

**Coherences and correlations in an ultracold Bose
gas**

by

Heather Jean Lewandowski

B.S. Physics, Michigan Technological University, 1997

M.S. Physics, University of Colorado, 2001

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
Department of Physics
2002

This thesis entitled:
Coherences and correlations in an ultracold Bose gas
written by Heather Jean Lewandowski
has been approved for the Department of Physics

Eric A. Cornell

Deborah S. Jin

Date _____

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

Lewandowski, Heather Jean (Ph. D. Physics)

Coherences and correlations in an ultracold Bose gas

Thesis directed by Dr. Eric A. Cornell

The development of a new and simplified system to create a Bose-Einstein condensate (BEC) is presented, as well as experimental studies of two spin-state condensed and non-condensed samples. The first part of this thesis describes in detail our apparatus and experimental procedure for creating a BEC. We designed a system to create a BEC that is simple and robust enough to be constructed by someone outside of the field of atom cooling and trapping. Our system includes several novel features that reduce the complexity of the apparatus, which include mechanical transfer of atoms and a hybrid magnetic trap.

The second part of this thesis describes studies of two spin-state clouds. We studied the correlations and coherences of non-condensed clouds using Ramsey spectroscopy to gain an understanding of interatomic interactions. We were able to measure precisely the mean-field interactions between coherent particles, as well as determine mechanisms that preserve and destroy coherence. These experiments led us to the study of spin waves, in which scattering of indistinguishable particles gives rise to coherent spin oscillations. Using sensitive Ramsey spectroscopy, we were able to fully examine the spatio-temporal spin-state oscillations. Following the studies of normal cloud coherence, we went on to explore coherence effects in condensates.

Acknowledgements

First I would like to acknowledge my advisor Eric Cornell. He always amazes me with his physical insight and ability to be entranced by “shiny bits of physics,” which led us to a lot of great science. He has the right balance of knowing when to give advice and when to let us direct the experiment. I am grateful to have worked with such an inspiring physicist. I would also like to acknowledge Carl Wieman who has offered guidance throughout my graduate career.

I have had the privilege to work with many talented people. During his postdoc, Dwight Whitaker helped to construct the apparatus and always kept us going even during the two months of vacuum hell. I could not have asked for a better graduate student to have in the lab than Dave Harber. His unquenchable enthusiasm combined with his laid back attitude made working long hours in the lab truly fun. Another postdoc I worked with, Jeff McGuirk, brought his training in precision measurement to our lab. He really made me think about the experiment in a new way. He was also a great friend who kept me at least partially sane during the writing of this thesis. It was great to have so many other graduate students and postdocs to talk to and learn from over the years including Paul Haljan, Kurt Miller, Brian DeMarco, Neil Claussen, Ian Coddington, Peter Schwindt, Ying-Ju Wang, John Goldwin, David Hall, Brian Anderson, Stephan Duerr, Peter Engles, Shin Inouye and Tetsuo Kisimoto.

The construction of the apparatus would not have gone so quickly and smoothly if not for the JILA instrument and electronics staff. Hans Rohner, Hans Green, Blain

Horner, James Fung-a-Fat, and Terry Brown were always willing to help figure out a solution to the many problems I presented them.

Lastly I would like to thank my family for all of their support. My husband Don put up with a lot of late nights and was always supportive. My parents Kathleen and Dennis have always encouraged me through the years. I would not be writing this thesis if not for them.

Contents

Chapter

1	Introduction	1
1.1	Overview	3
1.2	Organization of this thesis	3
2	Experimental design and procedure	5
2.1	Experimental resources	8
2.2	Lasers	8
2.2.1	Frequency control and stabilization	9
2.2.2	Shutters	11
2.2.3	Spatial filtering	13
2.3	Vacuum system	14
2.3.1	Chamber design	14
2.3.2	Chamber construction	18
2.3.3	Chamber bake out	20
2.4	Magneto-optical trap	22
2.4.1	MOT alignment	24
2.4.2	MOT characterization	26
2.5	From MOT to Ioffe-Pritchard trap	27
2.5.1	Transfer into the quadrupole magnetic trap	27

2.5.2	Fluorescence imaging	30
2.5.3	Transfer from vapor cell to UHV region	33
2.5.4	Ioffe-Pritchard Magnetic Trap	35
2.5.5	Transfer between magnetic traps	40
2.6	Rf evaporation	41
2.6.1	Evaporation optimization	43
2.6.2	rf coil	46
2.7	Absorption imaging	46
2.7.1	Optical set up	47
2.7.2	Imaging in a non-uniform magnetic field	48
2.7.3	Producing and controlling microwaves	50
2.7.4	Imaging procedure	52
2.7.5	Focusing the image	56
2.7.6	Measuring the image magnification	57
2.8	Image analysis	57
2.8.1	Image processing	57
2.8.2	Image fitting	58
2.8.3	Calculating cloud parameters	59
2.9	Computer control	62
3	Driven two-level atom	65
3.1	Rabi flopping	65
3.2	Ramsey spectroscopy	66
3.3	Two-photon transition in ^{87}Rb	68
4	Spectroscopy	70
4.1	Spectroscopy in a magnetic trap	71
4.2	Cold-collision energy shift	75

4.3	Missing frequency chirp	79
4.4	Coherence of normal clouds	81
4.5	Coherence of condensates	86
4.6	Speculation on precision metrology	90
5	Spin waves	92
5.1	Theoretical description of spin waves	93
5.2	Description of system	96
5.3	Bloch-sphere representation	96
5.4	Experimental method for studying spin waves	97
5.4.1	Extracting longitudinal phase	97
5.4.2	Extracting transverse phase	98
5.5	Frequency of oscillation of spin waves	101
6	Coherence effects in a finite temperature condensate	104
6.1	Dressed states	105
6.2	Symmetries of the dressed-state system	107
6.2.1	Stability of dressed states	108
6.2.2	Separation of dressed states	111
6.2.3	Evaporation of dressed states	113
6.3	Decoherence induced melting of a finite temperature condensate	115
	Bibliography	118
	Appendix	
A	Everything you ever wanted to know about Rubidium	123

Tables

Table

2.1	Vacuum component temperature limits	21
2.2	Parameters for trapping and loading into a magnetic trap	29
2.3	Image timing	53
3.1	Breit-Rabi equation constants for ^{87}Rb	69
A.1	Properties of ^{87}Rb	123

Figures

Figure

2.1	Optical set-up for saturated absorption spectroscopy.	10
2.2	Saturated absorption spectra of the hyperfine structure on the Rubidium $5S_{1/2} \rightarrow 5P_{3/2}$ transition	11
2.3	Laser frequency control for the three lasers in our experiment.	12
2.4	Vacuum system layout.	15
2.5	Diagram of the MOT cell.	16
2.6	Getter assembly.	17
2.7	Schematic of the MOT optical layout.	23
2.8	Picture of the vertical-beam optics for the MOT.	24
2.9	Schematic of the MOT/quadrupole trap servo circuit.	25
2.10	Optical layout of fluorescence imaging.	31
2.11	Illustration of vignetting.	32
2.12	Illustration of the motion of the quadrupole trap coils.	34
2.13	Science cell region.	37
2.14	Ioffe-Pritchard magnetic trap (end on view).	38
2.15	Ioffe-Pritchard magnetic trap (side view).	38
2.16	Machine drawing of the Boron-nitride HIP-trap form.	39
2.17	Sample evaporation trajectory.	45
2.18	Optical system for absorption imaging.	48

2.19	Possible transitions with different probe beam polarizations	49
2.20	Microwave frequency chain.	51
2.21	Energy level diagram showing imaging transitions.	56
2.22	Computer control diagram.	62
3.1	Ramsey's method of separated oscillatory fields.	67
3.2	Energy level diagram for the ground state of ^{87}Rb	68
4.1	Spatial dependence of the transition frequency.	73
4.2	Differential Zeeman shift at low magnetic fields for the $ 1\rangle \rightarrow 2\rangle$ transition.	74
4.3	Measurement of the cold collision shift.	77
4.4	Testing the exchange symmetry between the $ 1\rangle$ and $ 2\rangle$ states.	79
4.5	A data set of Ramsey fringes probing for frequency shifts as a function of coherence.	80
4.6	Coherence times in a normal cloud.	82
4.7	The frequency shift across the axial direction of the cloud.	84
4.8	The transition frequencies at different positions along the cloud are plot- ted as a function of B_{bias}	85
4.9	Ramsey fringes of the center bin of a condensate.	88
4.10	Alternate method of taking coherence data.	89
5.1	Intuitive picture of spin rotation during collisions.	95
5.2	Creation of spin waves	98
5.3	False color plots of the longitudinal and transverse components of the Bloch vectors.	99
5.4	Bloch vector trajectories across the cloud.	100
5.5	Frequency and damping of spin waves as a function of density.	102
6.1	Energy level diagram of the dressed states showing the avoided crossing.	106

6.2	Transfer rate is shown as a function of Rabi frequency.	110
6.3	Component separation in a dressed state system.	112
6.4	Condensate Rabi flopping after evaporation in the dressed states.	114
6.5	Measured phase space density as a function of time.	116
A.1	Energy level diagram for ^{87}Rb	124
A.2	Branching ratios for ^{87}Rb	125

Chapter 1

Introduction

The creation of the field of ultracold atomic physics has opened up many new possibilities for studying the effects of quantum mechanics in theoretically accessible systems. Ultracold temperatures reduce the complexity of the system by reducing the possible states a system can occupy, thus allowing a more complete understanding of the system. An extreme case of a sample of ultracold atoms is a Bose-Einstein condensate (BEC), in which all of the atoms occupy a single quantum state.

The theory of Bose-Einstein condensation was first suggested by Bose [1] and expanded upon by Einstein [2] more than 75 years ago. However it was not until 1995 that a BEC was first produced in a dilute vapor. The time lag in the experimental realization of the predicted phase transition is not so incredible, once one realizes the extraordinary conditions that must occur in order to Bose condense a gas.

In everyday life, quantum statistics does not play a major role in our experiences because the thermal wavelength of particles is small compared to the interparticle spacing. However if a particle's wavelength is made comparable to the interparticle spacing, quantum statistics will become important. The relevant parameter that determines whether particles obey classical or quantum statistics is the phase-space density, which is the number density times the thermal de Broglie wavelength cubed. As this parameter approaches one, quantum statistics becomes relevant. For Rubidium gas at room temperature and typical vacuum density of 10^{-9} torr, the phase space density of the

sample is $\sim 10^{-20}$, which is twenty orders of magnitude away from the Bose-Einstein condensation transition. It was predicted that one would have to cool Rubidium vapor below 1 millionth of a degree above absolute zero to observe a BEC. Remarkably, with the advent of laser and evaporative cooling, this goal turned out to be obtainable. Now over 35 groups in the world have condensed eight different atomic species. This accomplishment was recently recognized by the awarding of the 2001 Nobel prize in physics “for the achievement of Bose-Einstein condensation in dilute gases of alkali atoms, and for early fundamental studies of the properties of the condensates” [3].

It is somewhat ironic that, after finally achieving the goal of producing a condensate, the studies in our group concentrated on the most basic system in quantum mechanics: the two-level system. We wanted to study the macroscopic quantum phase of a condensate, and how this phase is affected by both internal and external couplings. One cannot measure the absolute phase of a particle, but rather only the relative phase with respect to another particle. Therefore if we want to study the phase of a condensate, we must have another condensate as a phase reference. We meet this requirement by producing a condensate in two distinct spin states, which we can manipulate and measure by standard atomic-physics techniques. The ultimate goal was to study how a normal cloud interacting with a condensate affects the phase of the condensate. Before we began our condensate coherence studies, we felt we should understand fully normal cloud coherences and correlations. This initial study led to many significant results including the unexpected observation of spin waves and precision measurements of interatomic collisions and correlations. These studies are directly relevant to the field of atomic clocks, where interatomic collisions currently limit the accuracy. After an extensive study of normal clouds in two spin states, we began studies of pure-condensate coherence and the effects of the normal cloud on condensate coherence. The studies of condensate coherence are still ongoing with the hope that we may be able to understand the coupling between the normal cloud and the condensate. It is a pleasure

to note that a two-level system can still yield surprises, 75 years after the advent of quantum mechanics.

1.1 Overview

We originally set out with two goals in mind. First, we wanted to design and build a new type of BEC apparatus, which would be simpler and more robust than was currently in use in dozens of labs around the world. Second, we wanted to use our stable apparatus to study phase coherence of condensates.

In the seven years since their first observation, dilute vapor condensates have been studied extensively. In most cases the condensate properties themselves are the focus of the investigations. Relatively little work has been done using a condensate as a tool to explore questions in other fields. We feel that physicists in other fields such as condensed matter have a different and unique perspective on possible experiments that could make use of condensates. Therefore, developing a system that could be used for these purposes is worthwhile. The current experimental systems [4, 5, 6] were designed by people with a tremendous amount of knowledge and experience in experimental atomic physics, and until now producing a BEC without expertise in ultracold atom trapping has been a daunting task. We felt, however, that with some modifications to the current experimental design, any experimental physicist, regardless of discipline, could produce a BEC in their lab.

1.2 Organization of this thesis

This thesis is organized as follows. The first part of this thesis (Chapter 2) will describe most of the basic steps in detail on how to build a BEC apparatus. Many of the techniques described have been developed by others [7, 8, 9] over the last 20 years but are included so that this thesis may serve as a ‘recipe’ for creating a BEC. Besides the traditional methods we also describe several new features of our design, which include

a new method of atom transfer and a hybrid magnetic trap.

The second part of the thesis (Chapters 3-6) is devoted to the scientific studies of two spin-state ultracold and condensed atom samples. We started our study of two-component samples by spectroscopically measuring cold-collision properties, which include mean-field and Zeeman energy shifts, as well as correlations between the spin states. The basic theory behind the spectroscopy is described in Chapter 3, and the results of these studies are described in Chapter 4. During our spectroscopic measurements, we observed by accident an effect we initially called anomalous spin-state segregation. We observed the spin states separating spatially in the trap in an unexpected way. Eventually, with the help of several theorists, we understood the spin dynamics to be spin waves. Chapter 5 discusses the origin of and our extensive measurements of spin waves produced in our two-component cloud. Armed with the knowledge of coherences and correlations gained from our experiments of cold collisions and spin waves, we studied how loss of coherence affected two-component condensates. Chapter 6 describes our results on coherence effects on finite temperature condensates.

Chapter 2

Experimental design and procedure

A large fraction of the effort represented in this thesis was the development of a new and improved system for creating a BEC. This system was developed from scratch, with little more than a bare optical table as a starting point. Several of the components developed for this apparatus have proven so useful that several groups at JILA, and at universities throughout the United States, are now incorporating them into their experimental systems.

We recommend several sources that contain background information on a variety of subjects not covered in depth in this text. For a textbook description of atom cooling and trapping we recommend *Laser cooling and trapping* by Metcalf and van der Straten [10]. Two useful papers that describe diode lasers and saturated absorption spectroscopy are “Using diode lasers for atomic physics” [11] and “A narrow-band tunable diode laser system with grating feedback and a saturated absorption spectrometer for Cs and Rb” [12]. A third paper by Wieman and coworkers describes the components and the process for constructing a magneto-optical trap (MOT) for use in undergraduate laboratories [13]. A review of many ideas in evaporative cooling may be found in van Druten and Ketterle [14]. There is a good section on imaging cold clouds in “Making, probing, and understanding Bose-Einstein condensates” by Ketterle and colleagues [15].

Our design has many features that allow a condensate to be made even if the system is not particularly optimized. We use a MOT with large beams (5 cm diameter)

with a considerable amount of laser power (400mW) to collect a large number of atoms. The large number of atoms in our MOT is efficiently and reproducibly translated 45 cm to a final magnetic trap via a mechanical transfer mechanism. We are able to place the magnetic trapping magnets and coils very close to our atoms thus producing very strong confinement and high collision rate. All of these features insure that evaporation will work well and produce a condensate consistently.

One of the main concerns when designing any BEC apparatus is the need to optically collect many atoms and yet to have a long lifetime for the atoms in the magnetic trap. These two constraints require orders of magnitude different vapor pressures. It is not easy to change the vapor pressure in a vacuum system by two orders of magnitude in a reasonable time of seconds. There are two conventional ways of solving this problem: a double MOT apparatus [16] and an atomic beam [17]. The double MOT system collects atoms in a MOT in one region of the vacuum system, which has a high alkali vapor pressure, and transfers the pre-cooled atoms to a second MOT in an UHV chamber, with a pressure in the low 10^{-11} torr, for further cooling. The transfer is done essentially by pushing the atoms between MOTs with a laser beam. There are several disadvantages to this method. Two MOTs are necessary, which requires the system to use more laser power than a single MOT system. Also the optics and optical access needed for the second MOT restrict the space for the magnetic trap coils, so that to create a specified magnetic gradient requires considerable electrical power. Finally, the push beam and MOTs are very sensitive to optical alignment making optimum performance difficult to sustain. The other main method is to use a laser-cooled atomic beam from a Zeeman slower. Zeeman slowers are large (1-2 m) and require a high temperature oven. Systems with Zeeman slowers have only one MOT, but they suffer from the same optical access problems around the magnetic trap as does the double MOT design.

A design similar to ours, developed by the Hänsch group in Garching [18], uses a series of electromagnetic coils to move the atoms from one chamber to the other. The

atoms are initially transferred from the MOT into a magnetic quadrupole trap; then by ramping the current in successive sets of quadrupole coils lined along the transfer tube, the atoms are transferred between chambers. This design has the benefit of being able to move the atoms easily around a corner to reduce line-of-sight between the chambers, which reduces trap loss from background collisions. However, the ten overlapping sets of quadrupole coils require a large amount of power to run, take up considerable space around the apparatus, and require effort to design, construct, and optimize.

Our system uses a moving magnetic trap to transfer atoms between regions of the vacuum system. Magnetic coils are mounted on a linear track and translated from one section of the system to the other; this technique has many advantages. It is very easy to use. The track, motor, and controller are commercially available as a turnkey system. The transfer worked the first time and works essentially every time without any maintenance. As with the Garching group's design, our system does not need a second MOT in the UHV region, so we can place our magnetic trap close to the atoms and produce very strong confining fields.

The other new feature in our system is a hybrid magnetic trap. It uses strong permanent magnets to produce radial confining fields and low power electromagnetic coils to produce axial confinement and a bias field. The permanent magnets do not consume power and thus do not need to be actively cooled, as would electromagnetic coils producing the same field. The tight confinement from the permanent magnets insures that we have the necessary collision rate to evaporatively cool the atoms. Permanent magnets are however plagued by long-term stability problems. The longitudinal bias field, the only parameter for which stability is critical, is produced by servo-controlled electromagnetic coils, which are air cooled. Our design allows tight confinement from the permanent magnets and bias-field stability from the coils. The magnetic trap is designed to have tight enough confinement so that we can produce a condensate even if the system is not well optimized.

We will outline in detail the steps required to make a Bose-Einstein condensate using our particular experimental system. This Chapter begins by giving directions for setting up the necessary vacuum and optical systems. Next we describe the experimental procedure for pre-cooling atoms in a MOT, transferring the atoms to a magnetic trap, and evaporatively cooling them to create a condensate. There is also an extensive Section dedicated to imaging the condensate and extracting useful parameters from the images.

2.1 Experimental resources

Our experimental apparatus does not require a large amount of space compared to other BEC systems, particularly those with Zeeman slowers. Our lab contains three components: an optical table, electronics, and control computers. We have the vacuum system, lasers, and all of the optics on one 4'×8' optical table. The system is not congested, and we still have several square feet of free space on the table. In addition to the optical table, we have a canopy the size of the table, which is half full of electronics, and one full electronics rack. We also have two computers in the lab, which control the experiment.

2.2 Lasers

The experiment requires three different wavelengths of laser light: two to make the magneto-optical trap and one to image the condensate. We use a different laser for each of these tasks. We use a commercial external cavity diode laser (ECDL) [12], which is amplified by a single pass through a tapered amplifier chip for the MOT trapping beams, in a master-oscillator power-amplifier (MOPA) configuration. This system will nominally produce 500 mW of power at 780 nm. The MOT hyperfine repump and probe beams are supplied by two separate ECDLs, which each produce ~ 8 mW.

2.2.1 Frequency control and stabilization

All three lasers are locked to atomic transitions in ^{87}Rb using saturated absorption spectroscopy [12]. This type of frequency stabilization is standard practice and thus will not be described here in detail. The basic idea is to produce sub-Doppler spectral lines, which can be used as feedback to stabilize the laser. The optical layout for saturated absorption spectroscopy is shown in Fig. 2.1.

We lock each laser to the peak of an atomic transition. The frequency location of the peak of the transition is relatively insensitive to intensity and broadening effects, which would change the locking set point if the laser were locked instead to the side of the line. Unfortunately a servo can only lock to a region where there is a slope of the line to use as feedback. We solve this problem by generating a derivative of the saturated absorption signal. We modulate the frequency of the laser, by *either* modulating the laser current or the radio frequency (rf) driving an acoustic optic modulator (AOM). The modulation has a depth of 5 MHz at a rate of 50 kHz, which is slow enough for the AOMs to respond and fast enough to be above the bandwidth of the servo. The signal from the saturated absorption spectrometer is routed to a homemade lock-in detector which gives the derivative of the original transition lines. The derivative changes sign at the peak, and thus is a convenient error signal for our servo.

We use AOMs to shift the frequency of the light that is used in the experiment. The AOMs are driven by amplified voltage-controlled oscillators. A diagram of the optical set-up for the AOMs is shown in Fig. 2.3. A portion of the light from the MOPA is sent through a 120 MHz AOM, which is frequency modulated. The negative first order diffracted beam is used in the saturated absorption spectrometer. This scheme allows us to lock the laser to the peak of the $(2 \rightarrow 2', 2 \rightarrow 3')$ crossover line (peak A in Fig. 2.2) and have the trapping light red detuned by several natural linewidths from the $2 \rightarrow 3'$ cooling transition. The unprimed states refer to the $5S_{1/2}$ manifold and the




Figure 2.1: Optical set-up for saturated absorption spectroscopy. A small amount of light (0.1-1 mW) is split off from the laser and sent through a half inch thick piece of Lucite to produce two weak beams. The two reflected beams, about 1-2 mm wide, pass through the Rb vapor cell and are focused onto two photodiodes. The saturating beam that passes through the lucite is sent into the vapor cell counter-propagating with one of the weak beams.

primed states refer to the $5P_{3/2}$ manifold. The repump laser is locked directly to the $1 \rightarrow 2'$ transition. We modulate the current of the repump laser instead of the frequency of an AOM. The probe beam is sent through two AOMs in our design. One AOM is essentially a fast shutter, while the other allows us to modulate the frequency sent to the saturated absorption and not impart the frequency modulation onto the probe beam itself. We need to shift the frequency of the probe beam only a few MHz from the lock frequency. Therefore the two AOMs are set to almost the same frequency, and their center frequency of 260 MHz is arbitrary. We lock the probe laser to the peak of the $(2 \rightarrow 2', 2 \rightarrow 3')$ crossover transition.

Vibration isolation is important when operating ECDLs, in our case commercial New Focus Vortex lasers. Vibration can cause undesirable frequency noise at a level that the feedback may not be able to fully correct. We mount the repump laser on a piece of 1/4" thick sorbathane sheeting to reduce the effect of vibrations from the table. The probe beam laser is mounted directly on the table. Mounting it on sorbathane could cause the beam pointing to drift and thus decrease the coupling into the fiber. We also mount our shutters on the same material, so that vibrations induced by the




Figure 2.2: Saturated absorption spectra of the hyperfine structure on the Rubidium $5S_{1/2} \rightarrow 5P_{3/2}$ transition. Widths and relative heights of the peaks are affected by beam alignment, intensity, and polarization. These are only representative signals.

solenoids when they open or close is not transmitted to the optical table.

2.2.2 Shutters

We use both mechanical shutters and AOMs to control the timing of the laser light. The mechanical shutters provide an excellent extinction ratio but are slow and can have jitter in their timing. We can also use an AOM as a shutter by using a diffracted order. AOMs are fast, with response times on the order of a $1 \mu\text{s}$, but their extinction ratio is less impressive. For critical light pulses such as the probe pulse, we use both

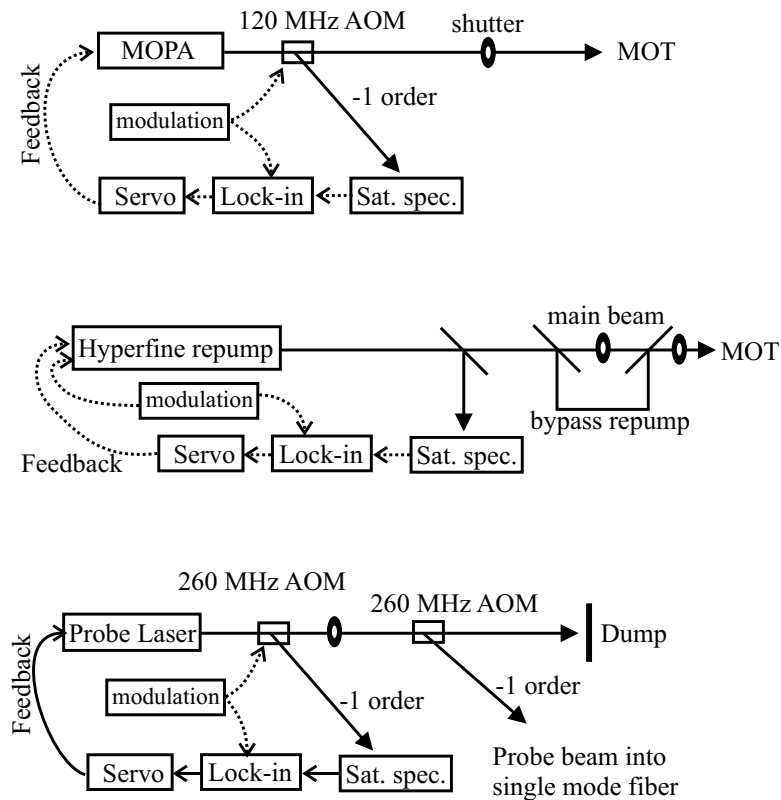


Figure 2.3: Laser frequency control for the three lasers in our experiment. For the less critical repump laser, the frequency modulation for the lock-in detection is applied to the laser frequency itself through current modulation. For the more critical trapping and probing beams, the frequency modulation is applied to the AOMs, and thus the frequency modulation is not on the light used on the atoms.

types of shutters. We create a pulse by first opening the mechanical shutter several ms before the start of the pulse, turning on the AOM for the desired pulse length, then closing the mechanical shutter several ms later. Mechanical shutters can cause noise on the lasers when they open or close through acoustical vibrations. We are careful to open the shutter early enough for the laser servo to recover from the acoustical noise before triggering the AOM for the probe pulse.

2.2.3 Spatial filtering

The spatial beam profile from diode lasers is not Gaussian and may contain high contrast stripes, which must be reduced before the light is used in the experiment. The output of a MOPA system can be nonuniform depending on the amplifier chip. A MOT does not require perfectly uniform beams but high contrast features will lower the trapping efficiency. We focus the MOPA beam, $f\#$ 16, through a large pinhole (50 - 100 μm), eliminating any high-frequency spatial modes while retaining most of the laser power ($\sim 75\%$). The MOT is largely insensitive to the spatial profile of the repump beam, so we do not spatially filter this light. The probe beam, on the other hand, must have a very uniform and flat intensity profile. We use a single-mode angle-polished polarization-preserving fiber to filter spatially the probe beam. It is critical to use both an angle-polished and polarization-preserving fiber to reduce temporal intensity fluctuations of the beam. The facets of a flat-polished fiber can form an etalon, which will produce high-frequency intensity fluctuations on the output. A non-polarization-preserving fiber will scramble the input polarization depending on the stress (thermal or mechanical) applied to the length of the fiber. These polarization fluctuations will be converted into intensity fluctuations when the light passes through a polarizer. One good method to align the input light polarization with the axis of the fiber is to first adjust the angle of the initial linear polarization until pure linear light is emitted from the fiber. This procedure may be too coarse to accurately align the axis, so the next step is to tap on or warm up with your hand the center of the fiber, so as to not change the coupling at the ends, and to watch the intensity fluctuations on the output after a polarizer. One can now more finely adjust the input polarization until a minimum of polarization fluctuations is observed on the output.

2.3 Vacuum system

2.3.1 Chamber design

The vacuum chamber is a simple single ion pump system that has a high vacuum vapor cell ($10^{-8} - 10^{-9}$ torr) and an ultra-high vacuum (UHV)(10^{-11} torr) science cell. Three vacuum pumps are used in the system, but only one of them is used on a continual basis. A diagram of the chamber is shown in Fig. 2.4. The turbo pump, connected to the system by an all-metal valve, is used only during initial pump down and bake out. The Ti-sublimation pump is turned on only every couple of years to remove extra Rb or H from the system. The workhorse pump is the 40 l/s ion pump, which pumps continuously on the chamber during normal operations. The ion pump is conduction limited for some atomic species and thus it may be larger than needed. However we would recommend the 40 l/s pump to ensure the system has adequate pumping speed. The valve on the ion pump was used only as a diagnostic tool during original testing of the system and would be removed if the system was reconstructed. The valve used to close off the turbo pump is necessary to remove the pump from the system after the bake out. The sealing surface of the valves and not the bellows should always face the vacuum side; this configuration reduces the surface area in the UHV chamber.

One thing to consider when deciding on Ti-sublimation pump placement is where the Titanium will deposit. The Ti-sublimation pump's filaments should be placed so that there is not a direct line of sight to any valve or pump. The Ti will coat all surfaces in direct line with the filaments and can cause a valve sealing or pumping problem. See Fig. 2.4 for position of the filaments in our system.

The differential pumping between the two chambers is maintained by placing a small aperture on each side of the bellows to reduce the conductance. The apertures are 5 mm diameter holes in the solid copper gaskets connecting the bellow to the system. With our design the pressure differential between the two chambers is about a factor




Figure 2.4: Vacuum system layout. The system is suspended with a series of clamps (not shown), so the center line of all the horizontal tubing is 18 cm above the optical table.

of 17 for Rb. A second ion pump could be added between the two cells to increase the pressure differential and allow a shorter MOT loading time without sacrificing science cell lifetime.

The conductance of gas through a tube and an aperture in the molecular-flow regime is

$$C_{tube} = \frac{3.81D^3}{L} \sqrt{T/M} \text{ liters sec}^{-1} \quad (2.1)$$

$$C_{aperture} = 3.7A \sqrt{T/M} \text{ liters sec}^{-1}, \quad (2.2)$$

where D and L are the diameter and length of the tube in cm [19]. A is the cross-sectional area of the aperture in cm^2 , T is the temperature in Kelvin, and M is the molecular weight of the gas. The conductance of a system can be found by adding the conductance of each individual part like capacitors in an electrical circuit as in the following: ($C_{parallel} = C_1 + C_2 + \dots$, $1/C_{series} = 1/C_1 + 1/C_2 + \dots$).

The science and MOT cells are cylindrical glass cells attached to glass-to-metal seals. The science cell is a 10 cm length of 1.4 cm outside diameter pyrex tubing with a window on one end and a glass-to-metal seal on the other. The MOT cell is a 20 cm length of 5 cm outside diameter pyrex tubing necked down on each end. One end is attached to a glass-to-metal seal, while a getter assembly is fused to the other end (Fig. 2.5).




Figure 2.5: Diagram of the MOT cell showing the getter assembly and glass-to-metal seal welded onto a knife-edge seal flange.

The getter assembly consists of a current feedthrough and two Rb dispensers (getters). The feedthrough, called a pin press, is a glass section with Tungsten pins inserted. The Rb getters are spot welded to the pins of the pin press. The Rb getters are a controllable source of Rb vapor. A getter is a small foil container of a Rb salt, which releases Rb when a moderate current of 2-6 A is run through the device. It is important to use a current limited power supply for the getters, because the resistance of the getters changes drastically as the temperature increases. The getter assembly can be seen in Fig. 2.6.

Special care must be taken with the getters to insure they will produce clean




Figure 2.6: Getter assembly showing two getters and reentrant glass with pins inserted.

Rb vapor. First, the getter material can easily absorb water, so we keep them under vacuum with desiccant during storage and flow only dry gas during the glass fusing process. Second, they release Rb as a double exponential function of temperature. Thus, we make sure that while the pin press is being fused to the cell, the getters are not heated significantly by the fusing torch. Avoiding moisture and heat, we are usually able to have getters that produce clean Rb vapor. However occasionally the getters will produce other products as well as Rb. In our current system when we turn on the getters the number of atoms in the MOT decreases, presumably due to contaminants being released from the getter. Our mode of operation is to turn on the getters at 3.5 A for 10 minutes to supply the MOT cell with a day's worth of Rb, and then to allow 10-20 min for the contaminants to pump out of the system before taking data.

A good question to ask is why we chose to use cylindrical glass cells instead of square cross-section cells. The main reason is it is easier to make cylindrical cells. Anyone with a small amount of glassblowing training can fabricate the cylindrical cells, whereas constructing square cells usually requires significantly more equipment and expertise. Another advantage to using a cylindrical MOT cell is the interference fringes

on the trapping beams have higher spatial frequency and are spaced less regularly than with a square cell. One can shake the glass MOT cell around by a 1 cm and see little movement of the MOT cloud. MOT alignment with a square cell can be more difficult because it is important to place the minimum of the fringes away from the center of the trap. This type of alignment requires more frequent adjustments and is more a trial-and-error process than simply overlapping the beams at the correct angles, as is required for a cylindrical cell. A cylindrical cell distorts the trapping beam, but this does not seem to be a large problem for us.

The cylindrical *science* cell, on the other hand, is less desirable. The probe beam is focused by the cell, which acts as two cylindrical lenses. This is generally not a problem for absorption imaging, but can be for phase contrast imaging. Phase contrast imaging requires placing a material in the fourier plane to shift the phase of the light. The large astigmatism induced by the cell requires placing a phase shifting line rather than a dot, thus making alignment more difficult. If we were to reconstruct our experiment, we would replace the cylindrical science cell with a commercial square cell, which is easy to obtain from companies who specialize in producing spectroscopic cells.

2.3.2 Chamber construction

Obtaining UHV pressures requires careful assembly of the vacuum components. The most important thing is to make sure all of the components are clean. We start the cleaning process by placing the submersible parts (no valves, pumps, or cells) in an ultrasonic cleaner with strong soap for 1 hour. If the valves are cleaned in the ultrasonic cleaner they must be re-greased before they are used. The long cleaning time allows the strong soap to remove residual oil from the factory. Typically when a vacuum part is baked in air it will become a golden color, which we assume is residual burnt factory oil. A one hour bath in a strong basic soap will remove this coating, and it will not return with subsequent air bakes. The long bath is not absolutely necessary and may

be reduced to a few minutes to just remove any particulates from the parts. After the ultrasonic bath the parts are rinsed with deionized water, acetone, and then methanol. Next the parts are baked in air for 4 hours at 400° C to drive off any residual material still on the parts. Once the parts have cooled they are wrapped in *oil-free* aluminum foil until assembly.

It is important to avoid contamination of the vacuum system during assembly. We always wear powder-free latex gloves and change them often. All knife edges and copper gaskets are wiped with ultra-pure methanol before installation to remove any factory residue.

We use silvered bolts on the knife edge flanges to reduce the possibility of bolts seizing in the flanges during the bake out. If silvered bolts are not available we place some molybdenum disulfide powder suspended in methanol on the threads of the bolts for lubrication. Suspending the molybdenum in methanol reduces the chance that it will fall into the vacuum system during assembly and become a contaminant.

After the entire vacuum system has been assembled it is pumped out and checked for leaks. We use a small turbo pump backed by a dry, four-stage diaphragm pump to initially pump out the chamber. We use a diaphragm pump rather than a standard oil-filled roughing pump, because the diaphragm pump does not contain any oil, which could leak into the system. Once the turbo pump has spun up to full speed we spray a small amount of ultra-pure methanol on all the flanges and cells. If there is a large leak, the pressure in the tubing connecting the turbo pump to the diaphragm pump, read by a thermocouple gauge, will change when methanol is applied. The system is pumped on overnight before the bake out is started. A carefully cleaned system, which does not contain any large leaks, should pump out overnight with the pressure reaching around several 10^{-8} torr, read from the ion pump.

2.3.3 Chamber bake out

The vacuum system must be baked at high temperatures under vacuum to remove contaminants to obtain UHV pressures. We bake most of the chamber at 300° C for several days. Before the bake out, we run about 6 A through the getters for 30 seconds to verify the presence of Rb, which can be seen by either laser absorption or fluorescence. The high current also degasses the getters.

The system can now be prepared for the bake. The first step in the bake out process is to wrap the glass cells with clean fiberglass cloth. The cloth will protect the cells from anything that may melt onto the cells during the bake. Next we place Teflon-coated thermocouples on the chamber at critical places such as the cells, glass-to-metal seals, and pumps. We then wrap the system with high-temperature resistive heater tape. It is not important to wrap the entire surface of the chamber with tape, but rather to have a constant tape-to-chamber surface area ratio. Heater tape is applied to different objects proportional to their surface area and not the mass of the object. The mass only defines the time constant for thermal equilibration, whereas the ultimate temperature is determined by the heat flow in and out of the region, which is proportional to the surface area. The tape should never overlap itself, or the intense heat will cause the tape to burn. Several short tapes are used to wrap the system so each section may be controlled independently. The turbo pump is not baked because it is not part of the final chamber and can not handle high temperatures. We do however bake the entire ion pump with the magnets in place. Typical maximum baking temperatures for different components are listed in Table 2.3.3. After the tapes are in place the chamber is wrapped with strips of fiberglass and then several loose layers of aluminum foil to insulate the system from the environment.

The system is slowly brought up to the final temperature over 6-8 hours. The ion

Vacuum element	Maximum baking temperature ($^{\circ}\text{C}$)
turbo pump inlet flange	120
ion pump magnets	250
ion pump body	400
ion pump cable	250
bakeable valve, open	450
bakeable valve, closed	300
Ti-sublimation pump	350
glass/metal seals	300

Table 2.1: Vacuum component temperature limits

pump is still off at this point. There is a large amount of material driven off the walls of the chamber during the initial several hours of the bake. We prefer to have the turbo pump, which has an almost infinite pumping lifetime, pump out initially rather than the finite lifetime ion pump. Each heater tape is powered by a variable AC transformer (variac) which controls the temperature of each section of the chamber. Generally we increase the temperature by at most 50°C per hour. Temperature gradients can apply significant stress to the system. We prefer to keep the temperature gradients to under 30°C across the glass cells and glass-to-metal seals, which are the most susceptible components to failure. Caution must be taken when the system is near 300°C because some parts of the system could overshoot in temperature due to long thermal time constants. During the bake 3 A are run continuously through the getters to clean them. If the getters are left off they will be the coldest part of the system, and contaminants will accumulate on them. We also run 25 A through the Ti-sublimation filaments during most of the bake. Throughout the warming up process, thermocouple readings, variac settings, and pressure readings are recorded to facilitate future bakes. Once the system is at the desired temperature we bake with just the turbo pump on for 12 hours. At this point we degas the getters and Ti-sublimation filaments. To degas the getters, we increase the current in each getter for 30 seconds to 5 A to drive off any surface contaminants. After the first 12 hours we turn on the ion pump and valve off the still-

running turbo pump. When the valve is above room temperature we close it only finger tight. Therefore, we do not allow the turbo pump to spin down until the system is back to room temperature, and the valve can be properly closed with a torque wrench. We allow the system to bake with the ion pump on for ~ 2 days, or until the pressure on the ion pump reads in the middle 10^{-9} torr. We cool the system down slowly over 4-6 hours. At this point the ion pump should read the lowest possible pressure, which is 10^{-10} torr for most pumps, if there is no leakage current. The ultimate test of the vacuum pressure is the lifetime of the atoms in a magnetic trap.

2.4 Magneto-optical trap

We have a large six-beam MOT. We are forced to use six beams and not retro reflect the beams because the cloud of atoms casts a shadow and would cause an intensity imbalance between the beams. Also the cylindrical cell distorts the beams as they pass through both sides of the cell. The diameter D , of the beams is approximately 5 cm. Large beams greatly increase the number of atoms in a MOT because the loading rate scales theoretically as $D^{\sim 4}$ for intensities even as low as half saturation [20].

The optical layout is shown in Fig. 2.7 with the vertical beam paths shown in Fig. 2.8. We use 5 cm diameter polarizing beamsplitting cubes and waveplates and 7.5 cm mirrors to accommodate the large beams. The repump beam enters the system via the backside of a cube; therefore the polarization of the repump beam will not be optimum when it enters the trapping region. This is not a problem because a Rb MOT requires very little repump power.

The MOT coils, which also serve as the quadrupole trap coils, are made of 24 turns of square hollow copper tubing coated with Kapton. The wire has a square cross-section of 4.15 mm on a side with a round 2.5 mm diameter hole in the center. The coils are cooled by running water through the center region of the wire. The wire is wound onto a phenolic spool and secured with epoxy. Phenolic was chosen as the




Figure 2.7: Schematic of the MOT optical layout. Some of the mirrors and waveplates for the vertical beams are not shown.

spool material because it will not support eddy currents when the current is abruptly changed in the coils. The inner diameter of the coils is 5 cm, and their centers are separated by 10 cm. The current in the coils, run in series, is controlled by a simple servo circuit (Fig. 2.9), which controls three power MOSFETs. We use a 580 Amp, 8 V switching power supply, run in voltage-controlled mode, to supply current to the MOT coils. The coil configuration produces a magnetic field gradient of 1 Gauss/cm/Amp along the axis of the coils.




Figure 2.8: Picture of the vertical-beam optics for the MOT omitted from Fig. 2.7. Figure shows the location of the track (blue arrow) and the vertical MOT beams (red arrows). The coils of the MOT/quadrupole trap are translated left, towards the science cell, out of the field of view of this photograph.

2.4.1 MOT alignment

The alignment of our MOT is not as sensitive as it would be for a MOT with 1" beams. We start aligning the MOT by placing an iris in the beam from the MOPA

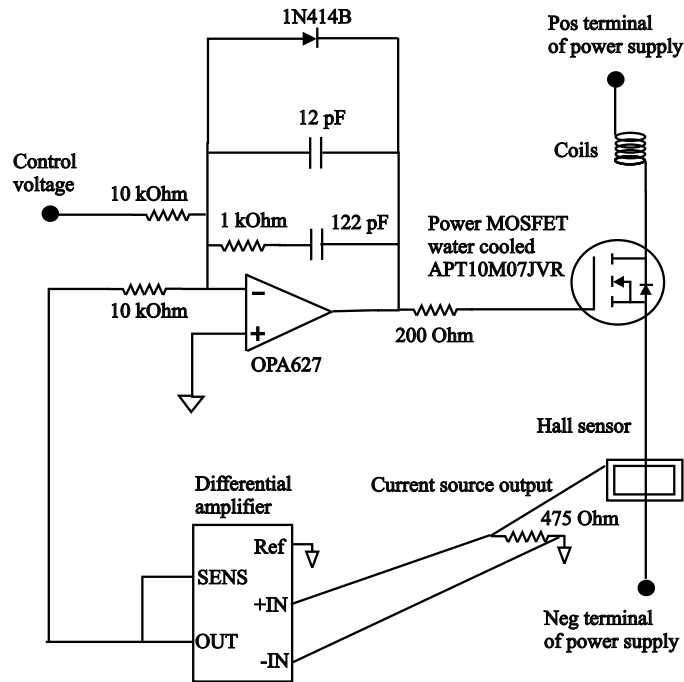


Figure 2.9: Schematic of the MOT/Quadrupole trap servo circuit. The circuit has standard PI loop gain, where the proportional gain comes from the feedback from the closed-loop Hall current sensor. We use 2/0 gauge welding cable to carry 300 A from the power supply to the coils and MOSFETs. The three MOSFETs are mounted on a water cooled copper plate to remove the heat generated by the MOSFETs.

before it is split into six. Closing the aperture allows us to align the centers of the beams. Once we get all the beams roughly aligned with respect to the magnetic coils and each other, we balance the power in the beams. We have $\lambda/2$ plates mounted on rotating mounts before every polarizing beamsplitting cube to adjust the power in each beam. We measure the power just before the beams enter the MOT cell and adjust the waveplates until the power is equal in each of the six beams to better than 10%. Next we open the iris and attempt to see a trapped cloud. A simple security camera can image the fluorescence from the cloud, which can be viewed on a monitor. It may take several hours with the getters running for the cell to be coated with Rb and a significant Rb vapor pressure to be established. Once a cloud is visible, an easy way to adjust finely the alignment is to reduce the beam size, optimize the cloud for number and roundness,

and repeat, all the time making sure the beam pairs are kept counterpropagating.

The position of the trapped cloud should be centered with the quadrupole magnetic trap to minimize energy gained by the cloud when it is transferred into the magnetic trap. To test this, we increase the current in the MOT coils until the cloud size is reduced greatly by the large magnetic field-induced detuning. The position of the small cloud, in a very large magnetic field gradient, is reliably at the center of the magnetic field, and thus the center of the magnetic trap. We then decrease the field, and adjust the the beam alignment and power balance until the cloud center is in the same location at high and low magnetic fields.

Another test for quality of the alignment is to look at optical molasses [21]. We quickly turn off the magnetic field of the MOT and look at the expanding cloud. If it moves rapidly in one direction this could be a sign of beam imbalance (from incorrect polarization or intensity splitting between the beam), poor alignment, or stray magnetic fields. The goal is have slow, spatially-uniform expansion during optical molasses. In traditional systems, it is necessary to use shim coils to cancel ambient magnetic fields and to have the atoms in the optical molasses expand uniformly. A feature of our system is the lack of need for shim coils. We transfer atoms to our quadrupole trap at a relatively high temperature, where the small reduction in energy from the shim coils would not make a great improvement in the phase space density of the magnetically trapped cloud.

2.4.2 MOT characterization

We calculate the number of atoms in the MOT by measuring fluorescence from the trapped cloud on a photodiode. We focus the fluorescence onto the photodiode with a single lens. The number of atoms in the MOT is

$$N = \frac{4\pi(\text{photodiode current})}{(\text{solid angle})(\text{responsivity})(\text{energy of a photon})(\text{R})(0.96)^k}, \quad (2.3)$$

where solid angle refers to the solid angle subtended by the collection lens, the responsivity refers to the current produced for a given power incident on the photodiode, and k is the number of uncoated glass surfaces between the atoms and the detector. R the photon scattering rate in photons/sec/atom, is

$$R = \frac{\frac{I}{I_s} \pi \Gamma}{1 + \frac{I}{I_s} + 4\left(\frac{\Delta}{\Gamma}\right)^2}, \quad (2.4)$$

where I is the total intensity of the six beams impinging on the atoms, I_s is the saturation intensity, which is 4.1 mW/cm² for random polarization for Rb. Γ is the natural linewidth of 6 MHz for Rb, and Δ is the detuning from resonance. In our experience, using the I_s appropriate for random polarization gives the most accurate number of atoms in a MOT.

We servo the MOT load to increase reproducibility in condensate number. We could load the MOT for a fixed time each shot. However, if the pressure of Rb changes throughout the day, the loading rate, and thus the total number loaded into the MOT will drift with time. We instead measure the voltage output from a photodiode collecting light from the MOT; when a particular value is reached we stop the loading and proceed to transfer atoms to the quadrupole trap. This also allows us to vary the resulting number of atoms in the final cloud by adjusting the MOT load level.

2.5 From MOT to Ioffe-Pritchard trap

2.5.1 Transfer into the quadrupole magnetic trap

There are three main steps to transferring atoms from a MOT into a magnetic trap: compressed MOT (CMOT) [22], optical pumping, and turn on of the magnetic trap. Our goal is to transfer the atoms into the quadrupole trap with the highest phase space density. When the atoms are caught in the magnetic trap, most of the resultant energy of the atoms comes from the added potential energy due to the Zeeman energy from magnetic field. The larger the cloud is when the magnetic trap is turned on the

greater the potential energy gained by the atoms. We can not adiabatically ramp on our magnetic trap from zero gradient, because at low magnetic gradients the trap center is significantly offset from the cloud center due to gravity. The center offset induces slosh in the trap, which turns into thermal energy. Therefore reducing the initial spatial extent of our CMOT cloud is more important than obtaining the coldest temperature.

Our CMOT stage consists of a MOT with greatly reduced repump laser power and large red detuning of the trapping laser. MOTs with large numbers of atoms have a maximum density of around 10^{11} atoms/cm³, which is limited by radiation pressure. The CMOT has the effect of reducing radiation pressure in the trap and thus creating a denser cloud of atoms. Reducing the repump power reduces the time the atoms spend in the state ($F=2$) resonant with the trapping light. Increasing the detuning of the trapping laser decreases the scattering rate and thus the absorption of reradiated photons.

We use a short CMOT stage in preparation for transfer to the magnetic trap. The CMOT has a much smaller loading rate than a regular MOT. Therefore we want to minimize the time spent in the CMOT stage and just go briefly to a CMOT configuration before the magnetic trap is turned on. The repump power is reduced from several mW to $50 \mu\text{W}$ for the CMOT stage. We have two separate repump beams entering the vapor cell. One beam is the main repump beam with several mW of power, and the other, which we call the bypass beam, has only $50 \mu\text{W}$ of power. Using two shutters we are able to have either full repump power or quickly reduce the power for the CMOT stage. We also jump the detuning of the trapping laser 50 MHz red of resonance. This frequency jump is accomplished by unlocking the laser, applying a voltage to the laser piezo electric tuner (PZT) for the CMOT and optical pumping stages, and then after the shutters are closed turning off the applied voltage and relocking the laser. We keep the magnetic field gradient constant at the MOT value during the CMOT stage. The entire CMOT stage lasts about 10 ms and is not very sensitive to changes in trapping

laser detuning on the order of 10 MHz.

We can characterize the atom cloud in the CMOT using fluorescence imaging. The position of the cloud in the CMOT may be very different from the position of the cloud in the MOT or magnetic trap because of beam imbalances or misalignment. We adjust the alignment and half-wave plates controlling the power in the beams to overlap the position CMOT with that of the magnetic trap, using the same high magnetic field gradient alignment technique used for the MOT/magnetic trap alignment. A large offset of the centers will increase the temperature of the magnetically trapped cloud.

After the CMOT stage we optically pump the atoms into the lower hyperfine ground state with arbitrary population in the magnetic sublevels. An atom excited to the $F' = 3$ state has about a 1 in 1000 chance of decaying to the $F = 1$; therefore the atoms will be pumped into the $F = 1$ state in less than one ms if the repump light is turned off. We trap only the $m_f = -1$ Zeeman state in the magnetic trap. Naively one might think that we trap only 1/3 of the atoms, surprisingly this is usually not the case. We can, in certain circumstances, get over 50% of the atoms in the right Zeeman state, depending on the MOT beam polarization. The population distribution in a MOT is not an easily controlled parameter but can be adjusted with small random tweaks of the MOT beams. We check to see that we are effectively pumping the atoms into the $F=1$ state by imaging the cloud with $F=2 \rightarrow F'=3$ light. If the atoms fluoresce, they have not been fully pumped into the $F=1$ state, and the optical pumping time must be increased.

Stage	Trapping/Repump	Detuning	Magnetic gradient	Time
MOT	On/3 mW	-2.5 Γ	8 G/cm	\sim 10 sec
CMOT	On/50 μ W	-10 Γ	8 G/cm	20 ms
Optical pumping	On/Off	-10 Γ	8 G/cm	1 ms
Magnetic catch	Off/Off	—	100 G/cm	200 μ s
Magnetic trap ramp	Off/Off	—	100 \rightarrow 250 G/cm	500 ms

Table 2.2: Parameters for trapping and loading into a magnetic trap

The atoms are now ready to be caught in the magnetic trap. As stated before we do not slowly ramp the magnetic field from zero due to the effect of gravity; we also do not want to turn on the magnetic trap at the highest gradient possible because this will add an excess amount of energy to the cloud. Our procedure consists of diabatically turning on the magnetic trap to a point where gravity has a minimal effect and yet the magnetic trap adds as little potential energy as possible. This catch point is determined empirically to be around 100 G/cm in the axial (vertical) direction. After the initial catch we adiabatically ramp the magnetic field gradient to 250 G/cm in 0.5 s. We optimize the transfer by imaging the atoms after they have been loaded into the magnetic trap, since we ultimately care about the temperature and number of atoms in the magnetic trap. We get $1-2 \times 10^9$ atoms at 250-400 μK in the fully compressed trap.

2.5.2 Fluorescence imaging

We need to image the atom cloud in the MOT cell for a variety of diagnostics, such as loading efficiency into and temperature in the magnetic trap, transfer efficiency to the science cell, and magnetic trap lifetime in different regions of the vacuum system. For these diagnostics it is not important to measure the absolute temperature and number of atoms in the trap but rather relative quantities. We use fluorescence imaging because, although it is far inferior to absorption imaging, it is easy to setup and gives us the information we require.

We optimize the transfer to the magnetic trap by imaging the cloud using fluorescence. We turn off the quadrupole trap, open the camera shutter, turn on the repump beam, flash the MOT beams for less than 1 ms, and image the fluorescence from the cloud onto a CCD camera. We image the atoms directly out of the magnetic trap without any expansion. We can not take an image too soon after the MOT stage because the fluorescence from the MOT will saturate the camera if the shutter is open. We therefore must wait several ms for the large camera shutter to open after the MOT stage before

we can take an image. The optical layout is shown in Fig. 2.10.




Figure 2.10: Top and side views of the optical layout for fluorescence imaging. The objective lens is apertured with an iris to 10 mm in diameter, and the second lens is 30 mm in diameter.

We calculate temperature of the cloud using the fitting parameters from a Gaussian fit. The cloud is not Gaussian in shape in a quadrupole trap. However, the cloud's profile is not far from Gaussian, and all we really want is a measure of the cloud that is monotonic with respect to size and fluorescence intensity. The temperature T and peak density n_0 are

$$T = \frac{2}{5} \frac{\mu_B g h}{k_b} B'_x \sigma_{FWHM} \quad (2.5)$$

$$n_0 = 1.27 \frac{N}{\sigma_{FWHM}^3}, \quad (2.6)$$

where g is the Landé g factor, and σ_{FWHM} is the full width at half max size of the cloud using the correct functional form of atoms in a spherical quadrupole trap. We calibrate the number N using the photodiode. Unfortunately our cylindrical glass cell causes some problems with imaging. The cell lenses the scattered light so that the size of the cloud is distorted by about 25% in the vertical direction. We get the temperature from the horizontal direction, on which the cylindrical cell has no effect.

Vignetting is such a common imaging systematic for fluorescence imaging that it deserves to be elaborated on here [23]. Vignetting occurs in a multiple lens system when imaging an extended object. It happens when there is more than one effective aperture

in the system. A common example is when a ray of light that passes through the objective lens does not make it through the second lens. The light that does not make it to the second lens comes primarily from the edge of the object as seen in Fig. 2.11. The decrease in imaged light from the edge of the cloud decreases the apparent size of the cloud.




Figure 2.11: Illustration of vignetting. Vignetting occurs when rays of light from the edge of an extended object are removed from the imaging system by a second aperture, in this case the second lens.

There are a few easy ways to check if an imaging system has vignetting issues. One way to check for a problem is to reduce the size of the first lens by a factor of two in diameter with an iris. This should decrease the total intensity of the image by a factor of four without changing the apparent width if vignetting is not a problem. One can also measure how close the system is to being affected by vignetting by reducing the size of the second lens with an iris. The size of the image will remain the same until the second lens begins to become an aperture in the system. There are several ways to eliminate vignetting: replace the second lens with a larger diameter lens, aperture the objective lens, or move the lens closer together.

2.5.3 Transfer from vapor cell to UHV region

We use moving magnetic coils to transfer the atoms from the relatively high pressure vapor cell to the UHV region, where we evaporatively cool to BEC. The quadrupole coils are mounted on a linear stage that is driven by a servo motor and controlled by a computer (Fig. 2.12). The maximum possible acceleration of the coils is about $1/3$ g, which is much less than the trapping force from the magnetic trap at 4 g; the atoms are therefore not heated any detectable amount. Also we do not see any atom loss from moving the atoms. We want to get the atoms out of the vapor cell as quickly as possible because collisions with the background gas limit the lifetime to 5-15 s depending on the Rb vapor pressure. However, we must slow down the coils as the trapped cloud enters the fringing fields of the Ioffe-Pritchard trap in order to adiabatically compress the cloud. We move the atoms out of the vapor cell, into the UHV region, and to within 4 cm from the center of Ioffe-Pritchard trap in about 1 s. We decelerate to a speed of 1 cm/s as the atoms enter the permanent magnetic region and are adiabatically compressed.

We purchased a commercial servo-linear track to move the coils from one end of the vacuum system to the other. The track itself is a table, mounted on a ground ball screw, which can accommodate higher speeds and has less backlash than a traditional lead screw. The servo motor in our system has reproducibility of $5 \mu\text{m}$, which is measured by a rotary encoder in the motor housing. The encoder signal, although sent through a shielded cable, is radiated from the encoder cable; typically we see 20 kHz spikes of $3 \mu\text{s}$ duration coming from the cable. So far this radiation has not caused any problems with other equipment or our ability to make a condensate.

The ability to move the atoms in the magnetic trap to different regions of the vacuum system allows us to measure the background pressure and also to localize possible places where stray light impinges on the the system. We need long, background-

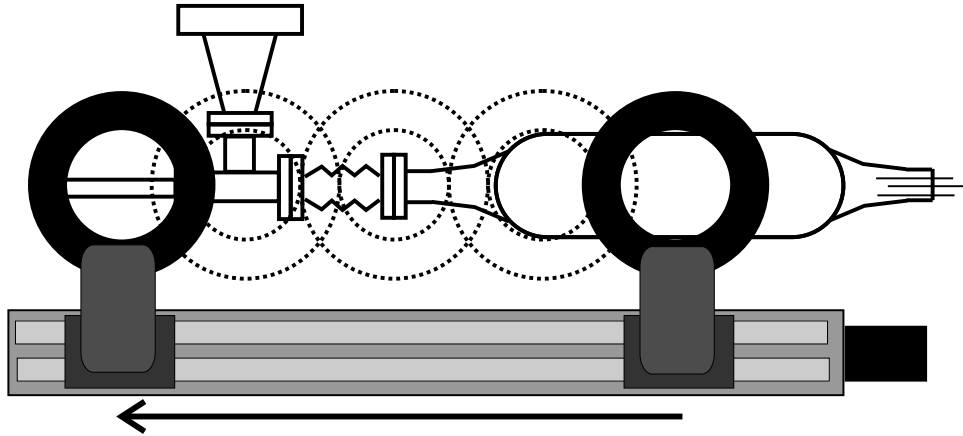


Figure 2.12: Illustration of the motion of quadrupole trap coils.

gas-limited lifetimes in the UHV region to be able to efficiently evaporate and form a condensate. We do not know the lower limit on the necessary lifetime, but we do know that our 170 s lifetime is much more than sufficient. We measure the lifetime in the magnetic trap by loading atoms into the quadrupole magnetic trap, moving the cloud to the desired position in the vacuum system, waiting some time, moving the atoms back to the MOT region, and imaging the cloud. We fit an exponential to the number of atoms remaining as a function of time. The exponential time constant gives us the lifetime in the trap.

There are two additional loss mechanisms that could reduce the lifetime in the magnetic trap. One is resonant light impinging on the atoms; the other is Majorana spin flips [24]. An atom absorbing a single photon has a large probability of falling back to an untrapped state and thus being ejected from the trap. There are two schools of thought on shielding laser light from the atoms in a magnetic trap. One is to place a box around the lasers and frequency stabilization optics; the other is to enclose the vacuum chamber and MOT optics. We choose to place a large (137 cm \times 124 cm \times 53 cm) box made from an aluminum frame with opaque lucite panels around the large MOT optics and vacuum system. This choice has the added benefit of also blocking room lights from

the experiment and thus allowing the experiment to be run with the room lights on.

The other loss mechanism is due to Majorana spin flips. Majorana or diabatic spin flips happen in a magnetic trap only when the trap has a zero of the magnetic field. Atoms can undergo a spin flip if the time rate of change of the magnetic field is greater than the Larmor frequency. In a quadrupole trap atoms which pass through an ellipsoid near the center of the trap can be lost due to spin flips to a non-magnetically trapped state. The lifetime associated with this loss rate is proportional to the square of the size of the cloud and is given by

$$\tau = \alpha l^2, \quad (2.7)$$

where α is determined experimentally to be

$$\alpha = 3.7(7) \times 10^4 \frac{s}{cm^2}, \quad (2.8)$$

and l is the radial half width half maximum of the cloud [24]. The loss rate due to spin flips is much smaller than the loss rate from background gas collisions for the typical cloud temperatures (200-400 μ K) we have in the quadrupole trap.

2.5.4 Ioffe-Pritchard Magnetic Trap

We have a hybrid Ioffe-Pritchard trap (HIP trap) which contains both permanent magnets and electromagnetic coils. Permanent magnets are useful because they produce large magnetic field gradients with no power consumption. However permanent magnets are sensitive to temperature fluctuations and thus can lead to instabilities if used to produce a bias field for a magnetic trap. The bias field is the most important parameter to stabilize because it determines the final evaporative cut and thus the temperature. In our trap the two permanent magnets produce a quadrupole field in the radial direction and no field along the axial or bias field direction.

The bias field and axial confinement are created by four electromagnetic coils. The outer two coils produce essentially all of the axial curvature, and the inner two

coils control the value of the bias field. Each pair of coils is run in series and controlled independently by a bipolar power supply. The power supplies internally servo the current to better than 1 part in 10^4 using an analog voltage set point supplied by a computer-controlled digital-to-analog converter. We increase the long term stability of our trap by continually running the operating current through our coils except for a 4 s period when the atoms are entering the HIP trapping region. The trap is therefore always at the same temperature even if our experimental timing changes. If we run operating current through the coils continually the coils reach a steady state temperature of 75°C . Although this temperature does not effect the operation of the magnetic trap it does increase the temperature of the glass cell, which it surrounds. Raising the temperature of the glass cell causes an undesirable increase in background pressure. We use forced air cooling to reduce the temperature of the coils from 75°C to 35°C . We have two air cooling ports fed by filtered compressed air as seen in Fig. 2.13. We do not use a fan to cool the trap because it adds noise to our trapping fields.

Our trap has the advantages of tight radial confinement from permanent magnets and also a stable bias field from well-servoed axial coils. Our trapping frequencies are (230, 230, 7) Hz with a 3 G bias field. The radial frequencies can be adjusted by changing the bias field. The radial frequency is

$$\nu = \frac{1}{2\pi} \sqrt{\frac{\mu_B m_f g_f}{m}} \frac{B'}{\sqrt{B_0}} \quad (2.9)$$

where μ_b is the Bohr magneton, g_f is the Landé g factor, m is the mass of Rb, B' is the field gradient, and B_0 is the bias field.

The trap configuration is shown in Figs. 2.14 and 2.15. The permanent magnets are $2'' \times 3/4'' \times 1/4''$ grade 35 Nd/Fe/B, which produce a gradient of 450 G/cm. Initially we used permanent magnets that produced a quadrupole gradient of 1200 G/cm. This gradient gave us 700 Hz radial trap frequencies at a bias field of 3 G. We found that having such tight confinement caused density dependent losses, which were so large




Figure 2.13: Science cell region showing the magnetic trap holder. Two stainless-steel endcaps on the ends of the Boron nitride form (white) are attached to a support structure behind the coil form. The microwave waveguide is shown on the left side of the picture directed towards the trapping region. Not shown is the objective lens on the back of the coil form.

that the final evaporation stage was not efficient, presumably due to inelastic collisions, and thus we produced smaller condensates with shorter lifetimes. With a sufficiently large number of atoms trapped initially, a larger transverse quadrupole gradient is not always better.

The outer(inner) coils are each 20(10) turns of 18 gauge magnet wire held in place with thermally conductive epoxy. In the normal configuration we run 13 A through the outer coils and 6.5 A through the inner coils. All of the trap components are mounted on a form made from Boron nitride. This ceramic has a high thermal conductivity, similar




Figure 2.14: Ioffe-Pritchard magnetic trap (end on view). The permanent magnets are at a 45° angle with respect to the horizontal axis so as to provide a magnetic field in the same direction as the magnetic field from the quadrupole trap used to transport the atoms. Four permanent magnets, magnetized through the thin dimension, will also work to provide a two dimensional quadrupole field with no field along the axial direction of the trap. However, using just two magnets magnetized through the thin direction will create a significant gradient along the axial direction of the trap.




Figure 2.15: Ioffe-Pritchard magnetic trap (side view) showing the permanent magnet and axial coil positions. Outer(inner) coils are each 20(10) turns of 18 gauge magnet wire and provide the axial confinement/bias.

to Aluminum, so that the heat generated from the coils can be removed. It also has a low coefficient of thermal expansion, smaller than stainless steel, which allows the position, and thus the trapping field, to remain constant as the trap holder changes temperature.

Boron nitride also allows microwaves through without attenuation for frequencies below 10 GHz. Being transparent to microwaves is important for our imaging procedure, in which we make transitions between hyperfine ground states, and for assorted scientific goals. Boron nitride has the consistency of a hard chalk but can be machined into simple shapes (Fig. 2.16).




Figure 2.16: Machine drawing of the Boron-nitride HIP-trap form.(all units in mm)

We choose to use a HIP trap in the experiment because of its great stability, but it is obviously not the only solution. A fully electromagnetic trap, such as the TOP trap would also work [24]. The main requirement for a trap is for it to have around 450 G/cm quadrupole gradient, which is not hard to achieve with electromagnetic coils close to the 1 cm diameter cell. The quadrupole gradient must be large to have an acceptable initial collision rate of at least a few Hz to begin evaporation. Initially the cloud is not in the harmonic region of the trap and is mostly confined by the quadrupole gradient. Therefore the quadrupole gradient defines the initial collision rate. A cloud

is in the harmonic region of the trap when the mean energy is less than one “bias field in energy”, $\mu_B g_f B_{bias}$, above the bottom of the trap; for a 3 G bias field a cloud is in the harmonic region when its temperature is below $30\mu\text{K}$.

2.5.5 Transfer between magnetic traps

We optimize the transfer between traps by transferring the atoms from the quadrupole trap to the HIP trap and back again. We then move the atoms back down to the MOT region and image them using fluorescence. We do this procedure instead of imaging in the HIP trap because we can not obtain accurate values for the temperature and number of hot clouds in the HIP trap. One reason it is nearly impossible to image in the HIP trap is that the cloud is too large in the axial direction for the image to fit on the CCD. Secondly there is a large magnetic-field-induced detuning across the radial direction of the cloud from the permanent magnetic field. Lastly the cloud’s optical depth is large in the magnetic trap, which leads to systematics in determining the number. The cloud’s temperature and number measured in the MOT cell after being brought back from the HIP trap are not a completely accurate representation of the parameters that existed in the HIP trap, but the comparisons are at least monotonic.

Transferring atoms between a quadrupole trap and an Ioffe-Pritchard trap can not be done completely adiabatically due to the relative directions of the magnetic fields in each trap. If the transfer is done correctly however, one can limit loss in phase space density to a factor of two to four during transfer. We start by slowly ramping down the quadrupole gradient (~ 500 ms) to a point where the cloud is approximately mode-matched in the axial direction; for our trap this corresponds to a gradient of 100 G/cm. Next we turn on the HIP axial coils discontinuously and turn off the quadrupole coils. We find that the timing of the traps turning on and off is not critical to within 10 ms. After the quadrupole trap is off, we quickly move the coils down to the MOT region in preparation for the next run.

2.6 Rf evaporation

Now that we have atoms in the HIP trap we can evaporatively cool them to degeneracy. The basic idea of evaporation is to remove atoms with more than the average energy of the cloud and allow the ensemble to equilibrate to a lower temperature through collisions [10]. We need an adequate elastic collision rate to have the sample reequilibrate before there is a large loss of atom number or a large increase in energy of the sample from inelastic collisions. The elastic collision rate is

$$\gamma = n(8\pi a^2)v_{rel}, \quad (2.10)$$

where the peak density for atoms in a quadrupole trap is

$$n = \frac{0.633 N}{\sigma_x \sigma_y \sigma_z}, \quad (2.11)$$

and the peak density in a harmonic trap is

$$n = \frac{1}{(2\pi)^{3/2}} \frac{N}{\sigma_x \sigma_y \sigma_z}. \quad (2.12)$$

The relative velocity is

$$v_{rel} = \sqrt{\frac{16k_b T}{3\pi m_{Rb}}}, \quad (2.13)$$

where N is the number of atoms, a is the s-wave scattering length, σ_i is the full width at half max of the cloud in the i dimension, and T is the temperature of the cloud.

There are three types of inelastic collisions we have to be concerned with during evaporation: one-, two- and three-body processes. One-body loss from collisions with background gas atoms will cause essentially only number loss and does not induce heating, because all atoms in the trap have about the same probability of removal. During the initial stages of evaporation one-body loss is the dominant factor because the density is low, inhibiting density-dependent collisions. As the density increases two- and three-body process become important. Two-body processes are significantly suppressed with a spin-polarized gas in the maximum angular momentum state of a

ground hyperfine state. Due to angular momentum conservation two atoms can not collide and spin flip to a different Zeeman level in the lower hyperfine state. The only way for atoms to change their spin is for one of the atoms to be excited to the upper hyperfine state, which requires 6.8 GHz of energy. An upper bound on the rate constant has been determined experimentally to be $1.6 \times 10^{-16} \text{cm}^3/\text{s}$ [25]. Two-body loss is seldom an issue for ^{87}Rb in the lower hyperfine state. Three-body loss happens when three atoms collide, two forming a molecule, and the other taking away the residual angular momentum. Three-body processes are proportional to density squared; the three-body rate constant has been measured to be $4.31(1.8) \times 10^{-29} \text{cm}^6/\text{s}$ for ^{87}Rb atoms [25] in the $|F = 1, m_f = -1\rangle$ state, and becomes a concern when the density approaches 10^{13}cm^{-3} . The three-body process not only causes atom loss but also heating because atoms are preferentially lost from the highest density region of the cloud, which corresponds to the atoms with the least energy in a magnetic trap. We call this processes anti-evaporation. A good rule of thumb for efficient evaporation is to have the ratio of elastic to inelastic collisions be ~ 100 . We define runaway evaporation as the situation where the collision rate increases as the cloud cools. Inelastic collisions always increase as the density increases, so without runaway evaporation one will reach a point where the elastic to inelastic collision ratio will drop below 100.

Our evaporation consists of applying rf radiation to the atoms to spin flip them to an anti-trapped state. The spin flipping is done in an energy dependent manner by taking advantage of the Zeeman shift from the magnetic trap. High energy atoms spend more of their time in high magnetic fields near the edge of the cloud compared to the low energy atoms, and thus the high energy atoms have a larger Zeeman energy splitting between magnetic sublevels. We can tune the rf to be resonant with the Zeeman energy splitting for atoms with the most energy, removing them from the trap. We continue to lower the rf, removing lower and lower energy atoms, all the time approximately maintaining thermal equilibrium. We want to remove the same fraction of energy from

the cloud for each evaporation frequency step if the collision rate is constant. This condition corresponds to an exponentially decreasing frequency ramp. As the collision rate changes so will the optimum exponential time constant. The functional form we use is

$$\nu(t) = (\nu_{start} - \nu_o)e^{-t/\tau} + \nu_o, \quad (2.14)$$

where ν_{start} is the frequency where we begin evaporating, ν_o is the frequency corresponding to the bottom of the trap, and τ is the exponential time constant of the ramp. The time constant is set by the collision rate and loss rate.

2.6.1 Evaporation optimization

We need several stages of evaporation, each with different parameters. Throughout evaporation both the elastic and inelastic collision rates change as well as the rf coupling to the atoms, thus we must adjust the evaporation time constant and the rf power for each stage. We decrease both the time constant, due to the increased collision rate, and the rf power delivered to the atoms, to avoid power broadening effects, as the atoms cool.

Power broadening of the rf knife will cause the evaporation to lose energy selectivity as the width of the knife becomes comparable to the temperature of the cloud. For the initial stages of evaporation we need more power to have the same large probability that an atom will spin flip as when the cloud is cold, because the atoms initially have a larger velocity. The probability of being ejected from the trap scales as $[1 - \exp(-\text{rf power}/\text{velocity})]$. As the cloud approaches zero temperature one must be very careful about not applying too much rf power. Another concern with setting the rf power is coil or amplifier resonances. The rf coil may have a self-resonant frequency in the frequency range spanned by the evaporation. An easy way to check for resonances is measure the rf power delivered to the atoms using a small pick-up coil placed near the the evaporation coil. The rf power from the synthesizer may have to be drastically

reduced near a resonance to avoid power broadening.

We break up the evaporation into enough stages so that we decrease the temperature by a factor of two or three with each stage; this criterion sets the start and stop frequency for each stage, points A-E in Fig. 2.17. We start with a time constant of 5 s. In general the time constant of the evaporation ramp should be about a factor 10-20 greater than $1/\text{collision rate}$. We do not actually ramp continuously the frequency of the rf for the initial stages but instead have discrete steps sent to the synthesizer through GPIB. Typically a single GPIB command will take between 30-50 ms to be received and executed; therefore we send a new frequency command every 50 ms. The discrete nature of the frequency ramp is not a problem for the initial stages when each step is small compared to the temperature of the cloud, but is a problem in the last stage of evaporation. For the last stage we sometimes use a programmable frequency synthesizer, which can ramp the evaporation frequency phase continuously. The extra synthesizer is not necessary but will produce larger condensates.

We want to optimize collision rate for each stage of the evaporation. If we image the cloud right after the stage we are optimizing the cloud may not be in equilibrium due to a too-rapid cut and appear to have an anomalously large collision rate. Imaging a cloud out of equilibrium can systematically misrepresent the collision rate. However if we add an additional evaporation stage before imaging, we can circumvent this problem. The second stage will not be as efficient if the cloud is out of equilibrium, and thus the cloud will have less number after the second stage.

Optimizing the initial stages of evaporation in our trap is difficult because of our inability to image hot clouds. As stated before, we can not obtain an accurate temperature or number of atoms in our cloud when the temperature is above $1\mu\text{K}$ because of the magnetic field gradients. However, we have created an optimization procedure for the first stages of evaporation that is acceptable. We start with two stages as seen in Fig. 2.17. The parameters of the first stage, segment A-B, are varied,

and the second stage, segment B-C, parameters are kept constant. We image the cloud at point C and maximize the peak OD by changing the parameters for segment A-B. Even with imperfect imaging, the peak OD measured after ramp B-C is monotonic in the true equilibrated collision rate produced by ramp A-B. Next we add a stage C-D and optimize segment B-C and so on. It is important to iteratively adjust the time constant and rf power because they are coupled. Luckily the initial evaporation is not very sensitive to the parameters of the cut so this procedure works well.




Figure 2.17: Sample evaporation trajectory with four segments shown. We typically use 8 segments, each providing a factor of 2-3 decrease in $(\nu - \nu_0)$

The final stages of evaporation are more critical than the first stages. However for the last stages we can accurately determine the temperature and density of the cloud. The procedure for extracting these numbers from the images will be discussed in Section 2.8. Except when optimizing the very last stage, we characterize a given stage by optimizing number in the cloud after an additional stage. It is easy to walk the parameters in the wrong direction, especially with the final stage; one tends to have too short a time constant and too much rf power. We reduce the rf power 15 db from the first to the last stage. Generally we can change the rf power by 5-10 db in the upper stages and 3 db in the later stages without observing a significant change in evaporation

efficiency. When optimizing the last few stages it is also important to remeasure the trap bottom ν_0 , which can be more accurately determined now that one has a cold cloud.

2.6.2 rf coil

We use a simple single-loop coil to deliver rf to the atoms for evaporation. We typically evaporate from 40→2 MHz. The large range of frequencies we use restricts us from impedance matching the coil to gain better coupling. Because we are in the near-field limit of the radiation for all evaporation frequencies, one can think of the rf as just an oscillating magnetic field. Only the component of the magnetic field produced by the coil perpendicular to the atoms' quantization field will cause transitions between the Zeeman states. Therefore we place the coil directly outside the glass cell with the axis of the loop perpendicular to the bias field to maximize the coupling to the atoms when they are cold. Hot clouds will have atoms in the quadrupole field with quantization axes in every direction, and there will be small regions in the cloud that are not affected by the rf. This does not appear to pose a serious problem, however. The coil must not be placed closer than one radius to any electrically conductive object; the conductive object will reduce the flux return path and thus the magnetic field produced. The size of our loop is about 1 cm in diameter, made from 18 gauge magnet wire. We want to minimize the enclosed area of the loop that is not near the atoms to maximize the field at the atoms, so the loop is soldered directly to a RG 175 cable, which is powered by a 5 Watt rf amplifier.

2.7 Absorption imaging

We image clouds in the HIP trap using laser absorption. We illuminate the cloud with resonant light, atoms scatter photons out of the beam, and we focus the shadow cast by the atoms onto a charge coupled device (CCD) array. The amount of light

absorbed gives the column optical density (OD) at a particular position in the cloud, which can be used to calculate the number and temperature of the sample. Optical density is defined by Beer's law and is given by

$$I = I_0 e^{-OD}, \quad (2.15)$$

where I is the intensity impinging on the camera and I_0 is the initial intensity of the beam.

2.7.1 Optical set up

The imaging optics are shown in Fig. 2.18. We use a probe beam that has been filtered spatially by a single mode fiber. The probe beam is expanded to a diameter of 1 cm so that the intensity across a 100 μm condensate is nearly constant. The light first passes through a polarizer and then through a $\lambda/2$ plate so that we can adjust the angle of the linear polarization. We place a 4 mm² aperture made from electrical tape on the magnetic trap form to eliminate excess light that could scatter onto the camera from defects in the glass cell. The shadow of the atoms is focused onto the camera with two lenses. We use, as the objective, a 1 cm diameter gradient index singlet lens, which is a diffraction limited optic. The objective lens is mounted directly on the trap coil form to collect the largest possible solid angle. We use a 30 mm diameter achromat doublet as the second lens. We use an achromat, not for its reduction in chromatic aberrations, but for its low spherical aberrations when oriented correctly. Our CCD camera, which is no longer commercially available, is a front illuminated CCD array with pixels 9 μm on a side. The entire array is 720 X 560 pixels and has 16 bit resolution. The quantum efficiency at 780 nm is around 35%.




Figure 2.18: Optical system for absorption imaging. Light from a single mode fiber is put through a polarizing beamsplitting cube (PBS) and a half waveplate to adjust the polarization. After the light passes through the cloud, the image is focused onto a CCD camera by two lenses.

2.7.2 Imaging in a non-uniform magnetic field

Imaging atoms in a strong quadrupole magnetic field presents many problems. Our goals are to i) expand the cloud above the resolution limit of our imaging system, ii) image through the radial direction so we can observe dynamics along the axial direction, and iii) extract useful parameters from the image without systematic effects. These goals are difficult to obtain, because first we can not turn off the quadrupole field created by the permanent magnets and allow the cloud to ballistically expand as is usually done. Second imaging in the radial direction causes the bias (quantization) axis to be perpendicular to the propagation direction of the probe beam, and thus does not allow us to drive purely $\sigma+$ or $\sigma-$ transitions, which is desirable because it would give us a cycling transition and thus a large signal-to-noise ratio. The last obstacle to overcome is the magnetic field gradient, which causes a spatially varying energy shift due to the Zeeman effect. Therefore, we can not apply light which is resonant with the entire cloud. The spatially varying detuning causes the image to have systematically the wrong width and intensity.

We have found ways to reduce or eliminate all of our imaging problems. We expand the cloud by transferring the atoms to an anti-trapped state and allowing them

Propagation direction wrt quantization axis	Polarization	Transitions
Parallel	Right circular	$\sigma+$
Parallel	Left circular	$\sigma-$
Parallel	Linear	$\sigma+, \sigma-$
Perpendicular	Circular	$\sigma+, \sigma-, \pi$
Perpendicular	Linear along quantization field	π
Perpendicular	Linear perpendicular to quantization field	$\sigma+, \sigma-$

Figure 2.19: Possible transitions with different probe beam polarizations

to fall off of the potential created by the magnetic trap. We use a microwave adiabatic rapid passage (ARP) to transfer coherently the atoms from the $|1, -1\rangle$ to the $|2, -2\rangle$ state [26]. ARP is best described in the dressed state picture, which is discussed in Chapter 6.

To ARP the atoms from one bare state to another we turn on a coupling field far off resonance, ramp the frequency slowly, compared to the Rabi frequency, through resonance, and then turn off the field. This will coherently transfer the atoms between the two bare states.

The other problem to solve is the incorrect imaging polarization. We would like to drive a cycling transition from the $|2, -2\rangle$ to the $|3', -3\rangle$ state. If we choose our probe beam polarization linear and perpendicular to the bias field we can drive not only the desired transition but also the $|2, -2\rangle \rightarrow |3', -1\rangle$ transition, which is obviously not a cycling transition. However we image in 100 G bias field which breaks the degeneracy of these two transitions by 31 linewidths and allows us to have a cycling transition. Table 2.19 lists the possible transitions for the different probe beam propagation directions

and polarizations.

Naively one might expect that at maximum only half of the light could be absorbed, because we think of linear light as an equal amount of $\sigma+$ and $\sigma-$ light. However in the atoms' frame these two polarizations are not the correct basis and are actually coupled. Absorption of this type is typically known as the Voigt effect. In fact all of the light can be absorbed by the atoms, and the only effect of the direction of the polarization is a reduction in the line strength by a factor of two.

Lastly, we reduce the Zeeman detuning across the cloud by increasing our bias field to 100 G; this is easily accomplished by reversing the current in the inner coils. The gradient adds in quadrature with the bias field thus reducing the spatial variation of the magnetic field from 2.4 G to 0.1 G for a typical expanded radial cloud radius of 100 μm . The residual variation in magnetic field across the cloud changes the imaging detuning by 140 kHz, which is much less than a natural linewidth.

2.7.3 Producing and controlling microwaves

We use microwaves for manipulating the state of the atoms during our imaging procedure as well as during our spectroscopic measurements. We use two different systems to produce microwaves. One system is a commercial microwave synthesizer, which, although flexible and easy to use, has a frequency stability of only 5×10^{-12} in a second. This stability is acceptable for ARP during our imaging procedure, but not quite for long interrogation times during Ramsey spectroscopy. The other system we use to produce microwave frequencies is a home-assembled frequency chain as diagrammed in Fig. 2.20. This frequency chain has frequency stability better than 10^{-12} in a second.

Our frequency chain begins with a reference signal from a global positioning system (GPS) disciplined oscillator. This device provides only long-term frequency accuracy and stability, where long term refers to a time scale on the order of hours. The GPS-referenced oscillator removes long-term drift of a Wenzel 10 MHz phase-locked

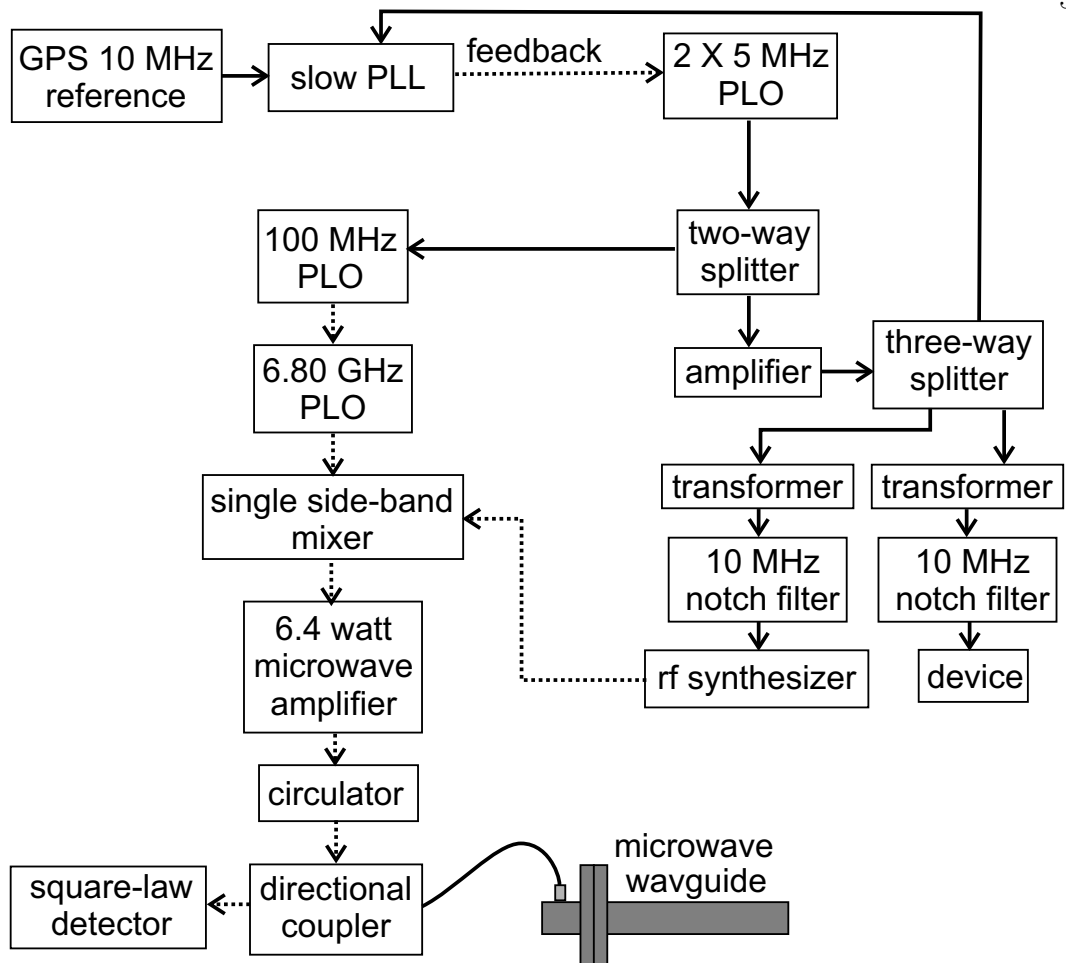


Figure 2.20: Microwave frequency chain showing all of the components we use to produce variable frequency microwaves at 6.8 GHz. The solid arrows indicate transmission of 10 MHz signals. See text for further description.

oscillator (PLO), via an external, slow-response phase-locked loop (PLL), which has a bandwidth of about 1 s. The Wenzel 10 MHz signal is split, with one portion sent to a Wenzel 100 MHz PLO and the other portion sent to another splitter. We use the signals from the second splitter to phase lock other rf devices in the experiment. We gain short term frequency stability from the 10 MHz PLO and mid-range frequency stability from the 100 MHz PLO. The cross-over region between the two PLOs is around 5 kHz. The 100 MHz signal is multiplied up to 6.80 GHz and phase-locked to a 6.8 GHz crystal oscillator by a Microlambda device. We require the ability to adjust the frequency of the

microwaves, so we mix adjustable-frequency rf and the 6.80 GHz signal with a single sideband mixer (image reject mixer). We use a single- instead of a double-sideband mixer because it does not produce significant power in the unwanted sideband; thus it allows greater power output of the desired sideband by the amplifier. The microwave signal is then amplified by a 6.4 Watt solid state amplifier. We place a circulator on the output of the amplifier to reduce reflected power coming back to the amplifier. Also for diagnostic purposes, we place a directional coupler after the circulator, such that the directional coupler splits off a small amount of the reflected power and sends it to a square-law detector. The square-law detector produces a dc signal proportional to the microwave power.

The microwaves are transmitted to the atoms by a sawed-off waveguide, which is placed near the HIP trap, directed along the axis of the trap. We use low-loss semi-flexible microwave cables to transfer the microwaves from the amplifier to the waveguide to maximize the power delivered to the atoms.

2.7.4 Imaging procedure

We start the imaging procedure with the atoms in the $|1, -1\rangle$ state. We ARP the atoms to the $|2, -2\rangle$ using microwaves. The microwave field must be swept phase continuously for atoms to remain in one dressed state. The easiest way to have a phase-continuous sweep is to hold the microwave frequency constant and ramp the bias magnetic field. We typically start about 1 MHz away from resonance and sweep through in 0.3 ms; our Rabi frequency is around 100 kHz. Next we jump the bias field to 48 G, and wait for the anti-trapped atoms to expand. If we expanded in our normal 3 G bias field trap the atoms would expand too rapidly into the anharmonic the region of the trap, thus making it difficult to calculate the effect of this expansion. If we jump directly to a 100 G bias field, the atoms would expand too slowly and fall under gravity, once again into the anharmonic region of the trap. The intermediate field keeps the

Time (ms)	Event
0	Camera triggered, probe beam shutter open
9	Centering coil on
10	Microwave on, bias field ramp for ARP started
10.3	bias ramp stopped, microwave off
10.3	Bias field jumped to 48 G for expansion
10.3 + Expansion time	Bias field jumped to 100 G for imaging
10.4 + Expansion time	Probe beam AOM on
10.402 + Expansion time	Probe beam AOM off
20	Probe beam shutter closed

Table 2.3: Image timing

atoms in the harmonic region of the trap during the entire expansion. The atoms are also slightly sagged in the trap due to gravity, so when they begin to expand they are sitting on the side of the potential, which induces some asymmetry to the expansion. We correct for this sag by applying a small magnetic field to shift the center of the trap to the center of the cloud just before the expansion. We find the correct magnitude of the centering field by imaging the cloud after long expansion times and adjusting the added field until the cloud remains in the same position in the vertical direction.

After the cloud has expanded the desired amount, we jump the bias field to 100 G and flash the probe beam for 20 μ s. We use a short 20 μ s pulse for two reasons. First we do not want the atoms to be excited to the to the $|3', -1\rangle$ state and fall back to a dark state. We are only 200 MHz detuned from the $|2, -2\rangle \rightarrow |3', -1\rangle$ transition and therefore will drive transitions to this state, although with a very low probability. Second if the atoms in the cloud scatter many photons they will pick up enough momentum to move along the direction of the probe beam; this motion could blur the image or cause the atoms' transition frequency to change as they experience a different magnitude of the magnetic field.

We take two additional pictures for normalization purposes. One picture is taken with the probe beam on but with no atoms present I_{light} ; this gives our light image which we use to calculate percent absorption. We also take a third picture with the

probe beam off and the camera shutter open I_{dark} . This image will give a calibration of any stray light that does not come from the probe beam as well as camera dark current. We calculate the optical density of each pixel, which is given by

$$\text{OD}_{meas} = \ln \left(\frac{I_{light} - I_{dark}}{I_{atoms} - I_{dark}} \right). \quad (2.16)$$

There are two common systematics that should be addressed with any absorption imaging system. One is the effect of a maximum observable optical density. Any light collected by the camera that can not be absorbed by the atoms will reduce the observed OD. Two usual culprits are off-resonant light and probe-beam light that does not pass through the cloud. A good way to check how much of the probe beam is far off resonant is to send the probe beam through a heated Rb vapor cell and measure the percent absorbed. Some diodes have a broad pedestal of light that is not in the main frequency mode of the laser and thus can cause a reduction in the observed OD. The second reason for a low maximum observable OD is probe light which is reflected off of objects and is indirectly scattered onto the CCD. We place a small aperture in front of the cell to reduce scattering light from the cell onto the CCD. We observe a maximum OD (OD_{sat}) around 2.8, even for clouds for which the actual OD is much greater. This OD_{sat} value is acceptable, but we must correct for it in our analysis of the image. The modified OD, taking out the effect of OD saturation, is

$$\text{OD}_{mod} = \ln \frac{1 - e^{-\text{OD}_{sat}}}{e^{-\text{OD}_{meas}} - e^{-\text{OD}_{sat}}}. \quad (2.17)$$

We measure the maximum observable OD by creating a large cloud and expanding it for 1 ms. The center of the cloud will have a flat top where the OD is saturated at the maximum value. If the correction factor between OD_{meas} and OD_{mod} is too large, the potential for error increases. We increase the expansion of the cloud until $\text{OD}_{meas} < \text{OD}_{sat}/2$.

The other systematic with absorption imaging is the effect of probe beam intensity

saturation. The actual OD is

$$\text{OD}_{actual} = \text{OD}_{mod} + (1 - e^{-\text{OD}_{mod}}) \frac{I}{I_{sat}}, \quad (2.18)$$

where I is the intensity of the probe at the position of the cloud and I_{sat} is 1.6 mW/cm² for Rb on a cycling transition. I_{sat} is increased by a factor of two for our case because the line strength is reduced for our imaging polarization. We like to minimize the correction factor, so we work at less than 1/10 of saturation intensity.

The resonant frequency for the imaging transition can be calculated easily because both the initial and final states are maximum angular momentum states. F and m_f are therefore always good quantum numbers, and the frequency splitting between the two state is $\Delta\nu = \mu_B B$. We confirm we are on resonance by taking a transition lineshape, which involves producing identical clouds, and probing them with different frequencies. We change the frequency of the first probe AOM and measure the peak optical depth. The resulting curve should be a Lorentzian with the natural linewidth, Γ . Measuring the natural linewidth with the expected center implies that many parameters are correct in the imaging system, including narrow laser linewidth, accurate calibration of magnetic fields, probe beam well below saturation, and correct control of probe frequency.

The lineshape can also be a useful diagnostic for probe laser frequency noise. Frequently the probe laser frequency may be affected by shutter-induced vibrations or current transients right before imaging. Therefore it is important to measure the noise on the laser during the imaging pulse. One can find the shot-to-shot standard deviation of the number of atoms imaged while the probe is tuned on resonance and contrast with that of a half linewidth off resonance. Comparing the two measurements rejects uncorrelated atom number fluctuations. A significant increase in shot-to-shot noise when the laser is tuned a half linewidth off resonance indicates probe laser frequency or magnetic field noise.

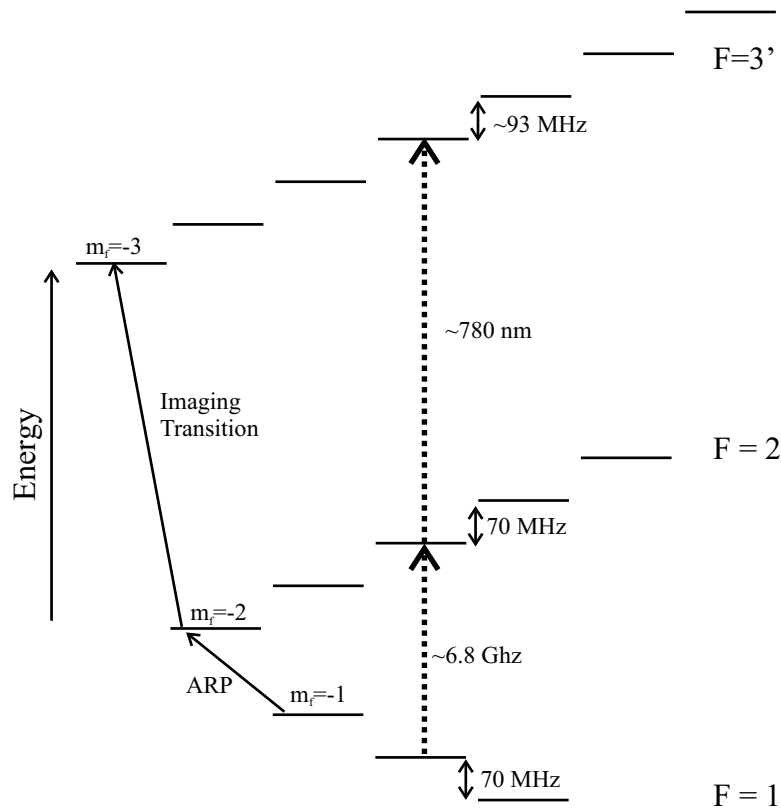


Figure 2.21: Energy level diagram showing imaging transitions in a 100 G magnetic field. The microwave ARP transition and optical imaging transition are shown with solid arrows.

2.7.5 Focusing the image

We focus the image onto the CCD camera by imaging a small cloud that has not expanded much. Before we focus the image we first take a lineshape to ensure we are on resonance. Imaging on resonance will minimize the cloud lensing the light, which leads to incorrectly focusing of the image. We also use a cloud with a peak OD of less than one to reduce effects of lensing. Once we have tuned the probe laser to the resonant frequency of the atomic transition, we adjust the position of the camera along imaging axis. The focus of the image will be at the minimum cloud width. We focus the image by measuring the width in the radial direction, which is not necessarily the focus in the axial direction of the cloud because of the astigmatism induced by the

cylindrical cell. If the cloud is exactly in the center of the glass cell the image would not be astigmatic because all the rays of light hit perpendicular to the glass and therefore are not refracted.

2.7.6 Measuring the image magnification

We measure the magnification of our imaging system by watching a cloud fall under the influence of gravity. We start with atoms in the $|1, -1\rangle$ state and perform an ARP to place them in the $|2, 0\rangle$ state, which is affected only slightly by the magnetic field gradient. We allow the atoms to fall for a varying time and measure their resulting position. A cloud's position as a function of time including the small acceleration due to the second order Zeeman shift is

$$z(t) = -\left(\frac{a}{c}\right) (\cos \sqrt{ct} - 1) + z_0, \quad (2.19)$$

where a is the acceleration due to gravity in pixels/ms², z_0 is the cloud's initial position, and c is

$$c = \left(\frac{4\pi\hbar}{m_{Rb}}\right) f B'^2, \quad (2.20)$$

where f is the second order Zeeman shift of 287.57 Hz/G², and B' is the magnetic field gradient. Fitting the position vs. drop time data will give a value for a , which can be used to find the magnification M , where M is given by

$$M = \frac{9.81\mu\text{m}/\text{ms}}{\sqrt{2} a}. \quad (2.21)$$

The factor of $\sqrt{2}$ is included because our imaging axis is at 45° with respect to gravity.

2.8 Image analysis

2.8.1 Image processing

After the three images have been downloaded to the computer we apply some image processing before fitting the images. We first calculate a measured OD for each

pixel using Eqn. 2.16. Occasionally we will get pixels that have anomalously high or low values due to noise or readout error. We remove these pixels by systematically going through the image array comparing nearest neighbor pixels. If there is a difference of 6, in units of OD, or greater the pixel is replaced by the average of the eight adjacent pixels. We perform the same procedure a second time, this time using a difference threshold of 0.8. After the spikes are removed, we apply corrections for saturation effects using Eqns. 2.17 and 2.18. Finally, depending on the size of the cloud, we bin the array. For most condensate images we use 2 X 2 binning, making a modest sacrifice in resolution in order to reduce calculation time during the fitting routine.

2.8.2 Image fitting

We use three different fitting routines depending on the degeneracy of the cloud [15]. For clouds above the condensation temperature we fit the image to a 2-D Gaussian. Clouds at finite temperature but with a condensate present we fit to a modified Gaussian plus a Thomas-Fermi distribution. In some cases, where the cloud has no detectable thermal fraction, we just use a Thomas-Fermi distribution.

Thermal clouds are fit using a standard 2-D Gaussian function which is

$$z_{Gauss} = |PeakOD| \exp\left(-\frac{1}{2} \frac{(x - c_x)^2}{s_x^2} - \frac{1}{2} \frac{(y - c_y)^2}{s_y^2}\right), \quad (2.22)$$

where the adjustable parameters are c_x , c_y , the cloud centers in both directions, s_x , s_y , the Gaussian widths in both directions, and the peak OD. The fitting is done using a Matlab script called from inside LabVIEW.

We use a slightly more complicated procedure to fit clouds that have both a condensate and thermal cloud present. The condensate portion of the image can be fit to a Thomas-Fermi profile, which is parabola integrated in one dimension. The thermal cloud is no longer an ordinary Gaussian when it is degenerate but is modified by Bose statistics. We first use only the pixels that have an OD from OD_{peak} to 70% OD_{peak} for

a fit to a Thomas-fermi profile given by

$$f_{TF} = |TFpeakOD| \text{MAX} \left[\left(1 - \left(\frac{x - c_x}{s_{TFx}} \right)^2 - \left(\frac{y - c_y}{s_{TFy}} \right)^2 \right)^{3/2}, 0 \right], \quad (2.23)$$

where s_{TFx} , s_{TFy} are the radii of the profile in each direction.

The Thomas-Fermi fit curve is subtracted from the image and the resulting image is fit to the modified Gaussian given by

$$f_{modGauss} = \frac{g_{\frac{3}{2}}(f_{Gauss})}{g_{\frac{3}{2}}(1)}, \quad (2.24)$$

where $g_{\frac{3}{2}}$ is

$$g_{\frac{3}{2}}(z) = z + \frac{z^2}{2^{3/2}} + \frac{z^3}{3^{3/2}} + \frac{z^4}{4^{3/2}} + \frac{z^5}{5^{3/2}} + \frac{z^6}{6^{3/2}} + \frac{z^7}{7^{3/2}} + \frac{z^8}{8^{3/2}} + \frac{z^9}{9^{3/2}} + \frac{z^{10}}{10^{3/2}} \dots \quad (2.25)$$

We use just the first 11 terms of $g_{\frac{3}{2}}$ in the fitting procedure. We also impose a constraint on the Gaussian fit to have the same center as the Thomas-Fermi fit.

2.8.3 Calculating cloud parameters

Once we fit the image and extract the fitting parameters we can calculate the properties of the cloud. The first step is to calculate the size of the cloud in the magnetic trap based on our anti-trapped expansion. The Boltzmann equation for the equilibrium distribution function F , which describes the anti-trapped expansion, is

$$\mathcal{L}F(x, v, t) = 0, \quad (2.26)$$

where F and \mathcal{L} are given by

$$F(x, v, t) = \frac{N m_{Rb}^3 \lambda_{dB}^6 |\bar{\omega}|}{h^6} \exp\left(-\frac{m_{Rb}}{2k_B T} (\bar{v}^2 + \bar{\omega}^2 \bar{x}^2)\right) \quad (2.27)$$

$$\mathcal{L} = \frac{\delta}{\delta t} + \bar{v} \cdot \nabla_{\bar{x}} - (\bar{\omega}^2 \bar{x}) \cdot \nabla_{\bar{v}}, \quad (2.28)$$

and the matrix of the trap frequencies is

$$\bar{\omega} = \begin{pmatrix} \omega_x & 0 & 0 \\ 0 & \omega_y & 0 \\ 0 & 0 & \omega_z \end{pmatrix}. \quad (2.29)$$

The solution to the Boltzmann equation, for each dimension of the cloud, is given by

$$\sigma(t=0) = \frac{\sigma_{meas} \omega}{\sqrt{\omega^2 + (\omega^2 + \omega_0^2) \sinh^2(\omega t)}}, \quad (2.30)$$

where σ_{meas} is the cloud size after expansion, ω is the harmonic trap frequency during the expansion, ω_0 is the original trapping frequency, and t is the expansion time. This treatment assumes that the initial position and velocity are uncorrelated and that the mean-field does not contribute significantly to the expansion. The mean-field contribution to the expansion is negligible in the axial direction, from which we calculate all size dependent parameters.

The expansion is dominated by the anti-trapping potential and not the mean field, so the expansion is essentially the same for the condensate and the normal component. We calculate the temperature of the cloud from the axial width because there is very little expansion in this direction, and the frequency of the potential does not change much during the expansion.

The following equations allow one to calculate the useful properties of normal clouds and condensates. The absorption cross-section A , is

$$A = (\text{branching ratio}) \left(\frac{3\lambda^2}{2\pi} \right) \times \frac{1}{1 + 4\frac{\Delta^2}{\Gamma^2}} \quad (2.31)$$

where λ is the wavelength of the transition, Δ is the detuning from resonance, and Γ is the natural linewidth. The number, temperature, peak number density, collision rate, relative velocity and phase space density of a normal cloud are respectively

$$N_{normal} = 4.88512 OD_{peak} s_{radial} s_{axial} A \quad (2.32)$$

$$T = \frac{m_{Rb}}{k_B} (\omega_{axial} s_{axial})^2 \quad (2.33)$$

$$n_{normal} = \frac{1}{(2\pi)^{3/2}} \frac{N a.r.^2}{s_{axial}^3} \quad (2.34)$$

$$C.R._{peak} = n (8\pi a^2) v \quad (2.35)$$

$$v = \sqrt{\frac{16 k_B T}{3 \pi m_{Rb}}} \quad (2.36)$$

$$PSD = n \left(\frac{h}{\sqrt{2\pi m_{Rb} k_B T}} \right)^3, \quad (2.37)$$

where $s_{radial/axial}$ is the Gaussian size as defined in Eqn. 2.22, ω_i is the trap frequency in the i dimension, $a.r.$ is the aspect ratio of the trap ($a.r. = \frac{\omega_{axial}}{\omega_{radial}}$), and a is the s-wave scattering length. The number, peak number density, and chemical potential of a condensate are given respectively by the following

$$N_{BEC} = \frac{2\pi}{5} O D_{peak} s_{TFaxial} s_{TFradial} A \quad (2.38)$$

$$n_{BEC} = \left(\frac{15}{8\pi} \right) \frac{N_{BEC} a.r.^2}{s_{axial}^3} \quad (2.39)$$

$$\mu_{BEC} = \frac{1}{2} \hbar \omega_{ho} \left(\frac{15 N_{BEC} a}{a_{ho}} \right)^{2/5}. \quad (2.40)$$

The harmonic oscillator length is

$$a_{ho} = \sqrt{\frac{\hbar}{m_{Rb} \omega_{ho}}}, \quad (2.41)$$

and the harmonic oscillator frequency is

$$\omega_{ho} = (\omega_{axial} \omega_{radial}^2)^{1/3}, \quad (2.42)$$

where $s_{TFradial/axial}$ is defined in Eqn. 2.23.

2.9 Computer control

We use two computers to run our experiment; one computer controls all of the timing, digital, analog, and GPIB commands, while the other is dedicated to running the camera. To initiate an image acquisition, the control computer externally triggers the camera, which in turn triggers the camera computer. It is useful to allocate these tasks to different computers so that the camera computer can analyze the data from the previous shot while the control computer moves on to the next shot. The preliminary analysis of each shot, which includes calculating cloud parameters, is completed in real time, greatly increasing the amount of data that can be compiled and digested in a day.

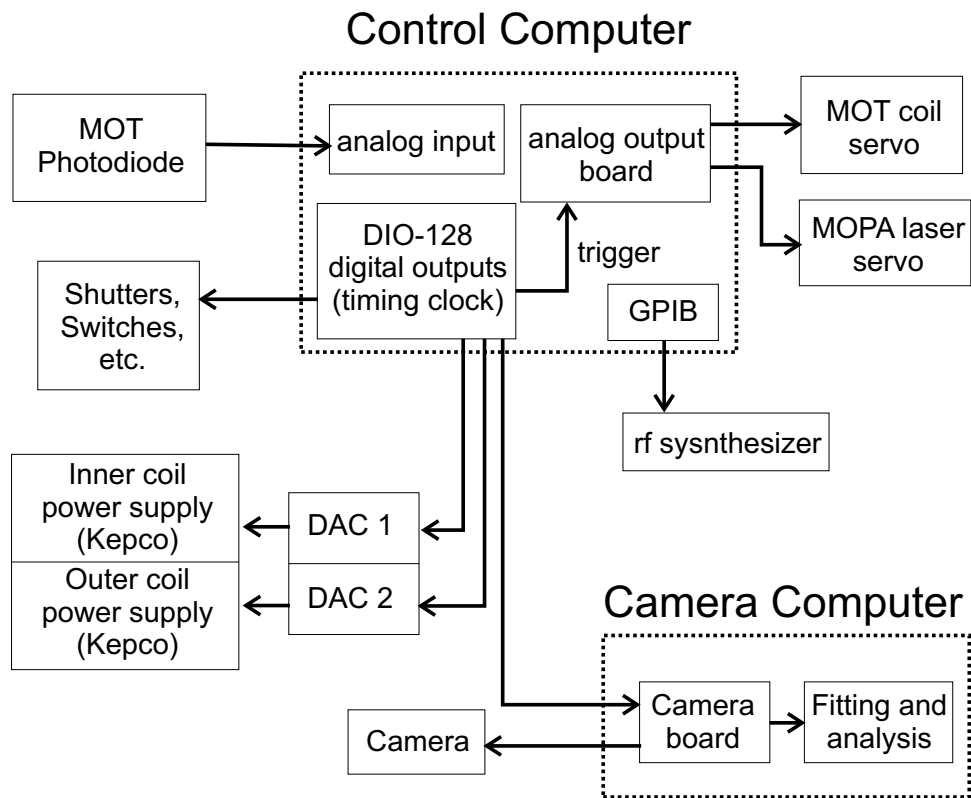


Figure 2.22: Computer control diagram showing the different computer boards and what they control.

A Bose-Einstein condensation experiment requires precise temporal control of

a variety of components. Most functions require timing resolution on the ms scale, but for certain key tasks, such as imaging, expansion, and microwave spectroscopy, we need timing on the microsecond scale. There are several basic types of outputs and inputs our control system needs to handle. We need TTL, analog, serial, and GPIB outputs. Most of the experiment is controlled by digital outputs, which control items such as shutters and switches. Our magnetic coils, both quadrupole and HIP trap, have servos that require analog voltage set points. The servo motor driving the track is controlled via a serial connection. We also have several instruments including rf and microwave synthesizers that use GPIB as the main mode of communication. Our only input port, excluding the camera, is an analog voltage from a photodiode that monitors the fluorescence from the atoms in the MOT. This input is fed into a multipurpose analog input board produced by National Instruments.

We have come up with a complete timing system that includes all the different I/Os. The programming software we chose is LabVIEW; it is a well known and easy to use software package. LabVIEW does have some limitations because it runs in a Windows environment. LabVIEW timing can vary by up to 10 ms shot-to-shot, because the operating system can interrupt the program at any time. Therefore we need another source to handle our precise timing. LabVIEW can still handle our imprecise events, such as GPIB commands, quadrupole coils current ramps, and track motion. Four computer boards and two external digital-to-analog converters (DACs) make up the control hardware.

We use a DIO-128 board produced by Viewpoint USA as the main clock in our system. It has 64 digital inputs and outputs and an internal oscillator that has 500 ns resolution [27]. We load an array of time stamps and port levels, which is the state (hi/low) of each digital port at each time stamp, into the first-in-first-out buffer on the board via LabVIEW. The buffer can hold up to 16000 words. When we want the timing sequence to start we send a trigger to the DIO-128 from within LabVIEW. From

that point on until the buffer is cleared, the board no longer communicates with the computer and thus is not susceptible to operating system interrupts.

We have two different methods to produce the analog output voltages required in the experiment. One way we produce an analog voltage is via a National Instruments analog output board, which resides inside the noisy environment of the computer. We use this board to control items that can handle the amount of noise produced, such as the quadrupole coil servos. The analog output board can store an array of voltage values and output them when triggered by the DIO-128. We create voltages with a lower level of noise using external 16 bit DACs digitally controlled by the DIO-128. The power supplies which drive the HIP trap coils are controlled with the low noise DACs.

Chapter 3

Driven two-level atom

The goal of this chapter is to describe how a two-level atom is affected by a coupling radiation field and a specific case of two hyperfine levels in ^{87}Rb coupled by a two-photon field. This chapter includes a basic description of Rabi flopping and Ramsey's method of separated oscillatory fields, which will be used as the method of spectroscopy for the following experiments. The goal of both of these methods of spectroscopy is to measure the energy difference between two atomic states. The choice of the two specific hyperfine levels and method of coupling them will also be discussed.

3.1 Rabi flopping

A two-level atom in an oscillating field can be described by a time dependent Hamiltonian, which is given by

$$H = \begin{pmatrix} \frac{\hbar\omega_0}{2} & \Omega_R e^{i\omega t/2} \\ \Omega_R e^{-i\omega t/2} & -\frac{\hbar\omega_0}{2} \end{pmatrix}, \quad (3.1)$$

where ω_0 is the the frequency difference between the two states, ω is the frequency of the driving field, Ω_R is the Rabi frequency, and t is time. The time-dependent Schrödinger equation with this Hamiltonian can be solved using the rotating wave approximation and the initial condition that the atom is in the ground state. The probability of being found in the excited state as a function of time is

$$P(t) = \left(\frac{\Omega_R}{\Omega'_R}\right)^2 \sin^2\left(\frac{\Omega'_R t}{2}\right), \quad (3.2)$$

where the effective Rabi frequency is

$$\Omega'_R = \sqrt{\Omega_R^2 + (\omega_0 - \omega)^2} \quad (3.3)$$

$$= \sqrt{\Omega_R^2 + \delta^2}, \quad (3.4)$$

and δ is the detuning from resonance. An atom in the field oscillates between the two states at the effective Rabi frequency, where the contrast of the oscillation is determined by the detuning. This sinusoidal population transfer is referred to as Rabi flopping.

Rabi's resonance method can be used to measure the energy splitting between two states. However, it has some limitations. Rabi's method probes the transition frequency while an atom is interacting with the coupling field, which can lead to frequency shifts from the AC Stark effect or additional phase noise. The frequency of the Rabi fringes depends not only on the frequency between the two states but also on the power of the driving field. Originally, in molecular beam experiments, the molecules would pass through a single region where they were illuminated by radiation and then be detected. It is reasonable to think that increasing the length of the interaction region l , would decrease the width of the spectral lines. However if the field is not uniform throughout the interaction region, increasing l will, at some point, lead to a broadening of the spectral lines. This broadening limits the resolution of Rabi spectroscopy. Ramsey's method of separated oscillatory fields was developed to overcome these limitations.

3.2 Ramsey spectroscopy

Ramsey's method uses two interaction regions separated by free propagation as seen in Fig. 3.1. Following a two-pulse Ramsey sequence, the probability of being in the excited state as a function of time, in the short intense pulse limit, with the approximations that $\delta \ll \Omega_R$ and $\tau \ll T$ is

$$P_t(T, \delta) = \frac{1}{2} + \frac{1}{2} \cos(\delta T), \quad (3.5)$$




Figure 3.1: Ramsey's method of separated oscillatory fields, where two pulses length τ are separated by free propagation length T .

where T is the free propagation time. A π pulse corresponds to 100% transfer between states, and a $\pi/2$ pulse creates an equal superposition of the two states. The two pulses do not need to be $\pi/2$ pulses. However this condition gives the highest signal-to-noise ratio and full contrast fringes. Note that the Ramsey fringe frequency only depends on the field detuning from resonance and thus on the energy difference of the atomic states and not on the power of the field, as long as $\tau\Omega_R$ is constant. The driving field must still be stable during the pulses; this is a shorter time and the constraint on field amplitude stability does not increase with increasing free propagation time.

Ramsey's method can be thought of as an interferometer in the following way. If an atom is initially in the ground state, the first $\pi/2$ pulse places the atom in an equal coherent superposition of the ground and excited states. The splitting is identical to the splitting by a beamsplitter in a optical interferometer. Next, during the propagation time, the relative phase of the two states evolves at a rate proportional to the energy difference between the states, and the coupling drive accumulates a phase ωT . The second $\pi/2$ pulse recombines the two states, interferometrically comparing the relative phase accumulated.

3.3 Two-photon transition in ^{87}Rb

We want to measure the phase properties of an ensemble of atoms to gain understanding of phase coherence and mechanisms that cause decoherence; these measurements will be described in depth in chapter 4. The absolute phase of an atom can not be measured but rather only the phase relative to another system. We study the phase properties of a system that is made up of two hyperfine ground states in ^{87}Rb . These states must have the same magnetic moment to first order to be trapped, overlapped spatially within the trap, and have nearly identical responses to magnetic field for spectroscopy.




Figure 3.2: Energy level diagram for the ground state of ^{87}Rb in a field of 3.2 G showing the relevant two-photon transition.

The two states we use are the $|1\rangle \equiv |F=1, m_f=-1\rangle$ and $|2\rangle \equiv |F=2, m_f=1\rangle$ states in the $5S_{1/2}$ manifold. They satisfy the requirement of having nearly identical magnetic moments. We can coherently transfer population between the two states via

$\frac{\mu_B}{h}$	1.399624624 (56) MHz/G
n_i	3/2
g_i	$0.9951414(10) \times 10^{-3}$ [28]
g_j	2.00233113(20)
ν_{hf}	6834.68261090434(3) MHz [29]

Table 3.1: Breit-Rabi equation constants for ^{87}Rb

a two-photon transition. We need two photons because the states differ by two units of angular momentum. We use a microwave photon (~ 6.8 GHz) and a rf photon (~ 3 MHz) to couple the two states as shown in Fig. 3.2. When the detuning from the intermediate state $|2, 0\rangle$ is large compared to the Rabi frequency for the one photon transition, the probability of making a transition to the intermediate state is small. In this limit the system can be thought of as just a two-level system with effective coupling given by the two-photon Rabi frequency, which is

$$\Omega = \frac{\Omega_{\mu wave} \Omega_{rf}}{2 \Delta_{IM}}, \quad (3.6)$$

where $\Omega_{\mu wave}$ is the Rabi frequency for the $|1, -1\rangle \rightarrow |2, 0\rangle$ transition, Ω_{rf} is the Rabi frequency for the $|2, 1\rangle \rightarrow |2, 0\rangle$ transition, and Δ_{IM} is the intermediate state detuning. The magnetic field dependence of this transition can be calculated using the Breit-Rabi formula, which is

$$\nu(B) = -\frac{\nu_{hf}}{2(2n_i + 1)} - g_i \mu_B m_f B \pm \frac{\nu_{hf}}{2} \sqrt{1 + \frac{4 m_f}{2 n_i + 1} x(B) + x^2(B)}, \quad (3.7)$$

where

$$x(B) = \frac{(g_j + g_i)\mu_B B}{\nu_{hf}}. \quad (3.8)$$

B is the magnitude of the magnetic field, and the constants for the equation are given in Table 3.1.

Chapter 4

Spectroscopy

This chapter, which was adapted from “Effect of cold collisions on spin coherence and resonance shifts in a magnetically trapped ultracold gas” published in Physical Review A [30], describes precision microwave spectroscopy of magnetically trapped, ultra-cold ^{87}Rb . We use this spectroscopy to characterize intra- and inter-state density correlations. The cold collision shifts for both normal and condensed clouds are measured. The results verify the presence of the sometimes controversial “factors of two,” in normal-cloud mean-field energies, both within a particular state and between two distinct spin species. One might expect that as two spin species decohere, the inter-state factor of two would revert to unity, but the associated frequency chirp one naively expects from such a trend is not observed in our data. We also measure coherence times of normal clouds as a function of potential inhomogeneity. Our simple model of decoherence, which ignores collisions, systematically underestimates the coherence times leading to the conclusion that collisions help to preserve coherence across a trapped sample.

In the cold collision regime, quantum statistical effects due to particle indistinguishability dominate scattering processes. The symmetrization requirement for identical particles in an ultra-cold Bose gas enhances the probability of finding two particles nearby, causing enhanced density fluctuations. At lower temperatures the statistical nature of the Bose gas causes the phenomena of Bose-Einstein condensation, where all

atoms in the condensate share the same wavefunction, suppressing the density fluctuations found in a noncondensed sample.

In other work, suppression of second-order density fluctuations in a condensate has been measured through analysis of the expansion energy of condensates [31, 32]. In a separate experiment, the suppression of third-order density fluctuations was probed by comparing the three-body loss rate of a condensate to that of a normal cloud [33]. The effect of cold collisions has also been measured as a density-dependent energy shift in atomic fountain clocks [34, 35, 36]. These shifts are quite small (~ 0.1 -10 mHz) due to the low densities at which the clocks operate, but are measurable because of their high precision. The uncertainty associated with these collisional shifts can be problematic; in fact the next generation of atomic fountain clocks are based on ^{87}Rb rather than ^{133}Cs because the collisional shift of ^{87}Rb is ~ 30 times smaller. In recent ultra-cold hydrogen experiments the cold collision shift provided the signature of Bose-condensation; below the BEC transition a large frequency shift of the 1S-2S transition was seen, reflecting the high density of the condensate [37].

4.1 Spectroscopy in a magnetic trap

We performed precision microwave spectroscopy on ultra-cold and condensed ^{87}Rb atoms confined in a magnetic trap. Due to the high densities achievable in a magnetic trap, the collisional energy shifts are 10^5 greater than those in ^{87}Rb atomic clocks, allowing a high-precision measurement of the shifts of the magnetically trappable states to be made with relative ease. The collisional shifts for both a normal and condensed sample were measured, providing a useful probe of the quantum statistics of the system. Additionally, magnetic confinement permits long interrogation times, allowing us to characterize temporal coherence of the normal cloud under various experimental conditions. Comparison of measured coherence times with a collisionless numerical simulation suggests that collisions preserve coherence in normal clouds.

Precision spectroscopy of trapped samples is difficult because atom trapping relies on spatial inhomogeneity of the atomic energy levels. Spatial inhomogeneity of the energy levels broadens the transition frequency, thus limiting the precision attainable through spectroscopy. In other work, this difficulty has been avoided by confining atoms in a blue-detuned optical dipole trap, where atoms spend little time interacting with the trapping fields [38]. In our work spatial inhomogeneity of the transition frequency was minimized through the use of a pair of energy levels which experience the same trapping potential. At a magnetic field of ~ 3.23 G, the $|1\rangle$ and $|2\rangle$ hyperfine levels experience the same first-order Zeeman shift. For a normal cloud at 500 nK, each energy level is Zeeman shifted by ~ 10 kHz across the extent of the cloud; however, at 3.23 G the *differential* shift of the two levels across the cloud is ~ 1 Hz. Compared to the differential Zeeman shift, the energy shift due to cold collisions is then a relatively large effect at high densities, making measurements of collisional shifts in this system possible. The small inhomogeneity allows for long coherence times, ~ 2 seconds and longer for low-density clouds, making this system attractive for precision measurements as well as for the study of condensate coherence in the presence of a thermal cloud.

A two-photon microwave-rf transition is used to transfer atoms between the $|1\rangle$ and $|2\rangle$ states. A detuning of 0.7 MHz from the $|2,0\rangle$ intermediate state provides a two-photon Rabi frequency of ~ 2.5 kHz. Ramsey spectroscopy of the $|1\rangle \rightarrow |2\rangle$ transition is performed by measuring the total number of atoms remaining in state $|1\rangle$ after two $\frac{\pi}{2}$ pulses separated by a variable time delay are applied [39]. The frequency of the resulting Ramsey fringes is the difference between the transition frequency ν_{12} and the two-photon drive frequency. We can measure local variations of ν_{12} (e.g. from magnetic or mean-field energy shifts) by detecting the number of atoms remaining in state $|1\rangle$ at specific spatial locations along the axis of the normal cloud Fig. 4.1a. We bin the cloud only in the axial direction because the collision rate is less than the radial trap frequency; therefore the atoms thermally average over the radial direction during

an experiment. For a particular peak density we measure the axial spatial variation in frequency by binning the cloud into several sections and determining a frequency for each bin Fig. 4.1b. This measurement is done at $B_{bias} = 3.23$ G to minimize the magnetic contribution to the inhomogeneity. The transition frequency vs. axial position along the cloud is shown in Fig. 4.1c.

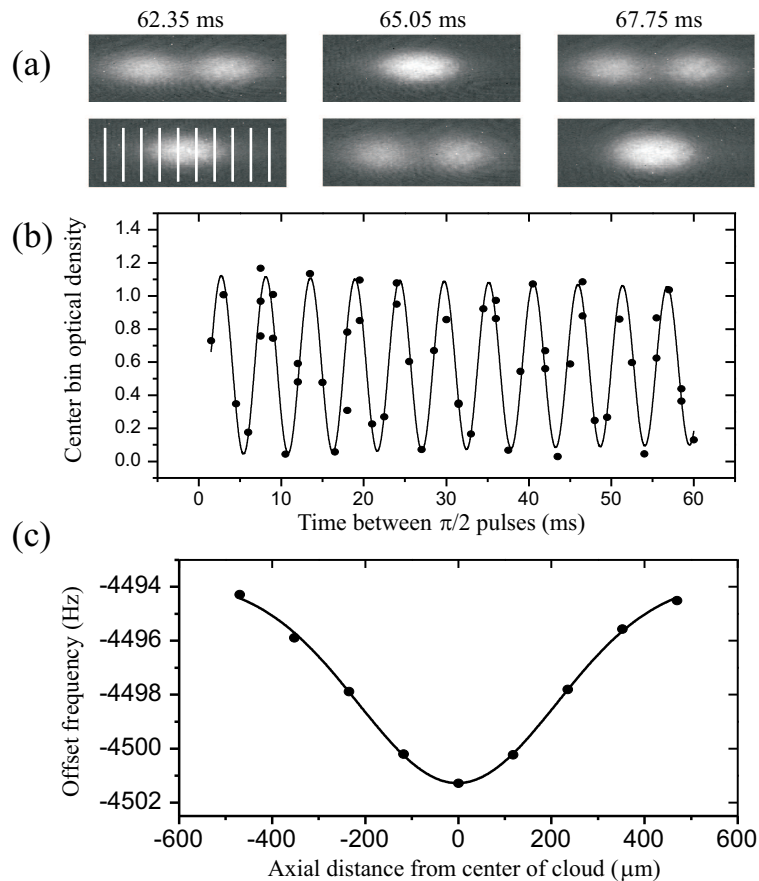


Figure 4.1: (a) Images of $|1\rangle$ atoms (upper) and $|2\rangle$ atoms (lower) after two $\frac{\pi}{2}$ pulses. The delay between the pulses is listed at the top. The edges of the cloud have become out of phase with the center. The white lines show how the cloud is divided up along the axial direction to measure the spatial dependence of the transition frequency. Each image is 1 mm on the horizontal (axial) dimension. (b) The center-bin $|1\rangle$ state optical density as a function of time between the two $\frac{\pi}{2}$ pulses. (c) The plot shows the spatial dependence of the transition frequency in the axial direction. Offset frequency is defined as the transition frequency - 6834682610.9 Hz. The solid line is a gaussian fit to the data. These data were taken at $B_{bias} = 3.23$ G.

By analyzing the spatio-temporal variations of ν_{12} , combined with the measured evolution of the $|1\rangle$ state after a single $\frac{\pi}{2}$ pulse, we were able to resolve spatially the evolution of spin waves [40], which will be discussed in depth in chapter 5. In order to perform measurements of ν_{12} insensitive to spin waves, one of the following two techniques was used. With one technique the entire cloud, rather than specific spatial locations, was monitored to average out the effects of spin waves. Alternatively, Ramsey spectroscopy was restricted to interrogation times short compared to the spin wave frequency [41].




Figure 4.2: Differential Zeeman shift at low magnetic fields for the $|1\rangle \rightarrow |2\rangle$ transition. The solid line is the predicted splitting from the Breit-Rabi formula. The inset plot expands the bias field region where most studies are performed.

One effect which shifts the transition frequency ν_{12} is the differential Zeeman shift. The Breit-Rabi formula predicts a minimum in ν_{12} at $B_0 = 3.228917(3)$ Gauss, thus the differential Zeeman shift between the $|1\rangle$ and $|2\rangle$ energy levels is first-order

magnetic field independent at $B = B_0$. The differential Zeeman shift about B_0 can be approximated as $\nu_{12} = \nu_{min} + \beta(B - B_0)^2$ [42]. Measuring ν_{12} for different magnetic fields allows us to calibrate our magnetic field from the expected dependence, see Fig. 4.2. By working in the vicinity of $B = B_0$ we greatly reduce spatial inhomogeneity of ν_{12} and also become first-order insensitive to temporal magnetic field fluctuations.

4.2 Cold-collision energy shift

A second effect which shifts ν_{12} arises from atom-atom interactions. In the s-wave regime, where the thermal de Broglie wavelength of the atoms is greater than their scattering length, atoms experience an energy shift equal to $\alpha \frac{4\pi\hbar^2}{m} an$, where α is the two-particle correlation at zero separation, n is the atom number density, a is the scattering length, and m is the atom mass. Therefore for a two-component sample the expected energy shift of each state is

$$\delta\mu_1 = \frac{4\pi\hbar^2}{m}(\alpha_{11}a_{11}n_1 + \alpha_{12}a_{12}n_2) \quad (4.1)$$

$$\delta\mu_2 = \frac{4\pi\hbar^2}{m}(\alpha_{12}a_{12}n_1 + \alpha_{22}a_{22}n_2), \quad (4.2)$$

where n_1 and n_2 are the $|1\rangle$ and $|2\rangle$ state density and

$$\alpha_{ij} = \frac{\langle \Psi_i^\dagger \Psi_j^\dagger \Psi_i \Psi_j \rangle}{\langle \Psi_i^\dagger \Psi_i \rangle \langle \Psi_j^\dagger \Psi_j \rangle}. \quad (4.3)$$

The shift of the transition frequency in Hz can then be written as

$$\begin{aligned} \Delta\nu_{12} &= \frac{2\hbar}{m}(\alpha_{12}a_{12}n_1 + \alpha_{22}a_{22}n_2 - \alpha_{11}a_{11}n_1 - \alpha_{12}a_{12}n_2) \\ &= \frac{\hbar}{m}n(\alpha_{22}a_{22} - \alpha_{11}a_{11} + (2\alpha_{12}a_{12} - \alpha_{11}a_{11} - \alpha_{22}a_{22})f), \end{aligned} \quad (4.4)$$

where $f = \frac{n_1 - n_2}{n}$ and $n = n_1 + n_2$.

For noncondensed, indistinguishable bosons, $\alpha = 2$ due to exchange symmetry; therefore $\alpha_{11}^{nc} = \alpha_{22}^{nc} = 2$ in a cold normal cloud (where the superscript c or nc refers to condensed or noncondensed atoms respectively). Distinguishable particles do not

maintain exchange symmetry, making $\alpha_{12}^{nc} = 1$ for an incoherent two-component mixture. However if a two-component sample is prepared by coherently transferring atoms from a single component, such as in Ramsey spectroscopy, then the excitation process maintains exchange symmetry, and we might expect $\alpha_{12}^{nc} = 2$ [43]. In this scenario the collisional shift should be calculated using $\alpha_{11}^{nc} = \alpha_{22}^{nc} = \alpha_{12}^{nc} = 2$, leading to a predicted frequency shift of

$$\Delta\nu_{12} = \frac{2\hbar}{m}n(a_{22} - a_{11} + (2a_{12} - a_{11} - a_{22})f). \quad (4.5)$$

This result can also be obtained by solving the transport equation [44, 45]. From spectroscopic studies [46] the three ^{87}Rb scattering lengths of interest have been determined to be $a_{22} = 95.47a_0$, $a_{12} = 98.09a_0$, and $a_{11} = 100.44a_0$, where a_0 is the Bohr radius. The frequency shift can then be written as

$$\Delta\nu_{12} = \frac{2\hbar}{m}a_0n(-4.97 + 0.27f). \quad (4.6)$$

If on the other hand the $|1\rangle$ and $|2\rangle$ states do *not* maintain exchange symmetry, such that $\alpha_{12}^{nc} = 1$, then the frequency shift would instead be

$$\Delta\nu_{12} = \frac{2\hbar}{m}a_0n(-4.97 - 97.82f). \quad (4.7)$$

These two models are clearly distinguished by the dependence of ν_{12} on f .

When we perform Ramsey spectroscopy with a pair of $\frac{\pi}{2}$ pulses, the populations of the $|1\rangle$ and $|2\rangle$ states are equal, and thus $f = 0$ during the interrogation time. From Eq. 4.4 it is apparent that with $f = 0$ the collisional shift is sensitive only to α_{ii}^{nc} and a_{ii} terms. For these measurements the bias field was set to B_0 , and the transition frequency was measured for a range of densities. To adjust density of the sample, the number of atoms in the initial MOT load was varied. All normal cloud data was taken at the same temperature of 480 nK, and all condensate data was taken with high condensate fractions in order to minimize effects due to the normal cloud. The density for the normal cloud was found by fitting Gaussian profiles to absorption images of the clouds

and extracting the number, temperature, and density. To measure condensate density, Thomas-Fermi profiles were fit to absorption images of the condensates and the total number, N_0 , in the condensates and the Thomas-Fermi radius along the long axis, Z , were extracted.




Figure 4.3: Measurement of the cold collision shift. Solid and open circles represent measurements of the normal cloud and condensate respectively. The solid line is a fit to the normal cloud data $\Delta\nu_{12} = 0.1(0.4) - 3.9(0.3)10^{-13}n$; the dashed line is a fit to the condensate data $\Delta\nu_{12} = -0.1(1.4) - 1.9(0.2)10^{-13}n$, where $\Delta\nu_{12}$ is in Hz and n is in cm^{-3} .

The results of this measurement are shown in Fig. 4.3. Comparing the collisional shift measured for the normal cloud to that measured for a condensate gives $\alpha_{ii}^{nc}/\alpha_{ii}^c = 2.1(2)$. If instead we *assume* $\alpha_{ii}^{nc} = 2$ and $\alpha_{ii}^c = 1$, then the data for both the condensate and normal cloud can be used to obtain a value for the difference in scattering lengths of $a_{22} - a_{11} = -4.92(28)a_0$, in agreement with values determined from molecular spectroscopy [46].

Many systematics can plague density measurements made through absorption imaging. In order to test independently our density calibration for both the nor-

mal and condensed samples, each of which can suffer from different errors, we used the Bose-condensation phenomenon. The density of normal clouds was tested through measurement of the critical temperature, and condensate density was tested with the Thomas-Fermi approximation. Assuming that disagreements are due only to errors in estimation of atom number, the worst case scenario, leads us to reduce normal cloud density by 11(4)% and increase condensate density by 11(3)%. Adjusting the cold collision shifts accordingly would yield a worst-case corrected value of $\alpha_{ii}^{nc}/\alpha_{ii}^c = 1.7(2)$ [47]. The adjusted normal and condensate density shifts can be combined as above to give a value for the difference in scattering lengths of $a_{22} - a_{11} = -4.85(31)a_0$; not significantly different from our unadjusted measurement.

The remaining significant systematic is atom loss due to $|2\rangle$ - $|2\rangle$ collisional dipolar relaxation. In order to minimize effects of this loss, interrogation times were kept as short as possible. Nevertheless for the highest density condensate measurements the $|2\rangle$ - $|2\rangle$ loss causes the total density to drop by 3% in 20 ms, the maximum interrogation time. For all other densities the loss was no larger than this, and in most cases much smaller. Finally, the $|1\rangle$ and $|2\rangle$ states begin undergoing spatial separation in the condensate after the first $\frac{\pi}{2}$ pulse [48]; however, the timescale for the separation is much longer than our 20 ms interrogation time.

Exchange symmetry between the $|1\rangle$ and $|2\rangle$ states can be tested by working at a fixed density and varying the relative $|1\rangle$ to $|2\rangle$ population by varying the length of the first Ramsey pulse [49]. In this case the first term in Eq. 4.4 will be constant and the measurement will test α_{12}^{nc} and a_{12} as well as the α_{ii}^{nc} and a_{ii} terms (see Eq. 4.6 and 4.7). To minimize systematics the interrogation times were kept short, making precise frequency determination difficult. Nevertheless, our measurement (Fig. 4.4) indicates $\alpha_{12}^{nc}/\alpha_{11,22}^{nc} = 1.01(2)$, where we have used the spectroscopically determined scattering lengths. This clearly indicates that exchange symmetry is maintained between the $|1\rangle$ and $|2\rangle$ states. A similar measurement was made on the




Figure 4.4: Testing the exchange symmetry between the $|1\rangle$ and $|2\rangle$ states. The transition frequency is measured as $f = \frac{n_1 - n_2}{n}$ is varied for a normal cloud at fixed peak density of $7 \times 10^{12} \text{ cm}^{-3}$ and temperature of 510 nK. The solid line is the fit, which yields $\alpha_{12}^{nc}/\alpha_{11,22}^{nc} = 1.01(2)$, which is to say, inter- and intra-state density correlations are quite accurately the same. The dotted line indicates the expected slope for $\alpha_{12}^{nc}/\alpha_{11,22}^{nc} = 1/2$.

$|F = 1, m_f = 0\rangle \rightarrow |F = 2, m_f = 0\rangle$ transition by Fertig and Gibble [35].

4.3 Missing frequency chirp

As a thought experiment, imagine distinct thermal populations of $|1\rangle$ and $|2\rangle$ atoms, separately prepared, then mixed together, with the energy of interaction (proportional to α_{12}^{nc}) measured for instance calorimetrically. Surely in this case the density fluctuations in state $|1\rangle$ and in state $|2\rangle$ would be uncorrelated, and α_{12}^{nc} would be determined to be 1, not 2. We lack the experimental sensitivity to make such a calorimetric measurement, and our Ramsey-fringe method of measuring energy differences obviously would not work for incoherent mixtures. We speculated, however, that if $\alpha_{12}^{nc} = 2$ for coherent superpositions, and if $\alpha_{12}^{nc} = 1$ for incoherent mixtures, then for partially decohered samples, α_{12}^{nc} would take on some intermediate value. So by performing a




Figure 4.5: A data set of Ramsey fringes probing for frequency shifts as a function of coherence. For this measurement normal clouds at a temperature of 480 nK and a peak density of $3.2 \times 10^{13} \text{ cm}^{-3}$ were used.

measurement similar to that in Fig. 4.4 we might expect to see a more negative slope for a partially decohered sample; alternatively a frequency chirp in the Ramsey fringes may be seen as the sample decoheres.

We probed the time evolution of α_{12}^{nc} in a way similar to Fig. 4.4; however rather than varying f we set $f \simeq 0.8$ then measured ν_{12} with long interrogation times, looking for a frequency chirp as the fringe contrast decreased. This method has the advantage that there is a relatively small $|2\rangle$ state population, so effects arising from $|2\rangle$ loss are minimized. Seven data sets were taken for this measurement; an example is shown in Fig. 4.5. By allowing a linear frequency chirp in the fit of the Ramsey fringes, the frequency shift can be constrained to $-0.2(3)$ Hz by the time the fringe contrast has

reduced to $1/e$ [50]. However if we hypothesize that α_{12}^{nc} goes from 2 to 1 linearly as fringe contrast goes from 100% to 0% we would expect a frequency shift of $-20(2)$ Hz as the fringe decayed, while the experimental limit is 40 times smaller. The lack of chirp disagrees with the traditional notion that the frequency of the transverse spin precession is a direct measurement of the energy difference between spin up and spin down. Clearly this appealing but unrigorous model is far too naive.

4.4 Coherence of normal clouds

Ramsey spectroscopy not only allows us to probe the energy difference between the two states but also permits the measurement of the coherence between the two states. Coherence measurements were performed using the same time domain method used to measure ν_{12} ; however, interrogation times were extended until the Ramsey fringe contrast was lost. The resulting data was fit to a $e^{-(t/\tau)^2}$ decay, where τ is the coherence time. Additionally, fitting allowed for both a loss in total atom number and a linear frequency chirp of ν_{12} . To ensure that the $1/e$ atom loss times were much longer than the coherence times, fractional transfers ($f \simeq 0.8$) to the $|2\rangle$ state were used for the high density data points. The results of coherence measurements for different magnetic fields and at three different densities are shown in Fig. 4.6.

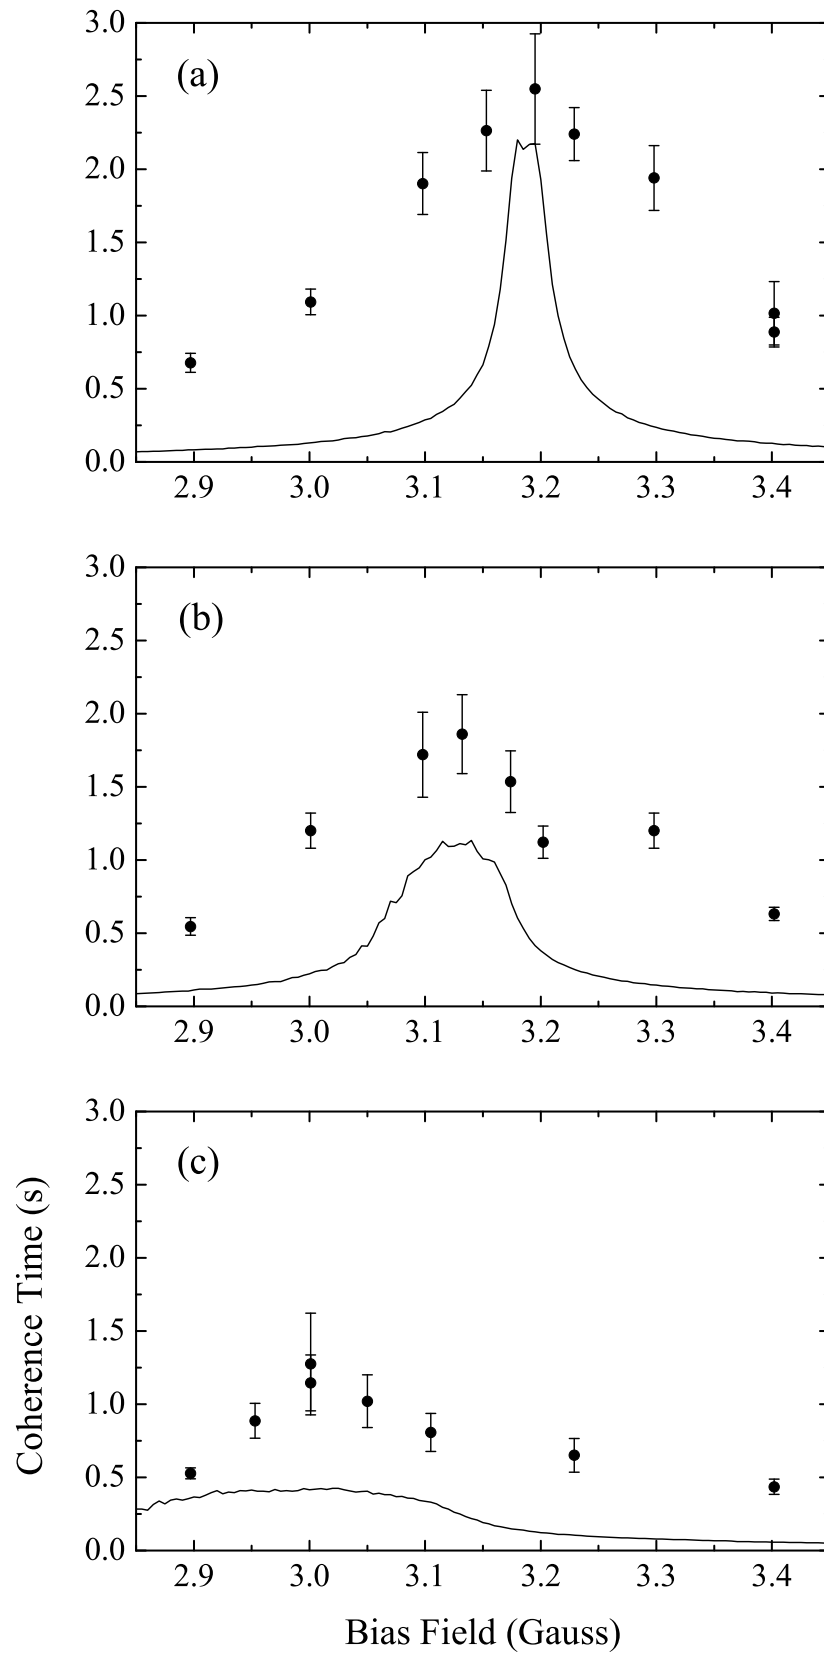


Figure 4.6: Coherence times in a normal cloud. The temperature was ~ 480 nK and the normal cloud peak density for each plot is: (a) $4 \times 10^{12} \text{ cm}^{-3}$; (b) $1.3 \times 10^{13} \text{ cm}^{-3}$; (c) $3.2 \times 10^{13} \text{ cm}^{-3}$. The solid line in each plot corresponds to the $1/e$ times obtained from a numerical simulation (see text). The noise in the simulation data is due to the finite number of particles used and the random initialization.

We expect that the primary source of decoherence is inhomogeneity in ν_{12} across the cloud, to which both the differential Zeeman and collisional shifts contribute. The collisional shift scales directly with density, and thus provides a Gaussian-shaped ν_{12} profile across the cloud. Inhomogeneity arising from the differential Zeeman effect depends on the bias field. The magnitude of the magnetic field near the bottom of the trap can be written as $B(z) = \frac{B''}{2}z^2 + B_{bias}$, thus the inhomogeneity due to the Zeeman shift is $\nu_{12} \propto \frac{B''^2}{4}z^4 + (B_{bias} - B_0)B''z^2$. By setting $B_{bias}(>, <, =)B_0$ the curvature of ν_{12} can be adjusted to be positive, negative, or nearly zero, respectively. This allows inhomogeneity due to the collisional shift to be roughly cancelled by an opposing Zeeman inhomogeneity Figs. 4.7 and 4.8 [51].

This effect of this cancellation can be seen in Fig. 4.6; as the density is increased from (a) to (c), the inhomogeneity induced by the collisional shift increases, so that a larger opposing Zeeman inhomogeneity is necessary for cancellation. Therefore the bias field for peak coherence time decreases as cloud density increases.

In an attempt to compare the measured decoherence times to the known spatial inhomogeneity of the transition frequency, we performed the following numerical simulation: consistent with a Maxwell-Boltzmann distribution, we randomly assign initial positions and velocities to ten thousand simulated atoms. Ignoring the effects of collisions, we calculate the three-dimensional trajectory of each atom for several simulated seconds, keeping track of the time integral of ν_{12} along the trajectory. At each point in time, we calculate the spatially integrated transverse magnetization and, as inhomogeneities cause this magnetization to wash out, find the time it takes the integrated transverse magnetization to reduce to $1/e$ of its original value [52]. This model should correctly account for the effect of motional averaging except that all collisional effects are explicitly excluded. The resulting modelled damping times are plotted as a solid line along with the experimentally measured damping times in Fig. 4.6. While the model does a reasonable job predicting the value of the bias field for which the coherence

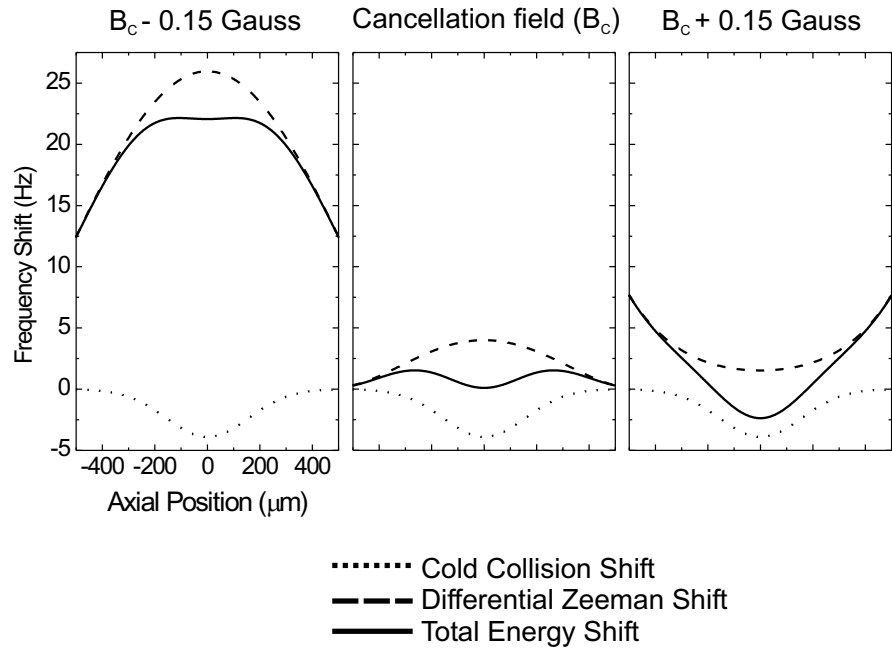


Figure 4.7: The frequency shift across the axial direction of the cloud is plotted for three different magnetic bias fields. The dotted curves indicate the the cold collisional energy shift, the dashed curves indicate the Zeeman energy shift, and the solid line is a sum of the two energy shifts. The cancellation field is shown in the center where there is a minimal variation of the frequency across the cloud. Increasing the bias field causes the energy shifts to have the same curvature and produces an overall larger variation of frequency across the cloud; reducing the bias field below B_C causes the Zeeman shift contribution to overwhelm the cold collision shift and thus dominate the energy vs. position profile.

time peaks, it consistently underestimates (in some cases by a factor of eight) the actual value of the coherence time. Our model neglects both velocity-changing collisions and the exchange-type collisions that lead to spin waves; it appears that these effects contribute significantly to preserving coherence across the trapped atom cloud.

The extreme aspect ratio of our trapped cloud complicates a proper quantitative analysis of the effects of collisions on coherence. Along the axial direction, other work has shown that the effects of spin-waves are to keep local magnetization across the cloud from straying too far from its spatially averaged value [see in particular Fig. 5.3(a) in chapter 5]. In the radial directions, the motional oscillation frequency exceeds the mean-




Figure 4.8: The transition frequencies at different positions along the cloud are plotted as a function of B_{bias} . The cloud has a peak density of $4 \times 10^{13} \text{cm}^{-3}$. Offset frequency is defined as the transition frequency - 6834682610.9 Hz . The frequency shift at 3.23 G (negligible magnetic contribution to the inhomogeneity) has a gaussian shape. At $B_{\text{bias}} = 3.37 \text{ G}$ the magnetic shift and the mean-field shift are in the same direction and therefore the inhomogeneity is across the cloud is increased. When the cloud is in a field of 2.83 G inhomogeneity due to the mean-field shift is almost completely cancelled by the magnetic shift. Data points are connected to guide the eye.

field exchange frequency in the cloud, and thus the effects of spin-waves are probably not relevant. On the other hand, velocity-changing collisions are likely important – in their absence, atoms with small transverse energies would stay near the axis of the trap, while atoms with large transverse energies would preferentially sample the larger magnetic fields and lower densities further from the axis. Velocity-changing collisions which re-randomize the transverse trajectories of these different classes of atoms (before they have a time to accumulate a radian or more of relative phase-difference) will serve to extend the coherence time of the sample. In the highest density data set presented [Fig. 4.6(c)] the mean elastic collision rate was 74 Hz , which should be compared to the

measured coherence times of around 0.5 to 1.2 seconds.

4.5 Coherence of condensates

We were also able to explore the coherence of nearly pure condensates. There are two standard data taking methods to measure coherence using Ramsey spectroscopy. One is to take equally spaced points and fit the data to an exponentially decaying sine wave (Fig. 4.9). This measurement technique is sensitive to phase noise because phase noise will artificially decrease the fitted coherence time. The functional form of the fringes is derived as follows. The same result can be more rigorously derived using a density matrix approach. The initial state after the first $\pi/2$ pulse is

$$|\psi\rangle = \frac{1}{\sqrt{2}} [|1\rangle + |2\rangle]. \quad (4.8)$$

After some time t the state is modified by loss in both states and phase evolution due to the energy difference between the states and can be written as

$$|\psi\rangle = \frac{1}{\sqrt{2}} [e^{-t/2\tau_1} |1\rangle + e^{-t/2\tau_2} e^{i\phi} |2\rangle], \quad (4.9)$$

where τ_i is the lifetime of state i . The approximation of an exponential loss rate is not entirely correct. The main loss rate is 2-body for the $|2\rangle$ state and 3-body for the $|1\rangle$. However the condensate density distribution is changing in time due to the mean-field separation effect and therefore the loss rate is just modelled as exponential. After some time t a second $\pi/2$ pulse is applied, producing the state:

$$|\psi\rangle = \frac{1}{2} [e^{-t/2\tau_1} (|1\rangle + |2\rangle) + e^{-t/2\tau_2} e^{i\phi} (-|1\rangle + |2\rangle)]. \quad (4.10)$$

From this state we can calculate the probability of being in the $|1\rangle$ state as a function of time, which is

$$P_{|1\rangle}(t) = \frac{1}{4} [e^{-t/\tau_1} + e^{-t/\tau_2} - 2 \cos \phi e^{-(t/2\tau_1 + t/2\tau_2)}]. \quad (4.11)$$

The $|1\rangle$ state has a much longer lifetime than the $|2\rangle$ state so the term $e^{-t/2\tau_1}$ in Eqn. 4.11 can be set equal to unity, and an exponential decoherence term can be added with time constant τ_{coh} . The probability of being in the $|1\rangle$ state then is

$$P_{|1\rangle}(t) = \frac{1}{4}[1 + e^{-t/\tau_2} - 2e^{-t/\tau_{coh}} \cos \phi e^{-t/2\tau_2}]. \quad (4.12)$$

The function we using in our fitting routine is

$$y(t) = A\frac{1}{4}[1 + e^{-t/\tau_2} - 2e^{-t/\tau_{coh}} \cos(2\pi(\nu + c t)t + \theta)e^{-t/2\tau_2}], \quad (4.13)$$

where the free parameters are an amplitude A , a $|2\rangle$ state lifetime τ_2 , a decoherence time τ_{coh} , a fringe frequency ν , a linear frequency chirp c , and a phase offset θ .

The other method for measuring coherence time, which is insensitive to phase noise, is to take data (~ 10 points) over a time interval equal to a single fringe at several different propagation times (Fig. 4.10a). Then the standard deviation Σ , is calculated for each fringe and fit to a decaying exponential in quadrature with a background noise level (Fig. 4.10b). Σ is

$$\Sigma(t) = \sqrt{\frac{1}{8}e^{-2t/\tau_2}}. \quad (4.14)$$

The functional form we use in the fitting routine is

$$y(t) = \sqrt{\frac{A^2}{8}e^{-2t/\tau_2} + y_0^2}, \quad (4.15)$$

where y_0 is the background noise level.

Measuring condensate coherence time is difficult because of the high loss rate in the $|2\rangle$ state and phase noise due to shot-to-shot number fluctuations. The high density in the condensate causes large dipolar relaxation rates and reduced lifetime of the $|2\rangle$ state. Therefore we lose signal rapidly as the $|2\rangle$ atoms are lost from the trap. We can reduce this loss rate by not applying a $\pi/2$ pulse but rather transfer only 20% of the atoms to the $|2\rangle$ state, which decreases the density and thus the loss in this state.




Figure 4.9: Ramsey fringes of the center bin of a condensate with peak density of $8.9(3)\times 10^{13}$ with a condensate fraction of 0.70(3). 24% of the condensate is initially placed in the $|2\rangle$ state. The fit gives a value for $\tau = 268(28)$ ms. (see text for discussion)

The other complication with measuring condensate coherence is phase noise induced by shot-to-shot number fluctuations. The cold-collision frequency shift for typical large condensates with densities of 10^{14} cm^{-3} is ~ 20 Hz. A 5% variation of the density will then cause a phase shift of 0.5 rad at 0.5 seconds. Therefore we must make condensates with small and repeatable density to measure coherence times with minimal phase noise. We create small stable condensates by first evaporating normally to produce a medium-sized condensate (10^5 atoms); then leave the rf on close to the edge of the condensate for several seconds. Any noise in the bias field will cause the rf knife to cut into the condensate. However the noise will average out over several seconds thus producing a small, consistently-sized condensates.

One example of a coherence measurement of a condensate is shown in Fig. 4.9. The fringe contrast and the mean decay at the same rate, which is not predicted by our




Figure 4.10: Alternate method of taking coherence data.(a) Ramsey fringes taken over one fringe at six different propagation times. (b) Standard deviation of each of the six fringe sets normalized by the mean of the data points within that set.

simple model. Our model predicts that, in absence of additional decoherence, the fringe contrast should decrease at half the rate of the mean. One might assume initially that there is just additional decoherence, but there are two reasons why we think this might not be the case. First the decay rate for both the contrast and the mean coincide with an independent measurement of the $|2\rangle$ state decay rate. Second the contrast and the mean decay at the same rate over a wide parameter space, which includes a variety of condensate fractions and number densities. One would think the additional decoherence would change as these parameters are varied.

4.6 Speculation on precision metrology

It is interesting to consider the usefulness of magnetically trapped atoms for precision metrology. Peak coherence times of approximately 2.5 s were realized with cold, low density samples. An interrogation time of 2.5 s provides a 0.2 Hz linewidth, which naturally leads one to consider using such a system for precision measurements. By working at $B_{bias} = B_0$, coherence times are slightly reduced. However, perturbations of ν_{12} due to magnetic field fluctuations become very small, on the order of 4 mHz for current typical experimental conditions. With careful design of the confinement coils, the current supply, and magnetic shielding of external fields, it should be possible to suppress fluctuations of the bias field below 1 mG, reducing Zeeman-induced frequency shifts to below 0.1 mHz. Perturbations of ν_{12} due to the cold collision shift are more significant; shot-to-shot density fluctuations will introduce frequency noise, so it is advantageous to work with the minimum possible density. However as atom number is reduced the maximum signal-to-noise ratio will decrease due to shot-noise [53]; therefore the optimum strategy is to work with an atom number such that the frequency uncertainty due to shot-noise is on the order of the uncertainty due to density fluctuations.

For example a normal cloud of 400 nK and a peak density $1.5 \times 10^{12} \text{ cm}^{-3}$ has 6×10^4 atoms. The shot-noise-limited signal-to-noise ratio is then 245:1. With an interrogation time of 1 second the single-shot statistical uncertainty is then 0.65 mHz; including the effects of decoherence and atom loss will increase this to approximately 0.9 mHz. Assuming that shot-to-shot number fluctuations are also shot-noise limited, then statistical uncertainty from the density shift is 0.84 mHz. Combining these gives a total single-shot uncertainty of 1.24 mHz; with our current 30 second evaporation time, the duty-cycle is such that an absolute precision of $6.8 \text{ mHz}/\sqrt{\text{Hz}}$ can be realized. This corresponds to a relative precision of $1 \times 10^{-12} / \sqrt{\text{Hz}}$, which in terms of measurement precision does not reach the level of atomic fountain clocks. This system however has

the advantage that small energy shifts can be measured in a compact, stationary spatial position; a 400 nK cloud occupies approximately only a $1040 \times 32 \times 32 \mu\text{m}$ region of space. It is certainly feasible to perform spectroscopy at this level within $100 \mu\text{m}$ of a surface, which might allow the measurement of short range atom-surface interactions.

We have demonstrated precise spectroscopy in an ultra-cold magnetically trapped gas. This permitted measurement of the cold collision shift in both a condensate and a normal cloud, allowing a probe of the quantum statistics of the system. Working at low densities minimizes the effect of the collisional shift, allowing long coherence times and precise determination of ν_{12} ; however the measurement of any quantity not related to atom-atom interactions will at some level be limited by this shift. An intriguing alternative would be to instead use a fermionic atom, which should have no collisional shift, and thus density induced frequency noise will not be an issue. On the other hand, the lack of collisions may also lead to more rapid decoherence: collisions appear to preserve our bosonic system from the decohering effects of spatial inhomogeneity.

Chapter 5

Spin waves

This Chapter is adapted from “Observation of anomalous spin-state segregation in a trapped ultra-cold vapor” and “Spatial resolution of spin waves in an ultra-cold gas” published in Physical Review Letters [51, 40].

Macroscopic collective behavior can arise in a quantum gas even above the onset of degeneracy. When indistinguishable atoms collide, the scattering can be significantly altered by the need to symmetrize the scattered waves. Examples of quantum scattering effects are seen in the strong polarization dependence of heat conduction and viscosity coefficients in spin polarized ^3He [54] and of thermalization rates in nondegenerate Fermi gases [55]. Quantum collisional effects can also produce spatio-temporal spin oscillations known as spin waves. In this Chapter we present spatially resolved images of coupled transverse and longitudinal spin perturbances propagating through a magnetically trapped atomic cloud. Frequencies and damping rates for the standing-wave excitations are studied as a function of density and temperature.

Although the concept of spin waves in ferromagnets dates back to Bloch’s predictions in 1930 [56], the theory of spin waves in liquids and in dilute gases was not formulated until much later [57]. The first evidence for spin waves (in the form of extra resonances in nuclear magnetic resonance spectra) occurred in spin-polarized hydrogen and soon after in polarized ^3He and dilute mixtures of ^3He in ^4He [58]. On a microscopic level, spin waves arise from interference effects in lowest partial-wave collisions. When

identical particles collide, the 180° backscattering event causes a scattered atom to propagate along a trajectory indistinguishable from that of either a forward scattered atom or an unscattered atom. When the two atoms are spin aligned, the backward scattering contribution is indistinguishable from the forward scattering contribution, which adds a factor of two to the mean-field collisional energy shift. When the spins are antiparallel, the backward scattering event is distinguishable, and the mean-field shift arises from the forward scattering event only. In the case of intermediate spin alignment, the backward scattering event is only partially distinguishable from the unscattered event, and one needs to add coherently two unparallel spin amplitudes, leading to a rotation of the spins of the scattered atoms. The cumulative effect of many such spin-rotating collisions is such that inhomogeneities of spin propagate like waves rather than diffusively. The characteristic frequency scale for the spin rotation effect is the exchange collision rate, $\omega_{\text{exch}} = 4\pi\hbar an/m$, where m , n , and a are the mass of the atoms, the number density, and the s-wave scattering length respectively.

5.1 Theoretical description of spin waves

The proceeding theoretical description follows a similar formalism to that of J. Williams and colleagues [45]. Ultimately we want to calculate the spin density as a function of position in the cloud and time. We start with the Hartree-Fock mean-field theory of a Bose gas, which leads to the collisionless Boltzmann equation. The solution to the Boltzmann equation predicts spin waves, but because it does not include collisions, it can not predict damping of the spin waves. The collision integrals will not be dealt with in this text but can be found in reference [59]. The Hamiltonian for two-level atoms in a trap is

$$H = \left(\frac{-\hbar^2}{2m} \nabla^2 + U_0(\mathbf{r}) \right) \hat{I} + \frac{\hbar}{2} \vec{\Omega}(\mathbf{r}) \cdot \vec{\sigma}, \quad (5.1)$$

where the first two terms are just the center of mass contribution including the kinetic energy and the potential, which is

$$U_0(\mathbf{r}) = U_{ext}(\mathbf{r}) + \frac{4\pi\hbar^2}{m} [a_{11} n_1 + a_{22} n_2 + \frac{a_{12}}{2}(n_1 + n_2)]. \quad (5.2)$$

U_{ext} is the magnetic trapping potential, and the other terms in Eqn. 5.2 are the mean field potential terms. For our experiments $4\pi\hbar^2 an/m \ll k_B T$ so the mean field terms can be ignored. The last term in the Hamiltonian is the contribution from the spin, where $\vec{\sigma}$ are the Pauli matrices. The coupling field, a combination of the drive field, energy difference from the mean field and Zeeman energy shifts as well as from the spin density, is

$$\vec{\Omega} = \vec{\Omega}' + \frac{2\pi\hbar^2 a_{12}}{m} \vec{S}(\mathbf{r}, t), \quad (5.3)$$

where

$$\vec{\Omega}' = [\Omega_x, \Omega_y, \Delta_{BR}(\mathbf{r}, t) + \Delta_{MF}(\mathbf{r}, t)], \quad (5.4)$$

where the components of Ω' are in Bloch space. The hyperfine splitting has been removed by taking the rotating wave approximation. Δ_{BR} is the differential Zeeman shift calculated from the Breit-Rabi formula, and the differential mean-field energy is

$$\Delta_{MF} = \frac{4\pi\hbar^2}{m} [2a_{11} n_1 - 2a_{22} n_2 + 2a_{12}(n_1 - n_2)], \quad (5.5)$$

where n_i is the density in state i . The factor of two in front of the a_{12} comes from the addition of the forward and backward scattering events as described previously in Chapter 4. Ω_i is the resonant Rabi frequency about the axis i and zero when the coupling field is off, such as after the $\pi/2$ pulse. The second term in Eqn. 5.3 is responsible for spin waves and comes from the extra factor of two in the mean-field interaction for identical particles. The spin density as a function of position in the cloud and time is

$$\vec{S}(\mathbf{r}, t) = \frac{\int d\mathbf{p} \times \vec{\sigma}(\mathbf{r}, \mathbf{p}, t)}{2\pi\hbar^3}, \quad (5.6)$$

where $\vec{\sigma}(\mathbf{r}, \mathbf{p}, t)$, the spin distribution function, can be obtained from the spin Boltzmann equation. The longitudinal component of the spin density S_z is the difference in

the population of the two states $[n_1(\mathbf{r}, t) - n_2(\mathbf{r}, t)]$. S_x and S_y are the real and imaginary parts of the internal coherence.

Now that the problem is defined, we can numerically integrate the kinetic equation to find the spin distribution function and subsequently the spin density, which was done by J. Williams and colleagues[45]. The kinetic equation, ignoring collisions, is

$$\frac{\partial \vec{\sigma}}{\partial t} + \frac{\mathbf{p}}{m} \cdot \nabla \vec{\sigma} - \nabla U_{ext} \cdot \nabla_p \vec{\sigma} - \vec{\Omega} \times \vec{\sigma} = 0, \quad (5.7)$$

where the second and third terms are the free streaming evolution, and the last term depicts the spin precessing about the local field, which is a combination of the driving field and the local spin.

Another slightly more intuitive way to understand the origin of spin waves is to think about the spins of each atom precessing about the total spin during a collision. When two atoms with different spins collide, their total spin must be conserved after the collision. However they can each precess about the total spin and thus have a different spin after the collision. This is depicted graphically in Fig. 5.1. The initial direction of




Figure 5.1: Intuitive picture of spin rotation during collisions. Atoms initially with different relative phases precess about the total spin during a collision.

the spin is dependent on the “history” or previous trajectory of each atom. Imagine a test atom coming from the center of the cloud with a certain spin; as it moves through the cloud its spin rotates because of its interaction with the average local spin of the cloud. A number of such interactions, because they are coherent, build up a coherent oscillation.

5.2 Description of system

The experimental apparatus and spectroscopy method used to prepare and probe the ultra-cold gas are described in Chapters 2-4. Typical atom cloud parameters for these spin wave studies are $n \sim 10^{13} \text{ cm}^{-3}$ and $T \sim 600 \text{ nK}$, several times the condensation temperature. The radial frequency is high enough that for the effects considered here it is sufficient to average the cloud over the radial dimensions, effectively reducing the cloud to a one dimensional density distribution. The atomic hyperfine states are $|1\rangle$ and $|2\rangle$, which together make up a pseudo-spin doublet [60]. The spin state is manipulated using the two-photon coupling transition described in Chapter 3.

The transition frequency $\Delta(z)$ is in general a function of the axial position z , due to the spatial inhomogeneity of the Zeeman shift and of the density-dependent mean-field shift. As stated earlier in Chapter 4, we can tune the bias magnetic field so as to cause the Zeeman shifts to cancel in whole or in part the mean-field shifts. As a consequence, the spatial curvature of the frequency, $\partial^2\Delta(z)/\partial z^2$, is a controllable parameter of the experiment which may be tuned in real time, even as a spin wave is propagating.

5.3 Bloch-sphere representation

To describe the effective two-level system, we use the language of the Bloch sphere (see [61]) and take the axis w as the “longitudinal” population inversion, and the axes u and v as the “transverse” coherences [62]. In a typical experiment, the atom-preparation


cycle concludes with the atoms initially all in the $|1\rangle$ state, along the w -axis (Fig. 5.2a). Applying a $\pi/2$ pulse starts the evolution of the wave by rotating all the spins to lie along the v -axis (Fig. 5.2b). When $\Delta(z)$ is set to have nonzero curvature, the spins along the long axis of the cloud begin to fan out in the u - v plane as the relative phase of the ($|1\rangle + |2\rangle$) coherent superposition develops a spatial dependence (Fig. 5.2c). Optionally at this point in the evolution, we can tune $\partial^2\Delta(z)/\partial z^2$ back to zero or leave it unchanged and watch the spins evolve in an inhomogeneous potential. In either case, the atoms' thermal motion now carries spin information back and forth along the length of the cloud, and spin-rotating collisions can tilt the spins up out of the u - v plane (Fig. 5.2d) [63]. Eventually, spin-currents develop that drive oscillations about an equilibrium spin distribution. Our driving potential, determined by the curvature of $\Delta(z)$, is even in z , and we predominantly drive the lowest order symmetric mode: a two-node standing wave.

5.4 Experimental method for studying spin waves

5.4.1 Extracting longitudinal phase

The longitudinal and transverse components of the Bloch vectors are measured as a function of time and space. The method for projecting the longitudinal component is similar to that described in [51] and is depicted graphically in Fig. 5.2. After an initial resonant $\pi/2$ pulse, the Bloch vectors are allowed to evolve for times up to $t = 800$ ms, and then the atoms are imaged, thereby projecting the Bloch vectors along the cloud onto the two spin states $|1\rangle$ and $|2\rangle$. The cloud is radially averaged and broken up into 25 equal bins along the axial direction for spatial resolution of the spin dynamics. The $|1\rangle$ and $|2\rangle$ populations are measured on separate experimental shots, and the longitudinal angle θ of the Bloch vector is extracted. ($\theta = \pi/2$ corresponds to equal populations, and $\theta = 0$ or π correspond to all the atoms in $|1\rangle$ or $|2\rangle$ respectively.) The evolution

a) Prepare ~ 600 nK normal cloud of atoms in state $|1\rangle$.



b) Apply a two-photon pulse to rotate the Bloch vectors into the u - v plane.



c) The v_{12} inhomogeneity causes the Bloch vectors to precess at different rates in the u - v plane.



d) Exchange collisions rotate the Bloch vectors out of the u - v plane.




Figure 5.2: Creation of spin waves (see text).

is traced out by repeating the experiment for many values of evolution time. Fig. 5.3a shows the time evolution of the longitudinal component of the spin across the cloud.

5.4.2 Extracting transverse phase

The transverse spin measurement consists of a two-pulse $\pi/2 - \pi/2$ Ramsey pulse sequence separated by an evolution time t [39]. For each time evolution time t and for each spatial bin along the trap axis, we measure the final density of the atoms in $|1\rangle$, resulting in a set of roughly sinusoidal (but with slightly time-varying phase and amplitude) Ramsey fringes for each location in space. The transverse phase and amplitude of the Bloch vector are extracted from the fringes at each time t for each radially-averaged section of the cloud by performing a sinusoidal fit over a small window




Figure 5.3: False color plots of the a) longitudinal and b) transverse components of the Bloch vectors across the cloud as a function of time for a density of $2 \times 10^{13} \text{ cm}^{-3}$ and a temperature of 800 nK (corresponding to a $380 \mu\text{m}$ full width, half maximum cloud size). The angles have been interpolated between bins and time steps of a) 10 ms and b) 3 ms.

centered about t in each spatial bin [64]. Fig. 5.3b shows typical phases extracted with this method. The spatial and temporal oscillations around the mean phase are the transverse projections of the spin waves.

The complete trajectories of the Bloch vector can be determined for each spatial bin by combining the transverse and longitudinal components. Normalizing the length of each vector to the length at time $t = 0$ removes any effects of decoherence (*e.g.* from inelastic loss and magnetic field fluctuations) and allows for a study of the equilibrium phases of the Bloch vectors as a function of the position across the cloud for various

$\Delta(z)$. Fig. 5.4 shows the trajectories of the reconstructed Bloch vectors in each spatial bin. Each bin has been offset by 1 rad in the transverse direction from the adjacent bins for clarity and to depict schematically the position across the cloud. The oscillations can be fit to extract spin wave frequency, amplitude, and damping rate.

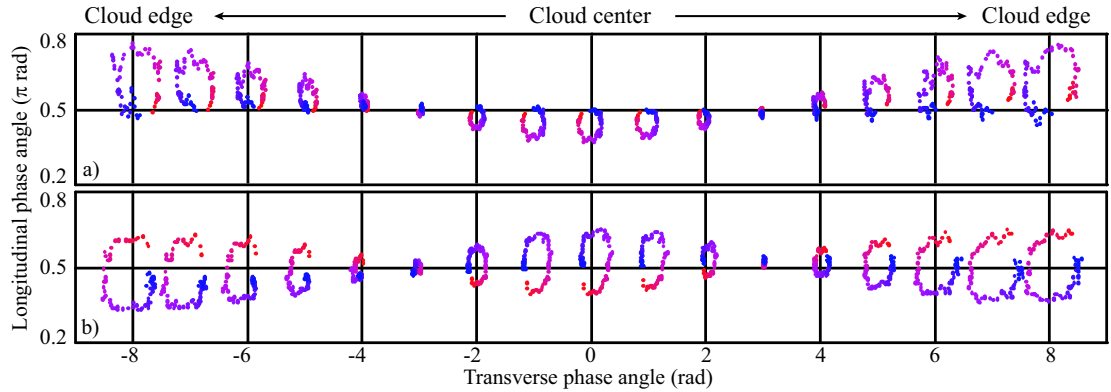


Figure 5.4: Bloch vector trajectories across the cloud for $T = 600$ nK and $n = 2 \times 10^{13}$ cm^{-3} . Each spatial bin (~ 40 $\mu\text{m}/\text{bin}$) is shifted by 1 rad from the adjacent bins (see text). In each bin, the path traced by the Bloch vectors is shown starting with the red points near time $t = 0$ and spiralling towards the blue points at $t = 0.3$ s as the Bloch vectors precess about the equilibrium position. a) Spin waves at $\partial^2\nu/\partial z^2 \approx 90$ Hz/mm^2 . The equilibrium position (center of precession) has a curvature to match the curved potential. b) Spin waves are excited for 70 ms at the same curvature, and then at $t = 0$ the bias magnetic field is ramped quickly so that $\partial^2\nu/\partial z^2 \approx 0$, giving an approximately flat equilibrium configuration near the v -axis for all spatial bins.

The steady-state Bloch vector arrangement is determined by the curvature of the potential. In order to equilibrate, the Bloch vectors across the cloud must arrange themselves so that the torque from the driving potential $\Delta(z)$ (which tends to spread the vectors in the u - v plane) is balanced by a torque from the local curvature in the spin field, which is spread out in the v - w plane. In the presence of a nonuniform potential, the equilibrium configuration of the Bloch vectors is curved longitudinally across the cloud to match the curvature of $\Delta(z)$, *i.e.* the Bloch vectors lie in the v - w plane but are not all collinear with the v -axis. However, if the bias magnetic field is chosen to cancel the density shift, then $\Delta(z)$ is predominantly flat across the cloud, and the equilibrium

corresponds to the Bloch vectors everywhere in the cloud collinear with the v -axis. These equilibrium configurations are illustrated in Fig. 5.4a and Fig. 5.4b, which show spin-wave evolution in curved and flat potentials respectively.

5.5 Frequency of oscillation of spin waves

Next, the density dependence of the spin wave frequency was studied. The spin wave frequency exhibits a strong density dependence, scaling roughly as $1/n$, shown in Fig. 5.5a. The linear regime for these excitations occurs only at extremely small amplitude, where spin waves are difficult to observe due to signal-to-noise limitations. In the nonlinear regime there is a dependence of frequency on amplitude. In order to remove this dependence, we excite spin waves over a range of amplitudes by controlling the curvature of the potential, $\partial^2\Delta(z)/\partial z^2$, and we extrapolate the frequency to the low amplitude limit for each density. The solid line is a numerical calculation predicting the frequency obtained from solving the one-dimensional Boltzmann transport equations [59]. The two dotted lines represent the weak and strong coupling limits, $\omega_{\text{exch}} \ll \omega_{\circ}$ and $\omega_{\text{exch}} \gg \omega_{\circ}$ respectively:

$$\omega_{\text{w}} = 2\omega_{\circ} - \frac{\omega_{\text{exch}}}{2} \quad (5.8)$$

$$\omega_{\text{s}} = \frac{k_{\text{eff}}^2(k_{\text{b}}T)}{m\omega_{\text{exch}}}. \quad (5.9)$$

where ω_{\circ} is the axial trap frequency and k_{eff} is an effective wave vector. The spatial wave vector $k = 2\pi/\lambda$ of the spin waves is studied by finding the distance $\lambda/2$ between the spatial nodes, *i.e.* the bins where the orientation of the Bloch vector remains constant. The wave vectors k and k_{eff} differ by a numerical factor of order unity due to the finite size of the sample. For all the work discussed here, the lowest order symmetric mode, $\eta = 2$, is excited by the symmetric inhomogeneous potential [65]. The wavelength of the $\eta = 2$ mode is observed to be directly proportional to the size of the trap. As predicted by Eqs. 5.8 and 5.9, there is no temperature dependence for the excitation

frequency, since we observe k to scale inversely as the size of the cloud, *i.e.* as $T^{-1/2}$. The numerical calculation agrees well with experiment (see Fig. 5.5a).




Figure 5.5: a) Frequency of the transverse spin wave component as a function of the peak density. Each frequency is an extrapolation to the limit of small amplitude spin waves to remove the dependence on driving amplitude. The solid line is the predicted dispersion relation, and the dotted lines show the weak and strong coupling limits. All data is for $T = 600$ nK except for the three tightly clustered points at $n = 2 \times 10^{19} \text{ m}^{-3}$ which correspond to $T = 350, 600,$ and 1000 nK. b) Damping rate as a function of density, including a numerical calculation (solid line) which includes elastic collisions and Landau damping.

The Bloch vectors initially precess about the equilibrium value with some initial amplitude. As the Bloch vectors precess, their trajectories spiral towards the equilibrium configuration due to damping of the spin wave excitation. The damping is predicted to be a combination of two effects: spin diffusion and Landau damping [66]. Spin diffusion dominates in the high density, or hydrodynamic, regime. For intermediate densities,

there is a maximum in the damping rate due to the combination of the two effects. It near this intermediate regime that these experiments are performed. We observe a relationship between spin wave lifetime and density which is good agreement with theory (see Fig. 5.5b).

We have observed collective spin excitations in a nondegenerate ultra-cold gas. Particle indistinguishability and the need to symmetrize collisions between like particles in coherent superpositions give rise to spatial and temporal oscillations of the Bloch vectors across a cloud of trapped atoms. We have separately imaged both transverse and longitudinal components of these spin waves and studied the effects of temperature and density on the spin wave frequency, paving the way for studying spin dynamics in a finite-temperature condensate system.

Our original motivation to explore spin dynamics in uncondensed clouds arose from our plan to extend to finite temperatures our earlier studies of spin coherence in a condensate as described in Chapter 4 and in particular to understand the extent to which an unpolarized normal cloud acts as a decohering thermal reservoir. The existence of spin waves complicates these studies. A pure condensate cannot support spin waves, but it will trivially separate due to small differences in the scattering lengths [48]. Spin oscillations in the thermal cloud can have a considerable effect on the spin dynamics of a finite temperature condensate, and conversely, having a small stationary cluster of condensed atoms will significantly perturb spin waves in a thermal cloud. Such a two component system is a rich system for studying spin dynamics in ultra-cold gases [66]. The next chapter will present our efforts to study a finite temperature condensate system.

Chapter 6

Coherence effects in a finite temperature condensate

Multiple spin-component Bose-Einstein condensates have permitted many studies of global symmetries and quantum statistical correlation effects in macroscopic quantum systems [67, 68, 25]. Interparticle correlations, or coherences, are fundamental properties in ultra-cold atomic gases as well as in the field of quantum optics. First-order coherences, referring to the deterministic nature of the phase of a system, can be seen in both fields. First-order coherence has been seen by overlapping two Bose-Einstein condensates and observing spatial density fringes [69], which is a counterpart to interference between two laser beams. We investigate further first-order correlations by watching finite-temperature condensate melting due to decoherence of the normal cloud. Using the Bose-Einstein phase transition as calibration, we examine the thermodynamics of the system to gain understanding of the correlations between atoms.

In this chapter we study how correlations affect the condensing and melting process in a condensate. We start by discussing the properties of the states in which we will study loss of coherence. We use two dressed states in ^{87}Rb , because these states enhance the natural symmetries of the system and allow us to more easily measure the correlations. Studies of the dressed states include measurements of stability of the states against transitions and experiments to determine the equilibrium spatial density distribution. Next we discuss the situation where we cool through the transition temperature with a partially coherent cloud to examine the symmetry of the condensed state, and

finally how loss of coherence leads to melting of the condensate.

6.1 Dressed states

The dressed-state formalism is useful to describe a two-state system coupled by a radiation field [26]. The energies of the bare states

$$E_{e/g} = \pm \frac{1}{2} \hbar \omega_0 \quad (6.1)$$

are shifted by the presence of the coupling field to

$$E_g = -\frac{1}{2} \hbar \omega_0 + N \hbar \omega \quad (6.2)$$

$$E_e = +\frac{1}{2} \hbar \omega_0 + (N - 1) \hbar \omega, \quad (6.3)$$

where ω_0 is the frequency difference between the two bare states, ω is the frequency of the coupling field, and N is the number of photons in the field. Solving for the eigenstates of the new Hamiltonian including the coupling field produces two eigenstates that are linear combinations of the bare states. These dressed states

$$|+\rangle = e^{-i(\Omega'_R - \delta - \omega_0)\frac{t}{2}} \left[\sqrt{\frac{1}{2}\left(1 + \frac{\delta}{\Omega'_R}\right)} |1\rangle + e^{-i\omega t} \sqrt{\frac{1}{2}\left(1 - \frac{\delta}{\Omega'_R}\right)} |2\rangle \right] \quad (6.4)$$

$$|-\rangle = e^{i(\Omega'_R - \delta - \omega_0)\frac{t}{2}} \left[-e^{-i\omega t} \sqrt{\frac{1}{2}\left(1 - \frac{\delta}{\Omega'_R}\right)} |1\rangle + \sqrt{\frac{1}{2}\left(1 + \frac{\delta}{\Omega'_R}\right)} |2\rangle \right] \quad (6.5)$$

are separated in energy by the effective Rabi frequency

$$\Omega'_R = \sqrt{(\omega_0 - \omega)^2 + \Omega_R^2}, \quad (6.6)$$

where Ω_R is the resonant Rabi frequency. A diagram of the dressed states as a function of detuning from resonance is shown in Fig. 6.1. When the coupling field is detuned far red of resonance the ground state is essentially the $|-\rangle$ state, and the excited state is essentially the $|+\rangle$ state. As the detuning approaches zero the dressed states become

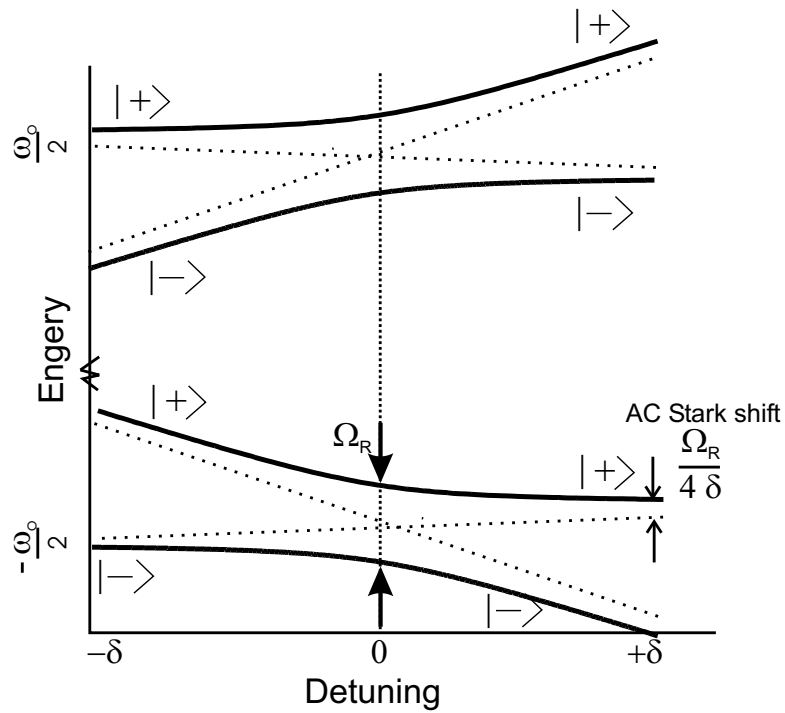


Figure 6.1: Energy level diagram of the dressed states showing the avoided crossing.

equal superpositions of the bare states. Finally, if the detuning increases in the blue direction the $|-\rangle$ state is now the excited state.

An atom initially in one dressed state will remain in that dressed state if the detuning is changed adiabatically. Adiabaticity is maintained if Υ is much greater than one, where Υ , the adiabaticity parameter, is

$$\Upsilon = \frac{\Omega_R^2}{\left| \frac{d\delta}{dt} \right|}, \quad (6.7)$$

where $\frac{d\delta}{dt}$ is the time rate of change of the detuning from resonance. The probability that an atom will make a transition to another dressed state is [70]

$$P = e^{-2\pi\Upsilon}. \quad (6.8)$$

6.2 Symmetries of the dressed-state system

Before looking at correlation effects, we must first understand the properties and symmetries of the dressed state system. The dressed states on resonance, being equal superpositions of the bare states, have convenient properties that are not present in the bare state system. These properties include identical collisional and loss rates as well as nearly identical trapping potentials.

A study of the symmetries includes the stability of atoms in a single dressed state. A collection of ultra-cold atoms in a binary mixture of two spin states has two global U(1) symmetries, which correspond to conservation of particle number in each species. However, when a coupling radiation field is tuned to the transition frequency between the spin states, one of the U(1) symmetries is broken. The symmetry is broken by the constant phase relation between the two species imposed by the field. The symmetry that remains corresponds to a conservation of total atom number. If the energy between the two states could be reduced to zero the system would be fully SU(2) symmetric. Such a system would be ideal to produce proposed Schrödinger cat states [71], but is excessively difficult with our system.

The experimental system is similar to the one described in Chapter 2. The system consists of a sample of atoms in the $|1\rangle$ hyperfine ground state of ^{87}Rb cooled to below 500 nK in the HIP trap. The dressing field consists of a microwave photon (~ 6.8 GHz) and a rf photon (3 MHz) that couples atoms to the $|2\rangle$ state through the intermediate state $|2,0\rangle$ [72]. Typical resonant Rabi frequencies are between 0.1-8 kHz and can be varied dynamically during an experimental run. The atoms can be prepared in a single dressed state by adiabatically ramping the dressing-field detuning from a large positive or negative value to zero. The atoms can also be prepared in an equal coherent superposition of the two dressed states by diabatically turning on a resonant dressing field.

We can selectively image atoms in either the dressed- or bare-state basis. To image atoms in the dressed basis we adiabatically ramp the dressing-field frequency far from resonance in either direction before turning off the field. For example ramping the frequency to a larger frequency transforms the $|+\rangle$ state into the $|1\rangle$ state and the $|-\rangle$ state into the $|2\rangle$ state. We can then use absorption imaging to probe the atoms only in the $|1\rangle$ state. We can also quickly turn off the dressing field while still on resonance to project the atoms into the bare-state basis, and therefore probe the coherence of the dressed states, which manifests itself as Rabi flopping in the bare states.

6.2.1 Stability of dressed states

We are unaware of any symmetry that implies conservation of atom number separately in each dressed state, and thus transitions between dressed states are not explicitly forbidden. We set out to observe transitions between dressed states because we thought transitions might be possible when the energy splitting of the dressed states is less than the thermal energy $k_B T$.

There are several mechanisms that can transfer population between dressed states. The most interesting type of transfer arises from collisions. During a collision the relative phase between the $|1\rangle$ and $|2\rangle$ bare parts of the wavefunction may acquire a phase shift, which results in population transfer in the dressed basis. Another mechanism of transfer is driven resonantly by the magnetic trapping fields. As atoms oscillate in the harmonic trap there appears, in the atoms' frame, to be an oscillating magnetic field. If the energy splitting between the dressed states, *i.e.* Rabi frequency, is tuned such that it is twice the trap oscillation frequency, the atoms have a probability of making a transition between dressed states. The radial frequency is large compared to the collision rate which allows an atom to oscillate freely in the trap for several cycles before a collision takes place and changes the effective oscillating magnetic field seen by the atom. This resonant transfer is an incoherent process, because the phase of the oscillating field for

each atoms depends on its position and velocity, which are distributed thermally. The populations in each state should therefore not oscillate coherently but relax instead to an equal mixture.

Besides collisionally and resonantly induced transfer, atoms can make a transition between dressed state via the Landau-Zener process [73, 74]. If an atom cannot adiabatically follow the torque applied by the dressing field it may transfer between dressed states. If $d\delta/dt > \Omega_R^2$, where δ is the detuning of the dressing field, there is a large probability an atom will undergo a Landau-Zener transition, as described in section 6.1. The detuning from the intermediate state varies spatially across the cloud due to the non-uniform density profile, leading to a spatially dependent mean-field energy shift, and also due to the Zeeman energy shift from the magnetic trapping fields. As the atoms oscillate in the trap they experience a time rate of change of detuning as they sample different densities and magnetic fields, which induces transfer.

We measure the rate of transfer between dressed states by adiabatically dressing the atoms in purely one dressed state on resonance and detecting atoms appearing in the other state. We start with the atoms in the $|-\rangle$ state and wait up to two seconds. We then ramp the drive detuning $10 \times \Omega_R$ in either direction to adiabatically convert one of the dressed states into the $|1\rangle$ state, which we image using absorption. We determine the transfer rate at different Rabi frequencies by extracting a linear slope of the fraction transferred as a function of time (Fig.6.2).

There are two features to note with the data in Fig. 6.2. First, there is a resonant increase in transfer where the Rabi frequency is twice the radial frequency for the reason described above. Second there is a background rate, which decreases as Rabi frequency increases, which is most likely due to Landau-Zener transitions. We can estimate the transfer probability per collision far from the resonance, above 3 kHz, to be no more than $4(2) \times 10^{-5}$. Although collisional transfer from the $|-\rangle$ to the $|+\rangle$ state is endothermic, it is energetically allowed because the atoms' thermal energy is larger than the splitting




Figure 6.2: Transfer rate is shown as a function of Rabi frequency. The inset graph highlights the resonant transfer at twice the radial trap frequency ($\nu_{radial} = 230.5\text{Hz}$); the solid line is a fit to a plausible functional form of a decaying exponential plus a Lorentzian. The error bars reflect the statistical uncertainty in the fitting. Initial cloud parameters are $T = 300\text{ nK}$ and density $= 2 \times 10^{13}\text{cm}^{-3}$.

between the dressed states. Collisional transitions maybe suppressed in ^{87}Rb because the triplet and singlet scattering lengths are nearly identical, which reduces the rate for atoms to spin-flip between hyperfine ground states during collisions. The small collisionally induced transfer rate implies an approximate symmetry. The system is not free to explore fully the entire $\text{SU}(2)$ space, which would be useful for verifying several theoretical predictions such as Mermin-Ho vortices [75]. For our purposes, the stability of atoms in the dressed states is fortunate. We want to compare the thermodynamics of atoms in an equal superposition of states to that of atoms in one dressed state. This would not be possible if the atoms, initially prepared in one state, would over short

times equilibrate to an equal populations in both states.

6.2.2 Separation of dressed states

Working with dressed states instead of bare states can also increase the spatial symmetry of a two-component condensate. A condensate in a superposition of two bare states will spatially separate in the trap to reduce the total energy of the system [48, 76]. The two spin states we work with in ^{87}Rb have slightly different mean field interactions, which drives the separation. In other work it has been shown that the $|1\rangle$ state, the state with a larger s-wave scattering length, will move to the lower density region of the condensate, and the $|2\rangle$ state will concentrate in the center of the cloud where the density is the highest [48].

The dressed states on resonance, on the other hand, are composed of equal superpositions of bare states. The mean-field interaction energy between the dressed states is equal and therefore does not drive spatial separation. However any significant spatial gradient in the differential potential between the dressed states, such as produced by a spatially inhomogeneous Rabi frequency, will cause separation. Atoms at different positions in the harmonic trap are detuned differently from the intermediate state, $|2, 0\rangle$, due to the Zeeman effect. The spatial variation of the intermediate state detuning results in a variation of the Rabi frequency across the cloud. Typically the Rabi frequency varies by at most a few percent across the entire condensate. However at large Ω_R , a few percent can mean 100 Hz difference, more than enough to drive spatial separation. The Rabi frequency inhomogeneity has the functional form of the magnetic trap: parabolic. When we operate with large Rabi frequencies ($>1\text{kHz}$) and negative intermediate state detuning, the dressed states separate with the $|-\rangle$ occupying the outer region of the condensate and the $|+\rangle$ in the center (Fig. 6.3). Changing the sign of the detuning causes the states to reverse positions in the trap. The goal for spectroscopy is to have the two states overlapped as much as possible in the trap. The dressed states in the

condensate remain completely overlapped, within experimental resolution, when the intermediate state detuning is increased to 1.4 MHz and the Rabi frequency is decreased 300 Hz as shown in Fig. 6.3 .

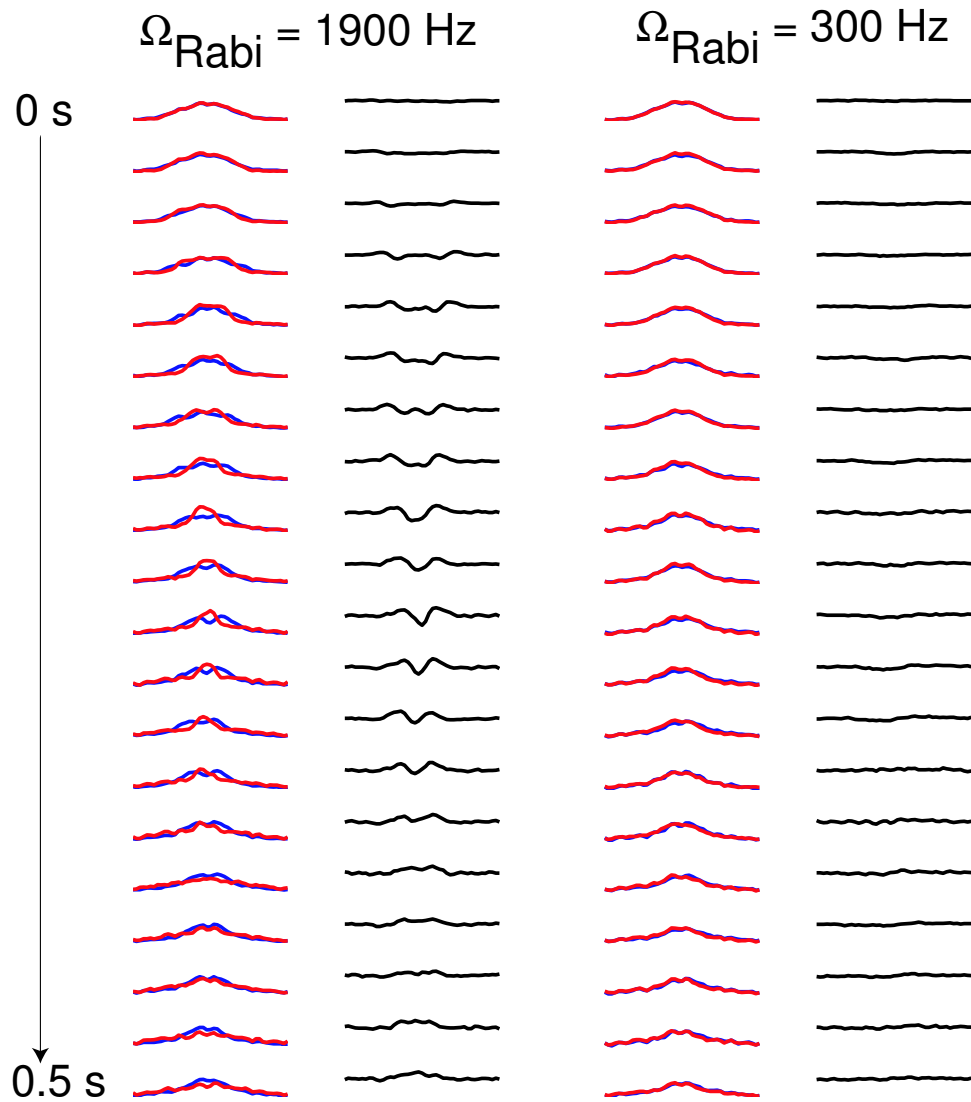


Figure 6.3: Normalized cross sections through the clouds comprised of $|+\rangle$ atoms in red and $|-\rangle$ in blue as a function of time; the difference between the cross sections is in black (second and fourth columns). There is significant separation of the two dressed states when the Rabi frequency is 1.9 kHz and $\Delta_{\text{IM}} = 0.7$ kHz. Decreasing the Rabi frequency to 300 Hz and increasing $\Delta_{\text{IM}} = 1.4$ MHz results in the clouds remaining overlapped for at least 0.5 seconds.

6.2.3 Evaporation of dressed states

We investigate the symmetries of the dress-state system also by condensing with the dressing field present. An interesting question to ask is how the condensation process is affected by the dressing field. Does the condensate form in the lowest energy state in the system or in an excited state such as a coherent superposition of dressed states? We examine these questions by starting with a cloud of atoms above the transition temperature, applying the resonant dressing field, and evaporatively cooling the cloud until a condensate forms.

We begin the experiment with a 50/50 fully-coherent superposition of $|+\rangle$ and $|-\rangle$ states, above the condensation temperature. During the evaporation the atoms dephase due to the spatial inhomogeneity of the Rabi frequency. The dephasing time is proportional to the magnitude of the Rabi frequency and is on the order of 100 ms for a Rabi frequency of 370 Hz. We cool the cloud in 120 ms using simultaneous microwave evaporation on both the $|1\rangle$ and $|2\rangle$ states. By the time the condensate forms the normal cloud has dephased significantly but not entirely. In this manner we can create a cloud with a condensate fraction up to 10%. After the condensate forms we investigate the state of the condensate by waiting a variable time before diabatically turning off the dressing field and imaging the atoms in the $|1\rangle$ state. We integrate the number through a column in the center of the cloud which includes most of the condensate and some of the normal cloud. The data in Fig. 6.4 show the number in the center of the cloud, which is 52(4)% condensate, Rabi flopping with 86(4)% contrast.

The high contrast fringes demonstrate that the condensate, which is condensed out of a normal cloud with residual coherence, forms in a superposition of dressed states. One might expect this result because the normal cloud retains some local coherence when the condensate forms. However, it is still surprising that the condensate is created in a state that cycles between the bare states, despite the powerful statistical mechanical

preference for groundstate condensation.

We are unable to fully answer the question of how correlations are lost, or what state a condensate is created in when the normal cloud has lost all correlations, because of the limited parameter space we can explore. By the time the condensate forms, the normal cloud has significantly decohered because of dephasing during the relatively long evaporation time. This limits the measurement to samples with already small correlations. On the other side, we can not form a condensate where the correlations have essentially disappeared. This inability to condense suggests that the correlations themselves could promote condensation. We can barely create a condensate when correlations are present because we lose a significant number of atoms due to dipolar relaxation in the $|2\rangle$ state. Therefore it is not surprising that we can not produce a condensate when the correlations are lost, because we would need a factor of 2 more atoms. Because of these limitations we also looked at the loss of correlations from the perspective of melting instead of condensing.




Figure 6.4: The density of atoms in the center of the cloud in the $|1\rangle$ state Rabi flopping as a function of hold time after evaporation. 52% of the central density consists of condensate atoms.

6.3 Decoherence induced melting of a finite temperature condensate

We use the statistics of the Bose phase transition to look at the correlations. A partially condensed system made up of fully coherent atoms in an equal superposition of two spin states is, in effect, composed of a single species of atom. If the noncondensed atoms dephase, due to an inhomogeneous potential, they become distinguishable particles and must be thought of atoms in two distinct populations. A condensate at finite temperature initially in equilibrium with a fully coherent normal cloud will presumably melt as the coherence is lost in the normal cloud. This can be understood by realizing the loss of coherence essentially reduces the effective number in the state by two ($N_{eff} = N/2$). The condensate must then melt to restore thermal equilibrium.

We observe condensate melting induced by decoherence in the normal cloud by initially creating a cloud in the $|1, -1\rangle$ state with about 45% condensate fraction. Next we dress the cloud in an equal coherent superposition of dressed states. We measure the temperature and density of the cloud over a time of 200 ms and calculate the apparent phase space density [77] of the normal cloud as the condensate melts. For a system in a single dressed state or a fully coherent superposition of states, the phase space density for the normal cloud should be equal to one, normalized by 2.61, as long as a condensate is present. If the system is an uncorrelated binary mixture of two states, the *total* normal cloud will appear to be at twice the phase space density. We also performed the same experiment with the sample in only one dressed state.

The one state data allows us to remove imaging and fitting systematics. There is a systematic error in calculating phase space density of the normal cloud when a condensate is present due to our fitting routine. We find the calculated phase space density increases linearly with increasing condensate fraction for the one component cloud, which should remain at a constant 2.61 when a condensate is present. We remove

this systematic by fitting a line to the phase space density vs. condensate fraction data, and then we remove this slope from the both the single and superposition state data. Also, when we took these data we did not have enough microwave power to transfer 100% of the atoms via ARP for imaging. Therefore, we could not detect all of the atoms, which led to an underestimate of the phase space density. We eliminated this problem by normalizing each phase space density data point by the average of the single state phase space density.

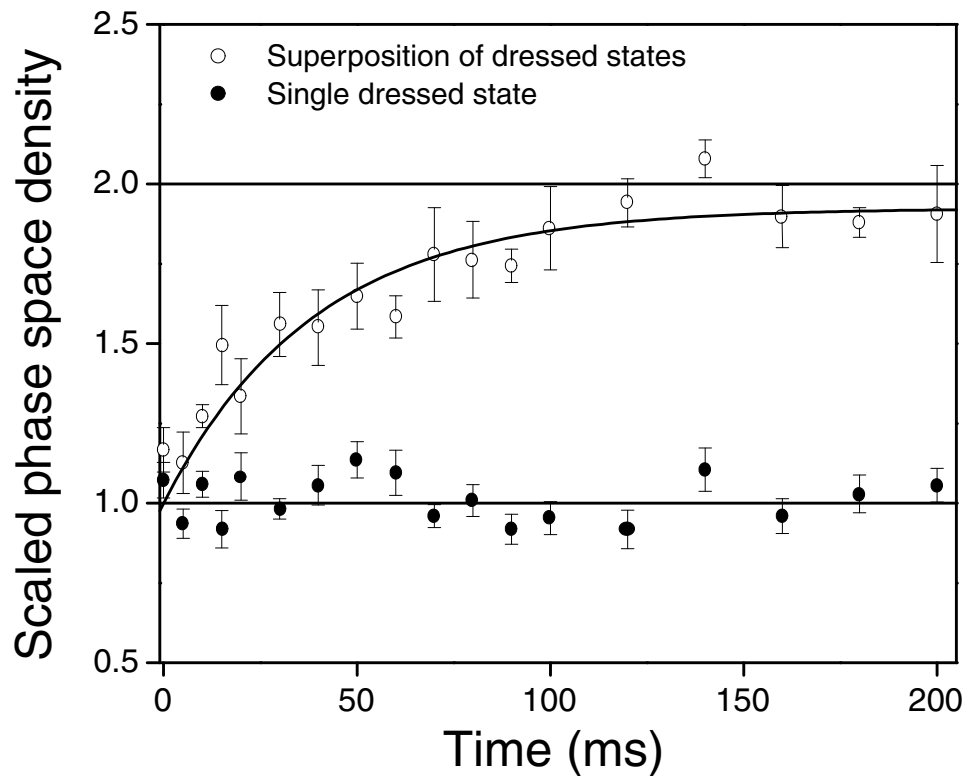


Figure 6.5: The measured phase space density (PSD), normalized by 2.61, as a function of decohering time. The solid circles are the PSD of the normal cloud in a single state, whereas the open circles represent the PSD of the normal cloud in an initially fully coherent superposition of dressed states. The superposition state PSD approaches two in a time scale of 40 ms, which is consistent with the normal cloud coherence time of 30(5) ms. Each point is a weighted average of nine independent measurements.

We are able to observe the apparent phase space density of the superposition case go from $1 \rightarrow \sim 2$ [Fig. 6.5]. The time scale for the correlations to go away is on the order

of 40 ms, as extracted from an exponential fit to the data in Fig. 6.5. One would expect the atoms to become uncorrelated on the same time scale as the decoherence, ignoring the thermal equilibration time of the system. We measure the coherence time of the thermal atoms at the edge of the cloud to give an estimate of the coherence time at the center of the cloud. We find the coherence time of the thermal atoms to be 30(5) ms which is consistent with the time for the correlations to disappear.

In the special case where the normal cloud has lost all correlations and can be thought of as a binary mixture of two states, the question arises what state describes the condensate? One possibility is that collisions into and out of the condensate from the normal cloud could create two condensates in a binary mixture. The two condensates would still have spatial coherence but would have lost internal coherence. There would be a distinct relative phase between the condensates in each dressed state but would vary from experimental shot-to-shot. In the end, any product of the decohering process would be a novel quantum state. Unfortunately we are currently unable to determine the state of the condensate after the decohering process, because by the time all of the correlations are lost there is very little to no condensate present due to the large loss rate in the $|2\rangle$ state.

In conclusion we have studied the properties of a two-state dressed system. Atoms in a dressed state show remarkable stability against relaxing to a different dressed state. This stability allows us to use these states to examine the effects of correlations on condensing and melting Bose-Einstein condensates. We were able to watch the correlations disappear and observe decoherence-induced melting in a finite-temperature condensate. The ultimate state of the condensate after all normal cloud correlations are gone would be interesting to study in the future.

Bibliography

- [1] S. N. Bose, *Z. Phys.* **26**, 178 (1924).
- [2] A. Einstein, *Sitzungsber Kgl. Preuss. Akad. Wiss.* **1924**, 261 (1924); *ibid.* **1925**, 3 (1925).
- [3] 2001 Nobel prize in physics citation.
- [4] M. Anderson, J. Ensher, M. Matthews, C. Wieman, and E. Cornell, *Science* **269**, 198 (1995).
- [5] K. Davis, M.-O. Mewes, M. Andrews, N. van Druten, D. Durfee, D. Kurn, and W. Ketterle, *Phys. Rev. Lett.* **75**, 3969 (1995).
- [6] C. Bradley, C. Sackett, J. Tollett, and R. Hulet, *Phys. Rev. Lett.* **75**, 1687 (1995), *ibid.* **79**, 1170 (1997).
- [7] E. Raab, M. Prentiss, A. Cable, S. Chu, and D. Pritchard, *Phys. Rev. Lett.* **59**, 2631 (1987).
- [8] A. L. Migdall, J. V. Prodan, W. D. Phillips, T. H. Bergeman, and H. J. Metcalf, *Phys. Rev. Lett.* **54**, 2596 (1985).
- [9] N. Masuhara, J. M. Doyle, J. C. Sandberg, D. Kleppner, T. J. Greytak, H. Hess, and G. P. Kochanski, *Phys. Rev. Lett.* **61**, 935 (1988).
- [10] H. J. Metcalf and P. van der Straten, Laser Cooling and Trapping (Springer-Verlag, 1999).
- [11] C. Wieman and L. Hollberg, *Rev. Sci. Instrum.* **62**, 1 (1991). **113**, 151 (1998).
- [12] K. MacAdam, A. Steinbach, and C. Wieman, *Am. J. Phys.* **60**, 1098 (1992).
- [13] C. Wieman, G. Flowers, and S. Gilbert, *Am. J. Phys.* **63**, 317 (1995).
- [14] W. Ketterle and N. van Druten, *Adv. At. Mol. Opt. Phys.* **37**, 181 (1996).
- [15] W. Ketterle, D. Durfee, and D. Stamper-Kurn, in Proceedings of the International School of Physics - Enrico Fermi, (IOS Press, 1999), p. 67.
- [16] C. Myatt, N. Newbury, R. Ghrist, S. Loutzenhiser, and C. Wieman, *Opt. Lett.* **21**, 290 (1996).

- [17] W. Phillips and H. Metcalf, *Phys. Rev. Lett.* **48**, 596 (1982).
- [18] M. Greiner, I. Bloch and T. Esslinger, *Phys. Rev. A* **63**, 031401(R) (2001).
- [19] J. H. Moore and M. A. Coplan, Building Scientific Apparatus, Second Edition (Perseus Books, 1991).
- [20] S. K. Kurt E. Gibble and S. Chu, *Opt. Lett.* **17**, 526 (1992).
- [21] P. D. Lett, R. N. Watts, C. E. Tanner, S. L. Rolston, W. D. Phillips, and C. I. Westbrook, *J. Opt. Soc. Am. B* **6**, 2084 (1989).
- [22] W. Petrich, M. H. Anderson, J. R. Ensher, and E. A. Cornell, *J. Opt. Soc. Am. B* **11**, 1332 (1994).
- [23] E. Hecht, Optics, Second Edition (Addison-Wesley, 1990).
- [24] W. Petrich, M. H. Anderson, J. R. Ensher, and E. A. Cornell, *Phys. Rev. Lett.* **74**, 3352 (1995).
- [25] E. Burt, R. Ghrist, C. Myatt, M. Holland, E. Cornell, and C. Wieman, *Phys. Rev. Lett.* **79**, 337 (1997).
- [26] C. Cohen-Tannoudju, J. Dupont-Roc, and G. Grynberg, Atom-Photon Interactions (Wiley-Interscience, 1992).
- [27] Viewpoint USA has recently begun to produce a similar board (DIO-64) with 64 outputs and an improved resolution of 50 ns. This new board is PCI rather than ISA and should be compatible with the existing software drivers written for the DIO-128.
- [28] E. Arimondo, M. Inguscio, and P. Violino, *Rev. Mod. Phys.* **49**, 31 (1977).
- [29] C. Salomon, Y. Sortais, S. Bize, M. Abgrall, S. Zhang, C. Nicolas, C. Mandache, P. Lemonde, P. Laurent, G. Santarelli, A. Clairon, N. Dimarcq, P. Petit, A. Mann, A. Luiten, and S. Chang, in Proceedings of the 17th International Conference on Atomic Physics (ICAP 2000), (American Institute of Physics, 2000), p. 23.
- [30] D. M. Harber, H. J. Lewandowski, J. M. McGuirk, and E. A. Cornell, *Phys. Rev. A* **66**, 053616 (2002).
- [31] M. J. Holland, D. S. Jin, M. L. Chiofalo, and J. Cooper, *Phys. Rev. Lett.* **78**, 3801 (1997).
- [32] W. Ketterle and H.-J. Miesner, *Phys. Rev. A* **56**, 3291 (1997).
- [33] E. A. Burt, R. W. Ghrist, C. J. Myatt, M. J. Holland, E. A. Cornell, and C. E. Wieman, *Phys. Rev. Lett.* **79**, 337 (1997).
- [34] K. Gibble and S. Chu, **70**, 1771 (1993).
- [35] C. Fertig and K. Gibble, *Phys. Rev. Lett.* **85**, 1622 (2000).

- [36] Y. Sortais, S. Bize, C. Nicolas, A. Clairon, C. Salomon, and C. Williams, Phys. Rev. Lett. **85**, 3117 (2000).
- [37] T. C. Killian, D. G. Fried, L. Willmann, D. Landhuis, S. C. Moss, T. J. Greytak, and D. Kleppner, Phys. Rev. Lett. **81**, 3807 (1998); D. G. Fried, T. C. Killian, L. Willmann, D. Landhuis, S. C. Moss, D. Kleppner, and T. J. Greytak, Phys. Rev. Lett. **81**, 3811 (1998).
- [38] N. Davidson, H. J. Lee, C. S. Adams, M. Kasevich, and S. Chu, Phys. Rev. Lett. **74**, 1311 (1995).
- [39] N. F. Ramsey, Molecular Beams (Clarendon Press, Oxford, 1956).
- [40] J. M. McGuirk, H. J. Lewandowski, D. M. Harber, T. Nikuni, J. E. Williams, and E. A. Cornell, Phys. Rev. Lett. **89**, 090402 (2002);
- [41] At $n = 10^{13} \text{cm}^{-3}$ $\nu_{\text{spin wave}} = 5$ Hz, so interrogation times of 20 ms are acceptable.
- [42] Where $\beta = 431.35957(9)$ Hz/G² and $\nu_{\text{min}} = 6834678113.59(2)$ Hz [51].
- [43] M. Ö. Oktel, T. C. Killian, D. Kleppner and L. S. Levitov, Phys. Rev. A **65**, 033617 (2002).
- [44] M. Ö. Oktel and L. S. Levitov, Phys. Rev. A **65**, 063604 (2002).
- [45] J. E. Williams, T. Nikuni, and C. W. Clark, Phys. Rev. Lett. **88**, 230405 (2002).
- [46] E. G. M. van Kempen, S. J. J. M. F. Kokkelmans, D. J. Heinzen, and B. J. Verhaar, Phys. Rev. Lett. **88**, 093201 (2002).
- [47] The Thomas-Fermi radius (and to a lesser extent the critical temperature) are of course interaction dependent, and therefore using measurements of these quantities to calibrate densities in order to characterize interaction strengths is somewhat circular. Our directly measured value of $\alpha_{ii}^{nc}/\alpha_{ii}^c = 2.1(2)$, on the other hand, is based predominantly on measured spatial extent of the cloud and on an absolute determination of the absorption depth, and therefore it is only weakly model dependent. We quote the “worst-case corrected value” of $\alpha_{ii}^{nc}/\alpha_{ii}^c = 1.7(2)$ in order to put a reasonable limit on the size of possible systematic errors, and not because we feel it is the preferred value.
- [48] D. S. Hall, M. R. Matthews, J. R. Ensher, C. E. Wieman, and E. A. Cornell, Phys. Rev. Lett. **81**, 1539 (1998).
- [49] When we changed the length of the first Ramsey pulse, the second pulse length was also changed such that the two pulses combined to form a π pulse, which gives us the maximum signal-to-noise.
- [50] The quoted chirp limit of -0.2(3) Hz includes a correction to account for a small frequency change associated with decay of total atom number during the measurement time. For each data set the number decay was found by allowing an exponential decay in the fit of the Ramsey fringes (not to be confused with the decay in Ramsey

fringe contrast associated with decoherence), resulting in a correction amounting to -0.74 Hz on average.

- [51] H. J. Lewandowski, D. M. Harber, D. L. Whitaker, and E. A. Cornell, Phys. Rev. Lett. **88**, 070403-1 (2002).
- [52] Fitting the decay of the integrated transverse magnetization to $e^{-(t/\tau)^2}$, the same functional form as we fit to the actual experimental data, gives values typically within 5%, and at worst 12%, of the $1/e$ time.
- [53] W. M. Itano, J. C. Bergquist, J. J. Bollinger, J. M. Gilligan, D. J. Heinzen, F. L. Moore, M. G. Raizen, and D. J. Wineland, Phys. Rev. A **47**, 3554 (1993).
- [54] See for instance Spin polarized quantum systems, Proc. of the Aussois Conference, J. Physique Colloq. **41**, C-7 (1980).
- [55] B. DeMarco, J. L. Bohn, J. P. Burke, Jr., M. Holland, and D. S. Jin, Phys. Rev. Lett. **82**, 4208 (1999).
- [56] F. Bloch, Z. Physik **61**, 201 (1930).
- [57] E. P. Bashkin, JETP Lett. **33**, 8 (1981). C. Lhuillier and F. Laloë, J. Physique **43**, 225 (1982); C. Lhuillier and F. Laloë, J. Physique **43**, 197 (1982); A. Leggett and M. Rice, Phys. Rev. Lett. **20**, 586 (1968), errata Phys. Rev. Lett. **21**, 506 (1968); K. Miyake, W. Mullin, and P. C. E. Stamp, J. Physique **46**, 663 (1985); L. P. Lévy and A. E. Ruckenstein, Phys. Rev. Lett. **52**, 1512 (1984), G. Tastevin, P. J. Nacher, M. Luduc, F. Laloë, J. Phys. Lett. **46**, 249 (1985).
- [58] B. R. Johnson *et al.*, Phys. Rev. Lett. **52**, 1508 (1984); P. J. Nacher, *et al.*, J. Phys. (Paris) Lett. **9**, L441 (1984); W. J. Gully and W. J. Mullin, Phys. Rev. Lett. **52**, 1810 (1984).
- [59] T. Nikuni, J. E. Williams, C. W. Clark, Phys. Rev. A **66**, 043411 (2002).
- [60] The two states in question are $|1\rangle \equiv |F = 1, m_F = -1\rangle$ and $|2\rangle \equiv |F = 2, m_F = 1\rangle$, which are pseudo-spin states since the system is not a true spin-1/2 system.
- [61] See for instance L. C. Allen and J. H. Eberly, Optical Resonance and Two Level Atoms (Dover, New York, 1956).
- [62] Alternately to the Bloch sphere picture, one can think of the \underline{u} , \underline{v} , and \underline{w} components as the expectation values of the pseudo-spin state \vec{S} : $\langle S_x \rangle$, $\langle S_y \rangle$, and $\langle S_z \rangle$ respectively.
- [63] It is important to realize that the difference in mean-field interactions between the two states, $(a_{22} - a_{11}) \approx 2.6 \text{ \AA}$, is extremely small in ^{87}Rb and is not the mechanism for spin wave propagation. For the purposes of these experiments $(a_{22} - a_{11})$ provides merely a contribution to the inhomogeneity in $\Delta(z)$. The total mean-field effect ω_{exch} is proportional to $a = (a_{11} + a_{22})/2 \approx 50 \text{ \AA}$.

- [64] The mean phase of the Bloch vector evolves at the average hyperfine splitting frequency of $\bar{\Delta} = 6.8$ GHz; the Bloch vector only moves slowly in the frame rotating at $\bar{\Delta}$. It is in this reference frame the effects of spin waves are considered. The local oscillator is set so that the Ramsey fringe frequency is much larger than the spin wave frequency, simplifying analysis of the “instantaneous” phase and amplitude of the Bloch vectors.
- [65] The quadrupole mode is excited with a frequency splitting $\Delta(z)$ that varies quadratically with distance from the trap center. The dipole mode could be driven by applying a linear gradient to the trapped atoms.
- [66] M. Ö. Oktel and L. S. Levitov, *Phys. Rev. Lett.* **83**, 6 (1999).
- [67] M. Naraschewski and R. J. Glauber, *Phys. Rev. A*, **59**, 4595 (1999).
- [68] H. J. Miesner *et al.*, *Phys. Rev. Lett.* **82**, 2228 (1999).
- [69] M. R. Andrews *et al.*, *Science* **275**, 627 (1997).
- [70] J. R. Rubbmark, M. M. Kash, M. G. Littman, and D. Kleppner, *Phys. Rev. A* **23**, 3108 (1981).
- [71] A. B. Kuklov and B. V. Svistunov, cond-mat/0202538.
- [72] M. R. Matthews *et al.*, *Phys. Rev. Lett.* **81**, 243 (1998).
- [73] L. D. Landau, *Z. Sowjetunion* **2,46** (1932).
- [74] G. Zener, *Proc. R. Soc. London A*, **137**, 696 (1932).
- [75] Tin-Lun Ho and Lan Yin, *Phys. Rev. Lett.* **84**, 2302 (2000).
- [76] B. D. Esry, C. H. Greene, J. P. Burke, and J. L. Bohn, *Phys. Rev. Lett.* **78**, 3594 (1997).
- [77] The apparent phase space density is not correct in the sense that we are counting all of the atoms in the normal cloud regardless of which state they are in. This is similar to a situation with two condensates of different atomic species, say Rb and Na condensates in the same trap. If the total number was measured, the phase space density would be calculated to be factor of two larger than if only one state was correctly considered.

Appendix A

Everything you ever wanted to know about Rubidium

Property	symbol	Value
cooling transition		$5S_{1/2}, F=2 \rightarrow 5P_{3/2}, F=3$
nuclear spin	I	$3/2$
wavelength in vacuum	λ	780.23 nm
mass	m_{Rb}	1.44×10^{-25} kg
lifetime of upper state	τ_n	27 ns
natural linewidth	$\Gamma/2\pi$	5.9 MHz
absorption cross section	σ_{abs}	290.7×10^{-15} m ²
saturation intensity(stretched transition)	I_{sat}	1.6 mW/cm ²
recoil temperature	T_{rec}	180 nK
recoil velocity	v_{rec}	0.59 cm/s
ground hyperfine splitting	ω_{hf}	6834.68261090434(3) MHz
$ 1, -1\rangle$ s-wave scattering length	a_{11}	100.44 a_0
$ 1, -1\rangle/ 2, 1\rangle$ s-wave scattering length	a_{12}	98.09 a_0
$ 2, 1\rangle$ s-wave scattering length	a_{22}	95.47 a_0

Table A.1: Properties of ⁸⁷Rb




Figure A.1: Energy level diagram for ^{87}Rb showing Landé g factors for each state.

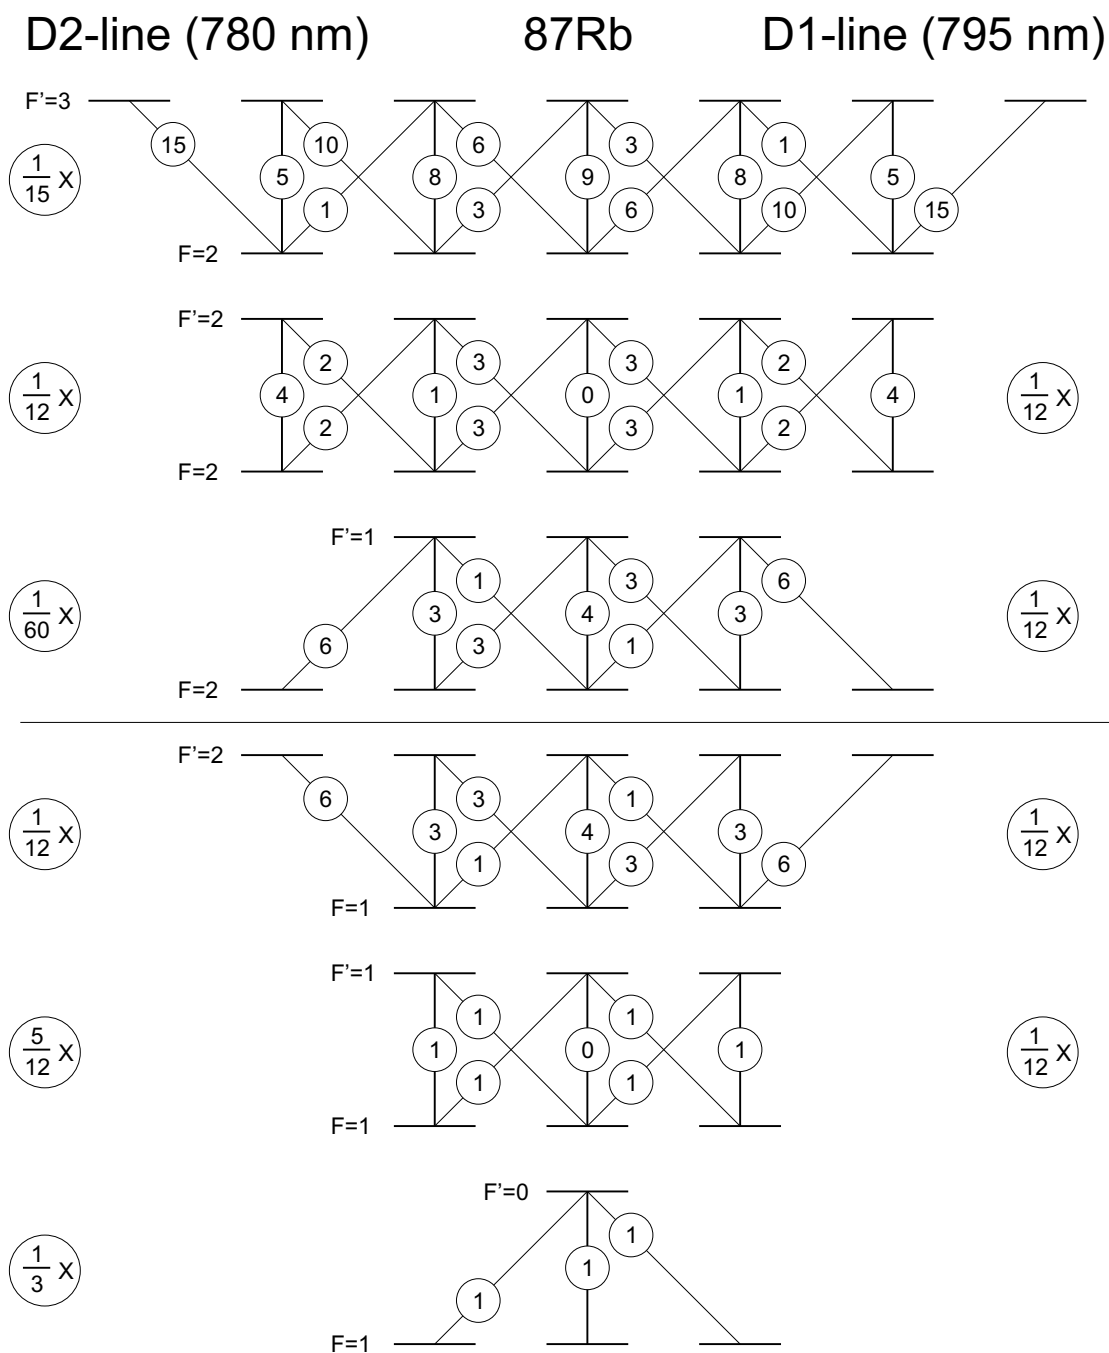


Figure A.2: Branching ratios for ^{87}Rb . Multiply by the circled number in the left(right) column to get the branching ratio for the D2(D1) line.