. A $\pi/2$ pulse rapidly converts the longitudinal magnetization of the condensate to transverse magnetization, which precesses at the Larmor frequency, here equal to $38 \pm 2\,\mathrm{kHz}$. After a variable hold time T_{hold}, a sequence of images is recorded. As the tipped spin Larmor-precesses about the bias field, the phase-contrast image intensities oscillate from bright to dark, accordingly as the spins are aligned or anti-aligned to the direction of propagation of the imaging beam. A composite image consisting of ≈ 40 images equally spaced by $\sim 50\,\mu\mathrm{s}$ (as in figure 3.5) is then uploaded to the computer. In order to account for imaging-beam inhomogeneity and for background light which may enter the imaging system, a second and third image are then recorded. The second is identical to the first, except that the atoms have been expelled from the trap, while in the third image the probe beam is turned off. The normalized, background-subtracted pixel-by-pixel signal is then

$$S(i,j) = \frac{I_1(i,j) - I_3(i,j)}{I_2(i,j) - I_3(i,j)}, \tag{3.42}$$

where $I_n(i,j)$ is the intensity of the (i,j)-th pixel in image n.

It is apparent from figure 3.5 that the images do, as expected, reveal the time-dependent relative phase between each spin population, that is, they show a sinusoidal oscillation due to Larmor precession.

Given that this is a measurement of coherence, it is natural to ask how long the coherence persists. This we ascertain by comparing Larmor precession images taken in two different ways, dubbed the "tip-and-hold" and the "hold-and-tip" measurements. The tip-and-hold measurement is precisely what has been described above, with a variable hold time intercalated between the Rabi pulse and the image pulse. The hold-and-tip measurement is, as its name suggests, the reverse: a variable hold time intervenes between the formation of the condensate and the Rabi pulse, with the imaging immediately following the Rabi pulse. The comparison of these two measurements should distinguish between loss of signal due to number loss or heating (hence diminished density) and signal loss that specifically affects transverse magnetism. The loss rates in the two cases correspond to the inverse T_1 and T_2 times of nuclear magnetic resonance. The signal obtained for each type of measurement is plotted in figure 3.6. It is striking that the decay of the Larmor precession signal in the two cases occurs at approximately the same rate, indicating that no decoherence more significant than number loss afflicts the condensate, in contrast to the case of a thermal cloud (see following section).

For a transversely spin-polarized condensate, one does not expect atomic collisions to be a source of decoherence (see section 4.3.1). The rôle of another prime suspect in decoherence, magnetic inhomogeneity, is more interesting. Figure 4.2 in the following chapter shows an image of a condensate to which has been applied a magnetic-field gradient of 113 mG/cm for 10 ms. The expected response of the system on short time-scales in the presence of a magnetic gradient is the development of a matching phase gradient

$$\frac{\mathrm{d}\phi}{\mathrm{d}z} = -\frac{1}{2}\frac{\mu_B}{\hbar}B' T_{\mathrm{hold}}, \tag{3.43}$$

where $B' \equiv \frac{\mathrm{d}|B|}{\mathrm{d}z}$. At longer times, one might expect, neglecting atomic collisions, that the trapping frequency would begin to enter the description. In the absence of collisions, the effect of the gradient is to displace the harmonic potential $\frac{1}{2}m\omega_z^2 z^2$ along the trap axis by an amount

$$\Delta z = \frac{m_F \mu_B B'}{2m\omega_z^2} \tag{3.44}$$

and to shift the minimum energy of the potential well for particles in state $|m_F\rangle$ by

$$\Delta E = -\frac{(m_F \mu_B B)^2}{8m\omega_z^2}. \tag{3.45}$$

The energy shift is of order $h \times 3\,\mathrm{Hz}$ for typical parameters at a gradient of 14 mG/cm, and the corresponding displacement is a few μm. After the sudden Rabi pulse, the $m_F = \pm 1$ components, which were no longer at the minima of their respective potentials, would begin to oscillate harmonically. Initially the $m_F = \pm 1$ components would accrue momentum relative to the $m_F = 0$ component, which would manifest itself as a winding of the magnetization, but after a quarter cycle the helical magnetization pattern would reach a maximum pitch and begin to unwind; after a full cycle, the spins would have returned to their initial configuration.

This is emphatically *not* what is observed. In fact, the presence of the collisional mean-field interaction dramatically alters the way in which the condensate responds to a magnetic-field inhomogeneity. As seen in figure 3.4(a), there is no evidence of rephasing at the trap period ($\sim$ 200 ms). Instead, the phase gradient merely continues to increase steadily for as long as we are able to measure it. Figure 3.4 confirms that the rate at which this phase gradient increases is simply proportional to the applied gradient. At long times without cancelled gradient, the phase varies rapidly along the axis of the condensate. To see the linear relationship between phase and position, one must "unwrap" the phase, i.e., choose the offset $2\pi n$ for some integer n. This is done by comparing the phase to that at neighboring positions, but for long times at finite gradient, when the spatial phase modulation period approaches the imaging resolution limit, this ambiguity in phase amplifies errors in determining the phase, since incorrect jumps of 2π can throw a given data point off considerably. This is the reason for the large error bars at long times in figure 3.4(a). This problem can be mitigated considerably by binning image pixels over smaller axial regions, as has been done in our more recent work, so that there are more measurements per spatial period.

Another way of stating this result is that the condensate does not appear to respond motionally to the magnetic gradient, but merely to develop a phase according to the local value of the magnetic field. This phenomemon can be understood in terms of the collisional interaction among atoms, which in very crude terms resists their being accelerated by the gradient. Indeed, in the Thomas-Fermi approximation, the mean-field density profile shifts the energy of a particle in the condensate by exactly the amount required to cancel the inhomogeneity of the trapping potential, so that the effective "potential" is flat. Because to good approximation the scattering lengths among all pairs states $|F = 1, m_F\rangle$ are equal, the same cancellation occurs for $m_F = 0$ atoms in an $m_F = 1$ condensate. But an admixture of $m_F = 0$ is simply a small rotation

$$e^{i\theta\mathbf{F}_y}\begin{pmatrix}1\\0\\0\end{pmatrix} \approx \begin{pmatrix}1\\\theta/\sqrt{2}\\0\end{pmatrix} + \mathcal{O}[\theta^2]. \tag{3.46}$$

As a result, excitations in the form of small rotations about the magnetization direction of a ferromagnetic condensate (so-called magnons) experience an effective potential which is flat, rather than harmonic, the consequence being that their period of oscillation is very long and that no rephasing of the magnetization in the presence of a magnetic-field gradient occurs.

More mathematically, this can be seen through a linear expansion of the condensate wave function in the Gross-Pitaevskii equation, as calculated by Ho [83] and Ohmi and Machida [84]. Indeed, the equations of motion for ψ_+ and ψ_0 are

$$\begin{aligned} i\hbar\frac{\mathrm{d}\psi_+}{\mathrm{d}t} &= \left(-\frac{\hbar^2}{2m}\nabla^2 + V\right)\psi_+ + c_0\left(|\psi_+|^2 + |\psi_0|^2 + |\psi_-|^2\right)\psi_+ \\ &\quad + c_2\left(|\psi_+|^2 - |\psi_-|^2\right)\psi_+ + c_2\left(\psi_0^*\psi_+\psi_-^*\psi_0\right)\psi_0 \\ i\hbar\frac{\mathrm{d}\psi_0}{\mathrm{d}t} &= \left(-\frac{\hbar^2}{2m}\nabla^2 + V\right)\psi_0 + c_0\left(|\psi_+|^2 + |\psi_0|^2 + |\psi_-|^2\right)\psi_0 \\ &\quad + c_2\left(\psi_+^*\psi_0 + \psi_0^*\psi_-\right)\psi_+ + c_2\left(\psi_0^*\psi_+ + \psi_-^*\psi_0\right)\psi_-. \end{aligned} \tag{3.47}$$

The zeroth-order solution may be taken to be spin-aligned with all population in $m_F = 1$, so that only ψ_+ is nonzero. The wave function can then be taken as the lowest-energy solution of the time-independent Gross-Pitaevskii equation

$$i\hbar\frac{\mathrm{d}\psi_+}{\mathrm{d}t} = \mu\psi_+ = \left(-\frac{\hbar^2}{2m}\nabla^2 + V\right)\psi_+ + (c_0 + c_2)|\psi_+|^2\psi_+. \tag{3.48}$$

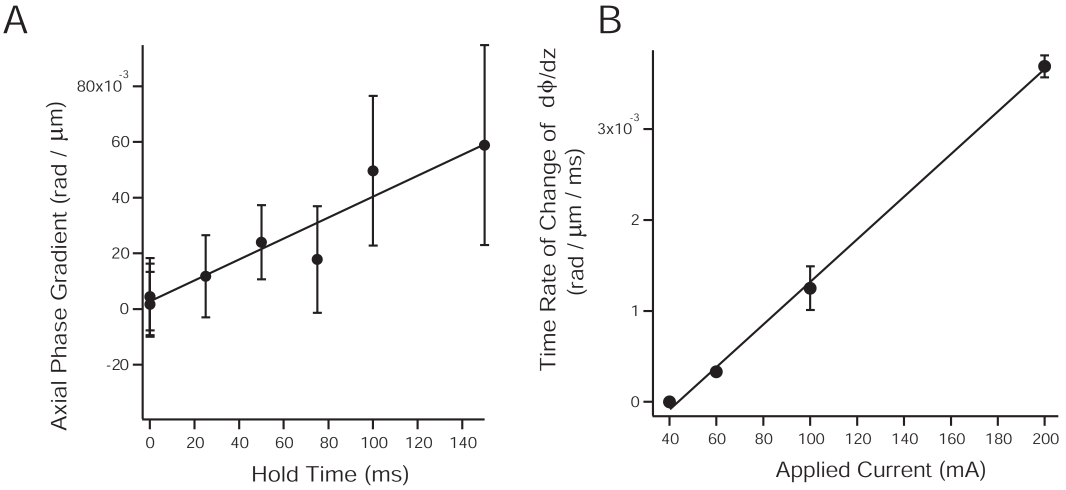

Figure 3.4. (A) Phase gradient vs. time. A current of 60 mA has been applied to the axial gradient coils, and images recorded after a variable hold time. The Larmor phase was obtained as a function of axial position and fitted to a straight line. The slope $\frac{\mathrm{d}\phi}{\mathrm{d}z}$ of this line is plotted against hold time. For this gradient, the spatial period of the phase modulation was approaching the imaging resolution at 200 ms. This is reflected in the large error bars at long times, and which are magnified by global phase ambiguity as described in the text. In (B), the slope obtained from (A) and from similar measurements at 40 mA, 100 mA, and 200 mA is plotted against the applied current, and is well fit by a straight line. At a current of 40 mA, the applied gradient cancelled the residual gradient from other sources, resulting in zero net gradient.

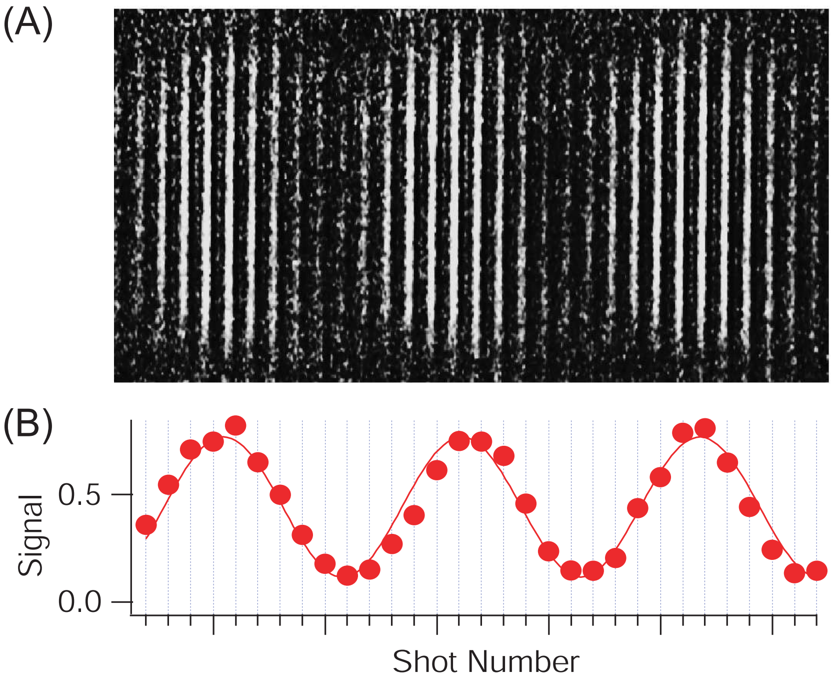

Figure 3.5. Larmor Precession Image

Now on top of this fixed background, consider a first-order excitation in the zero component ψ_0, which as noted above corresponds to local rotations of the spin direction. The linearized equation of motion of ψ_0 is

$$i\hbar\frac{\mathrm{d}\psi_0}{\mathrm{d}t} = \left(-\frac{\hbar^2}{2m}\nabla^2 + V\right)\psi_0 + (c_0 + c_2)|\psi_+|^2\psi_0 \tag{3.49}$$

which using equation (3.48) may be rewritten

$$i\hbar\frac{\mathrm{d}\psi_0}{\mathrm{d}t} = \left(-\frac{\hbar^2}{2m}\nabla^2 + \mu + \frac{\hbar^2}{2m}\frac{\nabla^2\psi_+}{\psi_+}\right)\psi_0. \tag{3.50}$$

This equation shows that the $m_F = 0$ admixture experiences an effective potential which is close to flat, since the term involving derivatives of ψ_+ is of order $\hbar^2/mL^2$ where L is the length of the condensate, except near the edge of the condensate. This energy is of order 1 pK, which is too small to play any part.

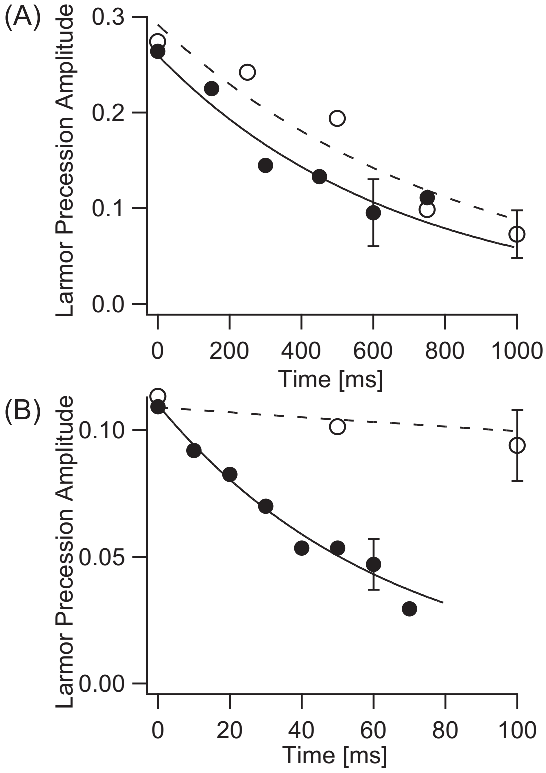

Figure 3.6. Decay of the Larmor precession signal for (a) a Bose-Einstein condensate and (b) a thermal cloud. Data from the "tip-and-hold" (filled circles) or "hold-and-tip" (open circles) methods are compared. For the Bose-Einstein condensate, the measured $1/e$ decay time of the Larmor precession was 670 ± 120 ms, close to the decay time of the "hold-and-tip" signal (830 ± 120 ms), indicating no significant source of decoherence. For the thermal cloud, the 65 ± 10 ms decay time of the Larmor precession signal was an order of magnitude shorter than the decay time of the "hold and tip" signal (1100 ± 500 ms).

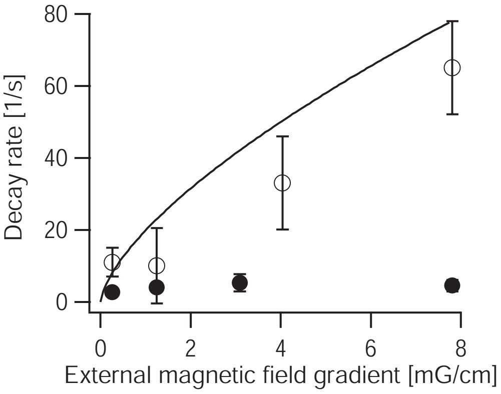

Figure 3.7. Decay of the Larmor precession signal for (a) a Bose-Einstein condensate and (b) a thermal cloud. Data from the "tip and hold" (filled circles) or "hold and tip" (open circles) methods are compared. For the Bose-Einstein condensate, the measured $1/e$ decay time of the Larmor precession was 670 ± 120 ms, close to the decay time of the "hold and tip" signal (830 ± 120 ms), indicating no significant source of decoherence. For the thermal cloud, the 65 ± 10 ms decay time of the Larmor precession signal was an order of magnitude shorter than the decay time of the "hold and kick" signal (1100 ± 500 ms).

3.5.3 Thermal Bose Gas

In contrast to the condensate case, transverse magnetization in a thermal non-degenerate spinor gas decays much faster than that in a BEC (figure 3.6). Moreover, while the Zeeman coherence in a BEC appears locally unaffected by magnetic inhomogeneity, the decoherence rate in a thermal gas depends strongly on the applied magnetic gradient (figure 3.7).

A simple model of this decoherence is obtained by noting that the collision rate of atoms in the optical trap ($\sim$ 200 Hz) is much higher than their axial trapping frequency (5 Hz). As a result, the motion of thermal atoms is diffusive, rather than ballistic. The axial distance travelled by a diffusing atom in time t is

$$d = l\sqrt{t/\tau_c}, \tag{3.51}$$

where $\tau_c = (2n\sigma_c v_{\rm th})^{-1}$ is the collision time and $l = v_{\rm th}\tau_c$ is the axial mean free path. Here $\sigma_c = 8\pi a^2$ is the collision cross section, n is the number density, and $v_{\rm th} = \sqrt{\frac{2k_BT}{\pi m}}$ is the one-dimensional average speed. The expression for τ_c contains an additional factor of 2 because the mean thermal speed in three dimensions (which determines the collision rate) is twice that in one dimension (which determines the mean axial distance travelled). Moreover, in a time t, atoms which have migrated a distance d will have accrued a relative phase compared to those which have remained stationary equal roughly to

$$\Delta\phi = \frac{1}{4}\frac{\mu_B B' d}{\hbar} t, \tag{3.52}$$

where the additional factor of two is an approximation to the average of the field experienced by the atom over its trajectory $\langle B \rangle \approx \frac{1}{2}\left(B'z + B'(z+d)\right)$. If one estimates the dephasing time as the time required for this mean relative phase to reach π, so that the phases associated with different trajectories begin to average to zero, then, combining equations (3.51) and (3.52) an explicit result is obtained for the dephasing time

$$t_\pi = \left(\frac{8\pi\hbar}{\mu_B B'}\right)^{2/3} \left(\frac{n\sigma_{\mathrm{c}}}{v_{\mathrm{th}}}\right)^{1/3}. \tag{3.53}$$

The decay rate $\Gamma_{\mathrm{LP}} \equiv t_\pi^{-1}$ determined from this simple theory is plotted versus magnetic gradient in figure 3.7 and compared to the experimentally measured decay.

It is worth noting that this simple picture does not take into account spin waves of the sort that were seen in pseudospin-$\frac{1}{2}$ gases [85–87] and some form of which should also exist in the spin-1 thermal Bose gas.

Chapter 4

Prospects for Magnetometry

4.1 Principle of the Measurement

As discussed in chapter 3, by imaging with an off-resonant circularly polarized probe beam, one can spatially resolve magnetization in a spin-1 spinor condensate with high contrast. A stronger dipole matrix element for the spin-up than for the spin-down state induces a larger phase shift for the former, and this phase is imaged by means of the phase-contrast technique. By recording a periodic sequence of such images, the temporal phase of the Larmor precessing transverse magnetization may be spatially mapped. If the Larmor frequency is stable to within one cycle over the interrogation time, it can be meaningfully compared to a local oscillator, and the phase measurement can be interpreted as a long-baseline frequency measurement. In essence, this is a Ramsey measurement, with the difference that the second $\pi/2$ pulse is replaced by a detection method which probes the Larmor phase and transverse magnetization directly.

Measured coherence times for Larmor precession in the condensate are of order 500 ms, and potentially larger in a longer-lived lower-density condensate. In a condensate of ≈ 1 million atoms, this would mean a limit on the measurement of magnetic field from atomic shot noise of approximately 1 nG or 100 fT on a single shot using the entire condensate as a single measurement channel. The rather long cycle time of a typical BEC apparatus would reduce the effectiveness of long-time integrations by a factor of $\sqrt{(T_{\mathrm{meas}} + T_{\mathrm{cycle}})/T_{\mathrm{meas}}}$, where T_{meas} is the measurement time and T_{cycle} the cycle time. but, as mentioned in the introduction, recent all-optical BEC-production techniques have dramatically lowered T_{cycle} to a few seconds [28–30], so that in a dedicated apparatus this may not be a large factor. This ideal value, though it represents a rather small field, is about two orders of magnitude short of the sensitivities of SQUID and atomic magnetometers [88–91]. Spinor-condensate magnetometry offers the offsetting merit, however, of very high spatial resolution ($\sim 5\,\mu$m). Similarly high resolution has been obtained with scanning SQUID microscopy [92, 93], but in contrast to such a scanning technique, a spinor-condensate magnetometer would record simultaneous information from each resolved picture element, which may be important in applications where spatial magnetization correlations at a particular time are desired. In addition, the spinor condensate technique is capable of measuring small spatial variations of the magnetic field on top of a large ($\sim$ 500 mG or more) uniform bias field with undiminished sensitivity, which is not true of the spin-exchange relaxation-free magnetometer of Kominis et al. [89]. It is also worth noting that, SQUID measurements inherently measure flux and must therefore rely on separate measurements of the sensor area to infer field. This area, moreover, may not remain constant over time. In contrast, the proposed spinor condensate magnetometry has the benefit of directly measuring the field in terms of fundamental constants, which is significant for applications in which the accuracy as well as the precision of the measurement is of importance.

4.2 Spatial Scale in Magnetometry

A commonly used figure of merit for a sensitive magnetometer is its noise floor in T/ $\mathrm{Hz}^{1/2}$. Since electromagnetic energy always depends quadratically on the fields, these units are straightforwardly related to a physically sensible spectral power density, e.g., in W/Hz This figure, however, fails to convey any information about the length scale at which the magnetometer operates. For many magnetometers, the intrinsic sensitivity can in principle be made arbitrarily high (i.e., the noise floor made arbitrarily low) by expanding the physical size of the magnetometer. Thus the reported noise floor may contain somewhat arbitrary factors related to the actual size of the magnetometer rather than to the intrinsic sensitivity of the measurement method. It is true that for practical purposes, the constraints on the attainable size of the magnetometer may be as significant as fundamental limitations, but both may be of interest. Of greater significance is the fact that for many applications, one wishes to measure spatial patterns of fields at a particular length scale. In order to take into account the sensitivity as a function of spatial scale, one may alternatively employ a figure of merit which explicitly depends on length, e.g., with units of $\mathrm{fT} \cdot \mathrm{cm}^{3/2}/\mathrm{Hz}^{1/2}$, as was done by Allred et al.[94]. The scaling with length of this figure may be understood by noting that a doubling of the volume of an atomic magnetometer should ideally double the number of atoms participating in the measurement, thereby reducing the statistical fluctuations by $\sqrt{2}$.

As estimated below, a spinor-BEC-based magnetometer ought to be able to achieve better than $4\,\mathrm{pT}$ resolution in a 1μm square pixel, limited by optical shot noise. This pixel represents a volume of $4 \times 10^{-12}\,\mathrm{cm}^3$, where the remaining dimension is the thickness of the condensate along the imaging direction of about $4\,\mu$m, so that the volume-invariant figure of merit becomes $8 \times 10^{-18}\,\mathrm{T} \cdot \mathrm{cm}^{3/2}/\mathrm{Hz}^{1/2}$ as compared to $\sim 10^{-15}\mathrm{T} \cdot \mathrm{cm}^{3/2}/\mathrm{Hz}^{1/2}$ for alkali vapor magnetometers based on nonlinear magneto-optical rotation [95] and the projected value of $10^{-17}\mathrm{T} \cdot \mathrm{cm}^{3/2}/\mathrm{Hz}^{1/2}$ for the spin-exchange relaxation-free magnetometer [94]. An imaging resolution better than a few millimeters, moreover, has yet to be demonstrated for these latter types of magnetometry. Thus for applications requiring reasonably high sensitivity and high spatial resolution, the spinor condensate system appears quite promising.

Because the proposed magnetometry is an imaging technique, it naturally records a two-dimensional map of magnetic fields. Therefore, an even more appropriate figure of merit would reflect the scaling of sensitivity with pixel size, the un-imaged dimension being held constant; this figure of merit would have units of $\mathrm{T} \cdot \mathrm{cm}/\,\mathrm{Hz}^{1/2}$, and is the most appropriate for comparison with scanning SQUID microscopy. Indeed, for typical values reported by Kirtley [90] of $10^{-10}\mathrm{T}/\mathrm{Hz}^{1/2}$ in a 7μm-square pickup, the two-dimensional figure of merit is $7 \times 10^{-14}\mathrm{T} \cdot \mathrm{cm}/\,\mathrm{Hz}^{1/2}$. For comparison, the same figure of merit for a spinor-condensate magnetometer is projected to be $4 \times 10^{-16}\mathrm{T} \cdot \mathrm{cm}/\,\mathrm{Hz}^{1/2}$, some two orders of magnitude better. It is worth noting that the scaling with pixel area compared in these two systems is rather different in

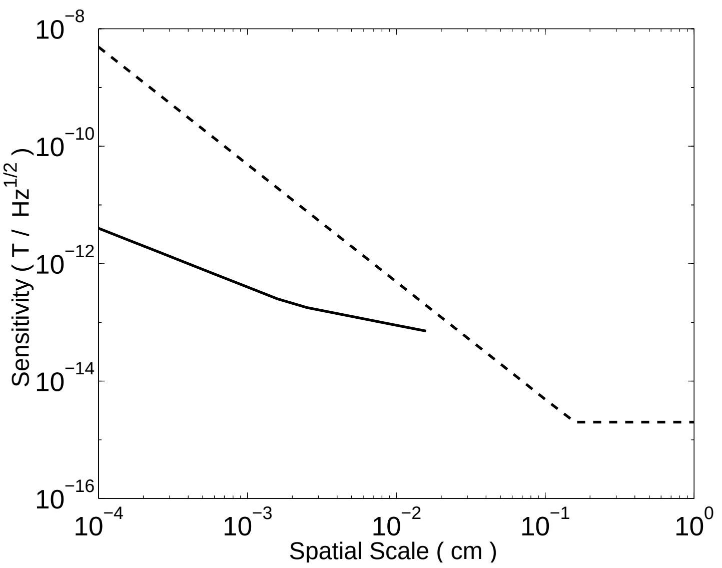

Figure 4.1. Comparison of projected sensitivity of spinor condensate magnetometry (solid line) and scanning SQUID microscopy (dashed line) vs. length scale. The spinor-magnetometry curve is not extended beyond the axial length of the condensate. For length scales smaller than the transverse size of the condensate, $\sim 20\mu$m, the scaling is L^{-1}, while for larger length scales the scaling is $L^{-1/2}$. SQUID sensitivity is limited to $\sim 2\,$fT by, e.g., Dewar magnetic-field noise.

physical origin. For the spinor condensate magnetometer, it results from the inclusion of more atoms in the pixel at larger pixel sizes, which implies a smaller shot noise. For the SQUID, however, the sensitivity is more naturally expressed in terms of the flux, e.g., as $\sim 10^{-6}\Phi_0/\sqrt{\mathrm{Hz}}$, where $\Phi_0 = h/2e$, and the scaling of the sensitivity with area is primarily due simply to the multiplicative conversion between flux and field. As a result the SQUID sensitivity should scale with the inverse area, while the spinor-magnetometry sensitivity should scale with the inverse linear dimension, within the limits set by the size of the sample. Because of this different scaling, the attempt to capture the sensitivity of each method in a single figure of merit is somewhat misleading. More properly one should look at the sensitivity as a function of length scale, as shown in figure 4.1. This figure emphasizes that at length scales shorter than a few hundred μm, the superior scaling of spinor-condensate magnetometry has the potential to render it superior to scanning SQUID microscopy. From figure 4.1 it is clear that the relatively strong dependence of SQUID sensitivity on sensor size is

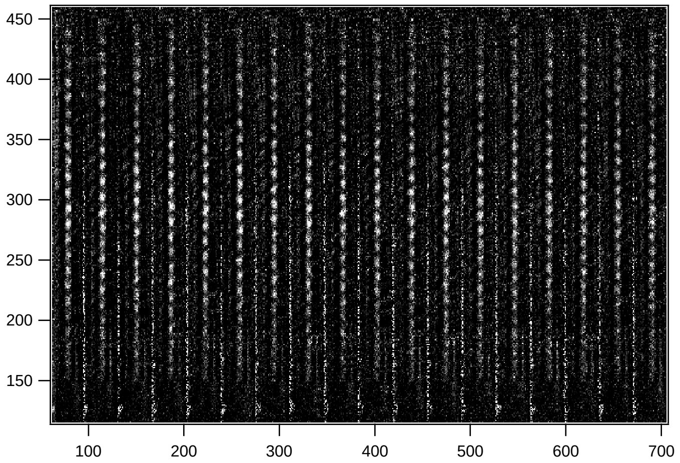

Figure 4.2. Wound-up condensate image. The Larmor-precessing BEC has been exposed to a gradient of 113 mG/cm for 10 ms. The resulting sequence of images spaced by 50 ms is shown. At each point, the signal exhibits Larmor precession across the image, but the phase of precession varies across the condensate, resulting in a strong spatial modulation of image brightness. The spatial period of this modulation is 12.5 μm for the image shown.

partly responsible both for the exceptionally good sensitivity of large-area SQUIDs and their relative weakness as short-distance magnetometers.

4.3 Resolution

The spatial resolution of spinor-condensate magnetometry is limited by the optical resolution of the imaging system. We have calibrated the resolution of our imaging system by imaging condensates whose spins have been tipped transverse to the magnetic field and allowed to evolve for a relatively short time (10 ms) in the presence of a variable magnetic field gradient. The evolution time is chosen so that the largest field gradient of interest (150 mG/cm) acting on the $m_F = \pm 1$ spin components for 10 ms will produce a displacement of not more than 1 μm in the worst-case scenario that the atoms were accelerated as free particles. As discussed in chapter 3, the gradient induces a linearly varying phase with slope

$$\frac{\mathrm{d}\phi}{\mathrm{d}z} = \frac{\mu_B B'(z)}{2\hbar} T_{\text{hold}} \tag{4.1}$$

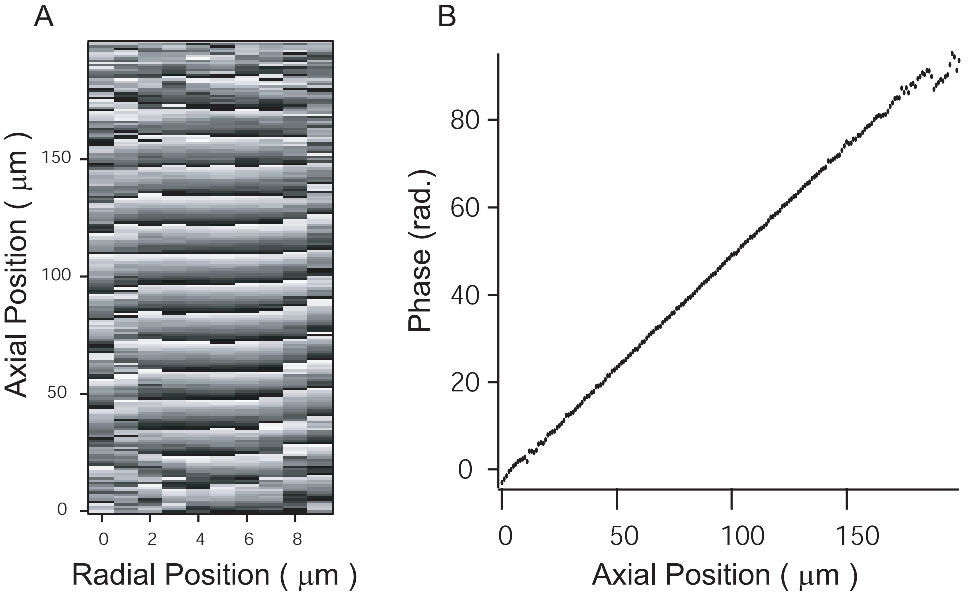

Figure 4.3. (A) Phase map of condensate. The image of figure 4.2 has been Gaussian-blurred with a $1/\mathrm{e}^2$ blur radius of $1.1\,\mu$m and Fourier-transformed pixel by pixel. The phase of the Fourier transform is shown in gray scale, with black representing $-\pi$ and white representing $+\pi$ radians. In (B), the phase along a single-pixel-wide vertical line through the phase map of (A) has been "unwrapped" and plotted against position to show the quality of the data and of the phase extraction. At the upper end, where the signal is weakest, a single phase disambiguation error is evident, while in the central portion residuals from the curve fit are ±200 mrad.

By fitting for the phase as a function of position, we can accurately infer the gradient at any current sufficiently small that the resolution of the imaging system does not affect the measurement. Several gradient-current points are then accurately fit to a straight line, which allows us to extrapolate to larger gradients, even when the imaging resolution affects our ability to determine them accurately from evidence internal to the image. Since we image to good approximation the component of the magnetization along the imaging direction, a linear phase gradient results in a sinusoidal modulation of the brightness of the image along the long axis of the condensate. By measuring the contrast of this modulation as a function of the applied gradient, we obtain a direct measurement of the response of the imaging system as a function of transverse wave number. The results of this measurement are shown in figure 4.4. Because the object being imaged is to good approximation sinusoidal, the interpretation of this measurement is particularly simple: to first approximation, transverse Fourier components of the scattered wave propagating out from the object with wave vectors of magnitude $2\pi\,\mathrm{rad}/6\,\mu\mathrm{m}$ or less are accepted by the imaging system and reconstituted at the image plane, while those whose magnitude exceeds this value are diffracted out of the imaging system altogether. The fact that the cut-off of figure 4.4 is not arbitrarily sharp is presumably attributable to spherical aberration of far off-axis rays. The value of $6\,\mu\mathrm{m}$ is the abscissa of the point at half-maximum contrast. This resolution is close to the diffraction-limited value of $k_\perp = k_\| R_\mathrm{lens}/L\mathrm{lens} = 2\pi/5.5\,\mu\mathrm{m}$ expected for a lens of radius $R_\mathrm{lens} = 2\,\mathrm{cm}$ at a distance of $L_\mathrm{lens} = 14\,\mathrm{cm}$ from the object.

As a result, with the current imaging system, the spatial resolution for magnetometry is approximately $6\,\mu\mathrm{m}$. The difficulty and expense of improving this number increase rapidly with diminishing length scale; a factor of two or three improvement in a dedicated machine with custom-made lenses, however, is certainly feasible. The number of atoms contributing to the measurement, of course, scales with the length, so that without accompanying improvements in atom number, the advantage to resolving ever smaller length scales will rapidly be offset by the increased atom shot noise. Moreover, for real-world applications the minimum length scale of interest will in general be no smaller than the minimum separation between the source of the magnetic field and the measuring device. Since the condensate requires ultra-high

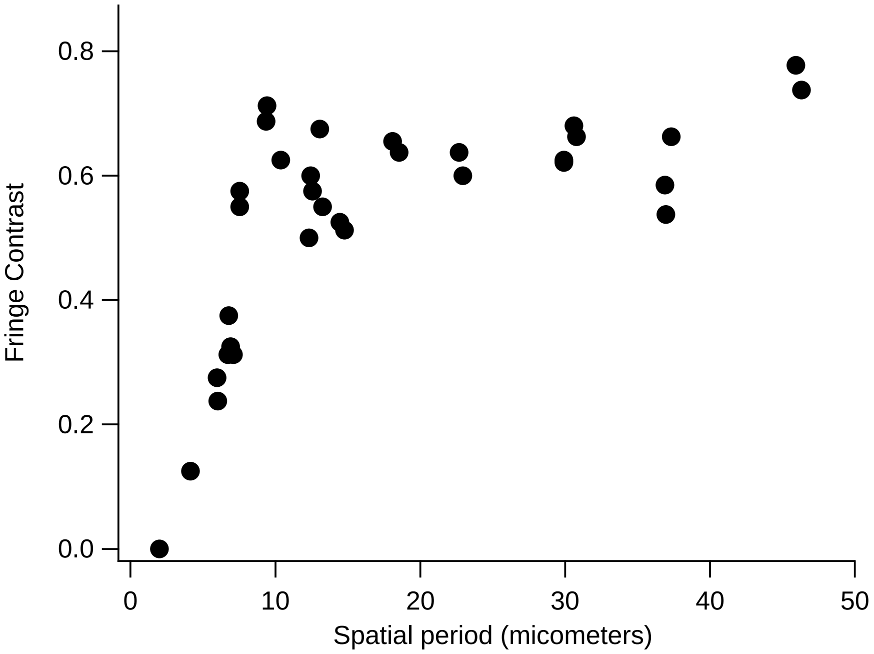

Figure 4.4. Calibration of imaging-system resolution. Condensates were forced to Larmor precess in a variable gradient (0 to 150 μm) for 10 ms, inducing a helical magnetization pattern, which, projected on the imaging direction, resulted in a sinusoidal modulation of the image brightness along the axis of the condensate. The contrast of this modulation is plotted versus its spatial periodicity. The contrast is reduced for high-frequency modulation, dropping to 50% at 6 μm.

vacuum for its survival, while many, if not most, ordinary materials have outgassing rates incompatible with ultra-high vacuum, one likely scenario for practical application of such a magnetometer would involve a thin membrane or window separating the optically trapped condensate from the sample whose magnetic field is to be measured as in [92, 93]. We have performed no detailed study on the feasibility of such a membrane, but a robust vacuum-maintaining membrane much thinner than $2-3\,\mu$m seems prima facie implausible. We note in passing, however, that such a membrane, with the use of an appropriately designed rough-vacuum airlock to contain samples, might not be obliged to withstand more than a small fraction of atmospheric pressure. It is also worth pointing out that the effects of a material surface, dielectric or conducting, on the condensate at tens of μm separation are likely to be deleterious, as observed by Lin et al. [96] and calculated by Rekdal et al.[97].

4.3.1 Absence of Mean-Field Shift

The atoms of a typical alkali-atom BEC are quite dilute ($n \sim 10^{14}\,\mathrm{cm}^{-3}$) by ordinary standards, but because of their large polarizability, the low-temperature cross sections are rather large. For instance, in the case of ^{87}Rb , the S-wave cross section is $\sigma_{\mathrm{coll}} = 2.5 \times 10^4 a_{\mathrm{Bohr}}^2$, much larger than one might naïvely suppose from the size of the atom. As a result, a ^{87}Rb BEC, while sufficiently weakly-interacting to make it theoretically rather tractable, is strongly interacting enough to render its use for many precision measurements difficult. Indeed, any measurement (for instance, the measurement of time in an atomic clock) which relies on the stability of the energy levels and spectral lines of isolated atoms is likely to run into systematic energy/frequency shifts on the order of the interaction energy. This energy is typically a few kHz and is dependent on the atom number or density, which is usually not a well-controlled parameter in these experiments, making these shifts particularly deleterious. The presence of such shifts, along with the small atom number obtainable in most alkali BECs, is a major reason that condensates, whose coherence and velocity spread would otherwise make them near-ideal atomic samples for metrology, have hitherto not been employed to a significant degree in atomic clocks or other precision measurements. A seeming exception is the experiment of Gupta et al.[98], in which the recoil frequency of sodium was measured at the 7ppm level, but in this experiment the condensate had been released from its trap and allowed to expand for 15 ms to lower the atomic density. The use of a freely expanding sample, however, while useful in this instance, means sacrificing the possibility of long interaction times, and so is not desirable for many precision measurements.

In proposing the use of spinor condensates for precise measurement of magnetic fields, then, it is important to emphasize that such frequency shifts are absent on the Larmor frequency, at least to lowest order. To see that this is so, we consider a spinor condensate that is spin-polarized, with its spin tipped transverse to a uniform static magnetic field. In a frame which rotates with the Larmor-precessing spins at the single-atom Larmor frequency, the magnetization of the condensate is clearly quasi-stationary, as would be true for any weakly-interacting ensemble of atoms. For a ferromagnetic condensate like ^{87}Rb , however, a stronger statement can be made; namely, that this spin-polarized state is the many-particle ground state of the system, and is consequently stable.

More formally, consider the second-quantized Hamiltonian for this system, $H = H_{\text{free}} + H_{\text{int}} + H_{\text{B}}$, where H_{free} and H_{int} are given by equations (3.1) and (3.2) respectively and

$$H_{\text{B}} \equiv g_F \mu_B \int \mathrm{d}^3\mathbf{x} \Psi^\dagger(\mathbf{x})\, \mathbf{B} \cdot \mathbf{F}\, \Psi(\mathbf{x})$$

Typical spin-dependent energies are on the order of $h \times 1\,\text{Hz}$, while for for the work described here $g_F \mu_B |B|/h$ is between 40 kHz and 100 kHz. Thus it is desirable to extract the rapidly varying evolution at the Larmor frequency by going to the interaction picture with respect to the Hamiltonian H_{B}, which is equivalent to going to the co-rotating frame. In the interaction picture, the field operator Ψ_I obeys the equation of motion

$$i\hbar \frac{\partial \Psi_I}{\partial t} = [\Psi_I, \tilde{H}^I_{\text{free}} + \tilde{H}^I_{\text{int}}] \tag{4.2}$$

where $\tilde{H}_{\text{free}}$ is simply H_{free} evaluated at the interaction-picture fields Ψ_I and $\Psi_I^\dagger$ and likewise for $\tilde{H}_{\text{int}}$, or in other words

$$\begin{aligned} \tilde{H}_{\text{free}} &= e^{-iH_{\text{B}}t} H_{\text{free}} e^{-iH_{\text{B}}t} = H_{\text{free}}[\Psi_I; \Psi_I^\dagger] \\ \tilde{H}_{\text{int}} &= e^{-iH_{\text{B}}t} H_{\text{int}} e^{-iH_{\text{B}}t} = H_{\text{int}}[\Psi_I; \Psi_I^\dagger] \end{aligned} \tag{4.3}$$

Note that the magnetic field Hamiltonian H_{B} no longer appears in equation (4.2). If we choose $\hat{\mathbf{z}}$ as the magnetic field direction, then in the interaction picture, the total spin projection along any axis, but in particular, say, the transverse $\hat{\mathbf{x}}$ direction is a conserved quantity. Indeed this total spin projection

$$\mathbf{F}_x^{\mathrm{tot}} \equiv \int \mathrm{d}^3\mathbf{x}\Psi^\dagger(\mathbf{x})\mathbf{F}_x\Psi(\mathbf{x}) \tag{4.4}$$

commutes with the Hamiltonian. This is clear from the consideration that the only allowed spin-changing collisions are $0 + 0 \leftrightarrow 1 + (-1)$. As a result, the state in which all atoms have the maximum spin projection along $\hat{\mathbf{x}}$ is a stationary state in the rotating frame, which is to say that the spins do not deviate from the single-particle Larmor frequency.

This result has not hitherto depended on the fact that ^{87}Rb is ferromagnetic, and indeed it is true that the fully magnetized state in a polar condensate is also a stationary state, although not the ground state. In the presence of relaxation mechanisms (e.g., through spin exchange with the thermal fraction), however, it is valuable to make use the true ground state of the system, which should offer relative immunity from thermally driven relaxation. The statement that the Larmor frequency is density independent may also be rephrased as stating that for the case of uniform transverse magnetization there is no phase diffusion of the type discussed by Javanainen et al.[99], in which different terms in the expansion of a product state evolve differently, leading to a spread in relative phase. That this does not occur in the spinor condensate is a straightforward consequence of rotation-invariance.

The fact that fully transversely magnetized state suffers no mean-field shifts is certainly promising, but if the promise of high magnetometric resolution is to be realized, the behavior of the system in the case of nonuniform magnetization is also of interest. An analysis of excitations around the uniform magnetization case was presented in chapter 3, section 3.5.2. This analysis is all within the S-wave approximation for interatomic collisions, and so neglects the possibility of dipolar relaxation. The effects of magnetic-dipole interactions are discussed in section 4.4.4.

4.3.2 Dynamic Range

An important parameter separate from the sensitivity of a magnetometer is its dynamic range, which specifies the ratio of the largest to the smallest measurable signals. For spatially resolved magnetometry, moreover, the dynamic range is a function of the length-scale over which the magnetic fields are to be measured, i.e., the pixel size of the magnetic-field map. The smallest measurable signal is determined simply by the sensitivity, that is, by the ultimate noise level of the device. The maximum signal that may be measured is subject to several potential limitations. For fields varying over short length scales, the maximum field that may be measured is determined again by the optical resolution, since a large spatially varying field will begin to induce unresolvable phase variations in the atomic magnetization. This limit is given by

$$B_{\mathrm{max}}(\lambda) = \frac{2\hbar\lambda}{\lambda_{\mathrm{res}}\mu_B T_{\mathrm{hold}}} \tag{4.5}$$

where λ is the wavelength of the spatial variation of the field and $\lambda_{\mathrm{res}} = 6\,\mu\mathrm{m}$. By making T_{hold} as short as possible, this quantity may be made quite large. Indeed, for $T_{\mathrm{hold}} = 2\,\mathrm{ms}$, one obtains

$$\frac{B_{\mathrm{max}}(\lambda)}{\lambda} = 190\mathrm{mG/\,cm}. \tag{4.6}$$

A more interesting limit is enforced by the requirement that the spatially varying magnetic field not cause undesirable spin excitations in the spinor BEC. This is expected to occur if the length scale over which the phase varies by $\sim \pi$ approaches the spin healing length, defined as the length scale at which the spin-dependent mean-field energy equals the kinetic energy:

$$\xi_{\mathrm{spin}} \equiv \frac{\hbar}{\sqrt{2mc_2 n}} \tag{4.7}$$

This limit on the spatial Fourier component of the measured field at wavelength λ is

$$B_{\mathrm{max}}(\lambda) \approx \frac{\lambda}{\mu_B T_{\mathrm{hold}}}\sqrt{2mc_2 n}. \tag{4.8}$$

For typical values of $c_2 n \approx 0.0046\mu \approx 0.55\,\mathrm{nK}$, where μ is the chemical potential, and again taking $T_{\mathrm{hold}} = 2\,\mathrm{ms}$ this yields the limit

$$\frac{B_{\mathrm{max}}(\lambda)}{\lambda} \approx 240\,\mathrm{mG/\,cm}. \tag{4.9}$$

This is less stringent, for these parameters, than equation (4.6)

A separate issue is the maximum uniform field that can be measured, as opposed to the differential or spatially varying fields considered above. For practical purposes, this maximum bias field is currently determined by the ability to produce imaging light pulses that are short compared to a single Larmor cycle. The present method of producing short pulses relies on an RF mixer, operated as a variable attenuator and driven by a Stanford Research Systems pulse generator; the mixer can rapidly turn on and off the RF power going to an acousto-optic modulator, thus acting as a rapid shutter. This method has easily produced rather square pulses down to 500 ns in duration, which, according to equation (3.41), allow one to image Larmor precession at up to 500 kHz. Some degree of improvement on this value is no doubt possible with the present technique, by focusing the probe beam tightly into the end of the AOM nearest the RF transducer, for instance, and much faster pulses are possible with an electro-optic modulator. For the moment, however, there is no strong motivation to push the Larmor frequency to higher values. It is worth pointing out that the duration of the required pulses scales inversely with the bias field, so that the required intensity scales linearly with the bias field. Thus the effective magnetic field imposed by the A.C. Stark shift from the probe beam scales linearly with the bias field as well. The sum of the true bias field and the effective optical field is a net effective field that is rotated by an angle from the true bias field. This angle by which the net effective field is rotated when the pulse is turned on remains constant as the bias field increases, and the amount by which the spin precesses around this field during the pulse also remains constant. Thus, although the effective optical magnetic field becomes quite large at high probe intensity, it acts for a very short time, so that the number of spin flips to be expected from this pulsing effective field remains small. For the estimated effective optical magnetic field of 5 kHz, a bias field of 100 kHz, and a 1 μs pulse duration, the calculated spin-flip probability is approximately 0.05%. Such optically induced spin flips have been looked for experimentally in Stern-Gerach-separated absorption imaging, but have not been observed within the measurement uncertainty of a few thousand atoms.

4.4 Limitations

4.4.1 Optical Shot Noise

The optical shot noise on an image may be estimated simply from equation (3.42) by noting that the signal is the ratio of the probe image and the bright-field image, neglecting dark-field subtraction

$$S = \frac{I_1}{I_2}, \tag{4.10}$$

with $I_{1,2}$ measured in camera photon counts, from which the r.m.s. error on S is estimated as

$$\frac{\Delta S}{S} = \sqrt{\frac{\Delta I_1^2}{I_1^2} + \frac{\Delta I_2^2}{I_2^2}}. \tag{4.11}$$

It is found experimentally that in order to avoid depleting the condensate through superradiant scattering, the number of photon counts on the camera must not exceed a certain value. Typically, one chooses a light level which will deplete no more than 10% of the atoms, which corresponds to about 30 counts/pixel on the camera at low camera gain for a detuning from the D1 $1 \rightarrow 2$ line of 220 MHz. For an average pixel in the condensate image, $S \sim 1.5$ and $I_2 \sim 30$, which leads to an estimate of $\Delta S \approx 0.39$, assuming $\Delta I_{1,2} = \sqrt{I_{1,2}}$. Technical noise is at the level of 1.3 counts/pixel, and is substantially smaller than optical shot noise for these parameters. The noise on the signal may be reduced by $\sqrt{2}$ by eliminating shot noise from the normalization (bright-field) image; this may be accomplished by increasing the intensity of the probe in the normalization image by a somewhat large factor and rescaling the final image by the inverse factor, or equivalently by maintaining a running average of bright-field images from shot to shot on the experiment. This has recently been implemented, and has successfully removed relevant levels of noise from the bright-field image. Noise on the dark-field image has also been eliminated by subtracting an averaged dark-field image, rather than subtracting it pixel by pixel. If stray light is kept to a minimum, so that the dark-field is dominated by readout noise, then this procedure should work well. With these noise-reduction techniques, for the same parameters, fluctuations in the phase-contrast signal are lowered to $\Delta S \approx 0.28$.

Somewhat higher light levels may be used if a longer time interval is allowed to intervene between successive exposures of kinetics frames, so that the motional coherence in the condensate responsible for superradiance is more strongly damped. The superradiance coherence time has been measured to be around 500 μs, but we have observed in earlier work that multi-pulse superradiance decays more rapidly with pulse spacing than would naïvely be predicted by a two-pulse measurement. In practice, we have successfully sampled at a rate as low as 5 kHz. At lower frequencies, chirping of the Larmor frequency due, it is supposed, to stray 60 Hz magnetic fields complicates interpretation of the images.

The error in the phase of a sinusoidal fit can be calculated, in the limit that deviations of data from the fit are small (so that the difference can be linearized in deviations of ϕ from its expected value). It is

$$\Delta\phi = \frac{\Delta y/A}{\sqrt{N_f}\sqrt{\langle\cos^2(kx_i+\phi_0)\rangle}}.$$

If we assume several points per cycle or randomly distributed points, the $\langle\cos^2\rangle = \frac{1}{2}$, so that

$$\Delta\phi = \frac{\Delta y}{A}\sqrt{\frac{2}{N_f}},$$

where A is the amplitude, N_f is the number of frames or data points and Δy is the error in each data point on the sine curve. The same result is obtained for a discrete Fourier transform on N_f points, again in the limit of small deviations of data from the sine curve. Taking $A = 0.5$ and $N_f = 30$ yields $\Delta\phi \approx 0.14$ radians. For a measurement time of 1 s, which has been achieved in a large-volume trap, this would imply a frequency resolution of 14 mHz, which corresponds (at 700 kHz/G) to 40 nG for one pixel. If the entire condensate were used as a single-channel magnetometer, this would imply an optical-shot-noise-limited performance of $40\,\text{nG}/\sqrt{4000} \approx 640\,\text{pG}$ or 64 fT, where 4000 is the approximate number of pixels in the condensate.

4.4.2 Optimum Probe Frequency for Measuring Magnetization

In general, the signal obtainable with phase-contrast imaging of an atomic vapor scales linearly with the noise as the imaging frequency is varied. This is for the reason that the largest allowable imaging intensity is limited by the scattering rate, so that the intensity I scales as the detuning squared, $I \propto \delta^2$. Optical shot noise is proportional to the square root of the intensity, and therefore the fractional shot noise scales inversely with δ. The phase-contrast signal, on the other hand, is proportional to the index of refraction of the atomic medium, and therefore also scales inversely with δ. Consequently the lower noise possible at larger detuning is offset by the lower signal, and the signal-to-noise ratio (SNR) is approximately constant. This argument does not hold, however, for multilevel atoms, however, and so it was desirable to compute the signal-to-noise frequency dependence. The signal here is taken as the Larmor-precession amplitude, which for our condensates is approximately $(S_{\max} - S_{\min})/approx 0.5$ at a detuning δ_2 from the $1 \rightarrow 2$ transition equal to 230 MHz. This signal in general is slightly nonsinusoidal, containing a quadrupolar contribution at twice the Larmor frequency, but for present purposes we will use only the spin-up signal and the spin-down signal (referred to the imaging direction) to define this peak-to-peak signal. Explicitly, these are

$$S_\pm = 3 - 2\cos\phi_\pm - 2\sin\phi_\pm, \tag{4.12}$$

where

$$\begin{aligned} \phi_+ &= \frac{1}{2}\text{O.D.}\left(\frac{1}{2\tilde{\delta}_2}\right) \\ \phi_- &= \frac{1}{2}\text{O.D.}\left(\frac{1}{12\tilde{\delta}_2} + \frac{1}{12\tilde{\delta}_1}\right), \end{aligned} \tag{4.13}$$

O.D. is the resonant optical density, unreduced by Clebsch-Gordan factors, and (δ_1, δ_2) are the detunings from the $1 \rightarrow 1$ and $1 \rightarrow 2$ D1 transitions. The Larmor-precession amplitude is then taken to be approximately $(S_+ - S_-)/2$.

To compute the noise, we first compute the scattering rate out of the $F = 1$ three-component ground state $\xi^{\text{in}}(\theta) = \left(\cos^2\theta, \sqrt{2}\sin\theta\cos\theta, \sin^2\theta\right)^T$, which represents a magnetized ($m = +1$) spin tipped by an angle θ relative to the imaging axis; the final state ξ_{out} in the scattering process is taken to be arbitrary, both in terms of the optical polarization and of atomic Zeeman states (in $F = 1$ and $F = 2$), for a total of $3 \times (3 + 5) = 24$ components. The off-resonant scattering rate is

$$R_{m_{\text{in}}}(\theta) = \frac{\gamma}{2}\frac{I}{I_{\text{sat.}}}\left|\sum_{F',m} \frac{\xi^{\text{out}*}_{F_f,m_f,m_{\text{out}}} C^*_{F',m_f+m_{\text{out}};F_f,m_f} C_{F',m+m_{\text{in}};1,m} \xi^{\text{in}}_m(\theta)}{\tilde{\delta}_{F'}}\right|^2 \tag{4.14}$$

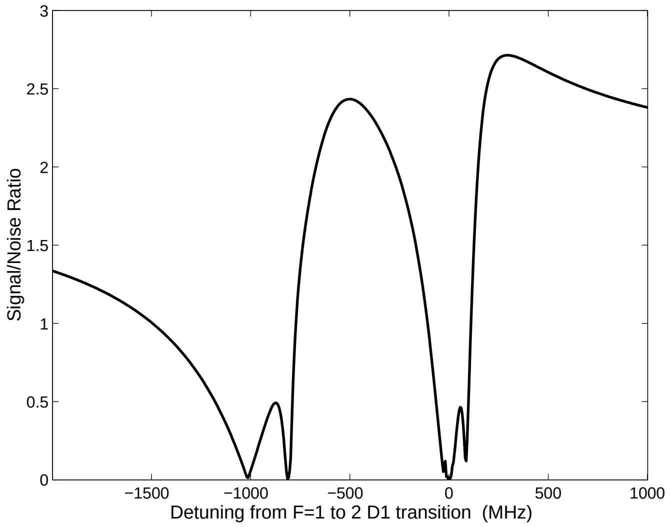

Figure 4.5. Signal-to-noise ratio (SNR) plotted against detuning. The signal is taken to be half the difference between spin-up and spin-down signals, normalized to 1 at the measured value of $\approx -200\,\mathrm{MHz}$. Noise is the sum in quadrature of technical noise (three CCD counts per pixel) and optical shot noise (normalized to 0.31 CCD counts per pixel at $-200\,\mathrm{MHz}$). The precession-averaged scattering rate is held constant, enforcing smaller intensity and higher shot noise nearer the $F' = 1$ line ($-812\,\mathrm{MHz}$) and the $F' = 2$ line ($0\,\mathrm{MHz}$). The magnetization signal vanishes at around $-1000\,\mathrm{MHz}$. A maximum of around 2.4 occurs near $-500\,\mathrm{MHz}$, as compared to the value of ~ 1.7 at $-200\,\mathrm{MHz}$. Another optimum occurs at $\sim +295\,\mathrm{MHz}$, with SNR≈ 2.7.

where m_{in} is the magnetic quantum number of the imaging laser (e.g., $+1$ for σ_+ light), m_{out} is the corresponding quantity for the spontaneously emitted photon, and the C's are composite Clebsch-Gordan coefficients, i.e., dipole matrix elements normalized to the cycling transition so that $C_{F'=3,m'=3;F=2,m=2} = 1$.

Experimentally, it is observed, as noted above, that the limitation on probe power is enforced by superradiant Rayleigh or Raman scattering from the condensate, so that the correct procedure is to compute not the total scattering rate to all final states but the largest possible scattering rate to a particular final state. This process will have the largest gain and will dominate because of mode competition [100]. Equation (4.14) may be written as a quadratic form in ξ^{out},

$$R_{m_{\mathrm{in}}}(\theta) = \frac{\gamma}{2}\frac{I}{I_{\mathrm{sat.}}}\xi^{\mathrm{out}\dagger} M \xi^{\mathrm{out}}, \tag{4.15}$$

where M is a 24×24 matrix, depending on m_{in} and θ, and whose 24 indices label the states $|F_f, m_f, m_{\mathrm{out}}\rangle$ in some definite but arbitrary order. The explicit form of M

is simple to obtain from (4.14), but is notationally difficult and will be omitted from the present discussion. Diagonalizing M at any particular detuning and selecting the largest eigenvalue will determine the maximum scattering rate and therefore the dominant superradiant gain. Averaging this gain with respect to θ is then expected to reproduce approximately the loss experienced by a Larmor-precessing condensate, although it should be emphasized that the general problem of superradiance from a non-stationary spin-state distribution is quite non-trivial will not be considered here. Figure 4.5 displays the signal-to-noise ratio calculated numerically by these means. The number of photon counts per pixel N_c on the camera that are allowed at any given detuning is normalized to the experimentally determined value of 30 (on the camera's low-gain setting) at a detuning of approximately $-200\,\mathrm{MHz}$, and its relative magnitude is then determined at all other frequencies by computing the maximum scattering rate averaged over θ, as described above. The noise is then taken to be

$$\Delta N_c \approx \frac{S_+ + S_-}{2} \frac{1}{N_c} \sqrt{N_c + (\Delta N)^2_{\text{tech.}}},$$

where $\Delta N_{\text{tech.}}$ is the readout noise of the camera, measured at around 3 counts at low camera gain, and $(S_+ + S_-)/2$ is the approximate magnitude of the average phase-contrast signal (see section 4.4.1), i.e., the factor by which the incident intensity is enhanced in the raw image. This corresponds to a value of $\Delta N_c \approx 0.31$ at the fiducial detuning of $-200\,\mathrm{MHz}$. Noise added in the image processing is ignored, since it is possible to avoid most such noise by averaging the dark field and using a high-intensity bright-field image. It is seen in figure 4.5 that there is an optimum in detuning around $-500\,\mathrm{MHz}$ with SNR of 2.4 relative to the $F = 1 \longrightarrow F' = 2$ D1 transition and another near $+295\,\mathrm{MHz}$ with SNR=2.7. This represents a possible improvement relative to $-200\,\mathrm{MHz}$ (SNR ≈ 1.7) of approximately 60%. These results, it is worth pointing out, appear somewhat sensitive to the details of the numerical assumptions that enter into them, although the general structure of figure 4.5 appears robust. In particular, the subject of superradiant gain competition presents many subtleties. There have been indications in our experiments that the general superposition states ξ_{out} are not, in fact, the preferred superradiant states, perhaps because of increased loss due to decoherence among states with differing magnetic moments. Moreover, as mentioned above, a state which has high gain on one probe pulse need hardly present the same high gain to a probe pulse which arrives several Larmor cycles later;

this effect may or may not be adequately taken into account by normalizing to the experimentally determined maximum allowable intensity at $-200\,\text{MHz}$. For all of these reasons, it is desirable to confirm this SNR dependence experimentally.

4.4.3 Atomic Shot Noise

Although the condensate wave function is smooth, in any given experiment the distribution of the finite number of atoms will display a statistical character. Since the number of atoms is much larger than the number of pixels in the condensate image, the distribution of atoms among pixels should be very accurately Poissonian, and in particular the root-mean-square fluctuations of the number per pixel should satisfy $\Delta N = \sqrt{\langle N \rangle}$. For the case where 10^6 atoms are distributed among 4000 pixels, this implies that $\langle N \rangle \approx 250$, $\Delta N \approx 16$, $\frac{\Delta N}{N} \approx 0.06$. For a mean phase-contrast signal of 1.5, the atomic shot-noise limit of $\frac{\Delta S}{S} = \frac{\Delta N}{N}$ yields $\Delta S \approx 0.1$. This value is approximately three times smaller than the corresponding value obtained from the optical shot noise, implying that optical shot noise enforces a more strict limitation on the ultimate magnetic-field resolution than does atomic shot noise. The atomic shot noise, however, is a more fundamental limit, and may prove important if improvements are made to allow higher signal-to-noise imaging. A more precise treatment is given in section 4.6.

4.4.4 Self-Field

Although the primary interactions among atoms are collisional, as described above, it is of course true that a spin-polarized cloud of atoms possesses a macroscopic magnetic moment and produces its own field.

Approximating the condensate as a cylindrical bar magnet consisting of 10^6 atoms each with magnetic moment $-\frac{1}{2}\mu_B$, we readily obtain an estimate of the magnitude of this field.

If the magnetic moments are aligned along the condensate axis, the magnetic field satisfies

$$\mathbf{H} = \frac{1}{\mu_0}\mathbf{B} - \mathbf{M} = 0 \tag{4.16}$$

from which $\mathbf{B} = \mu_0 \mathbf{M}$. For typical radial radii of $10\,\mu$m and axial radius of $150\,\mu$m, the magnetization is approximately $M_z = 10^{-4}$ J/T-m^3, which implies a field of $B_z \approx 1.2\,\mu$G, equivalent to a Larmor frequency of 0.84 Hz.

To lowest order, this dipolar interaction should not perturb Larmor precession, since for a uniformly and transversely magnetized atomic cloud the self-field is also transverse, so that to first order it does not change the Larmor frequency. For measurements of non-uniform magnetic fields, however, this protection will no longer fully apply, and so magnetic-dipole interactions are expected to be an important systematic. Current work on Bose condensates of chromium[101], whose magnetic moment is exceptionally large ($6\mu_B$), may soon help to elucidate the effects of long-range and anisotropic interactions in a quantum-degenerate gas.

4.4.5 Effects of Quadratic Zeeman Shift

The argument of section 4.3.1 for the absence of a mean-field shift to the Larmor precession frequency was strictly valid only for times or fields over which the quadratic Zeeman shift is negligible. Indeed, one can no longer transform away the bias field by going to a rotating frame when the quadratic Zeeman shift is substantial, because the rate of phase accumulation between the $m_F = 1$ and $m_F = 0$ components no longer equals the frequency between the $m_F = 0$ and the $m_F = -1$ components. Thus a rotating frame cannot simultaneously cancel both phases.

In the presence of the quadratic Zeeman shift, denoted q, the transversely magnetized mean-field ground state is modified to

$$\Phi = \begin{pmatrix} \frac{1}{2}\left(1 - \frac{q}{2|c_2|n}\right)^{1/2} \\ \frac{1}{\sqrt{2}}\left(1 + \frac{q}{2|c_2|n}\right)^{1/2} \\ \frac{1}{2}\left(1 - \frac{q}{2|c_2|n}\right)^{1/2} \end{pmatrix} \tag{4.17}$$

This is not, however, the state that is prepared when an $m = -1$ condensate is subjected to a quarter-cycle Rabi pulse; the prepared state is, up to a rotational phase, $(\frac{1}{2}, \frac{1}{\sqrt{2}}, \frac{1}{2})^T$. Consequently, in the rotating frame, the state that has been prepared is not stationary.

The stability of this non-stationary state can be investigated by means of the linearized Gross-Pitaevskii equation (GPE), as in section 3.5.2. Here we write the condensate wave function as the sum of Φ, which is stationary in the rotating frame, and a second function χ, which represents deviations from the stationary state. We note that the state (4.17) can be expanded for small q as

$$\Phi \approx \begin{pmatrix} \frac{1}{2} \\ \frac{1}{\sqrt{2}} \\ \frac{1}{2} \end{pmatrix} - \frac{q}{4|c_2|n} \begin{pmatrix} \frac{1}{2} \\ -\frac{1}{\sqrt{2}} \\ \frac{1}{2} \end{pmatrix} \tag{4.18}$$

and that the two column vectors in the above equation are the $m = +1$ and $m = -1$ eigenvectors of the spin operator F_x. As a consequence, the state prepared by the Rabi pulse is a small admixture of $m = -1$ into the transverse mean field state (4.17). This can be analyzed, following Ho and Ohmi et al.[83, 84], by retaining terms of linear order in χ in the GPE. Since Φ is stationary, it satisfies

$$\begin{aligned} \mu\Phi = &\left(-\frac{\hbar^2}{2m}\nabla^2 + V\right)\Phi \\ &+ c_0\left(\Phi^\dagger\Phi\right) + c_2\left(\Phi^\dagger\mathbf{F}\Phi\right)\cdot\mathbf{F}\Phi + q(F_z^2 - 1)\Phi \end{aligned} \tag{4.19}$$

The time-dependent GPE for $\Phi + \chi$, where $\chi \propto (\frac{1}{2}, -\frac{1}{\sqrt{2}}, \frac{1}{2})^T$, is

$$i\hbar\frac{\mathrm{d}\chi}{\mathrm{d}t} = \left(-\frac{\hbar^2}{2m}\nabla^2 + \sum_j \frac{\hbar^2}{2m}\Phi_j^{-1}\nabla^2\Phi_j + \mu\right)\chi, \tag{4.20}$$

where considerable simplification has resulted from substituting for $\Phi^\dagger\Phi$ using (4.19) and from exploiting the relations $\Phi^\dagger\chi = 0$ and $\Phi^\dagger\mathbf{F}\chi = 0$. This shows that the excitation present from applying a $\pi/2$ pulse at a nonzero quadratic Zeeman shift is not unstable, but on the contrary behaves as a particle in a box, to the extent that the derivatives of Φ are negligible.

4.5 Test via A.C. Stark Shift

As noted earlier, achieving high magnetic-field sensitivity on the present BEC apparatus is precluded by the impracticability of magnetic shielding. A common solution when measuring fields in a noisy environment is to employ a gradiometer, that is, to use one portion of the field sensor as a co-magnetometer for another. Thus, if one could apply a local magnetic field to a sub-portion of the condensate, one could measure the field difference between that portion and the rest of the condensate with high accuracy, even if the fluctuating bias field were ill-determined. Unfortunately for present purposes, however, the condensate is several centimeters from the nearest point at which such a source of local field could be placed, so that it is impossible without breaking vacuum to apply a magnetic field on the $10\,\mu$m length scale. Fortunately, however, it is possible to apply an effective magnetic field optically using the spin-dependence of the A.C. Stark shift.

Neglecting hyperfine structure, the A.C. Stark shift for σ_+ light acting on the $|J=\frac{1}{2}, m_J=\frac{1}{2}\rangle$ and $|J=\frac{1}{2}, m_J=-\frac{1}{2}\rangle$ states is

$$\begin{aligned}\Delta E_{+\frac{1}{2}} &= \frac{\hbar\gamma^2}{8}\frac{I}{I_{\rm sat}}\left(\frac{1}{\delta_{D2}}\right)\\ \Delta E_{-\frac{1}{2}} &= \frac{\hbar\gamma^2}{8}\frac{I}{I_{\rm sat}}\left(\frac{2}{3\delta_{D1}}+\frac{1}{3\delta_{D2}}\right).\end{aligned} \tag{4.21}$$

In order to mimic the effect of a magnetic field, one would like the $m_J = \pm 1$ states to be shifted by equal and opposite amounts, so that there is no net shift. This is the opposite of the requirement on an optical trap for spinor condensates, where it is desired to have confinement independent of Zeeman sublevel. The net shift $(\Delta E_{+\frac{1}{2}} + \Delta E_{-\frac{1}{2}})/2$ is zero when $2\delta_{\rm D1} + \delta_{\rm D2} = 0$, or at the wavelength $\lambda = \frac{3\lambda_{\rm D2}\lambda_{\rm D1}}{2\lambda_{\rm D2}+\lambda_{\rm D1}}$. For wavelengths in air of $\lambda_{\rm D2} = 780.0\,$nm and $\lambda_{\rm D1} = 794.8\,$nm, this implies that cancellation of the net shift occurs at the wavelength 789.8 nm.

The effective magnetic field is then given by the relation

$$\mu_B B_{\rm eff} = \frac{1}{2}(\Delta E_{+\frac{1}{2}} - \Delta E_{-\frac{1}{2}}) = \frac{\hbar\gamma^2}{8}\frac{I}{I_{\rm sat}}\left(\frac{1}{\delta_{D2}}\right). \tag{4.22}$$

At the cancellation wavelength, $\delta_{\rm D2} = -3.00 \times 10^{13}$rad/s and $\gamma = 2\pi \times (6\,\text{MHz})$ so that $\mu_B B_{\rm eff}/h = -0.94\,\text{Hz}\,(I/I_{\rm sat})$. If the A.C. Stark laser is a gaussian beam of $1/e^2$ intensity radii x_0 and y_0, then its intensity is $2P/\pi x_0 y_0$, so that the effective field may be written even more explicitly, using $I_{\rm sat} = 1.67\,$mW, as

$$\mu_B B_{\rm eff}/h = 358\,\text{kHz}\left(\frac{P}{\text{mW}}\right)\left(\frac{10\,\mu\text{m}}{x_0}\right)\left(\frac{10\,\mu\text{m}}{y_0}\right) \tag{4.23}$$

If, as is the case on the current experiment, the A.C. Stark beam is incident at an angle to the true magnetic bias field, then of course component of the effective field that adds linearly (as opposed to in quadrature) to the bias field is the field of equation (4.23) times the cosine of the angle between bias field and laser beam.

Preliminary experiments have been performed, and have seen a local Gaussian-shaped shift in the phase of Larmor precession. At long times, however, the background phase of Larmor precession, i.e., the phase distribution in the absence of the A.C. Stark beam, is rather non-uniform and fluctuates considerably. We have accordingly implemented a spin-echo scheme, which has succeeded in reconstituting clouds with a deviation in phase of not greater than 100 mrad at the echo time, for total evolution times of 500 ms. A remaining difficulty is the motion of the condensate in the trap, which is difficult to damp out and which causes phase inhomogeneity that cannot be removed by the spin-echo technique.

4.6 Possibility of Spin Squeezing

Although measurements of magnetization have not to date been limited by atomic shot-noise, forthcoming reductions in optical shot noise may make the more fundamental atomic shot noise limit accessible. The phase-contrast signal depends both on the total number of atoms in each pixel and on the net magnetization of the atomic sample along the imaging direction, so that it is important to distinguish between shot noise in these two quantities. For concreteness, let us consider an ensemble of atoms in some small region (e.g., a pixel) that are Larmor precessing in the plane of the probe beam. Such a state may be written as

$$\frac{1}{\sqrt{N!}}\big(\cos^2(\theta/2)a_1^\dagger + \sqrt{2}\cos(\theta/2)\sin(\theta/2)a_0^\dagger + \sin^2(\theta/2)a_{-1}^\dagger\big)^N|0\rangle, \tag{4.24}$$

where $a_m^\dagger$ is the creation operator for state m and quantization is along the imaging axis. From this state, by the trinomial expansion, the probability of measuring $\{l, m, n\}$ atoms in states $\{1, 0, -1\}$ is

$$P_{l,m,n} = \delta_{N,\, l+m+n}\frac{N!2^m}{l!m!n!}\left(\cos\theta/2\right)^{4l+2m}\left(\sin\theta/2\right)^{4n+2m}, \tag{4.25}$$

where δ is the Kronecker delta function. The mean value of the dimensionless magnetization $\langle M_z\rangle \equiv N\langle F_z\rangle$ is $N\cos\theta$, and the r.m.s. deviation of the measured component of the magnetization is readily computed to be

$$\Delta M_z = \sqrt{\frac{N}{2}} \sin\theta. \tag{4.26}$$

Equation (4.26) exhibits the shot noise in the magnetization signal, assuming the number of atoms N to be constant. It is noted that this result may also be derived, although perhaps with less physical transparency, in the collective pseudospin formalism and the Heisenberg picture. The shot-noise fluctuation in the total number of atoms in this region, on the other hand, is simply $\sqrt{N}$, and is always present, regardless of the phase of Larmor precession. The phase-contrast signal is given by

$$S = 1 + \alpha N + \beta M_z,$$

where α and β are constants, as seen in section 3.4.2. As a result, the r.m.s. shot noise in the phase-contrast signal in this situation is

$$\begin{aligned} \frac{\Delta S}{S} &= \frac{\sqrt{(\alpha + \beta\cos\theta)^2 \Delta N^2 + \beta^2 \Delta M_z^2}}{1 + \alpha\langle N\rangle + \beta\langle M_z\rangle} \\ &= \sqrt{N}\frac{\sqrt{(\alpha + \beta\cos\theta)^2 + \frac{1}{2}\beta^2(\sin^2\theta)}}{1 + \alpha N + \beta N \cos\theta}. \end{aligned} \tag{4.27}$$

Typical values for a single pixel are $N \approx 250$, $\alpha N \approx 0.5$, and $\beta N \approx 0.5$. For these values, one obtains a maximal value $\Delta S \approx 0.03$, which is within a small factor of current optical shot-noise levels. It is worth noting, however, that atomic shot noise, unlike optical shot noise, is filtered out by the finite imaging resolution, so that properly one should consider the shot noise on the minimum resolvable pixel, rather than on a single physical pixel. Since the resolution is approximately 6 μm, this increases the number of atoms under consideration by a factor of ~ 36, and therefore decreases the resolvable atomic shot noise by a factor of ~ 6. Optical shot noise should be similarly reduced when averaged over the larger region, although imaging at the shot-noise limit becomes more challenging at larger length scales because of technical noise.

Could such technical limitations be overcome, however, and could the density-dependent atomic shot noise be ignored, a measurement taken at $\theta = \pi/2$, in the notation of the preceding paragraphs, would project the state of the atoms in a given pixel into a spinor Fock state of the form

$$(l!m!n!)^{-1/2} \left(a_{-1}^\dagger\right)^l \left(a_0^\dagger\right)^m \left(a_{+1}^\dagger\right)^n |0\rangle,$$

where the populations $\{l, m, n\}$ are distributed according to equation (4.25). This is a state of definite $\hat{\mathbf{z}}$-magnetization, oriented approximately along $\hat{\mathbf{x}}$. As a consequence, measuring the magnetization at this point would collapse the phase of Larmor precession to a definite value. Probing the $\hat{\mathbf{z}}$-magnetization again close to an integer number of Larmor cycles later would then near-deterministically yield the same value as that obtained on the first measurement. Such a phase-squeezed state would then be ideal for measuring small accumulations of Larmor phase, since scatter in the measurement due to number-phase uncertainty would be greatly reduced. Unfortunately, however, the phase-contrast imaging method employed does not isolate shot noise due to the distribution of atoms among spin states from shot noise due to random distribution of atoms among pixels. Consequently the projective measurement of magnetization described above would not in fact be realized. In principle, this can be circumvented by taking a prior measurement of the distribution of atoms among pixels when the spins are oriented along the imaging axis ($\theta = 0$), at which time there is no magnetization shot noise, but only atom-number noise due to the stochastic distribution of atoms among pixels. After this measurement, the number of atoms in each pixel would be a known quantity, at least on time scales over which atoms do not move appreciably. Fluctuations in signal due to magnetization noise could then be measured at each pixel relative to this baseline population map of the pixels in the condensate.

Chapter 5

Theory of the Periodically Dressed Condensate

This chapter will discuss work on periodically dressed BECs and a possible application to metrology; this work was presented in the publications:

- *J. Higbie and D. M. Stamper-Kurn "Periodically Dressed Bose-Einstein Condensate: A Superfluid with an Anisotropic and Variable Critical Velocity." Phys. Rev. Lett.* **88***, 090401 (2002)*
- *J. Higbie and D. M. Stamper-Kurn "Generating macroscopic-quantum-superposition states in momentum and internal-state space from Bose-Einstein condensates with repulsive interactions." Phys. Rev. A* **69***, 053605 (2004)*

Contrary to the case of spinor condensates, where the coupling among spin states is given by the fundamental properties of the particular isotope employed, one can also consider systems where the interactions or other properties of the system are manipulable by the experimenter. Considerable interest has centered in recent years, for instance, on the use of Feshbach resonances to control atom-atom interactions, particularly in ultracold fermions [102–104]. A different way of manipulating the properties of the atomic ensemble is to modify the properties of individual atoms via the light-atom interaction. A Bose condensate formed of the modified atoms should then exhibit novel characteristics, effectively broadening the class of systems in which Bose condensation may be studied.

5.1 Single-Atom Theory

The particular system in which we would like to implement this program is an atom with two stable ground states, $|a\rangle$ and $|b\rangle$, and an excited state $|e\rangle$. The atom is illuminated by lasers of frequencies ω_1 and ω_2 and wave vectors $\mathbf{k_1}$ and $\mathbf{k_2}$ (see figure 5.1), with polarizations chosen so that the beams connect only states $|a\rangle$ to $|e\rangle$ and $|b\rangle$ and $|e\rangle$ respectively. The energy difference between the ground states is $\hbar\omega_0 \equiv \hbar(\omega_b - \omega_a)$. We define likewise the common detuning $\Delta \equiv \omega_1 - \omega_{ea}$, where ω_{ae} is the resonant frequency of the $|a\rangle$ to $|e\rangle$ transition, and the Raman detuning $\delta \equiv \omega_1 - \omega - 2 - \omega_0$. The detuning Δ is assumed to be sufficiently large that spontaneous emission from the state $|e\rangle$ is negligible. Since the two-photon Rabi frequency scales as $1/\Delta$ (see below), while the spontaneous-emission rate scales as $1/\Delta^2$, the relative importance of spontaneous emission may be made almost arbitrarily small.

The states of this system may be enumerated by specifying the internal state of the atom, its momentum state, and the number of photons in each laser beam, N_{ω_1} and N_{ω_2}. The states, and their energies, are:

$$\begin{array}{lc} \underline{\text{State}} & \underline{\text{Energy}} \\ |a, \mathbf{q} - \mathbf{k}/2; N_{\omega_1}, N_{\omega_2}\rangle & 0 \\ |b, \mathbf{q} + \mathbf{k}/2; N_{\omega_1} - 1, N_{\omega_2} + 1\rangle & -\hbar\delta \\ |e, \mathbf{q} + \mathbf{k}_1/2 + \mathbf{k}_2/2; N_{\omega_1} - 1, N_{\omega_2}\rangle & -\hbar\Delta \end{array} , \tag{5.1}$$

where the difference wave vector has been defined at $\mathbf{k} \equiv \mathbf{k}_1 - \mathbf{k}_2$. For simplicity, the number of photons in each beam is taken to be well-defined. This assumption is not essential to the following development, which could be carried through with coherent states as well. N_{ω_1} and N_{ω_2} are, however, required to be much larger than one.

Since the three-level atomic system is closed, and since the numbers of photons in the two laser beams can change only by converting a photon of type 1 to one of type 2 while simultaneously changing the atomic state from $|a\rangle$ to $|b\rangle$, and the three atom-plus-photon states listed above couple only among themselves. The Hamiltonian for this driven-atom system may be written in 3×3 matrix form (in the basis of equation

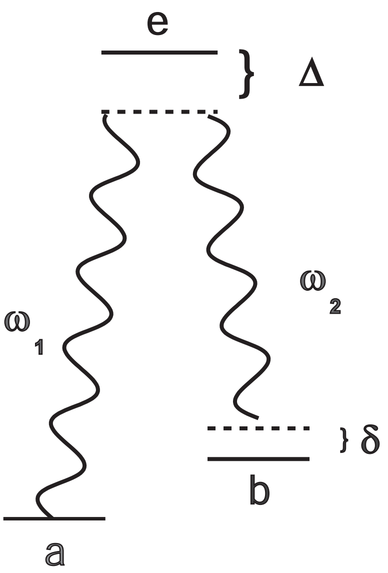

Figure 5.1. Levels in Raman transition. Ground states a and b are connected via lasers of frequencies ω_1 and ω_2 to an excited state e. The excited state detuning is Δ, assumed to be large. The detuning from Raman resonance is δ.

(5.1)) as

$$H = \frac{\hbar}{2}\begin{pmatrix} \frac{\hbar(\mathbf{q}-\mathbf{k}/2)^2}{2m} & 0 & \Omega_{\rm ae} \\ 0 & \frac{\hbar(\mathbf{q}+\mathbf{k}/2)^2}{2m} - 2\delta & \Omega_{\rm be} \\ \Omega^*_{\rm ae} & \Omega^*_{\rm be} & \frac{\hbar(\mathbf{q}+\mathbf{k}_1/2+\mathbf{k}_2/2)^2}{2m} - 2\Delta \end{pmatrix}, \tag{5.2}$$

where $\frac{\hbar\Omega_{\rm ae}}{2} \equiv \langle a|\mathbf{d}\cdot\mathbf{E}|e\rangle$ and $\frac{\hbar\Omega_{\rm be}}{2} \equiv \langle b|\mathbf{d}\cdot\mathbf{E}|e\rangle$ are the single-photon Rabi frequencies for lasers 1 and 2 on the $|a\rangle$ to $|e\rangle$ and $|b\rangle$ and $|e\rangle$ transitions. Here the lasers are treated simply as classical driving fields.

Since the detuning Δ is large enough that no population resides in the level $|e\rangle$, it is possible to eliminate this level from the dynamics and consider an effective Hamiltonian for the two ground states, with an effective coupling between them which encapsulates the relevant three-level physics. This general procedure is described, for instance, in Cohen-Tannoudji citecohe92. It can be iterated to arbitrarily high order, but for present purposes the first order is sufficient.

One defines a Hermitian matrix S, given to first order in Ω/Δ by

$$S \approx \frac{1}{2\Delta}\begin{pmatrix} 0 & 0 & -i\Omega_{\mathrm{ae}} \\ 0 & 0 & -i\Omega_{\mathrm{be}} \\ i\Omega^*_{\mathrm{ae}} & i\Omega^*_{\mathrm{be}} & 0 \end{pmatrix}. \tag{5.3}$$

The unitary transformation $U \equiv e^{iS}$ then breaks the Hamiltonian into block-diagonal form, $H' = UHU^\dagger$, where, to the same order,

$$H' = H_2 \oplus H_1,$$

where

$$H_2 = \begin{pmatrix} \frac{\hbar(\mathbf{q}-\frac{\mathbf{k}}{2})^2}{2m} + \frac{|\Omega_{\mathrm{ae}}|^2}{4\Delta} & \frac{\Omega_{\mathrm{ae}}\Omega^*_{\mathrm{be}}}{4\Delta} \\ \frac{\Omega^*_{\mathrm{ae}}\Omega_{\mathrm{be}}}{4\Delta} & \frac{\hbar(\mathbf{q}+\frac{1}{2}\mathbf{k})^2}{2m} + \frac{|\Omega_{\mathrm{be}}|^2}{4\Delta} - \delta \end{pmatrix} \tag{5.4}$$

and

$$H_1 = \hbar\left(\frac{\hbar(\mathbf{q}+\frac{1}{2}\mathbf{k}_1 - \frac{1}{2}\mathbf{k}_2)^2}{2m} - \Delta + \frac{|\Omega_{\mathrm{be}}|^2 + |\Omega_{\mathrm{ae}}|^2}{4\Delta}\right). \tag{5.5}$$

In other words the 3×3 matrix H' is reduced to a 2×2 matrix H_2 and a 1×1 matrix H_1. Consequently, if the initial state is in the 2×2 manifold corresponding to the ground states, then subsequent evolution will maintain it in this subspace. For convenience, we absorb the A.C. Stark shifts $\frac{|\Omega_{\mathrm{ae}}|^2}{4\Delta}$ and $\frac{|\Omega_{\mathrm{be}}|^2}{4\Delta}$ into the definition of the energies of the two states (or, equivalently, into δ) and define the two-photon Rabi frequency $\Omega \equiv \Omega_{\mathrm{ae}}\Omega^*_{\mathrm{be}}/2\Delta$.

Thus, the effective Hamiltonian for intra-ground-state dynamics becomes

$$\begin{aligned} H_{\mathrm{eff}} &= \frac{\hbar}{2}\begin{pmatrix} \frac{\hbar}{m}(\mathbf{q}-\mathbf{k}/2)^2 + \delta & \Omega \\ \Omega^* & \frac{\hbar}{m}(\mathbf{q}+\mathbf{k}/2)^2 - \delta \end{pmatrix} \\ &= \frac{\hbar^2}{2m}\left(\mathbf{q}^2 + \frac{\mathbf{k}^2}{4}\right) + \left(\frac{\hbar\delta}{2} - \frac{\hbar^2\mathbf{k}\cdot\mathbf{q}}{2m}\right)\sigma_z + \frac{\hbar\Omega}{2}\sigma_x, \end{aligned} \tag{5.6}$$

where σ_i are the standard Pauli matrices. From the latter form, the eigenvalues may be read off as

$$\hbar\omega^{\pm}_{\mathbf{q}} = \frac{\hbar^2}{2m}\left(\mathbf{q}^2 + \frac{\mathbf{k}^2}{4}\right) \pm \frac{\hbar}{2}\sqrt{\Omega^2 + \left(\frac{\hbar\mathbf{q}\cdot\mathbf{k}}{m} - \delta\right)^2}. \tag{5.7}$$

This dispersion relation is plotted for $\Omega = 0$ and $\Omega = 2\pi \times 5\,\mathrm{kHz}$ in figures 5.2 and 5.3, and is seen to have some rather interesting features. When $\delta = 0$,

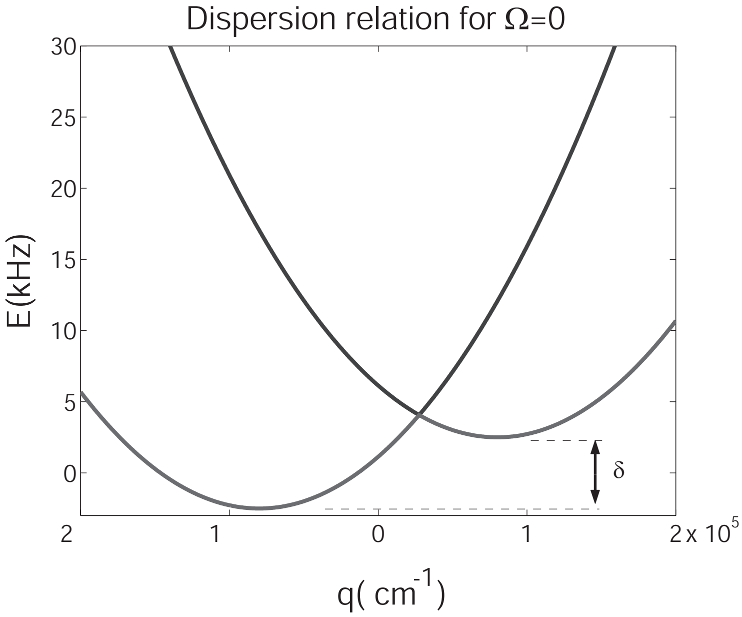

Figure 5.2. Dispersion relation with zero Rabi frequency Ω, consisting of two parabolae offset in momentum by k and in energy by $\hbar\delta$.

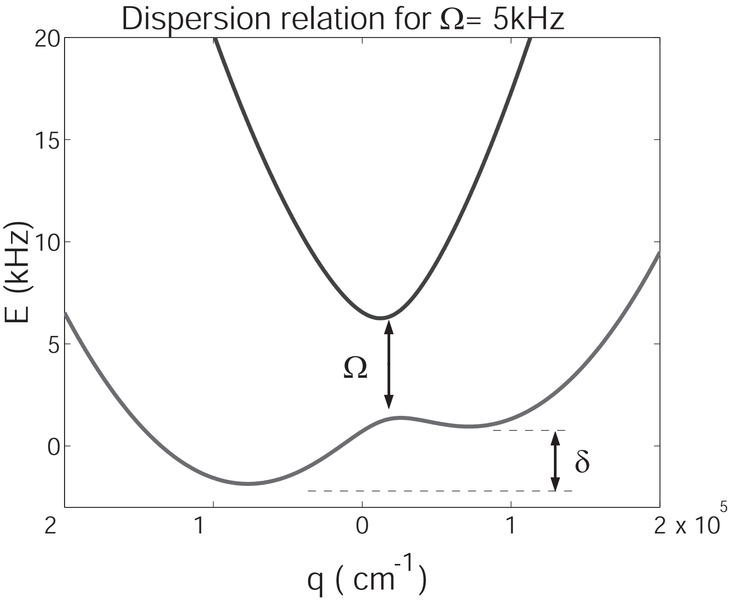

Figure 5.3. Dispersion relation for finite Ω, modified from the case shown in figure 5.2 and displaying an avoided crossing at intersection.

the dispersion relation is symmetric about $q_z = 0$. The upper branch always has a quadratic minimum at $q_z = 0$, and the effective mass (for motion along the z-axis) of the $|+\rangle$ state is

$$m_-^* = \frac{m}{1 + \frac{\hbar k^2}{2\Omega m}}.$$

The lower branch ($|-\rangle$ state) develops a pair of minima at $q_z = \pm\frac{k}{2}\sqrt{1 - \left(\frac{2\Omega m}{\hbar k^2}\right)^2}$ for $\Omega < \frac{\hbar k^2}{2m}$ where $k = |\mathbf{k}_1| = |\mathbf{k}_2|$, giving a reflection symmetry to the ground state. The effective mass of particles in these two symmetric wells is

$$m_-^* = \frac{m}{1 - \frac{4m^2\Omega^2}{k^4}}.$$

If $\Omega > \frac{\hbar k^2}{2m}$, then the lower branch has a single minimum with an effective mass of

$$m_-^* = \frac{m}{1 - \frac{\hbar k^2}{2\Omega m}}.$$

The effect of adding a small (compared to $\frac{2\hbar k^2}{m}$) nonzero δ is to raise one minimum with respect to the other and make it more shallow, i.e., reduce its curvature. For δ significantly larger than $\frac{2\hbar k^2}{m}$ the upper minimum in the lower branch disappears altogether.

The eigenstates of H_{eff}, corresponding to the eigenvalues 5.7, are

$$\begin{aligned} |+\rangle_{\mathbf{q}} &= \cos\frac{\theta_{\mathbf{q}}}{2}\,|a, \mathbf{q} - \mathbf{k}/2\rangle + \sin\frac{\theta_{\mathbf{q}}}{2}\,|b, \mathbf{q} + \mathbf{k}/2\rangle \\ |-\rangle_{\mathbf{q}} &= \cos\frac{\theta_{\mathbf{q}}}{2}\,|b, \mathbf{q} + \mathbf{k}/2\rangle - \sin\frac{\theta_{\mathbf{q}}}{2}\,|a, \mathbf{q} - \mathbf{k}/2\rangle \end{aligned} \tag{5.8}$$

where $\cot\theta_{\mathbf{q}} \equiv \frac{\delta}{\Omega} - \frac{2\hbar\mathbf{k}\cdot\mathbf{q}}{m\Omega}$. The photonic portion of the states has been suppressed for brevity, but the superpositions are of atom-photon states as in equation (5.1). It is also assumed in this expression that Ω is real, which is equivalent to choosing a particular relative phase for the Raman beams, or to redefining the phase of states a and b.

5.2 Bogoliubov Treatment of Many-Body Calculation

The single-particle states and dispersion relation described above set the stage for a many-body treatment of the periodically dressed condensate system. Indeed, one can define creation operators for the states of equation (5.8) as $\pi^\dagger_{\mathbf{q}}|0\rangle = |+\rangle_{\mathbf{q}}$ and $\mu^\dagger_{\mathbf{q}}|0\rangle = |-\rangle_{\mathbf{q}}$, where $|0\rangle$ is the 'vacuum' or no-particle state. Since both the atoms and photons constituting these states are bosonic, the superposition states are also bosonic in character, so that $\pi_{\mathbf{q}}$ and $\mu_{\mathbf{q}}$ satisfy

$$\begin{aligned} \left[\pi_{\mathbf{q}}, \pi^\dagger_{\mathbf{q}'}\right] &= \delta_{\mathbf{q},\mathbf{q}'} \\ \left[\mu_{\mathbf{q}}, \mu^\dagger_{\mathbf{q}'}\right] &= \delta_{\mathbf{q},\mathbf{q}'}, \end{aligned} \tag{5.9}$$

while commutators of $\pi_{\mathbf{q}}$ with $\mu_{\mathbf{q}}$ or $\mu^\dagger_{\mathbf{q}}$ vanish. Written explicitly in terms of creation operators $a^\dagger_{\mathbf{q}}$ and $b^\dagger_{\mathbf{q}}$ for states $|a\rangle$ and $|b\rangle$, the dressed-state creation operators are

$$\begin{aligned} \pi^\dagger_{\mathbf{q}} &= \cos\frac{\theta_{\mathbf{q}}}{2} a^\dagger_{\mathbf{q}-\mathbf{k}/2} + \sin\frac{\theta_{\mathbf{q}}}{2} b^\dagger_{\mathbf{q}+\mathbf{k}/2} \\ \mu^\dagger_{\mathbf{q}} &= \cos\frac{\theta_{\mathbf{q}}}{2} b^\dagger_{\mathbf{q}+\mathbf{k}/2} - \sin\frac{\theta_{\mathbf{q}}}{2} a^\dagger_{\mathbf{q}-\mathbf{k}/2}. \end{aligned} \tag{5.10}$$

In terms of these operators, the two-component position-space field operator may be written

$$\Psi(\mathbf{r}) = \frac{1}{V^{1/2}} \sum_{\mathbf{q}} e^{i\mathbf{q}\cdot\mathbf{r} - i\sigma_y\theta_{\mathbf{q}}/2} \begin{pmatrix} \pi_{\mathbf{q}} \\ \mu_{\mathbf{q}} \end{pmatrix}, \tag{5.11}$$

where V is the quantization volume.

Suppose that we are able to form a Bose condensate in the lowest-energy single-particle state, $|-\rangle_{\mathbf{q}}$. Then, as in the treatment of the weakly-interacting BEC by Bogoliubov [105], the macroscopic occupation of the ground state will modify the effective potential experienced by atoms not in the condensate. Or, in the language of condensed matter physics, the elementary excitations in the presence of the BEC will be different from those in its absence.

The analysis proceeds from the many-body Hamiltonian

$$H = \int \mathrm{d}^3\mathbf{r}\,\Psi^\dagger H_{\text{eff}}\Psi + H_{\text{int}}, \tag{5.12}$$

where

$$H_{\text{int}} = \frac{2\pi\hbar^2 a}{m} \int \mathrm{d}^3\mathbf{r}\left(\Psi^\dagger\cdot\Psi\right)^2, \tag{5.13}$$

with a the S-wave scattering length and H_{eff} is the effective light-atom Hamiltonian of equation (5.6). We have already diagonalized the noninteracting portion, which is no more complicated for N atoms than for one. In second-quantized notation it may be written

$$H_{\text{free}} = \sum_{\mathbf{q}} \left(\hbar\omega^{-}_{\mathbf{q}} \mu^{\dagger}_{\mathbf{q}} \mu_{\mathbf{q}} + \hbar\omega^{+}_{\mathbf{q}} \pi^{\dagger}_{\mathbf{q}} \pi_{\mathbf{q}} \right). \tag{5.14}$$

The interaction Hamiltonian, describing the effects of atomic collisions, is conveniently written in terms of the spatial Fourier transform of the density operator

$$n_{\mathbf{q}} \equiv \int \mathrm{d}^3\mathbf{r} \left(\Psi^{\dagger} \cdot \Psi \right) e^{i\mathbf{q}\cdot\mathbf{r}}, \tag{5.15}$$

resulting in the momentum-space version of the collision term

$$H_{\text{int}} = \frac{2\pi\hbar^2 a}{mV} \sum_{\mathbf{q}} (n_{\mathbf{q}} n_{-\mathbf{q}} - N). \tag{5.16}$$

It is, of course, not necessarily the case that atoms in state $|a\rangle$ and state $|b\rangle$ should collide with the same cross-section, as these expressions assume. For the hyperfine states of ^{87}Rb , however, this is a good approximation, valid to the 1% level. Moreover, the relative simplicity of the expressions under this assumption is justification enough, at least for the purposes of gaining insight into the behavior of the periodically dressed condensate system. We have also neglected collisions which change the internal state of the colliding atoms. This is likewise a good approximation for ^{87}Rb, given experimental observation of long lifetimes of spin mixtures [67].

Substituting the expression (5.11) for Ψ into $n_{\mathbf{q}}$ yields an explicit form for $n_{\mathbf{q}}$ in the dressed-state basis:

$$n_{\mathbf{q}} = \sum_{\mathbf{q}'} (\pi^{\dagger}_{\mathbf{q}'+\mathbf{q}} \quad \mu^{\dagger}_{\mathbf{q}'+\mathbf{q}}) e^{i\sigma_y(\theta_{\mathbf{q}'+\mathbf{q}} - \theta_{\mathbf{q}'})/2} \begin{pmatrix} \pi_{\mathbf{q}'} \\ \mu_{\mathbf{q}'} \end{pmatrix}. \tag{5.17}$$

Following the Bogoliubov approximation, we assume that the lowest-energy single-particle state must be macroscopically occupied with occupation number N_0. Calling the wave vector of this lowest-energy state $\mathbf{Q}$ (i.e., $\mathbf{Q}$ minimizes the dispersion relation (5.7), one may reasonably replace $\mu_{\mathbf{Q}}$ and $\mu^{\dagger}_{\mathbf{Q}}$ by the c-number $\sqrt{N_0}$ and retain only terms of order N_0 or higher in the interaction. At this level of approximation, the momentum-space density becomes

$$n_{\mathbf{q}} = \begin{cases} N_0 & , \quad \mathbf{q} = \mathbf{Q} \\ \sqrt{N_0} \Big(\pi^{\dagger}_{\mathbf{Q}+\mathbf{q}} \sin\Delta_{\mathbf{q}} + \pi_{\mathbf{Q}-\mathbf{q}} \sin\Delta_{-\mathbf{q}} & , \quad \mathbf{q} \neq \mathbf{Q} \\ \quad + \mu^{\dagger}_{\mathbf{Q}+\mathbf{q}} \cos\Delta_{\mathbf{q}} + \mu_{\mathbf{Q}-\mathbf{q}} \cos\Delta_{-\mathbf{q}} \Big) & \end{cases}, \tag{5.18}$$

where $\Delta_{\mathbf{q}} = (\theta_{\mathbf{Q}+\mathbf{q}} - \theta_{\mathbf{Q}})/2$.

The Hamiltonian $H_{\text{int}} + H_{\text{free}}$ can be diagonalized by means of a Bogoliubov-style transformation, generalized to the two-component case. We note that a similar generalization of the Bogoliubov method was previously made by Bassichis [106]. This Hamiltonian is quadratic in the creation and annihilation operators π and μ, and consequently should be exactly diagonalizable. In order to be able to interpret the eigenvalues as an energy spectrum, however, it is necessary that the transformation yield true creation and annihilation operators, i.e., that the commutation relations of these operators be preserved. To reduce the problem of determining the appropriate transformation to a matrix diagonalization, we write the ladder operators as a four-component vector

$$v = \left(\mu^{\dagger}_{\mathbf{Q}+\mathbf{q}}\,,\, \mu_{\mathbf{Q}-\mathbf{q}}\,,\, \pi^{\dagger}_{\mathbf{Q}+\mathbf{q}}\,,\, \pi_{\mathbf{Q}-\mathbf{q}}\right)^{T} \tag{5.19}$$

and define the "metric" matrix, in analogy with the Lorentz group,

$$\Lambda \equiv \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix},$$

in terms of which the ladder-operator commutation relations can be written $[v^{\dagger}_i \Lambda_{ij}, v_k] = \delta_{jk}$. If v transforms as Uv for some matrix U, then $v^{\dagger}\Lambda$ should transform as $v^{\dagger}\Lambda U^{\dagger}$, so that the commutation relations will be preserved if and only if $U^{\dagger}U = 1$, i.e., if U is unitary. Diagonalization of matrices by unitary operators is familiar territory; accordingly, we write the many-body Hamiltonian, discarding a constant offset and keeping terms of order N or higher, as

$$H = \frac{1}{2}\sum_{\mathbf{q}} v^{\dagger}_{\mathbf{q}} \Lambda (H_0 + \mu M) v_{\mathbf{q}}. \tag{5.20}$$

Here, $\mu = 4\pi\hbar^2 N_0/mV$ is the chemical potential, and the matrix H_0 is equal to

$$\hbar \begin{pmatrix} \omega^{-}_{\mathbf{Q}+\mathbf{q}} & 0 & 0 & 0 \\ 0 & -\omega^{-}_{\mathbf{Q}-\mathbf{q}} & 0 & 0 \\ 0 & 0 & \omega^{+}_{\mathbf{Q}+\mathbf{q}} & 0 \\ 0 & 0 & 0 & -\omega^{+}_{\mathbf{Q}-\mathbf{q}} \end{pmatrix}. \tag{5.21}$$

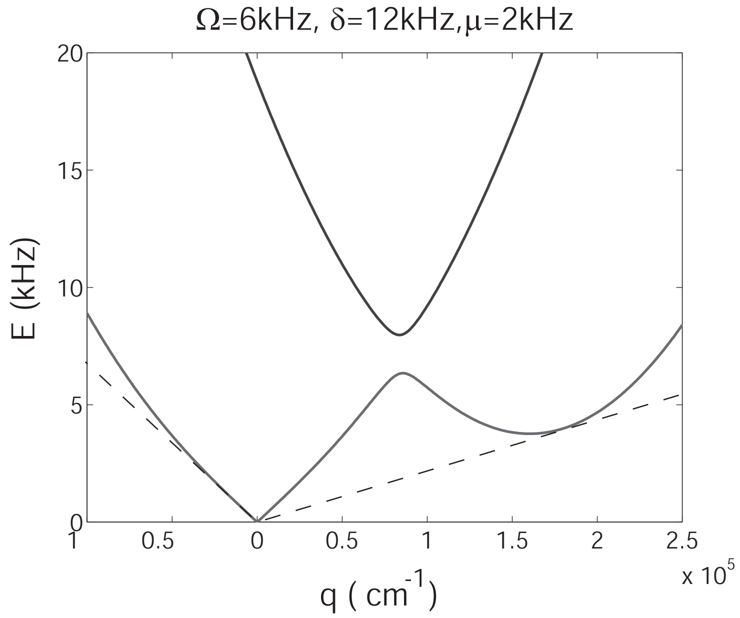

Figure 5.4. Dispersion relation for quasiparticle excitations of the condensate. In the neighborhood of the condensate, i.e., near $q = 0$, the energy displays a linear or phonon-like dependence on quasi-momentum q. The slope of the dashed lines determines the Landau critical velocity.

The interaction matrix M, furthermore, is equal to $g\Lambda x x^T$ with

$$x = (\cos\Delta_{\mathbf{q}}\,,\, \cos\Delta_{-\mathbf{q}}\,,\, \sin\Delta_{\mathbf{q}}\,,\, \sin\Delta_{-\mathbf{q}})^T\,.$$

Thus we have merely to diagonalize $H_0 + \mu M$ in the normal fashion to determine the dispersion relation of the quasiparticle excitations of a periodically dressed condensate. This is easily done numerically; the results of this diagonalization are shown in figure 5.4. It is seen that the lower minimum is no longer quadratic, as in the single-particle case, but instead has become linear or V-shaped around the condensate momentum $\mathbf{q} \approx \mathbf{Q}$, just as in the original treatment of Bogoliubov.

5.3 Superfluidity of Periodically Dressed Condensate

The quasiparticle dispersion relation is reminiscent of that of ^{4}He in its general form, possessing a linear dependence on wave vector in the neighborhood of the minimum (i.e., long-wavelength phonons) and a quadratic minimum at a nonzero difference wave vector. This secondary minimum is similar to the roton dip in the dispersion relation of superfluid ^{4}He, although the similarity is purely on the kinematic level, and should not be taken to imply any degree of sameness on the microscopic level.

Emboldened by this similarity, we may seek to apply the arguments of Landau [107] for the existence of superfluidity. According to the famous Landau criterion, a massive impurity of mass M and velocity $\mathbf{V}$ traveling through a fluid can lose energy to the fluid only if there exist energy- and momentum-conserving processes which can account microscopically for the frictional dissipation. Landau likewise reasoned that the fluid, being fundamentally a quantum-mechanical system, must possess elementary excitations, that is, low-lying excited states above the many-body ground state, and that the dissipative processes in question could be analyzed as the creation of these elementary excitations by the massive impurity. In this line of reasoning, if the dispersion relation of the elementary excitations is such that for momentum $\mathbf{p}$ the excitation energy is $\epsilon(\mathbf{p})$, then conservation of energy and momentum result in the relation

$$\frac{1}{2}M\mathbf{V}^2 = \epsilon(\mathbf{p}) + \frac{1}{2}M\left(\mathbf{V} - \frac{\mathbf{p}}{M}\right)^2, \tag{5.22}$$

or, in other words, if

$$\epsilon(\mathbf{p}) \geq \mathbf{V}\cdot\mathbf{p} - \frac{\mathbf{p}^2}{2M}, \tag{5.23}$$

then the impurity has insufficient energy to create the excitation while conserving momentum. This criterion is thus that if

$$|\mathbf{V}| \leq \min\frac{\epsilon(\mathbf{p}) + \mathbf{p}^2/2M}{|\mathbf{p}|}, \tag{5.24}$$

no dissipation can occur. The mass M is commonly taken to infinity to describe macroscopic objects moving through the fluid. For proposed experiments with the periodically dressed condensate, however, the case where M is equal to the mass of atoms in the fluid is also of interest, and so the criterion has been left general.

A convenient graphical method (see figure 5.4) for determining when the condition (5.24) is satisfied is to plot both $\epsilon(p) - p^2/2M$ and the auxiliary linear function $f(p) \equiv V\cdot p$ versus p. If the slope V of f is sufficiently small that $f(p)$ and $\epsilon(p)$ intersect only at the origin, then (5.24) is satisfied, and the largest value V_{Landau} of V

for which this condition holds true occurs when the line $f(p)$ is tangent to the curve $\epsilon(p)$.

5.4 Periodically Dressed Condensates and the Double-Well Problem

The concept of entanglement has long been recognized to lie at the heart of the sometimes strange and counterintuitive behavior of quantum mechanics [108, 109]. More recently, numerous research groups have begun to turn the sorts of nonclassical correlations present in measurements on entangled particles into practical tools, with applications in quantum computation and communication. Proposals have also been made for using correlated many-body states to enhance the statistical sensitivity of frequency and time measurements[110], analogous to the use of squeezed light to improve phase measurements.

To date, experimental realizations of controlled entanglement in atomic-physics systems have involved small numbers of particles. In contrast, superposition states of superconducting currents have been achieved involving large numbers of electrons, but with very short coherence times. As a result, it is appealing to attempt to achieve atomic entangled states with larger numbers of atoms, in order to benefit from the improved scaling of sensitivity with atom number.

Many interacting bosons in a double well are a paradigmatic system for the study of many-particle entanglement, as in, e.g., [111, 112]. As shown in section 5.5, the ground state of the double-well system is very different depending whether the interatomic interactions are attractive or repulsive. The repulsive case results in a relative-number-squeezed ground state, while the attractive case results in a pair of near-degenerate maximally entangled states, which are appealing for use in enhancing the sensitivity of atomic clocks (see section 5.6). Unfortunately, it is difficult to create large stable condensates with attractive interactions, since in general attractive interactions render a condensate unstable to collapse, although attractive condensates up to a certain critical size are stable. The periodically dressed condensate provides a way around this difficulty, by allowing the double well to appear in momentum space, where the interatomic interaction is (in an appropriate limit) attractive, even while it remains repulsive and hence stable in position space.

5.4.1 Raman Resonance and External Potential

To highlight the superfluid nature of the periodically dressed condensate, the foregoing discussion has all been of the homogeneous system. In order to consider the application of the periodically dressed BEC to investigation of the double-well system, however, it is useful to consider the effects of an inhomogeneous trapping potential. Such a potential is a necessity for the investigation of any but very short-time-scale phenomena in practice, since in its absence an atomic gas would fall under gravity and expand due to interatomic repulsion.

Accordingly, one considers the Hamiltonian (5.14) specialized to the case of Raman resonance ($\delta = 0$) , with the additional potential term

$$H_{\mathrm{trap}} = \int \mathrm{d}^3\mathbf{r}\Psi^\dagger(\mathbf{r})\frac{1}{2}m\omega_{\mathrm{t}}^2\mathbf{r}^2\Psi(\mathbf{r}). \tag{5.25}$$

For $\delta = 0$ the lower branch of the dispersion relation (5.7) resembles a double well, but in momentum space (see figure 5.3). To exploit this similarity with the well studied double-well problem, we write the entire Hamiltonian in momentum space, with the correspondence $x \leftrightarrow i\frac{\partial}{\partial q}$.

The trapping potential in momentum space and in the dressed-state basis takes the form

$$H_{\mathrm{trap}} = -\frac{1}{2}m\omega_{\mathrm{t}}^2\sum_{\mathbf{q}}\begin{pmatrix}\pi_{\mathbf{q}}^\dagger & \mu_{\mathbf{q}}^\dagger\end{pmatrix} e^{i\sigma_y\theta_{\mathbf{q}}/2}\nabla_{\mathbf{q}}^2\left[e^{-i\sigma_y\theta_{\mathbf{q}}/2}\begin{pmatrix}\pi_{\mathbf{q}}\\ \mu_{\mathbf{q}}\end{pmatrix}\right], \tag{5.26}$$

where an integration by parts has been performed and the boundary terms neglected. This is justifiable for the spatially localized atomic gas envisioned for these experiments, whose wave function and its derivatives must vanish at large distances.

If the energy of the atoms in the sample is small compared to the splitting between the upper and lower dressed states, which for large Ω is equivalent to their being small compared to Ω, then they may effectively be considered to populate the states on the lower branch of the dispersion relation. In this case, one can neglect the terms involving $\pi_{\mathbf{q}}$ and $\pi^{\dagger}_{\mathbf{q}}$ in equation (5.26). In this case, the noninteracting Hamiltonian simplifies to

$$H = \sum_{\mathbf{q}} \mu^{\dagger}_{\mathbf{q}} \left(\frac{1}{2} m\omega_t^2 \left(i\nabla_{\mathbf{q}} \right)^2 + \omega_q^- + \frac{m\omega_t^2}{8} \left(\hat{k} \cdot \nabla_{\mathbf{q}} \theta_{\mathbf{q}} \right)^2 \right) \mu_{\mathbf{q}}. \tag{5.27}$$

The first two terms are readily interpreted as momentum-space "kinetic" and "potential" energy. The remaining term may be considered as a gauge potential, resulting from the momentum-dependent change-of-basis. Evaluated explicitly, it is

$$\hat{k} \cdot \nabla_{\mathbf{q}} \theta_{\mathbf{q}} = \frac{\frac{2\hbar k}{m\Omega}}{1 + \left(\frac{2\hbar \mathbf{k}\cdot\mathbf{q}}{m\Omega} \right)^2}. \tag{5.28}$$

When only the lower dressed state is considered, the importance of this term is minor, representing only a small change in the shape of the double-well potential $\omega^-_{\mathbf{q}}$.

This Hamiltonian may be made to appear even simpler by an appropriate choice of units; if energy is measured in terms of the recoil energy $\hbar^2 k^2 / 2m$ and the dimensionless coordinate $q = k \cdot x$ is introduced, the single-atom Hamiltonian takes the form

$$\mathcal{H} = -\frac{1}{M} \frac{\partial^2}{\partial x^2} + W(x), \tag{5.29}$$

where $M = \left(\frac{\hbar k}{m\omega_t} \right)^2$ and W is the rescaled momentum-space double-well potential. This form is useful because the effective mass M, which governs the strength of tunnelling, is exhibited explicitly in terms of system parameters. However, the more conventional units will continue to be used for the following discussion, in order not to divorce the analysis any more than necessary from the physical system in which it is realized.

The Hamiltonian (5.27) is now of the standard double-well form, with the momentum $\mathbf{q}$ playing the rôle of position. It may be diagonalized with by considering the solutions of the Schrödinger equation

$$\left(\frac{1}{2} m\omega_t^2 \left(i\nabla_{\mathbf{q}} \right)^2 + \omega_q^- + \frac{m\omega_t^2}{8} \left(\hat{k} \cdot \nabla_{\mathbf{q}} \theta_{\mathbf{q}} \right)^2 \right) \phi(\mathbf{q}) = E\phi(\mathbf{q}). \tag{5.30}$$

These solutions come in near-degenerate even-odd pairs. It is assumed that the energy of the atoms is small compared to the splitting between the lowest and next-to-lowest pairs of eigenstates (which is of order $\hbar\omega_t$), but not compared to the splitting between the quasi-degenerate pair. This assumption is not guaranteed to be respected in an experimental implementation, though for small condensates in a tight trap it should be feasible. Even if interatomic interactions cause substantial numbers of oscillator states to be occupied, however, it will still be the case that there exists a lowest-energy near-degenerate pair of even and odd states which may be calculated self-consistently, e.g., by the Gross-Pitaevskii equation. This is the approach of Montina and Arrechi [113], as well as of Javanainen and Ivanov [114].

One can expand the operator $\mu_{\mathbf{q}}$ in terms of the even and odd states, truncating the expansion to the populated states $\phi_E(q)$ and $\phi_O(q)$

$$\mu_{\mathbf{q}} = \phi_E(\mathbf{q})\mu_E + \phi_O(\mathbf{q})\mu_O, \tag{5.31}$$

in terms of which the truncated version of Hamiltonian (5.27) is diagonalized

$$H = E_E \mu_E^\dagger \mu_E + E_O \mu_O^\dagger \mu_O. \tag{5.32}$$

For the double-well problem, it is convenient to define "left" and "right" wave functions as

$$\begin{aligned} \phi_L(\mathbf{q}) &\equiv \frac{\phi_E(\mathbf{q}) - \phi_O(\mathbf{q})}{\sqrt{2}} \\ \phi_R(\mathbf{q}) &\equiv \frac{\phi_E(\mathbf{q}) + \phi_O(\mathbf{q})}{\sqrt{2}} \end{aligned} \tag{5.33}$$

and corresponding creation operators defined by

$$|L\rangle = \mu_L^\dagger|0\rangle \qquad \text{or} \qquad |R\rangle = \mu_R^\dagger|0\rangle. \tag{5.34}$$

The non-interacting Hamiltonian, expressed in the left-right basis, is

$$H_0 = -\frac{\mathcal{J}_0}{2}\left(\mu_R^\dagger\mu_L + \mu_L^\dagger\mu_R\right), \tag{5.35}$$

where $\mathcal{J}_0 = E_O - E_E$ is the tunnelling rate (equal to the even-odd energy splitting).

To describe the effects of interactions one may write the Fourier-transformed density operator of equation (5.15) in terms of μ_L and μ_R

$$n_{\mathbf{q}} = A_{\mathbf{q}}N + B_{\mathbf{q}}\mu_L^\dagger\mu_R + B_{-\mathbf{q}}\mu_R^\dagger\mu_L, \tag{5.36}$$

where $A_{\mathbf{q}}$ and $B_{\mathbf{q}}$ are overlap integrals among the different states, with a weighting determined by the mixing angles at each momentum $\mathbf{q}$:

$$\begin{aligned} A_{\mathbf{q}} &= \sum_{\mathbf{l}} \phi_L^*(\mathbf{l}+\mathbf{q}/2)\phi_L(\mathbf{l}-\mathbf{q}/2)\cos\left(\frac{\theta_{\mathbf{l}+\mathbf{q}/2} - \theta_{\mathbf{l}-\mathbf{q}/2}}{2}\right) \\ B_{\mathbf{q}} &= \sum_{\mathbf{l}} \phi_L^*(\mathbf{l}+\mathbf{q}/2)\phi_R(\mathbf{l}-\mathbf{q}/2)\cos\left(\frac{\theta_{\mathbf{l}+\mathbf{q}/2} - \theta_{\mathbf{l}-\mathbf{q}/2}}{2}\right). \end{aligned} \tag{5.37}$$

The function $A_{\mathbf{q}}$ is peaked at $q = 0$, while $B_{\mathbf{q}}$ is peaked at $q \approx k$.

In terms of these coefficients, the collision term $H_{\text{int}} = \frac{2\pi\hbar^2 a}{mV} \sum_{\mathbf{q}} \left[n_{\mathbf{q}} n_{-\mathbf{q}}\right]$, although nonlocal in momentum space, has the relatively simple form (in the two-mode approximation)

$$\begin{aligned} H_{\text{int}} &= -UN_R N_L + \frac{\mathcal{J}_1}{2}\left(\mu_R^\dagger \mu_L + \mu_L^\dagger \mu_R\right) \\ &\quad + \mathcal{J}_2\left(\mu_R^\dagger \mu_R^\dagger \mu_L \mu_L + \mu_L^\dagger \mu_L^\dagger \mu_R \mu_R\right) \end{aligned} \tag{5.38}$$

where

$$\mathcal{J}_1 = \frac{8\pi\hbar^2 aN}{mV} \sum_{\mathbf{q}} A_{\mathbf{q}} B_{\mathbf{q}} \qquad\qquad \mathcal{J}_2 = \frac{2\pi\hbar^2 a}{mV} \sum_{\mathbf{q}} B_{\mathbf{q}} B_{-\mathbf{q}}$$

$$U = -\frac{4\pi\hbar^2 aN}{mV} \sum_{\mathbf{q}} B_{\mathbf{q}}^2, \tag{5.39}$$

with $N_L = \mu_L^\dagger \mu_L$ and likewise for N_R. An additive energy offset depending only on $N = N_L + N_R$ has been suppressed. The pair-tunnelling coefficient $\mathcal{J}_2$ is small, provided the left and right wave functions are reasonably well localized, since $B_{\mathbf{q}}$ is centered at $\mathbf{Q}$ while $B_{-\mathbf{q}}$ is centered at $-\mathbf{Q}$. For subsequent analysis, this term will thus be neglected. A simplified analytical form of these expressions for the case when the left and right states may be approximated as gaussian functions is given in Higbie et al.[42].

5.5 Ground States of the Double-Well Problem

And, finally, to complete the reduction of the Hamiltonian to a standard and easily manipulable form, one may define angular momentum operators

$$\begin{aligned} J_z &= \frac{1}{2}(\mu_R^\dagger \mu_R - \mu_L^\dagger \mu_L) \\ J_+ &= J_-^\dagger = \mu_R^\dagger \mu_L. \end{aligned} \tag{5.40}$$

These operators obey the usual angular-momentum commutation relations. The eigenstates of J_z are labeled by the population imbalance m_J and denoted by

$|J = N/2, m_J\rangle$, whose meaning is that $N_R = N/2 + m_J$ atoms are in the right well and $N_L = N/2 - m_J$ in the left.

In terms of the J operators, the Hamiltonian may be written

$$H = -\frac{\mathcal{J}}{2} J_x + U J_z^2, \tag{5.41}$$

where

$$\mathcal{J} = \mathcal{J}_0 - \mathcal{J}_1, \tag{5.42}$$

and again offset terms dependent only on N have been discarded.

From the Hamiltonian (5.41), it is easy to compute the many-body ground state in the extreme cases that either the tunnelling term or the interaction term dominates. Indeed, if $\mathcal{J} \gg U\sqrt{N}$, then the ground state approaches the highest-weight eigenstate of J_x, i.e., the state with all atomic pseudospin vectors aligned along x direction:

$$\frac{1}{\sqrt{2^N N!}} \left(\mu_R^\dagger + \mu_L^\dagger\right)^N |0\rangle. \tag{5.43}$$

Note the $\mathcal{J}$ is assumed positive, since the symmetric (even) ground state should be lower in energy than the antisymmetric (odd) state.

If, in the other extreme $UN \gg \mathcal{J}$ so that interactions dominate, then the ground state depends on the sign of U. If $U > 0$, then the interaction energy is minimized when $m_J = 0$, that is, when $N/2$ atoms are in each well,

$$\frac{1}{(N/2)!} \left(\mu_R^\dagger \mu_L^\dagger\right)^{N/2} |0\rangle. \tag{5.44}$$

Squeezed states of this sort have been observed in optical lattices, e.g., by [115], and similar physics occurs in the Mott-insulator state [116]. A partially squeezed state is illustrated in figure 5.5.

On the other hand, if $U < 0$, then there are two degenerate ground states, with $m_J = \pm N/2$, corresponding to all the atoms being on the left or all on the right. The parity eigenstates that can be formed in this degenerate subspace (and in general the energy eigenstates when the tunnelling is nonzero, although small) are then

$$\frac{1}{\sqrt{N!}} \left(\mu_R^\dagger \pm \mu_L^\dagger\right)^N |0\rangle. \tag{5.45}$$

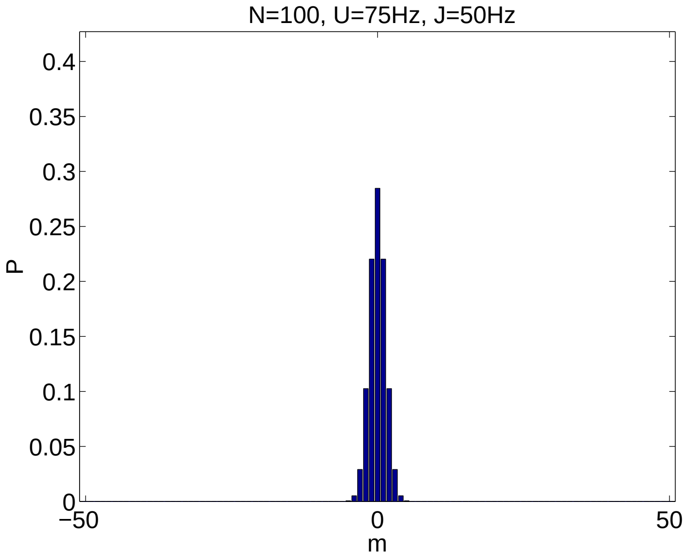

Figure 5.5. Squeezed state number distribution. The quantity m labels the population imbalance between left and right wells. For attractive interactions equal populations are favored, resulting in a narrowed distribution around $m = 0$.

These states are directly analogous to Schrödinger's famous $|\text{Alive}\rangle + |\text{Dead}\rangle$ cat state[108], since they contain a quantum superposition of macroscopically (or mesoscopically) distinct states. The energy difference between them $\hbar\Delta_{\text{sc}}$ will determine the macroscopic tunnelling time. It may be computed by an explicit diagonalization of the Hamiltonian. In figure 5.6, the probability distribution for a partial Schrödinger-cat state is shown, in which the probability distribution has bifurcated but in which there remains some joint probability for simultaneous detection in the two wells.

It is now possible to show that one can make good on the promise of section 5.4, that repulsive interactions in position space correspond to attractive interactions in momentum space, since positive scattering length a implies $U < 0$ according to equation (5.39), resulting in the cat-states of equation (5.45). This in intuitively reasonable, since position space repulsion tends to broaden the condensate wave function, which narrows it in momentum space.

Moreover, equations (5.39) and (5.7) highlight an important characteristic of the periodically dressed condensate for the exploration of many-body physics in the double well, namely the wide tunability of parameters. A short list of the readily tunable parameters and the corresponding effects follows:

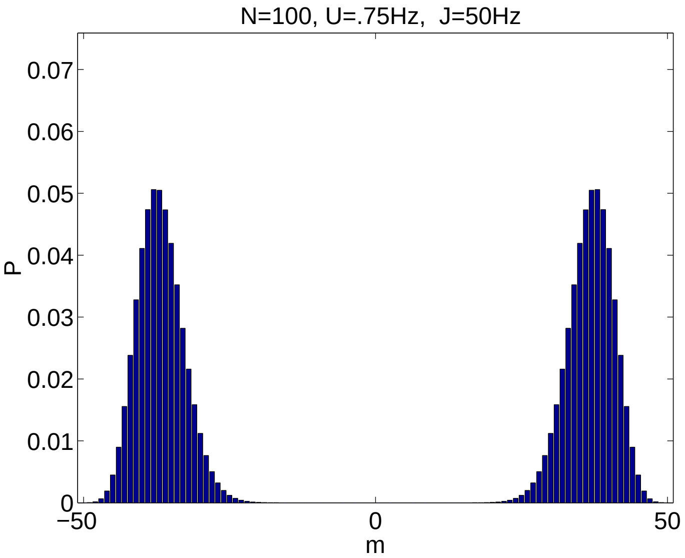

Figure 5.6. Partial Schrödinger-cat-state histogram. Repulsive interatomic interactions energetically favor having all atoms in one well. The ground state is a symmetric superposition of states with large population imbalance m.

- $\mathbf{k}$: The magnitude of the Raman momentum transfer may changed by allowing the Raman beams to intersect at different angles according to $|k| = 2k_0 \sin\beta$ where β is the angle between the beams, between a minimum of 0 and a maximum of $2k_0$. This has the effect of changing the separation between the two wells in momentum space, which primarily changes the tunneling rate $\mathcal{J}$. Dynamic tuning of this parameter, however, is challenging.

- Ω: The two-photon Rabi frequency may be varied by changing the power in either Raman beam or, within limits, the common detuning Δ. It can be varied smoothly between zero and approximately the recoil frequency $\hbar\mathbf{k}^2/2m$, at which point it deforms the double well out of existence. Changing Ω has the effect of raising and lowering the barrier between wells, as well as to a lesser extent changing the interactions by changing the character of the dressed states.

- N: The number of atoms participating is also not difficult to adjust within limits; in other words, it is not always easy to have more atoms, but fewer atoms may readily be obtained by loading the MOT for less time (see chapter 2) or by allowing trap loss to operate for a longer time. Changing N has the effect of changing the density (particularly in the small- condensate limit), which changes the interaction parameter U.

- ω_t: The trap frequency is generally dynamically variable over a fairly large range by changing the magnetic-trapping currents or by changing the power of an optical trap. Increasing the trap frequency decreases the effective momentum-space mass of the atoms, both increasing the tunnelling rate and increasing the density, which increases interactions. The effect on the tunnelling rate, however, is exponential and should dominate.

Another important advantage which this schemes for producing Schrödinger-cat states affords is the ability to map such states adiabatically into superpositions of internal states or of momentum states, which could then be easily detected either spectroscopically or by time-of-flight separation.

Indeed, if the Rabi frequency Ω were slowly decreased, the state $|L\rangle$ would adiabatically be converted into the state $|a\rangle$, while $|R\rangle$ would be adiabatically converted into $|b\rangle$. If $|a\rangle$ and $|b\rangle$ were states belonging to different hyperfine manifolds, then it would be easy to image them separately and thus to count their numbers.

5.5.1 Feasibility of Schrödinger-Cat States

The most serious experimental challenge in creating macroscopic superposition states is not merely creating a system in which the many-body ground state is such a state; it is, rather, a result of the same reason that such states are not observed in daily life. That is to say, it is interaction with the environment in the form of uncontrolled entanglement, or decoherence, that ultimately limits the number of particles that can be put into such a superposition state. As an illustration of the stringency of the requirements on decoherence for a Schrödinger-cat state, if even single particle is "measured" by the environment in the perfect state of equation (5.45), then the entire superposition would be collapsed.

Examples of such decohering measurements would be spontaneous emission from the Raman beams, number loss from the trap, collisions with thermal atoms or with background gas. To prove the existence of the superposition states (5.45), it is de-

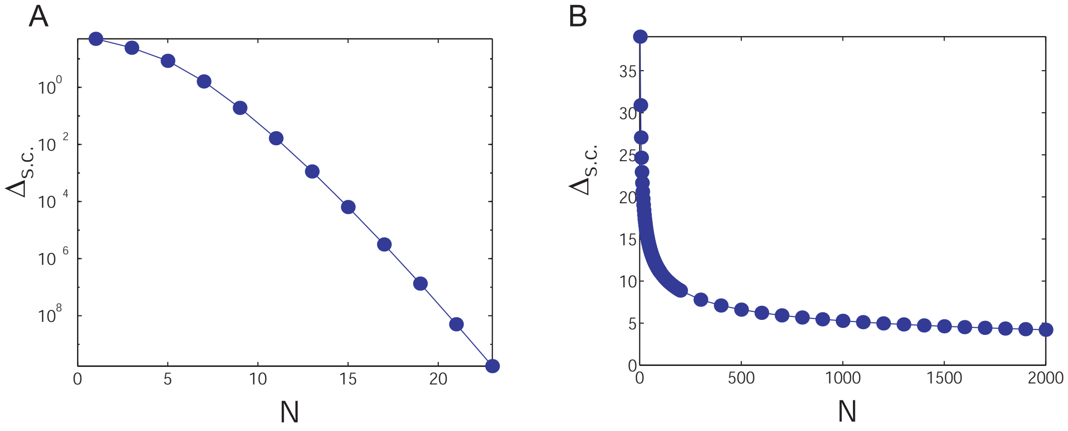

Figure 5.7. Pure and partial Schrödinger cat-state energy splitting vs. N. In both (A) and (B) the single-particle tunneling rate has been taken to be $\mathcal{J} = 50\,\mathrm{Hz}$. In (A) the interaction parameter U has been chosen according to the strict cat-state criterion $U = \mathcal{J}/\sqrt{N}$, and in (B) the partial-cat-state criterion, $U = \mathcal{J}/N$ has been used to determine U. In each case, the many-body Hamiltonian, for each value of N, was diagonalized and the splitting between the two lowest eigenvalues extracted. This splitting, $\Delta_{\mathrm{S.C.}}$, is plotted against N for the two cases. It is clear that the many-atom tunneling times rapidly become prohibitively long in case (A), while in case (B), the tunneling rates remain in a reasonable experimental range for hundreds to thousands of atoms.

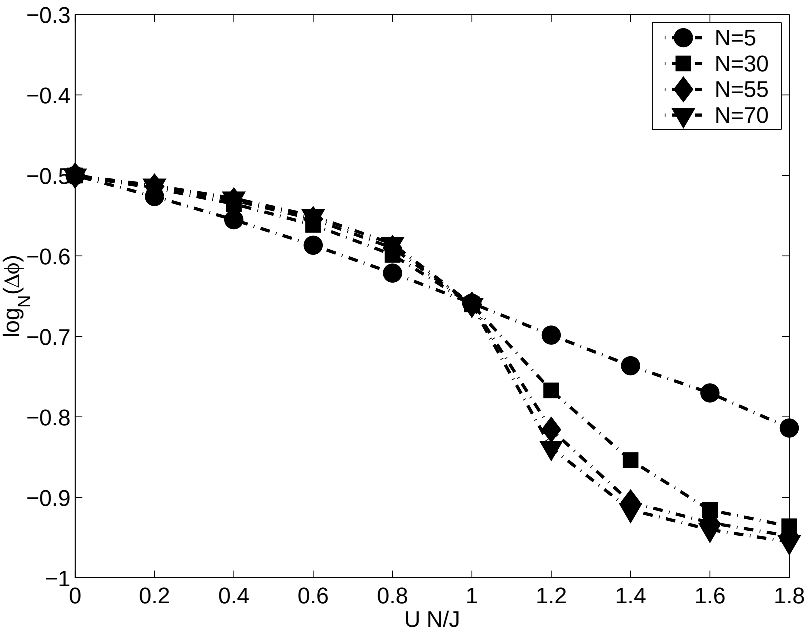

Figure 5.8. Clock enhancement from partial cat state. For weak interactions ($U \approx 0$) the uncertainty in phase $\Delta\phi$ scales as shot-noise, $\Delta\phi \propto N^{-1/2}$. Beyond the phase-transition to partial cat states (bifurcation in the probability distribution), around $UN/\mathcal{J} \approx 1$ the scaling improves rapidly, approaching the Heisenberg limit $\Delta\phi \propto N^{-1}$. A value of $\mathcal{J} = 100\,\mathrm{Hz}$ was used for this calculation.

sirable to see oscillations from the left well to right and back. This many-atom tunnelling rate is a rapidly decreasing function of the number of particles (see figure 5.7. These two effects, the increasingly deleterious effect of decoherence and the lengthening of the tunnelling time as N increases, militate against the observation of pure Schrödinger-cat states at large N.

A more hopeful alternative is to work not with the extreme cat-like distribution, in which all atoms are either in the left state or in the right, but rather with partial Schrödinger-cat states, in which the distribution is bimodal, with a peak on the left and one on the right, but with some finite width as in figure 5.6. Such states are substantially more forgiving than the pure cat states, since many atoms must be measured by the environment before the state is fully collapsed. As shown in figure 5.8, the enhancement in the scaling of measurement sensitivity with N (discussed in more detail in section 5.6) is largely realized even for partial cat states.

5.6 Application to Sub-Shot-Noise Interferometry

In a conventional microwave atomic clock based on Ramsey spectroscopy, atoms interact with two resonant microwave cavities (or in a fountain, the same one twice) driven by a single local oscillator. Each pulse of microwaves experienced by the atoms effects a $\pi/2$ Rabi pulse. The pulses are separated in time by a period of free evolution during which the atomic states simply accumulate phase according at their respective Bohr frequencies. If $|a\rangle \leftrightarrow (1,0)^T$ represents the initial state of the atoms and $|b\rangle \leftrightarrow (0,1)^T$ the state to which the microwave transition connects the initial state (these need not be the $|a\rangle$ and $|b\rangle$ of the preceding sections), then the matrix representing the $\pi/2$ pulse looks like an optical beamsplitter

$$U_1 = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -i \\ -i & 1 \end{pmatrix}. \tag{5.46}$$

The free evolution operator, assuming a difference in energies between the states of δ_{hf} since usually they are states of different hyperfine manifolds, may be written

$$U_{\mathrm{free}} = \begin{pmatrix} e^{i\delta_{\mathrm{hf}}T/2} & 0 \\ 0 & e^{-i\delta_{\mathrm{hf}}T/2} \end{pmatrix}. \tag{5.47}$$

When the atoms reach second microwave cavity, they experience another $\pi/2$ pulse, but with a relative phase $\phi = \delta_{\mathrm{osc}}T$ accumulated by the microwave oscillator of frequency δ_{osc} during the time T: $U_2 = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -ie^{-i\phi} \\ -ie^{i\phi} & 1 \end{pmatrix}$ A final state-selective ionization or fluorescence signal is then proportional to

$$|\langle b|U_2 U_{\mathrm{free}} U_1|a\rangle|^2 = \cos^2\left[\frac{(\delta_{\mathrm{hf}} - \delta_{\mathrm{osc}})T}{2}\right]. \tag{5.48}$$

Since in a conventional clock the atoms are uncorrelated, N atoms undergoing the same evolution will end up in the factorizable state

$$\frac{1}{\sqrt{N!}}[a^\dagger \cos\theta + b^\dagger \sin\theta]^N|0\rangle = \cos^N\theta \sum_{n=0}^{N} \frac{\sqrt{N!}\tan^n\theta}{\sqrt{(N-n)!n!}}|N-n, n\rangle, \tag{5.49}$$

where

$$\theta \equiv (\delta_{\mathrm{hf}} - \delta_{\mathrm{osc}})T/2. \tag{5.50}$$

On any given iteration of the experiment, a determination of $N_b = n$ will be made, with the probabilities of any particular result being given by the squared coefficients of equation (5.49). An illustration of simulated measurements on such an atomic state is shown in figure 5.9(a) for $N = 100$ as a function of phase. From this binomial distribution, it is straightforward to compute the mean and r.m.s. deviation. They are

$$\begin{aligned} \langle N_b \rangle &= N \sin^2 \theta \\ \Delta N_b &= \sqrt{N} \sin\theta \cos\theta. \end{aligned} \tag{5.51}$$

From these the r.m.s. deviations of θ are easily determined:

$$\Delta\theta = \frac{\Delta N_b}{\frac{\mathrm{d}N_b}{\mathrm{d}\theta}} = \frac{1}{2\sqrt{N}}, \tag{5.52}$$

and finally, the frequency and time uncertainty from a single iteration of the experiment are given with reference to equation (5.50) by

$$\left|\frac{\Delta t}{t}\right| = \left|\frac{\Delta \nu}{\nu}\right| = \frac{1}{2\pi\nu T\sqrt{N}}. \tag{5.53}$$

This scaling of the signal-to-noise ratio as $N^{-1/2}$ is generic to schemes involving uncorrelated atoms, distributed either according to the binomial distribution (as here) or according to the closely related Poisson distribution.

Suppose, now that instead of acting on each atom independently, the $\pi/2$ pulses acted directly on the N-body states (such a pulse has been christened a "magic beamsplitter" by Lee et al. [117]). The action of such a beamsplitter or $\pi/2$ pulse in the basis $\{|a\rangle^N, |b\rangle^N\}$ (in the interaction-dominated limit of equation (5.41)) would then be

$$U = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -ie^{i\phi} \\ -ie^{-i\phi} & 1 \end{pmatrix}. \tag{5.54}$$

Now a sample of atoms initially prepared in $|a\rangle^N$ would be mapped onto

$$\frac{1}{\sqrt{2}}\big[|a\rangle^N - i|b\rangle^N\big], \tag{5.55}$$

taking the initial beamsplitter phase to be zero. After evolving for time T, the state would become

$$\frac{1}{\sqrt{2}}\big[|a\rangle^N - ie^{-iN\delta_{\mathrm{hf}}T}|b\rangle^N\big]. \tag{5.56}$$

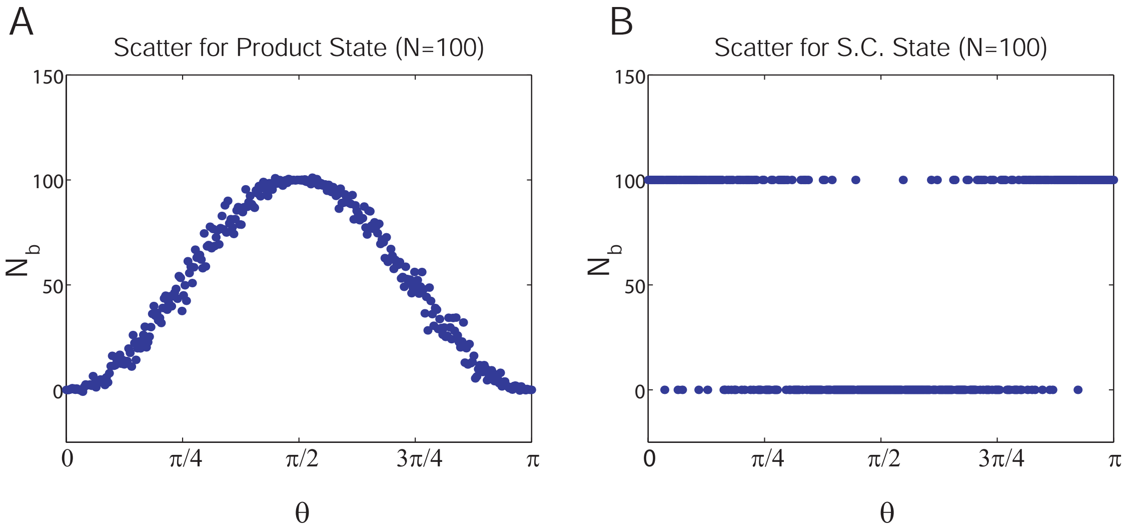

Figure 5.9. Comparison of projected measurements on uncorrelated atoms (A) to measurements on perfectly correlated Schrödinger cat-type many-atom state (B). The mean value of atoms in state 'b' $\langle N_b \rangle$ is the same for (A) and (B), but the cat state exhibits maximal fluctuations about the mean.

And after a second magic $\pi/2$ pulse the probability of being in state $|b\rangle^N$ would be $P_b = \cos^2\left(\frac{NT(\delta_{\text{hf}}-\delta_{\text{osc}})}{2}\right)$

If such a magic beamsplitter could be realized, it would dramatically affect the statistics of the final measurement. Proceeding in an analogous manner to the uncorrelated case, one can compute the moments of the distribution defined by the state

$$\cos\theta|a\rangle^N + \sin\theta|b\rangle^N. \tag{5.57}$$

One obtains the mean number of 'b' atoms $\langle N_b \rangle = N\sin^2\theta$, the r.m.s. deviation of this number $\Delta N_b = N\sin\theta\cos\theta$, and the corresponding r.m.s. deviation of the angle θ defined by

$$\theta = \pi N(\nu_{\text{hf}} - \nu_{\text{osc}})T, \tag{5.58}$$

that is to say,

$$\Delta\theta = \Delta N_b \left(\frac{\mathrm{d}N_b}{\mathrm{d}\theta}\right)^{-1} = \frac{1}{2}. \tag{5.59}$$

Thus the time uncertainty is given by

$$\left|\frac{\Delta t}{t}\right| = \frac{1}{2\pi\nu TN}. \tag{5.60}$$

Note that the scaling is now N^{-1} rather than $N^{-1/2}$. This improvement comes about from the fact that although $\Delta\theta$ larger for the Schrödinger-cat states by $\sqrt{N}$, these states accumulate phase N times as fast, so that time is effectively being measured against a more finely graded ruler. The combination of these two effects accounts for the net $\sqrt{N}$ improvement in scaling. An illustration of simulated measurements taken on such a cat state is shown in figure 5.9(b); atoms are detected either in one state or the other, but the probability varies sinusoidally, as in figure 5.9(a).

Thus far the magic beamsplitter has been hypothetical, but in fact it should be possible to realize such a beamsplitter with the Raman-resonant periodically dressed condensate. The beamsplitting pulse would proceed as follows: atoms would begin in the state $|a\rangle^N$. One would ramp up the Rabi frequency Ω adiabatically with respect to itself ($\frac{\mathrm{d}\Omega}{\mathrm{d}t} \ll \Omega^2$) and to the trap frequencies, but rapidly relative to the many-atom tunnelling time, yielding the state $|L\rangle^N$. After a quarter period of the many-body oscillation (i.e., a time of $\pi/2\Delta_{\mathrm{sc}}$), this state will have oscillated into the macroscopic superposition state $\left(|L\rangle^N + |R\rangle^N\right)/\sqrt{2}$, at which point adiabatically turning off the Raman dressing beams will map the state back to a macroscopic superposition of the undressed states $\left(|a\rangle^N + |b\rangle^N\right)/\sqrt{2}$. Thus this process is precisely the magic beamsplitter described above.

5.7 Stabilizing a Schrödinger-Cat Clock

In a conventional atomic clock, the number of atoms measured in, say, the upper hyperfine state on each cycle of the experiment has a small fractional deviation, approximately $1/\sqrt{N}$, so that a reasonably reliable estimate of the amount by which the local oscillator frequency has drifted may be made on that cycle and the appropriate change in its frequency made to bring it back into line with the atomic frequency. In contrast, the corresponding signal from a cat-state would exhibit fluctuations in the atom number of order 100%, making it almost a random guess on each cycle of the experiment whether to retard or to accelerate the local oscillator. At first glance, then, one might doubt the practicality of designing the appropriate feedback system to stabilize the local oscillator to the atomic frequency. In fact, however, such a

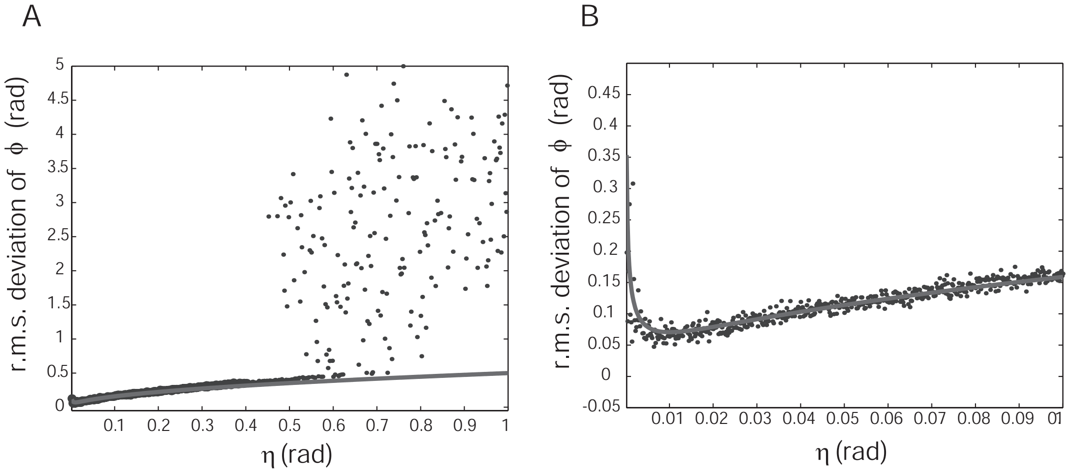

Figure 5.10. Root-mean-squared deviations in the phase ϕ accumulated by the clock on each experimental cycle versus feedback step size η. Solid curve in (A) and (B) is theoretical curve from equation (5.67); solid circles are simulation results based on recursion relation (5.62). The technical noise parameter is here chosen to be $\xi = 10\,\mathrm{mrad}$. As seen in (A), for values of $\eta \gtrsim 0.45$ the feedback is unstable, while for smaller values the scheme is stable. In (B) detailed plot of the small-η region is shown, revealing a minimum in the fluctuations when the feedback step size equals the r.m.s. technical noise per cycle, $\eta = \xi$.

feedback system appears quite possible, and should not be the limiting factor in any implementation of such a clock.

With each new batch of atoms that emerges from its second magic $\pi/2$ pulse, one will in the perfect Schrödinger-cat case measure either all atoms in state $|a\rangle$ or all in state $|b\rangle$. The probabilities of these two events are given by

$$\begin{aligned} P_\mathrm{a} &= \frac{1+\sin(2\phi)}{2} \approx \frac{1}{2} + \phi \\ P_\mathrm{b} &= \frac{1-\sin(2\phi)}{2} \approx \frac{1}{2} - \phi, \end{aligned} \tag{5.61}$$

where $\phi = \frac{\pi}{4} + \theta$, with θ as defined in equation (5.58). The assumption that ϕ can be kept small may or may not be borne out in the subsequent development; it is convenient for analysis, but not necessary. If on a given run of the experiment the detector clicks "b" then it is marginally more probable that $\phi < 0$, so that statistically it makes sense to retard the local-oscillator frequency slightly so as to increase ϕ on the next iteration. Suppose that the amount by which ϕ is increased, then, is a small quantity η so that on iteration $j+1$ the phase ϕ_{j+1} is related to that on the previous iteration j by

$$\phi_{j+1} = \phi_j + \eta x_j + \xi_j, \tag{5.62}$$

where the quantity x_j is conditional upon the results of measurement on iteration j.

$$x_j = \begin{cases} -1 & \text{if a measured} \\ +1 & \text{if b measured} \end{cases} \tag{5.63}$$

and ξ_j is random technical noise on shot number j, for instance detector noise in the final atom counting. From this description, it is possible to derive a recurrence relation for the average value of ϕ

$$\langle\phi_{j+1}\rangle = \langle\phi_j\rangle + \eta\langle x_j\rangle = \langle\phi_j\rangle\,(1-2\eta)\,. \tag{5.64}$$

Which has the solution

$$\phi_j = \phi_0(1-2\eta)^j \approx \phi_0 e^{-2\eta j}, \tag{5.65}$$

showing that the phase ϕ damps exponentially to zero with a time scale of approximately $1/2\eta$ iterations. One can also examine the fluctuations of ϕ in the steady state by taking the mean square of equation (5.62):

$$\langle\phi_{j+1}^2\rangle = \langle\phi_j^2\rangle + \eta^2\langle x_j^2\rangle + \langle\xi_j^2\rangle + 2\eta\langle\phi_j x_j\rangle. \tag{5.66}$$

In the steady state, $\langle\phi_{j+1}^2\rangle = \langle\phi_j^2\rangle$. One can explicitly evaluate $\langle\phi_j x_j\rangle = -2\langle\phi^2\rangle$ and use the fact that $\langle x_j^2\rangle = 1$ to obtain the mean square phase deviation

$$\langle\phi^2\rangle = \frac{\eta^2 + \langle\xi^2\rangle}{4\eta}. \tag{5.67}$$

This analytical value for the noise in the phase per cycle $\Delta\phi$ is compared to a simulation result based on the recursion relation (5.10) in figure 5.10, showing good agreement. It is seen that the diffusive, random-walk character of the feedback favors the use of a small value of η, so as not to walk away from the stabilized point too quickly. The desirability of rapid damping, on the other hand, favors a large value of η. In the absence of technical noise $\langle\xi^2\rangle$, the diffusion dependence is stronger than the damping dependence since at the stable point, damping plays no rôle. However, effective damping is required to recover from kicks to the system suffered from technical noise. As a result, an intermediate choice of η is optimal, and the r.m.s. phase $\Delta\phi$ attains a minimum value of $\sqrt[4]{\frac{\langle\xi^2\rangle}{4}}$ at the point $\eta^2 = \langle\xi^2\rangle$, as illustrated in figure 5.10.

5.8 Condensate-Number Fluctuations

Probably the most serious systematic afflicting the scheme outlined above is its dependence on the number of atoms N. Indeed, since the splitting Δ_{sc} is a strong function of N, the required duration of the magic beamsplitter pulse $\pi/2/\Delta_{\mathrm{sc}}$ to make it accurately a $\pi/2$ rotation will likewise change with N. In most existing ultracold-atom experiments, N is not a very well controlled parameter. Realistic proposals for control of atom number on the few-atom level have been made [118], and deterministic delivery of small numbers of atoms via an optical "conveyor belt" has been achieved [119]. Moreover, schemes exist for measuring numbers on the order of 100 [120], through which control of similar numbers might also be possible. It is also plausible that use of the Mott-insulator state, in which an integer number of atoms inhabits each site of an optical lattice, may be used to control atom numbers more effectively than was previously possible, particularly if variations in lattice occupation due to inhomogeneity of the trapping beams are eliminated, e.g., by superimposing the sort of repulsive box potentials used by the Raizen group in place of the usual inhomogeneous trapping potentials.

Bibliography

[1] W. E. Lamb and R. C. Retherford, Phys. Rev. **72**, 241 (1947).

[2] J. P. Carrico, E. Lipworth, P. G. H. Sandars, T. S. Stein, and M. C. Weisskopf, Phys. Rev. **174**, 125 (1968).

[3] S. A. Murthy, J. D. Krause, Z. L. Li, and L. R. Hunter, Phys. Rev. Lett. **63**, 965 (1989).

[4] B. C. Regan, E. D. Commins, C. J. Schmidt, and D. DeMille, Phys. Rev. Lett. **88**, 071805 (2002).

[5] K. Abdullah, C. Carlberg, E. D. Commins, H. Gould, and S. B. Ross, Phys. Rev. Lett. **65**, 2347 (1990).

[6] E. D. Commins, S. B. Ross, D. DeMille, and B. C. Regan, Phys. Rev. A **50**, 2960 (1994).

[7] H. Marion, F. P. D. Santos, M. Abgrall, S. Zhang, Y. Sortais, S. Bize, I. Maksimovic, D. Calonico, J. Grunert, C. Mandache, P. Lemonde, G. Santarelli, P. Laurent, A. Clairon, and C. Salomon, Phys. Rev. Lett. **90**, 150801 (2003).

[8] S. Bize, S. A. Diddams, U. Tanaka, C. E. Tanner, W. H. Oskay, R. E. Drullinger, T. E. Parker, T. P. Heavner, S. R. Jefferts, L. Hollberg, W. M. Itano, and J. C. Bergquist, Phys. Rev. Lett. **90**, 150802 (2003).

[9] A. T. Nguyen, D. Budker, S. K. Lamoreaux, and J. R. Torgerson, Phys. Rev. A **69**, 022105 (2004).

[10] E. Peik, B. Lipphardt, H. Schnatz, T. Schneider, C. Tamm, and S. G. Karshenboim, Phys. Rev. Lett. **93**, 170801 (2004).

[11] M. Fischer, N. Kolachevsky, M. Zimmermann, R. Holzwarth, T. Udem, T. W. Hansch, M. Abgrall, J. Grunert, I. Maksimovic, S. Bize, H. Marion, F. P. D. Santos, P. Lemonde, G. Santarelli, P. Laurent, A. Clairon, C. Salomon, M. Haas, U. D. Jentschura, and C. H. Keitel, Phys. Rev. Lett. **92**, 230802 (2004).

[12] T. E. Chupp, R. J. Hoare, R. A. Loveman, E. R. Oteiza, J. M. Richardson, M. E. Wagshul, and A. K. Thompson, Phys. Rev. Lett. **63**, 1541 (1989).

[13] C. J. Berglund, L. R. Hunter, J. D. Krause, E. O. Prigge, M. S. Ronfeldt, and S. K. Lamoreaux, Phys. Rev. Lett. **75**, 1879 (1995).

[14] J. J. Bollinger, J. D. Prestage, W. M. Itano, and D. J. Wineland, Phys. Rev. Lett. **54**, 1000 (1985).

[15] J. D. Prestage, J. J. Bollinger, W. M. Itano, and D. J. Wineland, Phys. Rev. Lett. **54**, 2387 (1985).

[16] P. Wolf, S. Bize, A. Clairon, A. N. Luiten, G. Santarelli, and M. E. Tobar, Phys. Rev. Lett. **90**, 060402 (2003).

[17] D. Brown, D. Budker, and D. P. DeMille, *Towards an improved test of Bose-Einstein statistics for photons* (AIP, ADDRESS, 2000), No. 1, pp. 281–287.

[18] T. W. Hänsch and A. L. Schawlow, Optics Communications **13**, 68 (1975).

[19] A. Ashkin, Phys. Rev. Lett. **40**, 729 (1978).

[20] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell, Science **269**, 198 (1995).

[21] K. B. Davis, M.-O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle, Phys. Rev. Lett. **75**, 3969 (1995).

[22] M. A. Kasevich, E. Riis, S. Chu, and R. G. DeVoe, Phys. Rev. Lett. **63**, 612 (1989).

[23] G. Santarelli, P. Laurent, P. Lemonde, A. Clairon, A. G. Mann, S. Chang, A. N. Luiten, and C. Salomon, Phys. Rev. Lett. **82**, 4619 (1999).

[24] D. S. Weiss, B. C. Young, and S. Chu, Phys. Rev. Lett. **70**, 2706 (1993).

[25] P. Cladé, E. De Mirandes, M. Cadoret, S. Guellati-Khélifa, C. Schwob, F. Nez, L. Julien, and F. Biraben, ArXiv Physics e-prints (2005).

[26] Y. Sortais, S. Bize, C. Nicolas, A. Clairon, C. Salomon, and C. Williams, Phys. Rev. Lett. **85**, 3117 (2000).

[27] F. P. D. Santos, H. Marion, S. Bize, Y. Sortais, A. Clairon, and C. Salomon, Phys. Rev. Lett. **89**, 233004 (2002).

[28] M. D. Barrett, J. A. Sauer, and M. S. Chapman, Phys. Rev. Lett. **87**, 010404 (2001).

[29] T. Weber, J. Herbig, M. Mark, H.-C. Nägerl, and R. Grimm, Science **299**, 232 (2003).

[30] T. Kinoshita, T. Wenger, and D. S. Weiss, Phys. Rev. A **71**, 011602 (2005).

[31] S. E. Maxwell, N. Brahms, R. deCarvalho, D. R. Glenn, J. S. Helton, S. V. Nguyen, D. Patterson, J. Petricka, D. DeMille, and J. M. Doyle, Phys. Rev. Lett. **95**, 173201 (2005).

[32] K. Gibble and B. J. Verhaar, Phys. Rev. A **52**, 3370 (1995).

[33] C. Fertig and K. Gibble, Phys. Rev. Lett. **85**, 1622 (2000).

[34] S. Gupta, Z. Hadzibabic, M. W. Zwierlein, C. A. Stan, K. Dieckmann, C. H. Schunck, E. G. M. van Kempen, B. J. Verhaar, and W. Ketterle, Science **300**, 1723 (2003).

[35] H. Katori, M. Takamoto, V. G. Pal'Chikov, and V. D. Ovsiannikov, Phys. Rev. Lett. **91**, 173005 (2003).

[36] M. Takamoto, F.-L. Hong, R. Higashi, and H. Katori, Nature **435**, 321 (2005).

[37] R. B. Laughlin, Phys. Rev. B **23**, 5632 (1981).

[38] N. Byers and C. N. Yang, Phys. Rev. Lett. **7**, 46 (1961).

[39] D. A. Lidar, I. L. Chuang, and K. B. Whaley, Physical Review Letters **81**, 2594 (1998).

[40] J. M. Higbie, L. E. Sadler, S. Inouye, A. P. Chikkatur, S. R. Leslie, K. L. Moore, V. Savalli, and D. M. Stamper-Kurn, Phys. Rev. Lett. **95**, 050401 (2005).

[41] J. Higbie and D. M. Stamper-Kurn, Phys. Rev. Lett. **88**, 090401 (2002).

[42] J. Higbie and D. M. Stamper-Kurn, Phys. Rev. A **69**, 53605 (2004).

[43] D. Steck, Rubidium ^{87}Rb D Line Data, `http://steck.us/alkalidata`, 2001.

[44] D. Durfee, Ph.D. thesis, Massachussetts Institute of Technology, 1999.

[45] H. J. Lewandowski, D. M. Harber, D. L. Whitaker, and E. A. Cornell, J. Low Temp. Phys. **132**, 309 (2003).

[46] D. E. Pritchard, Phys. Rev. Lett. **51**, 1336 (1983).

[47] A. L. Migdall, W. D. Phillips, J. V. Prodan, T. H. Bergeman, and H. J. Metcalf, Phys. Rev. Lett. **54**, 2596 (1985).

[48] T. Bergeman, G. Erez, and H. Metcalf, Phys. Rev. A **35**, 1535 (1987).

[49] J. D. Weinstein and K. G. Libbrecht, Phys. Rev. A **52**, 4004 (1995).

[50] M.-O. Mewes, M. R. Andrews, N. J. van Druten, D. M. Kurn, D. S. Durfee, and W. Ketterle, Phys. Rev. Lett. **77**, 416 (1996).

[51] C. V. Sukumar and D. M. Brink, Phys. Rev. A **56**, 2451 (1997).

[52] W. H. Wing, Prog. Quant. Electr. **8**, 181 (1984).

[53] W. Ketterle and D. E. Pritchard, App. Phys. B **54**, 403 (1992).

[54] M. S. Ioffe and R. I. Sobolev, J. Nuc. Energy C **7**, 501 (1965).

[55] D. E. Pritchard, Phys. Rev. Lett. **51**, 1336 (1983).

[56] W. Petrich, M. H. Anderson, J. R. Ensher, and E. A. Cornell, Phys. Rev. Lett. **74**, 3352 (1995).

[57] U. Ernst, A. Marte, F. Schreck, J. Schuster, and G. Rempe, Europhys. Lett. **41**, 1 (1998).

[58] C. J. Myatt, E. A. Burt, R. W. Ghrist, E. A. Cornell, and C. E. Wieman, Phys. Rev. Lett. **78**, 586 (1997).

[59] T. Esslinger, I. Bloch, and T. W. Hänsch, Phys. Rev. A **58**, R2664 (1998).

[60] H. F. Hess, Phys. Rev. B **34**, 3476 (1986).

[61] S. Chu, J. E. Bjorkholm, A. Ashkin, and A. Cable, Phys. Rev. Lett. **57**, 314 (1986).

[62] D. M. Stamper-Kurn, M. R. Andrews, A. P. Chikkatur, S. Inouye, H.-J. Miesner, J. Stenger, and W. Ketterle, Phys. Rev. Lett. **80**, 2027 (1998).

[63] S. Sydoriak, E. Grilly, and E. Hammel, Phys. Rev. **75**, 303 (1949).

[64] J. Bardeen, L. Cooper, and J. Schrieffer, Phys. Rev. **108**, 1175 (1957).

[65] D. D. Osheroff, R. C. Richardson, and D. M. Lee, Phys. Rev. Lett. **28**, 885 (1972).

[66] A. J. Leggett, Rev. Mod. Phys. **47**, 331 (1975).

[67] C. J. Myatt, N. R. Newbury, R. W. Ghrist, S. Loutzenhiser, and C. E. Wieman, Opt. Lett. **21**, 290 (1996).

[68] J. Stenger, D. M. Stamper-Kurn, M. R. Andrews, A. P. Chikkatur, S. Inouye, H.-J. Miesner, and W. Ketterle, J. Low Temp. Phys. **113**, 167 (1998).

[69] D. M. Stamper-Kurn, H.-J. Miesner, A. P. Chikkatur, S. Inouye, J. Stenger, and W. Ketterle, Phys. Rev. Lett. **83**, 661 (1999).

[70] H.-J. Miesner, D. M. Stamper-Kurn, J. Stenger, S. Inouye, A. P. Chikkatur, and W. Ketterle, Phys. Rev. Lett. **82**, 2228 (1999).

[71] M.-S. Chang, C. D. Hamley, M. D. Barrett, J. A. Sauer, K. M. Fortier, W. Zhang, L. You, and M. S. Chapman, Phys. Rev. Lett. **92**, 140403 (2004).

[72] W. Zhang, D. Zhou, M. Chang, M. Chapman, and L. You, Phys. Rev. A **72**, 031602 (2005).

[73] H. Schmaljohann, M. Erhard, J. Kronjäger, M. Kottke, S. van Staa, J. J. Arlt, K. Bongs, and K. Sengstock, J. Mod Opt. **51**, 1829 (2004).

[74] T. Kuwamoto, K. Araki, T. Eno, and T. Hirano, Phys. Rev. A **69**, 063604 (2004).

[75] T.-L. Ho, Phys. Rev. Lett. **81**, 742 (1998).

[76] C. K. Law, H. Pu, and N. P. Bigelow, Phys. Rev. Lett. **81**, 5257 (1998).

[77] T.-L. Ho and S. K. Yip, Phys. Rev. Lett. **84**, 4031 (2000).

[78] L. P. Pitaevskii, Sov. Phys. JETP **13**, 451 (1961).

[79] E. P. Gross, Nuovo Cimento **20**, 454 (1961).

[80] T. L. Gustavson, A. P. Chikkatur, A. E. Leanhardt, A. Gorlitz, S. Gupta, D. E. Pritchard, and W. Ketterle, Phys. Rev. Lett. **88**, 020401 (2002).

[81] M. R. Matthews, B. P. Anderson, P. C. Haljan, D. S. Hall, M. J. Holland, J. E. Williams, C. E. Wieman, and E. A. Cornell, Phys. Rev. Lett. **83**, 3358 (1999).

[82] F. Zernike, Month. Not. Roy. Astro. Soc. **94**, 0377 (1933).

[83] T.-L. Ho, Phys. Rev. Lett. **81**, 742 (1998).

[84] T. Ohmi and K. Machida, J. Phys. Soc. Jpn. **67**, 1822 (1998).

[85] N. P. Bigelow, J. H. Freed, and D. M. Lee, Phys. Rev. Lett. **63**, 1609 (1989).

[86] H. J. Lewandowski, D. M. Harber, D. L. Whitaker, and E. A. Cornell, Phys. Rev. Lett. **88**, 070403 (2002).

[87] J. M. McGuirk, H. J. Lewandowski, D. M. Harber, T. Nikuni, J. E. Williams, and E. A. Cornell, Phys. Rev. Lett. **89**, 090402 (2002).

[88] D. Budker, D. F. Kimball, V. V. Yashchuk, and M. Zolotorev, Phys. Rev. A **65**, 055403 (2002).

[89] I. Kominis, T. Kornack, J. Allred, and M. Romalis, Nature **422**, 596 (2003).

[90] J. R. Kirtley, Ann. Rev. of Mat. Sci. **29**, 117 (1999).

[91] C. Granata, A. Monaco, C. di Russo, M. P. Lissitski, and M. Russo, J. Mag. Mag. Mat. **272**, (2004).

[92] F. Gruhl, M. Muck, M. von Kreutzbruck, and J. Dechert, Rev. Sci. Inst. **72**, 2090 (2001).

[93] L. E. Fong, J. R. Holzer, K. K. McBride, E. A. Lima, F. Baudenbacher, and M. Radparvar, Rev. Sci. Inst. **76**, 053703 (2005).

[94] J. Allred, R. Lyman, T. Kornack, and M. Romalis, Phys. Rev. Lett. **89**, 130801 (2002).

[95] D. Budker, D. F. Kimball, S. M. Rochester, V. V. Yashchuk, and M. Zolotorev, Phys. Rev. A **6204**, 043403 (2000).

[96] Y.-J. Lin, I. Teper, C. Chin, and V. Vuletić, Phys. Rev. Lett. **92**, 050404 (2004).

[97] P. K. Rekdal, S. Scheel, P. L. Knight, and E. A. Hinds, Phys. Rev. A **70**, 013811 (2004).

[98] S. Gupta, K. Dieckmann, Z. Hadzibabic, and D. E. Pritchard, Phys. Rev. Lett. **89**, 140401 (2002).

[99] J. Javanainen and M. Wilkens, Phys. Rev. Lett. **78**, 4675 (1997).

[100] S. Inouye, A. P. Chikkatur, D. M. Stamper-Kurn, J. Stenger, D. E. Pritchard, and W. Ketterle, Science **285**, 571 (1999).

[101] A. Griesmaier, J. Werner, S. Hensler, J. Stuhler, and T. Pfau, Phys. Rev. Lett. **94**, 160401 (2005).

[102] C. A. Regal, C. Ticknor, J. L. Bohn, and D. S. Jin, Nature **424**, 47 (2003).

[103] M. W. Zwierlein, C. A. Stan, C. H. Schunck, S. M. Raupach, A. J. Kerman, and W. Ketterle, Phys. Rev. Lett. **92**, 120403 (2004).

[104] C. A. Regal, M. Greiner, and D. S. Jin, Phys. Rev. Lett. **92**, 040403 (2004).

[105] N. N. Bogoliubov, J. Phys. (USSR) **11**, 23 (1947).

[106] W. H. Bassichis, Phys. Rev. **134**, A543 (1964).

[107] L. D. Landau, J. Phys. (USSR) **5**, 71 (1941).

[108] E. Schrödinger, Naturwissenschaften **23**, pp. 807 (1935).

[109] A. Einstein, B. Podolsky, and N. Rosen, Phys. Rev. **47**, 777 (1935).

[110] D. J. Wineland, J. Bollinger, W. M. Itano, F. L. Moore, and D. J. Heinzen, Phys. Rev. A **46**, R6797 (1992).

[111] J. Javanainen, Phys. Rev. Lett. **57**, 3164 (1986).

[112] G. J. Milburn, J. Corney, E. M. Wright, and D. F. Walls, Phys. Rev. A **55**, 4318 (1997).

[113] A. Montina and F. T. Arecchi, Phys. Rev. A **67**, 23616 (2003).

[114] J. Javanainen and M. Y. Ivanov, Phys. Rev. A **60**, 2351 (1999).

[115] C. Orzel, A. K. Tuchman, M. L. Fenselau, M. Yasuda, and M. A. Kasevich, Science **291**, 2386 (2001).

[116] M. Greiner, O. Mandel, T. Esslinger, T. W. Hänsch, and I. Bloch, Nature **415**, 39 (2002).

[117] H. Lee, P. Kok, and J. P. Dowling, Jour. Mod. Opt. **49**, 2325 (2002).

[118] R. B. Diener, B. Wu, M. G. Raizen, and Q. Niu, Physical Review Letters **89**, 070401 (2002).

[119] S. Kuhr, W. Alt, D. Schrader, M. Müller, V. Gomer, and D. Meschede, Science **293**, 278 (2001).

[120] S. Leslie, N. Shenvi, K. R. Brown, D. M. Stamper-Kurn, and K. B. Whaley, Phys. Rev. A **69**, 043805 (2004).