

Bose-Einstein Condensates with Tunable Atom-atom Interactions:

The First Experiments with ^{85}Rb BECs

by

Jacob Lyman Roberts

B. S., University of Notre Dame, 1994

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

2001

This thesis entitled:
Bose-Einstein Condensates with Tunable Atom-atom Interactions:
The First Experiments with ^{85}Rb BECs
written by Jacob Lyman Roberts
has been approved for the Department of Physics

Carl Wieman

Deborah 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.

Roberts, Jacob Lyman (Ph. D., Physics)

Bose-Einstein Condensates with Tunable Atom-atom Interactions: The First

Experiments with ^{85}Rb BECs

Thesis directed by Distinguished Professor Carl Wieman

In this thesis, the creation of magnetically trapped atomic gas Bose-Einstein Condensates (BECs) with tunable atom-atom interactions is described. The atom-atom interactions are tuned by adjusting an external magnetic field through the phenomenon of a Feshbach resonance. The strength of the interactions can be tuned to any desired value, and the interactions can be made to be repulsive or attractive. Since many properties of BECs are determined by the strength of the atom-atom interactions, such as the spatial size, the collective excitation frequencies, and the stability of the BECs, the ability to tune the interactions allows a wide range of experiments to be performed.

The work in this thesis details the characterization of a Feshbach resonance in cold collisions between ^{85}Rb atoms. Both the elastic and inelastic collision rates are measured as a function of the external magnetic field, showing variations of four or more orders of magnitude as the external magnetic field is changed by a few gauss. Using the knowledge gained about the variation of the elastic and inelastic collisions near the Feshbach resonance, a stable BEC of up to 18,000 ^{85}Rb atoms was created through evaporative cooling. The spatial size of these BECs responded as predicted by mean-field theory as the strength of repulsive interactions was changed by tuning the external magnetic field. The mean-field description is

predicted to have a limited range of validity as the strength of the repulsive interactions is increased further, however, and in a different experiment we observed deviations from the mean-field predicted values of collective excitation frequencies of the BECs as the strength of the interactions was increased. Work was also performed on BECs with attractive interactions. For strong enough attractive interactions, the BECs are predicted to become unstable, and the onset of this instability as a function of the number of atoms in the BEC and the strength of the attractive interactions is reported here. The dynamics of the unstable BECs as they collapsed are also described in this work.

Dedication

I dedicate this thesis to my wife, Sarah, in recognition of her love and support,
and to my daughter, Abigail, for continually reminding me of the joy of discovering
something new.

Acknowledgements

While this work did occur in a vacuum, I certainly did not do it alone. I would like to take this customary opportunity to sincerely acknowledge those who have contributed to both directly to the work presented in this thesis and also to my education and life in Boulder over the past seven years.

First and foremost I would like to acknowledge my advisor, Carl Wieman, for providing me with a terrific graduate education. He is an excellent scientist who is innovative and possesses a good intuition. Carl is also honest and direct, and he has shown me how important those qualities are for performing good research. I am not sure how important a good sense of humor is for performing research, but I appreciated Carl's throughout my time here.

I also acknowledge Eric Cornell's contributions to my professional development and directly to the work presented in this thesis. It has been good to work with him, and I have been continually impressed with his ability to generate new ideas for directions of research and his skills in critically analyzing any scientific problem.

Neil Claussen has been a good colleague and friend of mine over the past several years. I have truly enjoyed working with him. I appreciate all of the hard work that he has put into the experiments described in this thesis. His skill at maintaining and improving the experimental apparatus certainly allowed me to

graduate sooner than I otherwise would have. I am sure that his thesis will be excellent and that lots of interesting new physics will be discovered using ^{85}Rb BECs.

It has been good my good fortune to work with the two post-docs who have worked on these experiments during my tenure. Simon Cornish was greeted very rudely by the experimental apparatus when he arrived but we managed to tame it well enough to get good results. His insights and pleasant nature were greatly appreciated. It was great to have Liz Donley join me for the last stages of my thesis work, and good to see her again after the first few years we spent as graduate students together. She was doing good work as soon as she arrived. It is nice to know that with Neil and Liz working on the experiment, it will be in capable hands in the future.

I would like to acknowledge the members of my thesis defense committee: Carl Lineberger for always having a kind word when I run into him in the halls, Chuck Rogers for his instruction on how to be a TA, Debbie Jin for her useful comments in group meetings and for taking the time to read this thesis critically, and Chris Greene for useful comments on some of the theory presented in this thesis, productive collaboration on much of the atomic collision physics presented here, and lots of help answering questions throughout the years.

It was a pleasure to collaborate with Jim Burke on many different aspects of atomic collision physics. Understanding the experiments was a lot easier with his help in interpreting the collision theory implications of the data collected. It was good to have him as a colleague here at JILA, and I appreciated his continuing

interest in our work after he left JILA. I also enjoyed interacting with and getting to know other theorists at JILA: John Bohn, Jinx Cooper, and Murray Holland.

Marco Prevedelli and Josh Milstein both worked on ^{85}Rb BEC experiments for brief amounts of time while I was here, and their contributions are appreciated. Undergraduates Nick Cizek, Alex Kane, and John Obrecht also contributed to the ^{85}Rb experiments.

From the time of my life that I was involved with the measurement of parity nonconservation in Cesium, I would like to acknowledge both Chris Wood and Steve Bennett. Chris taught me a lot about what it takes to perform a good precision measurement, and that sometimes there is no substitute for hard work. Steve was also good to work with, and I appreciated his patient explanations. Even though I did not spend as much time working with Dong Hyung Cho, I learned a lot from him and have enjoyed seeing him again off and on through the years.

In between parity nonconservation and ^{85}Rb BEC, I worked directly with Kristan Corwin for a brief period of time. However, I worked indirectly with Kristan for six years and was the better for the experience. She had a great knack for asking the right questions about any experimental problem and I value our friendship greatly.

I would like to acknowledge Eric Burt, Rich Ghrist, and Chris Myatt not only as colleagues, but also because they were responsible for assembling the apparatus used in this thesis work. Almost everything has broken at least once in the past several years, but many of the basic components are still in use.

JILA has been a wonderful place to work, and I have been blessed with a large number of good friends and coworkers. Jason Ensher is one who has been a great friend and coworker, not only helping me with my work at JILA but also bailing me out on numerous occasions when I have let an attacker slip by on the soccer field. It has been good having Stephan Duerr across the hall, both as a friend and also for when I needed to discuss a physics problem with someone. David Hall, Brian Demarco, Mike Matthews, and I managed not only to interact well while at JILA but were also able to avoid killing each other while traveling in Italy, although I don't want to think about what would have happened if it had taken another 20 minutes to find the hotel in Milan. Other people I would like to mention at this point are Eric Abraham, Brian Anderson, Mike Anderson, Brett Esry, Paul Haljan, Tetsuo Kishimoto, Simon Kuppens, Heather Lewandowski, Zheng-Tian Lu, Kurt Miller, Dirk Müller, Nate Newbury, Brad Paul, Harold Parks, Peter Schwindt, Michelle Stephans, Kurt Vogel, and Dwight Whittaker.

One of the other things that makes JILA a special place to work is the friendliness and quality of the support staff. I appreciate all of the help and education that I received from the JILA instrument shops (especially Dave Alchenberger, Hans Green, and Blaine Horner), electronics shop (especially Terry Brown and James Fung-a-Fat), scientific communications office (especially Marilee DeGoede and Laurie Kovalenko), and supply office (especially Maryly Doyle, Ed Holliness and Brian Lynch). Thanks also go to the many administrative support personnel at JILA for keeping everything running smoothly, and in particular to

Krista Beck, Fran Haas, Karen Melcher, and Barb Tennis for lots of assistance to me personally.

Many other people outside JILA have significantly enriched my life in Boulder, and while they have not contributed directly to the work in this thesis, I would like to acknowledge their friendship and support. The parishioners and staff at St. Thomas Aquinas helped make Boulder feel like home almost right away and I will always treasure the friendships formed there. Torrie Derbach, Kim King, John Metz, Jim Mooney, Sarah Parks, Orion Poplawski, and Heather Robinson have been good friends as well as soccer teammates over the past several years.

Finally, I would like to acknowledge the love and support of my wife, Sarah. She has helped me achieve all that I have achieved, and put up with the many long hours and days away required to complete the work presented here. I appreciate the sacrifices that she has made.

Table of Contents

Chapter I: Introduction.....	1
1.1 Overview of the Thesis	1
1.1.1 Introduction.....	1
1.1.2 Scattering Theory and Feshbach Resonances (Chapter II).....	3
1.1.3 Bose-Einstein Condensation and Mean-Field Theory (Chapter III).....	3
1.1.4 Tools of the Trade: ^{85}Rb Bose-Einstein Condensation Apparatus (Chapter IV).....	6
1.1.5 Resonant Magnetic Field Scattering in Cold ^{85}Rb (Chapter V).....	7
1.1.6 Magnetic Field Dependence of Ultracold Inelastic Collisions Near a Feshbach Resonance (Chapter VI).....	7
1.1.7 Stable ^{85}Rb Bose-Einstein Condensates with Widely Tunable Interactions (Chapter VII).....	8
1.1.8 Collective Excitations and an Initial Search for Beyond Mean-field Effects (Chapter VIII).....	9
1.1.9 Controlled Collapse of Bose-Einstein Condensates with Attractive Interactions (Chapter IX).....	9
1.1.10 Dynamics of Collapsing Condensates and Future Experiments (Chapter X).....	9
1.1.11 Notation.....	10
1.1.12 Prior Publication of Work in this Thesis.....	10
Chapter II: Scattering Theory and Feshbach Resonances	11
2.1 Introduction.....	11
2.2 Basic Scattering Definitions.....	12
2.3 The One-dimensional Scattering Problem---Using the Time-Independent Schrödinger Equation to Determine the Scattering Amplitude.....	14
2.3.1 Motivation.....	14
2.3.2 Gaussian Wavepackets.....	14
2.3.3 Scattering from a Step Function.....	15
2.4 Scattering in Three Dimensions.....	19
2.4.1 Comparison to One-dimensional Case.....	19
2.4.2 Boundary Conditions in Three Dimensions.....	19
2.4.3 Partial Waves.....	20
2.4.4 Phase Shift and Cross Section.....	22
2.5 The Born Approximation.....	23
2.5.1 First Born Approximation.....	23
2.5.2 Range of Validity of the First Born Approximation.....	24
2.6 Scattering from a Square Well Potential.....	25
2.6.1 Motivation.....	25
2.6.3 Scattering Length and the Position of the Last Bound State.....	26
2.6.4 Zero collision cross section at finite collision energy.....	29
2.7 Square Well Scattering with another Internal Degree of Freedom---.....	29
2.7.1 General Description.....	29
2.7.2 Toy model Hamiltonian.....	32

2.8 Inelastic Scattering	36
2.9 Rb-Rb Interatomic Potential	37
2.9.1 Born-Oppenheimer Potentials	37
2.9.2 Atomic hyperfine interaction.....	38
2.9.4 Matching the Calculated Born-Oppenheimer Potential to Measured Scattering Rates... 39	
2.9.5 Magnetic Dipole Interaction and Second-Order Spin Coupling	40
Chapter III: Bose-Einstein Condensation and Mean-Field Theory	41
3.1 Overview	41
3.1.1 Summary of the Chapter.....	41
3.2 The Phase Transition to Bose-Einstein Condensation in the Non-interacting Limit.....	42
3.2.1 Thermodynamics	42
3.2.2 Coherence Properties.....	45
3.3 Effect of Atom-atom Interactions on Bose-Einstein Condensates	46
3.3.1 Hartree-Fock Theory	46
3.3.2 Many-body Physics Formulation.....	50
3.3.3 Limitations to the GP equation.....	53
3.3.4 The Thomas-Fermi Limit	54
Chapter IV: Tools of the Trade: ^{85}Rb Bose-Einstein Condensation Apparatus	57
4.1 Overview	57
4.2 ^{85}Rb Atomic Structure.....	58
4.2.1 Ground and First Excited States	58
4.2.2 Ground State Hyperfine Structure and Energy Dependences at High Magnetic Fields .. 58	
4.3 Magneto-Optic Traps	64
4.3.1 Doppler Cooling	64
4.3.2 Sub-doppler Cooling	65
4.3.3 Magneto-Optic Traps (MOTs).....	68
4.3.4 Loading a MOT from a Vapor.....	73
4.4 The Double-MOT System.....	74
4.4.1 The Role of the MOT	74
4.4.2 Vacuum System.....	75
4.4.3 Baking Out the Vacuum System and Vacuum Failures Since 1995.....	78
4.4.4 Collection MOT	80
4.4.5 Transfer from Collection MOT to Science MOT	82
4.4.6 Science MOT.....	86
4.4.7 Control of Magnetic Fields, Laser Light, and Laser Frequency	91
4.4.8 CMOT and Molasses, Part I	93
4.4.9 Optical Pumping.....	94
4.5 Magnetic Traps	99
4.5.1 Principle of Operation	99
4.5.2 Ioffe-Pritchard Traps	100
4.5.3 Baseball Trap.....	103
4.5.4 Gravitational Sag and Typical Trapping Potentials.....	105
4.6 Classical Motion of Atoms in a Magnetic Trap.....	108
4.6.1 Classical Orbits.....	108
4.6.2 The Boltzmann Distribution Function and Thermal Equilibrium.....	109
4.6.3 Collective motions of the cloud in the absence of collisions	111
4.6.4 Ballistic Expansion.....	111

4.6.5 Sloshing.....	113
4.6.6 Breathing modes.....	113
4.6.7 Adiabatic Compression and Expansion.....	115
4.7 Optimizing the Loading of the Magnetic Trap from the MOT.....	117
4.7.1 CMOT and Molasses, Part II.....	117
4.7.2 Optical Pumping Optimization.....	121
4.7.3 MOT Shims and Sloshing.....	123
4.7.4 Other Spin States.....	125
4.8 Control of the Magnetic Trap Fields.....	125
4.8.1 Magnetic Field Requirements.....	125
4.8.2 MOSFET Banks.....	126
4.8.3 Parallel Baseball and Bias Coils.....	126
4.8.4 Magnetic Trap Current Control.....	130
4.8.5 Servo Optimization.....	137
4.9 Forced Radio-frequency Evaporation.....	137
4.9.1 Simple Picture of Radio-frequency Evaporative Cooling.....	137
4.9.2 Additional Details about Forced Rf-Evaporation.....	140
4.9.3 Coupling the Rf Field into the Science Cell.....	143
4.10 Calibration of the Magnetic Trapping Parameters.....	144
4.10.1 Calibration of the Trap Bias Magnetic Field.....	144
4.10.2 Calibration of the Magnetic Trap Frequencies.....	147
4.11 MOT Recapture and Absorption Imaging.....	150
4.11.1 General Description of Atom Number Diagnostics.....	150
4.11.2 Advantages and Disadvantages of MOT Recapture with Respect to Absorption Imaging.....	150
4.11.3 Optical Depth and Column Density.....	151
4.11.4 Preparation of Trapped Cloud for Imaging and Production of Imaging Light.....	154
4.11.5 Imaging the Atoms' Shadow onto a Camera.....	156
4.11.6 Lensing Effects From Dense Atomic Samples.....	163
4.12 Calibration of the Imaging System.....	164
4.12.1 Calibration of the Magnification.....	164
4.12.2 Probe Laser Frequency Stability and Single Mode Behavior.....	166
4.12.3 Measurement of the Probe Intensity and the Doppler Shift.....	168
4.12.4 Optical Pumping Effects in Absorption Imaging.....	171
4.12.5 The Maximum Observable Optical Depth and Forward Scattering.....	173
4.12.6 Quantization Field Orientation.....	174
4.12.7 Resolution Limit.....	175
4.12.8 Estimate of Absorption Imaging Accuracy.....	177
Chapter V: Resonant Magnetic Field Elastic Scattering in Cold ^{85}Rb	180
5.1 Introduction.....	180
5.1.1 External Electromagnetic Fields and Very-Low-Temperature Atomic Collisions.....	180
5.1.2 Motivation.....	182
5.1.3 Characterization of Atom-atom Interaction Potentials.....	183
5.2 Cross-dimensional Mixing and Thermal Equilibration.....	184
5.2.1 Elastic Collisions and Enforcement of Thermal Equilibrium.....	184
5.2.2 Calculation of the Equilibration Rate.....	186
5.2.3 Using the Equilibration Rate of the Aspect Ratio of Trapped Clouds to Measure the Elastic Collision Rate.....	195
5.2.4 Unitarity Limit and Thermal Equilibration.....	196
5.3 Creation of Nonequilibrium Trapped Clouds.....	197
5.3.1 Initial Trap Loading and Cooling.....	197

5.3.2 Bringing the Cloud out of Thermal Equilibrium	198
5.3.3 Anharmonic Mixing and the	198
5.4 1998 Search for the Feshbach Resonance	200
5.4.1 Relative Measurement of Elastic Collision Cross Sections	200
5.4.3 Determination of B_{peak} and B_{zero}	204
5.5 2000 Measurement of the Absolute Value of the Scattering Length	208
5.5.1 Inadequacies of the 1998 Experiment	208
5.5.2 Determination of B_{zero} from the BEC Stability Measurement	209
5.5.3 Measurement of the Absolute Size of the Elastic Scattering Cross Section	210
5.5.4 New Constraints on the Rb-Rb Interatomic Potential	213

Chapter VI: Magnetic Field Dependence of Ultracold Inelastic Collisions Near a Feshbach Resonance	216
6.1 Introduction	216
6.1.1 Motivation	216
6.1.2 Types of Inelastic Collisions in a Magnetic Trap	217
6.1.2.1 Spin-exchange Collisions	217
6.1.2.2 Magnetic Dipole Relaxation	218
6.1.2.3 Three-body Molecular Recombination	219
6.1.3 Loss Rate Equations for a Magnetically Trapped Sample	221
6.1.4 Heating Rates due to Inelastic Collisions	222
6.2 Measurement of the Inelastic Collision Rates	223
6.2.1 Sample Preparation	223
6.2.2 Measurement of the Density Evolution	223
6.2.3 Distinguishing Two- and Three-body Loss Rates	227
6.2.4 Character of the Loss Rates near the ^{85}Rb Feshbach Resonance	228
6.2.5 Temperature Dependence of the Loss Rates	228
6.2.6 Heating Rates	231
6.3 Comparison of Measured Inelastic Rates to Theory Predictions	231
6.3.1 Dipole Relaxation	231
6.3.2 Three-body Recombination	234
6.3.3 Energy Dependence of the Feshbach Peak	235
6.3.4 Conclusions	235
6.4 Appendix to Chapter VI: Polarization-Rotation Imaging	238
6.4.1 Why Choose Polarization-rotation Imaging for Non-destructive Imaging?	238
6.4.2 Arrangement of Polarization-rotation Imaging in our System	239
6.4.3 Performance of the Polarization-rotation Imaging	241

Chapter VII: Stable ^{85}Rb Bose-Einstein Condensates with Widely Tunable Interactions	244
7.1 Evaporative Cooling of ^{85}Rb	244
7.1.1 Difficulties Evaporatively Cooling ^{85}Rb	244
7.1.2 Elastic/Inelastic Collision Rates and Trap Frequency	245
7.1.3 Simple Model of Evaporative Cooling	250
7.1.4 General Evaporation Strategy for ^{85}Rb	254
7.1.5 Optimization of the Evaporation Trajectory	257
7.1.6 Comparison with Simple Model	263
7.1.7 Achieving Reproducible Evaporation Performance	265

7.1.7.1 Use of Inelastic Losses to Reduce Number Variation.....	268
7.1.7.2 Reproducibility Performance.....	269
7.2 Loss of Atoms From the Condensate When Ramping through the Feshbach Peak.....	269
7.2.1 Absorption Imaging and Ramping through the Feshbach Peak.....	269
7.2.2 Outward Impulse from the Magnetic Trap Turn-off.....	270
7.2.3 Loss of Condensate Atoms while Ramping through the Feshbach Peak.....	271
7.2.4 Comparison of our Observed Loss Rate to Coherent Loss Mechanisms.....	274
7.3 ^{85}Rb Condensates with Positive Scattering Lengths.....	277
7.3.1 Stability of the Condensates with Positive Scattering Length.....	277
7.3.2 Spatial Size of the Condensates as a Function of Magnetic Field.....	277

Chapter VIII: Collective Excitations and an Initial Search for Beyond Mean-field

Effects.....	286
8.1 Overview.....	286
8.2 Introduction to Collective Excitations.....	289
8.2.1 Condensate Excitations and the Gross-Pitaevskii Equation.....	289
8.2.2 Types of Excitations: Single-particle, Sound Wave, and Collective Excitations.....	290
8.2.3 The Non-interacting Limit.....	291
8.2.4 Collective Excitations with Mean-Field Interactions.....	295
8.2.5 Experimental Observations of Collective Excitations.....	299
8.2.6 Modes Frequencies between the Non-interacting and Thomas-Fermi Limit.....	300
8.3 Beyond The Mean-Field Theory.....	302
8.3.1 Pushing the Validity of Mean-field Theory.....	302
8.3.2 Calculations of the Energy of a Condensate Beyond the Mean-field.....	304
8.3.3 Finite Temperature Effects and Beyond-Mean-Field Theory.....	309
8.3.4 Predicted Magnitude of the Beyond-Mean-Field Shifts.....	310
8.4 Inducing Collective Excitations in ^{85}Rb	310
8.4.1 Changing the Scattering Length to Create Collective Excitations.....	310
8.4.2 Control of the Excitation Amplitude of Each Individual Mode.....	312
8.5 Measuring the Oscillation Frequencies.....	316
8.5.1 Experimental Method.....	317
8.5.1.1 Transverse Magnetic Field Imaging.....	317
8.5.2 Oscillation Frequency Signal-to-Noise Considerations.....	317
8.5.3 Condensate Heating and Loss Rates.....	319
8.5.4 Accounting for Number Variation.....	320
8.5 Results of the Oscillation Frequency Measurements.....	321
8.7 Distortions in the Measurement of Oscillation Frequencies.....	326
8.7.1 Measurement of the Axial Frequency at $a=0$	326
8.7.2 Measurement of the Oscillation Frequencies at Small a	330
8.8 Temperature-dependent Shifts in the Radial Oscillation Frequency.....	331
8.8.1 Frequency Shift as a Function of Time at High Loss and Heating Fields.....	331
8.8.2 Measurement of the Condensate Heating Rate.....	331
8.8.3 Evidence for Coupling between the Condensate and Thermal Cloud Motions.....	334
8.9 Estimation of the Magnitude of the Beyond-Mean-Field Frequency Shifts and Outlook for Future Measurements.....	336
8.9.1 Difficulties in Estimating the Beyond-Mean-Field Shift.....	336
8.9.2 Estimate of the Magnitude of the Beyond-Mean-Field Shift.....	339
8.9.3 Improving the Determination of the Magnitude of Beyond-Mean-Field Shifts with Additional Measurements.....	342

Chapter IX: Controlled Collapse of Bose-Einstein Condensates with Attractive Interactions.....	345
9.1 Attractive Interactions and Bose-Einstein Condensate Stability.	345
9.1.1 Bose-Einstein Condensate Mechanical Stability.	345
9.1.2 Theory Predictions for Condensate Stability.	347
9.1.3 Dynamics of the Condensate Instability.	351
9.1.4 Previous Work in ^7Li and Comparison to our Work.	351
9.1.5 Goals of this Measurement.	352
9.2 Measuring the Onset of Instability.	352
9.2.1 Creation of the Condensates to be Studied.	352
9.2.2 Ramping to $a < 0$ and the Measurement of the Number of Condensate Atoms.	353
9.2.3 Measuring the Temperature of the Condensates.	354
9.2.4 Monitoring the Number Variation and Drift.	355
9.3 Signature of the Collapse and Sharpness of the Transition from Stability to Instability.	356
9.3.1 Signature of the Collapse Event.	356
9.3.2 Sharpness of the Transition from Stability to Instability.	356
9.4 Fraction of Atoms Remaining in the Condensate After Collapse.	359
9.4.1 Fraction Remaining after Collapse.	359
9.4.2 Loss of Condensate Atoms at $a < 0$ in the Absence of Collapse and Unusual Statistics.	361
9.4.3 Increased Variation in the Collapsed Number.	363
9.5 Determination of the Stability Coefficient.	364
9.5.1 Measurement of the Onset of Instability as a Function of the Number of Atoms in the Condensate.	364
9.5.2 Comparison of the Measured Stability Coefficient to Theory Predictions.	370
9.5.3 Aside on more Sophisticated Analysis of the Collapse Data to Determine k	370
9.6 Search for Systematic Shifts in k due to Finite Temperature and Non-adiabatic Field Ramps.	371
9.6.1 Finite Temperature and the Onset of Instability.	371
9.6.2 Condensate Dynamics and the Onset of Instability.	373
9.7 Determination of B_{zero}	374
9.7.1 Improved Measurement of B_{zero}	374
Chapter X: Dynamics of Collapsing Condensates and Future Studies of ^{85}Rb Bose-Einstein Condensates	375
10.1 Overview.	375
10.2 Measuring the Collapse Dynamics.	375
10.2.1 Creating the Initial Condensates.	375
10.2.2 Ramping the Magnetic Bias Field to Change a	376
10.2.3 Imaging the Condensates.	376
10.3 Dynamics of the Collapse.	376
10.3.1 Number vs. τ_{evolve}	376
10.3.2 Contraction and Formation of Density Spikes.	379
10.3.3 Onset of Rapid Loss.	381
10.3.4 Explosion of	381
10.3.5 Remnant Condensate.	383
10.3.6 Missing Atoms.	385

10.3.7 Condensate Collapse After the Magnetic Trap Turnoff	385
10.4 Possible Future Experiments with ^{85}Rb BEC.	386
10.4.1 Revisiting the Region of Large Positive a	386
10.4.2 Thermodynamic Functions as a Function of Scattering Length.	387
10.4.3 Going Over the Feshbach Resonance Peak.	388
Bibliography	389

List of Figures

- Figure 2.1. This figure illustrates the elements involved in a scattering calculation. A beam of particles is incident from the left and then strikes the scattering potential, which is shown as the gray sphere. The potential deflects the direction of the particles in the incident beam, and a certain fraction of the particles will be scattered in the direction $d\Omega$, where they will strike a detector (the rectangle). In scattering theory, the precise motion of all the particles over their entire trajectory (including their motion in the region of the scattering potential) is not the end result of the calculation. Instead, the goal of a scattering calculation is to compute the quantity $d\sigma/d\Omega$, the fraction of atoms in the initial beam that are scattered into direction $d\Omega$ 13
- Figure 2.2. This figure illustrates the scattering of a wavepacket from a step potential. The incident wavepacket is shown at time $t=0$ and the reflected and transmitted wavepackets are shown after the incident wavepacket struck the step potential. The arrows indicate the direction of propagation of each of the wavepackets. 17
- Figure 2.3. The upper figure shows the functional form of the scattering length as a function of the depth of the attractive square well potential (V_0). Each divergence in the scattering length is associated with the formation of an $L=0$ bound state in the potential. The lower figure shows the square well potential at several different depths (solid line). The dotted line shows the energies of the $L=0$ bound states in the potential. From left to right, the scattering lengths shown in the lower figure are 4.55, 0.09, and 1.48 in the arbitrary units shown in the upper plot. 27
- Figure 2.4. This figure shows attractive square well scattering wavefunctions in the limit of zero collision energy for successively deeper depths of the square well. Actually, the wavefunction itself is not plotted; rather the function $u_E(r)$ that is related to the wavefunction as defined in the main text is what is shown. The range of the potential is set to be equal to 1 (note that the scale for r varies from plot to plot in the figure). The potential in (a) has no bound state and a scattering length of -0.56 . The potential in

(b) has one bound state near zero energy and a positive scattering length of 9.4. The potential in (c) has one bound state and a negative scattering length -1.84 . The position at which the wavefunction outside the potential ($r > 1$) would intercept the axis is equal to the scattering length. In this sense, the scattering length is equal to the position of the last zero in the wavefunction in the limit of zero collision energy. The solid line shows the wavefunction, while the dashed line is used to show where the wavefunction at ($r > 1$) would intercept the axis in the negative scattering length cases..... 28

Figure 2.5. Demonstration of the zero in collision cross section for a finite collision energy. The attractive potential used to generate this curve had several bound states and a negative scattering length. Similar structures can be observed with potentials with a positive scattering length, but the zeros in σ tend to be at higher energies..... 30

Figure 2.6. Feshbach resonance calculated using the toy model. In this case, the attractive tail of the potential was not included. The parameters used to calculate this curve were $V_0 = 125$, $R = 2$, $V_1 = 0$, $V_{in} = 8$, $V_{out} = 10$, $\Delta = 10$ 33

Figure 2.7. Feshbach resonance calculated in the toy model with the presence of the additional attractive tail on the Hamiltonian. The parameters for the calculation of the main curve in this figure were the same as in figure 2.6 except that $R_0 = 30$ and $V_1 = 0.00295$. The value of Δ for the curve in the inset was set to be 1, but otherwise the parameters were the same as those used to calculate the main curve. The fact that the resonance of the curve in the inset is narrower and the peak in the scattering length occurs at a different value of x demonstrates that both the width of the resonance and the location of the divergence depend on the value of the coupling parameter..... 34

Figure 3.1. The fraction of atoms in the lowest energy quantum state as a function of temperature is shown. This curve is calculated in the thermodynamic limit. The sudden increase in the fraction of atoms at $T = T_c$ is the phase transition in which a Bose-Einstein condensate is created..... 45

Figure 4.1. ^{85}Rb Hyperfine Structure of the Ground and $^3\text{P}_{3/2}$ Excited State. This figure diagrams the hyperfine structure of the two states that are involved in the laser transitions used in laser cooling and trapping, optical pumping and absorption imaging. All of the hyperfine level splittings are given in MHz. The trapping and hyperfine repump transitions used in the MOTs are shown explicitly. The lifetime of the excited state is 26.45 ns [51]. The level splittings are taken from Ref. [52] 59

Figure 4.2. Ground State Magnetic Sublevels. The ground state is (and excited states are) characterized by an additional quantum number, m_f ,

which is equal to the projection of the angular momentum F onto a particular axis. This axis is usually chosen to lie along the direction of the magnetic field at the atom's location. The hatched states indicate magnetic sublevels that increase in energy (through the Zeeman interaction) with increasing magnetic field at low magnetic fields. It is therefore possible to trap the hatched states in a purely magnetic trap. The reason that the $F=3$ positive m_f states are trapped while it is the negative $F=2$ m_f states that are trapped is that the Zeeman energy is determined primarily by the projection of the electron spin alone, and the electron spin is aligned with the nuclear spin in the $F=3$ state and opposite in the $F=2$ state 60

Figure 4.3. The energy of the $m_f = -2, -1,$ and 0 state as a function of magnetic field as determined by equation (4.1)..... 64

Figure 4.4. Frame (a) illustrates how the Doppler shift can lead to an increased scattering rate when an atom moves towards a red-detuned laser beam. The dashed line shows the frequency of the laser in the lab frame. The atom on the right moving into the beam sees the frequency of the beam shifted to the blue – towards the maximum scattering rate. The atom moving away from the beam sees the frequency shifted even farther to the red. Therefore, the atom moving into the beam scatters photons faster. Frame (b) shows the viscous force calculated from equation (4.4) with $I/I_{\text{sat}} = 1$, $\Delta = 10$ MHz and the linewidth and wavelength of the transition set to be equal to those of ^{85}Rb 66

Figure 4.5. Clebsch-Gordan Coefficients. The transition rate between different pairs of magnetic sublevels of the excited and ground states depends on the orientation of the electron spin and the polarization of the laser driving the transition. The matrix elements involved in the transition rate are proportional to Clebsch-Gordan coefficients that are determined from angular momentum sum rules, and these coefficients, multiplied by a factor of 100, are shown in this figure. Since the transitions are proportional to the matrix elements squared, to compare the transition rate between different pairs of magnetic sublevels the ratio of the square of the Clebsch-Gordan coefficients needs to be considered..... 69

Figure 4.6. MOT Confining Force. In one dimension, a linear gradient in a magnetic field combined with the polarization of counter-propagating trapping beams shown results in a spatially confining force. The example above is for the simplest state structure possible. As an atom moves from the center ($B=0$), Zeeman shifts in the excited state will cause the atom to scatter more photons from the beam toward which it had moved. This results in a restoring force pushing the atom back to the origin. 71

Figure 4.7. Simple schematic of the double-MOT vacuum system. Each arm of the science cell is round. The circular window on the science cell

shown in this side view (and its counterpart on the other side) is, unlike the other directions, at the end of an arm only a few mm long. Hence the science cell resembles a four-way cross. 76

Figure 4.8 Hexapole Guiding Fields. Six strips of permanent magnets oriented as shown are sent along the outside of the transfer tube (shown as the dotted circle). The fields from all the magnets cancel at the center of the guide and rise steeply near the edges. The lines with arrows show the general direction of the magnetic field 85

Figure 4.9. Schematic Diagram of the MOT Recapture Fluorescence Collection Optics. The pair of mirror bounces is necessary to bring the fluorescence light out from near the crowded region next to the science cell to a place where it can be more easily measured. Not shown in this diagram is the (5/16-inch diameter) mask placed over the first lens to better define the solid angle, among other things 89

Figure 4.10. Optical Pumping Scheme. The OP laser is represented by the solid black arrows and the trapping light used to pump atoms out of the $F=3$ state is represented by the dotted arrows. The detunings of the lasers from the excited hyperfine states is shown approximately and the $F'=4$ state is not shown. The goal of the pumping is to have all the atoms wind up in the $F=2$, $m_f = -2$ state. 95

Figure 4.11. Optical Pumping Efficiency as a Function of OP/Probe Laser Detuning. The optical pumping efficiency is defined as the ratio of the number of atoms magnetically trapped to the number of atoms remaining after the molasses stage. The vertical dotted lines show the positions of the $F=2$ to excited state transitions. The relative positions of the transitions are accurate but the overall location could be off by a few MHz, due to some uncertainty in the absolute frequency calibration. The absence of any pumping near the $F'=3$ transition is due to heating of the atoms from repeated photon scatters in the $F=2$ $m_f = -2$ ground state since that state is no longer dark. 97

Figure 4.12. Optical Pumping Efficiency as a Function of Pulse Time. Even though the initial rise of the higher intensity data is faster, the cloud takes about the same time to optically pump for both intensities. This behavior is not yet explained. The high intensity is a factor of 8 more than the $\sim 75 \mu\text{W}/\text{cm}^2$ low intensity. 98

Figure 4.13. Ioffe-Pritchard Trap Schematic. The arrows indicate the direction of the current through the coils and wires. 101

Figure 4.14. Schematic of Baseball coil. The shape of the baseball coil used for magnetic trapping is shown here along with the two bias coils. The

small gray circle indicates the minimum of the magnetic field. The resemblance between the baseball coil and the current bars and Helmholtz coil system described in Figure 4.13 can be seen by considering the direction of the current flowing in each part of the baseball coil. The current flowing along the parts labeled..... 104

Figure 4.15. Contour plots of a typical trapping potential (a) and its associated magnetic field (b). The contour lines are spaced 10 μK and 0.5 G apart respectively. The oval labeled..... 107

Figure 4.16 Anharmonicities turn sloshing into apparent heat. The top cartoon shows the trapped cloud oscillating around the center of the potential that is marked with the X. Two atoms in the cloud, A and B will oscillate in individual orbits, and A will have a larger amplitude than B even though they are both the same distance from the center of the trapped cloud. Since anharmonicities in the potential mean that the oscillation frequency depends on the size of an atom's orbit and A and B have different size orbits, after some time the two atoms will be oscillating more or less out of phase rather than in phase. The lower picture illustrates the relative position of A and B after they have shifted out of phase with each other. At this time, the range in orbit sizes in the trapped cloud will cause it to spread out over the entire region over which it was sloshing. Thus it appears that the cloud has increased in temperature. 114

Figure 4.17. Relative insensitivity of the trapping frequencies to initial baseball coil current. This figure demonstrates how increasing the amount of current in the baseball coil during the trap turn-cannot increase the trap frequencies from the typical conditions used by more than $\sim 25\%$, and that can only be achieved by using nearly twice as much current. This is because the magnetic moment of the atoms decreases as the bias field increases. The x axis is scaled to the usual trap turn on current. 120

Figure 4.18 Optimization of CMOT Stage Laser Detuning. This figure is typical of how the MOT parameters are adjusted for optimum performance. The filled circles show the number of atoms remaining in the trap after a fixed amount of evaporative cooling. The evaporative cooling is enough to reduce the temperature of the cloud to 2.2 μK . The open circles show the fraction of atoms that are loaded into the magnetic trap compared to the initial number of atoms that are loaded into the science MOT before the CMOT and subsequent stages. The decrease in transfer fraction is due to the fact that the losses during the CMOT stage increase as the detuning is increased. Presumably, the optimum CMOT detuning with respect to evaporation performance represents a trade-off between increasing the initial density of the sample (and cooling it a little bit during the CMOT stage) and losing atoms during the CMOT stage. For this data the CMOT stage was 30ms long, the typical time for that stage..... 122

- Figure 4.19. Baseball and Bias Coil Current Supply and Control Schematic. The FETs shown in the diagram are actually FET banks as described in the main text..... 127
- Figure 4.20. Schematic of the Old Arrangement of Coils and Current Controls [1]..... 128
- Figure 4.21. Schematic of the Baseball Coil Current Control Electronics. All of the individual components, except for the SRS Programmable Frequency Synthesizer which is a commercial product, will be detailed in the following figure. Each TTL signal comes from the Labview TTL output board. The bias coils are controlled by a similar chain of components, but without the synthesizer and its summing box and with no Bandwidth control TTL. 132
- Figure 4.22. Reference Box Circuit Diagram. All resistors are high-quality metal film. The 366.5 k Ω resistor shown in the box is actually several resistors in series that have that total resistance..... 133
- Figure 4.23 On/Off Box Circuit Diagram. The two OP amps after the LM399 are an artifact of the many changes that have been made to this circuit. In addition to turning the trap on and off, this circuit is also used to filter the high-frequency noise out of the Labview analog signal. 133
- Figure 4.24. Synthesizer Summing Box Circuit Diagram. 134
- Figure 4.25. Circuit Diagram for the servo. The switches shown are DG201s that are controlled via the bandwidth control TTL. The somewhat unusual configuration of capacitors and resistors on the lower left op amp is designed to deal with the coil's phase shift around 1 kHz described in the main text. From zero to 480 Hz the capacitors dominate and there is a 90° phase shift. At 480 Hz, the 10 k Ω resistor begins to dominate the 33 nF capacitor and the phase shift drops to 50° near 1 kHz. This reduction in phase shift is enough to allow the servo to run with a higher gain without the system oscillating..... 135
- Figure 4.26. The power spectrum of the noise with and without current flowing through the baseball coil is shown. The top plot is the hall probe signal, and since that is the error signal all that is demonstrated is how well the servo is working. The bottom plot shows the voltage drop across the coils. This is independent of the servo and so it is in some ways a better measurement of the true noise properties, but because of the inductance of the coil, the current flowing through the coil is not proportional to the voltage drop at frequencies higher than about 100 Hz. The DC signal sizes for the hall probe signal and voltage drop were 3.78 and 2.78 V respectively.

Not shown is a mild increase in the voltage drop noise near 11 kHz, the unity gain point of the servo. 138

Figure 4.27. Radio-frequency transitions leading to evaporation. The black arrows represent where the rf transition being driven is resonant with the atoms. Atoms move in the parabolic trapping potential (upper curve). The more energetic can reach the magnetic field where they will make a transition to an untapped state and be accelerated out of the trap. Thus the more energetic atoms can be selectively removed and the cloud cooled through evaporative cooling. 141

Figure 4.28. Magnetic Field Calibration. The point at which the two curves, ascending and descending, cross corresponds to the rf frequency that drives transitions at the bottom of the magnetic trap. Using the Breit-Rabi equation (4.1), the value of the magnetic field at the bottom of the trap can be computed. 146

Figure 4.29. Measurement of Sloshing Frequencies in the Axial and Radial Directions. The measurement of the axial trap frequency (top) and radial trap frequencies (bottom) was conducted at the same time. Each point represents a single measurement of the center of mass of the cloud while it sloshes in the trap. From the beat note pattern in the radial direction it is possible to extract all of three trap frequencies: 6.791(14) Hz axial and 17.259(8) and 17.493(17) Hz in the two radial directions. The last frequency corresponds to the direction that is not visible in the absorption imaging. 149

Figure 4.30. Calibration of the MOT Recapture Signal. The non-zero intercept is due to the creation during evaporation of hot atoms that are too diffuse to be detected by absorption but are still trapped (see Chapter VII for details). 152

Figure 4.31. Schematic of the Imaging Optics. 157

Figure 4.32. Depth of Focus Calculation. The depth of focus of the imaging system is calculated using geometric optics. The rays that originate at a position offset from the focus the imaging system are traced through the edge of the first lens. At the CCD array, these rays form a circle and the area of this circle can be calculated. The depth of focus is then the distance that the object is displaced from the focus when the area of the circle of rays on the CCD array is equal to the area of one pixel. This estimate will actually give a depth of focus that is slightly too small since not all of the rays from the object pass through the edge of the lens. Two-thirds of the lens radius is a more appropriate quantity since it is the average radius through which all of the rays pass. 158

Figure 4.33 Focusing the imaging system. The data in this graph are the apparent radial width of a sub-resolution limit cloud as a function of the position of the first lens in the imaging system. The probe laser frequency was set to be on resonance to avoid any lensing distortions. The focus point is taken to be the position of the minimum size of the radial width. 161

Figure 4.34. Using Gravity to Calibrate the Magnification of the Imaging Optics. By releasing the atoms from the magnetic trap and tracing the center of mass motion, from the measurement of the acceleration of the falling atoms due to gravity and the knowledge of g , the apparent size of a pixel (i.e. the magnification of the optics) can be determined. The black circles show the measured center-of-mass motion while the solid line is a fit to the data..... 165

Figure 4.35. Calibration of the Probe Detuning. The measured optical depth of a trapped cloud versus the detuning of the probe laser is used to determine the resonance of the $F=3 m_f = -3$ to $F'=4 m_f = -4$ transition in experimental units. The x axis shows the beat note between the probe laser and another laser locked to the peak of a saturated absorption transition. The variation in the points appears to be worse than it actually was; there was a drift in the size of the trapped clouds during the measurement that was easily removed when the center and width of the lineshape were being determined. The width of the transition from simple power broadening is expected to be 6.6 MHz, consistent with the measured width..... 167

Figure 4.36. Measured Optical Depth vs. Probe Intensity. This figure shows the effect of increasing probe intensity on the saturation of the atomic transition. By fitting the decrease in measured optical depth to equation (4.31), the intensity of the probe can be determined, calibrating the number of camera counts to the intensity of the probe at the location of the atoms. Each point consists of seven individual measurements, and the standard deviation of those seven measurements is used to calculate the errors shown. Note that the size of the fractional error decreases with increasing intensity. The atomic clouds used for this calibration were thermal clouds with a temperature of 115 nK. The probe laser was slightly detuned from resonance (1.1 MHz) when these data were taken. To avoid the potential systematic from optical pumping effects (subsection 4.12.4), only the four highest-intensity points were fit..... 169

Figure 4.37. Checking the Orientation of the Quantization Axis During Absorption Imaging. Each data point in this figure (filled circle) represents the ratio of the number of atoms measured at a particular up/down earth shim current to the number of atoms measured with that current equal to 0.2 A (open circle) during the absorption sequence. The decrease in measured atom number is due to the misalignment of the quantization axis caused by the application of the field from the up/down earth shim, which is

orthogonal to the main quantization field. The solid line is a parabola fit to the data. This parabola is a reasonable function since we expect that the measured number of atoms should be a maximum when the quantization axis is aligned along the probe direction. The orientation of the quantization axis was checked similarly in the other direction that is orthogonal to the main quantization field. The typical operating condition was 0.2 A, close to the peak of the curve, which fits to a maximum ratio of 1.0015. 176

Figure 5.1. This figure shows the approximate shape of the Feshbach resonance. It is calculated using the measured parameters for ^{85}Rb and equation (5.1). 181

Figure 5.2. Change in the Equilibration Coefficient κ as a Function of the Parameter β . The solid curve was calculated using equation (5.15). It shows that elastic collisions equilibrate the sample more efficiently when it is further out of equilibrium (higher β). The dashed curve uses equation (5.20) in addition to equation (5.15) to account for the effect of the unitarity limit on a sample whose initial temperature is 110 nK and whose scattering length $a = -342a_0$. These conditions correspond to the more precise 251 G point in subsection 5.5.3. 194

Figure 5.3. The points shown correspond to the measured normalized equilibration rates σ_N which are equal to the aspect ratio equilibration time divided by the density and average relative speed of the sample. The lines are calculations for the thermally averaged elastic cross section, not equilibration rates, and hence are not expected to fit the data points. 202

Figure 5.4. Figure 5.4(a) shows an expanded view of the region of Figure 5.3 near the peak of the Feshbach resonance, while Figure 5.4(b) does the same for the region near the minimum. 203

Figure 5.5 Constraints on the Rb-Rb Interatomic Potential. This figure compares the allowed v_{DT} , v_{DS} parameters space based on our work and measurements contemporaneous to our 1998 result. The large rectangle is the allowed range from Ref. [109]. The large diamond is the range from Ref. [91]. The small filled diamond is the range from our 1998 measurement. The large rectangle and large diamond were established using a value of $C_6 = 4550$ a.u. As is discussed in the main text, the position of the constrained region depends on C_6 , which is indicated by the small diamonds on either side of the filled diamond. These diamonds show the effect of 50 a.u. uncertainty in C_6 around the central value of 4700 (to which the filled diamond corresponds). The arrows indicate the sensitivity to changes in the resonance width and position in this parameter space. The small black square in the center of the small rectangle represents the constraints set by our more recent measurement and reflects the increase in precision of the measurement of B_{zero} . The small rectangle is another way of

representing the new constraint supplied by our more recent measurement. Unlike all of the other objects in this figure, the small rectangle does not correspond to a fixed value of C_6 . Rather, it is the total allowed range of v_{DT} , v_{DS} consistent with our measurement of B_{peak} and B_{zero} and the now smaller uncertainty in $C_6=4660\pm 20$ implied by our measurement of a_{bg} . In other words, the small rectangle is measuring the same constraint as the area between the two small open diamonds. 207

Figure 5.6. Relaxation of the Aspect Ratio to Equilibrium. The data shown were taken at $B = 169.7$ G. Each data point represents a single destructive absorption measurement of the cloud after it was released from the magnetic trap and allowed to expand for 1.6 ms. The solid line shows the fit to the data used to determine the elastic scattering cross section and hence the scattering length. The deviation of the data points from the fit around 3-4 seconds is due to the imperfect removal of the rolling motion. The lower plot shows the measurement of the individual axial and radial widths that went into determining the aspect ratio. 212

Figure 5.7. The value of a_{bg} implied by several aspect ratio relaxation measurements. This value was determined by dividing the s-wave scattering length determined at a particular magnetic field by the factor $\{1-(B_{\text{zero}}-B_{\text{peak}})/(B-B_{\text{peak}})\}$. The s-wave scattering lengths measured at 251.0 G were divided by an additional factor of 0.99 to account for the deviation of the approximate description of a vs. B form the full theory calculation. The error bars shown include the statistical errors plus the fit uncertainty due to the rolling systematic. The data represented by the open circles had a larger amount of rolling motion present than the data represented by the filled circles. There is no observable difference between the two sets of data. For clarity, the open circles are displaced 0.2G to the right of their true magnetic field value. The fit uncertainty is non-statistical and so the central values are clustered together more than the error bars would suggest. 215

Figure 6.1 Number of atoms versus time taken by polarization-rotation imaging at 159 G. Fits to both purely two-body and purely three-body inelastic rates are shown, illustrating the difficulty in separating two- and three-body loss processes. 224

Figure 6.2. Elastic Rate (upper plot) from Chapter V and total loss rates (lower) plot vs. magnetic field. The total loss rate is expressed as the sum of the two- and three-body loss as $K_2+\beta K_3$. Because of initial density differences caused by ramping across the peak, β has a different value above and below $B = 157$ G. The vertical lines represent the positions of the elastic rate maximum and minimum at $B = 155$ and 166 G, respectively. 226

Figure 6.3. The determination of the two-body inelastic rate for several B fields. The theory prediction from Ref. [119] is shown as a solid line. The

open circles (\square) are the two-body rates determined from the total loss from Figure 6.2 by assuming (not explicitly measuring) that the loss between 162 and 166 G is predominantly two-body. These points are added to aid comparison with theory. In addition to the statistical errors shown, there is another 20% systematic uncertainty due to the estimated error in our density determination. 229

Figure 6.4. The determination of the three-body inelastic rate. The points with the droplines are to be interpreted as upper limits on the three-body rate. The solid line here is a prediction from Ref. [122]. The open circles here are similar to those in Figure 6.3, only here we assume (but again do not measure) that the loss above 175 G is predominantly three-body. The error bars on the 250 G point are relatively small because a large amount of data was taken there. In addition to the statistical errors shown, there is another 20% systematic uncertainty due to the estimated error in our density determination. 230

Figure 6.5. Temperature dependence of the inelastic loss at 154.1 G. Here, the combined loss rate is shown by assuming that all of the loss at this field is due to a three-body process. All of the data on this plot were taken with approximately the same density. 232

Figure 6.6 Heating rate vs. magnetic field. The samples used to acquire this data had densities between $2-3 \times 10^{11} \text{ cm}^{-3}$. The heating rate is determined by dividing the temperature increase after the cloud had lost $\sim 25\%$ of the atoms by the time taken to lose that many atoms. Away from the zero, the heating rate is consistent with the rate expected from the density dependence alone. Near the minimum, the heating rate is higher than expected from the density dependence of the inelastic losses themselves, but near the minimum we are also more sensitive to heating rates from sources other than the inelastic collisions. 233

Figure 6.7 Conjecture log-log plot of K_3 vs. a^4 . This figure displays all of the data greater than 175 G assuming that the loss is due only to a three-body loss process. The slope of the best-fit line is 4.2, close to the predicted value of 4, but the quality of the fit is so poor (reduced $\chi^2 > 4$ with 2 degrees of freedom) that we certainly cannot claim that we have definitively observed an a^4 scaling. The intercept of this fit implies $C \sim 2200$ in equation (6.1). 236

Figure 6.8. Variation of measured loss rate of low-density clouds. The clouds used for this measurement had a density $\sim 10^9 \text{ cm}^{-3}$ and a temperature of approximately $40 \mu\text{K}$. Note that the zero is suppressed and the total variation is only 20%. 237

- Figure 6.9. Excited state magnetic sublevel energies vs. magnetic field. Each manifold of states is labeled by the value of m_j that is one of the quantum numbers that characterizes the eigenstates at high magnetic field. The individual levels within the four manifolds are further specified by the nuclear spin quantum number m_i . Moving up and down from one level to the next within a manifold changes the value of m_i by one. The m_i value for the highest and lowest states in each manifold is shown on the right. These quantum numbers are strictly good only at high values of magnetic field—at lower values of the field the eigenstates are a combination of the product states specified by $|m_j, m_i\rangle$. The arrow indicates the probe laser detuning suitable for polarization-rotation imaging. 240
- Figure 7.1. Elastic scattering cross section as a function of collision energy. This curve is calculated for $B=0$, and the behavior at fields far from the Feshbach resonance is similar..... 246
- Figure 7.2. The set of differential equations used in the simple model of evaporative cooling..... 253
- Figure 7.3. Optimum evaporation trajectory. The evaporation trajectory at 162.5 G is shown as the filled circles in this plot. The open circle represents the optical depth of the cloud of atoms when it is initially loaded. The cooling from ~ 45 to ~ 2 μK was performed at 251 G. The prediction of the optical depth as a function of temperature by the simple model is also shown as the solid line. To generate this curve, the following values were used in the model: $\eta=5$, $\kappa=1.6$, $K_2=5 \times 10^{-20} \text{ m}^3/\text{s}$, $a=167a_0$, $\tau=300 \text{ s}$, $\omega_r = 108.5$, $\omega_z= 42.4$. The value of κ was adjusted to fit the measured trajectory. 255
- Figure 7.4. Log plot of the frequency of the rf-knife as a function of time during the evaporative cooling at 162 G. Only stages 2-8 are shown. 81.209 MHz corresponds to the frequency that drives transitions at the bottom of the trap potential..... 260
- Figure 7.5. Comparison of Evaporation Trajectories at Different Magnetic Fields. At the peak of the 162 G evaporation trajectory, the bias field was ramped to a different value and the evaporation was then optimized at that new field. The 162 G curve was taken at a different time than the 162.5 G curve shown in Figure 7.3..... 264
- Figure 7.6. Evidence for an impulse imparted during trap turn-off. The filled triangles and circles represent the measured spatial size for condensates originally at 157 and 162 G respectively imaged 1.7 ms after the magnetic trap was turned off. The open symbols represent the measured widths when the condensates were imaged 5.7 ms after the trap turn-off. The observed size is larger for the longer expansion time data than for the shorter expansion time data, indicating that the condensates are expanding

outward outside the trap. The relative size increase is larger for the condensates that are originally smaller in spatial extent and are therefore denser. These data were taken soon after ^{85}Rb BECs were formed, before we were able to improve the reproducibility of the system or create condensates with more than a few thousand atoms, and so the number varied by a factor of three during the measurement. 272

Figure 7.7 Fractional of atoms lost following a rapid sweep of the magnetic field through the peak of the Feshbach resonance as a function of the inverse speed of the ramp. Data are shown for a condensate (\square) with a peak density of $n_{\text{pk}} \sim 3 \times 10^{12} \text{ cm}^{-3}$ and for a thermal cloud (\square) with a temperature $T = 430 \text{ nK}$ and a peak density of $n_{\text{pk}} \sim 0.5 \times 10^{12} \text{ cm}^{-3}$ 273

Figure 7.8. Fraction of atoms lost from a condensate following a rapid sweep over the Feshbach peak as a function of inverse ramp speed. The data shown here are similar to those displayed in Figure 7.7, only the peak density of these condensates is $n_{\text{pk}} \sim 0.2 \times 10^{12} \text{ cm}^{-3}$ 275

Figure 7.9 (color). False color images and horizontal cross sections of the condensate column density for various magnetic fields. The condensate number was varied to maintain a peak optical depth (OD) of ~ 1.5 . The magnetic field values are (a) $B = 164.9 \text{ G}$, (b) $B = 162.0 \text{ G}$, (c) $B = 158.1 \text{ G}$, (d) $B = 156.9 \text{ G}$, and (3) $B = 156.1 \text{ G}$ 279

Figure 7.10. Comparison of the shape of the condensate density and column density in the Thomas-Fermi limit along one spatial direction. 281

Figure 7.11. Measured condensate width versus the product of the scattering length and the number of atoms in the condensate, with that product raised to the $1/5$ power. From equation (7.20), the axial FWHM of a condensate is expected to scale as $Na^{1/5}$, and indeed the data do fall on a straight line. The slope of the line matches the prefactor in equation (7.20) to the precision of the measurement. The deviation from the fit line at small values of Na is discussed in the main text. 282

Figure 7.12. This figure shows a randomly selected subset of the data shown in Figure 7.11, only here the dependence of the width on the number and scattering length are plotted separately. It is difficult to make out details on this graph, but is possible to observe the spatial size increasing with N at a fixed a and the spatial size increasing with a at a fixed N 283

Figure 8.1. Oscillation of a condensate density in one dimension in a harmonic trapping potential in the non-interacting limit. The initial order parameter at $t=0$ was projected into the first 20 simple harmonic oscillator basis states, and from the evolution of the phase factor associated with each basis state the time evolution of the condensate density could be determined.

Since the initial condensate width is too large, it contracts as it evolves from $t=0$ 294

Figure 8.2. Oscillation of a condensate with a parabola-shaped initial density in the non-interacting limit. The time evolution of the condensate density was calculated in the same way as for Figure 8.1. The inset shows the presence of a spatial oscillation in the condensate density that exists at $t=\pi/2\omega$ (one half of the oscillation cycle). These ripples in the density were also present when 30 basis states were used to calculate the condensate evolution. While the size of the ripple is small, it can contain a significant fraction of the condensate atoms. For instance, if the condensate density shown was the same in all three dimensions, then 15% of the atoms would be contained in the ripples rather than the central feature. Notice that the central part of the order parameter is narrower at $t=\pi/4\omega$ than $t=\pi/2\omega$ 296

Figure 8.3. Normal modes of the $m=0$ collective excitations. The length of the arrows refers to the relative motion of the axial (horizontal) and radial (vertical) during the normal mode oscillation. Mode (a) is the higher-frequency $m=0$ mode, while (b) shows the motion of the other (lower-frequency) $m=0$ mode. 298

Figure 8.4. Comparison between the $m=0$ lower-frequency mode oscillation frequency calculated in an approximate model [148] developed by Perez-Garcia et al. and a full RPA calculation [41]. The approximate model is shown as the filled circles, while the RPA calculation is shown as a solid line. 303

Figure 8.5. Comparison of the total density with the beyond-mean-field shift correction and without. For simplicity, a spherical potential was assumed. The chemical potential between the two cases was set so that the number of atoms was the same in each case. The total density with the beyond-mean-field correction was calculated with equation (8.17) with $n_{pk}a^3$ set equal to 10^{-2} 307

Figure 8.6. Comparison of the observed radial and axial oscillation amplitudes to those predicted using the simple model described in the text. For these data the launch, intermediate, and final fields were 162.0, 158.3, and 156.5 G respectively. 315

Figure 8.7. This figure shows a typical measurement of radial oscillation frequency. There appears to be a small amount of damping, and this is reflected in the fit to the data. However, this damping may in fact have other causes, as described in section 8.5. These data were taken at a field of 158.7 G. 318

Figure 8.8. Dependence of the oscillation amplitude on condensate number. In this figure, the data have been binned by number into three sets and a sine wave was individually fit to each set. The two lower figures show the lack of any dependence of the measured frequency on the number and the oscillation amplitude of the fit to each set of data. The field at which these data were taken was 156.5 G. 322

Figure 8.9. Comparison of the measured oscillation frequencies to the mean-field and beyond-mean-field predicted frequencies. Most of the data points were taken over several oscillation periods and were therefore vulnerable to the temperature shift systematic discussed in the text. One data point taken over only two oscillation periods is also shown as the open circle, and the disagreement between that point and the others at the same field is additional evidence for the temperature dependent frequency shifts. The number of atoms used to measure the oscillation frequencies varied from field to field, and that accounts for the structure of the mean-field prediction curves. The dashed line shows the predicted oscillation frequencies from equation (8.16) for $B_{\text{zero}} = 154.9$, while the two dotted lines show the predicted frequencies as B_{zero} is shifted ± 0.4 G (see Chapter V). 323

Figure 8.10. Radial oscillation frequency measured at 155.7 G using a relatively short integration time. The average number of atoms in this set of data was about 800. 325

Figure 8.11. Measurement of the axial oscillation frequency near $a=0$ (165.85 G). Twice the axial frequency was 13.594 Hz, so the measured frequency does not agree with the expected value. The problem points, which are discussed in the text, are indicated by the arrows. 327

Figure 8.12. This figure shows the radial oscillation frequency measured at 156.1 G as a function of the time over which the oscillation was measured. For instance, if the radial oscillations were measured over times spanning 450 ms, the measured frequency was 35.27 Hz, while if the oscillations were only measured over times spanning 225 ms the frequency was measured to be 35.77 Hz. The decrease in measured frequency with increasing integration time indicates that the frequency of the oscillation is decreasing as the condensate oscillates at 156.1 G. 332

Figure 8.13. Measurement of the melting of the condensate at 156.1 G. The upper plot shows the measured condensate fraction as a function of time that the condensate spent at 156.1 G. No thermal component was originally visible, and so we estimate that the condensate fraction must therefore be greater than 90%. The lower plot shows the cross section of an absorption image used to measure the condensate fraction. The narrow component

corresponds to the condensate while the broader base is the thermal component..... 333

Figure 8.14. Radial oscillation measurement at 157.1 G. The measured widths are suggestive of the presence of a beat note in the oscillation, which would indicate that two frequencies were present in the oscillation. 335

Figure 8.15. Estimation of the beyond-mean-field shift in the presence of temperature-dependent fits. The data shown here are the data also displayed in Figure 8.12. The initial thermal fraction was assumed to be 10% for the fits shown here. The fitting function given in equation (8.22) is shown as the solid line, which is the best fit to the data, and the dotted line which represents the fit corresponding to the value of the intercept that is one standard deviation away from the best fit value. The parameter n is ~ 0.33 for the dotted line and so it corresponds to a temperature shift that changes linearly with the temperature of the cloud. 341

Figure 9.1. This figure shows the energy of a condensate whose density profile is fixed to be a gaussian as a function of the width of the condensate R (x-axis) and the value of the ratio $N|a|/a_{ho}$ (number label on each curve). Note that for the curve with $N|a|/a_{ho}$ equal to 0.7 there is no minimum in the energy as a function of R . The energy was calculated from equation (9.2).
..... 348

Figure 9.2. Comparison of Stable and Collapsed Condensates at $a < 0$. All of these images were taken with the mean-field expansion ramp described in the main text. Frame (a) corresponds to a stable condensate while frames (b) and (c) show collapsed condensates. All three images were taken under similar conditions, with the only difference being that in (b) and (c) $|a|$ was greater than in (a). The false color scale is not the same from image to image..... 357

Figure 9.3. Transition from stable to unstable condensates. This figure shows the fraction of atoms remaining as the magnetic field was ramped to higher magnetic fields (i.e. stronger attractive interactions). The hatched region indicates the expected 4 mG width due to the initial number variation. The initial number of atoms in the condensate was kept near 6400. At the field in the hatched region, the data fell into two groups and the individual points (\square) are displayed rather than the average of those points..... 360

Figure 9.4. This figure shows the truncated gaussian probability distribution used in analyzing the fraction of atoms lost after a ramp to fields with $a < 0$. The average of the initial number of atoms in the condensate before the ramp to $a < 0$ and the standard deviation of those points determine the location of the peak of the width of the overall gaussian distribution. The

probability distribution of the number of atoms in stable condensates is shown by the solid curve, while the probability distribution of the pre-collapse number of atoms in condensates that did collapse is shown as the dashed line. The critical number N_{crit} depends on the value of $|a|$ and therefore varies as a function of magnetic field. Both the stable and collapsed condensate distributions are normalized appropriately before being used in a calculation..... 362

Figure 9.5. Determining the onset of collapse as a function of the condensate number and the magnetic field. Figures 9.5 (a) and 9.5 (b) show data sets used to determine the stability coefficient via the boundary that divides the data between collapse (\square) and non-collapse (\square) events. In Figure 2(a), the data were concentrated in two field regions to accurately measure the slope of the boundary line. In Figure 2(b), the initial number was varied over a larger range to illustrate the function form of the onset of instability. Due to initial number variations, some collapse points appear to be on the wrong side of the stability boundary..... 366

Figure 9.6. This figure shows stable condensate numbers known to be on the verge of instability (see text) as determined from several data sets like those in figs. 9.6(a) and 9.6 (b). The fit to these points provides the value of the stability coefficient k 367

Figure 10.1. This figure illustrates several magnetic field ramps used to measure the condensate collapse dynamics. Rather than plotting the magnetic field as a function of time, the corresponding value of the scattering length is shown. The last change in the scattering length to ~ imaged 1.6 ms after this trap turnoff..... 377

Figure 10.2. Number of condensate atoms as a function of the evolution time at a magnetic field with attractive interactions (τ_{evolve}). These data were taken using the ramp shown in Figure 10.1(a) at a field with a value of $a = \square 30 a_0$. The initial number of atoms in the condensate is shown as N_0 and the number of atoms remaining after the completion of the collapse is given as N_{remnant} . The solid line shows the exponential fit used to determine the decay constant during the rapid loss of atoms during the collapse..... 378

Figure 10.3. Jet images for a series of τ_{evolve} values for the conditions of Figure 10.2. The evolution times were 2, 3, 4, 6, 8, and 10 ms (from a to f). Each image is $150 \times 255 \mu\text{m}$. The bar indicates the optical depth scale. The jets were longest (i.e. most energetic) and contained the most atoms at values of τ_{evolve} for which the slope of the loss curve (Figure 10.2) was greatest. A tiny jet is barely visible for $\tau_{\text{evolve}} = 2$ ms (image a), which is 1.7 ms before the onset of the rapid loss of atoms from the condensate. The

images also show how the number of condensate atoms decrease with time.
 380

Figure 10.4 (Color). Plot (a) shows a conceptual illustration of a focus of the atoms in the explosion. Plot (b) shows an image of a refocused explosion. The temperature in the axial direction (along the direction of the long axis of the refocused cloud) determined from this data is 62 nK. Plot (c) Shows the radially averaged cross section of plot (b) with a Gaussian fit (red line) to the burst energy distribution. The central 100 nm were excluded from the fit to avoid distortion in the fit due to the remnant and the thermal cloud. The latter is present in the pre-collapse sample and appears to be unaffected by the collapse. The dashed line indicates the fit to this initial thermal component. Ramps like the one shown in Figure 10.1(c) were used to collect these data..... 382

Figure 10.5 (Color). This figure demonstrates the oscillations of the remnant by plotting the radial width as a function of time at the $a < 0$ field. The solid line is a fit to the widths composed of two sine waves. The two frequencies of the sine waves correspond to twice the radial and axial frequencies within the measurement error. The two false color images show the peaked structures observed in the condensates as the remnant is oscillating..... 384

List of Tables

Table 4.1: ^{85}Rb atomic structure constants necessary for the Breit-Rabi equation.....	62
Table 4.2. The average number of scattering events required to pump an atom from a particular initial m_f state to the $m_f = -3$ state. Also shown is the length of time these photons scatters would take. The time column is how many photons an atom already in the $m_f = -3$ state would scatter on the $F=3, m_f = -3$ to $F'=4 m_f = -4$ transition in the time that it took to pump the atom from a different m_f state over. For example, and atom initially in the $m_f = 0$ state would be pumped to the $m_f = -3$ state in an average time equal to the amount of time it would take an atom initially in the $m_f = -3$ state to scatter an average 12.7 photons.	172
Table 4.3. Summary of Effects Considered in Absorption Imaging.....	178
Table 5.1. Summary of Measured Feshbach Resonance Parameters.....	214
Table 7.1. Summary of evaporation parameters used in our optimum evaporation trajectory. Stage 1 occurred at 251 G while the remaining stages are at 162 G.	261

Chapter I

Introduction

1.1 Overview of the Thesis.

1.1.1 Introduction. When I joined this experiment in 1997, my coworkers and I began working with an ^{87}Rb Bose-Einstein Condensate (BEC) machine [1] that had been converted into an experiment involving the other stable isotope of Rb, ^{85}Rb . Our primary goal when we started this work was to create BECs of ^{85}Rb atoms, a feat which had not been accomplished before. The reason that we wanted to create BECs of ^{85}Rb is that a Feshbach resonance was predicted to exist in ^{85}Rb . As a result of the Feshbach resonance, the ^{85}Rb - ^{85}Rb interatomic potential could be tuned sensitively by an external magnetic field. This meant that the elastic and inelastic collision rates could be changed dramatically through varying the value of the applied magnetic field. Of more interest to us, the ability to change the ^{85}Rb - ^{85}Rb interatomic potential meant that it was possible to change the interactions between atoms in a BEC. In principle, the strength of the interactions could be tuned to any desired value by an applied magnetic field. The interactions could also be tuned to be either attractive or repulsive.

Even though the interatomic interactions in trapped atomic gas BECs are weak enough that they can be treated accurately within a mean-field approximation, they play an important role in the properties of BECs such as: the average energy per particle in the condensate, the collective excitation spectrum describing the motion of the condensate in the magnetic trapping potential, the spatial size of the condensates in the magnetic trap, whether or not the condensates will be stable or will collapse due to attractive self-interactions, and others. The ability to alter the interparticle interactions in a BEC nearly at will through using the Feshbach resonance was therefore very exciting and would open the door to many novel tests of the behavior of BECs in magnetic traps.

While the precision of the theoretical description of alkali atom cold collisions had improved so that a Feshbach resonance was predicted to exist in ^{85}Rb collisions, the value of

magnetic field at which the Feshbach resonance would occur was still only predicted to within several tens of gauss. Also, when we began our work a Feshbach resonance had not been observed in cold atom-atom collisions. Therefore, our first experiments were designed to locate the magnetic field position of the Feshbach resonance and characterize the “strength” of the resonance --- how fast the collision properties and atom-atom interaction strength changed as the magnetic field was changed. The variation of the elastic collision cross section as a function of magnetic field was used to locate the position of the resonance and measure the resonance strength. Since these parameters are sensitive functions of the details of the Rb-Rb interatomic potential, their measurement provided improved constraints on the description of that potential. The variation of the inelastic collision rates near the Feshbach resonance was also determined by measuring the rate at which atoms were lost from the trap as a function of magnetic field and density of the trapped atom clouds.

Initially, it was hoped that the same techniques used to create BECs of ^{87}Rb would facilitate the creation of BECs of ^{85}Rb without substantial modification. However, initial attempts at cooling ^{85}Rb towards the BEC transition temperature in the same way as ^{87}Rb showed that the two atoms in fact behaved very differently from one another, and that creating an ^{85}Rb BEC would be difficult. Whereas a magnetically trapped gas of ^{87}Rb was cooled very efficiently using the technique of evaporative cooling, ^{85}Rb cooled under the same conditions performed poorly, with atoms lost rapidly from the gas well before the BEC transition temperature was reached. We realized that this behavior was due in part to high inelastic collisions rates in ^{85}Rb that led to loss of atoms from the magnetic trap and thus reduced the evaporative cooling efficiency. The knowledge of the elastic and inelastic collision rates as a function of magnetic field was therefore used to optimize the evaporative cooling and create a BEC of ^{85}Rb atoms.

Once the ^{85}Rb BECs had been created, we performed a series of different experiments with them. We varied the strength of repulsive interatomic interactions and observed the spatial size of the condensates shrink and expand in response to the changing interaction strength. We also measured the loss of atoms from the BEC as a function of the rate that the magnetic field was swept through the high inelastic collision rate region of the Feshbach resonance. The collective excitation frequencies were measured as a function of the strength of the atom-atom interactions to search for

predicted deviations from the mean-field description as the atom-atom interactions were tuned to be strong. Finally, the instability induced in BECs by attractive scattering lengths was studied in a controlled fashion. Both the onset of the instability as a function of number and interaction strength and the dynamics of these unstable BECs were studied.

The rest of this introduction details the organization of this thesis, as well as summarizing the work that is presented in each Chapter.

1.1.2 Scattering Theory and Feshbach Resonances (Chapter II). For cold collisions, a single parameter, the s-wave scattering length (a), can be used to describe the binary elastic scattering properties. The value of a is determined by the interatomic potential between the two atoms in the colliding pair, and in particular a is very sensitive to the position of the last bound state of the interatomic potential. The elastic scattering cross section is determined by the value of a and is proportional to a^2 . In ^{85}Rb , a Feshbach resonance was predicted to exist in collisions between two ^{85}Rb atoms in a magnetic sublevel state that could be magnetically trapped. Near a Feshbach resonance, the value of a can, in principle, be tuned by changing the external field to any value, positive or negative, large or small. Chapter II discusses the relationship between the interatomic potential and the s-wave scattering length, the origin of Feshbach resonances, and other cold collision properties.

1.1.3 Bose-Einstein Condensation and Mean-Field Theory (Chapter III). The ability to alter the cold elastic collision properties through the use of an external magnetic field was interesting to us from the perspective of cold collision physics, but measuring cold collision properties was not our primary motivation. The promise of such dramatic control of the atom-atom interactions through the use of an external magnetic field suggested that interesting experiments could be performed using BECs, since the interactions between the atoms in a BEC have an important influence on its behavior. The magnitude of the interaction energy is proportional to the s-wave scattering length, at least to first order. The ability to tune the s-wave scattering length meant that the interactions in a BEC could be controlled as well, making the interactions small or large, repulsive or attractive, opening the door to a wide range of possible experiments.

Before discussing how the interactions between atoms affect the properties of a BEC, I will briefly describe the phenomenon of Bose-Einstein condensation. In 1924, Satyendranath Bose [2] developed a statistical argument to derive the blackbody radiation spectrum. Albert Einstein extended this treatment to an ideal gas of monatomic atoms [3], and immediately noticed an interesting consequence of the theory: at low temperatures the number of atoms in states other than the lowest energy quantum state saturated. Einstein then postulated that a gas with a number of particles greater than the saturated value would have two components to it: one component would be a gas with the saturated number of particles, and the remainder of the particles would condense into the lowest quantum state of the system. Thus, the idea of a Bose-Einstein condensate was created. As a gas of particles described by these statistics is cooled below a critical temperature, it undergoes a phase transition characterized by discontinuities in various thermodynamic quantities such as the specific heat and also by a sudden macroscopic occupation of the number of particles in the lowest energy quantum state of the system.

The statistics developed by Bose and Einstein (referred to as Bose-Einstein statistics) is now recognized to apply only to particles with an integer spin (bosons). Other particles with half-integer spin (fermions), obey a different type of statistics, which is referred to as Fermi-Dirac statistics. Compound particles, such as ^{85}Rb which is composed of 85 nucleons and 37 electrons, obey the statistics that correspond to the total spin of all of the particles of which it is composed. Thus, ^{85}Rb is a boson.

The transition temperature at which BEC occurs is characterized by a quantity called the phase-space density. The phase-space density is equal to $n\lambda^3$ where n is the density and λ is the thermal deBroglie wavelength of the atoms, which characterizes the spatial spread of an atom's center of mass position due to the uncertainty principle. As the sample is cooled, λ becomes larger as the average momentum of the particles decreases. Once $n\lambda^3 \sim 1$, the wavefunctions of the individual atoms begin to overlap and at a critical value (2.612 for a homogenous gas) the phase-space available to bosons "fills up" and atoms begin to pile up in the ground state, signifying the onset of the BEC phase transition.

The well-known superfluid behavior of liquid Helium below $\sim 2\text{K}$ is associated with the formation of a BEC in the liquid [4]. In this system, however, the interactions between particles are very strong and the BEC cannot be described quantitatively by a theory that is based on the non-interacting gas and includes the atom-atom interactions as a perturbation. Macroscopic occupation of a single quantum state is also used to explain superconductivity and the superfluid properties of ^3He through the pairing of fermions. Again, the interactions in these systems are very strong and so a perturbative description of the particle-particle interactions cannot be used. One other condensed matter system of note is an exciton gas created in copper oxide through the application of a short laser pulse [5]. The excitons are long-lived enough that they can come to thermal equilibrium, and their effective mass is low enough that the BEC transition temperature can be reached using cryogenic methods. Unlike the other strongly interacting systems described above, the exciton gas is dilute in the sense that the interparticle spacing is much greater than the exciton radius. The momentum distribution of the exciton gas cannot be measured directly, and the presence of a BEC in the gas is deduced from other means, such as the recombination rate.

In 1995, BECs were created in magnetically trapped samples of dilute atomic gases. ^{87}Rb was condensed first here at JILA [6], and at MIT ^{23}Na [7] was condensed shortly thereafter. There was also evidence for BEC in ^7Li at Rice [8]. In all of these cases, laser cooling and trapping techniques [9] were used to collect and pre-cool a gas of alkali atoms. These atoms were then loaded into a purely magnetic trap and evaporative cooling was used to cool the gas to the BEC transition temperature. The method of evaporative cooling followed ideas originally developed in the effort to create a BEC in atomic hydrogen [10], which was eventually successful in 1998 [11].

The achievement of BEC in dilute gases is remarkable for several reasons. One remarkable feature from a technical perspective is the low temperature of the samples. For instance, we have observed temperatures below 3 nK in our samples. In the condensed matter systems mentioned in the previous subsections, the systems can be treated as being uniform in space (homogenous) and the phenomenon of BEC occurs only in momentum space. In the trapped gas systems, the inhomogenous potential means that the BEC exhibits features in both position and momentum space. It is therefore possible to measure the properties of the condensate by imaging the density

distribution directly through using a probe laser to illuminate the sample. This is useful as a tool to measure the condensates, but it is also satisfying to observe the behavior of a quantum mechanical object simply by looking at it. Additionally, the atomic gas condensates are dilute so that the effect of interatomic interactions can be included perturbatively in a mean-field approximation to a high degree of accuracy. The strength of the interatomic interactions can be determined independently from measurements not involving the properties of BECs and so quantitative comparisons between the theory and observed behavior of the condensates can be made.

The importance of the presence of interactions between atoms in the atomic gas condensates was readily demonstrated by experiments such as those measuring the density profile, average energy per particle, collective excitation frequencies, and stability of BECs [12]. The presence of atom-atom interactions in the condensates caused all of these quantities to be very different from their values calculated in the ideal (i.e. non-interacting) limit. The mean-field theory developed to describe such dilute condensates did an excellent job of describing the results of all the experiments performed, at least in the limit of zero temperature.

Given the importance of the atom-atom interactions in the properties of BECs, it is understandable why it would be desirable to tune the magnitude of these interactions through changing the value of a using a Feshbach resonance. Since the value of a can be changed, it is possible to experimentally measure some property of a BEC as a function of scattering length. It is also possible to investigate how sudden changes in a can influence the dynamics of BECs. The strength of the interactions can be increased to test the range of validity of the mean-field description of BECs, and regions where condensates cannot be formed (e.g. large negative scattering lengths) can be reached by creating a stable condensate and then changing the scattering length.

Chapter III summarizes Bose-Einstein statistics and the BEC phase transition. It discusses the formulation of the mean-field description of the condensate, both from a Hartree-Fock formalism and a many-body physics perspective. Also, the approximations that are made as part of the mean-field theory are discussed.

1.1.4 Tools of the Trade: ^{85}Rb Bose-Einstein Condensation Apparatus (Chapter IV).

The apparatus used to create and study ^{85}Rb BECs was the same as that used to create and study ^{87}Rb

even if the optical and magnetic trap parameters used to collect and cool the atoms were quite different. There were several steps necessary in any experiment we performed. ^{85}Rb atoms were collected and cooled in a magneto-optical trap, then loaded into a purely magnetic trap, evaporatively cooled to the desired temperature, and then usually manipulated by changing the magnetic field in some way. After that the atoms were released from the trap and measured using either absorption imaging or fluorescence detection. Chapter IV discusses how all of these things are performed in our apparatus, as well as describing the hardware of the apparatus itself. The calibrations of the magnetic field, magnetic trap, and imaging are also discussed in detail. Finally, the atomic ground and excited state structure of ^{85}Rb is described in Chapter IV.

1.1.5 Resonant Magnetic Field Scattering in Cold ^{85}Rb (Chapter V). Our first task was to locate the value of the magnetic field where the Feshbach resonance occurred. This was performed by measuring the time that trapped clouds of ^{85}Rb atoms took to reach thermal equilibrium. Since the equilibration occurred through elastic collisions, the equilibration time could be used to determine the elastic collision rate and hence the magnitude of the s-wave scattering length as a function of field. We observed the equilibration time to change by four orders of magnitude with a several gauss change in the value of the external magnetic field, clearly indicating the presence of a Feshbach resonance. The magnetic field location and width of the Feshbach resonance are sensitive functions of the exact details of the Rb-Rb interatomic potential, and so the measured location and width of the Feshbach resonance improved the knowledge of the Rb-Rb interatomic potential used to calculate cold collision properties. Chapter V details our equilibration measurement, specifying both the value of a as a function of magnetic field near the Feshbach resonance and discussing the constraints that are placed on the Rb-Rb interatomic potential as a result of this measurement.

1.1.6 Magnetic Field Dependence of Ultracold Inelastic Collisions Near a Feshbach Resonance (Chapter VI). Once our characterization of the elastic collision cross section near the Feshbach resonance was completed, we initially tried to create a BEC of ^{85}Rb atoms. Since a large elastic scattering rate is in general good for efficient evaporative cooling, we went to the magnetic field with the largest value of the s-wave scattering length and tried evaporatively cooling there. The

results were very unfavorable due to the presence of large heating and loss rates. From this we realized that the inelastic collisions that led to loss and heating rates of the trapped clouds of atoms needed to be measured as well as the elastic collisions in order to optimize the evaporation performance and create an ^{85}Rb BEC. Also, the characterization of the inelastic collision rates in non-condensed clouds of atoms would provide useful information in the future when performing experiments with BECs. If any loss was observed from a BEC, the inelastic collision rates must be known in order to ascertain if the loss is due to some mechanism associated with the physics of the BEC or whether it is simply due to the inelastic collisions that are occurring between the atoms in the cloud. Chapter VI contains our measurement of the inelastic collision rates near a Feshbach resonance. We found that the inelastic collision rates near a Feshbach peak had a complicated structure, but in general when the magnitude of the scattering length was large the inelastic collision rates tended to be large as well. The magnitude of the inelastic collision rates in ^{85}Rb was found to be orders of magnitude greater than in ^{23}Na , ^{87}Rb or ^7Li even at magnetic fields where the loss rates were near a minimum, suggesting that the evaporative cooling of ^{85}Rb would be difficult as compared to those atoms.

1.1.7 Stable ^{85}Rb Bose-Einstein Condensates with Widely Tunable Interactions (Chapter VII). Armed with the knowledge of how the elastic and inelastic collision rates scaled as a function of magnetic field, we proceeded to evaporatively cool the ^{85}Rb atoms as best we could. The value of the magnetic field determines both the elastic and inelastic collision rates and so there is no way to independently adjust the ratio of the two. Therefore, the value of the magnetic fields at which the evaporation occurred had to be carefully selected. We were eventually able to create BECs of ^{85}Rb , but because of the relatively high inelastic collision rates the number of atoms in the condensates was limited to about 15,000. The evaporative cooling never performed as well as for ^{87}Rb , but it was good enough to reach the BEC transition temperature.

Chapter VII describes evaporative cooling of ^{85}Rb atoms down to BEC temperatures. It also contains two other measurements. The number of atoms lost in rapid sweeps across the high loss region of the Feshbach resonance was measured. Also, the change in the spatial extent of the BEC ground state was measured as a function of scattering length with $a > 0$. Agreement was found

between the observations of the BEC ground state spatial extent and the predictions of mean-field theory over the range of values of a that were measured.

1.1.8 Collective Excitations and an Initial Search for Beyond Mean-field Effects (Chapter VIII) Once the elastic and inelastic collision rates were characterized as a function of magnetic field and the BECs of ^{85}Rb were created, the real fun began. One of our first measurements was to examine the collective excitation frequencies of condensates at large values of scattering length. It was predicted that the mean-field predictions for these frequencies would break down at large values of a and the measured frequencies would be different due to effects not included in the mean-field theory. We observed deviations from the mean-field predictions, but we found that the ability to tune the value of the scattering length does come with a catch. The magnetic fields with large values of a were also the fields with large loss and heating rates. As the BECs heated up, the collective excitation frequencies shifted as well, making the determination of the precise magnitude of the frequency shifts due to beyond-mean-field effects difficult. Chapter VIII details the collective excitation frequency measurements that we performed and what limits on the validity of the mean-field theory we were able to derive from them.

1.1.9 Controlled Collapse of Bose-Einstein Condensates with Attractive Interactions (Chapter IX). When the atom-atom interactions are attractive ($a < 0$), BECs are predicted to become unstable. The onset of instability (the BEC “collapse”) is also predicted to depend on the product of the number of atoms in the condensate and the magnitude of a , so by measuring the critical number of atoms above which the BECs became unstable as a function of field we were able to test the predicted stability criterion. We found that the onset of instability did indeed scale with the product of $|a|$ and the number of atoms in the BEC, but that the instability was occurring at a slightly smaller critical number than that predicted by theory. The details of this measurement are contained in Chapter IX.

1.1.10 Dynamics of Collapsing Condensates and Future Experiments (Chapter X). The last chapter (Chapter X) looks forward to future experiments with ^{85}Rb condensates. We have already performed measurements of the dynamics of BECs suddenly made unstable by a rapid change in the scattering length to an attractive ($a < 0$) value. We observed many interesting features,

such as an explosion of atoms from the collapsing BEC and an unexplained rapid loss of atoms from the BEC. This work on the dynamics of the collapse is still being developed, but the main features are discussed in Chapter X. Also included in Chapter X are possible future directions of research using ^{85}Rb BECs. There are many possible studies waiting to be undertaken in ^{85}Rb , and this thesis should serve as the opening chapter in an exciting set of experiments.

1.1.11 Notation. Most often in this thesis, Bose-Einstein condensates will simply be referred to as “condensates.” The gas of atoms trapped in the magnetic confining potential is usually referred to as a “cloud of atoms” or an “atom cloud,” but due to habit it may be referred to as the “gas of atoms” or the “trapped atoms” or even a “thermal cloud” if it does not contain a condensate. Quantities that are vectors are represented by bold characters in the main text. When vector quantities occur in an equation, an arrow above the vector quantity is used to indicate that it is a vector. Any value of magnetic field is quoted in gauss (G). Scattering lengths and other quantities associated with atom-atom collisions are often expressed in atomic units. For reference, the atomic unit of length is the Bohr (a_0), which is equal to 5.292×10^{-11} m. The s-wave scattering length is often expressed as “ a ” and the value of the magnetic field is often expressed as “ B .”

1.1.12 Prior Publication of Work in this Thesis. Several of the results that are described in this thesis have been published in various technical journals. The work in Chapter V appears in Refs. [13,14]; Chapter VI is in Ref. [15]; Chapter VII in [16]; Chapter IX in [17]; and parts of Chapter X in Ref. [18].

Chapter II

Scattering Theory and Feshbach Resonances

2.1 Introduction.

In the studies of the physics of Bose-Einstein Condensates (BECs), the collisions between atoms play a crucial role. Collisions are necessary for the creation of BECs through evaporative cooling, determine the lifetime of trapped cold clouds of atoms, and have a profound effect on most of the properties of the BECs themselves. Indeed, Bose-Einstein condensation in ^{85}Rb is interesting primarily because in ^{85}Rb it is possible to change the atom-atom collision properties in a controlled fashion. Because of the low temperatures of the atom clouds in magnetic traps, the description of collisions is simplified. In fact, a single parameter and a few phenomenological rates completely describe all of the relevant collision rates in a cold enough sample. It is the purpose of this chapter to sketch the relationship between the atom-atom interaction potential and these collision parameters.

In the trapped atom clouds that we study, the trapped clouds of atoms are dilute enough that binary collisions dominate (with one exception --- the three-body recombination rate can be significant). The collision between two atoms at these cold temperatures is described by non-relativistic quantum scattering theory. A full calculation of the scattering rate between two atoms with a sophisticated model of the interatomic potential is beyond the scope in this thesis. Instead, a “toy model” interatomic potential that has a much simpler structure is used to illustrate the basic physics of the real potential. After outlining the basics of non-relativistic quantum scattering theory, a square-well potential is used to illustrate the relationship between scattering and bound states. Next, a twist is added to the square-well potential to demonstrate the phenomenon of a Feshbach resonance [19,20,21], the ability to change the scattering rate between two atoms by changing an external field. Finally, the “toy model” is related to the actual Rb-Rb interatomic potential. Along

the way, the various quantities needed to describe collisions between cold atoms will be defined: the s-wave scattering length, and the phenomenological two- and three-body inelastic collision rates.

2.2 Basic Scattering Definitions.

The basic idea of scattering is illustrated in Figure 2.1. A particle, starting from far away, approaches a finite-ranged potential. After interacting with the potential, the particle can be deflected from its original trajectory. In classical mechanics [22], the trajectory of a particle can be described exactly as long as its initial conditions (position and velocity) are known. Often, however, the initial conditions are not known exactly. For instance, consider a monochromatic beam of particles incident on a target. The initial velocity of the particles would be known since the beam is monochromatic, but the impact parameters of the particles in the beam would be spread over a range if the beam possessed a transverse extent in space. When the initial conditions of each individual particle is not known or a range of initial conditions must be considered, the quantity of interest is the differential scattering cross section, $d\sigma/d\Omega$.

The differential scattering cross section is defined by the flux of particles reaching a detector situated both far from the scattering potential and in a direction $d\Omega$ divided by the flux of particles in the incident beam. In other words, $d\sigma/d\Omega$ is just the probability that a particle chosen at random from the initial beam will be one which scatters into direction $d\Omega$. The total cross section (σ) is determined by integrating $d\sigma/d\Omega$ over all 4π steradians and it is equal to the probability that a particle in the beam will be scattered by the potential. Classically, $d\sigma/d\Omega$ is determined by solving the equations of motion for all possible values of the impact parameter. The impact parameters that result in a particle scattering in the direction around $d\Omega$ are then known and the flux in direction $d\Omega$ can be determined by appropriately averaging over the impact parameter and initial velocity distribution of the incoming beam.

So far, only scattering of a particle off a potential, and not a particle off another particle, has been considered. It is possible to reduce the particle-particle scattering problem to a particle-potential problem by switching to the center-of-mass frame of the colliding particle pair and solving for the motion of the relative coordinates. This straightforward change is covered in most scattering texts [22], and I will not review it here.




Figure 2.1. This figure illustrates the elements involved in a scattering calculation. A beam of particles is incident from the left and then strikes the scattering potential, which is shown as the gray sphere. The potential deflects the direction of the particles in the incident beam, and a certain fraction of the particles will be scattered in the direction $d\Omega$, where they will strike a detector (the rectangle). In scattering theory, the precise motion of all the particles over their entire trajectory (including their motion in the region of the scattering potential) is not the end result of the calculation. Instead, the goal of a scattering calculation is to compute the quantity $d\sigma/d\Omega$, the fraction of atoms in the initial beam that are scattered into direction $d\Omega$.

In quantum mechanical scattering [23], the exact trajectory of a particle cannot even in principle be determined exactly since the impact parameter cannot be known exactly. In the quantum case, $d\sigma/d\Omega$ is the most that can be possibly determined, and so $d\sigma/d\Omega$ is the goal of the scattering calculations. By solving time-independent Schrödinger equation (TISE) with the appropriate boundary conditions, $d\sigma/d\Omega$ can be calculated for a given potential.

2.3 The One-dimensional Scattering Problem---Using the Time-Independent Schrödinger Equation to Determine the Scattering Amplitude.

2.3.1 Motivation. In learning about scattering theory as part of the understanding of the physics involved in our work with ^{85}Rb , I found many concepts to be non-intuitive and difficult to understand. The goal of this and the following sections is to build up a picture of how calculations of the phase-shift of time-independent wavefunctions are related to two particles scattering off one another. This picture is built up by considering simple systems first, and then adding complexity in each step. It is not meant to be a full or rigorous treatment, but rather one that demonstrates the relevant physics at each step and then makes analogies to more complicated situations. This section is intended for those who are not familiar with non-relativistic scattering theory, and the reader who is already familiar with scattering theory can skip to the last section for a discussion of the scattering properties of ^{85}Rb . The examples presented in sections 2.3 through 2.5 can be found in most textbooks that treat non-relativistic quantum mechanical scattering, but they are repeated here to establish a basis for discussing the physics of Feshbach resonances.

2.3.2 Gaussian Wavepackets. To begin the discussion of scattering, consider a particle moving in free space. For simplicity, only motion in one-dimension will be considered in this section. In quantum mechanics, a particle is described by a wavefunction. A reasonable choice for a particle wavefunction is a gaussian wavepacket peaked around a particular momentum. A wavefunction that describes such a wavepacket with width d , initial central position x_0 , and average momentum $p_0 = \hbar k_0$ at time $t=0$ is given by

$$\Psi(x, t = 0) = \frac{1}{\sqrt{d\sqrt{\delta}}} \exp\left(ik_0x - \frac{(x - x_0)^2}{2d^2}\right) \quad (2.1)$$

In order to examine the time evolution of this wavefunction, there are two choices. One is to simply solve the time-dependent Schrödinger equation with this wavefunction as an initial condition. The other method (which is of course completely equivalent) is to project the initial wavefunction onto the eigenstates of the TISE. Once that is done, a factor of $\exp(-iEt/\hbar)$, where E is the energy of the eigenstate, is associated with each eigenstate to trace the time evolution.

In free space, the one-dimensional Hamiltonian (H) consists only of the kinetic energy term $\hbar^2 k^2/2m$. Solving the TISE

$$H\Psi = -\frac{\hbar^2}{2m} \frac{\partial^2}{\partial x^2} \Psi = E\Psi \quad (2.2)$$

yields the eigenstates (each with eigenvalue k)

$$\phi_k(x) = \frac{1}{\sqrt{2\pi}} \exp(ikx) \quad (2.3)$$

Projecting the initial wavefunction onto these eigenstates and including the time evolution factor gives

$$\Psi(x, t) = \int dk \int dx \phi_k^*(x) \Psi(x, 0) \exp\left(\frac{-i\hbar k^2 t}{2m}\right) \quad (2.4)$$

The integral over dx is the one that determines the projection of the initial wavefunction onto the k-th eigenstate, and the integral over dk sums over the complete set of eigenstates. Performing the integral in equation (2.4), and then squaring the absolute value of the wavefunction to get the probability density of the particle results in

$$|\Psi(x, t)|^2 = \frac{2 d m}{\sqrt{\pi} \sqrt{4d^4 m^2 + \hbar^2 t^2}} \exp\left(\frac{-(x - x_0 - \frac{\hbar k_0 t}{m})^2}{d^2 + (\frac{\hbar t}{2 m d})^2}\right) \quad (2.5)$$

As time moves forward, the wavepacket moves to the right from position x_0 with a velocity $v = \hbar k_0/m$.

The wavepacket also spreads out due to the position/momentum uncertainty relation.

2.3.3 Scattering from a Step Function. A simple scattering process is the scattering of a particle off a step function in one-dimension [24]. Figure 2.2 shows the structure of the step-function potential. A wavepacket incident from the left will strike the potential, and part of the

wavepacket will continue to the right and part will be reflected back to the left. The calculation of the amplitudes of the reflected and transmitted wavepackets is analogous to finding $d\sigma/d\Omega$ for this simple scattering event. Throughout this discussion it is assumed that the incoming wavepacket is a gaussian sharply peaked around k_0 .

Again, it is possible to evaluate the time-dependent Schrödinger equation directly to find the time evolution of the wavepacket as it scatters off the step potential. Instead, the same strategy used in the previous section in equations (2.2)-(2.4) is used again. The initial wavefunction is projected onto the eigenstates of the TISE and then a phase factor $\exp(-iEt/\hbar)$ is associated with each eigenstate. The only difference is that the step potential has a different set of eigenstates than the free-particle case.

The eigenstates of the step-function Hamiltonian are simple to calculate, but some care must be taken in setting the boundary conditions (i.e. the behavior of the wavefunction far from the scattering potential). In fact, the appropriate choice of boundary conditions is very important in evaluating all scattering problems. Left of the step function ($x < 0$), it is possible to have wavepackets moving both to the right (the incident packet) and to the left (the reflected packet). Right of the step function ($x > 0$), it is only possible to have wavepackets moving to the right (the transmitted packet). There is no source for any wavepackets coming far from the right and striking the potential, and that fact is included in the boundary conditions.

The solutions of the TISE with the appropriate boundary conditions are

$$\varphi_k(x) = A \exp(ikx) + B \exp(-ikx), \quad k = \sqrt{\frac{2mE}{\hbar^2}}, \quad x < 0 \quad (2.6)$$

and

$$\varphi_k(x) = C \exp(ik'x), \quad k' = \sqrt{\frac{2m(E - V_0)}{\hbar^2}} = \sqrt{k^2 - k_c^2}, \quad x > 0 \quad (2.7)$$

$$B = \frac{(k - k')}{(k + k')} A, \quad C = \frac{2k}{(k + k')} A \quad (2.8)$$




Figure 2.2. This figure illustrates the scattering of a wavepacket from a step potential. The incident wavepacket is shown at time $t=0$ and the reflected and transmitted wavepackets are shown after the incident wavepacket struck the step potential. The arrows indicate the direction of propagation of each of the wavepackets, where $k>0$. Implicit in these definitions is that the incident wavepacket is sufficiently peaked around k_0 so that there are no significant components at $k < (2mV_0/\hbar^2)^{1/2}$. The term whose coefficient is A corresponds to a wave moving to the right with $x < 0$, and so it corresponds to the incident wavepacket. Likewise, the terms whose coefficients are B and C correspond to the reflected and transmitted wavepacket, respectively. The values of A , B , and C are fixed for a given value of k by matching the eigenfunctions and their first derivatives at $x=0$:

Just as was done in equations (2.2)-(2.5), the initial (incident) gaussian wavepacket is projected onto these eigenfunctions. By adding the phase factor $\exp(-iEt/\hbar)$ to each eigenstate, summing over the eigenstates results in the following evolution of the wavefunction

$$\begin{aligned} \Psi(x, t) &= N_1 A \exp\left(\frac{-(x - x_0 - \frac{\hbar k_0 t}{m})^2}{2(d^2 + \frac{\hbar^2 t^2}{d^2 m})}\right) + N_2 B \exp\left(\frac{-(x + x_0 + \frac{\hbar k_0 t}{m})^2}{2(d^2 + \frac{\hbar^2 t^2}{d^2 m})}\right), \quad x < 0 \\ &= N_3 C \exp\left(\frac{-(x - \frac{q}{k_0} x_0 - \frac{\hbar q t}{m})^2}{2(d^2 + \frac{\hbar^2 t^2}{d^2 m}) \frac{q^2}{k_0^2}}\right), \quad x > 0 \end{aligned} \quad (2.9)$$

where $q=(k_0^2-k_c^2)^{1/2}$ is the central momentum of the wavepacket when $x>0$, and N_1, N_2, N_3 contain phase factors ($|N_1|=|N_2|=|N_3|$) important only when there is significant overlap between the various components in the wavefunction (i.e. when the wavepacket is near the origin). B and C are evaluated in equation (2.8) at $k=k_0$. Equation (2.9) is not exact, but was derived in the limit that the initial wavepacket is sharply peaked around k_0 and x_0 goes to ∞ .

The three terms in equation (2.9) describe all of the components involved in the scattering event. The term associated with A is the incident wavepacket and it begins at x_0 at $t=0$ and then moves to the right. The term associated with B is the reflected wavepacket and begins at $-x_0$ and moves to the left. The last term associated with C is the transmitted wavepacket that begins at $x_0 * k_c/k_0$ and moves to the right. Since x_0 is large and negative, when $t=0$ the exponential factors associated with B and C are nearly zero (note that Ψ has different definitions depending on whether x is positive or negative). The initial wavefunction is simply that of a gaussian wavepacket moving to the right at velocity $v=\hbar k_0/m$. At very large times after the scattering from the potential has occurred ($t \gg x_0 m/\hbar k_0$), the exponential factor associated with A is nearly zero and the factors associated with B and C are nonzero. The wavefunction then describes the motion of the transmitted and reflected wavepackets after the scattering event has occurred. To compute the transmission and reflection probability, it is only necessary to integrate $|\Psi|^2$ at large times for $x<0$ and $x>0$, and compare that to the integration of $|\Psi|^2$ at $t=0$. Performing the integration shows that the probability

of transmission is $(q/k_0) \cdot (|C|^2/|A|^2)$ and the probability of reflection is $(|B|^2/|A|^2)$, where again equation (2.8) for B and C is calculated using $k=k_0$.

So, in the end the transmission and reflection probabilities only depend on the coefficients of the TISE eigenstates of energy $\hbar^2 k_0^2/2m$ when it is solved with the appropriate boundary conditions. It was not necessary to actually determine the full time evolution of the incident wavepacket. When an incident wavepacket is projected onto the eigenstates of the TISE, the overlap with any particular eigenstate is determined by the part of the eigenstate wavefunction that corresponds to the boundary condition for an incoming wavefunction. The amplitude of the boundary conditions corresponding to the outgoing wavefunctions then controls the future behavior of the wavepacket. While the computation of the transmission and reflection coefficients have only been derived in this particular example, this result is quite general. The scattering process, which is certainly a dynamic one, is well described in a time-independent formalism.

2.4 Scattering in Three Dimensions.

2.4.1 Comparison to One-dimensional Case. While the three-dimensional case is naturally more complicated than the one-dimensional case, the method of solving the scattering problem is still the same. Ideally, the TISE is solved with the appropriate boundary conditions and the resulting amplitudes associated with the parts of the eigenstate corresponding to the incoming and outgoing waves are used to determine the scattering probability $d\sigma/d\Omega$. In practice, there are two complications. First, the boundary conditions in three dimensions are more complicated than those in one dimension. Second, it is rarely possible to completely solve the TISE exactly, even for the simplest potentials. Therefore, approximation methods are employed. At low collision energies, appropriate for the work presented in this thesis, the scattering problem can be simplified significantly through the use of partial waves --- projecting the incoming and outgoing wavefunctions onto an angular momentum basis.

2.4.2 Boundary Conditions in Three Dimensions. In three dimensions, far from the potential, the TISE is solved with the following boundary condition:

$$\Psi(\vec{r}) = \exp(ikz) + f(\mathbf{q}, \mathbf{f}) \frac{\exp(ikr)}{r}, \quad r \rightarrow \infty \quad (2.10)$$

This wavefunction satisfies the free space TISE in three dimensions. The first term represents the incident particle and is simply a plane wave. As discussed in the previous section, an incident wavepacket can be built up in terms of incoming plane waves, and so the first term in equation (2.10) is associated with the incident wavepacket. The second term represents the outgoing wave resulting from the scattering off of the potential at the origin. The variables θ and ϕ represent the polar and azimuthal angle of \mathbf{r} with respect to the origin. The function $f(\theta,\phi)$ represents the scattering amplitude. Evaluating the probability density $|\Psi|^2$ of particles scattered in the $d\Omega$ (i.e. θ,ϕ) direction shows that $d\sigma/d\Omega$ is equal to $|f(\theta,\phi)|^2$. Therefore, solving for the scattering amplitude $f(\theta,\phi)$ is equivalent to solving the scattering problem. The wavefunction (2.10) is in general incomplete since it does not consider other internal quantum numbers, such as spin. The inclusion of other states can have significant effects on the scattering properties of a potential, as will be demonstrated in section 2.6. For now, structureless particles will be considered.

For the rest of this section, elastic scattering is considered in detail. When elastic scattering occurs, the internal states of the colliding pair are unchanged. Energy conservation then requires that the energy of the pair before and after the collision must be the same. Collisions that change the internal state of one or both of the colliding pair are called inelastic collisions. Inelastic collisions will be discussed less extensively and with less detail below.

2.4.3 Partial Waves. The plane-wave basis represented in equation (2.10) is not the only basis that can be used to describe the scattering wavefunction far from the scattering potential. For cold collisions, it is useful to switch to a partial wave basis [23]. In the partial wave basis, the implicit (k_x, k_y, k_z) plane-wave quantum numbers are replaced by the quantum numbers (E, L, M) where E is the energy, L is the angular momentum, and M is the projection of the angular momentum onto the z -axis. Both of these bases are complete.

The reason that it is advantageous to switch to the partial wave basis is that for sufficiently cold collisions, only a few (or even only one) partial waves contribute to $d\sigma/d\Omega$. To see why this is true, it is helpful to write the TISE radial equation for a spherically symmetric potential $V(r)$ [25]:

$$\left\{ -\frac{\hbar^2}{2\mathbf{m}} \frac{d^2}{dr^2} + \frac{L(L+1)\hbar^2}{2mr^2} + V(r) \right\} u_{E,L}(r) = E u_{E,L}(r) \quad (2.11)$$

where μ is the reduced mass, L is the angular momentum quantum number and $\Psi = u_{E,L}(r)Y_m^L(\theta, \phi)/r$. The second term in equation (2.11) represents an angular momentum barrier. For a finite-ranged potential $V(r)$, this angular momentum barrier will prevent high angular momentum (high L) partial waves from penetrating into the potential region. These partial waves cannot then contribute to the scattering amplitude.

For cold enough collisions, only the $L=0$, or s-wave, partial wave will contribute to the scattering cross section. One way to demonstrate this fact is to consider the classical value of the closest approach of a particle to the center of the potential, the impact parameter. For a typical case for ^{85}Rb , a collision energy of $100\mu\text{K}$ between particles with $L=2\hbar$, (d-wave partial wave, the second lowest angular momentum partial wave possible for two indistinguishable atoms) will have a closest approach x_{\min} , calculated using $L=\mu v x_{\min}$. For this case $x_{\min}=285 a_0$ (atomic units), much greater than the range of the Rb-Rb interatomic potential (see the last section in this chapter). So for cold temperatures, only s-wave scattering will contribute significantly to $d\sigma/d\Omega$.

To relate the partial wave basis to the scattering boundary conditions represented in equation (2.10), it is only necessary to expand the plane wave and outgoing spherical scattering wave in the (E,L,M) basis. I will not cover this transformation here. For a succinct description see Ref. [26]. The end result is that, for a spherically symmetric potential (and therefore no dependence on ϕ), the scattering amplitude $f(\theta)$ is

$$f(\mathbf{q}) = \frac{1}{k} \sum_{L=0}^{L=\infty} (2L+1) \exp(i\mathbf{d}_L) \sin(\mathbf{d}_L) P_L(\cos(\mathbf{q})) \quad (2.12)$$

The scattering amplitude is then described by a series of partial wave phase shifts δ_L . Probability conservation requires that all of the phase shifts be real quantities [26].

To illustrate how a phase shift can be related to a scattering amplitude, consider s-wave scattering off of a spherically symmetric potential. There are two possible solutions to the equation for $u_{E,0}(r)$ (2.11) with $L=0$ far from the scattering potential $V(r)$:

$$\begin{matrix} \sin(kr) \\ \cos(kr) \end{matrix}, \quad k = \sqrt{\frac{2mE}{\hbar^2}} \quad (2.13)$$

Since the radial function of Ψ is $u_{E,0}(r)/r$, only $\sin(kr)$ is regular at the origin. Therefore, in the absence of a scattering potential, $u_{E,0}(r)$ must be equal to $\sin(kr)$. In the presence of a scattering potential, the wavefunction must still be regular at the origin. However, far from the origin, the free-space wavefunction must be matched to the wavefunction that evolved in the presence of the scattering potential. Since the free-space wavefunction is not being matched at the origin, but rather at a distances beyond the range of the potential, it is no longer required that there be no $\cos(kr)$ component. In general, the wavefunction at large distances is $c_1\sin(kr)+c_2\cos(kr) = c(\sin(kr+\delta_0)) = c(\cos(\delta_0)\sin(kr)+\sin(\delta_0)\cos(kr))$. The phase shift δ_0 can be thought of as measuring the ratio of the incoming wavefunction to the scattered component of the wavefunction. If $V(r)=0$, then $\delta_0=0$ and only the incoming wave is present at large distances. As δ_0 increases, the wavefunction no longer matches only the incoming wave, indicating that scattering is occurring. When $\delta_0 = \pi/2$, the ratio between the scattered wave and the incoming wave must be a maximum.

2.4.4 Phase Shift and Cross Section. Once all of the partial wave phase shifts are known, the differential and total cross sections can be computed using (2.12). If only s-wave scattering contributes, then the differential and total cross sections are

$$\begin{aligned}\frac{d\mathbf{s}}{d\Omega} &= \frac{1}{k^2} \sin^2(\mathbf{d}_0) \\ \mathbf{s} &= \frac{4\mathbf{p}}{k^2} \sin^2(\mathbf{d}_0)\end{aligned}\tag{2.14}$$

In other words, s-wave scattering is isotropic and $d\sigma/d\Omega$ does not depend on θ . The dependence of the cross section on $\sin(\delta_0)$ can be understood in light of the discussion in the previous section---a factor of $\sin(\delta_0)$ is associated with the irregular ($\cos(kr)$) s-wave basis function. Note that equation (2.14) only applies to collisions between particles that can be distinguished from one another. See subsection 2.9.6 and Chapter V for the modifications that must be made for identical particle scattering.

The restriction that δ_0 be a real quantity means that the maximum cross section is limited by the wavevector k of the incoming wavefunction. This limit is known as the unitarity limit, and it can be a significant limit on the scattering rate even at low collision energies.

One other quantity of interest is the s-wave scattering length (a). It is defined as

$$\frac{-1}{a} = k \cot(\mathbf{d}_0), \quad \text{in the limit } k \rightarrow 0 \quad (2.15)$$

Plugging the definition (2.15) into the equation (2.14) and taking the limit that k goes to zero means that $d\sigma/d\Omega = a^2$ and $\sigma = 4\pi a^2$. In elastic scattering, δ_0 tends to zero as k goes to zero such that the s-wave scattering length is a finite constant. The significance of the s-wave scattering length will be discussed below when the square well potential is treated.

2.5 The Born Approximation.

2.5.1 First Born Approximation. An alternative approximation to the use of the partial wave expansion and the inclusion of only a few (or one) or the partial waves is the Born Approximation. In contrast to the partial wave expansion, this approximation tends to work much better as the collision energy increases and so it would seem to have limited utility in understanding cold collisions. It turns out, however, that solving the scattering problem within the Born Approximation is necessary to compute the interaction energy due to the interatomic potential in a BEC, and so I will present it briefly here.

The starting point, again, is the TISE. In this case, the free-space Hamiltonian ($H_0 = p^2/(2m)$) is explicitly separated from the scattering potential V . Ψ represents the full solution to the TISE and Φ represents the free space eigenstate, both with the proper scattering boundary conditions and.

$$(H_0 + V)\Psi = E\Psi \quad (2.16)$$

$$H_0\Phi = E\Phi \quad (2.17)$$

Equation (2.16) can be rewritten as [23,26]

$$\Psi = \Phi + \frac{V}{E - H_0 + i\epsilon} \Psi \quad (2.18)$$

where the small parameter ϵ will eventually be taken to the limit in which it is zero. It is put in only to handle the divergence in the denominator. Projecting equation (2.18) explicitly onto the position basis, applying the scattering boundary conditions (2.10), and then evaluating the resulting

expression far from the scattering potential results in the following integral equation for $\Psi(\mathbf{x})$

$$\Psi(\vec{x}) = \exp(ikz) + \frac{\exp(ikr)}{r} \left\{ \frac{-1}{4\mathbf{p}} \frac{2m}{\hbar^2} (2\mathbf{p})^3 \int d^3x' \frac{\exp(-i\vec{k}' \cdot \vec{x}')}{(2\mathbf{p})^{3/2}} V(\vec{x}') \Psi(\vec{x}') \right\}. \quad (2.19)$$

Equation (2.19) cannot be evaluated directly since $\Psi(\mathbf{x})$ appears under the integral sign on the right side of the equation. However, it can be evaluated approximately by assuming that the incoming wavefunction does not change significantly in the region of the scattering potential (i.e. the scattering potential is weak enough by some measure). In that case, Ψ under the integral sign can be replaced by the incoming wavefunction $\exp(ikz)$. Performing this replacement and using the definition of $f(\theta, \phi)$ in equation (2.10) results in the scattering amplitude

$$f^{(1)}(\mathbf{q}, \mathbf{f}) = \frac{-1}{4\mathbf{p}} \frac{2m}{\hbar^2} \int d^3x' \exp(i(\vec{k} - \vec{k}') \cdot \vec{x}') V(\vec{x}'), \quad (2.20)$$

where $\vec{k} = k\hat{z}$, $\vec{k}' = k\{\sin(\mathbf{q}) \cos(\mathbf{f})\hat{x} + \sin(\mathbf{q}) \sin(\mathbf{f})\hat{y} + \cos(\mathbf{q})\hat{z}\}$

This calculation of the scattering amplitude is commonly referred to as the first Born Approximation. It is possible to calculate the scattering amplitude beyond the first order by computing the variation of the incoming wavefunction in the region of the scattering potential to additional orders [26].

2.5.2 Range of Validity of the First Born Approximation. There are two ranges where the first Born Approximation describes the scattering amplitude well. One is high-energy scattering, where the energy of the incoming particle is much greater than the energy scale of the scattering potential. This consideration is not too relevant to very cold collisions. The first Born Approximation also works well at low collision energies as long as the scattering potential is weak enough. In that case, the second term in equation (2.19) is much smaller than the first and the first Born Approximation is a good one. The condition that defines a weak potential depends on the potential, but for an attractive potential it is the roughly the same as the condition that no bound state forms in the potential [23,26]. Therefore, a weak attractive potential is one that is not close to developing a bound state.

2.6 Scattering from a Square Well Potential.

2.6.1 Motivation. Now that the methods used to calculate $d\sigma/d\Omega$ have been outlined, it is useful to consider a potential for which analytical solutions are easy to derive in the low-collision-energy limit. One such potential is a spherically symmetric attractive square well potential, where the scattering potential is equal to $-V_0$ if $r < R$ and is zero otherwise. Only s-wave scattering will be considered in this section.

Not only will the solution of the square well scattering problem serve as an example for the method of solving the scattering problem outlined above, it will illustrate several relevant aspects of the physics of collisions between Rb-Rb atoms: the relationship between the s-wave scattering length and the position of the last bound state can be illustrated; the connection between the s-wave scattering length and the scattering wavefunctions can be discussed explicitly; and the minimum in cross section that can occur at finite collision energies can be demonstrated.

2.6.2 Solution of the Time-Independent Schrödinger Equation. The eigenfunctions of the TISE are simple to determine for the spherically symmetric attractive square well potential. Since the only scattering being considered is s-wave, $L=0$ and there is no dependence of the eigenfunctions on the angular variables (θ, ϕ) . The two possible solutions to the radial equation (2.11) are given in equation (2.13). In the region, $r < R$, only the regular solution is possible so that the wavefunction is finite at the origin. For $r > R$, either solution is possible. For a collision energy E ,

$$u_E(r) = \begin{cases} \sin(k'r) & , \quad k' = \sqrt{\frac{2\mathbf{m}(E + V_0)}{\hbar^2}}, \quad r < R \\ c_1 \sin(kr) + c_2 \cos(kr) = c \sin(kr + \mathbf{d}_0), & k = \sqrt{\frac{2\mathbf{m}E}{\hbar^2}}, \quad r > R \end{cases} \quad (2.21)$$

The constants c_1 and c_2 , or equivalently c and δ_0 , are determined by matching the value of $u_E(r)$ and its first derivative at $r=R$. The s-wave phase shift (δ_0) is then determined by a transcendental equation

$$k \cot(kR + \mathbf{d}_0) = k' \cot(k'R) \quad (2.22)$$

Once δ_0 is known, the s-wave scattering length (see equation (2.15)) and hence $d\sigma/d\Omega$ and σ are also known.

It is also possible to calculate the bound states with zero angular momentum of the square well potential as well. The only change is in the boundary conditions, where the eigenfunctions must go to zero as r goes to infinity and also be finite at the origin. The eigenfunctions are then

$$u_E(r) = \begin{cases} \sin(k'r), & k' = \sqrt{\frac{2m(E+V_0)}{\hbar^2}}, & r < R \\ c \exp(-kr) & , & k = \sqrt{\frac{2m|E|}{\hbar^2}}, & r > R \end{cases} \quad (2.23)$$

Since there is only one adjustable parameter, c , and there are two conditions that must be satisfied to match the wavefunctions at R , bound states will only exist for certain values of the energy E .

2.6.3 Scattering Length and the Position of the Last Bound State. There is a definite relationship between the position of the last bound state in the potential and the scattering length a . Figure 2.3 shows the scattering length of the square well as the depth of the well is increased. Initially, the scattering length decreases as the potential gets deeper. As the potential begins to get deep enough to form one bound state, the scattering length decreases more rapidly, until it diverges once the potential forms a bound state. As the potential continues to get deeper, the scattering length becomes positive, but then decreases until another bound state forms in the potential.

This behavior of the scattering length with respect to the last bound state in the potential is quite general. The reason that this is so can be seen qualitatively by examining the structure of the scattering wavefunction with E close to zero. Figure 2.4 shows scattering wavefunctions of nearly zero energy. Without the presence of a scattering potential, the wavefunction rises nearly linearly from a value of zero at the origin. When the potential is present, the wavefunction rises linearly far from the potential, but it no longer appears to intercept the r -axis at the origin and instead intersects at some finite value of r . This value is the s-wave scattering length, and so the s-wave scattering length represents the last zero in the scattering wavefunction. When a bound state exists with nearly zero energy, the wavefunction of the bound state outside the potential is nearly flat. Since inside the range of the potential the wavefunction of the bound state and a near zero energy scattering state are pretty much the same, the value and derivative of the two wavefunctions at R must also be nearly the

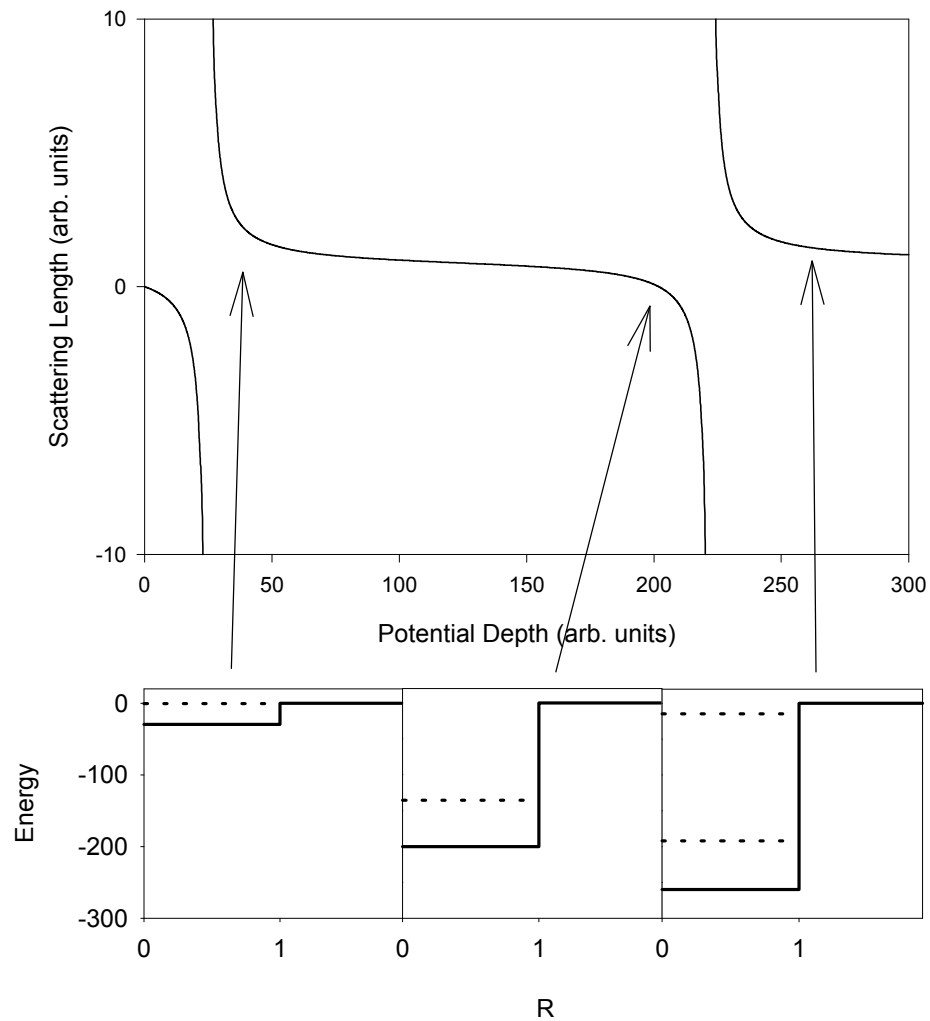


Figure 2.3. The upper figure shows the functional form of the scattering length as a function of the depth of the attractive square well potential (V_0). Each divergence in the scattering length is associated with the formation of an $L=0$ bound state in the potential. The lower figure shows the square well potential at several different depths (solid line). The dotted line shows the energies of the $L=0$ bound states in the potential. From left to right, the scattering lengths shown in the lower figure are 4.55, 0.09, and 1.48 in the arbitrary units shown in the upper plot.

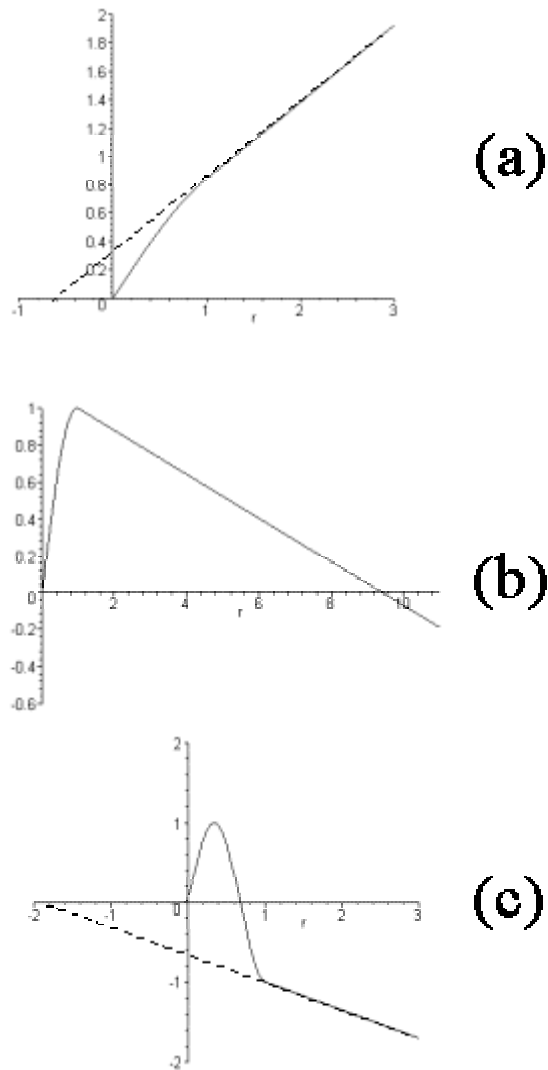


Figure 2.4. This figure shows attractive square well scattering wavefunctions in the limit of zero collision energy for successively deeper depths of the square well. Actually, the wavefunction itself is not plotted; rather the function $u_E(r)$ that is related to the wavefunction as defined in the main text is what is shown. The range of the potential is set to be equal to 1 (note that the scale for r varies from plot to plot in the figure). The potential in (a) has no bound state and a scattering length of -0.56 . The potential in (b) has one bound state near zero energy and a positive scattering length of 9.4 . The potential in (c) has one bound state and a negative scattering length -1.84 . The position at which the wavefunction outside the potential ($r > 1$) would intercept the axis is equal to the scattering length. In this sense, the scattering length is equal to the position of the last zero in the wavefunction in the limit of zero collision energy. The solid line shows the wavefunction, while the dashed line is used to show where the wavefunction at ($r > 1$) would intercept the axis in the negative scattering length cases.

same. The scattering wavefunction must also be flat outside R and therefore the scattering length must be large. This similarity between bound state and scattering wavefunctions at R allows the estimation of the depth of the highest bound state from the (positive) scattering length if $a \gg R$: the binding energy is approximately $\hbar^2/2\mu a^2$.

Even though the range of the potential, R , and the s-wave scattering length both have dimensions of length, they can be orders of magnitude different in size. Since the scattering properties are completely determined by the scattering length, two different potentials with the same scattering length will produce exactly the same scattering in the s-wave scattering limit. In the case of hard-sphere scattering, (V_0 is infinite at $r < R$, zero elsewhere), the last zero of the bound state must occur at R , and so the scattering length is equal to the radius of the hard sphere. A potential with a scattering length a behaves exactly the same as a hard-sphere of radius a . Therefore, the scattering length sets the effective length scale of the scattering potential.

2.6.4 Zero collision cross section at finite collision energy. Figure 2.5 shows the total cross section (σ) as a function of collision energy for a square well potential with a negative scattering length. At a finite collision energy the cross section goes to zero. This is because at a particular energy, the wavefunction at the boundary of the potential has exactly the same value and derivative as in the free-propagation case (i.e. in the absence of a scattering potential). The exact position of this zero depends on the details of the potential, but a similar behavior occurs in ^{85}Rb , which has important consequences for evaporative cooling (see subsection 7.1.1). This effect is sometimes called a “Ramsauer-Townsend minimum” in atomic scattering theory, especially when electron-atom scattering is being considered.

2.7 Square Well Scattering with another Internal Degree of Freedom--- “Toy Model” of a Feshbach Resonance.

2.7.1 General Description. So far, any different internal states of the colliding pair have been ignored. At first glance, the presence of any other internal states should not matter as long as elastic collisions are being considered, since in elastic scattering the initial and final internal states of the colliding pair must by definition be identical. However, internal states that are orthogonal to each other far from the scattering potential can be coupled by the scattering potential. Moreover, the

potential may not be the same for different internal states. This coupling due to the scattering potential will project an incoming wavepacket into a superposition of internal states in the region of




Figure 2.5. Demonstration of the zero in collision cross section for a finite collision energy. The attractive potential used to generate this curve had several bound states and a negative scattering length. Similar structures can be observed with potentials with a positive scattering length, but the zeros in σ tend to be at higher energies.

the scattering potential, and each state in the superposition will acquire a separate phase shift. When the wavepacket leaves the scattering region, the coupling will reproject the wavepacket onto the free-space internal state basis.

Because the potentials are different and different phase shifts are acquired, the colliding pair will in general scatter inelastically in addition to elastically. This is because the reprojected wavefunction will often project into internal states different from the initial state. However, this is not always the case. Away from the scattering potential, different internal states may have different energies, and if the collision occurs with the colliding pair initially in the lowest energy internal state and with a collision energy below all of the other internal state energies, then conservation of energy will prevent any inelastic scattering. Even in this case, the coupling to other internal states (so-called closed channels) due to the scattering potential can affect the elastic scattering, however.

In order to discuss this physics in a more concrete fashion, and to allow the analytic computation of bound states and scattering properties, the square well potential treated in the previous section is expanded by adding another internal state. When coupling between the two states is introduced, the scattering behavior changes radically. In order to allow more freedom in the selection of the scattering length in the model, an additional region of much weaker attractive potential is added outside R . This addition is justified by the real atom-atom potentials.

It is interesting to consider the case when different internal states respond to an external field in different ways. In this case, the external field can be used to change the position of the last bound state in the potential and can therefore dramatically alter the scattering properties of the potential. Such a change in scattering through the application of an external field is called a Feshbach resonance [19-21], and such a resonance in ^{85}Rb and its effect on elastic scattering, inelastic scattering, and the properties of BECs forms the subject matter of the rest of this thesis.

By using this “toy model” square well potential with two different internal states, the basic physics of the Feshbach resonance can be demonstrated. The relationship between the strength of the coupling and the “width” of the resonance can be shown, and the connection between the asymptotic bound states in each of the two scattering potentials and the position of the resonance can be shown explicitly.

2.7.2 “Toy model” Hamiltonian. The extra internal state is added to the square well scattering problem by treating the incoming particle as a two-component spinor. The two internal states then become spin “up” and spin “down.” A different depth square well is associated with each component. Only s-wave scattering is considered. The Hamiltonian for the model is

$$H = \begin{pmatrix} -\nabla_r^2 - V_0 & \Delta \\ \Delta & -\nabla_r^2 - V_0 + V_{in} + x \end{pmatrix} \Big|_{r < R} + \begin{pmatrix} -\nabla_r^2 - V_1 & 0 \\ 0 & -\nabla_r^2 - V_1 + V_{in} + x \end{pmatrix} \Big|_{R < r < R_0} + \begin{pmatrix} -\nabla_r^2 & 0 \\ 0 & -\nabla_r^2 + V_{out} + x \end{pmatrix} \Big|_{r > R_0} \quad (2.24)$$

where the Δ is the coupling term, V_0 is the scattering potential close to the origin, V_1 is the much weaker potential extended past R , V_{in} is the difference between the two spin states inside the scattering potential and V_{out} is the difference between the two spin states outside the potential. The variable x represents the change in energy due to an applied field, and only the “down” state shifts with the external field. V_{in} is defined to be positive, and the collision energy is always kept below V_{in} . The appropriate boundary conditions at large r are then

$$u_E(r) = \begin{pmatrix} \mathbf{a} \sin(kr + \mathbf{d}_0) \\ \mathbf{b} \exp(-k'r) \end{pmatrix} \quad \left. \begin{array}{l} k = \sqrt{E} \\ k' = \sqrt{V_{out} - E} \end{array} \right\} \quad r \rightarrow \infty \quad (2.25)$$

and the other boundary condition is that the wavefunction is finite at the origin.

Figure 2.6 shows the scattering length as a function of x for a selected set of parameters. Two features are of interest. First, the scattering length diverges for a particular value of the external field x . This is a Feshbach resonance. The divergence is due to the appearance of a bound state in the potential near zero energy as x is changed. Second, the scattering length is roughly equal to R away from the resonance. The reason for this is that the coupling (Δ) at R between the two different states projects the incoming “up” wavefunction into a “down” state that is not supported as a bound state by the “down” potential, and in that case the incoming and outgoing wavefunctions at R are matched only if the amplitude of the incoming wavefunction in the “up” state is nearly zero at R . The presence of a bound state near zero energy removes the difficulty in matching the incoming and outgoing wavefunctions and allows for the variation of the scattering length.

Figure 2.7 again shows the scattering length as a function of x , but in this case V_1 is not set to be equal to zero. The presence of the weak potential shifts the phase of the wavefunction and the

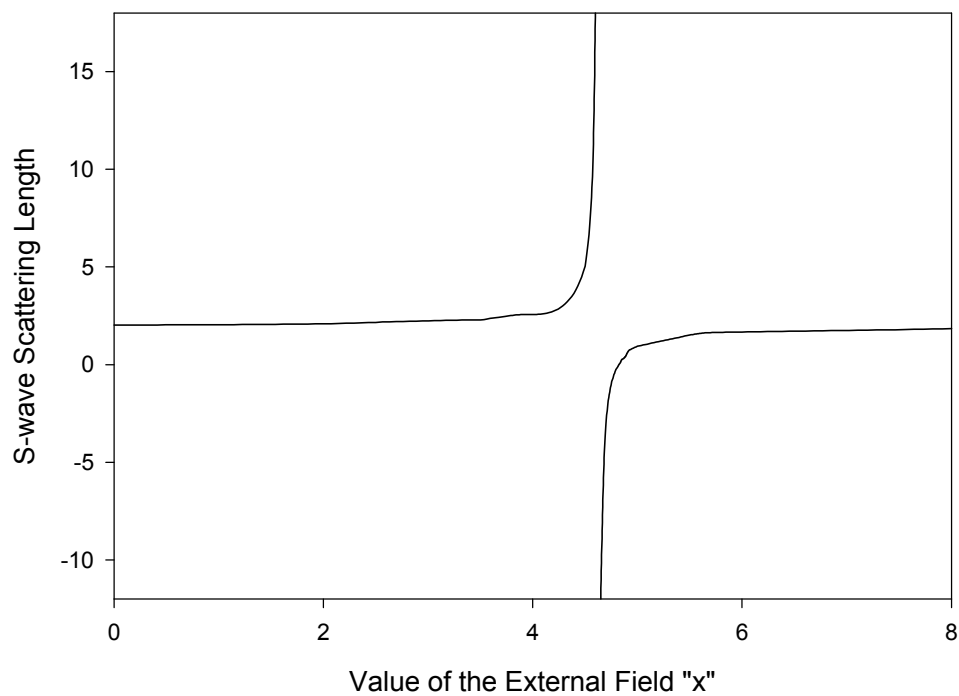


Figure 2.6. Feshbach resonance calculated using the “toy model.” In this case, the attractive tail of the potential was not included. The parameters used to calculate this curve were $V_0 = 125$, $R = 2$, $V_1=0$, $V_{in}=8$, $V_{out} = 10$, $\Delta = 10$.

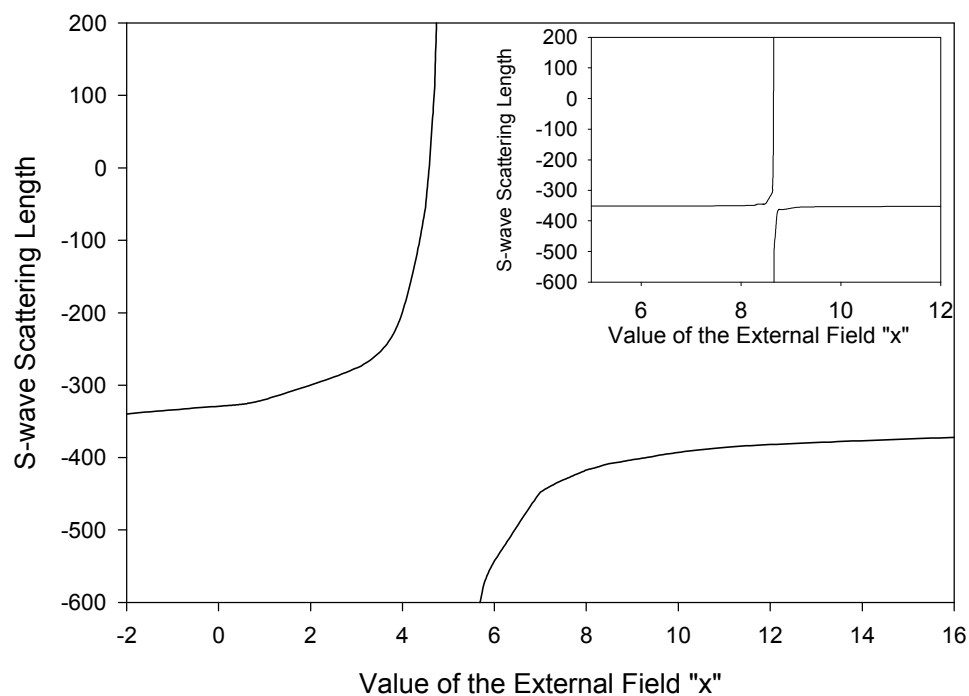


Figure 2.7. Feshbach resonance calculated in the “toy model” with the presence of the additional attractive tail on the Hamiltonian. The parameters for the calculation of the main curve in this figure were the same as in figure 2.6 except that $R_0 = 30$ and $V_1 = 0.00295$. The value of Δ for the curve in the inset was set to be 1, but otherwise the parameters were the same as those used to calculate the main curve. The fact that the resonance of the curve in the inset is narrower and the peak in the scattering length occurs at a different value of x demonstrates that both the width of the resonance and the location of the divergence depend on the value of the coupling parameter.

scattering length is no longer fixed to be roughly the same as the range of the potential. The relationship between the scattering length and the value of V_1 is not straightforward. Both the range of V_1 and its depth are important. The Feshbach resonance still occurs, with the scattering length diverging around a certain value of x . Also shown in Figure 2.7 is the dependence of the width of the Feshbach resonance on the strength of the coupling parameter Δ . The width of the Feshbach resonance is customarily defined as the difference in the value of x between the divergence of the scattering length and the value where the scattering length is equal to zero. From the calculations shown in Figure 2.7, the scaling of the width of the Feshbach resonance scales roughly as Δ^2 .

The position of the last bound state in the potential can be discerned from the scattering length, just as is the case for the simple potentials in section 2.6.3. While value of the external field is such that the scattering length is positive, there is a bound state just below zero energy, the top of the “up” potential. If the magnitude of the negative scattering length is large, then a bound state is almost able to form in the potential, but not quite.

The positions of the bound states in the asymptotic potentials (Δ set equal to zero) give information about the position of the last bound state in the coupled potential. If there is a bound state in the “down” potential near the collision energy, and its energy can be shifted relative to the collision energy via an external field, then the resonant enhancement of the scattering via a Feshbach resonance can occur. While the asymptotic potentials are useful in this way, there are important differences between the asymptotic potential and the coupled potential. First, there is a shift in the bound states of the actual potential due to the presence of the coupling Δ . Second, a true bound state is not possible if the energy of that bound state would be greater than zero, where zero energy is defined as the lowest energy of all the possible internal states far from the potential. The presence of a bound state in the asymptotic potential above zero energy is referred to as a quasi-bound state [25]. Because of the coupling between different internal states, any probability in this quasi-bound state would “leak out” through transitions to unconfined internal states. Quasi-bound states can certainly play a role in elastic and inelastic scattering, and it is possible to model them as eigenstates with a finite width (i.e. a decay of probability amplitude).

In summary, a Feshbach resonance occurs when three conditions are satisfied: there is a bound state in a closed channel that is near to the scattering energy; the energy of the states in the closed channel can be shifted by an external field in a way that is different from the incoming collision channel state; and there is some coupling between collision channels due to the potential.

2.8 Inelastic Scattering.

As mentioned in section 2.7.1, inelastic scattering occurs when the colliding pair changes its internal state as a result of the collision. Typically, this change in internal state results in a change of kinetic energy. For very cold collisions, the collision energy is usually smaller than any internal state energy splitting and so only inelastic collisions that increase kinetic energy are possible.

The calculation of inelastic scattering cross sections is similar to calculating the elastic cross section [25]. The boundary conditions need to be altered to include outgoing scattered waves (angular momentum projections of $\exp(ikr)/r$) in different internal states, and then a system of coupled differential equations in a form similar to equation (2.11) is solved to determine the amplitudes associated with each boundary condition in the TISE. These amplitudes can then be related to the differential cross section $d\sigma/d\Omega$ for all of the elastic and inelastic processes.

Unlike elastic scattering, where the cross section tends to a constant at low scattering energies (see the discussion after equation (2.15)), the cross section of inelastic rates tends to diverge at low collision energies as $1/k$. This leads to a different convention in expressing total scattering rates in a sample. The elastic scattering rate is expressed as a flux multiplied by a cross section, $\langle n \rangle \sigma \langle v \rangle$, where $\langle n \rangle$ is an average density and $\langle v \rangle$ is an average relative speed (which is proportional to k). Since the cross section diverges in inelastic collisions, the cross section is usually quoted as $K_2/\langle v \rangle$ so that the total scattering rate is $\langle n \rangle (K_2/\langle v \rangle) \langle v \rangle = \langle n \rangle K_2$.

In addition to two-body scattering events, three-body molecular recombination collisions are also significant. The three-body problem is much more complicated than the two-body problem and is not nearly as well developed. The current state of theory will be discussed in Chapter VI. Inelastic three-body cross sections also have a $1/k$ divergence and so $K_3 \langle n^2 \rangle$ is used to specify three-body collision rates similar to two-body inelastic collisions.

2.9 Rb-Rb Interatomic Potential.

2.9.1 Born-Oppenheimer Potentials. In order to calculate the elastic and inelastic scattering between Rb atoms, the time-independent Schrödinger equation (TISE) must be solved with the appropriate boundary conditions and with an accurate description of the Hamiltonian of the colliding pair of atoms [27]. Since two Rb atoms consist of a total of 74 identical electrons and two nuclei, all of which have a spin and magnetic moment, the full Hamiltonian is extremely complicated. One approximation that is used is the Born-Oppenheimer approximation [28]. The Hamiltonian can be split into a nuclear and electronic component, and the nuclei are frozen at a certain internuclear distance R . The electronic eigenstates are then computed as a function of R and that creates a set of potential energy curves. Since the electrons are so much lighter than the nuclei and can therefore adjust much more quickly to changes in the internuclear separation, the Born-Oppenheimer interatomic potentials are an accurate description of the atom-atom interaction potential.

When the nuclei are close enough to each other (about 20\AA), the overlap of the two electron clouds is an important consideration and molecular concepts must be used to describe the interactions of the colliding atom pair. The potential energy will depend greatly on whether the electrons are in a singlet or triplet spin configuration due to the Pauli exclusion principle, since in the triplet configuration the spatial part of the electron wavefunction must be anti-symmetric and so the overlap between electrons is minimized in that case.

When the nuclei are farther away, the energy due to the overlap of electrons decreases and the interatomic potential is then dominated by van der Waals, also called dispersion, forces. As the electrons move, small fluctuations occur in the charge density surround each atom. In turn, one atom can polarize the other atom by momentarily creating a dipole moment, and that in turn attracts the first. The dispersion potential can be expanded in a multipole expansion such as

$$V_{\text{dispersion}} = -\frac{C_6}{R^6} - \frac{C_8}{R^8} - \frac{C_{10}}{R^{10}} - \dots \quad (2.26)$$

where R is the internuclear separation and $1/R^6$ is the leading term for two alkali atoms colliding in the ground state.

2.9.2 Atomic hyperfine interaction. The overlap of the wavefunction of the outer electron in alkali atoms with the nucleus leads to an additional energy aside from the energy due to the electronic structure of the two colliding atoms. The magnetic field produced by a spinning nucleus couples with the magnetic moment of the electron leading to an interaction potential called the hyperfine interaction [29]. The orientation of the magnetic field of the nucleus and the magnetic moment of the electron depends on the spin of the nucleus and the electron respectively, and so the hyperfine energy only depends on the total angular momentum of the electron and the nucleus summed together. The ground state structure of ^{85}Rb discussed in section 4.2 contains the details of this interaction for that case. The energy scale of the hyperfine interaction is approximately 150 mK for ^{85}Rb .

2.9.3 Comparison of Rb-Rb Interaction Potential with the “Toy Model.” The energy scale of the hyperfine interaction, while much greater than the typical cold collision energy, is much smaller than the molecular energy terms discussed in the subsection 2.9.1. It is therefore only important when the two atoms in the colliding pair are far apart. However, while the energy of the hyperfine interaction depends only on the total electron and nuclear angular momentum, the energy due to the overlap of the two colliding atoms’ electron clouds depends on only the electron spin orientation, not on the nuclear spin orientation. This means that the interatomic potential will couple different hyperfine internal states (defined far from the region of electron overlap) at close range.

This is the same situation that was incorporated into the “toy model.” The point at which the overlap between electron clouds becomes important in Rb-Rb collisions corresponds to the position R in the toy model. Incoming colliding Rb atoms are projected from the hyperfine basis into the electron singlet-triplet basis and different singlet and triplet scattering phases are acquired. At “ R ”, the atoms are reprojected back into the hyperfine basis. While the true Rb-Rb potential is not as sharp as the toy model, there is a well-defined region where the singlet-triplet basis becomes a much more natural one than the hyperfine basis [30,31]. Just as was in the case with the toy model, in general the reprojected back onto the hyperfine basis would lead to inelastic scattering since the singlet and triplet phases are usually different. This process is called spin-exchange and can be an important consideration in atom-atom collisions.

If the overlap of the electron clouds was the only interaction, then the s-wave scattering length would likely be limited to be roughly equal to the radius at which the overlap became significant (in the absence of a bound state near zero energy). But the dispersion potential, equation (2.26), acts like the extra attractive part of the potential introduced into the toy model. The presence of this part of the potential allows the scattering length to take values other than “R”.

Feshbach resonances are also possible in the same way as they were in the toy model potential. Of course, the Rb-Rb potential is much more complicated, but the idea is still the same. If two atoms are colliding in a lower energy hyperfine state, then the potentials that exist between atoms in a higher energy hyperfine state may contain bound states close to the collision energy. The spin composition of the colliding state and the state in the higher energy hyperfine potential are different, and so the shift in energy due to the Zeeman effect with the application of an external magnetic field is also different. It is then possible to change the relative energy between the two potentials just as was demonstrated in the toy model, and a Feshbach resonance would exist in the atom-atom collisions. In ^{85}Rb , the Feshbach resonance that was studied in the work described in this thesis occurs between atoms colliding in the $F=2$, $m_f = -2$ hyperfine state. The closed collision channel responsible for the resonance is a bound state in the $F=3$, $m_f = -3$; $F=2$, $m_f = -1$ collision channel.

2.9.4 Matching the Calculated Born-Oppenheimer Potential to Measured Scattering Rates. In principle, it would be possible to calculate the Rb-Rb interatomic potential from just the Schrödinger equation and the total Hamiltonian of the two Rb atoms in the Born-Oppenheimer approximation. This is complicated, and different techniques are needed for different parts of the potential, including ab initio calculations and matching molecular spectroscopic data [27]. These calculations are not accurate enough to correctly predict ultracold collision scattering properties. Molecular spectroscopy of the upper bound states in the potential is difficult and extrapolating to the last bound states from the measurement of the lower bound states has proved to be unreliable. In order to match the potentials to scattering experiments, an additional adjustable parameter is added in the form of a correction to the inner wall of the singlet and triplet potentials. This parameter is adjusted until the ultracold collisions in a given experiment are reproduced. The position and width

of a Feshbach resonance are very sensitive to the exact details of the inner part of the potential and so the measurement of a Feshbach resonance tightly constrains the triplet and singlet scattering phases (see Chapter V).

2.9.5 Magnetic Dipole Interaction and Second-Order Spin Coupling. Two other interactions are important in the consideration of Rb-Rb scattering: the magnetic dipole interaction and second-order spin orbit coupling. These terms are not very important for elastic scattering, but they lead to significant inelastic scattering. The magnetic dipole interaction is the coupling between the magnetic dipole associated with the spin of one Rb atom's valence electron with the magnetic field produced by the magnetic dipole of a colliding Rb atom's valence electron. This interaction is not spherically symmetric and can couple incoming s-waves to outgoing d-waves (an inherently inelastic collision). It is relatively weak compared to the other interactions considered, but it does lift the degeneracy between the different spin projections in the triplet state.

This lifting of degeneracy is mimicked by an effect called the second-order spin-orbit interaction. If the electron clouds overlap enough, then a modification to the spin-spin interactions can occur due to indirect perturbative coupling to excited electronic molecular states. This is only important for heavy alkali atoms, and it is a significant contribution in ^{85}Rb with respect to the magnetic dipole interaction. Since the second-order spin interaction mimics the effect of the magnetic dipole interaction, the two produce the same effect and cannot be directly distinguished from one another.

2.9.6 Identical Particles and Spin Statistics. Throughout the discussion so far, the spin statistics and the indistinguishability of identical particles has been ignored. Both naturally occurring species of Rb are Bosons, since they have integer total spin. When two Rb atoms in the same state collide, their indistinguishability results in several factors of two that need to be considered in calculating the scattering rate. These details are covered in section 5.2.

Chapter III

Bose-Einstein Condensation and Mean-Field Theory

3.1 Overview

3.1.1 Summary of the Chapter. Bose-Einstein condensation is a phase transition that occurs in gases composed of particles with Bose statistics, which are those with integer spins. It happens at low temperature and is characterized by a macroscopic occupation of the lowest energy quantum state. This macroscopic occupation must happen as a consequence of thermodynamics, the identical nature of the particles in the Bose gas, and the fact that there are no restrictions on the number of particles in any quantum state (unlike Fermions). The first part of this chapter describes the thermodynamic considerations that lead to Bose-Einstein condensation and discusses the coherence properties of Bose-Einstein condensates (BECs), both in the absence of any atom-atom interactions.

The presence of atom-atom potentials (such as those discussed in the previous chapter) has important consequences for the properties of a BEC. In fact, the magnitude and sign of the interatomic potential determine, among other things, the size, collective excitation frequencies, stability, shape, and expansion energy of a BEC. The full solution of the Schrödinger equation is impractical for more than a few atoms, and so the atom-atom interactions are treated in mean-field theory. Here, mean-field theory is discussed briefly in the context of the Hartree approximation, and an equation for the ground state of a BEC with interactions is obtained.

Rather than using the Hartree language familiar to many atomic physicists, it is also possible (and equivalent) to use the formalism of many-body physics to describe BECs. This formalism is well suited to discuss the next-order corrections to the mean-field theory, and those corrections are discussed in a qualitative fashion. This is of particular interest in ^{85}Rb BECs since it is possible, via the Feshbach resonance, to achieve conditions where the mean-field theory is no longer expected to apply. Finally, the so-called Thomas-Fermi limit is discussed. This limit is

applicable for BECs with large mean-field interaction energy and is often very useful in estimating the properties of BECs.

It is not my intention to be as detailed in this chapter as in the previous chapter. Instead, my goal is to present the mean-field equations that describe BECs and qualitatively discuss their limitations. The interested reader is encouraged to consult the references for more detailed information. An excellent review of Bose-Einstein condensation is contained in Ref. [12].

3.2 The Phase Transition to Bose-Einstein Condensation in the Non-interacting Limit.

3.2.1 Thermodynamics. The phase transition to a BEC in a gas is the consequence of three things: there is no way to distinguish two identical particles from one another, there is no limit to the number of particles that can be in the same quantum state for particles exhibiting Bose statistics, and in thermodynamic equilibrium the entropy of the system must be maximized. In the treatment in this thesis, the grand canonical ensemble is used for ease of calculation, even though the microcanonical ensemble is more appropriate for trapped samples since the number of atoms and energy are fixed quantities. In the thermodynamic limit, which is letting the number of atoms go to infinity while holding the density constant, the choice of ensemble does not matter when central values are calculated. Also, the choice of statistical ensemble has been shown not to be significant for calculating the fraction of atoms in the BEC [32] aside from small effects due to the finite size of the sample. Given a total energy and number of particles, the entropy of a system of Bosons is maximized when the occupation number of each quantum state, labeled n , is [33]

$$f_n = \frac{1}{\exp\left(\frac{E_n - \mu}{k_B T}\right) - 1} \quad (3.1)$$

where E_n is the energy of the state, k_B is the Boltzmann constant, T is the temperature of the gas and μ is the chemical potential. The temperature and chemical potential can be considered as Lagrange multipliers that are used to satisfy the constraints of average energy and number of particles while maximizing the entropy. The temperature is related to the average energy per particle in the gas, and

the chemical potential is the change in energy of the sample when another particle is added. The factor of -1 in the denominator of equation (3.1) is associated with the Bose statistics of the gas.

In general, the chemical potential of the gas is negative, and the occupation number is positive and finite for all the possible quantum states. As the number of atoms in the gas is increased or the temperature decreased, the chemical potential must increase. Eventually, the chemical potential becomes nearly equal to the lowest possible quantum state. The denominator in equation (3.1) is then close to zero and the occupation of the lowest energy state becomes appreciable. At cold enough temperatures, nearly all the particles will be in the lowest energy state of the system. This macroscopic occupation of the lowest energy state of the system is referred to as Bose-Einstein condensation.

BEC is often considered in what is called the homogeneous case. In this case, there is no external potential present in the gas, and the Hamiltonian is simply the free-space Hamiltonian, $H=p^2/2m$. The lowest energy state is then $p=0$. For samples in the thermodynamic limit, the occupation of the $p=0$ state is a miniscule fraction of the gas unless the temperature of the sample is below a critical temperature. Once the gas is cooled below this critical temperature BEC occurs and the $p=0$ state contains a sizable fraction of the number of particles in the gas. There is no spatial signature of this condensation; in the homogeneous case BEC is a feature only in momentum space.

In the inhomogeneous case, BEC occurs in the presence of an external potential. In the case of BEC in trapped atomic gases the external potential is that of a simple harmonic oscillator. Ignoring any atom-atom interaction potential, the single-particle energy states of the system are then specified by three quantum numbers, (n_x, n_y, n_z) and the energy of a particular state is

$$E = \hbar \left\{ \left(n_x + \frac{1}{2} \right) \omega_x + \left(n_y + \frac{1}{2} \right) \omega_y + \left(n_z + \frac{1}{2} \right) \omega_z \right\} \quad (3.2)$$

where the trapping frequency in each direction is specified by ω . Unlike the homogeneous case, the lowest energy state in the inhomogeneous state has a finite spatial extent. In the absence of any interactions, the ground state is a three-dimensional gaussian whose widths are individually set in each direction by the strength of the trapping potential in that direction. Therefore, BEC does not

only occur in momentum space, but it is associated with an increase in density in the center of the trapping potential.

If the characteristic thermal energy scale $k_B T \gg \hbar \omega$, then the discreteness of the energy states can be ignored and E can be treated as a continuous variable. For most actual magnetic traps this is a reasonable assumption. For instance, for 10^6 ^{87}Rb atoms in trap with all the trap frequencies equal to $2\pi \times 60$ Hz, $k_B T_c = 94\hbar\omega$. In our ^{85}Rb traps with 10^4 atoms and $\omega_x = \omega_y = 2.56\omega_z = 2\pi \times 17$ Hz, $k_B T_c = \sim 20\hbar\omega$. The total number of particles in a gas in thermal equilibrium can then be calculated by integrating over all the possible energy states

$$N = N_0 + \int_0^\infty dE \left\{ \frac{E^2}{2\hbar^3 \mathbf{w}_x \mathbf{w}_y \mathbf{w}_z} \right\} \frac{1}{\exp\left(\frac{E - \mathbf{m}}{k_B T}\right) - 1}. \quad (3.3)$$

The energy has been shifted by removing the zero-point energy so that the lowest energy state is at $E=0$. The quantity in the brackets is the density of states and reflects the fact that as the energy increases there is an increasing number of different combinations of (n_x, n_y, n_z) which have the same energy. The number of atoms in the lowest energy state (the number of atoms in the BEC), $N_0 = (\exp(-\mu/k_B T) - 1)^{-1}$, has been separated from the integral since the assumption of a continuous distribution of energy does not properly account for the occupation of particles in the lowest energy state.

The chemical potential cannot be a positive quantity since the occupation number of the lowest energy state would then be negative. Solving the integral in equation (3.3) with $\mu=0$ results in the following relationship

$$(N - N_0) = 1.202 \left(\frac{k_B T}{\hbar} \right)^3 \left(\frac{1}{\mathbf{w}_x \mathbf{w}_y \mathbf{w}_z} \right) \quad (3.4)$$

The quantity $(N - N_0)$ is often referred to the ‘‘thermal component’’ of the gas. Since N_0 is a positive quantity, the maximum number of particles not in the lowest excited state has a limit that depends on the temperature. In other words, the number of atoms not in the lowest energy state saturates. By setting N_0 to zero in equation (3.4), the critical temperature, T_c , can be calculated in the thermodynamic limit

$$k_B T_c = 0.94 \hbar \mathbf{w}_{ho} N^{\frac{1}{3}}, \quad \mathbf{w}_{ho} = (\mathbf{w}_x \mathbf{w}_y \mathbf{w}_z)^{\frac{1}{3}} \quad (3.5)$$

and the fraction of the atoms in the condensate is

$$\frac{N_0}{N} = 1 - \frac{T^3}{T_c^3}. \quad (3.6)$$

This relationship between temperature and condensate number is shown in Figure 3.1. Comparisons between this prediction and the measured fraction are in reasonable agreement, especially when small effects due to the finite size of the sample are taken into account [34].




Figure 3.1. The fraction of atoms in the lowest energy quantum state as a function of temperature is shown. This curve is calculated in the thermodynamic limit. The sudden increase in the fraction of atoms at $T=T_c$ is the phase transition in which a Bose-Einstein condensate is created.

3.2.2 Coherence Properties. Another way of looking at the phenomenon of Bose-Einstein condensation is to examine the de Broglie wavelength as compared to the interparticle spacing. The thermal de Broglie wavelength ($\lambda=h/(2\pi mkT)^{1/2}$) is the scale over which the atoms exhibit coherent, or wavelike non-classical behavior. When the interparticle spacing is greater than the de Broglie

wavelength the particles do not exhibit coherent behavior and behave like classical particles. As the product $n\lambda^3$ (n is the density) increases, the overlap of de Broglie waves increases and quantum effects become more important. The transition to a BEC is associated with $n\lambda^3 \sim 1$. As coherent objects, BECs have a definite phase [35], and so therefore exhibit interference phenomena [36], particle correlations different than a thermal distribution [37], and superfluid behavior such as vortices [38].

3.3 Effect of Atom-atom Interactions on Bose-Einstein Condensates

3.3.1 Hartree-Fock Theory. In the absence of any atom-atom interactions, the wavefunction of the ground state ($T=0$) of a gas of N Bosons in a harmonic trapping potential is simply

$$\Psi(\vec{r}_1, \vec{r}_2, \dots, \vec{r}_n) = \alpha(\vec{r}_1) \alpha(\vec{r}_2) \dots \alpha(\vec{r}_N) \quad (3.7)$$

where $\psi(\mathbf{r})$ is the single-particle ground state of the potential. The inclusion of interactions between the atoms modifies the ground state (and excited state) energies and also modifies the ground state wavefunction. Since the atom-atom interactions in general depend on the relative coordinates of all the particles in the sample, the complete solution of the many-body system is extremely difficult and impractical.

In trapped samples of atomic Bose gases, the atom-atom interactions are weak enough that they can be treated as a perturbation. Even though this is the case, the interactions can be a significant contribution to the energy of the system, completely able to dominate the kinetic energy, for instance. Even though the atom-atom interactions can be strong enough to be a significant contribution to the energy of the system, it is still possible to treat them as a perturbation since there is a separation in scales between the spatial size of the condensate and the range of the atom-atom interaction. See Refs. [12,39] for details.

Other approximations can also be made because the samples are dilute. First, only binary collisions contribute significantly and three- and more-body elastic collisions can be almost completely ignored, with the exception being three-body molecular recombination (subsection 6.1.2.3) which can lead to loss. Second, the independent particle approximation can be used. Rather

than trying to solve the wavefunction as a function of all the individual atom coordinates and therefore all possible relative coordinates, the wavefunction of a single atom is calculated as it moves through a mean-field created by the atom-atom interactions with all of the rest of the other atoms. Since all the atoms of the BEC are identical (and are Bosons), the motion of one atom in the ground state in the independent particle approximation must be the same as the rest and so the orbit and the potential must be solved in a self-consistent fashion (the Hartree-Fock approximation). The following discussion closely follows that in Ref. [40].

The total Hamiltonian with binary interactions included is

$$H = \sum_{i=1}^N H_0(\vec{r}_i) + \sum_{i<j} V(\vec{r}_i - \vec{r}_j) \quad (3.8)$$

where V is the atom-atom interaction potential and the single-particle Hamiltonian is

$$H_0(\vec{r}) = \frac{-\hbar^2}{2m} \nabla^2 + \frac{m}{2} (\mathbf{w}_x^2 x^2 + \mathbf{w}_y^2 y^2 + \mathbf{w}_z^2 z^2). \quad (3.9)$$

The trial wavefunction in the independent particle ansatz is

$$\Psi(\vec{r}_1, \vec{r}_2, \dots, \vec{r}_n) = \varphi_0(\vec{r}_1) \varphi_0(\vec{r}_2) \dots \varphi_0(\vec{r}_N) \quad (3.10)$$

where φ_0 is the single-particle orbital that is to be determined self-consistently. With this definition of the trial function and the Hamiltonian, the expectation value of the Hamiltonian becomes

$$\begin{aligned} E &= N \frac{\langle \varphi_0 | H_0 | \varphi_0 \rangle}{\langle \varphi_0 | \varphi_0 \rangle} + \frac{N(N-1)}{2} \frac{\langle \varphi_0 \varphi_0 | V | \varphi_0 \varphi_0 \rangle}{\langle \varphi_0 | \varphi_0 \rangle^2} \\ \langle \varphi_0 | H_0 | \varphi_0 \rangle &= \int d^3x \varphi_0^*(\vec{r}) H_0(\vec{r}) \varphi_0(\vec{r}) \\ \langle \varphi_0 \varphi_0 | V | \varphi_0 \varphi_0 \rangle &= \int d^3x' \int d^3x \varphi_0^*(\vec{r}') \varphi_0^*(\vec{r}) V(\vec{r}' - \vec{r}) \varphi_0(\vec{r}) \varphi_0(\vec{r}'). \end{aligned} \quad (3.11)$$

The $N(N-1)/2$ term comes from summing the interaction term over all of the possible pairs in the sample. The ground energy eigenstate can be determined variationally by requiring that the energy calculated in equation (3.11) be stationary with respect to changes in $\varphi_0(\mathbf{r})^*$. In other words $\delta E=0$ with respect to variations $\delta \varphi_0(\mathbf{r})^*$. After some algebra [40], it is found that this condition is satisfied if $\varphi_0(\mathbf{r})$ is an eigenstate of the equation

$$\begin{aligned}
 [H_0(\vec{r}) + V'(\vec{r})] \phi_0(\vec{r}) &= \epsilon_0 \phi_0(\vec{r}) \\
 V'(\vec{r}) &= \int d^3x' \phi_0^*(\vec{r}') V(\vec{r}' - \vec{r}) \phi_0(\vec{r}')
 \end{aligned}
 \tag{3.12}$$

Here it has been assumed that ψ_0 has been normalized so that $\langle \psi_0 | \psi_0 \rangle = 1$. The evaluation of the integral V' is somewhat tricky. It is important that the atom-atom interactions be evaluated within the independent particle approximation. For an extreme example, consider a hard-sphere potential (V is infinite if $(\mathbf{r} - \mathbf{r}') < R$, zero if $(\mathbf{r} - \mathbf{r}') > R$). Plugging this potential into the expression for $V(\mathbf{r} - \mathbf{r}')$ causes the integral to be infinite for any nonzero value of $\psi_0(\mathbf{r})$, regardless of the size of the hard-sphere radius R . This seems nonsensical since even in the limit of an infinitesimal radius R , V' would still be infinite. Or, consider an attractive atom-atom potential that has a zero scattering length. Even though collisions were not occurring in the sample, V' could still be quite large and negative.

The reason for these apparent difficulties is that in the independent particle approximation, no correlations between particles are allowed. In the case of the hard-sphere scattering, in reality there would be no two atoms that were closer to each other than the radius R . However, there is no way for the independent particle approximation to account for this. Should the hard-sphere interactions then just be ignored since there are no atoms closer than the radius R and therefore no contribution from the atom-atom potential to the energy? This is not correct either, and the reason is that the presence of the hard-sphere potential causes the wavefunctions of two neighboring atoms to bend so that there is no overlap between the two within the hard-sphere radius. The resulting curvature of the wavefunctions represents kinetic energy that must be included in the total energy of the system.

The solution to this difficulty is to consider a potential that closely reproduces the scattering properties of the actual atom-atom potential that will work in the independent particle approximation. In order for this to be the case, the scattering potential should not significantly alter the incoming wavefunction in a scattering event in the region of the potential. For instance, it is possible to create a weak repulsive potential that will have the same scattering length as a hard-sphere potential. Outside the range of the potential, the scattering wavefunctions in the two cases will be identical. Inside the range of the weak potential, there is hardly any difference in the

wavefunction from the incoming wavefunction and that difference can be ignored for sufficiently dilute gases. The ability to ignore the change of the wavefunction in the scattering region implies that the correlations between particle coordinates will not affect the energy of the sample and the independent particle approximation will accurately describe the mean-field energy due to the atom-atom interactions.

The condition that the incoming wavefunction not be significantly altered in the region of the atom-atom potential is precisely the same condition that guarantees the Born Approximation (see section 2.5) will be accurate. Therefore, in the calculation of the mean-field energy, the actual atom-atom interaction potential is replaced by a different potential. This substitute potential will not possess a bound state so it can be accurately described in the Born Approximation, and the potential will be selected so that the s-wave scattering length of the replacement potential, *calculated in the Born Approximation*, is the same as the real s-wave scattering length of the actual atom-atom potential. Since the typical length scale of a scattering event, the s-wave scattering length, is usually much smaller than the spatial extent of the BEC, the specific form of the replacement potential does not matter. A convenient choice is a so-called pseudopotential [39]:

$$V(\vec{r}' - \vec{r}) = \frac{4p\hbar^2 a}{m} \mathbf{d}^{(3)}(\vec{r}' - \vec{r}) \quad (3.13)$$

This pseudopotential has an s-wave scattering length of a when the scattering is calculated in the Born Approximation. The Dirac δ -function also simplifies calculations. Once this substitution has been made, the ground state wavefunctions should be calculated only within a perturbation theory and not exactly, since in three dimensions, a gas of δ -function particles do not collide. See Ref. [40] for more details.

Plugging equation (3.13) into equation (3.12) results in the following equation for the ground state of the condensate

$$\left\{ \frac{-\hbar^2}{2m} \nabla^2 + \frac{m}{2} (\mathbf{w}_x^2 x^2 + \mathbf{w}_y^2 y^2 + \mathbf{w}_z^2 z^2) + \frac{4p\hbar^2 a(N-1)}{m} |\vartheta_0|^2 \right\} \vartheta_0 = \mathbf{e}_0 \vartheta_0 \quad (3.14)$$

Again, $\langle \psi_0 | \psi_0 \rangle = 1$. The single-particle orbital that obeys this eigenvalue equation will describe the ground state of BEC. The energy of that state will be ϵ_0 , which is equal to

$$\int d^3x \phi_0^*(\vec{r}) \left\{ \frac{-\hbar^2}{2m} \nabla^2 + \frac{m}{2} (\mathbf{w}_x^2 x^2 + \mathbf{w}_y^2 y^2 + \mathbf{w}_z^2 z^2) \right\} \phi_0(\vec{r}) + \frac{4\pi\hbar^2 a}{2m} \int d^3x |\phi_0(\vec{r})|^4 \quad (3.15)$$

Notice the factor of two in the denominator of the last term as compared to the coefficient in equation (3.14). The excited states can also be calculated using familiar atomic physics techniques such as the Random-Phase-Approximation (RPA) [41].

The dependence of the mean-field interaction term on the s-wave scattering length has some interesting consequences. Even though two atom-atom potentials may be similar, the atom-atom interaction energy in a BEC for the two cases can be quite different. For instance, compare ^{87}Rb BECs with ^{85}Rb BECs at $B \sim 0$. In ^{87}Rb - ^{87}Rb collisions the scattering length is positive and so the mean-field interaction energy term is positive and the atoms in the BEC tend to repel each other. In ^{85}Rb - ^{85}Rb collisions, the only difference in the potential is due to the small effect of the change in the hyperfine structure and some small isotopic shifts and so on the potential looks very similar to the ^{87}Rb - ^{87}Rb case. However, for ^{85}Rb - ^{85}Rb collisions the scattering length is negative and so the mean-field interaction energy is negative and the atoms in the BEC attract each other, exactly opposite of the case in ^{87}Rb . Even a potential that is purely attractive can lead to repulsive interactions between atoms.

Of course, the mean-field treatment is not exact and must eventually break down. The conditions of validity are discussed in the subsection 3.3.3 and in Chapter VIII.

3.3.2 Many-body Physics Formulation. While the Hartree-Fock approximation may be more familiar to atomic physicists, the formalism of many-body physics can also be used to describe BECs with atom-atom interactions. Non-relativistic quantum field theory [42] has been extensively developed for use in a variety of condensed matter systems and has a rich history of describing BEC. In fact, this formalism is found in the vast majority of theory papers discussing BEC. After a brief description of the mean-field equation in the many-body formulation, corrections to the mean-field picture will be discussed.

Second quantization is employed in the formalism of many-body quantum field theory. The Hamiltonian is formulated in terms of creation and annihilation operators $\hat{\Psi}^\dagger(\mathbf{x})$, $\hat{\Psi}(\mathbf{x})$. The effect of the creation (annihilation) operator is to create (annihilate) a particle at position \mathbf{x} . These

operators operate on states $|n_0, n_1, \dots, n_x, \dots\rangle$ where the integers $n_0, n_1, \dots, n_x, \dots$ specify the number of particles which occupy each single-particle orbital $\psi_0(\mathbf{x}), \psi_1(\mathbf{x}), \dots, \psi_x(\mathbf{x}), \dots$, which form a complete, orthonormal basis themselves. The single-particle basis is usually chosen to be a solution of the Hamiltonian in the absence of particle-particle interactions. For instance, the ground state of a non-interacting Bose gas with N particles in a harmonic potential (a BEC) is $|N, 0, 0, \dots\rangle$ where $\psi_0(\mathbf{x})$ is just the single-particle ground state of the simple harmonic oscillator.

The basis states $|n_0, n_1, \dots, n_x, \dots\rangle$ with different occupation numbers are orthogonal to each other. Since $\hat{\Psi}(\mathbf{x})$ annihilates a particle, the expectation value $\langle \hat{\Psi} \rangle$ is normally zero. However, this is not the case when a Bose-Einstein condensate is present. In that case $\hat{\Psi}$ cannot annihilate the ground state, due to spontaneously broken symmetry in the ground state [42]. The expectation value $\langle \hat{\Psi} \rangle$ then acquires a finite value. The operator $\hat{\Psi}$ can then be decomposed as the sum $\Phi + \hat{c}$, such that $\langle \hat{\Psi} \rangle = \Phi$. The quantity Φ is a classical field having the meaning of an order parameter and is often referred to as the wavefunction of the condensate. The condensate density $n(\mathbf{r}, t)$ is then fixed through the relation $n(\mathbf{r}, t) = |\Phi(\mathbf{r}, t)|^2$. The quantity \hat{c} represents the quantum fluctuations of $\hat{\Psi}$ around its mean value. In strongly interacting systems, \hat{c} is not negligible and a description of $\hat{\Psi}$ is inaccurate if it is ignored. However, in dilute systems, \hat{c} is a small quantity and $\hat{\Psi}$ can be replaced by its mean value Φ with a high degree of accuracy. The behavior of the BEC can then be described in a mean-field treatment through the computation of the evolution of the parameter Φ .

In order to determine this mean-field evolution, it is necessary to first compute the time dependence of the operator $\hat{\Psi}$. The second-quantized Hamiltonian for an interacting gas (through the potential V_{int}) in an external potential, V_{ext} , is

$$\begin{aligned} \hat{H} = \int d^3x \hat{\Psi}^\dagger(\vec{x}) \left[\frac{-\hbar^2}{2m} \nabla^2 + V_{\text{ext}}(\vec{x}) \right] \hat{\Psi}(\vec{x}) + \\ \frac{1}{2} \int d^3x d^3x' \hat{\Psi}^\dagger(\vec{x}') \hat{\Psi}^\dagger(\vec{x}) [V_{\text{int}}(\vec{x}' - \vec{x})] \hat{\Psi}(\vec{x}') \hat{\Psi}(\vec{x}) \end{aligned} \quad (3.16)$$

where only binary collisions are considered. At this point, the actual particle-particle V_{int} will be replaced by a potential that can be treated accurately in the Born Approximation, just as in the

previous subsection. From this point on, the pseudopotential given in equation (3.13) will be used and its substitution will be justified below.

By applying the commutation relations for $\hat{\Psi}^\dagger$ and $\hat{\Psi}$ ($[\hat{\Psi}^\dagger(\mathbf{x}), \hat{\Psi}^\dagger(\mathbf{x}')] = 0$, $[\hat{\Psi}(\mathbf{x}), \hat{\Psi}(\mathbf{x}')] = 0$, and $[\hat{\Psi}(\mathbf{x}'), \hat{\Psi}^\dagger(\mathbf{x})] = \delta^{(3)}(\mathbf{x} - \mathbf{x}')$), the time evolution of $\hat{\Psi}$ is

$$\begin{aligned} i\hbar \frac{\partial \hat{\Psi}(\vec{r}, t)}{\partial t} &= [\hat{\Psi}, \hat{H}] \\ &= \left[\frac{-\hbar^2}{2m} \nabla^2 + V_{ext}(r) \right] \hat{\Psi}(\vec{r}, t) + \frac{4p\hbar^2 a}{m} \{ \hat{\Psi}(\vec{r}, t)^\dagger \hat{\Psi}(\vec{r}, t) \} \hat{\Psi}(\vec{r}, t) \end{aligned} \quad (3.17)$$

Taking the expectation value of both sides results in an equation for the time evolution of Φ

$$i\hbar \frac{\partial \Phi}{\partial t} = \left[\frac{-\hbar^2}{2m} \nabla^2 + V_{ext} \right] \Phi + \frac{4p\hbar^2 a}{m} \left\{ |\Phi|^2 \Phi + 2 \langle c^\dagger c \rangle \Phi + \langle cc \rangle \Phi^* + \langle c^\dagger cc \rangle \right\} \quad (3.18)$$

The term $\langle c^\dagger c \rangle$ is often referred to as the normal density and $\langle cc \rangle$ is referred to as the anomalous density. The normal density includes all of the atoms that are in the gas but are not included in the BEC mean-field wavefunction while the anomalous density is related the spatial correlation between the atoms (i.e. the probability of finding two atoms a given distance apart). The last term in equation (3.18) is a three-body term and is usually small enough to be ignored in dilute samples even when beyond-mean-field effects are being considered. Ignoring the small fluctuation terms results in the Gross-Pitaevskii (GP) [43] equation

$$i\hbar \frac{\partial \Phi(\vec{r}, t)}{\partial t} = \left[\frac{-\hbar^2}{2m} \nabla^2 + V_{ext} \right] \Phi(\vec{r}, t) + \frac{4p\hbar^2 a N}{m} |\Phi(\vec{r}, t)|^2 \Phi(\vec{r}, t) \quad (3.19)$$

The normalization of Φ has been changed from equation (3.18) $\int \Phi^* \Phi d^3x = N$ to $\int \Phi^* \Phi d^3x = 1$ for equation (3.19). The correlations between particle positions are included in the quantum fluctuation terms and so when they are ignored the result is a mean-field treatment. Just like in the previous section, care must be exercised when evaluating the interaction term in a mean-field theory, and that justifies the pseudopotential substitution made earlier. The eigenstates of equation (3.19) are determined by the solution to

$$\left[\frac{-\hbar^2}{2m} \nabla^2 + V_{ext} \right] \Phi(\vec{r}, t) + \frac{4p\hbar^2 aN}{m} |\Phi(\vec{r}, t)|^2 \Phi(\vec{r}, t) = \mu \Phi(\vec{r}, t) \quad (3.20)$$

where μ is the chemical potential. In practice, μ is adjusted so that the Φ has the proper normalization. The form of this equation is the same as that for equation (3.14) up to the order $1/N$, and so for large numbers both equations are equivalent. With equations (3.19) and (3.20), both the time-dependent and time-independent behavior of a BEC can be calculated in a mean-field approximation. The GP equation has been remarkably successful in describing properties trapped atomic gas BECs, such as the ground state size and shape, collective excitation frequencies, and rate of expansion of BECs in the absence of the trapping potential (see Ref. [12]). It has been studied extensively and is used as the starting point for many theoretical investigations. Throughout this work, equations (3.19) and (3.20) were used to predict the behavior of BECs.

3.3.3 Limitations to the GP equation. The reason for the success of the GP equation in describing the properties of trapped atomic gas BECs is that these samples are dilute. The diluteness parameter is na^3 where n is the density of the sample and a is the s-wave scattering length. As long as this is small, the mean-field description should be accurate.

That na^3 should be the diluteness parameter could be seen in several ways. In the independent particle approximation, the particles' wavefunctions in the region of the scattering potential are not accurate, and so the fractional volume of the sample where the wavefunction is not being treated exactly is proportional to na^3 . Since interpreting the mean-field energy as the first-order term in a perturbation expansion is reasonable, the ratio of the second-order term in the perturbation expansion can then be calculated. When this is done, the condition that the second order term is less than the first order term is precisely that na^3 must be much less than one [39]. Finally, it is possible to calculate the leading order expectation values of the fluctuation terms $\langle \hat{c}\hat{c} \rangle, \langle \hat{c}^\dagger \hat{c} \rangle, \langle \hat{c}^\dagger \hat{c} \hat{c} \rangle$ and it is found that they scale as $(na^3)^{1/2}$ [44] as does the leading correction term to the energy in a perturbation expansion [39]. These calculations are performed for static homogenous samples and require careful treatment because of the presence of divergences; normally a renormalization procedure must be used. Even though these calculations are performed in the homogenous case, the relevant excitations occur on a wavelength scale that is much smaller than the

size of a BEC in the typical inhomogeneous case and the homogenous calculations are expected to be reasonably accurate [44]. However, all of these calculations are performed in the static case, and the fluctuation terms are dynamical quantities in general. In principle, it is possible to calculate the time evolution of the fluctuation terms, even in the inhomogeneous case [45], but the treatment is quite complex and the theory is still being developed.

The GP equation is formulated in the limit of zero temperature, and so corrections are expected as the temperature of the gas increases. Just as calculating the beyond mean-field effects is complicated, including finite temperature effects is difficult as well. In the mean-field picture, the GP equation can be extended to finite temperature by treating the thermal component as a static bath [46], but treating the full dynamics of the thermal cloud/BEC system is much more difficult [47,48].

The difficulty in calculating beyond-mean-field effects and difficulty in calculating finite temperature effects have a similar origin. The total density of atoms in the gas is given by the expectation value $\langle \hat{\Psi}^\dagger \hat{\Psi} \rangle$, and that is equal to $\Phi^* \Phi + \langle c^\dagger c \rangle$ --- the BEC density plus the normal component. At zero temperature $\langle c^\dagger c \rangle$ is not equal to zero due to the presence of atom-atom interactions, and the presence of a normal component at zero temperature is referred to as quantum depletion. At finite temperature, the thermal atoms also form a normal component. Since the normal component is not considered in a mean-field treatment, mean-field theory is expected to be accurate only if the depletion is small, and so for strong enough interactions or large enough temperatures the mean-field treatment must break down.

Chapter VIII details our first efforts to measure corrections to the GP equation. Additional discussion on beyond mean-field effects and collective excitations is contained there where it is relevant to the measurement.

3.3.4 The Thomas-Fermi Limit. For samples with a large enough mean-field interactions, there is a simplification that can be made. The relative contribution of the mean-field interaction term is characterized by the dimensionless ratio Na/a_{ho} , where a_{ho} is the mean simple harmonic oscillator length

$$a_{ho} = \sqrt{\frac{\hbar}{m\mathbf{w}_{ho}}} \quad (3.21)$$

For strong repulsive interactions and large numbers of atoms, this ratio can be made to be very large, well over 100 in fact. Even though Na/a_{ho} is very large, na^3 can still be, and in fact usually is, a small quantity and so the mean-field treatment is still valid.

For a large value of Na/a_{ho} the kinetic energy can be ignored since its contribution will be small. Ignoring the kinetic energy term is referred to as the Thomas-Fermi limit. Equation (3.19) and (3.20) then involve no spatial derivatives of the order parameter. The ground state of a BEC in a harmonic potential can be easily determined to be

$$n(\vec{r}) = N|\Phi(\vec{r}, t)|^2 = \left(\frac{m}{4\mathbf{p}\hbar^2 a}\right) (\mathbf{m} - V_{ext}) = \left(\frac{m}{4\mathbf{p}\hbar^2 a}\right) \left(\mathbf{m} - \frac{m}{2}(\mathbf{w}_x^2 x^2 + \mathbf{w}_y^2 y^2 + \mathbf{w}_z^2 z^2)\right) \quad (3.22)$$

in the Thomas-Fermi limit (the normalization $\int \Phi^* \Phi d^3x = 1$ is still being used). In equation (3.22), the density is taken to only have meaning where it is positive. Elsewhere, the density is considered to be zero. The shape of a BEC in a harmonic potential is therefore a parabola in Thomas-Fermi limit instead of a gaussian as in the non-interacting limit. The chemical potential is determined by the normalization of Φ to be

$$\mathbf{m} = \frac{\hbar \mathbf{w}_{ho}}{2} \left(\frac{15Na}{a_{ho}}\right)^{\frac{2}{5}}. \quad (3.23)$$

The Thomas-Fermi limit can also be used to simplify the time-dependent GP equation (3.19). For instance, oscillations induced in the rms widths ($b_{x,y,z}$) of the parabolic-shaped BEC in the Thomas-Fermi limit will evolve much like [49]:

$$\ddot{b}_i + \mathbf{w}_i^2 b_i = \left(\frac{15\hbar^2 Na}{7^{5/2} m^2}\right) \frac{1}{b_x b_y b_z} \quad (3.24)$$

In ^{85}Rb , the Thomas-Fermi limit is sometimes applicable and sometimes not. Since using the Feshbach resonance can alter the value of the scattering length, the quantity Na/a_{ho} can be made

small or large regardless of the value of N . When the interactions are strong (i.e. a is large) the Thomas-Fermi approximation simplifies the description of the condensate.

Chapter IV

Tools of the Trade: ^{85}Rb Bose-Einstein Condensation Apparatus

4.1 Overview

This chapter is devoted to the description of the apparatus used to make all of the measurements that are detailed in the following chapters. In order to study ^{85}Rb Bose-Einstein condensates (BEC), the BEC itself must first be created, and that involves cooling a gas of ^{85}Rb atoms down from room temperature ($\sim 300\text{K}$) to only a few nanokelvin. This cooling is accomplished in a two-stage process. The first step uses Magneto-Optic Traps (MOTs) to collect and laser cool the ^{85}Rb atoms down to temperatures around $45\ \mu\text{K}$. At that point, the MOT is turned off and the atoms are loaded into a purely magnetic trap. Forced radio-frequency evaporative cooling is then used to cool the gas, all the way down to BEC temperatures if desired. Once the sample has been created, experiments are performed by changing the bias magnetic field in the magnetic trap in a controlled fashion. The atoms are then released from the trap and imaged using absorption imaging or counted using fluorescence detection. Each of these four steps is treated in this chapter in the same format. General principles are discussed first, and then the physical description, control, and calibration of the actual apparatus is described.

The apparatus that we used for these measurements is referred to as JILA II, the same machine used for several experiments in ^{87}Rb Bose-Einstein Condensation [1]. Since I was not involved with the experiment when it was built, I refer the reader to Ref. [1] for details concerning construction and design criteria. Instead, I will focus on how the apparatus was optimized and operated to create BEC in ^{85}Rb . This machine holds the distinction that it is the only one that has ever been used to create BECs in two separate species of atoms. In principle, it should be able to be converted from an ^{85}Rb BEC machine to an ^{87}Rb BEC machine in a matter of hours. While the

machine has been used to evaporatively cool both species of Rb within a half-hour of each other, we have not yet demonstrated the ability to create different species BECs on the same day.

4.2 ^{85}Rb Atomic Structure

4.2.1 Ground and First Excited States. In order to better discuss the laser cooling, optical pumping, magnetic trapping and absorption imaging, the atomic structure of ^{85}Rb will be outlined in this section. ^{85}Rb has 37 electrons, 37 protons and 48 neutrons. All but one of the electrons are in closed-shell orbits and so there is one optically active electron and the relevant excited state spectrum consists of excitations of that single electron. The nucleus has a spin $I=5/2$ which leads to hyperfine structure in the ground and excited states. The quantity F is simply the total angular momentum of the nucleus and the electrons and is therefore equal to $I + J$ in the standard angular momentum notation [50].

Figure 4.1 shows the ground state and first excited states. In this experiment, we used transitions on the D2 line at 780 nm. Because of the selection rule $\Delta F=0,\pm 1$ in electric dipole transitions and the magnetic sublevel structure, the transitions from the ground state $F=3$ level to the excited state $F'=4$ is a cycling transition. The transition from the ground state $F=2$ to excited state $F'=1$ transition is not closed, because for a particular polarization of laser light there is always a coherent superposition of magnetic sublevels in the $F=2$ ground state that cannot be coupled to the $F'=1$ excited state. Many photons can be scattered on the $F=3$ to $F'=4$ cycling transition before an off-resonant transition to the $F'=3$ state occurs, allowing the atom to possibly decay to the $F=2$ ground state and break the cycle. The cycling transition is used both for the absorption imaging and laser cooling and trapping. Transitions between other hyperfine ground and excited state levels are used to optically pump the atoms into a particular spin orientation.

4.2.2 Ground State Hyperfine Structure and Energy Dependences at High Magnetic Fields. Figure 4.2 shows the hyperfine structure of the ground state (see section 2.9.2 for a description of the hyperfine interaction). This structure is important for these experiments in two ways. First, it is the $F=2$ $m_f = -2$ state that possesses the Feshbach resonance and so it was necessary to prepare a gas of atoms polarized in that state. Second, magnetic trapping uses the Zeeman shift to create the trapping potential. In order for an atom to be trapped, it is necessary that the energy of the




Figure 4.1. ^{85}Rb Hyperfine Structure of the Ground and $^3P_{3/2}$ Excited State. This figure diagrams the hyperfine structure of the two states that are involved in the laser transitions used in laser cooling and trapping, optical pumping and absorption imaging. All of the hyperfine level splittings are given in MHz. The trapping and hyperfine repump transitions used in the MOTs are shown explicitly. The lifetime of the excited state is 26.45 ns [51]. The level splittings are taken from Ref. [52]




Figure 4.2. Ground State Magnetic Sublevels. The ground state is (and excited states are) characterized by an additional quantum number, m_f , which is equal to the projection of the angular momentum F onto a particular axis. This axis is usually chosen to lie along the direction of the magnetic field at the atom's location. The hatched states indicate magnetic sublevels that increase in energy (through the Zeeman interaction) with increasing magnetic field at low magnetic fields. It is therefore possible to trap the hatched states in a purely magnetic trap. The reason that the $F=3$ positive m_f states are trapped while it is the negative $F=2$ m_f states that are trapped is that the Zeeman energy is determined primarily by the projection of the electron spin alone, and the electron spin is aligned with the nuclear spin in the $F=3$ state and opposite in the $F=2$ state.

state increase with increasing magnetic field (i.e. that the atom is in a so-called weak-field seeking state). This means that only a limited number of states can be trapped.

For low magnetic fields, the Zeeman shift can be expressed very simply as $E = g_F \mu_B B m_F$, where μ_B is the electron magnetic moment, and g_F is a factor that depends on the hyperfine state. In ^{85}Rb , $g_F = 1/3$ in the $F=3$ state and $-1/3$ in the $F=2$ state. As the magnetic field is increased, this simple relationship no longer holds. The hyperfine interaction is diagonal in the total angular momentum F basis. The Zeeman energy depends almost exclusively on the electron orientation since the electron magnetic moment is much larger than the nuclear magnetic moment, and so the Zeeman interaction is diagonal in the product spin basis where the nuclear and electron spins are specified separately. Therefore, the product spin basis is the most appropriate for large magnetic fields where the Zeeman energy dominates the hyperfine interaction energy, while the total angular momentum basis is the most appropriate for small magnetic fields where the reverse situation is true. At intermediate fields, the energy eigenstates are a mixture of these two bases.

The Feshbach resonance in ^{85}Rb occurs at a field in the intermediate regime (155 G). Because the ground state is an S state with no electron angular momentum, it is possible to derive a simple analytical expression for the energies of the sublevels at any field, known as the Breit-Rabi equation [50]. While the total angular momentum is not a good quantum number at high magnetic fields, the sum of the electron spin (m_s) and nuclear spin (m_i) is always a good quantum number. For any given value of $m_F = m_s + m_i$, except when $m_F = \pm 3$, there are two possible combinations of m_s and m_i , and the one with positive m_s corresponds to the upper hyperfine state. For example, the $F=2$, $m_F = -2$ state corresponds to $m_s = -1/2$, $m_i = -3/2$. The other $m_F = -2$ state, the $F=3$, $m_F = -2$ state, has $m_s = 1/2$, $m_i = -5/2$.

The Breit-Rabi equation that determines the energy of these states (i.e. all the states except when $m_F = \pm 3$) is

$$E_{m_q = m_i + m_s = m_i \pm \frac{1}{2}} = -\frac{\hbar n_{hfs}}{2(2I+1)} + g_I \mathbf{m}_B m_q B \pm \frac{\hbar n_{hfs}}{2} \sqrt{1 + \frac{4m_q x}{2I+1} + x^2}, \quad (4.1)$$

$$x = \frac{(g_J - g_I) \mathbf{m}_B B}{\hbar n_{hfs}}$$

There is some potential confusion about the sign of g_I , the nuclear gyromagnetic factor. Since the nucleus has a positive charge, the nuclear contribution to the energy shift due an external magnetic field is opposite that of the electron. Sometimes that change in sign is included in g_I and sometimes it is included in the Hamiltonian used to derive equation (4.1). Sometimes the nuclear magnetic moment is used and sometimes g_I is defined with respect to the Bohr magneton (electron magnetic moment). Just so things are clear, the Zeeman shift used to derive (4.1) is

$$H_{Zeeman} = (g_J J + g_I I) \mathbf{m}_B B \quad (4.2)$$

and so the opposite sign of the nucleus is included in the factor g_I . The various quantities in equation (4.1) are as follows for ^{85}Rb [53] in the ground state:

Term	Symbol	Value
Ground State Hyperfine Splitting	ν_{hfs}	3.035732439 GHz
Electron Gyromagnetic Ratio	g_J	2.00233113
Nuclear Gyromagnetic Ratio	g_I	-0.000 293 640
Nuclear Spin	I	5/2
Bohr Magnetron	μ_B	9.274009×10^{-28} J/G
Planck's Constant	h	6.626069×10^{-34} J

Table 4.1: ^{85}Rb atomic structure constants necessary for the Breit-Rabi equation.

The energies vs. magnetic field for three states, the $F=2$, $m_f=-2$, $-1,0$, are shown in Figure 4.3. At low fields, both the $m_f=-2$ and -1 states are trapped. As the magnetic field increases, the energy of both states initially increases linearly, but the rate of increase slows and eventually both energy vs. magnetic field curves roll over. Since the magnetic moment is given by the derivative of the energy vs. magnetic field, the magnetic moments both of these two states are decreasing as the magnetic field increases. At about 400 Gauss (G), the $m_f=-1$ state will no longer be magnetically trapped since the sign of energy vs. magnetic field has changed and the atoms in that state will become strong-field seeking. Even before this happens, the weaker magnetic moment will make the atoms in the $m_f=-1$ state difficult to trap.




Figure 4.3. The energy of the $m_f = -2, -1,$ and 0 state as a function of magnetic field as determined by equation (4.1).

4.3 Magneto-Optic Traps (MOTs)

4.3.1 Doppler Cooling. In an absolute sense, the bulk of the cooling of the room temperature sample of ^{85}Rb is performed by using MOTs. Much of this cooling comes through the mechanism of Doppler cooling. Doppler cooling relies on the mechanical effects of an atom scattering photons. If a near-resonant laser beam is directed onto a single two-level atom, the atom will absorb and emit photons, scattering them at a rate [54]:

$$R = \frac{2p\Gamma}{2} \left(\frac{I/I_{sat}}{1 + \frac{I}{I_{sat}} + 4\left(\frac{\Delta}{\Gamma}\right)^2} \right) \quad (4.3)$$

where Γ is the natural linewidth defined as $\Gamma=(2\pi\tau)^{-1}$ where τ is the natural lifetime of the excited state (26.45 ns for the ^{85}Rb $^3\text{P}_{3/2}$ excited state), I_{sat} is the saturation intensity defined as $4\pi^2\hbar c/6\tau\lambda^3$ where c is the speed of light and λ is the wavelength of the light, and Δ (in Hz) is the detuning of the laser from the transition. This scattering of multiple photons will have an effect on the mechanical motion of the atom. Each photon carries a momentum with it equal to $(h\nu)/c = \hbar k$. When the atom absorbs the photon, momentum conservation requires that it increase its velocity along the laser beam by $\hbar k/m$ where m is the mass of the atom. After absorbing a photon, the atom will spontaneously decay from the upper state, and when it spontaneously decays, it will emit a photon in a random direction. Over many absorption and emission cycles, the average momentum kick of all the spontaneously emitted photons will be zero. Since the momentum kicks received from absorbing the photons from the laser beam will not average to zero, the net effect is an acceleration of the atom in the direction of the laser beam.

Doppler cooling is most easily described in one dimension. If two nearly resonant counter-propagating beams are directed onto a two-level atom, a velocity-dependent force will result. Even if the two laser beams are the same frequency, the frequencies experienced by a moving atom are different due to the Doppler shift. Cooling can be arranged by detuning the lasers to the red (lower frequency) of the resonant transition. In that case, the atom will scatter photons at a higher rate from the laser beam towards which it is heading than from the other laser beam. Momentum kicks from

that laser that the atom is moving towards will slow the atom, cooling it. This situation is illustrated in Figure 4.4. The viscous force coefficient can be calculated from the scattering rate by including the Doppler shift v/λ and the photon recoil momentum:

$$\frac{F}{v} = \frac{\mathbf{p}\Gamma\hbar k}{v} \left(\frac{I/I_{sat}}{1 + \frac{I}{I_{sat}} + 4\left(\frac{\Delta - v/I}{\Gamma}\right)^2} - \frac{I/I_{sat}}{1 + \frac{I}{I_{sat}} + 4\left(\frac{\Delta + v/I}{\Gamma}\right)^2} \right) \quad (4.4)$$

Doppler cooling is limited in two important ways. First, the act of shining the laser light on the atoms does introduce a heating rate, often referred to as statistical heating. Perhaps the easiest way to see the origin of this heating rate is to consider a gas of atoms where all the atoms are initially at rest. If on-resonant laser light illuminates on the sample from all directions, the atoms will begin to move. There will be no net motion introduced, but on average an atom will acquire a speed from the random-walk scattering and emission of photons. This statistical heating leads to a temperature limit on a Doppler-cooled sample referred to as the Doppler limit [55] $T_D = \hbar\Gamma/2k_B$, about 140 μK for Rb. The second limitation on Doppler cooling is that an atom that is moving sufficiently fast will not feel a significant viscous force, and so Doppler cooling will only have an effect on atoms that are already moving slowly.

4.3.2 Sub-doppler Cooling. The previous discussion has so far neglected the internal structure of the atoms being cooled. For most alkali atoms, their internal structure leads to mechanisms that can cool the atoms to temperatures significantly colder than the Doppler limit. The temperature that can be ultimately achieved is only a few times the recoil limit, when again statistical heating will limit the temperature of the sample. The recoil limit is given by the temperature associated with a velocity equal to a recoil of a single photon, $T_r = (\hbar k)^2 / (mk_B)$. For ^{85}Rb this corresponds to a limit of $T_r = 400 \text{ nK}$.

These additional mechanisms are generally referred to as polarization-gradient cooling. As the name implies, the cooling is related to the change in polarization in space that occurs when two lasers with different polarizations are overlapped. There are two distinct mechanisms that have comparable effects, a Sisyphus cooling mechanism and a mechanism which will be referred to as

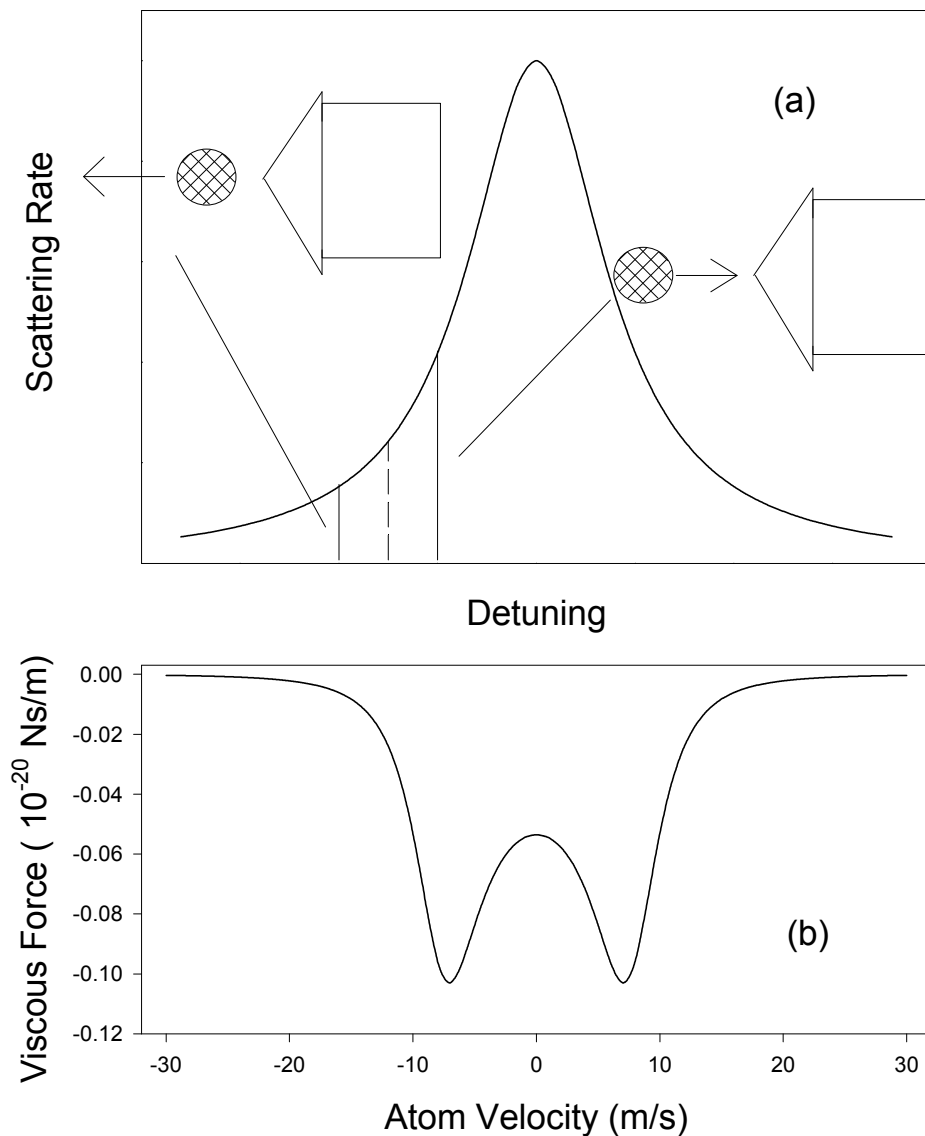


Figure 4.4. Frame (a) illustrates how the Doppler shift can lead to an increased scattering rate when an atom moves towards a red-detuned laser beam. The dashed line shows the frequency of the laser in the lab frame. The atom on the right moving into the beam sees the frequency of the beam shifted to the blue – towards the maximum scattering rate. The atom moving away from the beam sees the frequency shifted even farther to the red. Therefore, the atom moving into the beam scatters photons faster. Frame (b) shows the viscous force calculated from equation (4.4) with $I/I_{\text{sat}}=1$, $\Delta=10$ MHz and the linewidth and wavelength of the transition set to be equal to those of ^{85}Rb .

“induced orientation” cooling. Both are treated in detail in Ref. [56]. Again, these mechanisms are considered in one dimension.

The Sisyphus cooling mechanism occurs in the so-called lin \perp lin polarization configuration, where two counter-propagating laser beams of orthogonal linear polarization are overlapped. The polarization of the laser light as a function of position along the axis of the two beams changes from circularly polarized with one handedness (σ^+), to linearly polarized, back to circularly polarized with the opposite handedness (σ^-), to linearly polarized, back to the original polarization, and so on. In this section, the handedness of a circularly polarized light beam will be specified by the transitions that it drives in atoms that are exposed to it, rather than the direction the electric field winds around. A σ^+ polarization drives $\Delta m = +1$ transitions and a σ^- beam drives $\Delta m = -1$ transitions. Atoms with different magnetic sublevels, like ^{85}Rb , will experience different couplings to the excited state depending on the orientation of the atom and the polarization of the laser light (see Figure 4.5). The presence of the laser light leads to light shifts, which are shifts in the energy of the magnetic sublevels. These light shifts are proportional to the coupling strength between the ground and excited state, and so changing the polarization of the light will lead to different light shifts.

As an atom in a particular magnetic sublevel moves, it will slow down and speed up in response to the light shifts induced by the change in polarization. It will spend the most time at the laser polarization where it is moving the slowest, which is when the light shift is the most positive in energy. Since it spends the most time at the polarization with the largest positive energy light shift, it is most likely to scatter a photon at that polarization. If the atom spontaneously decays to the same level from which it was excited, it picks up a random momentum kick and nothing else happens. If it decays to a different magnetic sublevel, however, the light shift of that level will be lower than the one that from which it was excited. The spontaneously emitted photon in that case has more energy than the one that was absorbed, and therefore the atom must lose energy and slow down on average. Another way of looking at this is that when an atom is excited from the top of a light shift potential and then decays to a different state, it then has to move “uphill” again. Since it is more likely to

scatter from the top of a light shift potential than the bottom, it spends more time moving “uphill” than moving “downhill” and is cooled by this process.

The second polarization-gradient cooling process is sometimes referred to as “induced orientation” cooling. This cooling results from the optical pumping of the sample. In this case, the two counter-propagating beams have circular polarization, and one is polarized σ^+ , the other σ^- . The polarization of the light at any point in space is then linear, and the direction of this linear polarization varies in a helical pattern. Since the laser field polarization is linear at any given point in space (driving $\Delta m=0$ transitions in the frame of the laser’s polarization at a given point in space), an atom that undergoes several photon scattering events will be optically pumped into a magnetic sublevel state that has a symmetric distribution around $m_l=0$. As the atom moves, the direction of the linear polarization changes, and in the frame of the laser field the magnetic sublevel state will no longer be symmetric. The Clebsch-Gordan coefficients (see Figure 4.6) result in the atom scattering many more photons from the laser towards which the atom is moving. The lasers again pump the atom into a symmetric distribution, and it is found that the photons necessary to do so come overwhelmingly from the beam towards which the atom is moving, and that results in a viscous force that slows and therefore cools the atom.

Both of these polarization-gradient cooling methods have about the same size of viscous force, and the size of this viscous force is much greater than the Doppler cooling so that the temperature of the atoms can be cooled to a few times the recoil temperature. Since polarization-gradient cooling works by scattering at least a few photons within a laser wavelength, atoms that are moving much faster than several cm/s will not be affected by polarization-gradient cooling. Therefore the range of atom velocities that can be cooled is much more limited than for Doppler cooling. Also, these polarization-gradient cooling mechanisms rely on the atom spin orientation, and so the presence of a magnetic field, which causes the atom spin orientation to precess around the field, will have an effect on the polarization-gradient cooling. Strong enough magnetic fields will disrupt the polarization gradient cooling entirely.

4.3.3 Magneto-Optic Traps (MOTs). The cooling discussed so far will apply a viscous force to the atoms, but it does not have any dependence on position and so will not localize the




Figure 4.5. Clebsch-Gordan Coefficients. The transition rate between different pairs of magnetic sublevels of the excited and ground states depends on the orientation of the electron spin and the polarization of the laser driving the transition. The matrix elements involved in the transition rate are proportional to Clebsch-Gordan coefficients that are determined from angular momentum sum rules, and these coefficients, multiplied by a factor of 100, are shown in this figure. Since the transitions are proportional to the matrix elements squared, to compare the transition rate between different pairs of magnetic sublevels the ratio of the square of the Clebsch-Gordan coefficients needs to be considered.

atoms at any point in space. When laser cooling is done in without any spatial confining force in the manner described in the previous subsection, it is referred to as an optical molasses. However, the internal structure of the atoms can be used to not only create a viscous force, but also a trapping potential that will trap the atoms at a particular location in space. The way that this is done is through using magnetic fields in addition to the laser fields to create the atom trap. Hence the name Magneto-Optic Trap [57,58]. The general idea is again best illustrated in one dimension. A magnetic field is created that varies linearly in space. The ground and excited magnetic states will be shifted in energy by the Zeeman effect. Two red-detuned counter-propagating beams, one with σ^+ polarization and one with σ^- , illuminate the atoms. Because the σ^+ polarization drives only $\Delta m=+1$ transitions, the magnetic field (B) can be arranged so that when the atoms move away from the point where $B=0$ towards the σ^+ beam, that beam is shifted into resonance and the σ^- is shifted further out of resonance. Therefore, more photons will be scattered from the σ^+ beam than the σ^- beam and the atom will be pushed back towards the point where $B=0$. In the same arrangement of fields and lasers, atoms that move towards the σ^- will also be pushed back to $B=0$. This general behavior is described in Figure 4.6.

A real three-dimensional trap uses all of the ideas presented in the one-dimensional examples discussed above. Three orthogonal sets of counter-propagating red-detuned beams with σ^+ and σ^- polarization are detuned to the red and sent into a vacuum chamber. A magnetic field that varies linearly in space is applied, and the orientation of the polarizations of the laser beams are selected appropriately to create a confining force at the position where $B=0$ [59]. Atoms that are moving slow enough are cooled by Doppler cooling and then captured by the confining potential. Once confined, the atoms continue to be cooled, both by Doppler cooling and also by polarization-gradient cooling. The combination of these viscous forces and the confining magneto-optic forces creates a cold cloud of atoms trapped near $B=0$.

Besides the laser used to provide the light necessary for the trapping and cooling of atoms in the MOT, for alkali atoms another laser is required. This laser is referred to as the hyperfine repump laser. The cycling transition is used for the trapping light, but as was mentioned section 4.2.1, occasionally off-resonant transitions do occur and then the atom can spontaneously decay to




Figure 4.6. MOT Confining Force. In one dimension, a linear gradient in a magnetic field combined with the polarization of counter-propagating trapping beams shown results in a spatially confining force. The example above is for the simplest state structure possible. As an atom moves from the center ($B=0$), Zeeman shifts in the excited state will cause the atom to scatter more photons from the beam toward which it had moved. This results in a restoring force pushing the atom back to the origin.

the lower hyperfine state. Atoms in the lower hyperfine state will no longer scatter photons from the trapping laser and would be lost from the trap if they were allowed to remain in that state. In ^{85}Rb , this happens infrequently compared to the number of scatters from the trapping beams and a weak laser can be used to pump the atoms up from the lower hyperfine state back to the upper using the transition shown in Figure 4.1.

Given the σ^+ and σ^- polarization of the beams, it is tempting to say that only the “induced -gradient cooling occurs in the sample. This is only strictly the case when the atoms are moving only along one of the beam axes. When moving in any other direction, the change of laser light polarization with position is quite complex and so both polarization-gradient cooling mechanisms cool the sample. Given the complexity of the light polarization, it is perhaps surprising that the simple one-dimensional models provide reasonable descriptions of what is happening in the MOT.

The creation of a BEC relies on increasing the phase-space density, $n\lambda^3$, of a gas from its initial value at room temperature. While MOTs are a reasonable first step, they cannot be used to create BECs in alkali atoms due to inherent limitations in the MOT. The limitations to temperatures have been discussed already --- temperatures of only a few times the recoil limit can be obtained with the methods described above. There are also limitations to the density as well. Once many atoms are trapped in the sample, the cloud of atoms starts to become optically thick, and it becomes likely that a photon scattered from the laser beams will be scattered again on the way out of the trap. This rescattering of photons leads to a repulsive force (reradiation force) between the atoms, which limits the density of the sample to less than $\sim 10^{11} \text{ cm}^{-3}$ [60]. The presence of the laser light also introduces density-dependent light-assisted losses in the MOT as well [61], and this is another limitation on the density of the atoms in the MOT. At a temperature of $3T_r$ and a density of 10^{11} cm^{-3} , the gas of atoms is still three orders of magnitude in phase-space density away from the BEC transition. This simple analysis ignores the fact that this density and this temperature are not simultaneously achieved in MOTs, since the rescattering of photons also increases the temperature of the atoms in the MOT as well as limiting their density. For a detailed experimental treatment of

the maximum phase-space density in MOTs as well as the relationship between density, laser power, laser detuning, laser size, and other variables see Ref. [62].

4.3.4 Loading a MOT from a Vapor. Since Doppler cooling does not work very well for atoms that are moving sufficiently fast, it is necessary to have a source of cold atoms that can be loaded into (captured by) a MOT. One possible source is a Zeeman slower [63], a device that uses laser cooling to slow a thermal atomic beam to speeds that can be captured by a MOT. A much simpler scheme, and the one that is used in our apparatus, is to load atoms from the low-velocity tail of the Maxwell-Boltzmann distribution of speeds of atoms in a vapor cell [58]. Loading from a vapor results in loading times comparable to loading with a Zeeman slower, and the ease of operation of a vapor cell MOT more than makes up for any disadvantage.

The rate at which the MOT will load atoms from the vapor is proportional to the partial pressure of the background gas of the species that is to be trapped. The presence of this background vapor also leads to a loss of atoms from the MOT due to collisions of the cold MOT atoms with the hot background gas atoms, and so the loss rate depends on the partial pressure of the trapped species as well. The loading of the MOT at low pressures can be then described by a simple equation

$$\frac{dN}{dt} = R - \frac{N}{\tau_{BG}} \quad (4.5)$$

where N is the number of atoms in the MOT, τ_{BG} is the inverse of the collision rate between background atoms and MOT atoms, and R is a loading rate that depends on the available laser power, size of the laser beams, steepness of the magnetic field gradient, and other factors [58]. If the background vapor pressure is dominated by the atoms that are being trapped, the number of the atoms in the MOT will be a steady-state number, $N_{ss}=R\tau_{BG}$, since both R and τ_{BG} are then proportional to the vapor pressure. In order to achieve reasonable loading times, the pressure in the MOT chamber should be high enough that the MOT loading time is a few seconds. However, these same background atoms will result in loss during evaporative cooling, and so for evaporative cooling a low pressure is desired. These competing factors led to the creation of a Double-MOT system [1]. One MOT at relatively high vapor pressures collects atoms from the vapor. These atoms are then

transferred to another MOT in a vacuum chamber that is maintained at much lower pressure. In the lower pressure chamber, the atoms are then loaded into a magnetic trap for evaporative cooling.

If indeed Rb is the dominant component of the background vapor, then the pressure (P) can be estimated using the relationship

$$\frac{1}{\tau_B} = n_b \sigma v_b \quad (4.6)$$

where $n_b = P/k_B T$ is the density of the Rb atoms in the background vapor, v_b is the average speed of those atoms and σ is the collision cross section appropriate to collision energies at room temperature. Using a value for σ of $2 \times 10^{-13} \text{ cm}^2$ from Ref. [58] (even though this cross section was determined for Cs, Rb should be similar), equation (4.6) implies a two second lifetime τ_B corresponds to 3×10^{-9} torr.

4.4 The Double-MOT System.

4.4.1 The Role of the MOT. Since a MOT cannot be used to create a BEC, its purpose in our apparatus is to collect and initially cool a gas of atoms so that they can be loaded into a purely magnetic trap. The MOT conditions should then be adjusted to produce the best initial conditions in the magnetically trapped sample. What are these conditions for the best loading into the magnetic trap? One criterion that has often been used is phase-space density. Since BEC will be achieved once the phase-space density is increased to a particular value, it makes some sense to use set the MOT so that the initial phase-space density in the magnetic trap is as high as possible. It has been our experience in ^{85}Rb that phase-space density can be increased rapidly and efficiently through the use of evaporative cooling once the atoms have been loaded into the magnetic trap. Rather than try to simply maximize the initial phase-space density of the atoms loaded from the MOT to the magnetic trap, it is better to improve the conditions needed for efficient evaporative cooling (e.g. high initial elastic collision rate) at the expense of some phase-space density. This is almost the same as maximizing the initial phase-space density (the two depend slightly differently on the temperature of the sample). The criteria for ^{85}Rb then are that the gas should be both dense and cold, although the temperature is a little more important than ^{87}Rb or ^{23}Na due to the dependence of the ^{85}Rb elastic scattering cross section on collision energy (see section 7.1.1).

In order to maximize the MOT performance, the laser detuning, hyperfine repumping, and magnetic field gradient of the MOT used to directly load the magnetic trap are manipulated in several steps. First, the MOT that is used to load the magnetic trap is configured to efficiently collect a large number of atoms. Next, the laser detuning, hyperfine repumping, and magnetic field gradient are altered to increase the density and optimize the cloud shape of the gas trapped in the MOT. Finally, the magnetic gradient is turned off and polarization-gradient cooling is used to lower the temperature in the sample. Once these steps have been accomplished, there is one more step that is necessary; the atoms need to be optically pumped so that they are loaded into the trap in the proper ground state magnetic sublevel.

4.4.2 Vacuum System. A schematic of the vacuum hardware is shown in Figure 4.7. The Collection MOT chamber is a stainless steel conflat six-way cross. Five two-inch (1.5-inch clear aperture) diameter windows are mounted on five ends of the arms of the cross to allow five of the six MOT beams to enter the chamber. The other arm of the cross is connected to a four-way cross. One arm of this four-way cross has a two-inch window to allow the last MOT beam to enter the chamber. Through the third arm, an 11 L/s ion pump is connected to maintain a low ($\sim 10^{-9}$ torr) pressure in the chamber. The fourth arm is connected to a bakeable UHV valve that opens to a conflat where additional pumps can be attached for use when the vacuum system is being baked out (the “pump out station”). In addition, there are five other small ports in the chamber. Two of these ports are used to allow monitoring of the collection MOT through an inexpensive CCD camera and a photodiode that monitors the MOT fluorescence. Two other ports are associated with the transfer of atoms from the collection MOT to the science MOT. One is connected to the transfer tube, which is described below, and the other is a window that allows a laser beam to be sent down the transfer tube. The last port is connected to the supply of Rb. Solid Rb was enclosed in a glass ampoule and then placed into a stainless steel tube that was in turn connected to the collection vacuum chamber. The ampoule was broken under vacuum to release the Rb, and the partial pressure of Rb in the chamber is then controlled through heating the outside of the stainless steel tube that contains the ampoule. A simple needle valve can be closed to isolate the Rb source from the rest of the chamber. Now Rb getters are considered to be a more convenient source of Rb vapor for loading the MOT as




Figure 4.7. Simple schematic of the double-MOT vacuum system. Each arm of the science cell is round. The circular window on the science cell shown in this side view (and its counterpart on the other side) is, unlike the other directions, at the end of an arm only a few mm long. Hence the science cell resembles a four-way cross.

opposed to solid Rb in an ampoule since their output is more easily controlled and more directional [64]. The construction of this system predated the use of getters for this purpose, and in any case has functioned well for six years and counting.

The Science MOT chamber is a glass cell that resembles a four-way 1-inch diameter cross with one end that is flared open (see Figure 4.7). The open end is connected (via glass frit) to a stainless steel “T.” The end of the “T” opposite the cell has a 2-inch diameter window to allow one of the MOT beams into the chamber. The other end of the “T” is connected to a stainless steel section that contains a titanium sublimation pump. At the end of that section, a 60 L/s ion pump is connected. Two additional ports extend off the glass cell. One is open and is connected to the transfer tube. The other port is sealed and is used as an observation port. Five orthogonal surfaces of the glass cell have been made to be flat, providing 1-inch diameter windows for the MOT beams.

These two chambers are connected via a 40 cm long, 1.1 cm diameter stainless steel tube referred to as the transfer tube. Bellows are connected on each end of the transfer tube to relieve any stress due to the mounting of the collection and science chambers. This long thin tube only has a small conductance, approximately 0.4 L/s for air, 0.2 L/s for Rb in the molecular flow regime [65] and allows for differential pumping between the collection chamber and the science chamber. Using the ion pump speeds listed above and making simple estimates about the conductance from the pump to the center of the MOT chambers implies that a factor of on the order of 100 between the pressure in the collection chamber and the pressure in the science chamber. This estimate of the differential pumping is may be too low, however, since the estimate of the conductance of the transfer tube is probably too high because the molecular flow regime is not applicable. The dominant source of vapor pressure in the collection MOT is Rb, and any collision of a Rb atom with the wall of the transfer tube would likely result in the Rb atom sticking to the wall, which is not the situation associated with the molecular flow regime. However, once the transfer tube does get coated with a layer Rb, the subsequent Rb atoms that strike the wall are not likely to stick and the molecular flow regime may be the correct one to use after all. In any case, the differential pumping does seem to produce a pressure difference on the order of a factor of 100.

4.4.3 Baking Out the Vacuum System and Vacuum Failures Since 1995. The amount of Rb in the ampoule was estimated to last ten years [1], and one of the advantages of the Double-MOT design is that there is no reason to crack the vacuum system for routine maintenance. As the pumps continue to remove residual outgassing materials from the surfaces in the vacuum system, the vacuum should simply get better and better as time goes on. This was the expectation. The reality, however, was that since the original bakeout in 1995, the system has been let up to air very briefly once and the system has been baked out again on two other occasions.

Baking out the system was complicated by our inability to easily remove the magnetic trapping coils. These coils were put onto the science cell after the initial bakeout, and they are sheathed in insulating material that cannot withstand temperatures much above 125° C. At temperatures near 125° C the insulating material would outgas and coat the science cell windows, which would in turn degrade the MOT performance. Also, it was not practical to remove the magnets on the science chamber ion pump.

In order to perform the later bakeouts, a turbopump was connected to the pump out station and all of the optics that were close to the chambers were removed. The ion pumps were wrapped with nichrome wire surrounded by ceramic beads, and then several layers of aluminum foil were used to cover the pumps. For the 11L/s pump, the magnets were removed and it was heated to over 300° C. The larger 60L/s pump was lifted ¼ inch off of the optical table to reduce the thermal the magnets on. Since the nichrome wire was heated using variacs, the combination of the magnets near the large ion pump and the AC current caused the wires to vibrate and care had to be taken that they did not short out. The larger sections of stainless steel were wrapped with nichrome wire as well. Heater tape and then aluminum foil was used on the rest of the apparatus except for the glass cell. The temperature of the sections covered by heater tape (transfer tube, upper MOT chamber) was maintained at about 100° C. The glass cell was heated by a standard heat lamp (and again aluminum foil was used to insulate) to 110° C. The system was baked for a few days at a time. After the bakeout, it was necessary to heat the Rb ampoule more than normal in order to produce the same vapor pressure in the collection MOT, a process sometimes referred to as “curing” the collection chamber.

There were two reasons for the additional bakeouts. One of the reasons, while exciting at the time it happened, is not very interesting. The window on the CCD viewing port of the collection chamber, for no apparent reason, suddenly sprung a leak. Since someone was present in the lab at the time, it was possible to shut down the ion pumps before any damage resulted and hook up a turbopump to the pump out station to hold a vacuum. The system was then briefly let up to atmosphere by filling it with dry nitrogen that had been boiled off liquid N₂. The faulty window and collection chamber ion pump were then replaced in a few minutes and the system was sealed, pumped back down and then baked out.

The other two bakeouts occurred as a response to a slow deterioration of the vacuum in the science chamber. The vacuum in the science chamber is most easily characterized by the lifetime of a low-density cloud of atoms trapped in the magnetic trap (see Chapter VI). After the initial 1995 bakeout, the lifetime was about 700 seconds. By 1997, it had dropped to 20 seconds. This was the most dramatic decrease that we observed, although it is not uncommon for the vacuum to degrade slowly over time. The reason for the large drop was traced to a failure of the collection chamber ion pump. Turning off the ion pump improved the vacuum in the science chamber, so instead of pumping the ion pump must have been outgassing. It is interesting to note that the failed pump had a large amount of surface current.

Since the correction of this problem with the collection chamber ion pump and through other subsequent bakeouts, we have observed that the pressure in the collection chamber limits the pressure in the science chamber, but the relationship is not exactly proportional and there is a delay between changes in the collection chamber and changes in the science chamber. Certainly, maintaining relatively high vapor pressures in the collection MOT for days at a time causes a measurable decrease in the lifetime of magnetically trapped samples in the science MOT over several days. Running with a lower pressure in the collection chamber over days will also improve the magnetic trap lifetimes. Flashing the titanium sublimation pump has never been correlated with an improvement in science chamber pressure, but gently heating the glass cell with a heat lamp (about 35-40° C) overnight has been observed to lead to improvement.

A typical collection MOT fill time is around two seconds, and a typical magnetic trap lifetime is 400 seconds, a ratio of 200. From equation (4.6) this implies a pressure of $\sim 1 \times 10^{-11}$ torr for the science cell, although the loss rates from collisions with unpumped background atoms for a MOT and a magnetic trap are in general different and so this is only an order-of-magnitude estimate. This is about the ratio than would be suggested by the differential pumping due to the ion pumps alone. Also, in the past there was a time when the magnetic trap lifetimes were 3000 seconds, a ratio of 1500 (these lifetimes have not been recovered since the apparatus came up to air). It appears that the titanium sublimation pump is providing no useful additional pumping, because when the lifetime of the magnetically trapped samples decreases, flashing the titanium sublimation pump does not seem to have any effect.

I speculate that the surfaces of the science cell are effectively pumping Rb. Alkali atoms have a high affinity for sticking to surfaces [66], so this supposition is reasonable. The day-length time delays between an increase in the pressure of the collection MOT and a subsequent decrease in the magnetic trap lifetime in the science cell are also reasonable given our experience with curing the collection MOT after a bakeout. Since the surface will stop pumping Rb atoms away once a layer of atoms has already been adsorbed, over time the surface pumping would become less effective. Heating the cell overnight would drive Rb from the surface and so it would improve the science cell vacuum as observed. If this scenario is the true cause of the vacuum deterioration over time, then running at lower pressures in the collection chamber and lightly baking the science cell overnight should reverse any decrease in magnetic trap lifetime.

After the bakeout that occurred after the system was let up to air, the magnetic trap lifetime was initially only 200 seconds. Over the course of several more days of heating just the science cell, the lifetime improved to nearly 600 seconds.

4.4.4 Collection MOT. The collection MOT is a standard three-beam retro-reflected MOT [1]. Two different external-cavity stabilized diode lasers in the Littrow configuration [67] provide the trapping and hyperfine repump laser light, although only part of the hyperfine repump laser light is used for the collection MOT. Efficient collection of atoms by the MOT requires large beam diameters and high laser intensities [58]. The diode laser used for the trapping light, currently a

Hitachi model HL 7851G, only produces 20 mW of output laser power as measured after the external cavity grating (which is how all of the laser powers in this chapter will be specified). That is weak compared to many other laser systems. The loading rate R in equation (4.5) depends more strongly on the beam diameter than on the intensity of the laser and so the beams are shaped by a pair of lenses to be collimated to almost fill the 2-inch windows of the collection chamber. The hyperfine light is combined with the trapping light through using a polarizing beam-splitting cube, and the hyperfine light is shaped by a pair of lenses to match that of the trapping light. After the light has been collimated, it is split into three directions using large-area beamsplitters. $\lambda/4$ waveplates are used to set the proper polarization of the MOT trapping beams.

The frequency of the hyperfine pump laser, a Sharp LT025MD0 with 7.5 mW of output power, is locked to the $F=3$ to $F'=3$ transition (see Figure 4.1) by using a standard saturated absorption set-up as a frequency transducer [68]. The hyperfine laser current is dithered at 10 kHz and a lock-in amplifier is used to create an error signal that is fed back to the grating PZT to hold the laser frequency at the peak of the saturated absorption signal corresponding to the desired transition. In contrast, the frequency of the trapping laser is stabilized at a detuning from the cycling transition by using a DAVLL system [69]. Rather than using the saturated absorption lines, this laser frequency control scheme uses a magnetic field to induce a rotation of the polarization of the linearly polarized laser light as it passes through a vapor cell. The magnitude of the rotation of the polarization depends on the frequency of the laser and so this can be used as a frequency transducer. Unlike saturated absorption signals, the polarization rotation signal has a width that is related to the Doppler width of the gas in the vapor cell, and so it allows a much larger range of frequencies over which the laser can be locked (over 100 MHz compared to about 10 MHz for saturated absorption). The error signal generated from the DAVLL is fed back to both the PZT and the current of the trapping laser to increase the bandwidth of the frequency control and to allow fast changes in the laser frequency.

In order to generate the magnetic field needed for the collection MOT, a set of anti-Helmholtz coils are used. Anti-Helmholtz coils are simply a Helmholtz coil pair with the current direction reversed in one of the coils so that the two coils tend to cancel each other's magnetic field

between them. The closer the coils are to the trapping region, the less current that is needed to produce a given gradient. Therefore, the coils are wrapped directly around two of the arms of the six-way cross. They are on average ~ 3 cm away from the center of the trapping region, 1 cm in extent, 3.6 cm in diameter, and consist of 10 turns. Magnet wire is used to make the coils, and a 10 A supply at full output is sufficient to supply the current for the anti-Helmholtz pair. The resulting gradient is 8 G/cm. In addition to the anti-Helmholtz coils, there are three shim coils that are used to shift the location of the zero of the magnetic field and hence the location of the MOT. The three shim coils are 8 cm in diameter, have between 50-150 turns of magnet wire, and are located about 6 cm from the center of the collection chamber. Each shim coil is oriented along an axis that is orthogonal to the other two coils.

The maximum number of atoms trapped in the collection MOT when it is allowed to fill completely is estimated to be roughly 2×10^8 from the fluorescence of the MOT. While loading the atoms from the vapor, the detuning of the trapping laser is 9.5 Mhz to the red of the transition. The amount of hyperfine repump power is 2.4 mW. The $1/e$ MOT loading time (see equation (4.5)) depends on the Rb vapor pressure, and is usually set to be about two seconds by controlling the Rb ampoule temperature and hence the Rb pressure. The MOT is operated nearly in the Rb-pressure dominated regime [58], so that increasing the Rb vapor pressure reduces the loading time but does not much change the number of atoms in the MOT.

4.4.5 Transfer from Collection MOT to Science MOT. Once the collection MOT has captured Rb atoms out of the vapor, they are pushed from the collection chamber to the science cell via the transfer tube. The atoms are accelerated down the transfer tube using a laser beam referred to as the push beam. Magnetic guiding is used to prevent the atoms from striking the wall of the transfer tube, and the atoms are slowed and trapped by another MOT in the science chamber. After one load is sent in this fashion, the collection MOT is turned back on to capture more atoms and the cycle is repeated, typically about 50 times. The time that the collection MOT is allowed to fill between pushes is typically 200 msec.

The creation of the push beam is accomplished through a liquid crystal switch (model LV300-OEM from Diplaytech, Inc.) that controls the direction of the trapping laser [70]. The liquid

crystal switch is simply a waveplate whose birefringence can be near zero or $\lambda/2$ depending on the electric field applied to the crystal. A polarizing beam-splitting cube is used right after the liquid crystal switch and so depending on the polarization of the laser light hitting the cube, the laser light will be directed through the optics that shape the MOT beams or through another path to form the push beam. Changing the state of the switch has the dual effect of turning off the collection MOT and turning on the push beam. At the same time the state of the switch is changed, the frequency of the trapping laser is changed from a frequency good for loading to one closer to resonance (6.8 MHz detuned), which is good for pushing. Simply changing the set point used in the DAVLL system changed the frequency of the laser.

There is a slight gloss in the previous subsection that can be pointed out now. The collection MOT is not just a three-beam retro-reflected MOT; it has one other beam in addition. Even when the liquid crystal switch is set to the MOT state, there is some leakage light along the push beam path. This leakage light is part of the collection MOT as well, and it tends to move the center of the MOT towards the transfer tube. We have found that having this extra beam can be beneficial, and so we have made no efforts to eliminate it.

The push beam is aligned through one of the ports in the collection chamber so that it passes through the transfer tube. Before entering the chamber, the linearly polarized push beam goes through a $\lambda/4$ waveplate. This is because the push beam actually has two effects. One is to accelerate the atoms out of the MOT and down the transfer tube, and the other is to spin-polarize the atoms. The atoms will scatter photons from the push beam and accelerate until they either make an off-resonant transition and wind up in the lower hyperfine state or get Doppler shifted out of resonance, reaching speeds measured at 12 m/s. Since the transfer tube is long and thin, many of the atoms accelerated out of the MOT would normally strike the wall of the transfer tube and get stuck there. To prevent this from happening, magnetic guiding is used. The circular polarization of the push beam spin-polarizes the atoms so that their magnetic moments are aligned in a direction so that they can be magnetically guided through the transfer tube.

The magnetic guiding is accomplished through a hexapole field created by a set of six permanent magnets that are strung down the transfer tube. The shape of the field produced is shown

in Figure 4.8. Magnets made of rubber impregnated with rare earth ferromagnetic materials (refrigerator magnet material) are used for the permanent magnets. This material has the advantage that it is cheap and easy to work with. The size of the field at the surface of the magnets that we used is about 600 G, which is strong enough to stop an atom with a velocity of a few meters per second.

To optimize the loading of the science MOT we began by achieving the most efficient transfer of atoms from the collection MOT to science MOT, performed by maximizing the number of atoms in single load pushed down the transfer tube as measured by the fluorescence of this single load in the science MOT. Geometrical considerations were useful in getting close to the optimum adjustment of all the collection MOT parameters. The collection MOT beams were centered on the collection chamber windows and the retro-reflected beams were nearly overlapped with the incoming beams. Exact overlap resulted in optical feedback that reduced the laser's frequency stability. After this initial alignment, the alignment of the MOT beams and the shim coils were used to position the center of the collection MOT right above the transfer tube. The push beam was then aligned so that it passed directly down the thin transfer tube. Once this was done, the alignment of the push beam, alignment of the MOT beams, and the shim fields were walked until the size of a single pushed load was maximized. The frequencies of the trapping laser used for the MOT and the push were also independently optimized, as were the orientations of the push waveplate and the liquid crystal shutter. After this optimization, it was sometimes necessary to deflect the push beam slightly if it was going straight down the transfer tube and striking the science cell MOT directly. This slight offset in angle had a negligible ($< 3\%$) effect on the number of atoms pushed down the transfer tube in a single shot.

The efficiency of transfer for a single load of atoms is estimated to be 15% with an uncertainty of around a factor of two due to the relatively poor calibration of the number of atoms in the collection MOT. This push transfer efficiency was found to be sensitive to the distance between the end of the hexapole guiding magnets and the center of the science cell MOT. Shortening that distance from 16 cm to 7 cm resulted a factor of two increase in push transfer efficiency. The extension of permanent magnets towards the science cell has some disadvantages, however. The

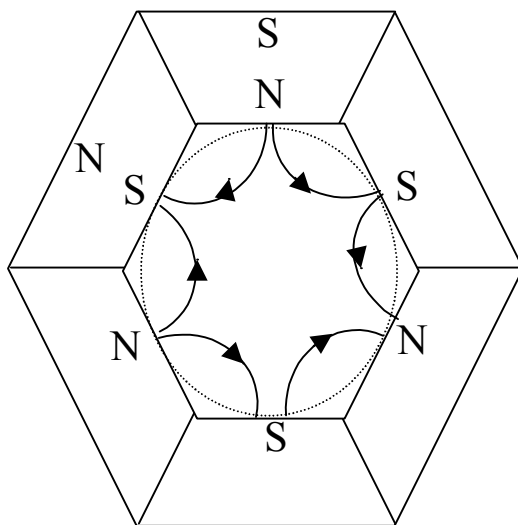


Figure 4.8 Hexapole Guiding Fields. Six strips of permanent magnets oriented as shown are sent along the outside of the transfer tube (shown as the dotted circle). The fields from all the magnets cancel at the center of the guide and rise steeply near the edges. The lines with arrows show the general direction of the magnetic field.

fringing fields from the permanent magnets can disrupt polarization-gradient cooling, hampering the cooling efficiency of the science cell MOT. More importantly, these fringing fields affect the magnetic trap, and the fringing fields from the permanent magnets are not stable at the required mG level. This led to complications in the evaporation described in detail in Chapter VII. The current distance of 7 cm represents a compromise between the push transfer efficiency and the mitigation of these other effects. One way to avoid this trade-off would be to use electromagnets for the guiding, but to generate fields of 600G currents of ~ 50 amps through coils with 75 turns and therefore a fair amount of resistance would be required, and the additional complexity involved was deemed to not be worth the gain.

The real figure of merit is not actually the single push transfer efficiency but the push transfer rate, the number of atoms that move from the collection to science chamber per unit time. The push transfer rate is a product of the collection MOT loading rate R and the push transfer efficiency. Currently, we have a rate of about 3×10^7 atoms/sec. We have spent some time optimizing the push transfer efficiency and cannot seem to improve it past 15%. The MOT loading rate can be increased through increasing the Rb vapor pressure, with the consequences to the vacuum discussed in subsection 4.4.3, or through increasing the amount of laser power available for the trapping beams. This last option would involve increasing the complexity of the laser system used but could result in a significant improvement in the experiment duty cycle. One potential drawback of more power is the sensitivity of the liquid crystal to high intensities--- the push/trap switching scheme may need to be altered.

4.4.6 Science MOT. Unlike the collection MOT, the science MOT is a true six-beam MOT; no retro-reflections are used. The advantage of the six-beam configuration versus retro-reflected configurations is that when the number of atoms trapped in the MOT becomes large, the cloud becomes optically thick and the trapping beam laser light is attenuated. With retro-reflection, this leads to an imbalance in power between trapping beams. The laser that provides the trapping light is again an external-cavity stabilized diode laser in the Littrow configuration. The model of this laser is a 50 mW Spectra Diode Labs # SDL-5400-G1 laser diode producing 11mW useful output power. This particular laser has been remarkably good, performing for over seven years. Its

frequency is controlled using a DAVLL system similar to the collection MOT trapping laser. The trapping beams are expanded and collimated to 1-inch diameter by a lens pair. A series of polarizing and non-polarizing beam splitters are used to divide the trapping light into six beams of approximately the same power. Again, $\lambda/4$ waveplates were used to set the polarization of the trapping beams. The hyperfine repump light comes from the same laser as the one used for the collection MOT. The main science MOT hyperfine repump beam (0.4 mW) is combined with the trapping light by a polarizing beam-splitting cube. In addition to this beam, light is also picked off of the upper hyperfine light path (60 μ W) and sent to the science MOT as well. This additional beam is referred to as the auxiliary hyperfine pump.

The frequency of the science MOT trapping laser is manipulated often during the preparation of the gas for loading into the purely magnetic trap, and the wide frequency range of the DAVLL system makes these changes in frequency particularly simple. Rather than using AOMs, it is possible to change the frequency of the science MOT trapping laser over a range of more than 100 MHz simply by changing the frequency lock servo set point. Since there is feedback to the diode laser current, changes in the set point then change the laser frequency in less than 100 μ s. The main disadvantage to the DAVLL system is that unlike a peak lock, small drifts can occur in the DC level of the DAVLL signal and the frequency of the laser will then drift as well. These drifts are particularly large when the room temperature is changing rapidly. In order to eliminate the effects of drifts, the frequency of the science MOT laser is compared to a laser locked to saturated absorption peak, which is less susceptible to drifts, by performing the following procedure. The hyperfine pump laser is unlocked from the normal repump transition and then locked to the crossover saturated absorption peak approximately 60 MHz away from the cycling transition. The hyperfine pump light is then beat with the science MOT trapping light on a fast photodiode (New Focus DC-125 MHz Low Noise Photoreceiver), and the beat frequency then gives the frequency of the trapping laser, allowing for the correction of any drifts that may have occurred.

A pair of anti-Helmholtz coils provides the magnetic field gradient for the MOT. Since it is more important that the purely magnetic trap coils are closer to the cell than the MOT coils, the anti-Helmholtz coils for the science MOT are outside the baseball and bias coils. The two anti-

Helmholtz coils are each 12 cm in diameter, 1 cm wide, have 15 turns of magnet wire each and are located 3 cm from the center of the science cell. The coils are wired in series and a 30A power supply is used to supply the current. For normal loading of atoms transferred from the collection MOT, a current of 17 A implying a gradient of 7.6 G/cm is used. While the coils only get warm to the touch during normal operation without cooling, it was discovered that the heat being produced by the anti-Helmholtz coils was causing a drift in the very sensitive evaporative cooling due to changes in temperature of the nearby magnetic trapping coils. Air-cooling the coils with a fan removed the heat they generated very efficiently and solved this problem.

In addition to the anti-Helmholtz coils, there are three shim coils, all aligned in orthogonal directions as in the collection MOT. These shim coils are similar in construction to the collection MOT shim coils, except that their current control is more sophisticated as will be discussed below. Each shim is actually a double set of coil. The second set of shim coils, referred to as earth shims, are also made from magnet wire, but they consist of less turns of wire and their current is controlled independently from the three main shim coils. In fact, the current in the earth shims is never turned off during the course of the experiment. They are present to null out any ambient fields, such as the earth's magnetic field. The current through the earth shims is not servoed, but large resistors placed far from the main apparatus are used to limit the current and protect against current changes due to changing temperatures of the earth shims. There is also one other coil, the optical pumping (OP) coil. This coil is used to provide a quantization axis for the optical pumping and the absorption imaging.

The measurement of the science MOT fluorescence was a diagnostic used to determine the number of atoms in the MOT. A lens with a focal length of 6 cm was placed directly on the end of the observation port, approximately 7 cm from the center of the MOT. To provide a known collection solid angle, a mask with a 5/16-inch diameter hole centered on the lens was also included between the lens and the trapped atoms. After several mirror bounces and a couple of other lenses, the fluorescence collected by the lens from the atoms scattering photons from the MOT beams was focused onto a photodiode (see Figure 4.9). The focus of this system of lenses was checked by removing the photodiode and replacing it with a CCD camera. The CCD camera was moved until




Figure 4.9. Schematic Diagram of the MOT Recapture Fluorescence Collection Optics. The pair of mirror bounces is necessary to bring the fluorescence light out from near the crowded region next to the science cell to a place where it can be more easily measured. Not shown in this diagram is the (5/16-inch diameter) mask placed over the first lens to better define the solid angle, among other things.

the MOT was in focus, and then the photodiode was replaced at the same position as the CCD array. The number of atoms for small samples can be calculated using the formulas in Ref. [62] and by using the absorption imaging to calibrate the fluorescence signal. The calibration achieved by the direct calculation is 2.0×10^6 atoms per nanoamp of photodiode current, not terribly far off from the more trustworthy absorption imaging calibration 2.5×10^6 atoms per nanoamp, especially when the uncertainties in the determination of the solid angle and trapping laser intensity are considered. For large numbers of trapped atoms, the trapping beams become attenuated and it is difficult to determine the precise number from fluorescence. However, relative measurements still have some value and the number can at least be estimated from the calibrations presented above. Of course, dirt on any of the optics in the fluorescence lens system can reduce the signal size and it is important to keep the optics clean. After ignoring the cleaning of one particular mirror for several months, the fluorescence signal dropped by 10%.

In order to load as many atoms as possible in the magnetic trap, it is desirable that as many atoms as possible are loaded into the science MOT. The number of atoms in the science MOT has been observed to be limited by two separate effects. The MOT beams are not spatially filtered and have spatial structure that creates interference patterns in the MOT. Since the density of the atoms in the MOT is limited due to reradiation effects, adding more atoms makes the MOT physically larger. Once the MOT gets to be a certain spatial extent, atoms are able to find an “escape route” where the interference in the MOT beams has caused there to be a direction where there isn’t sufficient trapping or repump power and the atoms are able to leak out in that direction and escape the MOT. Usually this limitation can be removed by realigning the MOT beams slightly to change the interference pattern. The second limitation is due to the finite lifetime of atoms in the MOT. The lifetime of the atoms in the science MOT is determined partly by the background gas pressure in the science chamber and partly by the density-dependent loss mechanisms present in a MOT [61]. Since the science MOT loading rate is fixed at the push transfer rate, eventually the loss rate from the MOT will equal the loading rate and the number of atoms in the science MOT can no longer increase. In this case, the maximum number could be increased by improving the push transfer rate. Conservatively, the maximum number of atoms in the MOT depends on the alignment of the beams

but is usually about 7.5×10^8 , where this number was determined without taking into account the effects due to absorption of the trapping light by the large cloud. The detuning of the trapping laser from the cycling transition that is best for the push transfer is 10.0 MHz.

4.4.7 Control of Magnetic Fields, Laser Light, and Laser Frequency. In order to prepare the gas to be loaded into the magnetic trap from the science MOT, the anti-Helmholtz and various shim magnetic fields, laser light, laser detuning, and the magnetic trap itself all must be controlled with timing that is reproducible to better than the millisecond level. This same sort of control is necessary for the MOT recapture diagnostics, and even better timing is required for the absorption imaging and optical pumping laser pulses. The timing during the loading and imaging sequence is controlled by a computer using Labview (version 4.1) and an output board (National Instruments AT-DIO-32F) that has 32 TTL outputs. These TTL signals are then used to control the timing of the MOT and magnetic trap loading and imaging sequences. A GPIB board is also used to control a programmable function generator (SRS DS345) and a radio-frequency synthesizer (HP 8656B) that are used for controlling the magnetic trap and the forced evaporation respectively. The discussion of the magnetic trap control and radio-frequency evaporation will be deferred to subsections 4.8.4 and 4.9.3.

The current Labview-based timing system is adequate, but there are some obvious improvements that could be made. There are two timing modes that are used. In one mode, Labview is used to create a string of sequential configurations of all of the 32 TTL outputs. The state of the TTLs is then updated on a computer clock signal from this string of configurations every 0.1 milliseconds. Within this sequence, the timing is reproducible to about the 1 microsecond level. This mode is not practical for any time period longer than 10 seconds (100,000 updates) given the time that computer takes to create and load the sequence of the TTL configurations. For longer time periods, the built-in software delays included in Labview are used. These built-in delays are only reproducible at the millisecond level, and so they cannot be used for timing-sensitive applications. A clear improvement would be to program a timing board so that all of the individual updates were not necessary.

Shutters (Uniblitz LS series) are used to control the laser light being sent to the collection and science chambers. Since all but one of the laser beams are shaped using a pair of lenses, there is usually a focus of the laser beam where the shutter can be placed. The laser without a focus, the OP/probe laser, is coupled into an optical fiber and so its beam is always small and the shutter can be placed anywhere. This allows a rapid shut-off of the laser light (faster than 0.1 ms, as fast as 10 μ s for lasers with a smaller focus). Using the more precise Labview timing scheme, it was possible to achieve a timing precision of ~ 10 μ s using these shutters alone. In order to get this precision, a small focus of the beam was necessary and it was also necessary to move the shutter so that the beam passed near the edge of the aperture rather than through the center. We also found that it was necessary to tilt all of the shutters slightly so that no light was reflected back into the laser, causing frequency instability through optical feedback.

Two of the lasers have additional controls on their light level. The collection MOT trapping laser has the liquid crystal shutter described above. The laser used for optical pumping and probing (OP/probe), described below, requires very short pulses that cannot be achieved using a mechanical shutter. An acousto-optic modulator (AOM) is used both to allow these shorter pulses and to control the light intensity since that laser is used for two different functions. The timing for these pulses is controlled by using an SRS DG535 pulse generator that sends TTL signals to rf switches to turn the AOM on and off and control the intensity of the light pulse. Two rf switches in series are closed to allow the rf to the AOM to turn it on, while another switch sends the rf power through one attenuator for the optical pumping segment and a different attenuator for probing. The rf switches used to control the AOM are not perfect and so a mechanical shutter is used to back up the AOM control to prevent “leakage” light from the AOM.

Associated with many of the TTL controls is a “level box,” which is an LM399 voltage reference that is connected to a potentiometer that serves as a voltage divider. An analog switch controlled by a TTL signal is connected to the output so that the voltage is produced only when the TTL is high. The laser detunings are controlled by feeding signals from the level boxes to the servo set point. The magnetic fields are controlled in one of two ways. For the collection MOT fields and the OP coil, the amount of current is controlled by the power supply providing it, and the level box

signal is used to saturate the gate of a FET used as a switch. For the science MOT shims and anti-Helmholtz coil, the current in each coil is sent through a FET and a $0.5\text{-}\Omega$ resistor, which serves as a transducer. The FET gate voltage was controlled by a simple servo to match the $0.5\text{-}\Omega$ resistor drop to a reference signal produced by a level box. This arrangement allowed the fields to be adjusted to different values by summing different level boxes together to produce the reference voltage. The simple servos were put in primarily to simplify the control of the levels and to correct for DC temperature drifts that occur in the FETs; the AC noise on the coils was not large enough to be a concern.

4.4.8 CMOT and Molasses, Part I. The conditions for best capture of atoms in the science MOT are not the same as the conditions that produce the best loading into the magnetic trap. After the desired load size has been transferred from the collection MOT, the science MOT magnetic field gradient, hyperfine repump power, and laser detuning are modified to increase the density and decrease the temperature of the gas of atoms. In this section, the manipulations that were performed on our particular MOT are discussed, but the discussion of the choice of the optimum parameters is delayed until subsection 4.7.1, after the magnetic trap has been discussed. In general, it is desirable to increase the density in the MOT and decrease the temperature. The density increase is achieved in a stage referred to as the CMOT (compressed MOT) stage [71], while the temperature is decreased in an optical molasses stage through polarization-gradient cooling.

The primary limitations to the density are the losses in the MOT, which limit the number, and the reradiation force, which limits the size. Since decreasing the scattering of photons from the trapping laser will decrease the reradiation force, both increasing the trapping laser detuning from resonance (see equation (4.3)) and reducing the hyperfine repump power so that the atoms spend more time in the $F=2$ state accomplish a reduction in the reradiation force. Both of these changes come at the cost of the MOT continuing to contain the atoms in the cloud, since a significant increase in detuning or decrease in hyperfine repump power leads to a loss of atoms from the MOT. Still, it is possible to produce a transient increase in the density of the sample, and that is good enough in preparing the cloud for loading into the magnetic trap.

During the CMOT stage, the hyperfine repump intensity is reduced from $80 \mu\text{W}/\text{cm}^2$ to $12 \mu\text{W}/\text{cm}^2$. This is accomplished by shuttering closed the main science MOT hyperfine repump path, leaving only the auxiliary hyperfine pump beam shining on the cloud. The detuning of the trapping laser from the cycling transition is increased somewhat, from 10 MHz to 13 MHz. These conditions are maintained for 30 ms to allow the MOT to settle to its new, higher density. During that time, 10% of the atoms are lost from the cloud. The rms size of the cloud decreases to 75% of its pre-CMOT size, to an average density of $\sim 1 \times 10^{10} \text{ cm}^{-3}$.

In addition to changing the hyperfine repump power and the trapping laser detuning, the anti-Helmholtz magnetic field gradient was reduced as well, from 8 G/cm to 4 G/cm. This decreases the effective trap spring constant confining the atoms to the center of the MOT, but it also improves the region in which polarization-gradient cooling is effective [71]. We found that changing the magnetic field gradient only had a weak effect on the loading of the large science MOT samples over the range of 4 to 11 G/cm.

After the CMOT stage, the atoms are then cooled in an optical molasses. The anti-Helmholtz and shim fields were shut off and the detuning of the trap laser was changed to 56 MHz. This molasses cooling stage was necessary for two reasons. First, the temperature of the gas in the CMOT was greater than $120 \mu\text{K}$. Second, we found that shutting off magnetic fields while the trap laser light is present can produce a large amount of transient heating through magneto-optic forces. After 4 ms in molasses, the temperature of the gas was cooled to about $15 \mu\text{K}$. The disadvantage of this molasses stage is that the cloud expands in the absence of the MOT confining forces, reducing its density.

4.4.9 Optical Pumping. After the molasses stage, it is necessary to spin-polarize the sample into the desired spin state, the $F=2, m_f = -2$ magnetic sublevel. This involves two processes: pumping all of the atoms out of the $F=3$ hyperfine state, and pumping all of the atoms into the $F=2, m_f = -2$ sublevel. The optical pumping scheme and the levels involved are shown in Figure 4.10. No hyperfine repump light is allowed to shine on the atoms. The trapping laser detuning is shifted from the molasses detuning to 20 MHz to the blue of the $F=3$ to $F=3'$ transition. This serves to pump all of the atoms out of the $F=3$ hyperfine state.




Figure 4.10. Optical Pumping Scheme. The OP laser is represented by the solid black arrows and the trapping light used to pump atoms out of the $F=3$ state is represented by the dotted arrows. The detunings of the lasers from the excited hyperfine states is shown approximately and the $F'=4$ state is not shown. The goal of the pumping is to have all the atoms wind up in the $F=2, m_f = -2$ state.

The laser light used to polarize the spin of the atoms comes from the OP/probe laser. This laser is another external-cavity stabilized diode laser, model Sharp LT024BD0. Its frequency is stabilized using a DAVLL system. The output of the laser is coupled through an AOM, and the first-order diffraction beam is then coupled through a polarization-maintaining optical fiber. The optical fiber is used to increase the pointing stability of the laser beam and to ensure that the beam is of a good spatial quality. The AOM is used to control the OP/probe laser intensity and allow the creation of sub-millisecond pulses. The OP/probe laser is linearly polarized, and a polarizing beam-splitting cube is used to couple the probe laser along one of the science MOT beam paths. When the OP/probe light passes through the MOT $\lambda/4$ waveplate, it becomes circularly polarized.

After the detuning of the MOT trapping laser is switched, the OP coil is switched on in 0.2 ms to produce a 3 G magnetic field. The direction of the field is chosen so that the OP/probe laser has a σ^{\square} polarization. A pulse of $\sim 75 \mu\text{W}/\text{cm}^2$ OP light is turned on for 900 μs . This laser is detuned between the F=2 to F=1' and F=2 to F=2' transition, a few MHz from the F=2 to F=2' transition. The detuning was determined by optimizing the number of atoms that were pumped into the F=2 $m_f = -2$ state. A curve showing the pumping efficiency is shown in Figure 4.11. The theory of the optical pumping is that atoms would scatter σ^{\square} photons and move to more negative m_f states. Occasionally, they would spontaneously scatter to the F=3 state, and the trapping laser would pump them back to the F=2 state. Once the atoms were in the F=2, $m_f = -2$ state, they would no longer scatter any photons, since there is no transition for σ^{\square} light from this state. The atoms would all pile up in that state, and then they would be loaded into the magnetic trap.

The reality is that we have not been able to achieve more than 80% of the available atoms in the $m_f = -2$ state. The reasons for this are as of yet unknown [72]. I speculate that since the Clebsch-Gordan coefficients are such that for light tuned between the F=2 to F=1' and 2' transitions, it is difficult to pump the atoms from the positive m_f states to negative m_f states without going through the F=3 upper ground state. The depolarizing nature of the MOT light used to pump atoms from the F=3 back to the F=2 state may be a source of the limitation in our optical pumping. We do have some data indicating the nature of the problem. Figure 4.12 shows that while a higher probe

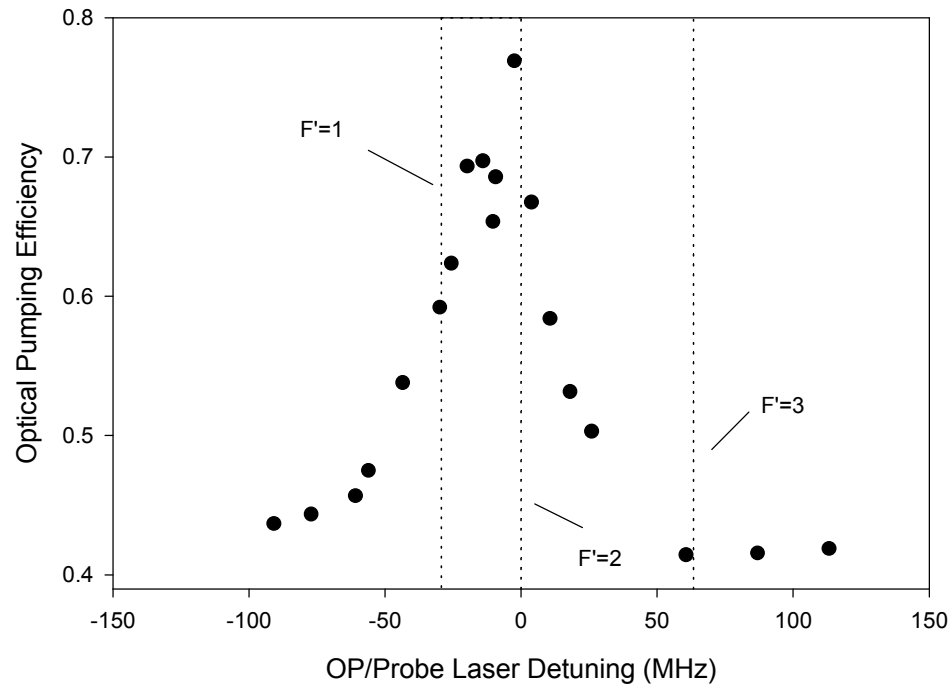


Figure 4.11. Optical Pumping Efficiency as a Function of OP/Probe Laser Detuning. The optical pumping efficiency is defined as the ratio of the number of atoms magnetically trapped to the number of atoms remaining after the molasses stage. The vertical dotted lines show the positions of the F=2 to excited state transitions. The relative positions of the transitions are accurate but the overall location could be off by a few MHz, due to some uncertainty in the absolute frequency calibration. The absence of any pumping near the F'=3 transition is due to heating of the atoms from repeated photon scatters in the F=2 $m_f = -2$ ground state since that state is no longer dark.

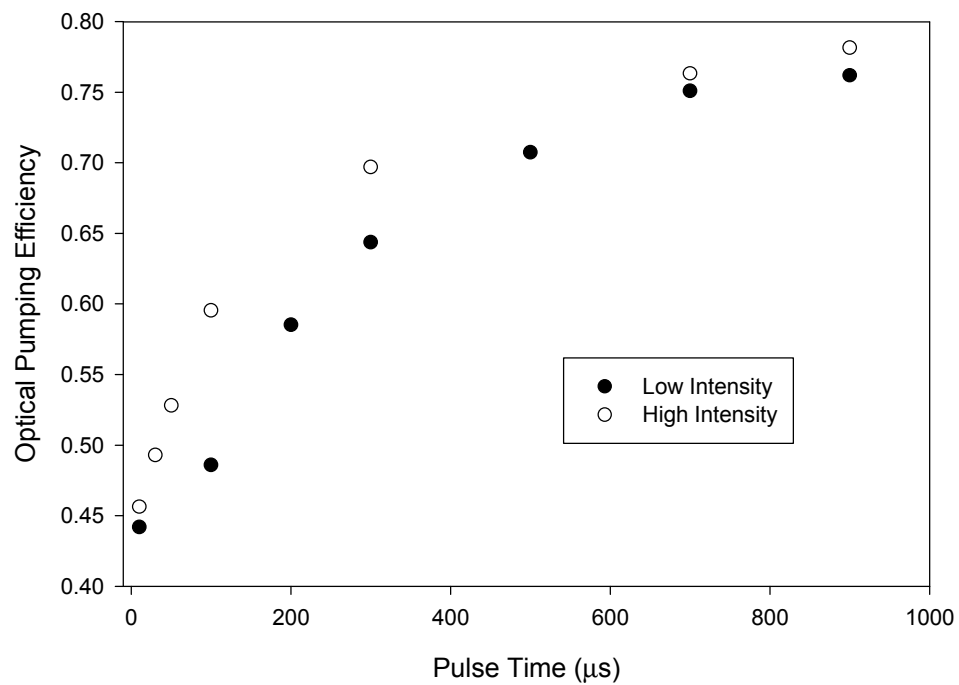


Figure 4.12. Optical Pumping Efficiency as a Function of Pulse Time. Even though the initial rise of the higher intensity data is faster, the cloud takes about the same time to optically pump for both intensities. This behavior is not yet explained. The high intensity is a factor of 8 more than the $\sim 75 \mu\text{W}/\text{cm}^2$ low intensity.

intensity pumps atoms into the $m_f = -2$ state faster initially, the rate levels off so that it takes the same time to pump the cloud regardless of intensity. When the OP/probe laser is detuned from the optimal detuning, the cloud also pumps up on the same time scale, and the ultimate efficiency is smaller. Varying the detuning of the trapping laser used to pump the atoms out of the $F=3$ state has no effect, even when tuned ± 60 MHz from the $F=3$ to $F'=3$ transition. Changing the OP coil magnetic field does not have much of an effect, either. It is possible to completely deplete the $m_f = -2$ state if the OP coil magnetic field is reversed, since reversing the magnetic field changes the OP laser light from σ^- to σ^+ . This maximum efficiency does not depend on the number of atoms in the cloud. The final fact is that the maximum efficiency of the optical pumping is determined by the initial MOT and molasses stages. The MOT trapping laser can be aligned and the molasses shims tweaked so that even without additional optical pumping light, the atoms tend to be aligned towards the negative m_f sublevels. Just pumping the atoms out of the $F=3$ state will then produce some atoms in the $F=2$, $m_f = -2$ state. We have observed that the OP light seems to only be able to double that number at best. This means that good optical pumping efficiency is achieved only for proper alignment of the MOT and molasses. The optimization of the OP process will be discussed in subsection 4.7.2 below.

In order to make sure that the atoms are pumped out of the $F=3$ state completely, the OP light is turned off 200 μs before the trapping laser light is. After the atoms have been optically pumped, they are ready to be loaded into the magnetic trap. Before discussing the specifics of our apparatus, a general description of magnetic trapping will be presented.

4.5 Magnetic Traps.

4.5.1 Principle of Operation. Magnetic traps rely on the magnetic moment of the atoms in order to confine them in a potential. As long as the energy of an atom increases with increasing magnetic field, like the $F=2$ $m_f = -2$ and -1 states (see Figure 4.3), it is possible to trap the atom at a minimum in the magnetic field. At first glance, it would seem to be impossible to form a magnetic field minimum in free space, due to Earnshaw's theorem in electrostatics [73]. This is true when the magnetic field is considered as a vector, but not when only the amplitude of the field is considered

(i.e. $\nabla \vec{B} \neq 0$ while $\vec{\nabla} |\vec{B}| = 0$). If the field is strong enough and the atom is not moving too fast, the magnetic moment of the atom will adiabatically follow the field and its Zeeman energy will depend only on the amplitude of the field, allowing weak-field seeking states to be trapped.

For sufficiently small fields, the atoms will no longer adiabatically follow the field direction, and then a so-called Majorana transition can occur and the atom will transfer to a different magnetic sublevel. The Larmor frequency of an atom in even a small magnetic field is on the order of MHz, and so only for very small magnetic fields (less than one gauss) will the atoms no longer adiabatically follow the field direction. While Majorana transitions were a limiting factor in early attempts in evaporative cooling [74], in our experiments the magnetic field amplitude is rarely less than 140G, which corresponds to a Larmor frequency of 123 MHz. In order for an atom moving with an energy of 10 μ K to have a reasonable probability of undergoing a Majorana transition, the field gradient would have to be well over 10^{11} G/cm, several orders of magnitude larger than what we can achieve in our magnetic traps.

4.5.2 Ioffe-Pritchard Traps. A diagram of a Ioffe-Pritchard trap [75] is shown in Figure 4.13. This trap can be used to confine atoms near its center, and it has the advantage that the minimum magnetic field in the center of the trap is non-zero. The magnetic field from the four current bars is simple to compute in the limit that they are infinitely long. When only the magnetic field from these current bars is considered (i.e. the Helmholtz coil pair is ignored), the magnetic field is zero along the z-axis. If an atom moves away from the center in any radial direction, the magnitude of the magnetic field will increase with distance, as the atom gets closer to one or two of the current bars and farther from the others. Therefore, the current bars provide a restoring force in the radial direction on an atom in a weak-field seeking state. The Helmholtz coil pair provides the axial confinement, because an atom moving away from the center will see a stronger magnetic field since it is closer to one of the Helmholtz coils than the other. In addition to the axial confinement, the Helmholtz pair also produces a gradient in the radial direction, but that will be assumed to be small for the sake of this discussion, and the Helmholtz coils will be assumed to produce only a field along the z-axis that is uniform over the radial directions. This assumption is quite good as long as the radius of the Helmholtz coils is large enough.

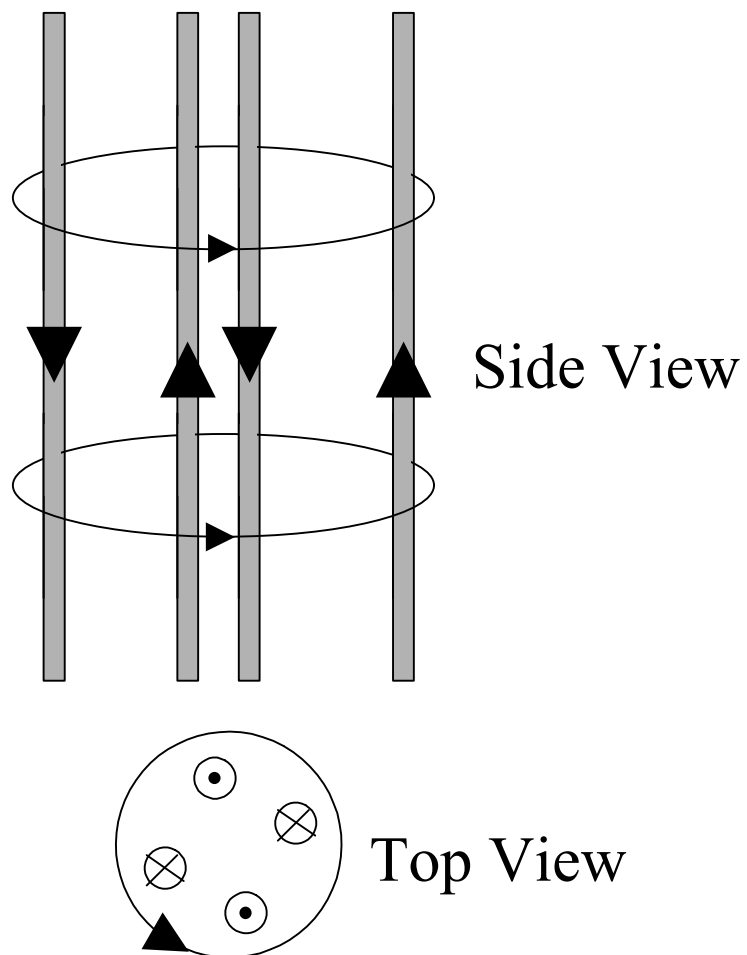


Figure 4.13. Ioffe-Pritchard Trap Schematic. The arrows indicate the direction of the current through the coils and wires.

With these simple considerations, it is possible to calculate the magnetic field near the center of the trap. If R is the distance of the current bars from the center of the trap, I_{cb} is the size of the current bar current, I_{hh} is the Helmholtz coil current, z_0 is the distance of each Helmholtz coil from the trap center, and a_0 is the radius of the Helmholtz coils, then the magnetic field is

$$\vec{B} = \begin{pmatrix} \frac{\mathbf{m}_0 I_{cb}}{2\mathbf{p}} \left(\frac{-y}{((x-R)^2 + y^2)} + \frac{(y-R)}{((y-R)^2 + x^2)} + \frac{-y}{((x+R)^2 + y^2)} + \frac{(y+R)}{((y+R)^2 + x^2)} \right) \hat{x} \\ \frac{\mathbf{m}_0 I_{cb}}{2\mathbf{p}} \left(\frac{(x-R)}{((x-R)^2 + y^2)} + \frac{-x}{((y-R)^2 + x^2)} + \frac{(x+R)}{((x+R)^2 + y^2)} + \frac{-x}{((y+R)^2 + x^2)} \right) \hat{y} \\ \frac{\mathbf{m}_0 a_0^2 I_{hh}}{2} \left\{ \left((z_0 + z)^2 + a_0^2 \right)^{\frac{3}{2}} + \left((z_0 - z)^2 + a_0^2 \right)^{\frac{3}{2}} \right\} \hat{z} \end{pmatrix} \quad (4.7)$$

Again, this expression is not exact, but is it valid in the limit that the changes in the magnetic field in the radial direction due to the current bars are much greater than those due to the Helmholtz coils. For the baseball coil that we use to create our magnetic trap, equation (4.7) does not do an accurate job of describing the magnetic fields, but it illustrates the principle of operation of a Ioffe-Pritchard-type magnetic trap. In our magnetic trap, the axial and radial confinement can both be provided by the baseball coil alone, so a separate set of Helmholtz coils is not necessary. This is quite different from the current bars and Helmholtz coils configuration. However, the baseball coil does produce a magnetic field that is directed along a single axis, much like the Helmholtz coils modeled in equation (4.7), and radial confinement by the baseball coil is provided in much the same way as the current bars do. Therefore, the magnetic field produced by the baseball coils is qualitatively similar to that in equation (4.7), and so that equation is useful for describing the principles of the operation of our magnetic trap. From the expression for \mathbf{B} in equation (4.7), the absolute value of the magnetic field can be calculated. Keeping only the leading terms in an expansion around the trap center $x=y=z=0$, the magnetic field can be written

$$B \cong B_0 + \frac{4\mathbf{m}_0 I_{cb}^2}{\mathbf{p}^2 R^4 B_0} r^2 + \frac{B_0(12z_0^2 - 3a_0^2)}{(z_0^2 + a_0^2)^2} z^2, \quad B_0 = \frac{\mathbf{m}_0 a_0^2 I_{hh}}{(z_0^2 + a_0^2)^{\frac{3}{2}}} \quad (4.8)$$

where $r^2 = x^2 + y^2$. In the following discussion, it will be assumed that $z_0 > a_0/2$ so that potential is confining in the z direction. The energy of an atom in this magnetic field can be written $E = \mu_{\text{mr}} B$,

where μ_{mf} is the magnetic moment of the atom's magnetic sublevel. The following relation then holds

$$E = \mathbf{m}_{mf} \cdot \mathbf{B} = \mathbf{m}_{mf} B_0 + \frac{4\mathbf{m}_0 I_{cb}^2 \mathbf{m}_{mf}}{\mathbf{p}^2 R^4 B_0} r^2 + \frac{B_0(12z_0^2 - 3a_0^2) \mathbf{m}_{mf}}{(z_0^2 + a_0^2)^2} z^2 = E_0 + \frac{m}{2} (\mathbf{w}_r^2 r^2 + \mathbf{w}_z^2 z^2) \quad (4.9)$$

$$\mathbf{w}_r = \sqrt{\frac{\mathbf{m}_{mf} 8\mathbf{m}_0}{m\mathbf{p}^2}} \frac{I_{cb}}{R^2 \sqrt{B_0}} \quad \mathbf{w}_z = \sqrt{\frac{\mathbf{m}_{mf} a_0 \mathbf{m}_0 (24z_0^2 - 6a_0^2)}{m(z_0^2 + a_0^2)^2}} \sqrt{I_{hh}}$$

and it can be seen that around the trap center, the magnetic field produces a harmonic trapping potential. Since we do not use four current bars plus a Helmholtz coil pair to provide our trapping field, equation (4.9) is not particularly useful for calculating the frequencies in our magnetic trap. However, the radial (ω_r) and axial (ω_z) frequencies in our trap scale with current, bias magnetic field, and radius of the trapping coils in the same way as equation (4.9). The radial frequency is proportional to the current and inversely proportional to the square root of the bias magnetic field at the bottom of the trap. The dependence on the square root of the bias field comes about from the fact that the magnetic field from the current bars adds to the orthogonal bias magnetic field in quadrature. The larger the bias field, the smaller the fractional change in B due to the current bar magnetic fields. The axial frequency is proportional to the square root of the current. The other scaling that is notable is the dependence of the trapping frequencies on the distance of the current bars and coils from the center of the trap. Both the radial and the axial frequencies increase rapidly as the current bars and Helmholtz coils are brought closer to the center of the trap. In order to use smaller power supplies, it is then desirable to make the coils as small as possible.

The magnetic trapping potential from a Ioffe-Pritchard trap is not strictly harmonic, and so there are anharmonic terms present as well. The importance of the anharmonic terms is determined by the trapping region sampled by the atoms in the trapped cloud, and for cold enough clouds (see subsection 4.5.3), these terms are small and the magnetic trap is well approximated by a purely harmonic potential. Some of the qualitative effects of these anharmonic terms will be discussed in the next section.

4.5.3 Baseball Trap. The magnetic trap that we use is a modified Ioffe-Pritchard design referred to as a “baseball” trap. Figure 4.14 shows the baseball coil geometry. The baseball coil can




Figure 4.14. Schematic of Baseball coil. The shape of the baseball coil used for magnetic trapping is shown here along with the two bias coils. The small gray circle indicates the minimum of the magnetic field. The resemblance between the baseball coil and the current bars and Helmholtz coil system described in Figure 4.13 can be seen by considering the direction of the current flowing in each part of the baseball coil. The current flowing along the parts labeled “b” in corresponds to the current bars, while the current flowing along the paths marked “a” approximates the Helmholtz coils. be wrapped from a single wire, making its construction relatively simple. The baseball trap gets its name because the coil has the same shape as the seams of a baseball. In addition to the baseball coil, there is a set of Helmholtz coils, referred to as the bias coils, that are used to control the bias magnetic field in the trap. When current is flowing just through the baseball coil, both radial and axial confinement are provided. The typical bias magnetic field produced by the baseball coil itself can be quite large, on the order of several hundred gauss. In order to get stronger confinement in the radial direction, the bias coils are set up to oppose the baseball coil field bias field.

The baseball coil that we use is composed of 16 turns of hollow 1/8 inch copper refrigerator tubing. The tubing is sheathed in heat-shrink insulation, which was responsible for problems with baking out the chamber (section 4.4.3). All the voids between the wires have been filled with epoxy for mechanical stability. The coil fit snugly over the glass science cell, and the coil was also (probably unwisely) epoxied to the cell as well. The coil geometry allows for access for all of the MOT laser beams. The hollow tubing conveniently allows the water-cooling of the coil, which is necessary given the heat generated by the large (~ 200 A) currents running through it. A booster pump is used to push 1.2 L/min. through the coils at 240 psi. In the course of typical operation, the temperature of the coils would only rise a few $^{\circ}\text{C}$ due to the 200 A of current running through them. The total electrical resistance of the baseball coil was measured to be $0.019\ \Omega$, and the inductance was measured to be $25\ \mu\text{H}$. This inductance was determined by varying the voltage drop across the baseball coil sinusoidally and determining the frequency where the current to voltage ratio changed by a factor of $(2)^{1/2}$.

The bias coils were each also constructed out of 8 turns of the same copper tubing, sheathed in the same insulating material as the baseball coils. The mean radius of these coils is 4.5 cm, and they are each placed 4 cm apart, roughly centered on the center of the baseball coils. Again, epoxy was used to fill all the gaps in the coils for stability. The electrical resistance of the pair of coils is $0.016\ \Omega$. The bias coils are also water cooled, with the water flowing in parallel to the baseball coil for greater cooling efficiency.

Several changes have been made to the coil system since its use in Ref. [1], but these changes will be detailed below in subsection 4.8.3.

4.5.4 Gravitational Sag and Typical Trapping Potentials. In order to calculate the trapping potentials produced by the baseball and bias coils, numerical techniques had to be used. The reason that a small expansion around the field at the center of the coils (i.e. the point of minimum B) cannot be used is that in our case the trapped atoms are not near the center of the coils. In the absence of gravity, the atoms would be confined to the center of the coils, but the relatively heavy ^{85}Rb atoms also respond to gravity, and so they are pulled down from the center of the trap,

typically by about 2 mm. In order to combat this gravitational “sag,” the trap is oriented with one of its strong radial directions opposing gravity (i.e. the axial direction is perpendicular to down). For stronger radial confinement, sag would be less of an issue, but the magnetic fields of interest to us limit the radial confinement possible.

For purely harmonic potentials, the sag is easily calculated to be g/ω_r^2 . The curvature of the field, and hence the trapping frequency, would not change regardless how far from the center of the coils the atoms were. The actual baseball/bias coil trap potentials are not strictly harmonic, but a minimum in the total energy of the atoms (magnetic plus gravity) can still exist. It is always possible to expand the potential around that minimum, and the leading terms will be quadratic in the coordinates and so the potential can still be approximated as harmonic. Anharmonic terms will also be present, but the smaller the spatial size of the trapped cloud (and correspondingly the lower the temperature) the less important these terms will be.

In order to numerically calculate the precise trapping potential, the baseball coil was modeled as a set of four semi-circular wires. The bias coils were also included and the current through both the baseball and bias coils was then specified. Each component of the magnetic field was calculated at a particular point in space by using the Biot-Savart law and numerically integrating along each of the current paths. The contributions from all of the currents were then added together, and the absolute value of the magnetic field was calculated. In this way, the absolute value of the magnetic field was computed at all the desired points in space. The Breit-Rabi equation (4.1) was then used to determine the magnetic energy of the atoms, and the gravitational energy was also included. Figure 4.15 shows a typical potential for $F=2$, $m_f = -2$ calculated this way, and also shows the magnetic field across the same region in space.

One of the features in Figure 4.15 is that the magnetic field lines cut horizontally through the trapping region, so that in the region of the trapped atoms, the magnetic field seems to vary mostly along the vertical direction. This is a consequence of the relatively weak trap. By calculating the magnetic moment of the $F=2$, $m_f = -2$ ^{85}Rb atoms at 160 G (the center of the bias field in Figure 4.15), the magnetic field must have a linear gradient of about 26 G/cm to support the atoms against gravity. The curvature of the magnetic field is 155 G/cm². A 10 μK cloud would

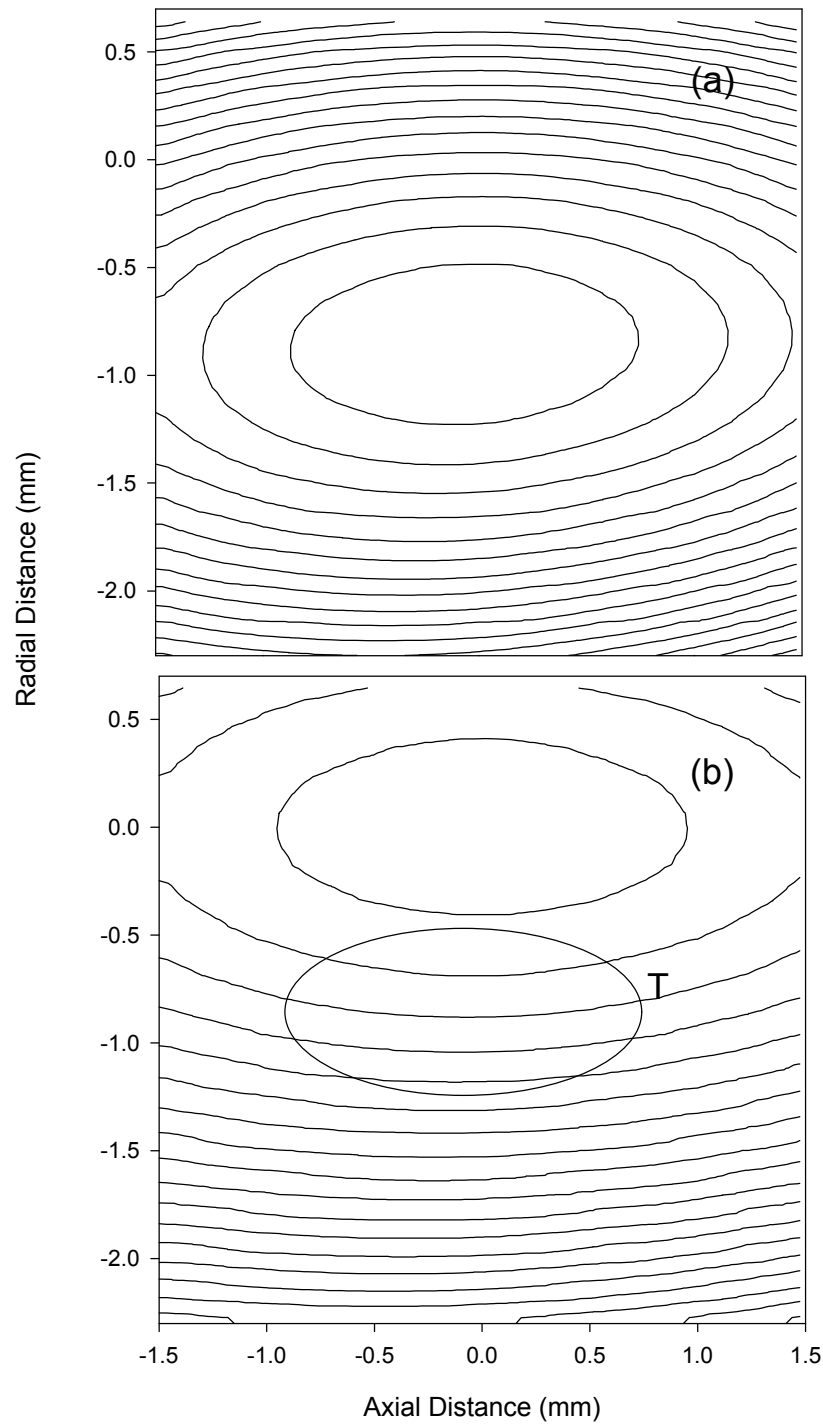


Figure 4.15. Contour plots of a typical trapping potential (a) and its associated magnetic field (b). The contour lines are spaced $10 \mu\text{K}$ and 0.5 G apart respectively. The oval labeled “T” shows the location of the interior trapping contour on the magnetic field plot. These plots were calculated using the numerical method discussed in the text. For this case, the radial trapping frequency is $2\pi \times 17.1 \text{ Hz}$ and the axial is $2\pi \times 7.6 \text{ Hz}$.

have an rms size of 0.029 cm in this trap, meaning that the curvature accounts for a spread in magnetic field, defined in this case (and differently elsewhere in the thesis) as $(\langle B^2 \rangle - \langle B \rangle^2)^{1/2}$, across the cloud of 0.13 G while the gradient accounts for 0.73 G. These two quantities are added in quadrature to get the total spread. Therefore, the variation in magnetic field across the cloud is due almost entirely to the linear gradient that is offsetting gravity and not to the curvature of the trap. This means that the spread in the magnetic field across the cloud can then be computed fairly accurately by simply multiplying the width of the cloud by 26 G/cm.

The calculation of the trapping potentials matched the measurement of the actual trapping potentials such that the trapping frequencies matched only to the 10% level. The reasons for this are several: the baseball and bias coils are composed of several turns of wire and therefore have a finite thickness that was not included in the calculation, the size, spacing, and location of the coils are not known to better than 5%, the baseball coil does not have the exact shape of six semi-circles, the bias coils are not mounted perfectly perpendicular to the baseball coil axis, the real coils have leads that are not included in the calculation, and the bias coils are not exactly centered on the baseball trap axis. In principle, these position and angle offsets along with the thickness of the wires could be included in the numerical calculation and measurements of the actual trap frequencies could be used to constrain all of these additional parameters. However, no need to do so has arisen.

4.6 Classical Motion of Atoms in a Magnetic Trap.

4.6.1 Classical Orbits. In a purely parabolic potential in the absence of elastic collisions between atoms, each atom in the trapped sample exhibits simple harmonic oscillator motion in each of the three directions. Since the potential is parabolic, the motion in the three directions is separable. As a consequence of the simple harmonic motion, the radial (axial) position of each individual atom at time intervals spaced by one radial (axial) trap period will be the same.

In reality, the potential is not purely parabolic and one effect of the anharmonic terms is to break the separability of the three spatial dimensions. Motion in one direction will be slowly converted to motion in an orthogonal direction by some of the anharmonic terms. In other words, the anharmonicities will cause the orbits of the atoms to slowly precess. The precession rate is

determined by the difference in the trapping frequencies between the two directions being considered. The magnitude of the precession, how far the orbit of the atom tilts from the original direction, is in general quite small unless the difference in trap frequencies is also small.

Another effect of the anharmonic terms is that the oscillation frequency depends on the size of the atom's orbit, which is important in considering any center-of-mass motion of the atom cloud. The anharmonicities cause any of center-of-mass motion of the cloud to be effectively converted into heat. The details of this conversion will be discussed below after the collective motion of the atoms in the cloud has been discussed.

The presence of elastic collisions in the trapped cloud of course means that individual atoms will have their orbits changed as the result of a collision event. One effect of these elastic collisions is to randomize the velocity distribution of atoms in the trapped cloud. After many collisions, the cloud must then be driven to a state of maximum entropy by the elastic collisions, and so their presence eventually brings the cloud into thermal equilibrium. This thermalizing effect of elastic collisions is treated in detail in Chapter V, but thermal equilibrium will be discussed in the following subsection.

4.6.2 The Boltzmann Distribution Function and Thermal Equilibrium. In order to describe the collective behavior of a gas, the Boltzmann distribution function is used. As its name suggests, the Boltzmann distribution function, usually denoted $f(\mathbf{x}, \mathbf{v}, t)$, specifies the distribution of velocities and positions of the atoms in the gas. The number of atoms in a volume d^3x around position \mathbf{x} with a velocity range d^3v around velocity \mathbf{v} at time t is $f(\mathbf{x}, \mathbf{v}, t)$. The Boltzmann distribution function is therefore a distribution function in phase-space, and in fact the phase-space density of the sample is directly related to $f(\mathbf{x}, \mathbf{v}, t)$ [33]. From the $f(\mathbf{x}, \mathbf{v}, t)$ it is possible to calculate all the standard thermodynamic functions, such as the number, energy and entropy of the gas. For instance, the number of the atoms in the gas is simply

$$N = \int d^3x d^3v f(\vec{x}, \vec{v}, t). \quad (4.10)$$

The distribution function of a classical gas in thermal equilibrium is determined by requiring that the occupation fraction of each orbit with energy E is $\exp(-(E-\mu)/k_B T)$ [33]. The temperature (T) is related to the total energy of the gas, and the chemical potential μ is adjusted so

that the equation (4.10) is satisfied. For an axially symmetric potential, the Boltzmann distribution function is

$$f(\vec{x}, \vec{v}) = \frac{Nm^3 \omega_r^2 \omega_z}{8\mathbf{p}^3 k_B^3 T^3} \exp\left(-\frac{mv^2}{2k_B T} - \frac{m}{2k_B T} (\omega_r^2 r^2 + \omega_z^2 z^2)\right) \quad (4.11)$$

where N is the number of atoms, ω_r and ω_z are the radial and axial trap frequencies (in radians/sec), and m is the mass of an atom. In order to find the density of atoms as a function of position $n(\mathbf{x})$, it is only necessary to integrate $f(\mathbf{x}, \mathbf{v})$ over all possible velocities

$$n(\vec{x}) = N \omega_r^2 \omega_z \left(\frac{m}{2\mathbf{p}k_B T}\right)^{\frac{3}{2}} \exp\left(-\frac{m}{2k_B T} (\omega_r^2 r^2 + \omega_z^2 z^2)\right). \quad (4.12)$$

This can also be written as

$$n(\vec{x}) = \frac{N}{(2\mathbf{p})^{3/2} \mathbf{s}_r^2 \mathbf{s}_z} \exp\left(-\frac{r^2}{2\mathbf{s}_r^2} - \frac{z^2}{2\mathbf{s}_z^2}\right) \quad \mathbf{s}_i = \frac{1}{\omega_i} \sqrt{\frac{k_B T}{m}} \quad i = r, z. \quad (4.13)$$

The prefactor in equation (4.13) is the peak density at the center of the trap potential. Equation (4.13) shows that the spatial size of the trapped cloud is proportional to the square root of its temperature. The colder the cloud, the smaller it is. Once the density as a function of position is known, the average density can also be computed

$$\begin{aligned} \langle n \rangle &= \frac{1}{N} \int d^3x n(\vec{x})^2 \\ &= \frac{N}{8\mathbf{p}^{3/2} \mathbf{s}_r^2 \mathbf{s}_z} \end{aligned} \quad (4.14)$$

Other quantities of interest are the total energy

$$E_{tot} = \int d^3x d^3v \left(\frac{1}{2} mv^2 + \frac{1}{2} m(\omega_r^2 r^2 + \omega_z^2 z^2) \right) f(\vec{x}, \vec{v}) = 3Nk_B T \quad (4.15)$$

the rms velocity

$$v_{rms} = \sqrt{\frac{1}{N} \int d^3x d^3v (v^2) f(\vec{x}, \vec{v})} = \sqrt{\frac{k_B T}{m}} \quad (4.16)$$

and the average relative speed

$$\langle v \rangle = \int d^3x \int d^3v d^3v' |\vec{v} - \vec{v}'| f(\vec{x}, \vec{v}) f(\vec{x}, \vec{v}') = \sqrt{\frac{16k_B T}{\rho m}}. \quad (4.17)$$

Again, all of these quantities are computed in thermal equilibrium.

4.6.3 Collective motions of the cloud in the absence of collisions. The time evolution of the Boltzmann distribution function is controlled by the Boltzmann equation [76]

$$\frac{\partial f}{\partial t} = - \left(\vec{v} \cdot \vec{\nabla}_x + \frac{\vec{F}}{m} \cdot \vec{\nabla}_v \right) f + J(f) \quad (4.18)$$

where $\vec{\nabla}_x$ indicates the gradient of $f(\mathbf{x}, \mathbf{v}, t)$ with respect the position coordinates and $\vec{\nabla}_v$ is the gradient with respect to the velocity coordinates. The last term accounts for the net flux into a phase-space coordinate due to collisions. Collisions are not treated in detail in this section, and this topic will be discussed in more detail in Chapter V.

In order to determine the time evolution of a distribution function, it is necessary to solve the differential equation (4.18) with an initial distribution function $f(\mathbf{x}_0, \mathbf{v}_0)$ establishing the boundary conditions. One simple way that this can be done in the absence of collisions is to compute the individual atom motion $\mathbf{x}(t)$ and $\mathbf{v}(t)$ with \mathbf{x}_0 and \mathbf{v}_0 as an initial condition. Once the motion as a function of time is known, it is possible to invert the equations, solve for \mathbf{x}_0 and \mathbf{v}_0 in terms of $\mathbf{x}(t)$ and $\mathbf{v}(t)$ and then plug these expressions back into the initial distribution function in order to compute $f(\mathbf{x}, \mathbf{v}, t)$. This method will be used to determine three collective classical motions of the trapped gas: ballistic expansion, sloshing, and breathing.

4.6.4 Ballistic Expansion. It is often desirable to release the gas from the magnetic trap and let it expand. This has the effect of reducing the optical thickness of the cloud and increasing the spatial size of the cloud for imaging. It also can be used to measure the energy of the sample. In addition, it may be impossible to image the cloud immediately after turning off the magnetic trap, and so the evolution of the cloud outside the trap needs to be calculated to accurately determine the size of the cloud in the magnetic trap. Once the trap is turned off, the cloud will expand due to the kinetic energy of its particles. Of course, without the trapping potential, the cloud of atoms will also fall under the influence of gravity.

For a particle with an initial position \mathbf{x}_0 and velocity \mathbf{v}_0 , the position and velocity as a function of time outside the magnetic trap will be $\mathbf{x}(t)=(x_0+v_{x0}t, y_0+v_{y0}t - \frac{1}{2}gt^2, z_0+v_{z0}t)$ and $\mathbf{v}(t)=(v_{x0}, v_{y0} - gt, v_{z0})$ where down is in the direction of the negative y axis. Inverting these equations it is possible to solve for \mathbf{x}_0 and \mathbf{v}_0

$$\begin{aligned}\vec{x}_0 &= \left(x - v_{x0}t, y - v_{y0}t + \frac{1}{2}gt^2, z - v_{z0}t \right) \\ \vec{v}_0 &= (v_x, v_y + gt, v_z)\end{aligned}\quad (4.19)$$

Plugging these expressions into $f(\mathbf{x}_0, \mathbf{v}_0)$

$$f(\vec{x}, \vec{v}, t) = \frac{N}{8p^3 s_r^2 s_z v_{rms}^3} \exp \left(-\frac{(x - v_x t)^2}{2s_r^2} - \frac{(y - (v_y + gt)t + \frac{1}{2}gt^2)^2}{2s_r^2} - \frac{(z - v_z t)^2}{2s_z^2} \right. \\ \left. - \frac{v_x^2}{2v_{rms}^2} - \frac{(v_y + gt)^2}{2v_{rms}^2} - \frac{v_z^2}{2v_{rms}^2} \right) \quad (4.20)$$

Of particular interest is the evolution of the density over time. This is determined by integrating equation (4.20) over velocity

$$n(\vec{x}, t) = \frac{N}{(2p)^{3/2} s_r(t)^2 s_z(t)} \exp \left(-\frac{x^2}{2s_r(t)^2} - \frac{(y + \frac{1}{2}gt^2)^2}{2s_r(t)^2} - \frac{z^2}{2s_z(t)^2} \right) \quad (4.21)$$

$$s_r(t)^2 = s_r^2(1 + w_r^2 t^2), \quad s_z(t)^2 = s_z^2(1 + w_z^2 t^2)$$

The two widths σ_r and σ_z are defined in equation (4.13). Equation (4.21) then describes the density of the cloud as it expands after being released from the trap. The width of the cloud will increase in all directions, the peak density will decrease, and the center of the cloud will fall due to gravity. If the trap frequencies and expansion time are known, then the initial size of the cloud in the trap can be determined from the measured widths of a ballistically expanded cloud.

It is interesting to note the behavior of the observed aspect ratio as a function of expansion time. The aspect ratio is defined as the ratio of the axial to radial width (in this thesis at any rate --- other groups define the aspect ratio as the ratio of the radial to axial width). In the trap the aspect ratio is equal to the ratio of trapping frequencies ω_r/ω_z , and for our case $\omega_r > \omega_z$ and so the aspect

ratio is greater than one. As the cloud expands ballistically, the smaller radial direction expands faster and the aspect ratio decreases. At long times, the aspect ratio approaches one. This is expected since the velocity distribution in the trapped cloud is isotropic and at long expansion times the initial velocity distribution will determine the density distribution.

4.6.5 Sloshing. If the center of mass of the cloud is not at the center of the trap, then the cloud will exhibit oscillatory center-of-mass motion. Using the same techniques as in the consideration of the ballistic expansion, it is possible to calculate the evolution of the density of a cloud that is offset initially by x_{cm} in the x direction, y_{cm} in the y direction and z_{cm} in the z direction with widths that are equal to their thermal equilibrium values:

$$n(\vec{x}, t) = \frac{N}{(2\pi)^{3/2} s_r^2 s_z} \exp\left(-\frac{(x - x_{cm} \cos(\omega_r t))^2}{2s_r^2} - \frac{(y - y_{cm} \cos(\omega_r t))^2}{2s_r^2} - \frac{(z - z_{cm} \cos(\omega_z t))^2}{2s_z^2}\right) \quad (4.22)$$

The center of the cloud oscillates as if it were a single particle in the trapping potential. This property makes observation of sloshing a useful way of measuring the trap frequencies.

The presence of anharmonic terms in the potential will lead to the washing out of this collective center-of-mass motion. Atoms at different points in the cloud will have different amplitude oscillations (see Figure 4.16). Since the anharmonicities cause orbits with different amplitudes to have different frequencies and there are a range of amplitudes within the cloud, eventually the cloud will be spread out over the entire region over which it was sloshing. Since the width of the trapped cloud increases, these anharmonic terms can be said to convert sloshing motion into heat, and for all practical purposes this is indeed what happens.

4.6.6 Breathing modes. It is typical that immediately after loading the magnetic trap the initial potential energy and kinetic energy of the trapped cloud will not be in equilibrium. In this case, the widths of the cloud will oscillate as the oscillations of the individual atoms convert potential to kinetic energy (and vice versa) once every quarter trap cycle. We call this the “breathing” mode. If it is still assumed that the cloud has a gaussian distribution of both kinetic and potential energy, it is possible to again determine the time evolution of the density of the cloud by solving for each individual orbit in terms of its initial conditions x_0 and \mathbf{v}_0 , inverting those equations,




Figure 4.16 Anharmonicities turn sloshing into apparent heat. The top cartoon shows the trapped cloud oscillating around the center of the potential that is marked with the X. Two atoms in the cloud, A and B will oscillate in individual orbits, and A will have a larger amplitude than B even though they are both the same distance from the center of the trapped cloud. Since anharmonicities in the potential mean that the oscillation frequency depends on the size of an atom's orbit and A and B have different size orbits, after some time the two atoms will be oscillating more or less out of phase rather than in phase. The lower picture illustrates the relative position of A and B after they have shifted out of phase with each other. At this time, the range in orbit sizes in the trapped cloud will cause it to spread out over the entire region over which it was sloshing. Thus it appears that the cloud has increased in temperature.

and plugging the inverted equations into the initial Boltzmann distribution just as was done in subsection 4.6.4. Considering the density evolution of just a single dimension, it is found that

$$n(x, t) = \frac{A}{\left\{ v_0^2 \left(\frac{1}{2} - \frac{1}{2} \cos(2\omega t) \right) + \omega^2 s_0^2 \left(\frac{1}{2} + \frac{1}{2} \cos(2\omega t) \right) \right\}^{\frac{1}{2}}} \exp \left(- \frac{\omega^2 x^2}{2 \left\{ v_0^2 \left(\frac{1}{2} - \frac{1}{2} \cos(2\omega t) \right) + \omega^2 s_0^2 \left(\frac{1}{2} + \frac{1}{2} \cos(2\omega t) \right) \right\}} \right) \quad (4.23)$$

where v_0 is the initial rms value of the initial velocity distribution and s_0 is the initial rms value of the initial position distribution. A is just a constant that will need to be adjusted to get the number correct. Since the three directions in a harmonic potential are separable, a similar equation will describe the motion in each direction. If $v_0 = \omega s_0$, then the cloud will not exhibit any breathing motion, and this is precisely the condition that is true in thermal equilibrium. Note that the width of the cloud oscillates at twice the trap frequency, which is expected from the symmetry of the trapping potential.

Elastic collisions will cause the breathing mode to damp out and then energy associated with the breathing motion will just be turned into heat in the cloud. Since the breathing mode is not a thermal equilibrium state, as the breathing motion damps out, the entropy of the cloud will increase (which means that the phase-space density will decrease), and so the breathing motion is undesirable from the point of maximizing the phase space density of the cloud. Anharmonicities will appear to have a similar effect as the elastic collisions in damping the breathing motion of the trapped cloud.

4.6.7 Adiabatic Compression and Expansion. Slowly changing the trap frequencies can be used to alter the density and temperature of atoms in the cloud. For instance, adiabatic compression is often used in ^{87}Rb and ^{23}Na to increase the elastic collision rate for better evaporative cooling. There are two time scales that determine whether the change in the trap frequencies can be considered to be adiabatic. One time scale is the inverse of the trap frequencies themselves (the trap periods), and the other time scale is set by inverse of the elastic collision rate in the sample. Usually, the trap period is much shorter than the time scale set by the elastic collision rate, and so it is possible to change the trap in a way that is adiabatic with respect to the trap periods but not with respect to elastic collisions. In light of this fact, the adiabatic change of trap frequencies is considered first in the absence of collisions, and then the effect of elastic collisions are considered.

The adiabatic adjustment of the trap frequency and its effect on a particle's orbit in a harmonic potential is treated in many standard classical mechanics texts [22]. There is an adiabatic parameter of the motion, I where $I\omega=E$. If ω is changed slowly enough, I does not change. This means that the energy of the particle will increase or decrease in a way that is proportional to the change in the trap frequency. The criteria that determines if the change in trap frequency is slow enough is the requirement that

$$\frac{d\mathbf{w}}{dt} \left(\frac{1}{\mathbf{w}^2} \right) \ll 1. \quad (4.24)$$

Since the trapping potential is separable, in the absence of any collisions the individual orbits of the atoms in the cloud will respond to adiabatic changes in any of the trap frequencies just like a single particle in a harmonic potential. Adiabatically increasing any particular trap frequency will increase the energy in that direction. For example, consider a cloud originally in thermal equilibrium in a trap with frequencies ω_r and ω_z . The radial frequency is then increased adiabatically to ω_r' . The cloud can be described as a distribution with two different effective temperatures, T_r and T_z . Originally $T_r=T_z=T_0$. Since the energy in any one of the three directions is $Nk_B T$, after the adiabatic ramp $T_z=T_0$ and $T_r=\omega_r'/\omega_r$. The radial width, which is proportional to $T_r^{1/2}/\omega_r$, will decrease, and the rms velocity in the radial direction will increase.

If the radial and axial effective temperatures are different, elastic collisions will thermalize the sample and cause energy to be transferred from the more energetic direction(s) to the other direction(s). This will occur until the energy is redistributed so that the cloud is in thermal equilibrium. Measuring the time the cloud takes to equilibrate is one way of determining the elastic collision rate.

If the frequencies are changed much more slowly than the elastic collision time scale, then the sample will never be out of thermal equilibrium. In order for this condition to be satisfied, the change in trap frequency ($d\omega/dt$) must be smaller than the square of the elastic collision rate. Since the cloud is always in thermal equilibrium, changes in the trap frequencies that are adiabatic with respect to the collision rate preserve the entropy (and phase-space density) of the sample. The entropy of the (classical) cloud is [33]

$$\begin{aligned}
S &= -k_B \int f \log(f) d^3x d^3v \\
&= -k_B N \left\{ \log \left(\frac{N \mathbf{w}_r^2 \mathbf{w}_z^2}{8 \mathbf{p}^3 k_B^3 T^3} \right) - 3 \right\}
\end{aligned} \tag{4.25}$$

and so for changes in the trap frequencies adiabatic with respect to the elastic collision rate will preserve the ratio $\omega_r^2 \omega_z / T^3$.

For a concrete example, consider a doubling of the radial trap frequency with a cloud of initial temperature T_0 . If the frequency change is adiabatic with respect to the trap frequencies but not the elastic collision rate, then after the cloud thermalizes after the change the temperature of the cloud will be $(2+2+1)/3 = 1.67T_0$. The entropy will have increased. If the frequency change is made slowly enough compared to the elastic collision rate, the temperature increases to $1.59T_0$ and the entropy will not change.

4.7 Optimizing the Loading of the Magnetic Trap from the MOT.

4.7.1 CMOT and Molasses, Part II. Since BEC only occurs at a certain value of the phase-space density, one criterion for the optimization of the science MOT parameters is the maximization of the initial phase-space density in the magnetic trap. However, we have found that evaporative cooling can efficiently cool the cloud and increase the phase-space density, and so it is desirable to sacrifice some phase-space density in favor of a better initial elastic collision rate that leads to efficient evaporation. In practice, the MOT parameters were adjusted so that the number of atoms in the magnetic trap at a selected temperature reached after some evaporative cooling was maximized.

The scaling of the elastic collision rate and phase-space density with the temperature and density of the cloud being loaded into the magnetic trap are useful in understanding the process and limitations involved in loading the magnetic trap from the science MOT. The elastic collision rate scales as $\langle n \rangle \sigma \langle v \rangle$, and the phase-space density is proportional to $\langle n \rangle T^{-3/2}$. For both of these quantities then, it would seem that it would be desirable to maximize the density of the atoms in the cloud at the end of the MOT/CMOT/Molasses/Optical pumping (MCMO) stage for loading into the magnetic trap. By reducing the hyperfine repump power and detuning the trap laser by 25 MHz or more, it was possible to dramatically decrease the size of the MOT during the CMOT stage, making

it a factor of two smaller, albeit at the cost of almost half of the atoms in the cloud. However, it turns out that this large increase in density could not be used to improve the magnetic trap loading due to the magnetic trapping potential that we could produce, so we used the strategy of detuning the trap laser to only 13 MHz and reducing the hyperfine repump power not quite as far. This latter strategy did not increase the density as much but only 10% of the atoms were lost, and that lead ultimately to better evaporation performance.

The fundamental reason that we could not take advantage of the high density after the CMOT stage was that the density of the cloud of atoms in the magnetic trap is not determined by the initial density of the cloud loaded in, but rather the total energy of the cloud. Since the magnetic trap is turned on suddenly, a potential energy is imparted to the cloud due to the trapping potential. If a dense cloud is loaded into the trap and the kinetic energy in the cloud is greater than the potential energy, the cloud will expand and exhibit breathing motion as detailed in section 4.6.6. Elastic collisions (and anharmonicities in the trapping potential) will bring the cloud to thermal equilibrium. For a cloud after the MCMO stages with an initial kinetic energy corresponding to a temperature T_0 and a spherically-symmetric spatial size σ , the energy of the cloud after being loaded into the magnetic trap is

$$PE = \frac{m}{2} (2w_r^2 + w_z^2) \mathbf{s}^2, \quad KE = \frac{3}{2} Nk_B T_0 . \quad (4.26)$$

The temperature of the cloud in the trap will eventually be $PE+KE=3Nk_B T$. The only advantage of a large initial density is therefore a reduction in the potential energy imparted to the gas when the magnetic trap is turned on. If the kinetic energy of the cloud is much higher than the imparted potential energy at the time of loading, which is the case for the high-density CMOT parameters, then the eventual density of the cloud in the magnetic trap will be determined by the kinetic energy. The initial large density will not contribute to the density in the magnetically trapped sample.

One way to preserve the an initially high density is to make the imparted potential energy match the initial kinetic energy by increasing the trapping frequencies. While this will result in a higher-density trapped sample, this is somewhat counter-productive since it will raise the temperature of the cloud in the trap. However, the trapping frequencies can be changed

adiabatically, and so if increasing the trapping frequencies leads to an increase in phase-space density in the trap, the trap can be adjusted to any desired configuration after the atoms are loaded. There is a problem, however, that prevents us from increasing the trapping frequency.

For reasons of simplicity of design, only the baseball coils are turned on when the atoms are loaded into the magnetic trap. Without the bias coils, there is no independent control of the bias field in the magnetic trap. This limits the range of possible trapping frequencies available for loading the atoms into the trap. The current through the baseball coil can be increased, but since the bias field will increase as well, the magnetic moment of the $F=2$ $m_f=-2$ atoms will decrease. This offsets the increase in the curvature of the trap and the trapping frequencies do not change significantly over a large range of currents (see Figure 4.17). The trapping frequencies of the trap that we use to load the atoms are $2\pi \times 9.2$ Hz in the radial and $2\pi \times 5.6$ Hz in the axial direction. In principle, the radial direction could be strengthened by using the bias coils as well as the baseball, but the axial direction would still be relatively weak. Given the approximately spherical shape the cloud has before being loaded into the trap, the matching into a high aspect ratio trap would still be poor and an initially high density will not be maintained.

Another way to match the initial kinetic energy and the imparted potential energy is to cool the cloud. The cloud will not exhibit any breathing motion, and so it will be loaded in at a higher phase-space density. Since we cannot easily adjust the trapping frequencies, this is in fact what we do. The molasses stage is used to remove kinetic energy and cool the cloud from temperatures over $120 \mu\text{K}$ to $15 \mu\text{K}$. While the molasses is cooling the cloud, there is no longer a confining force present and the cloud will expand and its density will decrease. For an average temperature of $60 \mu\text{K}$ during the 4 ms molasses stage, the cloud would expand 0.3 mm. This expansion size is added in quadrature with the initial size of the sample. So, for clouds with an initial rms width of 0.6 mm or greater this is not a big effect, but it does limit the maximum achievable density since it will be impossible to decrease the size below 0.3 mm.

In addition to the time spent in the molasses, the atoms must also be optically pumped into the proper state, and this takes just under 2 more ms, allowing the cloud to expand further. With this additional expansion, any initially high density is diluted and there is no significant advantage

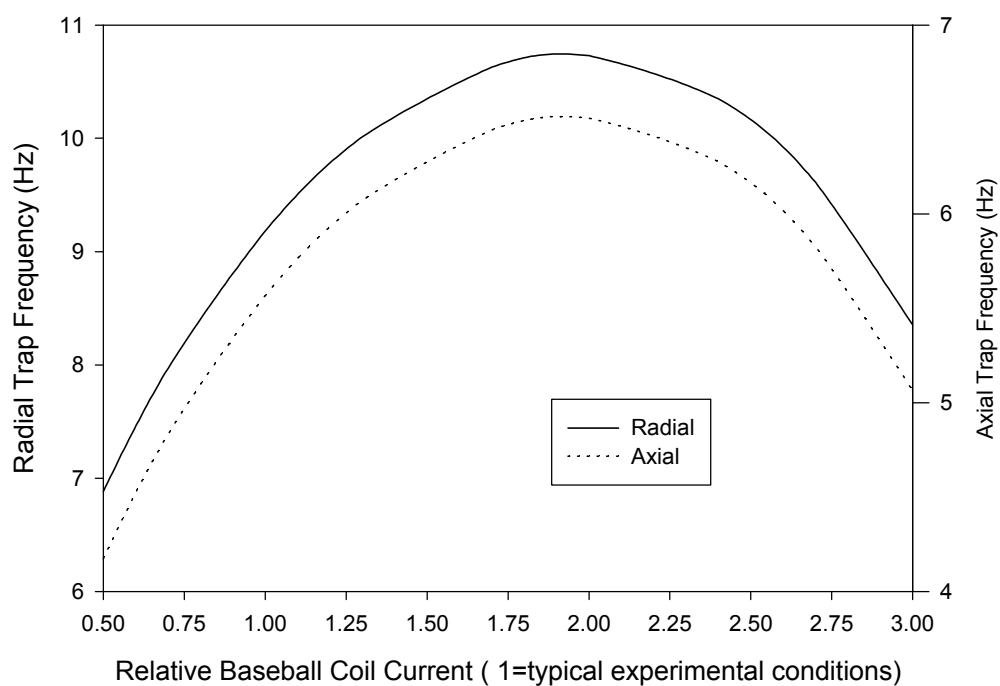


Figure 4.17. Relative insensitivity of the trapping frequencies to initial baseball coil current. This figure demonstrates how increasing the amount of current in the baseball coil during the trap turn cannot increase the trap frequencies from the typical conditions used by more than $\sim 25\%$, and that can only be achieved by using nearly twice as much current. This is because the magnetic moment of the atoms decreases as the bias field increases. The x axis is scaled to the usual trap turn on current.

gained from adjusting the CMOT parameters to produce a dense sample. In fact, the atoms lost during the CMOT stage configured to produce very high initial densities ($> 3 \times 10^{10} \text{ cm}^{-3}$) offset any gains. Figure 4.18 shows the optimization of the evaporation performance with respect to the CMOT detuning. It is beneficial to detune slightly from the science MOT loading detuning of 10 MHz, presumably to increase the density of the cloud some, but detuning further does not benefit the evaporation performance.

4.7.2 Optical Pumping Optimization. It was also necessary to optimize the optical pumping. The figure of merit used was the optical pumping efficiency, the ratio of the number of magnetically trapped atoms to the number of atoms available after the MCMO preparation. The density of the cloud in the magnetic trap is directly proportional to the fraction of post-MCMO atoms that are successfully pumped to the desired $F=2 \ m_f = -2$ state. As was mentioned in subsection 4.4.9, the limitations to the optical pumping are not understood. Since it seems that the OP laser beam can only double the optical pumping efficiency from simply pumping the atoms out of the $F=3$ state, the alignment of the MOT and molasses is very important in order to achieve efficient optical pumping. Great pains were taken to tweak both the alignment of the MOT trapping beams and the earth shims (the earth shims are the only coils on during the molasses stage). The MOT trapping beams are fairly well balanced with the power in each pair of trapping beams being: up/down, 1.6/1.5 mW; north/south, 1.8/1.7 mW; and east/west, 1.8/1.9mW.

We found that the best starting point for the MOT alignment was to overlap the six beams as best as possible with each other and the center of the MOT coils. Once this was done, the earth shims were adjusted to make the molasses look as uniform as possible. Any stray magnetic field would tend to make the atoms accelerate out in a single direction. Once the molasses looked good, the earth shims were adjusted very slightly to improve the optical pumping efficiency. Next, the alignment of the MOT trapping beams were adjusted to improve the optical pumping efficiency. Each of the six beams going to the MOT had at least two mirror bounces, and each mirror could be adjusted in both the horizontal and vertical direction. These 24 knobs were each individually changed to get the best optical pumping efficiency possible. Usually, a few rounds of this

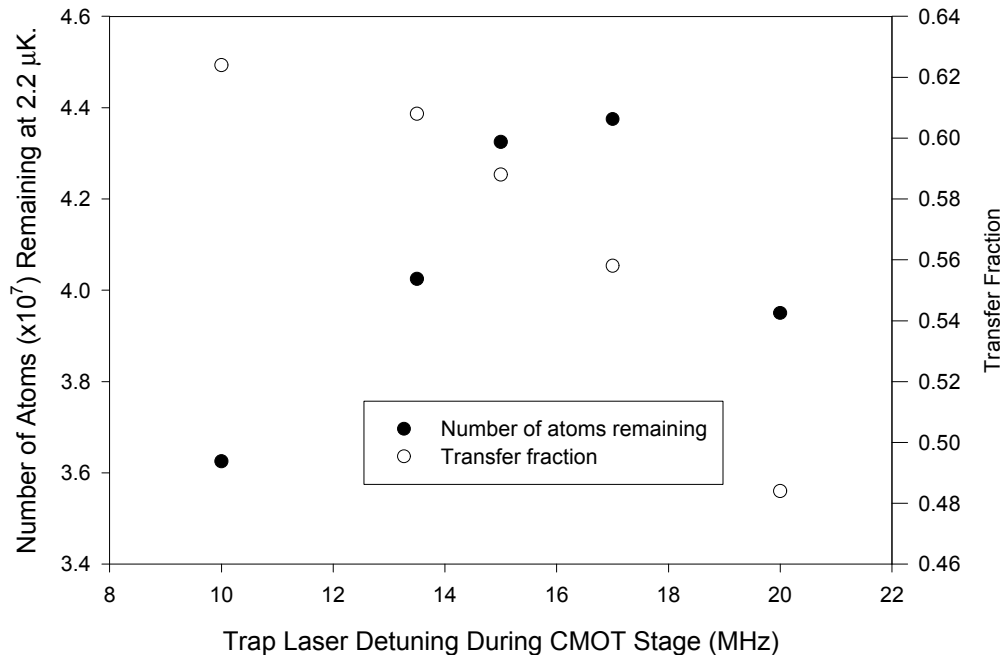


Figure 4.18 Optimization of CMOT Stage Laser Detuning. This figure is typical of how the MOT parameters are adjusted for optimum performance. The filled circles show the number of atoms remaining in the trap after a fixed amount of evaporative cooling. The evaporative cooling is enough to reduce the temperature of the cloud to 2.2 μK . The open circles show the fraction of atoms that are loaded into the magnetic trap compared to the initial number of atoms that are loaded into the science MOT before the CMOT and subsequent stages. The decrease in transfer fraction is due to the fact that the losses during the CMOT stage increase as the detuning is increased. Presumably, the optimum CMOT detuning with respect to evaporation performance represents a trade-off between increasing the initial density of the sample (and cooling it a little bit during the CMOT stage) and losing atoms during the CMOT stage. For this data the CMOT stage was 30ms long, the typical time for that stage.

optimization were necessary until the optical pumping efficiency was maximized. Once the laser alignment had been tweaked, the earth shims were adjusted again for good optical pumping efficiency. While the optimization of the MOT laser alignment did not seem to degrade the MOT performance, the final earth shim settings were often a compromise between optical pumping efficiency and efficient molasses cooling. The alignment of the OP/probe beam was adjusted as well to maximize the optical pumping efficiency.

Unfortunately, the MOT laser alignment is not stable but does drift slowly over time. We believe that this is most likely due to a small change in the angle of the diffraction grating [67] used to stabilize the laser frequency. Temperature changes cause the angle of the grating necessary for the desired frequency to drift, leading to day-to-day variations. We found that it is sufficient to tweak one mirror shortly after the laser every other day or so to maximize the optical pumping efficiency, and that seems to correct this long-term drift.

Once the lower MOT has been aligned well, we would observe optical pumping efficiencies of up to 80%. The optical pumping efficiency when the atoms were only pumped out of the $F=3$ state (no OP light) was as good as 45%, representing a large bias towards the negative m_f states in the MOT/molasses stages.

There was an additional unexplained behavior in the OP sequence. For the intensities of OP light that were typically used ($50 \mu\text{W}/\text{cm}^2$), there was only a few μK of heating due to the optical pumping. As the intensity of the laser was increased, the atoms would heat up much more. This would be expected if more photons were scattered when the intensity was increased, but we found that this heating occurs even if the OP laser power/OP pulse time product was kept constant, which should naïvely keep the number of photon scatters constant.

4.7.3 MOT Shims and Sloshing. Another consideration is the minimization of the sloshing of the atoms in the magnetic trap. This minimization means that the cloud in the CMOT needs to be centered on the center of the magnetic trapping potential and that the magnetic trap should be turned on quickly to minimize the time the atoms have to accelerate under the influence of gravity. The trap turn-on is discussed in the following subsections detailing the control of the magnetic trap. Centering the CMOT cloud on the center of the magnetic trap potential is performed by using the

MOT shim coils. Fields created by these coils move the position of the $B=0$ point in the MOT and CMOT and therefore shift the position of the cloud. Since the main anti-Helmholtz magnetic field and laser detunings are changed between the MOT and CMOT stages, the shims are set to run at one current for the MOT stage and a different current for the CMOT stage. During the MOT stage the shims are adjusted for good push transfer efficiency while during the CMOT stage the shims are responsible for centering the CMOT cloud on the center of the magnetic trap.

Two methods are used to minimize the slosh. Absorption imaging (see section 4.11) was used to look at the center-of-mass motion of the cloud and the CMOT shims were adjusted to minimize the motion. The cloud can only be imaged in two directions in this way, but it is possible to measure the sloshing of the cloud in the third dimension by using the MOT fluorescence. The third dimension is visible to the MOT fluorescence light collection. By placing a razor blade in the collection path so that it blocked a fraction of the fluorescence light, any oscillatory center-of-mass motion of the cloud would appear as an oscillating fluorescence signal as the cloud moves behind then past the blocking razor blade. The other way that the sloshing was minimized was to simply evaporatively cool the cloud to a fixed temperature and then adjust the shims so that the number of atoms remaining in the trap at that chosen temperature was maximized.

Of the two methods, measuring the center-of-mass motion is a more accurate way of minimizing the slosh. Also, if any changes were made to the magnetic trap loading timing or current control, then the atoms could possibly fall before being loaded into the magnetic trap and acquire a velocity that would lead to slosh that could not be removed through adjusting the position of the cloud in the CMOT. Measuring the center-of-mass motion ensured that such a velocity-induced slosh was not introduced by any changes that were made. However, measuring the center-of-mass motion requires the imaging system to be in a low-magnification configuration and the insertion of the razor blade in the MOT fluorescence collection path. It is much easier to optimize the evaporation performance at a fixed temperature since less changes had to be made to the apparatus, and so this latter method was used to periodically check the slosh as part of maintaining good evaporation performance.

4.7.4 Other Spin States. The $F=2$ $m_f = -2$ state is not the only state that can be trapped in a magnetic trap. By just examining magnetic moments, the $F=2$, $m_f = -1$ state and several of the $F=3$ states appear trappable as well. The optical pumping process is not perfect, and so some contamination of atoms in these other states in the trap is expected. This turns out to not be a large concern, however, due to the importance of the sag in the magnetic trap. Atoms with different magnetic moments will have different trapping frequencies and therefore different amounts of sag in the magnetic trap. Since the atoms are loaded into a trap at ~ 200 G, atoms in the $F=2$, $m_f = -1$ state will have their magnetic moments weakened significantly and will fall away from the center of the $F=2$, $m_f = -2$ potential where they were initially loaded. As the $F=2$, $m_f = -1$ atoms fall, the magnetic field increases and their magnetic moment gets weaker and with the kinetic energy that they pick up while falling they eventually leave the trap.

There are also a few $F=3$ atoms in various magnetic sublevels that are left from the OP sequence, but since the center of their potentials is not the same as the $F=2$, $m_f = -2$ potential due to gravitational sag they slosh in large orbits in the trap. They could be a potential source of heat if the $F=3$ atoms collided with the $F=2$, $m_f = -2$ cloud, but the anharmonicities in the trap make these collisions unlikely after a few oscillations. Experimentally, we have found no difference in evaporation performance when the number of $F=3$ atoms more than doubles. The only state that seems likely to be a concern is the $F=3$, $m_f = 2$ state, since at zero field it has the same magnetic moment as the $F=2$, $m_f = -2$ state. However, at 200 G the $F=3$, $m_f = 2$ state has a magnetic moment that is 38% greater than the $F=2$, $m_f = -2$ state. So, atoms in the $F=3$, $m_f = 2$ state are still far from the center of its potential and therefore those atoms will also slosh in large orbits in the trap.

4.8 Control of the Magnetic Trap Fields.

4.8.1 Magnetic Field Requirements. The currents required to produce the necessary bias magnetic field and trap curvature in our apparatus are reasonably large, on the order of 200 A. There are two requirements for these currents: they must be stable at the ppm level so that the forced evaporative cooling is reproducible, and they must be able to be changed quickly (~ 0.1 ms) both to turn off and on the magnetic trap, and to alter the bias magnetic field quickly so that the scattering length can be rapidly changed through the Feshbach resonance. These two requirements conflict

with each other in the sense that any servo usually requires making a trade-off in bandwidth and gain. In order to get around this difficulty, we have created an electronics system that can be switched from stable operation for evaporative cooling to fast operation for sudden current changes.

4.8.2 FET Banks. The current sent through the baseball coil is controlled by a set of 4 Advanced Power Technology model APT10M07JVR high-current field-effect transistors (FET) that are wired in parallel with each other. This set of FET transistors is referred to as a FET bank. An 8 volt 580 amp HP 6681A power supply is connected to the drain of the FET transistors, while the source is connected to the top of the baseball coil. The general arrangement of the FET banks, power supply and coils is shown in Figure 4.19. The amount of current that flows through the FET bank is determined by the gate-source voltage, and so controlling the gate-source voltage regulates the amount of current flowing through the FET bank and hence through the baseball coil. The current through the bias coils is controlled in a similar fashion with another FET bank.

The FETs in the bank are connected to a copper plate that is water cooled in order to remove the large amount of heat (up to 700 W) dissipated in the FETs themselves. We found that it was important to have good thermal contact between the FETs and the copper plate in order to prevent them from burning out. This was accomplished by applying a very thin layer of thermal grease between the FET and copper plate and then tightly bolting the FET to the copper plate.

4.8.3 Parallel Baseball and Bias Coils. In the ^{87}Rb work performed with this apparatus prior to this thesis work, the coils were not arranged in the way that they are shown in Figure 4.19, but rather as shown in Figure 4.20. Running the two sets of coils in series had the advantage of significant common-mode noise reduction in the value of the bias magnetic field. In the course of typical operation, the bias field was maintained at approximately 1 G by canceling the ~ 400 G bias field produced by the baseball coil with 399G from the bias coils. Any increase in the current through the baseball coil would be accompanied by an increase in current through the bias coil, and so the noise induced by current fluctuations in the trap bias magnetic field were suppressed by a factor of about 400.

The dummy coil was included in order to facilitate fast switching times. Rather than the power supply having to supply a large voltage to change the current in the inductive coils quickly,

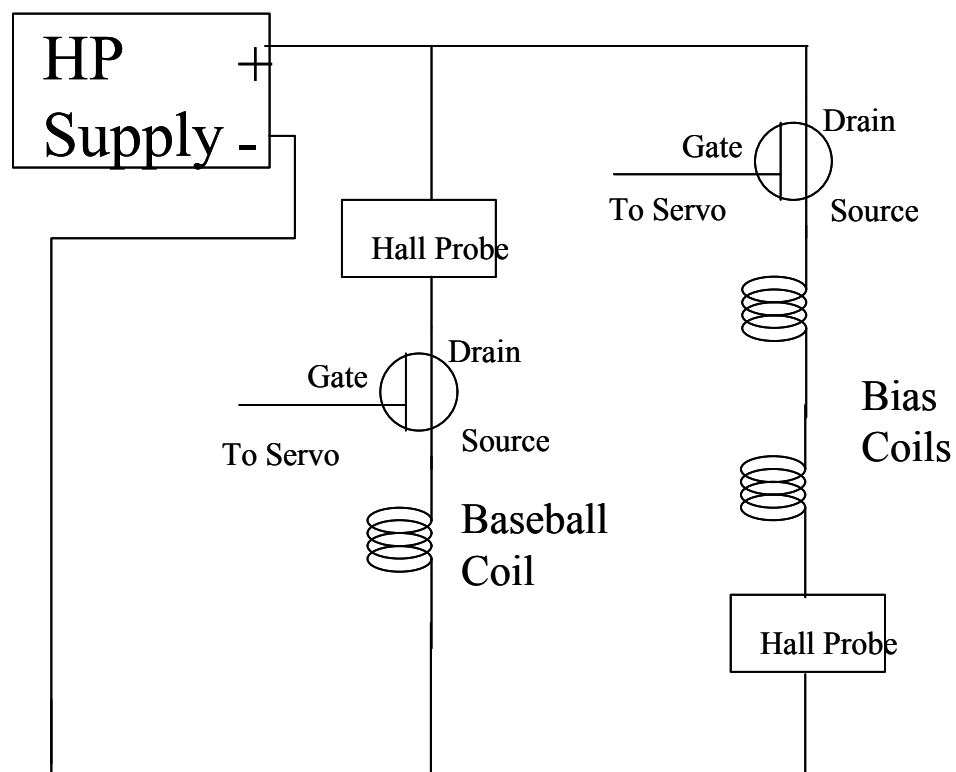


Figure 4.19. Baseball and Bias Coil Current Supply and Control Schematic. The FETs shown in the diagram are actually FET banks as described in the main text.

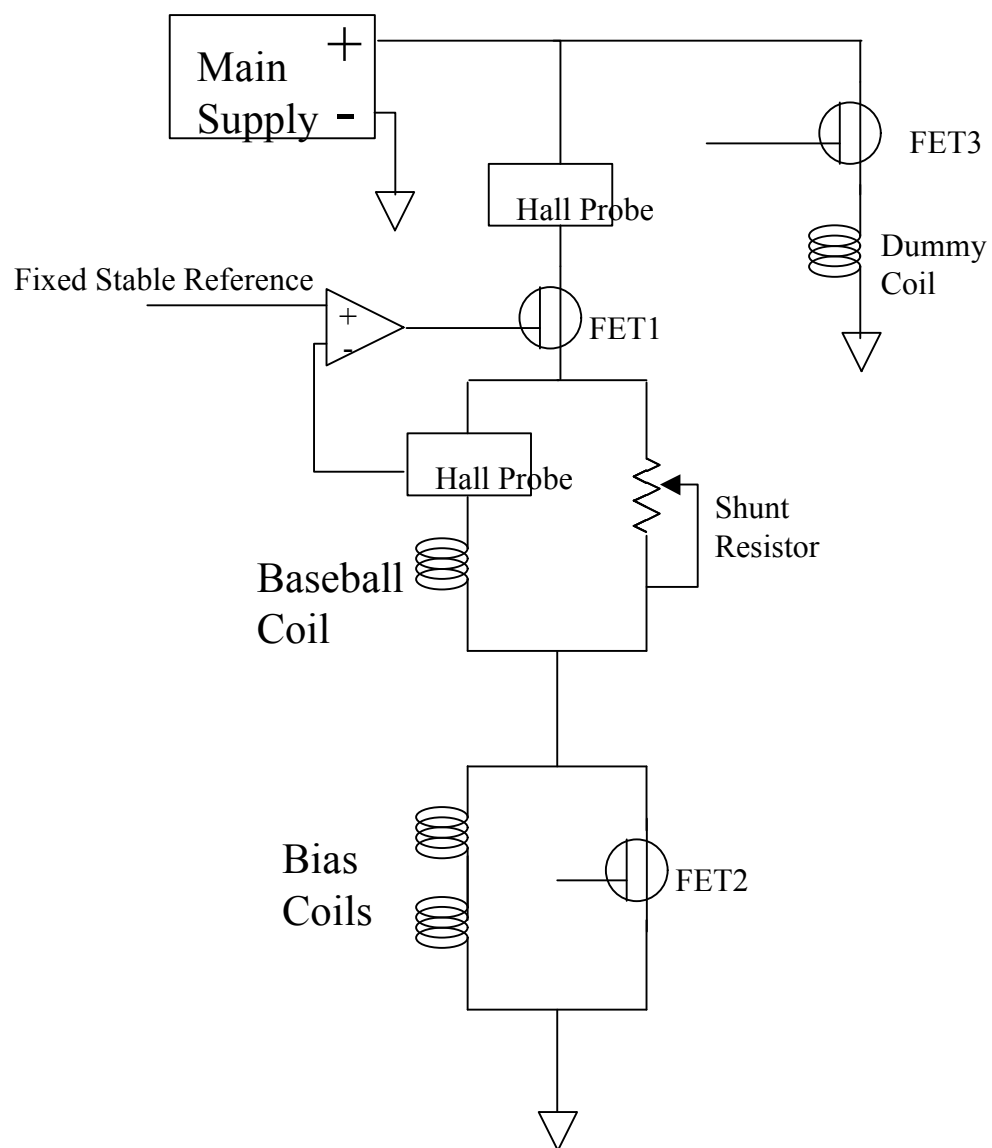


Figure 4.20. Schematic of the Old Arrangement of Coils and Current Controls [1].

the current could be rapidly switched to the dummy coil. This also had the apparent advantage of running the main power supply in a constant current mode rather than forcing it to turn on and off large amounts of current quickly.

Finally, the bias field in the magnetic trap was controlled using the shunt resistor shown in Figure 4.19. The overall current through the coils was regulated by FET1, but FET2 was operated in the completely closed configuration. By varying the resistance of the shunt resistor, a variable amount of current could be diverted around the baseball coil, allowing the bias magnetic trap field to be altered. Unfortunately, this arrangement was sensitive to changes in temperature since they led to changes in resistance of the coils and shunt resistor that did not scale proportionally, and slow drifts were a significant problem.

All three of these features have been changed for the work in ^{85}Rb : the coils are in parallel rather than series, and the dummy coils and shunt resistor have been removed. The primary reason for this is that rather than operating at small values of the bias magnetic field, the Feshbach resonance required us to work over a range of fields near 160 G. This made the shunt resistor impractical both because it would be difficult to change the bias field in the trap and because an unreasonably large amount of current would have to be diverted through it. Common-mode suppression would not be as effective at these higher magnetic fields, which obviated the need for a series coil configuration. Once the coils were no longer in series, diverting the current through the dummy coils became more complicated, requiring the coordination of three FET banks rather than two, so the dummy coil was eliminated.

In addition to these considerations, there were other problems with the series configuration as well. The series configuration worked well as long as the FET bank FET2 was run either completely open or closed. When FET2 was run in an intermediate regime, large parasitic oscillations would occur for certain values of the gate-source voltage. There was no servo present in the system when these oscillations were observed. Merely applying a DC voltage to the FET gates was enough to cause the oscillations. These oscillations had frequencies of a few kilohertz, and the size of the oscillations was such that the current was going to zero and then peaking past the desired value set by the gate-source voltage. The baseball and bias coils combined with the gate-source and

gate-drain (“Miller”) capacitances of the FETs means that the system had characteristics like an RLC circuit, and so the appearance of large oscillations was not necessarily surprising. Given that we wanted to perform experiments over a range of magnetic fields this situation was unacceptable. The parallel configuration of the coils eliminated this problem.

The dummy coils were also prone to these parasitic oscillations, although the range of operation was not as limited as when the coils were in a series configuration. The coordination of three separate currents would also have been complicated, and so the dummy coils were removed. The power supply must now supply voltage to rapidly turn the coils on and off, but the turn on and turn off times (described below) are adequate. The power supply has been operated for a couple of years under these conditions and has shown no problems resulting from these rapid current changes. An added advantage of the removal of the dummy coils is that the FET banks themselves are the first stage of our current servo control. Since a constant gate-source voltage should provide a constant current, variations in the power supply voltage are suppressed by the FET properties.

4.8.4 Magnetic Trap Current Control. In order to regulate the current through the baseball and bias coils, there must be some way of measuring that current. We use hall probes (FW Bell CL-300) to measure the current in both the baseball and the bias coils. These hall probes use the Hall effect to measure the magnetic field produced by the current in the wires leading to the baseball and bias coils. Originally, we used open loop hall probes that simply produced a signal that was proportional to the detected magnetic field. This signal was sensitive to changes in temperature, and the day-to-day reproducibility was quite poor. In fact, it was more stable for us to simply use the voltage drop across the coils to measure the current.

After a time, we switched to closed loop Hall probes. In the closed loop probes, the Hall sensor is placed in the gap of a C-shaped ferrite core. A current flowing through a wire wrapped around the ferrite core is used to produce a magnetic field in the gap that cancels the magnetic field due to the current flowing to the baseball/bias coils. A simple servo regulates the current around the ferrite core so that the field at the Hall sensor is zero. The canceling current is then measured to determine the baseball/bias current. Since the Hall probe is reading zero, the closed loop probes are much less sensitive to temperature changes. The bandwidth of the closed loop probes is much better

than the open loop probes as well (200 kHz is the specification). A high-precision low-temperature-shift resistor (Vishay model VHP) is used to convert the canceling current to a voltage that indicates how much current is flowing past the Hall probe.

A servo that uses the Hall probe voltage as an error signal is used to regulate the current through the coils. The error signal is compared to a reference voltage and the servo is designed to make the difference between those two signals be zero. The diagram of the circuits involved in the servo control system is shown in Figures 4.20-4.25. Probably the easiest way to explain the operation of all the components of the servo is to go through a typical trap operation cycle and describe the requirements at each stage.

The switching on and off of the trap currents is controlled by a TTL signal sent to the on/off control box. When the trap is to be held off, the on/off box sends a reference voltage to the servo that corresponds to a negative current (which cannot possibly be provided) forcing the servo to set the FET bank gate voltages to a negative value and opening all of the FETs. Once the trap is to be turned on, the reference voltage is switched to a positive current value as provided by reference box.

Once the atoms are ready to be loaded into the magnetic trap, it is important to turn the trap on quickly. Otherwise, the atoms will fall under gravity and slosh in the trap. This sloshing will effectively be turned into heat and therefore degrade the evaporation performance. Only the baseball coils are turned on initially. The trap is turned on with the capacitors in the baseball servo circuit switched out (see Figure 4.25). This improves the servo bandwidth at the expense of the DC gain, but the DC level is not crucial for the loading process. The DC gain is low enough that the offset available in the servo circuit is used to set the DC level to the appropriate value. With this system, the baseball current is turned on to 90% of its full value in 0.4 ms. We found that this turn-on time was sufficient to load the atoms with a minimum of sloshing. For slower turn-on times (above 1 ms), the sloshing of the atoms in the magnetic trap due to their initial velocity gain while falling degraded the evaporation performance significantly.

After the atoms have been loaded, the capacitors are then switched back in to improve the DC stability of the baseball current. The improved DC gain means that the offset voltage in the servo circuit is no longer important and the reference box determines the reference voltage. Initially,

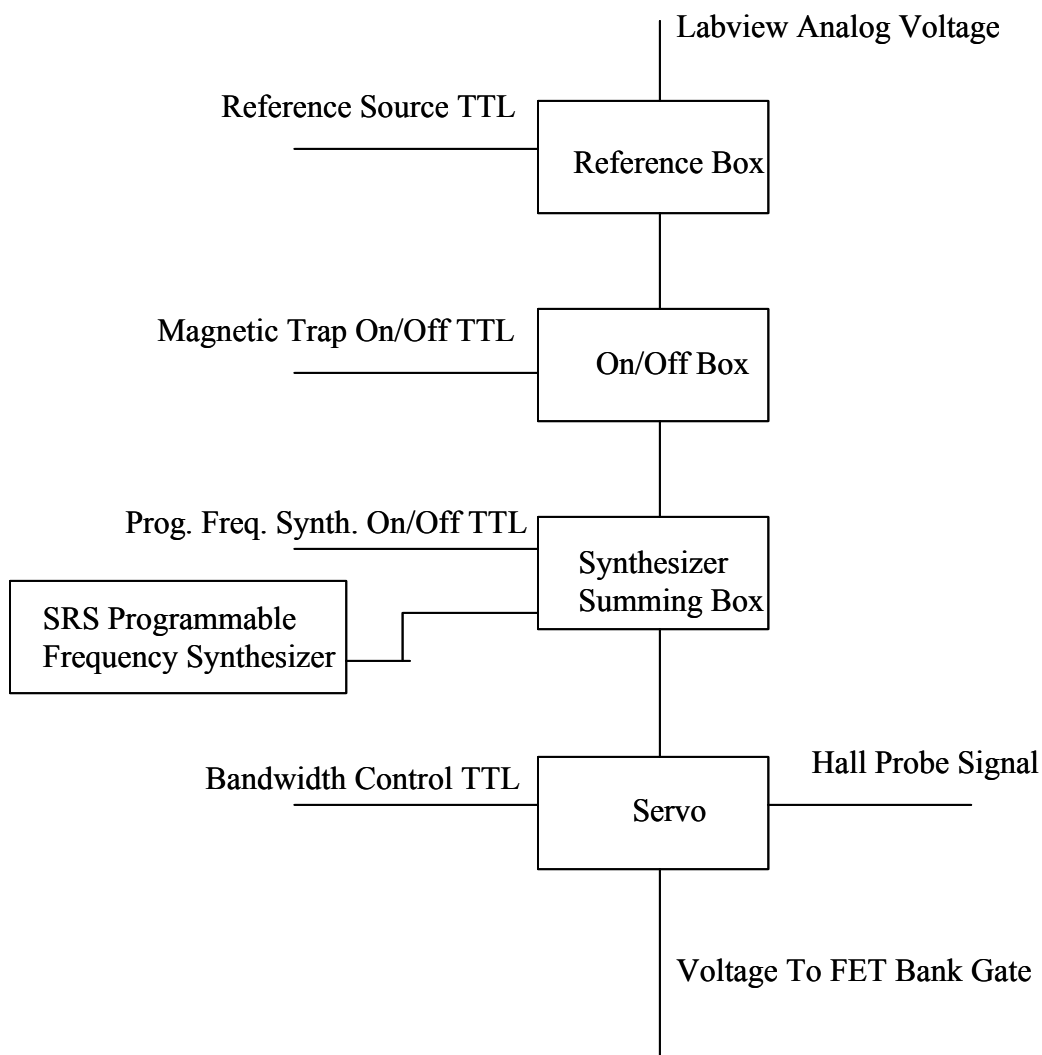


Figure 4.21. Schematic of the Baseball Coil Current Control Electronics. All of the individual components, except for the SRS Programmable Frequency Synthesizer which is a commercial product, will be detailed in the following figure. Each TTL signal comes from the Labview TTL output board. The bias coils are controlled by a similar chain of components, but without the synthesizer and its summing box and with no Bandwidth control TTL.

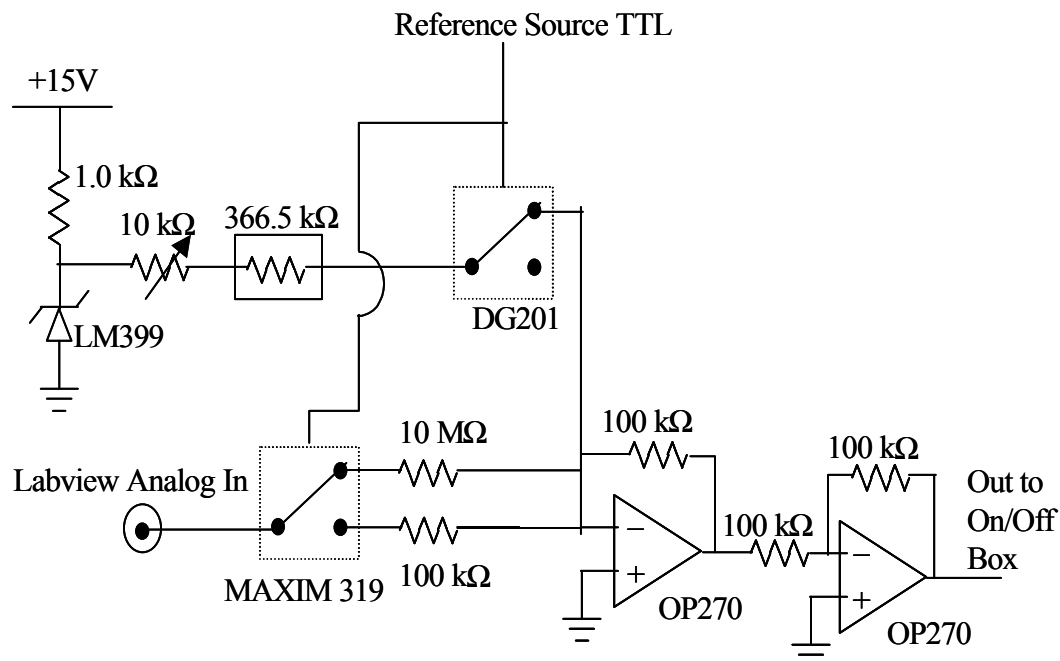


Figure 4.22. Reference Box Circuit Diagram. All resistors are high-quality metal film. The 366.5 kΩ resistor shown in the box is actually several resistors in series that have that total resistance.

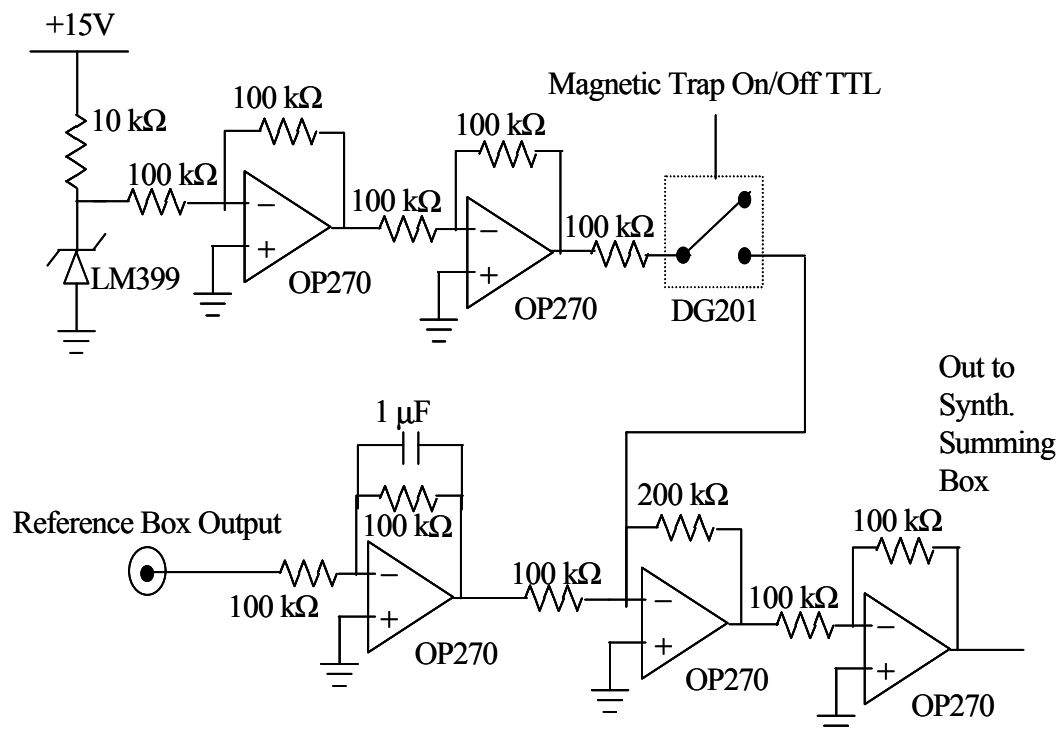


Figure 4.23 On/Off Box Circuit Diagram. The two op amps after the LM399 are an artifact of the many changes that have been made to this circuit. In addition to turning the trap on and off, this circuit is also used to filter the high-frequency noise out of the Labview analog signal.




Figure 4.24. Synthesizer Summing Box Circuit Diagram.

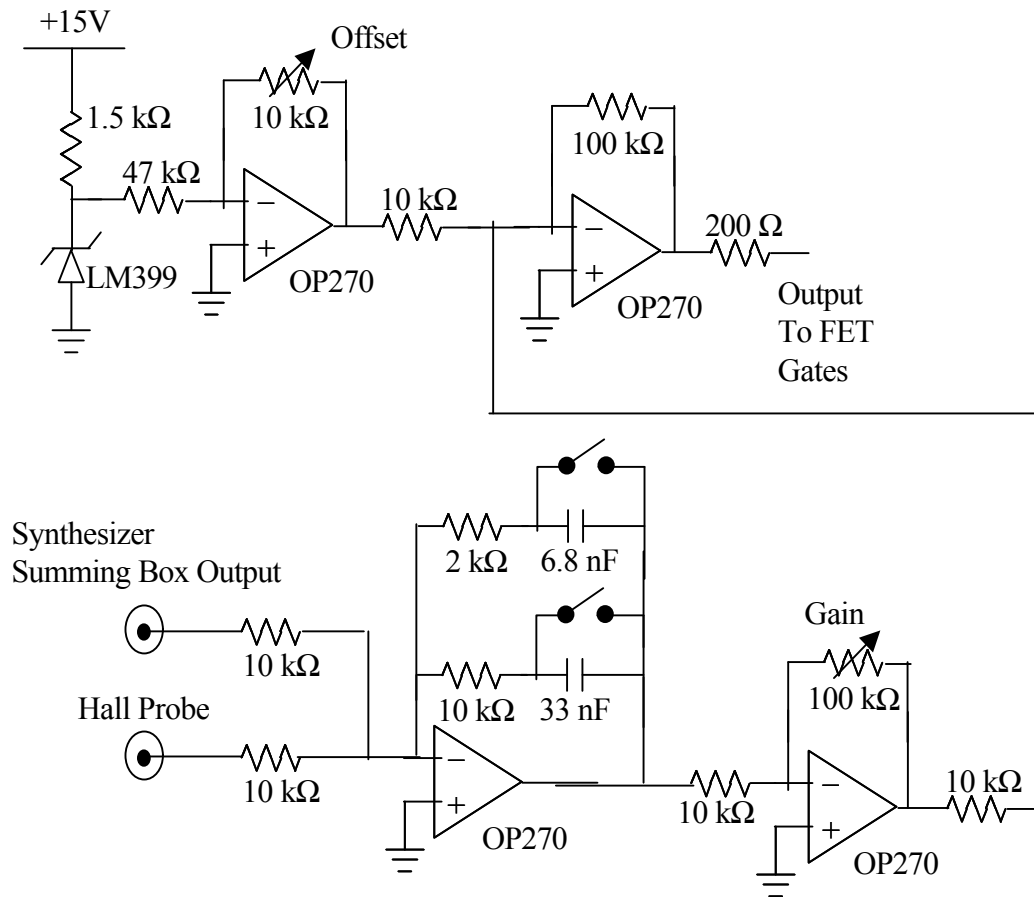


Figure 4.25. Circuit Diagram for the servo. The switches shown are DG201s that are controlled via the bandwidth control TTL. The somewhat unusual configuration of capacitors and resistors on the lower left op amp is designed to deal with the coil's phase shift around 1 kHz described in the main text. From zero to 480 Hz the capacitors dominate and there is a 90° phase shift. At 480 Hz, the $10\text{ k}\Omega$ resistor begins to dominate the 33 nF capacitor and the phase shift drops to 50° near 1 kHz. This reduction in phase shift is enough to allow the servo to run with a higher gain without the system oscillating.

the reference box is set so that the Labview analog output voltage controls the reference level. This output voltage can then be ramped to any desired value, controlling the baseball coil current. At this stage, the bias coil current, which is controlled in the same way as the baseball, is ramped on slowly using a second Labview analog output to control its reference voltage level. While these computer outputs have reasonable DC stability, they contain noise at higher frequencies. Therefore, these computer signals are filtered using a 100 ms time constant RC filter in the on/off box. This cleans up the reference signals at the expense of being able to change the currents quickly via the analog computer outputs.

Some evaporation is then performed to cool the atoms at 251 G. Once the atoms get cold enough, the trap bias magnetic field is changed to 162 G and that is where the remainder of the evaporative cooling takes place. The Labview DC stability is not adequate for evaporating to low temperatures, so the source of the reference voltage is then switched in the reference box. The computer analog output voltages are summed in at a gain of 1/100 with voltages produced by an LM399 stable voltage reference and a voltage divider comprised of high-quality resistors. Control is then still available through the computer analog voltages, but its speed is still limited and now its range is limited as well.

Often it is desirable to change the trap bias magnetic field rapidly once the atoms have been evaporatively cooled. For maximum flexibility and speed we use a programmable SRS DS345 frequency synthesizer. The output of the synthesizer is summed directly into the reference voltage at the servo circuit so that it is not filtered and therefore can be changed quickly. During evaporation, an analog switch is thrown open to prevent the synthesizer output from being summed in, because there is noise on the synthesizer output.

After the magnetic field manipulations are completed, the magnetic trap must be turned off quickly in order to image the atoms or recapture them in a MOT to measure their fluorescence. This is accomplished simply by changing on/off box state and overdriving the servo to a “negative” current. The trap is turned off in 160 μ s, with the bias current reaching zero \sim 20 μ s before the baseball current does.

4.8.5 Servo Optimization. In order to get as much DC gain as possible, the servo is designed as an integrator. Due to some delays in the system, however, a simple integrator was not optimal. We noticed that with proportional gain alone the current oscillated at 11 kHz when the servo gain was turned up sufficiently high. With a pure integrator, the servo oscillated at a frequency of ~ 1 kHz. Since the only difference between the two cases was the 90° phase shift in the integrator, we surmised that around 1 kHz there was a natural phase shift of over 90° but less than 180° . Therefore the response function of the servo circuit was changed to reduce the phase lag near 1 kHz.

By examining the reproducibility of the evaporative cooling, we can determine that the regulation of the DC current level must be better than 6 ppm. The AC performance of the servo is also difficult to determine, since the variations are relatively small. Figure 4.26 shows the hall probe signal and the voltage drop across the baseball coil while the current is on and while the current is off. The baseline of the hall probe signal corresponds to 4 ppm fractional noise, but from the plot it can be seen that there is little difference in the baseline when the current is on and when it is off. The dominant noise feature is the 60 Hz harmonic at 300 Hz with a 9.5 ppm amplitude. The voltage drop measurement indicates that the hall probe signal is an accurate representation of the noise, at least near 300 Hz, even though it is the error signal itself.

4.9 Forced Radio-frequency Evaporation.

4.9.1 Simple Picture of Forced Radio-frequency Evaporative Cooling. Once the atoms have been loaded into the magnetic trap, they are cooled through forced radio-frequency (rf) evaporative cooling. At thermal equilibrium, a small fraction of atoms have energies that are several times greater than the average energy of atoms in the cloud. If these atoms are removed from the cloud, then the average energy of the remaining atoms is reduced. Elastic collisions will then cause the sample to come to equilibrium at a colder temperature. If this process is repeated the cloud can be cooled at the expense of the number of atoms in the cloud.

Evaporative cooling can be very efficient. In the absence of any loss or heating rates, it would be possible to wait an infinite amount of time and then remove a single atom that contained all of the energy in the sample, a very efficient removal of energy. The primary limitation to the

Figure 4.26. The power spectrum of the noise with and without current flowing through the baseball coil is shown. The top plot is the hall probe signal, and since that is the error signal all that is demonstrated is how well the servo is working. The bottom plot shows the voltage drop across the coils. This is independent of the servo and so it is in some ways a better measurement of the true noise properties, but because of the inductance of the coil, the current flowing through the coil is not proportional to the voltage drop at frequencies higher than about 100 Hz. The DC signal sizes for the hall probe signal and voltage drop were 3.78 and 2.78 V respectively. Not shown is a mild increase in the voltage drop noise near 11 kHz, the unity gain point of the servo.

efficiency of the cooling is therefore set by the loss and heating rates of the trapped cloud. As these limiting rates are increased, atoms need to be removed via evaporation faster to remove energy more rapidly from the cloud. The more atoms that are removed, the lower the average energy of the atoms removed since the most energetic are taken first, and the less efficient the evaporation. A detailed discussion of evaporation in ^{85}Rb is presented in Chapter VII. The rest of this section concerns itself with the details of how the most energetic atoms are removed from the magnetically trapped sample.

In order to remove the most energetic atoms, it is necessary that there be some way to discriminate them from the lower energy atoms in the trapped cloud. The more energetic atoms have larger orbits, and therefore span a larger range of magnetic fields in a single orbit. Any magnetic field or energy sensitive transition that is narrow enough can then address only the highest energy atoms in the cloud. Radio-frequency (rf) transitions between different magnetic sublevels are an ideal candidate for such a narrow transition.

Time-dependent magnetic fields can cause transitions from one magnetic sublevel to another [77] via the Zeeman interaction. The frequency of the transition will depend on the energy difference between the two sublevels and is therefore dependent on the value of an external magnetic field. These transitions have weak oscillator strengths and in general long wavelengths and so their natural lifetime is enormous – on the order of years is typical. Therefore, the transitions are very narrow, precisely what would be required for evaporative cooling.

Forced-rf evaporation is most often considered in a strongly confining magnetic trap with a small bias field. Even though this is different from the situation in our magnetic trap, it will be discussed first since certain concepts are more easily explained. In this trap, the cloud would be located at the minimum of the magnetic field (gravitational sag would not be important). Also, the transition frequencies between different magnetic sublevels would be very nearly degenerate. An rf field applied at a selected frequency will then drive transitions between different sublevels only at a particular magnetic field. The frequency is chosen so that transitions are driven at a magnetic field that corresponds to an energy that is several times the average energy of the cloud. Energetic atoms will have orbits that cross this ellipsoid-shaped transition surface. Once these atoms reach that

surface they undergo transitions to other magnetic sublevels that are not trapped. The atoms then leave the trap, being evaporated away. This idea is illustrated in Figure 4.27.

This evaporation will cool the atoms sample, which will in turn cause the evaporation rate to decrease as fewer atoms with enough energy to reach the transition surface are produced by elastic collisions. To prevent this slowdown in evaporative cooling, the rf transition frequency is changed as the cloud cools to correspond to an energy that is a fixed distance above the average energy of the cloud. Hence the forced nature of the evaporative cooling. The temperature of the cloud will continue to decrease, as will the number of atoms in the sample. As long as the atoms evaporated are in the high-energy tail of the thermal distribution, they will carry away much more than the average energy and the cooling will be rapid and efficient.

The removal of energetic atoms makes the trapped cloud spatially smaller. Therefore, the act of removing atoms in this way is often referred to “cutting” the atoms away and the transition point where the atoms undergo sublevel transitions is referred to as “edge” of the “rf-knife.” The terms “carving” and “shaving” are also applied when atoms are removed faster than the cloud can equilibrate to thermal equilibrium. These terms will be used in other portions of this thesis.

4.9.2 Additional Details about Forced Rf-Evaporation. There is a gloss in the simple picture presented in the previous subsection. While rf transitions can only be driven at a frequency that matches the difference in energy between magnetic sublevels, it is not obvious exactly what transitions will be made. Unlike scattering optical photons that lead to transitions to excited electron states, the rf transitions are so narrow that they cannot be treated in a simple rate-equation model. Once an atom is in resonance with the rf field, it will undergo rapid Rabi oscillations between different magnetic sublevels.

Rather than considering transitions between the bare magnetic sublevels in the atom, the rf transitions can also be examined in the dressed state picture [77]. As an atom moves in the magnetic trap, the magnetic field at its location changes and so the resonance frequency of rf transitions also changes. The atom will move along a dressed state corresponding to the initial magnetic sublevel. Normally, an atom that could be evaporated would be in an orbit that would likely have its turning point near the rf transition surface. Since the atoms are usually trapped in the most extreme




Figure 4.27. Radio-frequency transitions leading to evaporation. The black arrows represent where the rf transition being driven is resonant with the atoms. Atoms move in the parabolic trapping potential (upper curve). The more energetic can reach the magnetic field where they will make a transition to an untapped state and be accelerated out of the trap. Thus the more energetic atoms can be selectively removed and the cloud cooled through evaporative cooling.

magnetic sublevel, the Rabi oscillations between different sublevels that begin to occur near the transition surface weaken the trapping potential.

Two possible transitions can then happen that make the atom leave the trap. The first is that the weakening of the potential combined with the energy of the atom in its orbit carries the atom past the transition surface. For a strong enough rf field it would then stay in the same dressed state and make an adiabatic rapid passage to an untrapped state. This state would then be expelled from the trap. Another possible transition can occur if the atom does not have enough energy in its orbit to cross the transition surface but does get close to that surface. The atom will slow in its motion as it approaches a turning point and can make a diabatic transition to another dressed state. The magnetic moment of the atom can change at that time, and if the magnetic moment becomes weaker it will move further from the center of the trap and go through the transition region. If it does not make an adiabatic rapid passage back to a trapped state it will then continue out of the trap.

If the atom does make a transition back to the trapped state, it will remain in the trap. The net effect of the rf field will then be to increase the energy of the atom's orbital motion. As the atom oscillates in the trap field, it will pass the rf transition surface. As long as adiabatic transitions are made, the atom will remain in the trap. It will change from a trapped state to an untrapped state as it passes the transition point on the way back in to the center of the trap, but it will be converted back into a trapped state on its way back out. Eventually, it will make a diabatic transition and that will lead to it being lost from the trap.

The evaporation picture is more complicated in the case where the cloud of atoms sags in the magnetic trap due to gravity. Since the magnetic field lines cut vertically across the cloud (see Figure 4.15), the evaporation surface is not an ellipsoid as it would be in a trap with no sag but is rather just a plane. In contrast to the non-sagged case, atoms can then only be removed if they have enough orbital energy in the vertical direction alone. Atoms with large orbital energy in the other two directions will not be affected by the rf field. This means that evaporation will occur only along one direction rather than in all three directions, and that leads to a loss of efficiency in the evaporative cooling rate (see Chapter VII for the details of evaporative cooling along one-dimension and Chapter V for how this effect is used to create out of thermal equilibrium clouds).

One last complication is that since our magnetic traps are operated at relatively high magnetic fields, the transitions between different magnetic sublevels are no longer degenerate. That is not so important since all of the different sublevels have different locations of minimum potential energy. An atom that undergoes a transition will leave the region of the trapped cloud and be accelerated towards a different point in space. It will then move in a high-energy orbit and the anharmonicities of the trap make it unlikely that the atom will return to the location of the cold trapped cloud. As it oscillates in the trap it will move through different values of B and it is then possible that the rf can cause the atom to make transitions to other, truly untrapped states and the atom will then leave the trap.

4.9.3 Coupling the Rf Field into the Science Cell. An HP 8656B frequency synthesizer synthesizes the radio frequencies. The output of this synthesizer is sent to a 10 W, +40 dBm rf amplifier, whose output is then sent to a coil through about 2 meters of coaxial cable and a 50 Ω resistor. The coil has a one inch inner diameter and is placed right up next to the science cell so that it is approximately 1.5 cm from the center of the magnetic trap. It is composed of a single loop of thick (~3mm) copper magnet wire. The frequency and output amplitude of the synthesizer is controlled using the Labview code via a GPIB interface. The GPIB interface has the disadvantage of introducing variable delays on the millisecond level, making it difficult to measure any precise trap evolution of the cloud between the time the rf was turned on and the time it was turned off.

The single thick loop of wire was used because of difficulty coupling the rf power to the loop. The impedance matching of the coil is poor – we couple less than 1% of the power coming out of the amplifier into the loop (the remaining 99% of the power is reflected back into the amplifier itself). At the frequencies at which we operate, around 80 MHz and 130 MHz, the single thick loop of wire had better overall coupling than a thin wire loop consisting of a single or several turns. We tried using a variable capacitor and resistor to impedance match the coil. Away from the trapping coils, we were able to tune the frequency response of the coil to our desired parameters, but once the loop was placed in the apparatus, right up next to the baseball coils and near seven other coils, the mutual inductances present changed the frequency response. It should be possible to improve the coupling by adjusting the coil *in situ*, but the rf amplifier has survived the reflected power and we do

not have indications that more power to the atoms would improve our evaporation performance since the optimum power for evaporative cooling is well below (by at least a factor of ~ 20) the maximum available power.

The amplitude of the rf magnetic field at the atoms at the highest power setting that was used was measured to be ~ 14 mG. The way that this was measured was to apply the rf field and then ramp the trap bias magnetic field. The rf frequency and ramp were chosen so that the atoms made an adiabatic rapid passage to from the $F=2$ $m_f = -2$ to the $F=2$ $m_f = -1$ state. The trap strength was also adjusted so that the $F=2$ $m_f = -1$ atoms were not trapped. By measuring the trap loss as a function of ramp speed, the Rabi frequency (Ω) of the rf field could be measured [78], and was found to be $2\pi \times 10$ kHz. This can be related to the field strength through the relation $\Omega = \mu B / 2\hbar$, where the factor of two comes from the Clebsch-Gordan coefficients between the magnetic sublevels involved.

4.10 Calibration of the Magnetic Trapping Parameters.

4.10.1 Calibration of the Trap Bias Magnetic Field. The measurements that we performed required knowledge of the bias magnetic field in the magnetic trap. Since the bias magnetic field is proportional to the baseball and bias coil currents, it was necessary to calibrate the change in magnetic field with the change in the trap currents.

The fact that the atoms sag in the magnetic trap means that the bias magnetic field is only approximately linear with respect to the trapping currents. As the trapping currents are changed, the trap frequencies can change as well. If the radial trapping frequency changes, then the position of the atoms in the magnetic trap will be altered. While the magnetic field produced by a coil is proportional to the coil current at each individual point in space, the change in the position of the trap potential minimum means that the bias magnetic field at the minimum of the trap potential will not be proportional to the current.

In practice, we rarely change the radial trap frequency by more than 10% even when we range over one hundred gauss. This translates into a non-linearity of the current versus the bias field at the atoms of only a few percent over that entire range. In the region near the Feshbach resonance, the situation is even better. In that region, we arrange the baseball and bias coil currents so that a change in the baseball current will not change the radial trap frequency. The reason that this is

possible is that the radial trapping frequency depends both on the baseball coil current and the bias magnetic field. With the proper bias coil current, increasing the baseball coil current increases the bias field while simultaneously increasing the magnetic field curvature such that those two changes cancel each other to first order and the trap frequency remains nearly constant. The calibration of the trap frequencies is discussed below, but the nonlinearity of the bias magnetic field with respect to the baseball current results in an apparent change of slope of less than 0.1% over 10 G.

For a selected baseball and bias coil current, the magnetic field is measured in the following way. A cloud of atoms is evaporatively cooled to a reasonably cold temperature (~ 100 nK) and then the rf-knife is ramped closer in toward the center of the cloud. The final position of the rf-knife is varied and the number of atoms that remain is recorded as a function of the final frequency of the rf-knife. This measurement is repeated, except that now the rf-knife is ramped in the other direction (e.g. is the rf-knife was coming at the atoms from below, the direction is reversed so that the rf-knife comes at the atoms from above). Since the magnetic field lines are nearly horizontal planes, the intersection of the two curves of number remaining vs. final rf frequency then determines the rf frequency that drives magnetic sublevel transitions at the center of the trap. The Breit-Rabi equation (4.1) can then be used to determine the magnetic field. An example of this measurement is shown in Figure 4.28. We estimate that the uncertainty in the calibration of a single field value with a measurement like that in Figure 4.28 is 20 mG.

The intersection of the two curves in Figure 4.28 does not occur with zero atoms remaining. This is a consequence of the finite time of the ramp and the rf power in the rf-knife. Weak rf power is used so that there is not enough time to drive all of the atoms from the trap in the time of the ramp. Using stronger rf power not only produces shifts and broadening in the rf transitions themselves, it can lead to systematic shifts since the rf can be rectified in the baseball and bias current servos to produce an offset in the magnetic field. Because of the finite width of the rf-knife, both lengthening the ramp time and increasing the rf power will move the intersection point to lower, even negative (i.e the two curves in Figure 4.28 would not cross), number. These shifts due to the rf power have been observed to be up to 0.2 G for rf powers two orders of magnitude greater than the power used in the measurement shown in Figure 4.28. To make sure that the power was low enough not to




Figure 4.28. Magnetic Field Calibration. The point at which the two curves, ascending and descending, cross corresponds to the rf frequency that drives transitions at the bottom of the magnetic trap. Using the Breit-Rabi equation (4.1), the value of the magnetic field at the bottom of the trap can be computed.

cause any error in the calibration, the same field was measured with several different rf powers spanning three orders of magnitude. When this was done it was observed that the lowest powers in the data set all agreed with each other and deviations from that value grew with increasing rf power. The value of the field implied by the lowest powers was then used to calibrate the field.

If there were only a linear magnetic field gradient in the vertical direction, the bias field at the position of the minimum of the trap potential (i.e. the center of the trap potential) and the average field in a cloud would be exactly the same. However, there is curvature in the magnetic field in addition to the linear gradient, and so the average field in the cloud and the bias field at the position of the minimum of the trap potential will not be the same, and the difference between the two quantities will depend on the spatial extent of the cloud of atoms in the trap. For 100 nK clouds (like those used for the calibration of the field), the average magnetic field in the cloud is within 4 mG of the magnetic field at the minimum of the trap potential.

4.10.2 Calibration of the Magnetic Trap Frequencies. The trap frequencies are best characterized by sloshing measurements. The magnetic trap frequencies are suddenly changed in order to induce a center-of-mass motion in the cloud. Slosh is easily excited in the vertical direction by rapidly changing the radial trap frequency. The change in radial trap frequency changes the location of the minimum of the trapping potential because of sag, and the cloud is accelerated toward the new potential minimum. The slosh was induced in the axial direction by applying a bias field from the OP coil, although whether this is due to the field from the OP coil itself or a transient from the turn-on and turn-off of that current is unclear. After the slosh is induced, the center of mass of the cloud was measured as a function of time in the trap by absorption imaging. The resulting oscillation frequency is then equal to the trap frequency in whichever direction the slosh is being measured.

Since the absorption imaging only shows two directions in the trap, it is straightforward to measure the sloshing in the two visible directions but not obvious how to measure the sloshing sensitively in the third. It turns out, however, that anharmonicities in the trapping potential couple the motion in one radial direction with the other direction, causing energy to be transferred between the two radial directions. Therefore the oscillation frequency of one direction shows up in the

center-of-mass motion of the other radial direction. Careful inspection of the radial oscillation showed that it was in fact composed of two frequencies beating against one another, and those two frequencies are the frequencies of the two radial directions.

The fact that the radial trap frequencies are different from one another was expected. The gravitational sag breaks the radial symmetry of the trap. Calculations using the model described in subsection 4.5.4 suggest that the difference between the two frequencies should be $\sim 1\%$, when it is measured to be 1.4%.

To reduce the effect of anharmonicities in the trapping potential on the center-of-mass motion of the cloud, small (i.e. cold) clouds and small slosh amplitudes were used to measure the trap frequencies. A typical set of sloshing data is shown in Figure 4.29. The clouds used for this measurement were under 50 nK and the size of the slosh that was induced was less than 80 μm . The trap frequencies were measured at two fields, 156.6 and 163.9 G and over that range the frequencies changed by less than 0.5% if the baseball current was adjusted to change the trap bias magnetic field (see subsection 4.10.1).

To check the sloshing frequency measurements, the trap frequencies were also measured by parametric heating [79]. A sinusoidal modulation was applied to the bias coil current at twice the trapping frequency. Applying a modulation to the baseball current did not do any good since that did not change the trap frequencies to first order. The sinusoidal modulation led to an increase in energy of the atoms in the cloud, and atoms were then heated up and driven from the trap. By measuring the trap loss as a function of the modulation frequency, the trap frequency can be deduced since the most heating, and therefore loss, will occur at precisely twice the magnetic trap frequency. This method is not as precise as the sloshing measurement, but the two methods for determining the trap frequency agree with each other within their errors. It is also possible to parametrically drive radial sloshing motion through modulating the radial trapping frequency and hence the minimum of the potential in space. The largest sloshing amplitude will occur when the modulation frequency is equal to the radial trap frequency. This measurement of the trapping frequency agrees with all of the others.




Figure 4.29. Measurement of Sloshing Frequencies in the Axial and Radial Directions. The measurement of the axial trap frequency (top) and radial trap frequencies (bottom) was conducted at the same time. Each point represents a single measurement of the center of mass of the cloud while it sloshes in the trap. From the beat note pattern in the radial direction it is possible to extract all of three trap frequencies: 6.791(14) Hz axial and 17.259(8) and 17.493(17) Hz in the two radial directions. The last frequency corresponds to the direction that is not visible in the absorption imaging.

4.11 MOT Recapture and Absorption Imaging.

4.11.1 General Description of Atom Number Diagnostics. After the evaporation and manipulation of the atom cloud is completed, it is necessary to extract information about the cloud. One piece of information is the number of atoms in the cloud. This can be measured in two ways: through MOT recapture and through absorption imaging. In MOT recapture, the magnetic field is turned off and the MOT fields and trapping and hyperfine repump lasers are turned on. The fluorescence of the sample is then measured as described in subsection 4.4.6. In absorption imaging, a laser beam is sent through the cloud of atoms and is then focused onto a CCD camera. The laser beam is either on or near the resonance of the cycling transition, and so the atoms absorb and scatter light out of the laser beam. The amount of light absorbed is related to the column density of atoms in the cloud. By examining the amount of light absorbed, the number of atoms can be determined. In addition, the shape of the shadow that is cast by the atoms can be used to determine the spatial size of the sample as well, and from this size the temperature and density of the sample can be inferred.

Both of these measurements are destructive and so once the atom cloud is measured using MOT recapture or absorption imaging it is then necessary to create another sample. The MOT is destructive to the sample because the atoms rapidly come to an equilibrium distribution in the MOT that is not related to the distribution in the magnetic trap. In absorption imaging, many photons are scattered and so kinetic energy is imparted to the atoms via the photon recoils, radically heating up the cold samples being imaged.

4.11.2 Advantages and Disadvantages of MOT Recapture with Respect to Absorption Imaging. MOT recapture has some advantages over the absorption imaging. The signal-to-noise is better for moderately large numbers of atoms in the trapped cloud, and so it provides a more reproducible measure of the number of atoms in the cloud. Large atom densities can lead to lensing effects in absorption imaging (see subsection 4.11.6 below) that can distort the number measurement, and MOT recapture is immune to those considerations. Finally, there is no data processing involved with the MOT recapture measurement; all that is involved is reading off the

signal coming from a photodiode. This makes it a simpler diagnostic to use than absorption imaging.

There are disadvantages to MOT recapture, however. There is no way to determine the spatial size of the magnetically trapped cloud from MOT recapture, and so no way to determine the density and temperature of the cloud. It is also difficult to use MOT recapture to measure very small numbers of atoms. This is partly a signal-to-noise issue. The background noise due to scattered light fluctuations makes detecting small signals difficult, although this problem can be circumvented by modulating the trap laser frequency and using phase-sensitive detection. A larger problem comes from the imperfect vacuum in the science cell. Once the MOT is turned back on to recapture the atoms it starts to load a small (typically $\sim 10^6$ atoms/sec) number of atoms from the background Rb vapor. The noise on the signal due to these atoms makes measuring signals difficult if the number is below 100,000 atoms. Finally, calibrating the number of atoms in the magnetic trap via MOT recapture is not very accurate. There are uncertainties (30%) in the effective Clebsch-Gordan coefficients that should be used [62], the measured beam intensities, and fluorescence light collection solid angle. Due to these uncertainties, the MOT is calibrated using the absorption imaging. The MOT signal and absorption signal are compared and the relationship between the two is used to calibrate the MOT signal via the absorption signal calibration. Figure 4.30 shows such a calibration.

4.11.3 Optical Depth and Column Density. If resonant or near-resonant laser light is sent through the cloud of atoms, the atoms will scatter light out of the beam. Given the scattering rate in equation (4.3), it is possible to calculate the absorption of near-resonant laser light. Assume that the number of atoms N in a volume element d^3x is $n \cdot d^3x$ (and so n is the density). The number of photons scattered per second is $N \cdot R$, where R is given by equation (4.3). The flux of photons entering the volume is $A \cdot I/h\nu$ (I is the intensity, A is the area, h is planck's constant and ν is the laser frequency) and the flux of photons exiting the volume is $A \cdot (I-dI)/h\nu$. Assuming the laser propagating along the y -direction, this leads to the relation




Figure 4.30. Calibration of the MOT Recapture Signal. The non-zero intercept is due to the creation during evaporation of hot atoms that are too diffuse to be detected by absorption but are still trapped (see Chapter VII for details).

$$\begin{aligned} \frac{A \cdot (I - dI)}{h\mathbf{n}} - \frac{A \cdot I}{h\mathbf{n}} &= \frac{-A \cdot dI}{h\mathbf{n}} = N \cdot R = n d^3 x R \\ \Rightarrow \frac{dI}{dy} &= -nh\mathbf{n}R \end{aligned} \quad (4.27)$$

This relation is often written

$$\frac{dI}{dy} = -n\mathbf{s}I, \quad \mathbf{s} = \frac{h\mathbf{n}R}{I} = \frac{h\mathbf{n}}{2t} \frac{1}{I + I_{sat} \left(1 + 4 \left(\frac{\Delta}{\Gamma} \right)^2 \right)} \quad (4.28)$$

where σ is equal to $3\lambda^2/2\pi = \sigma_0$ in the low-intensity limit. This specification of σ_0 is for a two-level system, and is also the value that should be used for circularly polarized light driving a closed transition from the stretched state in a spin-polarized sample of atoms (e.g. σ^{\square} light driving the $F=3$ $m_f = -3$ to $F'=4$, $m_f' = -4$ transition).

In order to determine the amount of light absorbed by a cloud of atoms, it is only necessary to integrate equation (4.28) along the direction of the laser beam with the initial condition that the intensity of the light is I_0 at $y = -$

three widths, σ_x , σ_y , and σ_z , the intensity of light transmitted (I) is

$$\begin{aligned} \int_{I_0}^I \frac{1 + I/I_{sat} + 4(\Delta/\Gamma)^2}{I} dI &= -\mathbf{s}_0 \int_{-\infty}^{\infty} n(x, z) \frac{1}{\mathbf{s}_y \sqrt{2p}} \exp\left(\frac{-y^2}{2\mathbf{s}_y^2}\right) dy \\ \left(1 + \frac{4\Delta^2}{\Gamma^2} \right) \ln\left(\frac{I}{I_0}\right) + \left(\frac{I - I_0}{I_{sat}} \right) &= -\mathbf{s}_0 n(x, z) = -\text{OD}(x, z) \quad (4.29) \\ \text{where } n(x, z) &= \frac{N \exp\left(\frac{-x^2}{\mathbf{s}_x^2} - \frac{z^2}{\mathbf{s}_z^2}\right)}{2p\mathbf{s}_x\mathbf{s}_z} \end{aligned}$$

The quantity $\text{OD}(x, z)$ is the optical depth at position (x, z) . It is proportional to the column density of the atoms in the cloud at each point in space. The meaning of the optical depth is illustrated by writing the transmitted intensity (I) in the limit of low intensity and zero detuning:

$$I = I_0 \exp(-\text{OD}(x, z)) \quad (4.30)$$

Therefore the optical depth is simply the number of absorption lengths at position (x,z) in this limit.

The observed $OD_m = \ln(I/I_0)$ is related to the actual OD as

$$OD_m = \frac{OD + (I - I_0)/I_{sat}}{1 + 4 \frac{\Delta^2}{\Gamma^2}} \quad (4.31)$$

If the cloud does have a gaussian shape, then the number is related to the two visible widths and the peak optical depth, which is the optical depth at $x=0, z=0$:

$$N = \frac{2p OD_{pk} \sigma_x \sigma_z}{\sigma_0} \quad (4.32)$$

where σ_x, σ_z and σ_0 are all defined above.

4.11.4 Preparation of Trapped Cloud for Imaging and Production of Imaging Light.

While it is possible to image the clouds in the magnetic trap, turning off the trap allows a more straightforward interpretation of the absorption signal and a better signal-to-noise ratio. The better signal-to-noise ratio is due to the ability to scatter many photons per atom on the cycling transition, whereas in the trap there is no cycling transition for atoms in the $F=2, m_f = \pm 2$ magnetic sublevel. When the magnetic trap is turned off, it is important to maintain the atoms' spin polarization. This is accomplished by turning on the OP coil to produce a 1 G field along the probe laser direction while the current in the baseball and bias coils is being shut off. The atoms adiabatically follow the field, and so the spin polarization is preserved.

The inductance of the coils complicates turning off the magnetic trap. The rapid change in the current is opposed by a back emf since the magnetic energy in the coils (~ 0.2 J in the baseball) has to be dissipated. In addition, eddy currents may be created in the pieces of metal near the magnetic trap. While 99% of the current turns off in 160 μ s, there is a magnetic field that persists for a little over 1 ms. This residual magnetic field varies from day to day, making any sort of compensation difficult. Therefore, after the magnetic trap is turned off, the atoms are allowed to expand outside the trap for at least 1.6 ms to let the trapping magnetic fields die away.

Since the atoms are trapped in the lower hyperfine state, it is necessary to pump them into the upper hyperfine state so that the cycling transition can be used for imaging the cloud. After the

atoms have been allowed to expand out of the magnetic trap, the hyperfine repump light used for the MOT is turned on 100 μs before the probe light is turned on. Not all of the repump light is used, since half the light is on a path through the optics that would blind the camera and is therefore blocked. The 100 μs lead-time is enough to very nearly completely hyperfine pump the sample. The number measured by absorption imaging shows less than a 5% decrease when the lead-time is shortened to 10 μs . It is important to keep this lead-time short since the atoms scatter photons in the course of being pumped. A photon recoil imparts 6 mm/s to an ^{85}Rb atom, and a couple of photon scatters are required on average to pump the atoms to the $F=3$ state. In order for the motion induced by the scattering not to distort the cloud shape, it is important that the atoms not have much time to move before being imaged.

The polarization of the hyperfine pump light is exactly opposite that of the MOT beams because of the way that the hyperfine pump light is coupled into the MOT beam paths. The repump light therefore drives $\Delta m=0,\pm 1$ transitions and depolarizes the atoms in the cloud. The atoms are pumped from the $F=2$ $m_f=-2$ state to a range of sublevels in the $F=3$ manifold. Since the branching ratio of the $F=3$ state is such that the atoms have a 55% chance to fall into the $F=3$ once excited, most of the atoms are hyperfine pumped after one or two scatters and remain toward the negative end of the $F=3$ sublevel manifold. Some fraction of the atoms will be pumped towards the positive m -levels, however.

The OP/probe laser described in subsection 4.4.9 produces the probe laser light. The probe light is σ^{\square} polarized, and its intensity is adjusted to be 320 $\mu\text{W}/\text{cm}^2$. After the probe light emerges from the optical fiber, it is then sent through an iris with an aperture of ~ 2 mm diameter. The small aperture was used to reduce forward scattering from dust and other scattering centers along the laser beam path that went around the shadow of the atom cloud to the camera (see subsection 4.12.5). The iris position was adjusted until the probe beam passing through the iris was centered on the point in space where the atoms were magnetically trapped. The AOM is used to control the probe pulse length, which was typically 40 μs , and the pulse intensity. The action of scattering photons during the imaging heats the atoms up and so short pulses are required in order to avoid distorting the shape of the atom cloud. The 40 μs pulse length represents a compromise between distorting the

atom cloud shape and an accurate measurement of number. Shorter pulses (10 μs) are used when the spatial size of the condensate is of primary interest rather than the number.

The laser is tuned close to or on resonance with the $F=3$ to $F'=4$ cycling transition. If the atoms are in the $m_f = -3$ sublevel, then if the laser is perfectly σ^{\square} polarized the atoms would scatter photons on a completely closed transition since the only state the $F'=4$ $m_f = -4$ can decay to is the original $F=3$ $m_f = -3$ level. The atoms would then scatter photons on this transition repeatedly until an off-resonant transition occurs allowing the atoms to return to the $F=2$ ground state. The atoms then have to be repumped from the $F=2$ state. On resonance, typically ~ 125 photons per atom are scattered in 40 μs . Any atoms not in the $F=3$ $m_f = -3$ state originally will be optically pumped to that state after several photon scatters. Any small imperfections in the probe polarization mean that the transition between magnetic sublevels is not perfectly closed, but atoms that make a transition to a different state will again be quickly optically pumped back.

4.11.5 Imaging the Atoms' Shadow onto a Camera. Once the probe laser passes through the atom cloud, the atoms will have scattered light out of the beam, appearing to cast a "shadow" in the beam. This shadow is magnified by a series of lenses and focused onto a CCD array in a camera (a Roper Scientific camera with a TEK 512x512 Back Illuminated CCD array). The series of lenses is shown in Figure 4.31. After the laser passes through the atoms, the shadow in the beam is collected and nearly collimated by the first lens. The second converging lens and the diverging lens focus the shadow onto the CCD array in the camera. The first lens is mounted so that a micrometer can be used to adjust its position.

This arrangement of lenses is designed to look at small objects. The magnification of the imaging system shown in Figure 4.31 is a factor of ten. The depth of focus is calculated in the way shown in Figure 4.32 and is found to be 40 μm . When larger clouds were studied, different lenses were used that had a smaller magnification and greater depth of focus.

The optical access to the cell is limited by the presence of the magnetic trapping coils, and so the windows used for the MOT light are also used for the imaging system. In order not to block the MOT beams, the camera used in absorption imaging cannot be in the MOT beam paths. Therefore, the probe beam path is deflected from the MOT beam path. To do this with a minimum




Figure 4.31. Schematic of the Imaging Optics.

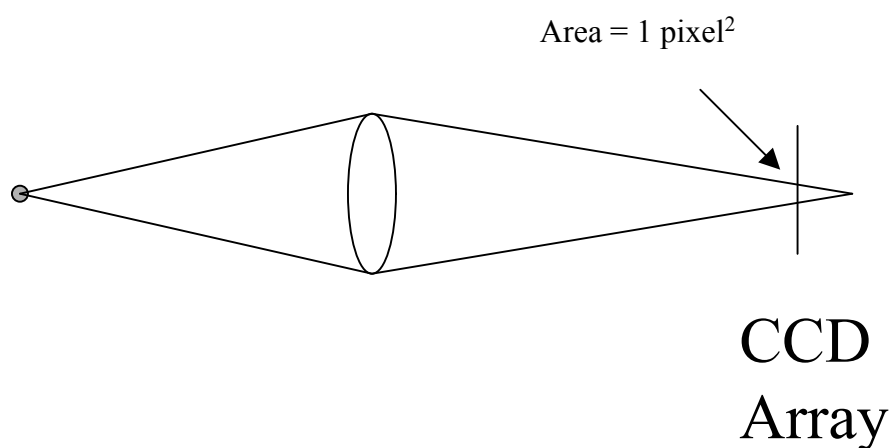


Figure 4.32. Depth of Focus Calculation. The depth of focus of the imaging system is calculated using geometric optics. The rays that originate at a position offset from the focus the imaging system are traced through the edge of the first lens. At the CCD array, these rays form a circle and the area of this circle can be calculated. The depth of focus is then twice the distance that the object is displaced from the focus when the area of the circle of rays on the CCD array is equal to the area of one pixel. This estimate will actually give a depth of focus that is slightly too small since not all of the rays from the object pass through the edge of the lens. Two-thirds of the lens radius is a more appropriate quantity since it is the average radius through which all of the rays pass.

loss of light, the probe beam and co-propagating MOT beam are set to have opposite polarizations and a polarizing beam splitting cube is used to deflect the probe beam into the camera. This limits how closely a lens can be positioned to the location of the atom cloud to ~ 10 cm. This consideration is important because often the clouds that are to be imaged are small and it is desirable to get as good an imaging resolution as possible. Also, to save space on the table the camera is mounted on a post above the table. The farther away from the atoms the first lens is, the farther the camera must be placed above the table for the same magnification. The diverging lens in the beam path (see Figure 4.31) is there primarily to prevent the camera from having to be put any higher above the table.

The diffraction limit of a 1-inch diameter lens placed 10 cm away is $3.7 \mu\text{m}$ [80]. We have tried several different makes of lenses for this initial lens. We used Melles-Griot achromat lenses, high-quality camera lenses, and achromat-meniscus lens pairs. The resolution limit was always about the same, $\sim 7.5 \mu\text{m}$ FWHM, which is worse than the diffraction limit, although the achromat-meniscus lens pair worked a bit better than the other lenses that we tried. There are several surfaces between the lens and the atoms, the glass window of the cell, a $\lambda/4$ waveplate and the polarizing beam splitting cube, and it could be imagined that these lead to a degradation of the resolution of the system. However, off-line tests showed that the performance of the lenses was still not at the diffraction limit when imaging small objects even in the absence of these surfaces, although the off-line tests indicated a resolution limit that was about 20% better than what was observed in the imaging of the atom clouds. While better resolution would be nice, we have evidence of features that are less than $1 \mu\text{m}$ that would be of interest in the BECs. Given that this is the approximately the wavelength of light that we are using to image the atoms, no improvements in resolution would allow us to see all the features that we would like to see.

The camera was positioned at the focus of the lenses coarsely by imaging the atoms in the science MOT. Since the science MOT is centered on the magnetic trap center, this is an adequate place to start. Absorption images of the MOT would have been tedious to take due to the slow readout and processing of the main CCD camera. The camera takes around one second to read out the entire array (less if only reading out part of the array) and it takes tens of seconds to process an image. Rather than taking absorption images of the MOT then, a slightly more elaborate procedure

was used. The MOT was imaged using a single lens onto a smaller CCD array (an inexpensive infra-red sensitive camera) connected to a standard video monitor. This smaller CCD array position was easily adjusted to bring some edge in the MOT into a sharp focus. A piece of paper with lines on it was then placed right in front of the main camera and the position of the main camera was adjusted until the lines of the paper were in focus on the smaller CCD array where the MOT had been focused before. This placed the main camera approximately at the focus of the lens system. For fine adjustments, small BECs were used.

Because of the Feshbach resonance, it is possible to set the atom-atom interactions to zero and then the BEC has the extent of the ground state of a simple harmonic oscillator. For our trapping conditions this is an rms size of $1.9 \mu\text{m}$ ($4.5 \mu\text{m}$ FWHM) in the radial direction and $3.0 \mu\text{m}$ ($7.0 \mu\text{m}$ FWHM) in the axial. The radial direction is below the resolution limit, and these small clouds are used to focus the lens system and to measure the resolution limit. The lens system is focused by adjusting the position of the first lens via the micrometer so that the cloud appears as small as possible. A plot of the radial size of the cloud vs. micrometer position is shown in Figure 4.33.

The standard absorption measurement of the cloud consists of three separate camera images. In the first image, the probe beam with the shadow cast by the atoms is recorded. The second image records just the probe beam without the atoms present. The last image is taken with the probe AOM off to measure the background light. The last image is subtracted from both the first and second image, and the ratio of the resulting intensities is used to produce a measured optical depth. This subtraction and division is done on a pixel-by-pixel basis. The processed pixels form a 2D array showing the optical depth (and hence the column density) as a function of position. The absorption image thus contains information about the spatial size of the cloud and the number of atoms in the cloud. Usually the processed image of optical depths was fit with a 2D gaussian via a χ^2 minimization to determine the radial and axial widths and the number of atoms in the cloud. The number of atoms cloud could also be calculated without any fit by summing all of the measured optical depths in all of the pixels.




Figure 4.33 Focusing the imaging system. The data in this graph are the apparent radial width of a sub-resolution limit cloud as a function of the position of the first lens in the imaging system. The probe laser frequency was set to be on resonance to avoid any lensing distortions. The focus point is taken to be the position of the minimum size of the radial width.

BECs are not expected in general to have precisely gaussian shapes, but were still fit with gaussians both for ease of processing and because the images of the BECs resembled gaussians most of the time. The widths and peak optical depth of the gaussians will still correlate to the actual number of atoms in the BEC and its spatial size. The number measured by the gaussian fits was compared to the sum of the optical depths in all of the pixels to determine a correction (which was typically 10% for BECs). A more flexible fitting function would be appropriate to remove the need for this correction, and improvements in the image processing are currently being made.

One difficulty with the images is the presence of interference fringes on the probe beam. The fringes form a regular sinusoidal interference pattern on the probe beam with an amplitude of up to 30% of the average intensity and an apparent spatial frequency of $\sim 15 \mu\text{m}$. The reason that this is an apparent spatial frequency is that the fringes are almost certainly not in focus. As long as the fringe pattern is stable, they present no difficulty. However, the fringe pattern appears to move slightly between the first and second images (it could be the camera moving), and when the difference and ratio of intensities is taken a residue of the improperly subtracted fringes then appears on the processed optical depth array, distorting the image of the atom cloud. Fortunately, as long as the spatial size of the cloud being imaged is different from the spatial frequency of the fringes, the gaussian fits tend to smooth out the effect of the fringe distortion in calculating the width and the number of atoms in the cloud.

These interference fringes were due to two sources. One was the window in front of the camera's CCD array, and centering the cloud of atoms on a part of the array where the fringe amplitude was a minimum could reduce the effect of those fringes. This did not eliminate all of the fringes, however, and the remaining fringes likely come from the science cell, waveplates, or lenses in the imaging path. The fringes in the unprocessed images had a contrast ratio of $\sim 7\%$, and this resulted in fringes with a typical size of 0.1 peak-to-peak in apparent optical depth due to the improper subtraction of the fringes from the unprocessed images. The contrast ratio of the fringes in the processed images varied from shot-to-shot, from no observable fringes to 0.2 peak-to-peak apparent optical depth.

Because of the regular nature of these fringes, it is possible to filter them out of the processed data digitally. The 2D array is first Fourier transformed to spatial-frequency space. The interference fringes then appear as a bright dot in the Fourier transform. Since the cloud spatial size and fringe spacing are different, it is possible to set all the spatial-frequency space around the interference fringe to zero. The data is then Fourier transformed back to real space and the original image is recovered minus the fringes. Fits to the data before and after the transform show that this does not significantly distort the measured cloud shape or number.

The signal-to-noise ratio of the imaging system is limited by the shot noise due to the finite number of photons impinging on each pixel. There is a background noise associated with the variation of stray light from the room being detected by the camera and the noise on the camera dark counts. This background noise was equal to an rms size of 2.6 counts per pixel for our typical conditions. The number of counts in a pixel is proportional to the number of photons that strike that pixel during the camera exposure time. The shot noise per pixel was equal to $\sim 0.5 * (\text{counts})^{1/2}$, and the intensity used corresponded to around 300 counts per pixel. There was no evidence of any component of the noise that was directly proportional to the number of counts. The signal-to-noise ratio would therefore improve with increased probe duration or increased probe intensity. However, increasing the probe intensity leads to undesirable saturation of the atomic transition and Doppler shifts as discussed below. Increasing the probe duration leads to Doppler shifts and a distortion of the cloud shape through heating. Therefore, the probe intensity and duration were chosen to comprise between these undesirable effects and the signal-to-noise ratio. As an example, the calculated signal-to-noise ratio due to the probe intensity noise in the determination of the number of atoms in a ~ 5000 atom BEC expanded to have a peak optical depth of 1.2 is 125 for these compromise conditions.

4.11.6 Lensing Effects From Dense Atomic Samples. When the probe laser is not tuned exactly onto resonance, in addition to absorbing the laser light, the cloud also possesses an index of refraction that is different from one [81]. The variation in density of the cloud then means that the index of refraction varies in space and so the cloud can act like a lens, bending the wavefront of the laser beam. If the cloud is in focus, nothing happens unless the lensing is so severe that some of the

atoms' shadow is bent outside the collection angle of the first lens. If the cloud is out of focus, then the lensing action can lead to serious distortions in the observed shape of the cloud. In fact, these distortions can be used as an alternate way to make sure the atoms are at the focus of the imaging system. If the cloud shape looks the same when the probe laser is blue-detuned as it does when the probe is red-detuned then the cloud is in focus, provided that the depth of focus is bigger than the cloud being imaged.

Normally we avoided these distortions by tuning the probe laser on resonance with the cycling transition, making the index of refraction of the cloud equal to one.

4.12 Calibration of the Imaging System.

4.12.1 Calibration of the Magnification. In order to determine the size of an imaged object, the magnification of the imaging system must be measured. In other words, when an image is formed on the CCD array, an individual pixel will have an apparent size that is simply the actual size of the pixel divided by the magnification of the lens system. This apparent size is referred to as the pixel size. The way that this is done is to release a cloud of atoms from the magnetic trap and watch it fall. The center-of-mass motion of the cloud is then recorded as a function of expansion time. Since the cloud accelerates under the force of gravity, the acceleration of the center of mass of the cloud can be used to calibrate the pixel size. Figure 4.34 shows a typical pixel size calibration measurement.

One potential systematic in this measurement is the presence of a magnetic field gradient that provides acceleration to the atoms in addition to the acceleration due to gravity. The gradients would have to be significant compared to 26 G/cm, and this is a fairly strong gradient and so it is not expected that there would be any stray fields that were even a few percent of this size. Even so, to be sure that this was not a problem the pixel size calibration was also checked in a different way. The trap frequencies were measured for two different radial trap strengths. Knowledge of the radial trap frequencies in both traps, the distance (in pixels) that the center of the trap moved, and the relationship between sag and radial trap strength provided a pixel size calibration that was in agreement with the normal gravity drop measurement.




Figure 4.34. Using Gravity to Calibrate the Magnification of the Imaging Optics. By releasing the atoms from the magnetic trap and tracing the center of mass motion, from the measurement of the acceleration of the falling atoms due to gravity and the knowledge of g , the apparent size of a pixel (i.e. the magnification of the optics) can be determined. The black circles show the measured center-of-mass motion while the solid line is a fit to the data.

4.12.2 Probe Laser Frequency Stability and Single Mode Behavior. The frequency of the probe laser needs to be known in order to calculate the scattering rate and hence the cross section. In addition, the short pulse time requires that the stabilization of the probe laser frequency be good from DC to nearly 40 kHz. The short pulse time does not allow any variations in this range to be averaged down during the pulse itself and so any variations will lead to shot-to-shot noise. The frequency of the laser is stabilized using the DAVLL system to provide an error signal, and the associated servo system stabilizes the laser frequency well except for a long-term DC drift in the DAVLL signal. This long-term drift is monitored by referencing a saturated absorption signal.

The frequency of the probe laser is the most critical of all the lasers in the apparatus, and so more time has been spent improving and characterizing its servo system. Both the PZT and laser current receive feedback from the servo system, and the PZT is used to compensate for long term drifts while the current is used to suppress faster variations. We found that some care was necessary in the design of the servo system. All op-amps have an intrinsic gain-bandwidth limit, and the inclusion of gain stages as large as 1000 slowed the servo system down significantly, limiting the bandwidth to a few kHz. Breaking up the gain stages into smaller pieces and using better op-amps (AD 829) improved the servo performance. The frequency noise was characterized by using a saturated absorption system separate from the DAVLL system. By measuring the total variation of the saturated absorption signal the frequency noise was measured to be less than 0.4 MHz peak-to-peak.

To monitor the long-term drift, part of the probe laser beam is overlapped with a beam from a reference laser that is peak-locked to a saturated absorption peak. The two combined beams are sent to a fast-response photodiode and the beat frequency (~ 60 MHz) between the two beams is measured. The drift in the DAVLL signal can then be monitored in real time and the servo set point is occasionally adjusted to compensate for this drift.

The beat frequency where the laser is in resonance with the cycling transition is determined by examining the peak optical depth of a cloud of atoms released from the trap as a function of probe laser frequency. Figure 4.35 shows such a measurement. Since absorption imaging is destructive, a




Figure 4.35. Calibration of the Probe Detuning. The measured optical depth of a trapped cloud versus the detuning of the probe laser is used to determine the resonance of the $F=3$ $m_f = -3$ to $F'=4$ $m_f = -4$ transition in experimental units. The x axis shows the beat note between the probe laser and another laser locked to the peak of a saturated absorption transition. The variation in the points appears to be worse than it actually was; there was a drift in the size of the trapped clouds during the measurement that was easily removed when the center and width of the lineshape were being determined. The width of the transition from simple power broadening is expected to be 6.6 MHz, consistent with the measured width.

sample must be created through evaporative cooling for each measured point. The location of the peak of this absorption curve is consistent with a field of 1 G, which is what is expected from the optical pumping coil. The 6.78(34) MHz width is consistent with the natural width of the cycling transition when power broadening is included (see the next subsection).

One other concern with respect to the frequency of the probe laser is that the laser should be single-mode. Probe light that is in a different mode is not near the cycling transition and therefore cannot be scattered by the atoms. This light will distort the optical depth measurement since it will not be properly subtracted away when the optical depth is calculated. The amount of light in other modes can be determined by using a standard Rb cell. The probe laser is tuned near the peak of the Doppler absorption profile, and then the temperature of the cell is increased. Eventually, the increased Rb density in the cell causes the absorption of all of the light that is in the mode near the Rb transitions. The remaining transmitted light then indicates the amount of light in the other modes (typically $\sim 2\%$). A correction to the optical depth measurement can be applied based on this measurement, but the effect of the light in other modes produces the same effect as forward scattering around the atoms described in subsection 4.12.5 below, and so the two corrections are usually lumped together.

4.12.3 Measurement of the Probe Intensity and the Doppler Shift. Because the atomic transition can be saturated, the probe intensity must be determined for an accurate calculation of the column density. The determination of the probe intensity by directly measuring the size and power of the probe beam is complicated by the presence of several absorbing surfaces between the atoms and the camera. Instead, the peak optical depth is measured as a function of probe intensity and from a fit to the saturation of the atomic transition the probe intensity is calibrated. Figure 4.36 shows such a measurement.

There are several features of interest of the curve in Figure 4.36. First, the peak optical depth does decrease as the power in the probe is increased, and the fit, assuming that the atomic transition is simply saturating, agrees with the data reasonably well. Second, the data are consistent with an increase in measured optical depth as the intensity is turned up from zero. This type of




Figure 4.36. Measured Optical Depth vs. Probe Intensity. This figure shows the effect of increasing probe intensity on the saturation of the atomic transition. By fitting the decrease in measured optical depth to equation (4.31), the intensity of the probe can be determined, calibrating the number of camera counts to the intensity of the probe at the location of the atoms. Each point consists of seven individual measurements, and the standard deviation of those seven measurements is used to calculate the errors shown. Note that the size of the fractional error decreases with increasing intensity. The atomic clouds used for this calibration were thermal clouds with a temperature of 115 nK. The probe laser was slightly detuned from resonance (1.1 MHz) when these data were taken. To avoid the potential systematic from optical pumping effects (subsection 4.12.4), only the four highest-intensity points were fit.

behavior is expected from optical pumping effects discussed in the next subsection. Finally, the variation in measured optical depth decreases by a factor of 1.5 as the probe power is increased.

There are a few reasons why the variation in the measured optical depth should decrease as the probe power is increased. The transition will be saturated and power broadened, and so the sensitivity to variations in the probe frequency should be reduced. Over the range of the probe power used, however, the saturation and broadening is not enough to explain the reduction in noise. The signal-to-noise ratio should improve as the probe intensity is increased, but the noise expected from the shot-noise of the photons counted is only 0.8 % at $I/I_{\text{sat}} = 0.2$ and so this does not explain the observed variation of 5%. The actual number of atoms is varying from shot to shot, but that variation is independent of the probe power. The most plausible explanation is that this variation is due to variations in the optical pumping of the sample. The distribution of atoms in the $F=3$ magnetic sublevels will vary and the number of photons required to pump the atoms over to the $F=3$ $m_f = -3$ state will vary as well. The more photons that are scattered, the better optically pumped the sample and so these variations should decrease with increasing probe power.

Rescattering of probe photons in an optically thick cloud and the acceleration of the atoms due to the momentum transfer from the probe beam will also lead to a decrease in the measured optical depth as the power of the probe is increased. The fact that the cloud is optically thick means that some of the photons scattered out of the probe beam will rescatter off some of the other atoms in the cloud. These rescattering events cause the atoms to spend more time in the excited state, effectively reducing the saturating intensity. The size of this estimated reduction is $\sim 25\%$ for a cloud with an optical depth of 1.2. The act of scattering photons also accelerates the atoms since the probe is directed from just one direction. Since 6 mm/s on average is imparted per photon scatter, after 100 scatters the atoms are moving at 60 cm/s and the Doppler shift v/λ is 0.8 MHz. Increasing the power in the probe beam increases the number of scattering events that then leads to larger Doppler shifts. This leads to a distortion of the measurement of the saturation intensity as well. Exactly on-resonance, the correction would only be $\sim 3\%$ for the 40 μs pulse length and the range of intensities covered in Figure 4.36. Since the probe detuning was slightly (1.1 MHz) off resonance, the correction of the saturation intensity is then $\sim 10\%$. The correction to the saturation intensity due

to the Doppler shift was estimated by integrating equation (4.3) over the length of the pulse with the Doppler shift due to the time-dependent velocity of the atoms included in the detuning term Δ . The time-dependent velocity at time $t=t_0$ was assumed to be equal to the number of photons scattered between $t=0$ and $t=t_0$ times the average velocity imparted by a single photon. This meant that the scattering rate R of equation (4.3) and the time-dependent velocity were coupled equations, and so they were solved as such. The number of photons scattered in the optically thick regions of the cloud were smaller and so the Doppler shift was less of an effect for those regions. Appropriately averaging over all of the optical depths in the cloud produced the correction estimates listed above. The peak optical depth of the clouds in the estimate was assumed to be 1.2.

Since the optical depth of the clouds used to measure the saturation intensity are roughly the same as the optical depth of the clouds that were studied in the course of our investigations, the effective saturation intensity measured by the data in Figure 4.32 can be used to find the actual column density with no additional adjustments. The effective saturation intensity can be compared to the probe intensity determined by the shot noise and the two agree to within 25%. This is a reasonable agreement given the effects of optical thickness of the clouds, the small Doppler shifts and the absorption of the probe light by the glass cell, waveplate, polarizing beam splitting cube, and lenses between the atoms and the camera CCD array.

4.12.4 Optical Pumping Effects in Absorption Imaging. As was alluded to in the previous sections, the distribution of atoms among the magnetic sublevels of the $F=3$ state has an impact on the measured optical depth. The reason for this is that the oscillator strength of a particular absorption transition is proportional to its Clebsch-Gordan (CG) coefficient (see Figure 4.3). From the $F=3$ $m_f = -3$ state to the $F'=4$ $m_{f'} = -4$ the CG coefficient is one, and so any other transition will have a smaller transition probability and therefore a smaller transition rate. This smaller scattering rate leads to an undercount of the number of atoms in the cloud.

After the atoms are pumped from the $F=2$ to the $F=3$ state in preparation for the imaging, they occupy a range of magnetic sublevels. The population of the atoms will be predominantly in the negative m_f levels since the atoms start in the $F=2$ $m_f = -2$ state, but they will certainly not all start in the $F=3$ $m_f = -3$ state. As the atoms scatter photons, they will be pumped by the σ^{\square}

polarization of the probe laser towards the $F=3$ $m_f = -3$ state. The CG coefficients are such that several photon scatters are required before the atoms are pumped into this state. The following table shows the average number of photons that must be scattered to pump the atoms to the $m_f = -3$ state and the amount of time that process takes.

<i>Initial m_f state</i>	<i>Avg. Number of Photon Scatters</i>	<i>Time (scaled to scattering rate from the $F=3$, $m_f = -3$ state</i>
-2	4.0	5.3
-1	5.9	9
0	7.2	12.7
1	8.1	17.7
2	8.9	26
3	9.5	50.4

Table 4.2. The average number of scattering events required to pump an atom from a particular initial m_f state to the $m_f = -3$ state. Also shown is the length of time these photons scatters would take. The time column is how many photons an atom already in the $m_f = -3$ state would scatter on the $F=3$, $m_f = -3$ to $F'=4$ $m_f = -4$ transition in the time that it took to pump the atom from a different m_f state over. For example, and atom initially in the $m_f = 0$ state would be pumped to the $m_f = -3$ state in an average time equal to the amount of time it would take an atom initially in the $m_f = -3$ state to scatter an average 12.7 photons.

The amount that an atom will be undercounted can be determined from the values in Table 4.2. For instance, assume that the probe pulse is on long enough to scatter 100 photons from the $m_f = -3$ state. If an atom begins in the $m_f = -2$ state, then the atom will on average scatter 4.0 photons while being pumped into the $m_f = -3$ state and then scatter an additional 94.7 photons (on average) and so it will only scatter 98.7 photons in total and will be undercounted by 1.3%. An atom initially in the $m_f = -1$ state will be undercounted by 3.1%, and an atom in the $m_f = 3$ state will be undercounted by 40% on average. The longer and more intense the pulse, the more photon scatters that will occur and the less the atoms will be undercounted. The times listed in Table 4.2 do not include the consideration of the role the ~ 1 G quantization field plays in removing the degeneracy of

all of the transitions, and so the pumping times listed are faster (by up to 30% for the most positive m_f states) than the true pumping rate.

The values quoted in Table 4.2 are only average quantities, and when the average time that it takes to pump an atom starts to become comparable to the pulse length, the amount of photons that atom scatters will be highly variable. The atoms farther from the $m_f = -3$ state will be more variable. If a scattering event occurs early in the pulse that pumps the atoms toward the $m_f = -3$ state then the atom can scatter a large number of photons. If it does not get pumped over until late in the pulse, not many photons will be scattered. Of course there are many atoms in the sample and so these variations will be averaged down, but we suspect that the variations in optical depth that we observed during the I_{sat} measurement are brought on in the variation of the optical pumping of the atoms during the imaging pulse. The exact calculation of the expected variability is complex and the size of the variability difficult to estimate. However, it does have the right time scale and intensity dependence to explain the observed variation.

Since the optical pumping effects will definitely cause an undercount of the atom sample and because we have observed variability that seems to indicate that the atoms have not been completely pumped to the $m_f = -3$ state early in the probe pulse, a correction is applied to the number of atoms measured. This correction is estimated to be 5%, based on the observed variability in the measured optical depth.

4.12.5 The Maximum Observable Optical Depth and Forward Scattering. The probe laser light can scatter not only from the atoms but also any dust or imperfections in any of the optical elements of the imaging system. This scattering will be strongly peaked in the forward direction, and it will limit the maximum observable optical depth of the cloud during absorption imaging. Scattered probe light from near the edge of the shadow cast by the atoms can fill in the shadow at the camera and make the optical depth appear smaller than it actually is. The smaller the spatial size of the cloud of atoms the easier it is for light to come from the edge and fill in the shadow at the center of the cloud (since the angle of the forward scattering would then be less). Using an iris in the probe beam to make the beam smaller reduces this effect dramatically, by a factor of 6 or more in the

amount of light that leaks into the shadow. Using a higher-magnification lens system makes the effect worse.

This effect is measured by creating a very optically thick cloud and then imaging it. The measured optical depth is observed to saturate around the center of the cloud, and the saturated value determines the amount of light that is scattered into the shadow of the cloud. The maximum observable optical depth is referred to as the “OD ceiling.” The OD ceiling decreases as the size of the cloud decreases. It is typically around 3.5 for thermal clouds and 3.0 for spatially large condensates (13 μm rms width). It has been observed to be as small as 2 for spatially smaller condensates. This effect is indistinguishable from the presence of light not in the main mode of the probe laser and so measuring the OD ceiling is really measuring both effects.

It is possible to model the OD ceiling as light that cannot be absorbed that is present in the probe beam. This model of the OD ceiling predicts that its presence will lead to an underestimate of the number of atoms that is typically around 10% for the large condensates discussed above. Based on this model, a correction is applied to the measured number to account for this effect.

4.12.6 Quantization Field Orientation. The magnetic sublevels only have meaning if they are defined with respect to some axis, referred to as a quantization axis. In our apparatus, that axis is set by the magnetic field created by the OP coil. The polarization of the probe laser is also defined with respect to this axis. In order for the probe to drive only $\Delta m = -1$ transitions, it must both propagate directly along the quantization axis and be perfectly circularly polarized. Otherwise, in the atoms’ frame it will be possible to drive $\Delta m=0, +1$ transitions as well, albeit with a lower probability. This would lead to an underestimate of the atom number for two reasons. The CG coefficients for the $\Delta m=0, +1$ transitions are not as large and so the photon scattering rate will be less than for $\Delta m = -1$ transitions alone. Also, the $\Delta m=0, +1$ transitions will optically pump the atoms away from the $m_f = -3$ state and while the atoms are being pumped back over they will scatter less photons.

The polarization of the probe is set by the birefringence of the $\lambda/4$ waveplate and its orientation since the probe light is linearly polarized before going through the waveplate. The waveplate birefringence was measured off-line and found to be within $1/30^{\text{th}}$ of $\lambda/4$, close to quarter

wave. The waveplate on the other side of the cell was measured at the same time and found to have a birefringence equally as close to quarter wave. The first waveplate's orientation was adjusted so that the optical pumping efficiency was maximized, which is sensitive to the depolarizing nature of imperfectly polarizing light. The probe light is extinguished to the 1% level by the polarizing beam-splitting cube, and so the polarization should be circular to that level. Calculations show that this 1% level is reasonable since small deviations of the birefringence from $\lambda/4$ and small deviations in the orientation of the waveplate scale with the deviation squared and so are small.

Rather than change the probe direction, it is easier to check whether or not the OP field was aligned well enough along the direction of the probe laser. This is done by measuring the optical depth as external fields were applied to make sure that it is near the maximum. The results are shown in Figure 4.37, and indeed the quantization field is well aligned. This result was confirmed by altering the strength of the quantization field. If there is some stray field perpendicular to the quantization field, then the field should become misaligned as its strength is adjusted downward and better aligned as the field strength is increased. After adjusting the probe frequency to compensate for the change in the atoms' resonance frequency as the quantization field was changed, we found no variation in the measured number with quantization field strength, again indicating that the field was well aligned.

4.12.7 Resolution Limit. The finite resolution limit of the optical imaging system makes it difficult to interpret the images of objects roughly the size of the resolution limit or smaller. Objects smaller than the resolution limit will appear to have a width close to the size of the resolution limit. Therefore, the spatial size of any object smaller than the resolution limit cannot be accurately determined.

The resolution limit also has an impact on the observed number of atoms. The diffraction from a small object will cause the "shadow" cast by the object to become larger than the size of the collection lens in the imaging system. The camera will then miss some of the shadow and the number of atoms will be undercounted. The apparent number decrease is not as severe as the limitation on the apparent width of the cloud. The size of the apparent number decrease can be estimated by Fourier-transforming the cloud shape and cutting out the spatial frequencies that will




Figure 4.37. Checking the Orientation of the Quantization Axis During Absorption Imaging. Each data point in this figure (filled circle) represents the ratio of the number of atoms measured at a particular up/down earth shim current to the number of atoms measured with that current equal to 0.2 A (open circle) during the absorption sequence. The decrease in measured atom number is due to the misalignment of the quantization axis caused by the application of the field from the up/down earth shim, which is orthogonal to the main quantization field. The solid line is a parabola fit to the data. This parabola is a reasonable function since we expect that the measured number of atoms should be a maximum when the quantization axis is aligned along the probe direction. The orientation of the quantization axis was checked similarly in the other direction that is orthogonal to the main quantization field. The typical operating condition was 0.2 A, close to the peak of the curve, which fits to a maximum ratio of 1.0015.

not be captured by the collection lens. For instance, a cloud with a size that is 70% of the resolution limit in the radial direction and 2.5 times that in the axial (the aspect ratio of our trap) is only predicted to be undercounted by 8%.

The decrease in measured number for small clouds can be determined by using the mean-field interaction to change the size of condensates created under similar conditions and then examining the apparent number of atoms in the condensates as a function of their spatial size. In this measurement, however, not only will the resolution limit of the optics play a role but the combination of higher peak optical depth and lower optical depth ceiling for the smaller condensates will also limit the measured number. It is difficult to separate the effect from the resolution limit from the optical depth limit. For condensates with a radial FWHM of 11 μm and an axial FWHM of 22 μm and peak optical depth of around 2, we see a $\sim 10\%$ decrease in the measured number over condensates with larger sizes, larger than what would be expected from the simple estimate of the apparent number decrease performed in the previous paragraph.

The resolution limit only plays a role when small clouds are being imaged. In order to avoid undercounting the number of atoms in a small cloud, the size of the cloud is expanded above the resolution limit. See Chapters IX and X for more details.

4.12.8 Estimate of Absorption Imaging Accuracy. The measurement of the number of atoms in a sample is notoriously difficult. The reason is that there are many things that can affect the absorption signal, like all those discussed in the preceding subsections. We have estimated the effects of all the considerations listed above on the number, but are not confident that we know all of the effects perfectly. Therefore, it is necessary to assign an uncertainty to our measurement of the number. Based on the typical size of the effects that we do know about, we estimate that our uncertainty in the number calibration is 10%. A summary of all of the effects considered is presented in the table below.

Absorption Imaging Consideration	Effect on the Absorption Imaging	Thesis Subsection, Relevant Equations
Imaging Optics Magnification Calibration	Necessary to convert a measured OD to a column density and hence an atom number. Determined with 1% precision.	Subsection 4.10. Equation(4.29), (4.32)
Probe Laser Frequency	Necessary to compute the scattering rate. Measured stability is ~0.4 MHz.	Subsection 4.12.2 Equation (4.3), (4.31)
Probe Laser Light in Non-resonant Modes	Limits the maximum observable optical depth. See optical depth ceiling.	Subsection 4.12.2
Lensing due to Dense Clouds	Distorts the shape of small, dense clouds. Avoided by imaging on-resonance.	Subsection 4.11.6
Probe Intensity	We typically operated the probe laser at 15-25% of the saturation intensity. The resulting saturation of the atomic transition must be included in the calculation of the scattering rate.	Subsection 4.12.3 Equation (4.31)
Doppler Shift and Photon Rescattering	These two effects reduce the effective saturation intensity by ~30% of the value calculated for a single non-moving atom. Will cause the apparent saturation intensity to depend mildly on the optical depth of the cloud.	Subsection 4.12.3
Optical Depth Ceiling	Forward scattering limits the maximum observable optical depth. A 10% correction to the measured atom number is applied to account for this effect.	Subsection 4.12.5
Optical Pumping	Time required to spin-polarize the sample during the imaging pulse reduces the scattering rate and increases the imaging noise. A 5% correction is applied to account for this effect.	Subsection 4.12.4
Quantization Field Orientation	Quantization magnetic field must be aligned along the probe axis in order to avoid undercounting. The field is aligned so that any error from this effect is nulled to better than 1%	Subsection 4.12.6
Probe Laser Polarization	Same issues as with the quantization axis. Errors nulled to better than 1%.	Subsection 4.12.6
Resolution Limit	Affects the measured spatial size and atom number of small clouds ($< 7 \mu\text{m}$ FWHM)	Subsection 4.12.7

Table 4.3. Summary of Effects Considered in Absorption Imaging.

This concludes the Chapter on the details of the experimental apparatus. Throughout the rest of the thesis, reference will be made back to the appropriate subsections of this chapter when any relevant calibrations of the system are being discussed.

Chapter V

Resonant Magnetic Field Elastic Scattering in Cold ^{85}Rb

5.1 Introduction

5.1.1 External Electromagnetic Fields and Very-Low-Temperature Atomic Collisions.

Very-low-temperature collision phenomena can be quite sensitive to applied electromagnetic fields. Several groups have altered inelastic collision rates in optical traps by applying laser fields [82]. There have been numerous proposals [20,21,83,84,85,86] for using laser and static electric and magnetic fields to influence cold collisions. Of particular interest to this work are the predictions made in Ref. [21] that there should be Feshbach resonances in collisions between cold ($\sim\mu\text{K}$) alkali atoms. These Feshbach resonances allow dramatic alterations of the inelastic and elastic scattering rates as a function of an external field. Since the collisions being considered are cold enough, the elastic collisions can be characterized by a single-parameter, the s-wave scattering length a (see subsection 2.4.4). Near a Feshbach resonance, a varies dispersively as a function of the applied external field, and can, in principle, take on any value from $-\infty$ to $+\infty$ (see Figure 5.1). Since the elastic scattering cross section is $8\pi a^2$, changing a will change the elastic scattering rate as well. The origin of these Feshbach resonances was discussed in Chapter II.

While attempts had been made to observe magnetic-field-dependent Feshbach resonances in ^{133}Cs and ^{87}Rb [87,88], Feshbach resonances in collisions between alkali atoms had not yet been observed when we began searching for one in ^{85}Rb . In retrospect, the failure of these early searches for Feshbach resonances is not surprising because there was an enormous uncertainty as to the precise magnetic field locations and widths of the predicted resonances. This large uncertainty was due primarily to the fact that the position and width of a Feshbach resonance is very sensitive to the details of the atom-atom interaction potential. After these initial searches, results of photoassociation spectroscopy greatly improved the knowledge of the alkali atom-atom interaction potentials, allowing these resonances to be predicted with much less uncertainty [89]. As a result,




Figure 5.1. This figure shows the approximate shape of the Feshbach resonance. It is calculated using the measured parameters for ^{85}Rb and equation (5.1).

several groups, including ours, independently observed Feshbach resonances between alkali atoms within the span of several months: in ^{23}Na [90], ^{85}Rb [91,92], and ^{133}Cs [93].

The methods used to observe the resonances were diverse. In ^{23}Na , the Feshbach resonance was observed with optically trapped Bose-Einstein Condensates (BECs). There were two signatures of a Feshbach resonance in ^{23}Na : a change in the expansion rate of a BEC released from the trap due to a change in scattering length, and a large loss rate that is potentially due to molecular formation [94,95,96,97]. Whatever the cause, the loss rates were large enough to preclude any studies of the BEC behavior near the divergence of the scattering length and in the region where $a < 0$. In ^{133}Cs , the samples were also optically trapped, but instead of using BECs over 30 Feshbach resonances were located by examining the elastic and inelastic scattering rates in thermal clouds. Characterizing more than 30 Feshbach resonances would be a tremendous task for us, but the all-optical method of creating cold samples in the Stanford work represents an improvement on the order of 100 in duty cycle over our work, resulting in a factor of 10^4 improvement in integration time. In other words, what would take us a month took them less than an hour. In the ^{85}Rb work in Ref. [91], a Feshbach resonance was inferred from the observation of a magnetic-field dependent change in photoassociation spectra of magnetically trapped thermal clouds. In our work, a Feshbach resonance in ^{85}Rb was measured by examining the elastic scattering cross section in magnetically trapped clouds, work that is detailed in the rest of this chapter.

5.1.2 Motivation. The fact that no Feshbach resonances had been observed in cold alkali collisions was certainly one reason that we began our search for one in ^{85}Rb , especially in light of the improved predictions for the resonance location and width. While these Feshbach resonances are interesting examples of cold collision physics in and of themselves, our primary motivation in fact lay elsewhere. The foremost reason we were searching for a Feshbach resonance was due to the importance of the scattering length in determining the properties of a BEC.

The role of the scattering length in the mean-field energy of a BEC was discussed in Chapter III. Many of the properties of a BEC depend on the mean-field energy, and so the ability to tune it through an external field opened up exciting possibilities for experiments in Bose-Einstein Condensation. The mean-field energy could be made to be very strong, and expected deviations

from mean-field theory could be tested (Chapter VIII). Also, the atom-atom interactions could be made to be attractive enough that the BEC was no longer stable and BECs in this unstable region could be studied (Chapter IX and X). The presence of a nearly resonant molecular state may also play an interesting role in the behavior of BECs near a Feshbach resonance as well [94-97]. In order to perform any of these experiments, however, it was necessary to first locate a Feshbach resonance to use to tune the s-wave scattering length. In ^{85}Rb , a prominent Feshbach resonance was predicted to exist at experimentally accessible fields, and this was the resonance that we found. From now on, I will refer to this particular Feshbach resonance as “the” Feshbach resonance.

Another reason that the elastic scattering was studied as a function of field is that the evaporative cooling of ^{85}Rb is also more complicated than that in ^{23}Na and ^{87}Rb . Knowledge of both the elastic and inelastic scattering rates as a function of magnetic field was required in order to cool the ^{85}Rb atoms to BEC temperatures (see Chapter VI and VII). Also, away from the Feshbach resonance ^{85}Rb was predicted to have a large negative scattering length, which meant that any more than a few (~ 70) atoms were expected to be unstable unless the scattering length could be changed.

Finally, the sensitivity of the Feshbach resonance to the interatomic interaction potential meant that a measurement of the location and width of the resonance would then tightly constrain the parameters used to characterize the interatomic interaction potential. In addition to simply improving the knowledge of the cold collisions potentials for their own sake, the s-wave scattering length as a function of magnetic field also needs to be known in order to interpret the results of BEC experiments near the Feshbach resonance.

5.1.3 Characterization of Atom-atom Interaction Potentials. The Rb-Rb interatomic potential is discussed in section 2.9. There are primarily three parameters that once known specify the atom-atom interaction potential for cold elastic collisions and allow the calculation of elastic scattering cross sections. These three parameters can be represented as the singlet scattering length (a_S), triplet scattering length (a_T), and the C_6 coefficient, which is the leading coefficient in the dispersion force expansion (see equation (2.26)) [30]. Instead of the parameters a_S , and a_T , alternative specifications of the triplet and singlet potentials can be used such as the triplet and singlet quantum defects [31] or the fractional bound state number [91]. In addition to these three

parameters, the additional terms in the dispersion force expansion (such as C_8 and C_{10}) also play a role, but the elastic scattering cross section is not as sensitive to these parameters.

The s-wave scattering length as a function of magnetic field (B) near the Feshbach resonance is also approximately characterized by three parameters: the position of the resonance, the width of the resonance, and the background scattering length (a_{bg}), which is the scattering length far from the Feshbach resonance. The position of the resonance is defined as the magnetic field (B_{peak}) where the divergence in scattering length is located. The width (often written as Δ) is defined as the difference in magnetic field between the peak and the where the scattering length is equal to zero (B_{zero}). In terms of these three parameters, the scattering length as a function of field is [30,90]

$$a = a_{bg} \left\{ 1 - \left(\frac{B_{zero} - B_{peak}}{B - B_{peak}} \right) \right\} \quad (5.1)$$

Since the interatomic potential is specified by three parameters, the measurement of a_{bg} , B_{peak} , and B_{zero} in turn determines a_T , a_S , and C_6 . The relationship between these two sets of parameters is complex, however. For our purposes, the three parameters a_{bg} , B_{peak} , and B_{zero} are most useful since usually our only concern is the value of a at a particular magnetic field. Once a_{bg} , B_{peak} , and B_{zero} are known, a_T , a_S , and C_6 can be determined from numerical calculations and the resulting potential used by theorists for any cold collision calculations requiring the interatomic potential.

Equation (5.1) is only approximate, but the approximation is better than 0.5% near the Feshbach resonance (155.3 to 180 G), which is sufficient precision for our purposes (and certainly smaller than the error in the determination of the various interatomic potential parameters). Over the range of magnetic fields from 180 G resonance out to 250 G the approximation is still good to ~1 %, except near another narrow width (~ 0.01 G) Feshbach resonance somewhere near 230 G (which has not been observed).

The three parameters a_{bg} , B_{peak} , and B_{zero} are each directly measurable experimentally, and this chapter details our measurement of these three quantities.

5.2 Cross-dimensional Mixing and Thermal Equilibration.

5.2.1 Elastic Collisions and Enforcement of Thermal Equilibrium. In a gas that is not initially in thermal equilibrium, elastic collisions will cause the gas to come to thermal equilibrium.

For any gas of atoms, the gas reaches thermal equilibrium once the entropy of the sample is maximized. Since elastic collisions cannot on average decrease the entropy of the gas, but can increase the entropy, after a sufficient number of elastic collisions the entropy will be maximized and the gas will have reached thermal equilibrium. The time that a gas takes to reach equilibrium is therefore determined by the elastic collision rate.

One feature of a magnetically trapped gas in thermal equilibrium is that it has an aspect ratio (axial width divided by radial width) that depends on the frequencies of the trapping potential. If the aspect ratio of the cloud is changed from this equilibrium value, then elastic collisions will force it back to the equilibrium ratio. The rate that it achieves this equilibration can then be used to determine the elastic scattering rate, and hence the elastic scattering cross section, of the collisions between the atoms in the gas [98]. This is in fact how we determine the s-wave scattering length as a function of magnetic field.

In principle, any two directions of the trap that do not have an equilibrium aspect ratio can be used to measure the equilibration rate, but the presence of anharmonic terms in the potential can obscure the equilibration rate between the two radial directions. Since the two radial frequencies are close to one another, it is possible for atoms' orbits to precess between the two directions (see subsection 4.6.1). While it is true that in the absence of elastic collisions the entropy of the cloud cannot increase, the precession of the atoms' orbit can mimic an equilibration of the gas. Atoms in higher-energy orbits precess out of one direction and into the other. If only the spatial distribution of a cloud with different initial average radial energies is measured, then the direction that the atoms have precessed into (out of) will look wider (narrower) and the gas will appear to be equilibrating even without elastic collisions. If the velocity distribution is measured carefully as well as the spatial distribution, then in principle the true equilibration can be separated from this anharmonic mixing. In practice, this would also be a difficult prospect since the spatial and velocity distributions would have to be measured precisely and that is far more difficult than simply measuring a width.

Measuring the equilibration rate between the axial and radial directions minimizes the difficulties due to the presence of anharmonic mixing. Because of the frequency differences between the axial and radial directions, the amplitude that an individual atoms' orbit can precess due

to anharmonicities is very small. Therefore, any change in the aspect ratio between the radial and axial directions is due to elastic collisions rather than anharmonic terms in the trapping potential. Anharmonic mixing still occurs in the radial direction however, and that can lead to complications in both creating samples that are out of thermal equilibrium and measuring the aspect ratio as the cloud equilibrates.

5.2.2 Calculation of the Equilibration Rate. Since the equilibration rate is proportional to the elastic scattering rate, it must be proportional to the elastic scattering cross section and therefore a^2 . If only the relative equilibration rate at different fields is being compared, then this is more or less all that needs to be known about the relationship between the equilibration time and scattering length. In order to determine the absolute value of the scattering length, however, the relationship between the collision rate and the equilibration time must be known. In other words, the average number of elastic collisions necessary to bring an atom to thermal equilibrium must be calculated.

The calculation of the average number of elastic collisions required to equilibrate an atom, from now on referred to as κ , is complicated by the presence of several factors of two. The elastic collision cross section is computed between two atoms of a particular collision energy, while in the trapped gas collisions are occurring at a range of energies and it is necessary to sum over those energies. In the case of a gas of spin-polarized bosons, the identical nature of the atoms and the spin statistics need to be accounted for without double-counting the number of collisions. Finally, the correlation function for bosons in a thermal state needs to be considered as well. Because of some confusion over these factors of two in the past, the calculation of κ is treated here in detail.

The calculation of κ is performed using Enskog's equation [76], which computes the rate of change due to elastic collisions of the average of any quantity of the gas using classical mechanics.

If χ is the quantity being considered, then Enskog's equation is written

$$\frac{d\mathbf{c}(\bar{\mathbf{x}}, t)}{dt} d^3x = \frac{1}{2} \int d^3v d^3v' f(\bar{\mathbf{x}}, \bar{\mathbf{v}}, t) f(\bar{\mathbf{x}}, \bar{\mathbf{v}}', t) |\bar{\mathbf{v}} - \bar{\mathbf{v}}'| \frac{d\mathbf{s}}{d\Omega} \Delta \mathbf{c}(\bar{\mathbf{v}}, \bar{\mathbf{v}}' \rightarrow \bar{\mathbf{q}}, \bar{\mathbf{q}}') d\Omega d^3x \quad (5.2)$$

where $\chi(\mathbf{x}, t)$ is the average value in the volume d^3x of the quantity χ at position \mathbf{x} and time t , $f(\mathbf{x}, \mathbf{v}, t)$ is the position and velocity distribution function, $d\sigma/d\Omega$ is the differential cross section, and $\Delta \mathbf{c}(\bar{\mathbf{v}}, \bar{\mathbf{v}}' \rightarrow \bar{\mathbf{q}}, \bar{\mathbf{q}}')$ is the change in the quantity χ when two atoms scatter from velocities \mathbf{v}, \mathbf{v}' to

\mathbf{q}, \mathbf{q}' . The integration over $d\Omega$ represents the integration over all of the possible final velocities \mathbf{q}, \mathbf{q}' given the initial velocities \mathbf{v}, \mathbf{v}' . Equation (5.2) is really just a sum over all the collisions at point \mathbf{x} weighted by the change in the quantity χ . The first term is the number of particles at \mathbf{x} in a volume d^3x , the $f(\bar{x}, \bar{v}', t)|\bar{v} - \bar{v}'|$ term represents the flux of particles that can collide with those at \mathbf{x} , $d\sigma/d\Omega$ sets the rate of atoms that scatter in the $d\Omega$ direction, and $\Delta\chi$ is the change in χ for a scatter in that direction.

In order to make sure that all of the factors of two are accounted for properly, I found it helpful to consider the scattering from two beams of particles incident on one another. Most scattering calculations are performed with this situation in mind, and so it is easier to get all of the factors of two straight. Since we have worked closely with Chris Greene and Jim Burke, Jr., I am following all of the conventions that Jim used in his thesis [27] so that our specification of the s-wave scattering length is the same. Consider two beams of indistinguishable, spinless particles with Bose statistics that have a uniform density throughout space (n_A and n_B). The center-of-mass frame is set to be the same as the lab frame so that one beam has a velocity $v_0/2$ and the other $-v_0/2$. The collision energy is low enough that only s-wave scattering need be considered.

Since the particles have Bose statistics, the scattering wavefunction must be even under an exchange of the two particles that have scattered off one another. In the center-of-mass frame this means that the differential cross section must be [23]

$$\frac{dS}{d\Omega} = |f(\mathbf{q}, \mathbf{f}) + f(\mathbf{p} - \mathbf{q}, \mathbf{f} + \mathbf{p})|^2 \quad (5.3)$$

For s-wave scattering, in which case $f(\theta, \phi)$ is independent of θ and ϕ , this results in a factor of four enhancement of the differential cross section as compared to distinguishable particle scattering. Equation (5.3) also means that only even angular momentum partial waves will have nonzero differential cross sections (due to the odd parity of odd angular momentum states). Therefore the two lowest partial waves for identical bosons are s- and d-wave scattering and p-wave scattering is not allowed. Note that in order for two particles to be indistinguishable it is not only necessary that both particles be the same species of atom or molecule, both particles must also be in the same

internal (i.e. spin) state. Since we are only concerned with elastic scattering in a spin-polarized gas of bosons, I will continue to concentrate on indistinguishable particle scattering.

While the differential cross section is enhanced by a factor of four from the indistinguishable to distinguishable case, the total cross section is not. The differential cross section represents the flux of particles through a detector at angle $d\Omega$ (see section 2.2). For distinguishable particles, it is possible to tell the two particles involved in a single scattering event apart, at least in principle, and so there will be a flux associated with each type of particle. For indistinguishable particles there is no meaningful way to detect two separate fluxes even though two particles are involved in every scattering event. In order to make indistinguishable particle total cross sections look more like distinguishable particle total cross sections, the total elastic scattering cross section is specified by integrating the differential cross section over 2π steradians, rather than integrating over all space (4π steradians). Note that this specification is a matter of convention; the number of particles which are scattering into an angle $d\Omega$ is rigorously set by equation (5.3). The enhancement of a factor of four in the differential cross section for indistinguishable particles then is reduced to a net factor of two enhancement in the total cross section. This remaining net factor of two enhancement is the reason that the total elastic scattering cross section for indistinguishable bosons is equal to $8\pi a^2$ rather than $4\pi a^2$ presented in equation (2.15).

One additional specification concerning indistinguishable particles must be made. The incoming flux can be defined with respect to one or both incoming beams. Following the convention adopted in Ref. [27], the incoming flux is defined with respect to a single beam. Returning to the two-colliding-beams example, the flux can then be considered to be $n_A v_0$ into a volume of size d^3x . The number of scattering centers in that volume is then $n_B d^3x$. Since the two colliding beams example being discussed is considered in the center-of-mass frame, the selection of the flux as $n_A v_0$ and the number of targets as $n_B d^3x$ was made arbitrarily. The flux and number of targets could have been equally as well specified by $n_B v_0$ and $n_A d^3x$ respectively. It is also possible to explicitly consider the collisions in a lab frame where one beam is moving and the other is stationary with no change in the calculation for the value of κ . From the discussion in the preceding paragraphs and equations (5.3), (2.15), and (2.15), the differential cross section is $4a^2$. Therefore,

the scattering rate into the angle $d\Omega$, which is equal to the incoming flux times the differential cross section times the number of scatters, is $(n_A v_0)(4a^2)(n_B d^3x)$. Integrating over all of space, the total number of particles that are scattered per unit time is $16\pi a^2 n_A n_B v_0 d^3x$.

The total kinetic energy in the z-direction of both beams at any point in space before any elastic collisions is easily computed to be $(n_A + n_B)mv_0^2/8 d^3x$, where the z-direction is taken to be along the beams' collision axis. After some particles have elastically scattered, this energy will be reduced, because the elastic scattering will convert velocities along the z axis to velocities along the x and y axes as well. Since the scattering is isotropic, the final collision energy will on be on average equally distributed between all three axes, and therefore on average a single scattered particle loses $2/3(mv_0^2/8)$ in z-direction kinetic energy. Combining this average loss rate with the scattering rate derived in the previous paragraph, the change in total z-direction kinetic energy per unit time is then

$$16\pi a^2 n_A n_B v_0 \frac{2}{3} \frac{mv_0^2}{8} d^3x = \frac{4\pi}{3} a^2 n_A n_B mv_0^3 d^3x \quad (5.4)$$

This same loss rate of kinetic energy from the z-direction can be computed using Enskog's equation (5.2). By matching the solution from Enskog's equation to the result in (5.4), all of the factors of two will be properly accounted for and then Enskog's equation can be used to compute more complex situations.

First, the distribution function needs to be specified for the two incident beams

$$f(\vec{v}) = n_A \mathbf{d}^{(3)}\left(\vec{v} - \frac{v_0}{2} \hat{z}\right) + n_B \mathbf{d}^{(3)}\left(\vec{v} + \frac{v_0}{2} \hat{z}\right) \quad (5.5)$$

For now the position and time labels have been dropped since they are not needed. $d\sigma/d\Omega$ is isotropic, but its amplitude will be allowed to float for now. The initial z-direction kinetic energy of a colliding pair is $mv_0^2/4$ (twice the energy of a single atom), while the final z-direction kinetic energy is $(mv_0^2/4)\cos(\theta)^2$, where θ and ϕ are the angles the final velocities make with the z axis. With these definitions, it is possible to evaluate equation (5.2) and find

$$\frac{dKE_z}{dt} d^3x = \frac{d\mathbf{s}}{d\Omega} \frac{2\pi}{3} n_A n_B mv_0^3 d^3x \quad (5.6)$$

Comparing equation (5.6) with equation (5.4) shows that the differential cross section should be set to be $d\sigma/d\Omega = 2a^2$ in order for Enskog's equation to calculate the change in z-direction kinetic energy.

Before calculating κ , there are two simplifying assumptions that are made. The first assumption is that the position and velocity distribution in the trapped sample can be described by three parameters, an effective temperature in each direction. This assumption is partly justified by the fact that the spatial distribution of atoms in the clouds in the experiment itself was gaussian even at early times, even though the energy in different directions is not the same. It is also reasonable to assume that the kinetic and potential energy in each direction are equal and therefore can be specified by a single parameter, since otherwise the cloud would exhibit breathing motion, which is unlikely to be created in the cloud the way we experimentally put the clouds out of thermal equilibrium and also is not observed. In fact, any deviations from a gaussian shape would lead to oscillations in the shape of the cloud, and such deviations are again not observed and so this first assumption is well satisfied.

The second assumption that is made is that the cloud is sufficiently dilute that an atom undergoes several oscillations in the trap between collisions. The reason for this assumption is that the elastic collisions in the cloud only randomize the energy associated with the velocity of the colliding atoms, not the potential energy associated with their position since that cannot be changed by a collision. However, since the kinetic energy of an atom is turned into potential energy as it moves in the harmonic potential (and vice versa), if the atom can make several oscillations before the next collision, then the potential and kinetic energy will be well "mixed" and the elastic collisions can be considered to be randomizing the total energy of the atom rather than just the kinetic energy.

In order to estimate the validity of this assumption, it is possible to use a very crude model that captures some of the physics involved in the second assumption. This simple model consists of a set of differential equations

$$\begin{aligned}
\frac{d^2 KE_r}{dt^2} &= -\frac{1.25}{\tau_{me}} \left(\frac{1}{2} \frac{dKE_r}{dt} - \frac{dKE_z}{dt} \right) - 4\mathbf{w}_r^2 \left(KE_r - \frac{1}{2} E_r \right) \\
\frac{d^2 KE_z}{dt^2} &= -\frac{1.25}{\tau_{me}} \left(2 \frac{dKE_z}{dt} - \frac{dKE_r}{dt} \right) - 4\mathbf{w}_z^2 \left(KE_z - \frac{1}{2} E_z \right) \\
\frac{dE_r}{dt} &= -\frac{1.25}{\tau_{me}} \left(\frac{1}{2} \frac{dKE_r}{dt} - \frac{dKE_z}{dt} \right) \\
\frac{dE_z}{dt} &= -\frac{1.25}{\tau_{me}} \left(2 \frac{dKE_z}{dt} - \frac{dKE_r}{dt} \right) \\
PE_i &= E_i - KE_i \quad i = r, z
\end{aligned} \tag{5.7}$$

where KE_r and PE_r are the total kinetic and potential energy in the two radial directions, KE_z and PE_z are the total kinetic and potential energy in the axial direction, and τ_{me} is the equilibration time in this model. This simple model captures two pieces of physics. First, energy will oscillate between kinetic and potential at twice the trap frequency, as is expected. Second, the transfer of energy from the radial to the axial direction occurs only through the equilibration of kinetic energies. For sufficiently slow τ_{me} , the equilibration rate is independent of τ_{me} . Once τ_{me} approaches the trap oscillation period, however, the finite time it takes for the kinetic energy to be changed to potential alters the actual equilibration rate. For a sense of scale, when $\tau_{me} = 0.08$ s in a trap with radial frequency $2\pi \times 17.5$ Hz and axial $2\pi \times 6.7$ Hz, there is an 8% change in the equilibration time in this crude model. This simple model also captures the fact that for fast equilibration times (i.e. times shorter than the trap periods) the energy of the cloud actually oscillates in a collective motion, although it does not get the frequencies correct at all [99]. The second assumption is expected to break down for the large scattering lengths near the peak of the Feshbach resonance, and the relaxation to thermal equilibrium will become limited by the time scale of the magnetic trap periods.

With these assumptions in place, it is now possible to evaluate Enskog's equation in order to determine the equilibration rate. For now it is assumed that the energy in the two radial directions is the same. With the two assumptions made above, it is possible to consider the quantity χ as the sum of the potential and kinetic energy in the z direction and to integrate equation (5.2) over all of space as well

$$\frac{dE_z}{dt} = \frac{1}{2} \int d^3x d^3v d^3v' f(\vec{x}, \vec{v}, t) f(\vec{x}, \vec{v}', t) |\vec{v} - \vec{v}'| \frac{d\mathbf{s}}{d\Omega} \Delta E_z(\vec{v}, \vec{v}' \rightarrow \vec{q}, \vec{q}') d\Omega$$

$$f(\vec{x}, \vec{v}, t) = \frac{N \mathbf{w}_r^2 \mathbf{w}_z m^3}{8 \mathbf{p}^3 k_B^3 T_r^2 T_z} \exp\left(\frac{-m}{2kT_r} (v_r^2 + \mathbf{w}_r^2 r^2) + \frac{-m}{2k_B T_z} (v_z^2 + \mathbf{w}_z^2 z^2)\right) \quad (5.8)$$

Now, the evaluation of equation (5.8) will be carried out in some detail. The position and velocities are separable, and so it is possible to integrate over the position

$$\frac{dE_z}{dt} = \frac{N^2 \mathbf{w}_r^2 \mathbf{w}_z}{128} \left(\frac{m^3}{\mathbf{p}^3 k_B^3 T_r^2 T_z}\right)^{\frac{3}{2}} \int d^3v d^3v' \exp\left(\frac{-m}{2k_B T_r} (v_r^2 + v_r'^2) + \frac{-m}{2k_B T_z} (v_z^2 + v_z'^2)\right) |\vec{v} - \vec{v}'| \frac{d\mathbf{s}}{d\Omega} \Delta E_z \quad (5.9)$$

Since the elastic collisions cannot change the center-of-mass motion of a colliding pair, it is useful make a change of variables from \mathbf{v}, \mathbf{v}' to \mathbf{vs}, \mathbf{vd} , which are defined so that

$$vs_i = (v_i + v'_i), \quad vd_i = (v_i - v'_i), \quad i = r, z \quad (5.10)$$

With this substitution, it is then possible to integrate over \mathbf{vs} , careful to include the factor of 8 associated with the change of variables

$$\frac{dE_z}{dt} = \frac{N^2 \mathbf{w}_r^2 \mathbf{w}_z}{128} \left(\frac{m^3}{\mathbf{p}^3 k_B^3 T_r^2 T_z}\right) \int d^3vd \exp\left(-\frac{m}{4k_B T_r} vd_r^2 - \frac{m}{4k_B T_z} vd_z^2\right) \sqrt{vd_r^2 + vd_z^2} \frac{d\mathbf{s}}{d\Omega} \Delta E_z \quad (5.11)$$

The $d\Omega$ integral is evaluated next. A little algebra shows that a pair of atoms colliding with z -velocities v_z, v_z' have an energy in the z -direction of $m(v_z^2 + v_z'^2)/4$. Again, the potential energy is ignored because it cannot be changed during a collision. Since the center-of-mass motion cannot be changed by the collision, the change in z -direction energy depends only on the square of the initial and final difference velocities (vd_z and vd_z'). It is again advantageous to make a change of variables, this time to spherical coordinates. In that case, the $d\Omega$ integral is

$$\int d\Omega \frac{d\mathbf{s}}{d\Omega} \Delta E_z = \int \sin(\mathbf{q}') d\mathbf{q}' d\mathbf{f}' (2a^2) \frac{m}{4} vd^2 (\cos^2(\mathbf{q}') - \cos^2(\mathbf{q}))$$

$$= 2\mathbf{p} a^2 m vd^2 \left(\frac{1}{3} - \cos^2(\mathbf{q})\right) \quad (5.12)$$

where the differential cross section has been set to $d\sigma/d\Omega = 2a^2$. The only thing left to do is integrate over difference velocities. The integration over ϕ can be carried out easily, and so what remains to be evaluated is

$$\frac{dE_z}{dt} = \frac{N^2 \mathbf{w}_r^2 \mathbf{w}_z m^4 a^2}{96 \mathbf{p} k_B^3 T_r^2 T_z} \int \sin(\mathbf{q}) d\mathbf{q} \int d\mathbf{v} d\mathbf{v}' (1 - 3 \cos^2(\mathbf{q})) \exp\left(-\frac{m \mathbf{v} d^2}{4k_B} \left(\frac{\sin^2(\mathbf{q})}{T_r} + \frac{\cos^2(\mathbf{q})}{T_z}\right)\right) \quad (5.13)$$

The $d\mathbf{v}d\mathbf{v}'$ integration over all relative velocities can be carried out

$$\frac{dE_z}{dt} = \frac{2N^2 \mathbf{w}_r^2 \mathbf{w}_z m a^2}{3\mathbf{p}(1-\mathbf{b})^2} \int \sin(\mathbf{q}) d\mathbf{q} (1 - 3 \cos^2(\mathbf{q})) \left(\frac{\sin^2(\mathbf{q})}{(1-\mathbf{b})} + \cos^2(\mathbf{q})\right)^{-3} \quad (5.14)$$

where $T_r = (1-\beta)T_z$. Before evaluating equation (5.14), the equation will be recast into a different form so that the role of the collision rate will be more apparent. Setting $\sigma_0 = 8\pi a^2$, $E_z = Nk_B T_z$, $3T_0 = T_z + 2T_r$, and using the expressions for the density (equation (4.12)) and average relative speed (equation (4.17)) with the relation $T_r = (1-\beta)T_z$ means that equation (5.14) is equivalent to

$$\frac{dT_z}{dt} = \frac{-1}{\hat{\epsilon}} \langle n \rangle \langle \mathbf{s}_0 \rangle \langle v \rangle (T_z - T_0) \quad (5.15)$$

$$\frac{1}{\hat{\epsilon}} = \frac{-1}{2\sqrt{\mathbf{b}(1-\mathbf{b})} \left\{ \sqrt{\mathbf{b} + (1-\mathbf{b}) \operatorname{arctanh}(\sqrt{\mathbf{b}})} \right\}} \int_{-1}^1 dx \frac{1-3x^2}{\left(\frac{(1-x^2)}{(1-\mathbf{b})} + x^2\right)^3}$$

where κ has now been explicitly defined. The expression for κ appears quite complicated as a function of β , but in fact it barely shows any dependence, even over a large range. In the limit of $\mathbf{b} \rightarrow 0$ the value for κ is computed to be 2.50, and for larger values of β , the values of κ are shown in Figure 5.2.

When calculating the equilibration rate, there is one additional factor of two that must be accounted for. The Bose statistics of the atoms in the cloud play other roles beyond what has been detailed so far. Many of those are only important at lower temperatures than were studied in the equilibration measurements discussed in this chapter. For instance, at lower temperatures $f(\mathbf{x}, \mathbf{v})$ would be noticeably different from a classical distribution and the enhancement due to Bose statistics of the scattering into a state that is already occupied would play a role in the equilibration rate. However, there is one effect that is relevant at the temperature at which our studies were conducted. In a thermal sample, indistinguishable bosons are twice as likely to be found near another boson as compared to distinguishable particles. In other words, the correlation function for indistinguishable bosons is different from distinguishable particles, and it is twice as likely that two indistinguishable bosons will be occupying the same region in space. The range of this enhanced




Figure 5.2. Change in the Equilibration Coefficient κ as a Function of the Parameter β . The solid curve was calculated using equation (5.15). It shows that elastic collisions equilibrate the sample more efficiently when it is further out of equilibrium (higher β). The dashed curve uses equation (5.20) in addition to equation (5.15) to account for the effect of the unitarity limit on a sample whose initial temperature is 110 nK and whose scattering length $a = -342a_0$. These conditions correspond to the more precise 251 G point in subsection 5.5.3.

probability is roughly the deBroglie wavelength of the particles. As long as the s-wave scattering length (i.e. the range of the scattering potential) is less than the deBroglie wavelength, the two-body elastic collision rate is expected to be enhanced by a factor of two due to this effect. For ^{85}Rb at 300 nK, the deBroglie wavelength is $\sim 6000 a_0$ (where a_0 is the atomic unit of length, the bohr radius), much larger than the scattering lengths that were typically measured. The reason that this factor of two is not included in the calculation of the scattering length from the atom-atom potential is that all of those calculations are done with the colliding atom pair in a coherent rather than a thermal colliding state. Since BECs are coherent state, this factor of two disappears for collisions between condensate atoms as has been verified experimentally [100].

With this final factor of two, the calculated equilibration rate agrees with the Monte-Carlo calculation in Ref. [101]. So far, the trapped clouds have been treated assuming perfect axial symmetry. When the clouds are actually brought out of thermal equilibrium in the trap, however, there is often a difference in energy between the two radial directions, up to as much as 30%. Explicit calculations show that using the average radial energy in the equilibration model does not significantly affect the value of κ that should be used at more than the $\sim 1\%$ level, even for the largest initial radial energy differences.

5.2.3 Using the Equilibration Rate of the Aspect Ratio of Trapped Clouds to Measure the Elastic Collision Rate. The quantities $\langle n \rangle$, $\langle v \rangle$, T_0 , and T_z can all be computed from the spatial widths (and number of atoms for $\langle n \rangle$) of a magnetically-trapped cloud, and so equation (5.15) can be used to compute the elastic scattering cross section. One way to do this is to bring the cloud out of thermal equilibrium and monitor the temperature of the cloud in one direction, as suggested by equation (5.15). A better way to measure the equilibration rate is to measure the relaxation rate of the aspect ratio of the cloud, however. There are a couple of reasons why monitoring the aspect ratio is better. First, we found that the individual widths of the cloud varied more than the aspect ratio, both from shot to shot and over longer periods of time. Second, the total energy of the cloud is not fixed and the sample can heat up over time, and the aspect ratio is less sensitive to this heating rate.

In order to relate the relaxation rate of the aspect ratio to the elastic scattering cross section, it is necessary to use equation (5.15) to determine the time evolution of the cloud from a certain set of initial conditions. Since κ is approximately constant, it can be treated as such for now in equation (5.15). If a linear heating rate α is included, the evolution of the radial and axial temperatures of a trapped gas are

$$\begin{aligned}\dot{T}_z &= -\frac{2}{3}\Gamma(T_z - T_r) + \mathbf{a} \\ \dot{T}_r &= -\frac{1}{3}\Gamma(T_r - T_z) + \mathbf{a} . \\ \Gamma &= \frac{2 \langle n \rangle \mathbf{s}_0 \langle v \rangle}{\hat{e}}\end{aligned}\quad (5.16)$$

If Γ is treated as a constant, the relaxation rate of the aspect ratio of the cloud can be computed from (5.16) and (4.13)

$$\frac{\mathbf{s}_z}{\mathbf{s}_r} = \frac{\mathbf{w}_r}{\mathbf{w}_z} \sqrt{\frac{1 + 2r \exp(-\Gamma t) + \mathbf{a} t}{1 - r \exp(-\Gamma t) + \mathbf{a} t}} \quad (5.17)$$

where $\sigma_{z,r}$ are the rms widths of the cloud in the axial and radial direction and the constant r is related to the initial aspect ratio AR_0 as

$$r = \frac{\mathbf{w}_z^2 AR_0^2 - \mathbf{w}_r^2}{\mathbf{w}_z^2 AR_0^2 + 2\mathbf{w}_r^2} . \quad (5.18)$$

In reality, Γ is not constant as a function of time. There are a few reasons for this. κ is not strictly constant over the evolution of the cloud. The density of the cloud can also change. Atoms are lost from the sample due to various loss processes, and that reduces the density. Also, since the widths of the cloud are changing in time, the shape of the cloud and therefore the density changes as well even in the absence of any loss of atoms. The average relative speed of the atoms also changes slightly as the velocity distributions in the axial and radial directions change. When the aspect ratio relaxation data are fit, all these considerations are included and the set of equations (5.16) are numerically integrated with the time dependences of Γ included.

5.2.4 Unitarity Limit and Thermal Equilibration. One other effect that must be included in the calculation of the elastic scattering cross section is the unitarity limit. The elastic cross section

depends on the energy of the colliding pair of atoms. In particular, the cross section cannot be greater than $8\pi/k^2$, where $\hbar k=p$. In any trapped cloud, the atoms collide with a range of energies and so the cross section for some of the collisions will be limited by the unitarity limit. The size of this effect can be estimated by using a cross section that scales with energy as

$$\mathbf{s} = \frac{\mathbf{s}_0}{(1 + (ka)^2)} \quad (5.19)$$

where a is the s-wave scattering length and $k=(\mu v/\hbar)^{1/2}$ with μ being the reduced mass of the colliding pair of atoms and v being the relative velocity. Not only does this cross section have the right form at high and low collision energies, it does a decent job of interpolating in between. This dependence can be included in Enskog's equation, and a new value for κ that depends on the average temperature of the trapped gas can be computed. The result of such a calculation is shown in Figure 5.2, and it can be seen that even at temperatures of ~ 100 nK, the unitarity limit still results in a correction of several percent to the measured relaxation time.

In addition to the energy dependence due to the unitarity limit, the positions B_{peak} and B_{zero} are also temperature-dependent quantities. For the cold temperatures used in this work, these shifts are predicted not to amount to more than ~ 0.2 G, and only that for the hottest temperatures studied. At much higher temperatures, such as 100 μK , B_{peak} and B_{zero} can shift by several gauss [102].

5.3 Creation of Nonequilibrium Trapped Clouds.

5.3.1 Initial Trap Loading and Cooling. The first step in the creation of the cold trapped atom cloud samples was to load the magnetic trap. This was done using the Double MOT system to load the Ioffe-Pritchard type baseball magnetic trap described in Chapter IV. As is usually the case in this work, the atoms were loaded in the $F=2$ $m_f = -2$ hyperfine ground state. Measurements of the elastic scattering near the Feshbach resonance were performed twice, once in 1998 and once in 2000. In our early work, not as many atoms were initially loaded into the magnetic trap as in our later work. In the early work about 70 million atoms were initially loaded into the trap instead of the 350 million that was typical of our later work. The main reason for the difference is simply the fact that the Double MOT system was improved over time. After loading the magnetic trap, evaporative cooling was used to cool the sample. Again, we were able to achieve much colder temperatures in

our later work (see Chapter VII for the evaporation details). In fact, in our later work we intentionally deoptimized the evaporative cooling by speeding up the initial stages to reduce the final density of the sample.

5.3.2 Bringing the Cloud out of Thermal Equilibrium. The rf-“knife” used for evaporative cooling (see subsection 4.9.1) was also used to create clouds that were out of thermal equilibrium. Since the atoms sag in the trap due to gravity (subsection 4.5.4), the lines of constant magnetic field are nearly horizontal planes across the cloud. This means that atoms are removed with the rf-knife only in the vertical direction. By bringing the rf-knife in faster than the clouds can equilibrate through elastic collisions, it is possible to remove energy selectively from the vertical direction of the cloud.

It was easiest to cut with the rf-knife faster than the equilibration time by tuning the bias magnetic field of the trap so that the s -wave scattering length in the cloud was near zero. Once we found the approximate location of B_{zero} , this is in fact what we did. Before that, we used low-density (10^9 atoms/cm³) samples that therefore had lower collision rates and cut quickly with the rf-knife.

5.3.3 Anharmonic Mixing and the “Rolling” Systematic. Initially, we had thought that we would remove energy from the vertical direction alone. However, anharmonic mixing between the two radial directions meant that once an energy imbalance was created between the two directions, energy could flow between them even in the absence of elastic collisions. Atoms with higher-energy orbits are precessing in the trap. In thermal equilibrium, there is a balance between the number of atoms precessing into and out of each direction in the trap and there is no way to observe this motion. Once atoms have been removed from the vertical radial direction, it is possible to see other atoms rotating into and out of that direction as they precess. The cycloid-like shape of the variation of the radial width led us to refer to this as “rolling” motion.

Since the rolling motion leads to a periodic change in the radial width of the cloud, it leads to a distortion in the measurement of the aspect ratio. To minimize the rolling motion, the rf-knife was held at a fixed frequency for a time longer than the period of the rolling motion. As the orbits of the energetic atoms precessed, they were then removed from the cloud and energy was removed not only from the vertical radial direction but the other radial direction as well.

This scheme for removing the rolling motion is not perfect, however. The rf-knife has a width in magnetic field due to power broadening. This means that it is always cutting into the cloud even if it is held at a fixed frequency. This continual cutting means that an energy imbalance is continuously being created and so it is not possible to wait and remove all of the rolling motion. Turning down the power of the rf-knife to reduce its broadening helps, but only at the cost of having to wait longer and longer to remove atoms from the cloud. If the cloud is not at $a=0$, eventually the cloud will equilibrate through elastic collisions faster than the energy imbalance can be created and the cloud cannot be brought out of thermal equilibrium. Even if the cloud is at $a=0$, the increased wait time results in a loss of atoms from the cloud due to inelastic collisions and collisions with unpumped background atoms and also causes the rate at which data can be accumulated to drop considerably. Usually, a 30 second rf-knife hold time was adopted at relatively low rf-power to reduce the rolling, and this was sufficient to remove most, but not all, of the rolling motion. Tailoring the rf-knife cutting ramp and the power of the rf-knife during the hold time would probably result in a better reduction of the rolling motion, but that additional complexity was not added at the time we performed our studies.

Aside from the periodic change in radial width that distorts the aspect ratio measurement, the rolling motion also complicates the measurement of the density of the cloud. In the absence of any anharmonic mixing, the radial direction that cannot be seen in the absorption imaging would have the temperature that the cloud had before the rf-knife cut in to produce the energy imbalance. If the rolling motion could be removed perfectly, then the initial energy in both radial directions would be the same. In our actual experiments we were between these two cases, but much closer to the latter situation. Energy conservation was used to infer the energy in the radial direction that was not visible. It is possible to measure the energy of the axial and visible radial direction from their respective spatial widths while the cloud was relaxing to equilibrium. Once the cloud was in thermal equilibrium, the total energy could be inferred from the two visible widths, and by using this value for the total energy and the observed widths as the cloud was relaxing to equilibrium the width of the third direction could be calculated. The presence of any heating complicated this procedure, but the

heating rate could be measured by measuring the increase in energy of the cloud once it was in thermal equilibrium.

5.4 1998 Search for the Feshbach Resonance

5.4.1 Relative Measurement of Elastic Collision Cross Sections. At the time that we began our work in ^{85}Rb , no Feshbach resonances had been observed in cold alkali atom collisions [103]. One of the goals of our work then was to simply locate the position of a Feshbach resonance as a function of magnetic field. In order to locate the fields B_{peak} and B_{zero} , only relative measurements were necessary, since only a maximum and minimum of the elastic scattering cross section needed to be located as a function of magnetic field. Performing a relative measurement also allowed us to avoid several difficulties with measuring the absolute scattering rates.

Measuring an absolute scattering requires the knowledge of the actual density of the sample, and therefore the measured number would have to be calibrated. That calibration is tricky enough (see section 4.12) that being able to avoid it was an advantage. More importantly, we did not have the same evaporation performance from our apparatus in 1998 as we did after improving it to create ^{85}Rb BEC. Our shot-to-shot reproducibility was not nearly as good in 1998, and more importantly we were limited in the densities and temperatures that we could achieve. This was important because the temperature of the cloud determined the relative importance of the unitarity limit and the spread in magnetic field across the trapped sample.

The range of magnetic field sampled by a trapped cloud depended on the temperature (T) of the cloud such that the FWHM spread of magnetic field in the cloud scales as $T^{1/2}$ and was 0.6 G at 0.5 μK for the trapping frequencies used in the 1998 measurement. Since the scattering length depends on magnetic field, the elastic collision cross section was then position-dependent in the trapped cloud. Also, the temperatures were hot enough that the unitarity limit was a significant consideration in the determination of the elastic scattering rate. This combination of position- and energy-dependence complicated the determination of the effective value of κ for relating the measured aspect ratio relaxation rate to the zero-temperature limit of the elastic scattering cross section, and we believed that a complex monte-carlo simulation was probably required to do so. Given the state of affairs at the time, we elected to avoid this additional complexity.

5.4.2 Measurement of the Relaxation Rate of the Aspect Ratio. After preparing cold magnetically-trapped atomic cloud samples in the manner described in Section 5.3 such that the radial energy was 0.6 that of the axial energy, we then ramped the bias field in the magnetic trap to a selected value and allowed the cloud to equilibrate for a fixed amount of time. The shape of the cloud was then determined using absorption imaging (section 4.11). By varying this equilibration time, the relaxation rate of the aspect ratio at the selected magnetic field was measured. This relaxation curve was fit to a simple exponential to determine an equilibration time τ_{eq} . For most of the data, only a few data points were measured for the τ_{eq} fit. The relaxation was not strictly exponential, but the small deviations from an exponential fit were far smaller than the relative change of τ_{eq} with magnetic field. The absorption imaging also allowed us to determine the density and the temperature of the sample. The magnetic field was calibrated in the manner described in subsection 4.10.1 to a precision of ± 0.1 G.

The measured equilibration time τ_{eq} depended both on B and T and varied from 0.15 s to nearly 2000 seconds. This fastest time is close to the point where we expect the equilibration rate to be limited by the trap frequencies ($\sim 2\pi \times 11$ Hz radial and $2\pi \times 5.5$ Hz axial). In order to better compare elastic scattering rates for samples at different temperatures and densities, we used a normalized equilibration rate $\sigma_{\text{N}} = (\langle n \rangle \langle v \rangle \tau_{\text{eq}})^{-1}$. This normalized rate has the units of a cross section, and would be equal to 0.8 times the true s-wave elastic scattering cross section σ_0 in the low-temperature limit. Figure 5.3 and Figure 5.4 display the measured value of σ_{N} as a function of magnetic field, and it is apparent that there is a clear maximum near 155 G and a clear minimum near 167 G. Even though we did not calculate the precise relationship between σ_{N} and σ_0 , the maximum and minimum of both quantities will occur at the same field in the limit of low temperature since the spread in magnetic field will be small and the unitarity limit will not be significant.

In addition to the measured values of σ_{N} , a theoretical calculation of the cross section is shown in Figure 5.3 and Figure 5.4 for two different temperatures. The calculation was performed by coworkers Jim Burke, Jr. and Chris Greene using the Rb-Rb potential derived in the manner




Figure 5.3. The points shown correspond to the measured normalized equilibration rates σ_N which are equal to the aspect ratio equilibration time divided by the density and average relative speed of the sample. The lines are calculations for the thermally averaged elastic cross section, not equilibration rates, and hence are not expected to fit the data points.

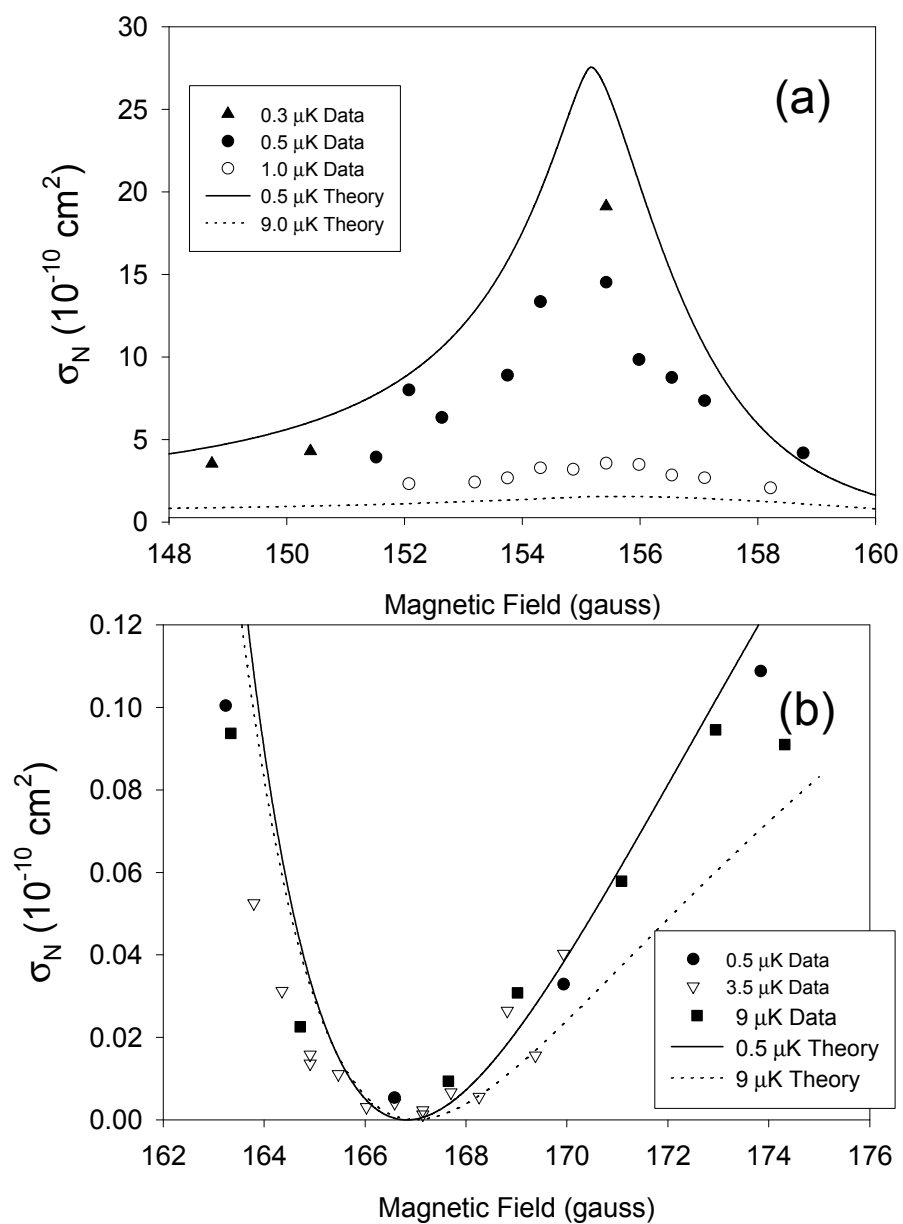


Figure 5.4. Figure 5.4(a) shows an expanded view of the region of Figure 5.3 near the peak of the Feshbach resonance, while Figure 5.4(b) does the same for the region near the minimum.

described below in subsection 5.4.3. The temperature dependence was calculated by performing the following average

$$\langle \mathbf{sv} \rangle = \frac{1}{N^2} \int d^3v d^3v' d^3x \mathbf{s}(E(\bar{v} - \bar{v}'), B) |\bar{v} - \bar{v}'| f(\bar{x}, \bar{v}) f(\bar{x}, \bar{v}') \quad (5.20)$$

$$\mathbf{s} = \frac{\langle \mathbf{sv} \rangle}{\langle v \rangle}$$

where $\langle v \rangle$ is the average relative speed, B is the bias magnetic field of the trapped cloud, $f(\mathbf{x}, \mathbf{v})$ is the Boltzmann distribution of a trapped cloud in thermal equilibrium at temperature T , and $\mathbf{s}(E(\bar{v} - \bar{v}'), B)$ is the theoretically calculated elastic scattering cross section as a function of magnetic field and the energy of the colliding pair. Because of the difficulty calculating κ , the theory curves are not expected to match the data, and are there only for rough comparison only.

Both the calculated cross section curves and measured equilibration data show the qualitative features. For temperatures below a few μK , there is a slightly asymmetric peak near 155 G, and the width and the height of this peak are strong functions of temperature. The reason for this height and width dependence is that for these large scattering lengths, the unitarity limit causes the effective elastic cross section to drop as the temperature of the sample increases. At 167 G there is a profound drop in the elastic scattering rate. This dip is quite asymmetric, but the shape is relatively insensitive to T , and at the bottom (B_{zero}) the rate is essentially zero.

5.4.3 Determination of B_{peak} and B_{zero} . Because the precise relationship between σ_0 and σ_N is not known, the desired quantities B_{peak} and B_{zero} could not be found by a detailed fit to all of the data. We determine them by fitting a simply smooth curve to only the few highest (or lowest) points at each temperature below 5 μK , and assigning error bars that span the values determined for all temperatures.

The 0.5, 0.75, and 1 μK data were used to determine B_{peak} . The location of the peak for each temperature can be determined to the same precision and the values agree to within 0.4 G. As was mentioned in the previous subsection 5.2.4, the location of the peak is predicted to depend on the energy of a colliding pair of atoms and hence on the temperature of the sample, but over the range of temperatures that we measured this shift is expected to be on the order of ~ 0.1 G, below the

assigned errors. The lack of any temperature dependence in the position of B_{peak} also justifies the assumption that σ_0 and σ_N reach their maximum values at the same magnetic field. The location of B_{zero} was determined using primarily the 3.5 μK data. For this temperature, a small shift (~ 0.2 G) in the location of B_{zero} due to the finite temperature was expected. This shift is due to the energy dependence of the Feshbach resonance and not the unitarity limit, which is negligible at the small values of a . We found $B_{\text{peak}} = 155.2(4)$ G and $B_{\text{zero}} = 166.8(3)$ G. We now believe that small errors were present in these determinations, and those will be discussed with respect to the 2000 measurement below. These small errors do not significantly effect the conclusions presented in the rest of this section.

The measurement of the location of B_{zero} was also complicated by the finite lifetime (~ 100 seconds) of the trapped clouds. It was possible to achieve equilibration times that were significantly longer than the lifetime of the atoms in the trap. This resulted in two things. First, a correction to the density had to be made as the clouds lost a significant fraction of their atoms during the time the equilibration was measured. Second, the signal-to-noise ratio dropped dramatically as we had to measure 1000-second equilibration times with only ~ 100 seconds available to study the relaxation. This situation can now be improved through using higher cloud densities and longer magnetic trap lifetimes than what were achieved in 1998. Eventually inelastic collisions proportional to the density will limit the sample lifetime as well (see Chapter VI).

Often, rather than using the variables B_{peak} and B_{zero} , the Feshbach resonance is described by its position and Δ , the width of the resonance. Again, Δ is simply defined as $B_{\text{zero}} - B_{\text{peak}}$ and from our measurement Δ is then found to be 11.6 (5) G. The location of the resonance is the same as B_{peak} .

5.4.4 Constraints on the Rb-Rb Interatomic Potential. The measured location of B_{peak} and B_{zero} can be used to constrain the parameters necessary to describe the Rb-Rb interatomic potential. The theoretical calculations described in this subsection were performed by collaborators Jim Burke, Jr. and Chris Greene. As described in subsection 2.9.4, a free parameter is added to the inner walls of both the triplet and singlet molecular potentials. These two free parameters are adjusted with the values of the dispersion coefficients C_6 , C_8 , C_{10} , etc. chosen beforehand. Using a

quantum defect theory (MQDT) [31] it is possible to calculate the location of B_{peak} and B_{zero} as a function of these free parameters and once the free parameters are set the quantities a_{T} and a_{S} that correspond to a particular location of the peak and zero can then also be calculated.

In ^{85}Rb , it is inconvenient to tabulate results as a function of a_{S} and a_{T} because both are close to lying on top of a divergent pole. Instead, the results of the measurement of B_{peak} and B_{zero} can be presented as constraints on a bound state phase (fractional bound state number) ν_{D} as in Ref. [91]. The singlet ν_{DS} or triplet ν_{DT} bound state phase corresponds to the triplet or singlet scattering length through the relation [104,105]

$$\nu_{\text{D}} = \frac{\cot^{-1}(a/a_{\text{ref}} - 1)}{\mathbf{p}}, \quad a_{\text{ref}} = \frac{\Gamma(5/4)}{\Gamma(3/4)\sqrt{8}} (2mC_6 / \hbar^2)^{1/4} \quad (5.21)$$

where m is the reduced mass of the colliding atom pair.

For a given value of the dispersion coefficients, and especially C_6 , adjusting the inner wall free parameters to match the measured values of B_{peak} and B_{zero} within the experimental uncertainty severely constrains the quantities ν_{DS} and ν_{DT} . However, the position of the peak and zero are not sufficient to constrain all three parameters C_6 , ν_{DS} and ν_{DT} [106]. In order to constrain the value of C_6 , an additional constraint was added to the theoretical potentials: they must predict the g-wave shape resonance in an energy range consistent with the measured value given in Ref. [107].

The results of this analysis are presented in Figure 5.5, along with a comparison to work that was performed contemporaneously [91]. This illustrates that accurate measurements of Feshbach resonances are an extremely precise method for determining properties of interatomic potentials. Our measurement allowed us to reduce the combined C_6 , ν_{DS} , and ν_{DT} parameter space by a factor of 80 as compared to existing measurements. As discussed in Ref. [91], the peak of the resonance depends mostly on C_6 and the sum $\nu_{\text{DS}} + \nu_{\text{DT}}$. The width of the resonance depends mostly on $\nu_{\text{DS}} - \nu_{\text{DT}}$. As Figure 5.5 shows, the allowed ν_{DS} , ν_{DT} parameter region is an extremely correlated function of C_6 . In particular, we find that the area of the allowed parameter region is independent of both C_6 and C_8 . The uncertainty in these dispersion coefficients simply shifts the position of the ‘‘diamond.’’ Converting the bound state phases back into scattering lengths, agreement was found with prior work [108,109, 110]. From the potentials determined in this work, the singlet and triplet




Figure 5.5 Constraints on the Rb-Rb Interatomic Potential. This figure compares the allowed v_{DT} , v_{DS} parameters space based on our work and measurements contemporaneous to our 1998 result. The large rectangle is the allowed range from Ref. [109]. The large diamond is the range from Ref. [91]. The small filled diamond is the range from our 1998 measurement. The large rectangle and large diamond were established using a value of $C_6 = 4550$ a.u. As is discussed in the main text, the position of the constrained region depends on C_6 , which is indicated by the small diamonds on either side of the filled diamond. These diamonds show the effect of 50 a.u. uncertainty in C_6 around the central value of 4700 (to which the filled diamond corresponds). The arrows indicate the sensitivity to changes in the resonance width and position in this parameter space. The small black square in the center of the small rectangle represents the constraints set by our more recent measurement and reflects the increase in precision of the measurement of B_{zero} . The small rectangle is another way of representing the new constraint supplied by our more recent measurement. Unlike all of the other objects in this figure, the small rectangle does not correspond to a fixed value of C_6 . Rather, it is the *total* allowed range of v_{DT} , v_{DS} consistent with our measurement of B_{peak} and B_{zero} and the now smaller uncertainty in $C_6 = 4660 \pm 20$ implied by our measurement of a_{bg} . In other words, the small rectangle is measuring the same constraint as the area between the two small open diamonds.

scattering lengths can be specified (in atomic units): $a_T(^{85}\text{Rb}) = -369 \pm 16$, $a_S(^{85}\text{Rb}) = 2400^{+600}_{-350}$, $a_T(^{87}\text{Rb}) = 106 \pm 4$, $a_S(^{87}\text{Rb}) = 90 \pm 1$. The value for C_6 was determined to be 4700(50) atomic units. Several other scattering observables predicted by the 1998 potentials agreed with previous measurements. Specifically, the 1998 potentials predicted a broad d-wave shape resonance [111] in ^{87}Rb , the scattering length ratio $a_{2,1}/a_{1,-1}$ [112], and the thermally averaged $|2,2\rangle + |1,-1\rangle$ inelastic rate constant [113]. Also, 10 of the 12 measured d-wave bound states were found within the 2σ error bars [109]. The 1998 potentials also predicted the locations and widths of additional Feshbach resonances besides the one observed: a narrow one with width ~ 0.01 G at 226 ± 4 G and a wider one

The difference between our measured peak and width compared to the results in Ref. [91], which were $B_{\text{peak}} = 164(7)$ G and $\Delta = 8.4(3.7)$ G, is most likely explained by the fact that the experiments in Ref. [91] occurred at much higher temperatures (~ 100 μK) and the shift in the position of the Feshbach resonance with the energy of a pair of colliding atoms was not taken into account [114].

5.5 2000 Measurement of the Absolute Value of the Scattering Length.

5.5.1 Inadequacies of the 1998 Experiment. While the 1998 measurement of the Feshbach resonance resulted in a dramatic improvement in the knowledge of the parameters that characterize the Rb-Rb interaction potential, the precision of the determination of the scattering length as a function of magnetic field was not good enough for some of our later experiments, particularly on the stability of BECs at negative scattering lengths. Also, it was suggested to us that the value of C_6 used in our previous work was incorrect [115]. The difference in the background scattering length, a_{bg} of equation (5.1) calculated with our original determination of C_6 and the value of C_6 from this new analysis was a whopping 50%. Even though this analysis was later found to not be correct, we decided that it was necessary to pin down the absolute size of the scattering length experimentally so that we would not have to rely so heavily on theoretical determinations of a as a function of B .

Also, in 2000 as opposed to 1998 most of the reasons that we avoided the measurement of the absolute size of the scattering length were no longer present. The imaging system had been

calibrated (section 4.12) to determine the number of atoms in the trapped cloud as part of our studies of BEC stability at negative scattering length. The improvements in evaporative cooling that we had made allowed us access to a much greater range of temperatures and densities, and so colder samples could be used. This reduced the uncertainties in the relation of the $T=0$ value of σ_0 to the observed aspect ratio relaxation rate, since the unitarity limit was not as significant for colder samples and the spread in magnetic field of the trapped samples was lower as well. The colder samples also improved our ability to calibrate the magnetic field and the improved evaporation led to much better shot-to-shot reproducibility.

Finally, as part of our experiment on the stability of BECs with attractive interactions, we were able to make a better determination of B_{zero} . This measurement did not agree with our 1998 results within the stated errors, and led us to reexamine our original conclusions. In the course of doing so, we found the presence of a few small errors in our published analysis. In order to both rectify these errors and improve the determination of the scattering length as a function of field, we revisited the measurement of the elastic scattering near the Feshbach resonance.

5.5.2 Determination of B_{zero} from the BEC Stability Measurement. The more precise measurement of the position of B_{zero} was a natural outgrowth of our work determining the stability of magnetically trapped BECs with attractive atom-atom interactions (Chapter IX). When large enough interactions (defined by Na) are present, the BECs will become unstable and collapse. This collapse criterion is given by the relation $N_{\text{cr}}|a|=\beta$, where N_{cr} is the number of atoms in the condensate and $\beta=1.41 \mu\text{m}$ for our magnetic trap. N_{cr} was determined as a function of magnetic field down to very small a . A simple extrapolation to the point where $N_{\text{cr}} = \square$, such as in Figure 9.5, determined the position of B_{zero} . The value of B_{zero} determined in this way is 165.85(5) G, a factor of six more precise than before.

This increase in precision is due to a few reasons. First, the stability measurement is proportional to a rather than a^2 , making it naturally more sensitive to the position of a minimum. Second, the BEC stability measurement did not suffer from the signal-to-noise difficulties associated with long equilibration times. Also, the colder samples allowed more precise magnetic field

calibrations. An additional advantage of using BECs is that shifts due to finite temperature are not a consideration.

The disagreement between the BEC stability determination of B_{zero} and the 1998 measurement has been traced to a couple of small errors. First, we originally used the wrong sign for the nuclear magnetic moment in the Breit-Rabi equation, equation (4.1), and that resulted in an error of 0.3 G in the magnetic field calibration. Second, the small shifts ($\sim 0.2\text{G}$) due to the energy dependence of the Feshbach resonance and finite spread of the magnetic field across the cloud were not included. Finally, the temperature-dependent asymmetric shape of the cross section versus magnetic field was not properly accounted for; the minimum was found assuming a symmetric function, which also resulted in a systematic shift. The error in the magnetic field calibration was also present in the determination of B_{peak} , and so that has to be adjusted as well. Because colder clouds were used to measure the peak location, several temperatures were tested, and the asymmetry was not as pronounced, the other considerations are not expected to apply to the peak.

5.5.3 Measurement of the Absolute Size of the Elastic Scattering Cross Section. Once again, the samples were prepared in the same way as detailed in Section 5.3, and again the rf-knife was used while the sample was near $a=0$ to reduce the radial energy to 0.6 that of the axial. Once this was done, the field was ramped to the desired value of bias magnetic field and the cloud was allowed to equilibrate for a fixed amount of time. The size, number, and shape of the cloud were measured using resonant absorption imaging. This procedure was then repeated for different equilibration times and the rate of the relaxation of the aspect ratio of the clouds to equilibrium was recorded.

The measured relaxation rate curves were then fit with a model based on equations (5.15) and (5.16), and the relationship between the temperature of a magnetically trapped cloud and the width of the cloud, equation (4.13). The density as a function of time was determined from the number and widths measured by the absorption images and the implied width determined by energy conservation as described above in subsection 5.3.3. Examining the heating rate of clouds that had reached thermal equilibrium and performing a simple linear fit determined the total energy as a function of time. The density, energy (and therefore $\langle v \rangle$), and κ all varied as a function of time

during the relaxation of the aspect ratio (as much as $\sim 10\%$), and these variations could be measured or calculated and were included in the fit to the data. From the fits, it was possible to extract the elastic scattering cross section σ_0 .

The functional form of the Feshbach resonance calculated from theory (approximated by equation (5.1)) is expected to be very accurate, and so the absolute value of the scattering length does not need to be determined everywhere but only at a single (or few) selected fields. Data were taken at $B = 168.0, 169.7,$ and 251.0 G. These fields were selected since they were expected to show different equilibration rates, but were still far enough from the Feshbach resonance so that the ~ 0.2 G FWHM spread of magnetic field across the sample should not distort the extraction of σ_0 from the fits. The difference between a averaged over all the collisions in the cloud and the square root of a^2 averaged over all collisions was less than 0.03% for the most sensitive field. The distance from the Feshbach resonance also ensures that the precise location of B_{peak} was not a limiting factor in analyzing the results. Across the range of magnetic fields measured, the equilibration times varied by a factor of about 30. The density of the samples was maintained near $2.5 \times 10^{10} \text{ cm}^{-3}$ and the temperature was kept near 130 nK. Figure 5.6 shows a set of aspect ratio relaxation data and the corresponding fit.

The uncertainty in the determination of a is dominated by the determination of the number and the fit to obtain σ_0 . The number determination relies on the calibration of our absorption system, which is estimated to be 10% . Since the relaxation rate scales as a^2 , this results in a 5% uncertainty in each determination of a . The uncertainty in the fit to obtain σ_0 is primarily caused by residual small oscillations in the radial direction due to the imperfect removal of the rolling motion. Although this limited the accuracy of the fit, the size of this imbalance was varied and no significant shifts were observed in the results were observed. Also, the period of the oscillations does not depend on the value of a and so the agreement between data sets of very different equilibration times also indicates that this effect is not distorting the results significantly. The measured value (in atomic units) of a is $-342(10)$, $-97(6)$, and $-60(4)$ at $B = 251.0, 169.7,$ and 168.0 G respectively. The listed uncertainty does not include the common uncertainty due to the atom number determination.




Figure 5.6. Relaxation of the Aspect Ratio to Equilibrium. The data shown were taken at $B = 169.7$ G. Each data point represents a single destructive absorption measurement of the cloud after it was released from the magnetic trap and allowed to expand for 1.6 ms. The solid line shows the fit to the data used to determine the elastic scattering cross section and hence the scattering length. The deviation of the data points from the fit around 3-4 seconds is due to the imperfect removal of the rolling motion. The lower plot shows the measurement of the individual axial and radial widths that went into determining the aspect ratio.

5.5.4 New Constraints on the Rb-Rb Interatomic Potential. Using the approximate value of the scattering length as a function of field given in equation (5.1) and the measured values of B_{peak} and B_{zero} , it is possible to determine a value of a_{bg} for each measured value of a (shown in Figure 5.7). The weighted average of all of these data gives $a_{\text{bg}} = -380(21) a_0$ where the uncertainty is dominated by the determination of the number of atoms in the cloud.

Now that a_{bg} has been measured along with B_{peak} and B_{zero} , it is possible to constrain the three parameters needed to describe the Rb-Rb interatomic potential for ultracold collisions. Using the same quantum defect method as in the 1998 measurement, the Rb-Rb potentials are adjusted until the theoretically predicted B_{peak} , B_{zero} , and a_{bg} match the observed values. Matching the Rb-Rb potential to the observed scattering lengths implies $a_{\text{T}}(^{85}\text{Rb}) = -332 \pm 18$, $a_{\text{S}}(^{85}\text{Rb}) = 3650^{+1500}_{-670}$, and $C_6 = 4660 \pm 20$, where all of the values are given in atomic units. Again, Ref. [106] is used for the value of C_8 and C_{10} . Allowing C_8 to vary by 10% would increase the uncertainties of a_{T} , a_{S} , and C_6 by about 25%. This variation in C_8 not only allows for an error in its calculation, but it also covers the uncertainty in the parameters that serve as the initial inputs to the calculation of the Rb-Rb potential [27]. The new constraints in the v_{DS} , v_{DT} plane are shown in Figure 5.5. This more precise determination of the C_6 coefficient agrees with both an *ab initio* calculation [116] and a recent photoassociation measurement [117].

With this new measurement we have now characterized the s-wave scattering length as a function of magnetic field near the Feshbach resonance to a precision around 5%. This characterization will not only be useful in future experiments in ^{85}Rb Bose-Einstein Condensation, but also represents the single most precise measurement of the Rb-Rb interatomic potential to date. To summarize the results of our studies, the following table is presented.

$$a = a_{bg} \left\{ 1 - \left(\frac{B_{zero} - B_{peak}}{B - B_{peak}} \right) \right\} \quad (5.1)$$

Quantity	Measured Value	Experiment (Thesis Section)
\mathbf{B}_{peak}	154.9(4) G	1998 Elastic Scattering with Calibration Correction (5.4)
\mathbf{B}_{zero}	165.85(5) G	BEC Stability with $a < 0$ (9.7)
\mathbf{a}_{bg}	-380(21) a_0	2000 Elastic Scattering (5.5)

Table 5.1. Summary of Measured Feshbach Resonance Parameters.




Figure 5.7. The value of a_{bg} implied by several aspect ratio relaxation measurements. This value was determined by dividing the s -wave scattering length determined at a particular magnetic field by the factor $\{1-(B_{zero}-B_{peak})/(B-B_{peak})\}$. The s -wave scattering lengths measured at 251.0 G were divided by an additional factor of 0.99 to account for the deviation of the approximate description of a vs. B from the full theory calculation. The error bars shown include the statistical errors plus the fit uncertainty due to the rolling systematic. The data represented by the open circles had a larger amount of rolling motion present than the data represented by the filled circles. There is no observable difference between the two sets of data. For clarity, the open circles are displaced 0.2G to the right of their true magnetic field value. The fit uncertainty is non-statistical and so the central values are clustered together more than the error bars would suggest.

Chapter VI

Magnetic Field Dependence of Ultracold Inelastic Collisions Near a Feshbach Resonance

6.1 Introduction

6.1.1 Motivation. The same reasons that led us to study the magnetic field dependence of elastic scattering rates also led us to study the magnetic field dependence of inelastic rates. First, since Feshbach resonances themselves had only been discovered relatively recently, the inelastic collision rates near them had not been yet been fully characterized, and so measuring the inelastic collision rates would itself contribute to the general understanding of the physics of Feshbach resonances in cold alkali atom collisions. Several predictions had been made before the Feshbach resonances were observed experimentally that they should not only enhance the elastic scattering rates, but would enhance both two-body dipole relaxation [118,119] and three-body recombination [120,121,122,123,124] inelastic collision rates as well, and we hoped to provide a test of these theoretical predictions. This enhancement of inelastic rates was observed in the work in ^{133}Cs [93], although the precise character (two- or three-body) of the inelastic rates was not definitively determined and there were no instances reported in that work where both the elastic and inelastic rates had been characterized for a single Feshbach resonance.

In the same way that knowledge of the s-wave scattering length is necessary to interpret experiments with BECs near a Feshbach resonance, it is also necessary to know the inelastic collision rates. Only if the ordinary dipole relaxation and molecular recombination rates are known is it possible to determine if any observed loss from the BEC is due to ordinary inelastic collisions or some other process that is characteristic of the BEC itself. In particular, the extraordinarily high loss rates observed near Feshbach resonances in ^{23}Na BECs [90] spurred several proposals for exotic coherent loss processes [94-97], even though the effect of the Feshbach resonance on the ordinary inelastic rates had not been determined and the calculations of the enhancement of the ordinary loss rates [94-97] had not been tested. This situation left many outstanding questions as to the nature of

inelastic losses near a Feshbach resonance. How large are the dipole relaxation and three-body recombination near the Feshbach resonance and how accurate are the calculations of those quantities? How much of the observed condensate losses in Ref. [90] are due to these more traditional mechanisms and how much arise from processes unique to condensates? How severe are the “severe limitations” [90] that inelastic loss put on the use of Feshbach resonances to change the s-wave scattering length in a BEC? Would these same processes be observed in ^{85}Rb as in ^{23}Na ? Part of the answer to these questions was a study of the inelastic collision rates near the ^{85}Rb Feshbach resonance.

Finally, the efficiency of evaporative cooling is limited by the ratio of elastic to inelastic collision rates. If this ratio is not high enough, then it will not be possible to create a BEC. Measuring the inelastic collision rates as a function of field was therefore a necessary step in the optimization of evaporative cooling of ^{85}Rb .

6.1.2 Types of Inelastic Collisions in a Magnetic Trap. The inelastic collisions in a magnetic trap are easily classified by the number of atoms involved in the collisions. Spin-exchange and dipolar relaxation are two-body processes, while molecular recombination is a three-body process. Each of these inelastic collisions is discussed briefly in the following subsections. For all of these processes, it turns out that the collision cross section is inversely proportional to the velocity of the atoms involved in the collision. This means that the collision rate is independent of the velocity of the atoms and only proportional to an inelastic rate coefficient and the density of the atoms in the sample squared (two-body) or cubed (three-body).

6.1.2.1 Spin-exchange Collisions. Due to the difference in the singlet and triplet scattering lengths of the interatomic potential, two atoms that come together in one set of magnetic sublevels can collide and then leave in a different set of magnetic sublevels (see subsection 2.7.1). If one or both of the post-collision magnetic sublevels is an untrapped state, these collisions will lead to loss from the magnetic trap. For arbitrary spin states, these loss rates are quite large ($\sim 10^{11}$ cm³/s) and would be a significant source of loss from the trap. This is not the case in our spin-polarized samples of ^{85}Rb in the $F=2$ $m_f = -2$ state. The spin angular momentum along the z-axis must be conserved in the collision, and since there is no $m_f = -3$ state in the lower hyperfine level and the

collision energies are much smaller than the ground state hyperfine splitting, the spin-exchange rate is suppressed because there are no final states into which the atoms can scatter while still maintaining angular momentum conservation. From this point on, the spin-exchange rate will therefore be ignored.

6.1.2.2 Magnetic Dipole Relaxation. While angular momentum conservation can prevent a change in spin state for atoms that both enter and exit the collision on an s-wave channel, it is still possible for the spins of the atoms to relax and that angular momentum be transferred to their orbital motion. Such an inelastic collision is referred to as dipole relaxation, since one source of it is the magnetic dipole interaction between the two colliding atoms. The magnetic dipole interaction is described in subsection 2.9.5 in a little detail, but it is simply the energy due to the fact that there is a magnetic dipole associated with each atom due to the spin of its valence electron and its nucleus and the two dipoles of a colliding atom pair exert a force upon one another. This interaction is not spherically symmetric, and so it can couple an s-wave entry channel to a d-wave exit channel leading to a change of spin state of the initial colliding pair (when decomposed into spherical tensors, there is no $L=1$ component in the magnetic dipole interaction and so it is not possible to couple an s-wave entry channel to a p-wave exit channel).

In a dipole relaxation collision of ^{85}Rb atoms in the $F=2$ $m_f = -2$ state, the atoms would exit the collision either both in the $F=2$, $m_f = -1$ state, or one in the $F=2$ $m_f = -2$ state and the other in the $F=2$, $m_f = 0$ state. In the high magnetic field traps that we use, the spin relaxation results in a change of the Zeeman energy of the colliding atoms, and about 5 mK of energy is transferred to the kinetic energy of each atom in the colliding pair. They will therefore leave the magnetic trap since the trap is not 5 mK deep for the $F=2$, $m_f = -2$ state and is even shallower for the $m_f = -1, 0$ states.

The magnetic dipole interaction is not the only mechanism by which spin relaxation can occur. In heavy atoms, the second-order spin-orbit coupling also plays a role (see subsection 2.9.5). The electronic wavefunctions of the colliding atom pair can perturbatively couple to excited electronic states, breaking the degeneracy of the electronic spin projection in the triplet part of the atom-atom interaction potential. This removal of spherical symmetry of the electronic part of the atom-atom interaction potential produces a term that has exactly the same form as the magnetic

dipole interaction and so the two effects are indistinguishable in atom-atom collisions. The total dipole relaxation rate will be determined by the sum of the effects from the magnetic dipole interaction and the second-order spin-orbit coupling.

The size of the second-order spin-orbit coupling was calculated for alkali atom collisions in Ref. [119]. The predicted inelastic rates were too small to explain the depolarization of observed in room-temperature optical pumping experiments, however [125]. Including relativistic effects explicitly rather than applying perturbation theory to nonrelativistic wavefunctions in a new calculation improved the values of the spin-orbit coupling [126].

6.1.2.3 Three-body Molecular Recombination. The two ^{85}Rb atoms involved in a collision are not in the lowest energy state of the interatomic potential; it is energetically favorable for the two atoms to scatter into a lower energy bound state of the molecular potential. However, it is not possible for only two atoms to collide and form a molecule and simultaneously satisfy both energy and momentum conservation. This is easily demonstrated in the center-of-mass frame. The net momentum of the two incoming atoms is equal to zero in that frame. If the atoms formed a molecule, they would both have to move away from the center of the collision in the same direction and that would violate momentum conservation. If three atoms collide, then two can form a molecule while the third satisfies energy and momentum conservation, and therefore molecular formation is intrinsically a three-body process.

The calculation of molecular formation rates, called three-body recombination rates, is much more complicated than the two-body process. Not only is the calculation made more difficult by the simple fact that two relative coordinates have to be considered instead of one, there is in general a very large number of possible exit channels, over 1000 in the case of ^{85}Rb [122]. Progress is being made, but the precise calculation of three-body recombination rates is still in its early stages. Simple scaling arguments have been used, however, to relate the three-body recombination rate to the s-wave scattering length of the two-body interatomic potential. Since three bodies are involved, the recombination rate, in units of sec^{-1} , must be proportional to the density n as n^3 , and it is possible to write [120]

$$\mathbf{n}_{rec} = \mathbf{a}n^3, \mathbf{a} = C \frac{\hbar a^4}{m} \quad (6.1)$$

where v_{rec} is the three-body recombination event rate, a is the s-wave scattering length and m is the mass of an atom in the colliding pair. The coefficient C is then dimensionless.

In order to simplify the calculation of the three-body recombination rate, the full interatomic potential is not used. Instead, recombination only into the bound state closest to threshold is considered [120,123,124] or model potentials with only a few bound states are used [122]. This appears to be a reasonable assumption since ultracold recombination events are expected to couple mostly to the bound states closest to the collision threshold [121]. Most of the theoretical analyses are for positive scattering length, and the predictions for the value of C range from 3.9 in Ref. [120] to a prediction that C should oscillate as a function of the scattering length between zero and 67.9 [124]. In Ref. [122] both positive and negative scattering length potentials are considered, and the a^4 scaling is still predicted to hold in the negative scattering length case, with $C \sim 169$.

Three-body recombination, like dipole relaxation, will cause a loss of atoms from the magnetically trapped sample. The molecule that is formed may or may not have a magnetic moment that will allow it to be trapped. For the purposes of our detection system, however, whether the molecule is trapped or not makes no difference since it will be insensitive to both MOT recapture and absorption imaging (section 4.11). The third atom may or may not be lost from the trap. It may undergo a change in spin state as a result of the collision and would then likely be untrapped. Even if it does not undergo a change in spin state, it will pick up kinetic energy from the binding energy released in the molecular formation. The scale of this kinetic energy will be on the order of at least tens of μK , except near the Feshbach resonance peak where the binding energy of the last bound state can become quite small. If the third atom is not lost from the magnetic trap, then it could be counted in the MOT recapture signal, but since the absorption imaging measures the spatial extent of the original cold (~ 600 nK) cloud of atoms it is less likely that this third atom would be counted in the absorption image. Comparing the loss rates calculated by the imaging and the MOT recapture will give some indication of whether this third atom is remaining in the trap or not.

6.1.3 Loss Rate Equations for a Magnetically Trapped Sample. Since both dipole relaxation and three-body recombination result in the loss of atoms from the magnetic trap, measuring the rate and density dependence of the trap loss allows the determination of the inelastic scattering rates. The loss from the trap is expressed as three components: the one-body loss rate $1/\tau$, the two-body loss rate K_2 , and the three-body loss rate K_3 . The one-body loss rate is due to collisions of cold trapped atoms with room temperature unpumped atoms in the imperfect vacuum in the science cell. For a gas with a density n , the loss rate is then expressed as

$$\frac{dn}{dt} = -K_2 n^2 - K_3 n^3 - \frac{n}{\tau} \quad (6.2)$$

Therefore, a magnetically trapped gas has a total loss rate of

$$\frac{dN}{dt} = -K_2 \langle n \rangle N - K_3 \langle n^2 \rangle N - \frac{N}{\tau} \quad (6.3)$$

Equation (6.3) was simply obtained by summing equation (6.2) over all space while using equation (4.12) for the density of atoms in the sample. The quantity $\langle n \rangle$ is the average density of the atoms in the cloud given by equation (4.13), and $\langle n^2 \rangle$ is defined as

$$\langle n^2 \rangle = \frac{1}{N} \int n(\vec{x})^3 d^3x = \frac{8}{\sqrt{27}} \langle n \rangle^2 \quad (6.4)$$

Another way of writing equation (6.3) is

$$\frac{dN}{dt} = -K_2 \frac{N^2}{V} - \frac{8}{\sqrt{27}} K_3 \frac{N^3}{V^2} - \frac{N}{\tau} \quad (6.5)$$

$$V = 8p^{\frac{3}{2}} \sigma_r^2 \sigma_z$$

where $\sigma_{r,z}$ are the radial and axial rms widths of the trapped cloud. Since the cloud heats up during the measurement of the losses, the quantity V increases as a function of time and that evolution needs to be included to determine K_2 and K_3 accurately. Equation (6.5) makes this effect of the heating more explicit while in equation (6.3) it is contained in the time evolution of the density.

Some care needs to be taken in order to compare K_2 and K_3 to the theoretically predicted quantities. For identical particles with Bose statistics, the effect of the correlation function (subsection 5.2.2) in a thermal sample needs to be taken into account. This results in a factor of $2! = 2$

increase in the two-body rates as compared to a coherent sample and a factor of $3!=6$ increase in the three-body rates. However, in order to prevent an over-counting of the number of collisions in the sample, it is necessary to divide by a factor of 2 and 6 for two-body and three-body rates respectively when the collisions occur between identical particles, canceling the correlation function enhancement. Often, the theory quotes event rates, how many inelastic collisions occur in a unit of time, while the quantities K_2 and K_3 here specify loss rates, the number of atoms lost per unit time. Two atoms are lost in a two-body collision and three atoms are lost in a three-body collision and so again these factors must be taken into account.

6.1.4 Heating Rates due to Inelastic Collisions. Not only do the inelastic collisions cause loss from a magnetically trapped atom cloud, they also cause it to heat up. There are two reasons for this heating. First, the inelastic collisions are exothermic, producing atoms with much greater kinetic energy than they originally possessed. These atoms may collide on the way out of the sample and that may result in energy being deposited in the cloud. The other way that the atoms can cause the sample to heat up is through the density dependence of the losses. Since the cloud of atoms in the magnetic trap is most dense at its center, and the colder atoms in the sample spend more time in the central region of the cloud, the inelastic collisions will preferentially remove the colder atoms from the sample. It is straightforward to calculate the average energy of the atoms lost from the cloud by an appropriate integral:

$$\begin{aligned}\bar{E}_{2\text{-body}} &= \frac{1}{N^2} \int f(\vec{x}, \vec{v})^2 \left\{ \frac{m\mathbf{w}_r^2 r^2}{2} + \frac{m\mathbf{w}_z^2 z^2}{2} + \frac{mv^2}{2} \right\} d^3x d^3v = \frac{9}{4} k_B T \\ \bar{E}_{3\text{-body}} &= \frac{1}{N^3} \int f(\vec{x}, \vec{v})^3 \left\{ \frac{m\mathbf{w}_r^2 r^2}{2} + \frac{m\mathbf{w}_z^2 z^2}{2} + \frac{mv^2}{2} \right\} d^3x d^3v = 2k_B T\end{aligned}\tag{6.6}$$

where $f(\mathbf{x}, \mathbf{v})$ is the thermal equilibrium distribution function (equation (4.11)). Since the average energy of an atom in the trapped sample is $3k_B T$, equation (6.6) explicitly shows that the average energy of the atoms lost due to inelastic collisions is less than the average energy of the atoms in the cloud.

6.2 Measurement of the Inelastic Collision Rates.

6.2.1 Sample Preparation. The cold atom clouds that were used to study the inelastic collision rates were created by loading the Ioffe-Pritchard Baseball magnetic trap from the Double-MOT system as described in Chapter IV. Initially, 3×10^8 $F = 2$ $m_f = -2$ ^{85}Rb atoms at 45 μK were loaded into the magnetic trap. Forced radio-frequency evaporative cooling was used to increase the density of the cloud of atoms while decreasing its temperature. Because of the high ratio of inelastic/elastic collision rates in ^{85}Rb and the dependence of that ratio of the magnetic field and temperature, trap conditions must be carefully chosen to achieve efficient cooling (Chapter VII). We evaporated to the desired temperature and density (typically 3×10^{11} atoms/cm³ and 500–700 nK) at a final field of $B = 162$ G.

6.2.2 Measurement of the Density Evolution. After the desired density and temperature had been achieved, we adiabatically changed the bias magnetic field to a selected value and measured the density (i.e. number of atoms and size of the cloud of atoms) as a function of time. There were two ways in which the density of an atom cloud was measured. The first was using absorption imaging (section 4.11), which is destructive. We also used non-destructive polarization imaging, which is discussed as the end of this chapter (section 6.4). In both cases, the cloud was imaged on the camera CCD array to determine the spatial size and number of atoms in the cloud. The nondestructive method allowed us to observe the time evolution of the number and spatial size of a single sample. The destructive method required us to prepare many samples and observe them after different delay times, but had the advantages of better signal to noise for a single image and larger dynamic range. Figure 6.1 shows a set of nondestructive imaging data. In addition to these techniques, we also measured the number of atoms using MOT recapture.

We measured the time dependence of the number of the atoms and the spatial size of the trapped clouds at several different values of B . We assigned a 10% systematic error to our density determination, based primarily on the error in our measurement of number. The value of B was calibrated in the usual way (subsection 4.10.1). The magnetic field width of the clouds scaled as $T^{1/2}$ and was 0.39 G FWHM at 500 nK.




Figure 6.1 Number of atoms versus time taken by polarization-rotation imaging at 159 G. Fits to both purely two-body and purely three-body inelastic rates are shown, illustrating the difficulty in separating two- and three-body loss processes.

In each data set, we observed the evolution of the sample while a significant fraction (20% – 35%) of the atoms was lost. In addition to the number loss, the volume of the trap also increased with time due to heating, and that volume increase (typically 50%) was fit to a polynomial --- usually a straight line was sufficient within our precision. This heating rate varied with B and scaled roughly with the inelastic rates (see subsection 6.2.6). The number (N) of a function of time was then fit to equation (6.5) with the time-dependence of the volume included. Also included in equation (6.5) is the background loss rate τ , which was typically 450 seconds and independently determined by measuring the loss from low-density clouds for long (~ 400 seconds) times. The fractional change in the number of atoms measured by MOT recapture agreed on average with the change measured by absorption imaging or polarization-rotation imaging to within a few percent, even after some inelastic collisions had occurred. This suggests that a large majority of the inelastic collisions resulted in all of the atoms being completely lost from the magnetic trap, not just lost from the original cold atom cloud.

The change in the combined inelastic rates as a function of B around the Feshbach resonance is shown in Figure 6.2. All of the inelastic data shown in Figure 6.2 were taken with initial temperatures near 600 nK. The points less than 157 G were taken with initial densities within 10% of $1 \times 10^{11} \text{ cm}^{-3}$, while the points higher than 157 G were taken with initial densities that were $2.7 \times 10^{11} \text{ cm}^{-3}$. The decrease in initial density below 157 G was due to ramping through the high inelastic loss region after forming the sample at 162 G. Because of this density variation, we have used a different weighting parameter β in the sum of the two- and three-body rate constants for the two regions in Figure 6.2. These rates are generally quite large as compared to ^{23}Na and ^{87}Rb , although they are less than those of ^{133}Cs .

Also shown on Figure 6.2 is the elastic rate determined previously (see Chapter V). The shape of the inelastic rate vs. B roughly follows that of the elastic rate vs. B . In particular, the peak of the inelastic rate occurs at 155.1 ± 0.5 G, within the error of the position of the elastic peak at

. The correlation of the peak of the inelastic and elastic scattering rates is predicted both for two- and three-body inelastic collisions. The reason for this is that the Feshbach resonance is a resonance in the entrance channel of the collisions, and so all the rates, inelastic and elastic, are




Figure 6.2. Elastic Rate (upper plot) from Chapter V and total loss rates (lower) plot vs. magnetic field. The total loss rate is expressed as the sum of the two- and three-body loss as $K_2 + \beta K_3$. Because of initial density differences caused by ramping across the peak, β has a different value above and below $B = 157$ G. The vertical lines represent the positions of the elastic rate maximum and minimum at $B = 155$ and 166 G, respectively.

enhanced. Just as is the case for the elastic rates, the inelastic rates around the Feshbach resonance vary by orders of magnitude. The peak in the inelastic rate is much less symmetric, however. Another interesting feature is that the loss rates not only increase near the elastic rate peak but decrease near its minimum. The field where the loss is a minimum, 173.5 ± 2.5 G is higher than the minimum of the elastic rate at 165.85 ± 0.05 G.

6.2.3 Distinguishing Two- and Three-body Loss Rates. Even though two- and three-body inelastic rates have different density dependencies, it is surprisingly difficult to separate them. Figure 6.1 shows how a purely two-body or purely three-body loss curve fits equally well to a typical data set. An excellent signal-to-noise ratio or, equivalently, data from a large range of density, are required to determine whether the loss is three-body, two-body, or a mixture of the two. To better determine the density dependence, we decreased the initial density by up to a factor of 10 for several B values. For fields with relatively high loss rates, the signal-to-noise ratio was adequate to distinguish between two- and three-body loss rates. However, where the rates were lower this was not the case.

The reason that two- and three-body rates could not be distinguished at the fields with lower loss rates is the presence of the background loss rate τ . Decreasing the density by an order of magnitude causes the overall loss of atoms from the cloud to be dominated by this background rate, obscuring the loss rate due to inelastic collisions. If it were possible to increase the density by an order of magnitude, then it would be possible to distinguish between the two different loss processes, but the maximum density is limited by the evaporation performance.

Once data had been collected (at the higher loss rate fields) at both the higher and lower densities, each individual set of data was fit to equation (6.5). For each individual set of data, the initial number of atoms was treated as a free parameter, but the coefficients K_2 and K_3 were fixed across all the data at the same field, both high- and low-density. The two parameters K_2 and K_3 were then varied until the standard fit parameter χ^2 summed from all of the data points was minimized [127] to determine the best-fit value for K_2 and K_3 . Once that was done, K_2 and K_3 were then individually varied while minimizing all the other fit parameters (i.e. the initial number of atoms in each set of data and the other loss constant) until the value of χ^2 increased by one to determine the

statistical errors on the K_2 , K_3 determinations. The variation of individual data points taken under the same conditions was used to determine the individual uncertainty on each single data point. The values of χ^2 were statistically reasonable, indicating that equation (6.5) was a reasonable model of the loss rates.

6.2.4 Character of the Loss Rates near the ^{85}Rb Feshbach Resonance. Figure 6.3 shows the two-body inelastic rate determined from the density-varied data and Figure 6.4 shows the same for the three-body rate. The character of the inelastic loss clearly changes as one goes from higher field to lower field (right to left on Figures 6.2-4). On the far right of the graph, at $B = 250$ G, the inelastic rate is dominated by a three-body process. From $B = 250$ G down to $B = 174$ G the inelastic rate decreases to a minimum and then begins to increase again. Near the minimum, we could not determine the loss character, but at 162 G the losses are dominated by a two-body process at this density. At 158 G, the two-body process is still dominant and rising rapidly as one goes toward lower field. However, by ~ 157 G it has been overtaken by three-body recombination that is the dominant process between 157-145 G. At lower fields, both two-body and three-body rates contribute significantly to the total loss at these densities.

While the following is not confirmed, a reasonable interpretation of this structure is that the three-body rate that dominates at fields much higher than the Feshbach resonance decreases as the Feshbach resonance is approached from higher fields. As the three-body loss decreases, it becomes smaller than the two-body rate for the density of the atom clouds studied. In this interpretation, the density of the atom clouds that are measured determines the minimum of the combined loss rate, and it represents the point at which the three-body and two-body loss rates are equal. The two-body rates climb as the peak of the Feshbach resonance is approached from higher field, and it dominates the loss rate at the measured densities for some time until the three-body recombination rate, increasing more rapidly as the peak is approached from higher field, overtakes it to dominate the losses at the peak of the resonance. This interpretation is consistent with the general theory predictions for the two- and three-body losses near the Feshbach resonance.

6.2.5 Temperature Dependence of the Loss Rates. Along with varying the density, the initial temperature was varied at $B = 145, 156, 160,$ and 250 G. There was no significant rate change




Figure 6.3. The determination of the two-body inelastic rate for several B fields. The theory prediction from Ref. [119] is shown as a solid line. The open circles (\square) are the two-body rates determined from the total loss from Figure 6.2 by assuming (not explicitly measuring) that the loss between 162 and 166 G is predominantly two-body. These points are added to aid comparison with theory. In addition to the statistical errors shown, there is another 10% systematic uncertainty due to the estimated error in our density determination.




Figure 6.4. The determination of the three-body inelastic rate. The points with the droplines are to be interpreted as upper limits on the three-body rate. The solid line here is a prediction from Ref. [122]. The open circles here are similar to those in Figure 6.3, only here we assume (but again do not measure) that the loss above 175 G is predominantly three-body. The error bars on the 250 G point are relatively small because a large amount of data was taken there. In addition to the statistical errors shown, there is another 20% systematic uncertainty due to the estimated error in our density determination.

in the loss rate for temperatures between 400 and 1000 nK for the 160 and 250 G points. The combined loss rate increased by a factor of 8 at 156 G from 1 μ K to 400 nK, and by a factor of 2 at 145 G. Figure 6.5 shows the temperature dependence of the combined loss rate at 145.1 G as a function of temperature. This temperature dependence near the peak is expected in analogy to the temperature dependence of the elastic rates [23]. The fact that the loss rates near the peak exhibit both a two- and three-body character plus temperature dependence makes interpreting the data challenging.

In order to account for the combination of this temperature dependence and the heating of the sample that occurs while the density evolution is being measured, a naïve temperature dependence is included in the K_3 coefficient for the few points near the peak. It is assumed that the three-body coefficient decreases from its value at the initial temperature (~ 600 nK) as T^{-2} , based on the behavior of the elastic scattering unitarity limit (equation (5.20)) and the predicted scaling of the three-body rate equation given in equation (6.1). This temperature dependence ignores the fact that near the peak, the scattering length will eventually become larger than the deBroglie wavelength of the atoms in the cloud. Once this happens, the enhancement in loss rates due to the nature of the correlation function (subsection 5.2.2) will be reduced, causing a further decrease in the loss rate as the temperature of the cloud increases. The T^{-2} temperature scaling modifies the measured three-body rate by $\sim 20\%$, and this is included in the data displayed in Figures 6.2-6.4. Since this estimate is simplistic, the uncertainty on the points near the peak was increased adding 20% in quadrature to the statistical uncertainty. This temperature dependence did not affect our conclusion that the three-body losses dominate at the peak of the Feshbach resonance.

6.2.6 Heating Rates. As mentioned in subsection 6.2.2, the sample heated up during the density evolution at a rate that scales roughly with the inelastic scattering rate. Figure 6.6 shows the heating rate measured at several selected fields.

6.3 Comparison of Measured Inelastic Rates to Theory Predictions.

6.3.1 Dipole Relaxation. A calculated dipole relaxation loss rate is compared with the data in Figure 6.3. This was calculated using the ^{85}Rb atom-atom interaction potential including the magnetic dipole interaction and the second-order spin-orbit coupling and finding the scattering




Figure 6.5. Temperature dependence of the inelastic loss at 145.1 G. Here, the combined loss rate is shown by assuming that all of the loss at this field is due to a three-body process. All of the data on this plot were taken with approximately the same density.




Figure 6.6 Heating rate vs. magnetic field. The samples used to acquire this data had densities between $2\text{-}3 \times 10^{11} \text{ cm}^{-3}$. The heating rate is determined by dividing the temperature increase after the cloud had lost $\sim 25\%$ of the atoms by the time taken to lose that many atoms. Away from the zero, the heating rate is consistent with the rate expected from the density dependence alone. Near the minimum, the heating rate is higher than expected from the density dependence of the inelastic losses themselves, but near the minimum we are also more sensitive to heating rates from sources other than the inelastic collisions.

amplitude such that two $F=2$, $m_f = -2$ collide inelastically and exit in a spin state different than the incoming one [128]. Keeping in mind the significance of error bars on a log plot, the agreement is reasonably good. In the region between $B = 160$ and 167 G where the determination of K_2 is not complicated by three-body loss, the agreement is particularly good.

The predicted dipole relaxation rates exhibit a dependence on the energy of the two atoms involved in the collision, like the unitarity limit in elastic collisions. Away from the peak where the loss rates are small, this is not a significant effect for the typical temperatures of the atom clouds used to measure the inelastic rates. Closer to the peak, however, the energy dependence limits the dipole relaxation rate. The predicted dipole relaxation rate scaled as

$$K_2(T) = \frac{K_2(T=0)}{(1 + \alpha E)} \quad (6.7)$$

where E is the energy of the colliding pair and α is a constant that depends on the value of the magnetic field. I determined the value of α for each field from the calculation supplied in Ref. [128]. In order to generate the theory line on Figure 6.3, this energy dependence was accounted for by averaging the energy-dependent two-body rate equation (6.7) over the range of collision energies in a cloud in thermal equilibrium that had the same temperature as the clouds that were measured.

The calculation of the magnitude of the dipole relaxation rate was performed using an empirical coefficient for the second-order spin-orbit interaction that was used to match the rate observed in the optical pumping measurements (subsection 6.1.2.2). The agreement between the predicted rate and our observed rates in the region that we were best able to measure the two-body coefficient ($B = 160$ to 167 G) then represents additional evidence that the initial calculation of the second-order spin-orbit coupling was indeed too small [119] and that the empirical coefficient not only fits the room-temperature optical pumping data but our measurement of the ultracold two-body loss rate as well.

6.3.2 Three-body Recombination. In Figure 6.4, we also show a prediction of the three-body recombination rate. The predicted a^4 dependence of the loss rate was combined with the measured value of a as a function of magnetic field (Chapter V) to generate the theory curve in Figure 6.4. The value used for the coefficient C in equation (6.1) was taken from Ref. [122], using

the predicted value for C when the scattering length is negative and for when the scattering length is positive in the appropriate regions of magnetic field. Temperature dependence has not been included in the prediction. Qualitatively, the main features of the predicted three-body recombination rate match the data. The three-body rate decreases as a decreases and it increases rapidly where a diverges at $B = 155$ G. From $B = 155$ to 166 G, a is positive and the three-body loss is much smaller than it is at B fields with comparable negative a , as expected from theory. This overall level of agreement is reasonable given the difficulties and approximation in calculating three-body recombination rates.

Because of the loss of signal-to-noise ratio as the inelastic collision rate decreased, we could not definitively determine that the three-body losses were scaling as a^4 . Figure 6.7 shows all of the points higher than 175 G assuming that all of the observed loss is due to three-body collisions. The fit to the a^4 function is not particularly good, but it does not appear possible to rule out an a^4 scaling.

6.3.3 Energy Dependence of the Feshbach Peak. While measuring the background loss rate (τ), we found that the background rate appeared to decrease near the Feshbach resonance, indicating that the loss rate was enhanced for some fields. Figure 6.8 shows the decay rate of the low-density clouds used to make the measurement of τ as a function of magnetic field. We now believe that this increase in the loss rate at some fields is due to the fact that the density of the clouds is not low enough that the large losses near the peak of the resonance cannot cause some loss through inelastic collisions. The fact that this effect is largest near 162 G rather than 155 G is a signature of the energy dependence of the position of the Feshbach resonance peak.

6.3.4 Conclusions. The ^{85}Rb Feshbach resonance has a profound effect on the two- and three-body inelastic collision rates, changing them by orders of magnitude. Far from the resonance at 250 G, the two-body and three-body loss rates were measured to be

$$\begin{aligned} K_2 &= (1.87 \pm 0.95 \pm 0.19) \times 10^{-14} \text{ cm}^3/\text{s} \\ K_3 &= (4.24_{-0.29}^{+0.70} \pm 0.85) \times 10^{-25} \text{ cm}^6/\text{s} \end{aligned} \quad (6.8)$$

respectively. Here the first error listed in the statistical uncertainty and the second is the systematic uncertainty due to the number determination. Near the resonance, the two- and three-body rates




Figure 6.7 Conjecture log-log plot of K_3 vs. a^4 . This figure displays all of the data greater than 175 G assuming that the loss is due only to a three-body loss process. The slope of the best-fit line is 4.2, close to the predicted value of 4, but the quality of the fit is so poor (reduced $\chi^2 > 4$ for 5 degrees of freedom) that we certainly cannot claim that we have definitively observed an a^4 scaling. The intercept of this fit implies $C \sim 2200$ in equation (6.1) for a negative scattering length.




Figure 6.8. Variation of measured loss rate of low-density clouds. The clouds used for this measurement had a density $\sim 10^9 \text{ cm}^{-3}$ and a temperature of approximately $40 \mu\text{K}$. Note that the zero is suppressed and the total variation is only 20%.

change by orders of magnitude. The dependence of the inelastic rates on magnetic field is similar in structure to the dependence of the elastic rate: the maxima in both rates occur at the same field, while the minima are close but do not coincide. The total loss is a complicated mixture of both two- and three-body loss processes. They have different dependencies on field so both have field regions in which they dominate.

6.4 Appendix to Chapter VI: Polarization-rotation Imaging.

6.4.1 Why Choose Polarization-rotation Imaging for Non-destructive Imaging? Since we work at relatively high magnetic fields, polarization rotation imaging [129] is a natural choice for non-destructive imaging. Other methods of non-destructive imaging, such as dark ground [130] and phase contrast imaging [131], involve placing a small object in the Fourier transform plane of the imaging system to effect the light that did not go through the atom cloud in a way that is different from the light that did go through the cloud. This means that these methods require that the cloud being imaged be small so that the image of the cloud is big in the Fourier transform plane. For the size of clouds that we were measuring, it was difficult to achieve this separation of scales when we tried dark-ground imaging. Also, small fringes and scattering patterns from dust in the probe beam did not pass through the object in the Fourier transform plane either and complicated the interpretation of the dark-ground images.

Polarization rotation imaging takes advantage of the fact that the high bias magnetic field in the trap breaks the degeneracy of the excited state magnetic sublevels. It is then possible to select a probe laser detuning that is in between two transitions that are driven by different polarizations. Not only do atoms scatter light, they also can impart a phase shift to it [81]. If a probe beam that is composed of these two different polarizations passes through a gas of atoms, the different phase shift experienced by each polarization will result in a change in the polarization of the probe beam. The magnitude of the phase shift depends on the column density of the atoms through which the light passes and so the amount of polarization change can be used to determine the column density of the cloud. The magnitude of the phase shift scales as Δ^{-1} , where Δ is the detuning from the atomic transition, while the scattering rate scales as Δ^2 . Because the scattering rate drops much faster with

Δ than the phase shift, for large detunings it is possible to measure the column density of the cloud while scattering very few photons on average, leaving the cloud largely undisturbed.

6.4.2 Arrangement of Polarization-rotation Imaging in our System. Since the probe laser already passes through a polarizing beam-splitting cube, that cube is a natural choice as a way to measure the polarization of the laser light. Normally, the polarization of the laser light is selected so that the transmission of through the cube is maximized. By inserting a $\lambda/2$ plate in the probe beam path, it was possible to arrange the polarization of the probe laser so that no light would be transmitted through the cube. We were able to achieve an extinction ratio of $\sim 0.1\%$. With this arrangement, the probe light was circularly polarized in the region where the atoms were trapped. Any change in the polarization of the probe light would result in more light being transmitted through the polarizing beam-splitting cube.

The energy dependence of the excited states in ^{85}Rb is shown in Figure 6.9. The states are labeled by their nuclear spin and electron total angular momentum (spin plus angular momentum) projections (m_i, m_j) that specify the magnetic sublevel eigenstates at very high magnetic fields (where the hyperfine interaction can be regarded as a small perturbation). At finite magnetic fields, the eigenstates are composed of a superposition of $|m_i, m_j\rangle$ product states, and it is possible to drive transitions from the $F=2, m_f = -2$ ground state to the several of the states indicated in Figure 6.9 while still satisfying the electric dipole selection rules. When the probe laser is detuned as shown, a different phase was accumulated between the two polarization components of the circularly polarized probe beam and the polarization of the beam is changed.

The precise calculation of the phase shift from the atomic structure of ^{85}Rb is quite complicated. One complication is that the direction of the magnetic trap bias magnetic field is perpendicular to the probe beam direction. If the direction of the bias field of the trap is labeled by the z-axis and the probe beam propagates along the x axis, then the polarization of the probe beam is initially proportional to $i\hat{y} + \hat{z}$. The z component of the polarization drives $\Delta m = 0$ transitions, while the y direction will drive a coherent superposition of $\Delta m = 1$ and -1 transitions. Evaluating this coherent superposition is complex and simple rate equations cannot be used for the detunings typically used for non-destructive imaging. The transitions are not closed, and so both the multiple




Figure 6.9. Excited state magnetic sublevel energies vs. magnetic field. Each manifold of states is labeled by the value of m_j that is one of the quantum numbers that specifies the eigenstates at high magnetic field. The individual levels within the four manifolds are further specified by the nuclear spin quantum number m_i . Moving up and down from one level to the next within a manifold changes the value of m_i by one. The m_i value for the highest and lowest states in each manifold is shown on the right. These quantum numbers are strictly good only at high values of magnetic field—at lower values of the field the eigenstates are a combination of the product states specified by $|m_j, m_i\rangle$. The arrow indicates the probe laser detuning suitable for polarization-rotation imaging.

excited states and appropriate decay rates to the ground states must be included in the calculation of the transition strength. Also, the fact that the transitions are not closed precludes the evaluation of the transition rates in the steady state. Another complication is that when an ^{85}Rb atom's optically active electron is excited in the y direction, the magnetic field will cause it to start moving along the x direction as well. This means that the electric dipole polarization of the atomic sample will be along all three directions: x, y, and z and the evaluation of Maxwell's equations in the slowly varying envelope approximation [81] will be correspondingly more complicated as well.

Rather than attempt this complicated calculation, we merely assumed that an atom in the trap induced some phase shift between the y and z polarization without calculating exactly what it would be. From the general considerations of the polarization of the probe laser and the orientation of the atoms it was possible to make an intelligent guess as to what probe detuning would be required for this to be so. Whatever the phase shift of an individual atom, the total phase shift induced in the probe polarization would be proportional to the column density of the atoms, at least in the absence of any significant photon scattering rate. The number of atoms thus measured was then calibrated by comparing the polarization rotation signal to the MOT recapture or absorption imaging measurement of the number of atoms in the cloud.

With the arrangement of the polarization of the probe beam so that no light is transmitted in the absence of any atoms, any intensity observed when atoms are present in the trap is given by

$$I(y, z) = I_0(y, z)(1 - \cos(\phi_0 n(y, z)))^2 \quad (6.9)$$

where ϕ_0 is the phase shift induced by a single atom, $n(y, z)$ is the column density of the cloud at position (y, z) and $I_0(y, z)$ is the intensity of the incident probe light. The atoms then appear as a bright intensity signal on the CCD array of the camera. For small phase shifts, the intensity of transmitted light is proportional to the square of the column density.

6.4.3 Performance of the Polarization-rotation Imaging. The polarization-rotation imaging did allow nearly non-destructive measurement of the thermal atom samples while in the trap. Typically, about 1% of the atoms in the sample were lost due to the $\sim 40 \mu\text{s}$ probe pulse used and the sample did heat up by several nK, but this loss and heating was easily accounted for and it was a great advantage not to have to create several samples to map out the density evolution. The

analysis of the signal was complicated by the presence of a spatial variation in the polarization of the laser beam, however. These polarization “stripes” appeared to be due to interference in the $\lambda/4$ waveplates used to set the MOT polarization and perhaps other optical elements in the system. The stripes had a magnitude of near 0.1 radians, producing a significant shift in the measured phase shift due to the atoms in the magnetic trap. The spatial structure of these stripes was easily measured, and in order to measure their phase amplitude so to correct for these stripes, long exposure time (~ 500 μs) images were taken of a trapped atom cloud. These long exposure times were destructive but produced very high signal-to-noise images. The cloud measured with these long exposures was fit to a gaussian and the deviations from this fit used to determine the phase amplitude of the stripes. Once the phase amplitude and spatial structure of the polarization stripes were known, corrections could be included in the fits to the non-destructive data to determine the cloud shape and number of atoms. The validity of these corrections was demonstrated by the fact that the measurement of the K_2 and K_3 coefficients was the same for the polarization rotation and destructive absorption imaging methods.

The signal-to-noise ratio achievable with the polarization-rotation imaging was not nearly as good what as could be obtained with absorption imaging. This is not particularly surprising since the large detuning of the probe laser used in the polarization-rotation measurement leads to a reduction in the signal that can be obtained from each atom. Therefore, only samples with a sufficient column density (equivalent to an on-resonance absorption peak optical depth of at least ~ 10 or more) could be usefully measured with polarization-rotation imaging. Moving the probe laser detuning closer to resonance produced very confusing signals. By changing the frequency of the probe, we determined that $\sim 25\%$ of the light always passed through the cloud. Near resonance, the phase-shift induced by the atoms in the cloud exceeded 2π , meaning that light and dark polarization signals were superimposed on the absorption shadow. The spread of the magnetic field across the sample was enough that we could not be sure that the shape would not be distorted even when the probe was tuned to the resonance of the transition. Finally, the absorption of only one of the polarization components of the light meant that the absorption was also producing a polarization rotation in addition to a straightforward absorption signal. These confusing signals limited this

method's utility for destructive in trap imaging. More control over the polarization of the probe laser beam, the optics used to measure the polarization, and the ability to change the direction the probe propagates through the sample would allow better imaging of the atoms in the trap. Making these changes, however, would involve substantial changes in the present imaging system, and the relative lack of optical access in the science cell means that the path the probe laser follows must also be able to be used for the science MOT as well.

Chapter VII

Stable ^{85}Rb Bose-Einstein Condensates with Widely Tunable Interactions

7.1 Evaporative Cooling of ^{85}Rb

7.1.1 Difficulties Evaporatively Cooling ^{85}Rb . As was mentioned in subsection 4.9.1, the ratio between the elastic collision rate needed for efficient evaporative cooling and the rate of loss and heating determines the efficiency of the evaporative cooling. In ^{85}Rb , the elastic rate is a function of field and can be greatly enhanced or suppressed through the use of the Feshbach resonance (Chapter V). The inelastic collisions rates that lead to loss and heating of the trapped sample are also a function of field, however, and the change of these rates with magnetic field is similar to changes of the elastic rate (Chapter VI). This was clearly demonstrated by our first attempts at evaporative cooling after the characterization of the elastic scattering rates as a function of magnetic field. We tried cooling at the peak of the Feshbach resonance and the results were a rapid loss of atoms without significant cooling. This was our first indication that the inelastic collisions could dominate the elastic collisions and therefore a careful selection of magnetic fields at which to evaporatively cool would be necessary, especially since the inelastic collisions rates were in general orders of magnitude higher than those in ^{87}Rb and ^{23}Na (the two species in which a BEC was first created in a dilute trapped gas).

In addition to the relatively large inelastic rates, there is another evaporative cooling drawback in the scattering properties of cold ^{85}Rb . In ^{87}Rb and ^{23}Na , adiabatic compression (subsection 4.6.7) is used to improve achieve efficient evaporative cooling. This adiabatic compression consists of increasing the frequencies of the magnetic trap to increase the density and temperature of the sample and hence the elastic collision rate, which is proportional to $\langle n \rangle \sigma \langle v \rangle$, where as usual $\langle n \rangle$ is the average density (equation (4.12)), $\langle v \rangle$ is the average relative speed (equation (4.17)) and σ is the elastic scattering cross section (equation (2.14)). The elastic collision rate is usually increased by ~ 100 using adiabatic compression. In contrast, there are several reasons

why adiabatic compression cannot be used in ^{85}Rb , and so the two orders of magnitude increase in the elastic collision rate are lost. First, away from the Feshbach resonance, σ drops rapidly with increasing collision energy, as is shown in Figure 7.1. The increase in temperature during adiabatic compression therefore reduces the elastic collision cross section, leading to no significant increase in the elastic collision rate. Second, the combination of the location of the Feshbach resonance at high magnetic bias field (~ 155 G), the scaling of the radial trap frequencies with bias field (equation (4.9)), and the finite range of both our trap coil cooling and coil power supply means that it is impossible for us to achieve large trap frequencies in the vicinity of the resonance. Finally, the scaling of the relative elastic and inelastic collision rates on trap frequency, detailed below, indicate that weaker magnetic traps are in general more favorable for evaporative cooling; adiabatic compression increases both the elastic and inelastic scattering rates, and it does so in a way that affects their ratio unfavorably.

One other difficulty comes about because the weak magnetic traps that are used with ^{85}Rb allow the atoms to sag substantially due to gravity. This means that the surfaces of constant magnetic field form horizontal planes across the cloud and energy is removed by the rf-knife from only one direction instead of all three (see subsection 4.5.4). This change in dimensionality of the energy removal is aptly referred to as one-dimensional evaporation. One-dimensional evaporation leads to a loss of efficiency in evaporative cooling, since atoms with much greater energy than the average energy in the cloud will not be removed in two out of the three directions. This loss in efficiency combined with the other difficulties in evaporating ^{85}Rb made it highly uncertain whether or not an ^{85}Rb BEC could be created.

7.1.2 Elastic/Inelastic Collision Rates and Trap Frequency. The goal of evaporative cooling is ultimately to increase the phase-space density of the trapped cloud of atoms until the BEC transition is reached. Part of the optimization of our evaporative cooling is to choose an appropriate strength for the trap potential. Unlike many of the other parameters in our evaporative cooling sequence, the strength of the trap potential is not easy to optimize. This is because after any change in the trap parameters we have found that we need to optimize all of the other parameters involved in




Figure 7.1. Elastic scattering cross section as a function of collision energy. This curve is calculated for $B=0$, using the methods presented in Ref. [31] and the cold atoms collisions calculator in Ref. [132]. The behavior at fields far from the Feshbach resonance is similar.

the evaporation and that takes one day or more. Therefore, the trap potential strength was selected on the basis of the following general arguments.

For convenience, several equations are repeated here: the BEC transition temperature

$$k_B T_c = 0.94 \hbar \omega_{ho} N^{\frac{1}{3}}, \quad \omega_{ho} = (\omega_x \omega_y \omega_z)^{\frac{1}{3}}; \quad (3.5)$$

the average density of atoms in the magnetic trap

$$\begin{aligned} \langle n \rangle &= \frac{1}{N} \int d^3 x n(\vec{x})^2 \\ &= \frac{N}{8 \mathbf{p}^{3/2} \mathbf{s}_r^2 \mathbf{s}_z} = \frac{N \omega_r^2 \omega_z}{8 \left(\frac{\mathbf{p} k_B T}{m} \right)^{3/2}}; \end{aligned} \quad (4.14)$$

the average relative speed of the atoms in the trap

$$\langle v \rangle = \int d^3 x \int d^3 v \int d^3 v' |\vec{v} - \vec{v}'| f(\vec{x}, \vec{v}) f(\vec{x}, \vec{v}') = \sqrt{\frac{16 k_B T}{\mathbf{p} m}}; \quad (4.17)$$

the elastic collision rate

$$\Gamma = 2 \langle n \rangle \mathbf{s} \langle v \rangle; \quad (5.16)$$

(the factor of two is included because of the correlations in a thermal cloud of identical bosons) and the loss rates

$$\frac{dN}{dt} = -K_2 \langle n \rangle N - K_3 \langle n^2 \rangle N - \frac{N}{\mathbf{t}}. \quad (6.3)$$

All of these quantities depend on the temperature of the atoms in the trapped cloud and the trap frequencies (ω) either explicitly or through a dependence on the density of the atoms in the trap.

Adiabatic compression can be used to increase elastic collision rate while maintaining the phase-space density of the cloud. Since during an adiabatic compression the quantity ω_{ho}/T is fixed, it is straightforward to determine the scaling of the elastic scattering rate for a fixed phase-space density assuming that the elastic cross section is independent of temperature

$$\Gamma \propto \omega_{ho}^2 \quad (7.1)$$

However, if the cross section falls as $1/T$ then Γ scales only as ω_{ho} . Not only does the elastic collision rate increase for stronger traps, the higher BEC transition temperature means that any

absolute heating rate becomes a less important consideration. The steep scaling of equation (7.1) and higher absolute value of the transition temperature seem quite attractive from the standpoint of creating a BEC, but the inelastic collision rates do not scale so favorably.

If three-body recombination dominates the loss rate, then the ratio of elastic/inelastic scattering is

$$\frac{6\sqrt{3} \mathbf{p} \mathbf{s} k_B^2 T^2}{N \mathbf{w}_{ho}^3 m^2 K_3} \quad (7.2)$$

where MKS units are used for all of the quantities. Clearly, increasing the trap frequencies will cause this ratio to get worse: an adiabatic compression will cause this ratio to scale as $1/\omega_{ho}$. Also, as the cloud gets colder, this ratio becomes more and more unfavorable. For 10^7 atoms a factor of 100 away in phase-space density from the BEC transition, a tightly confining trap of $2\pi \times 200$ Hz radial by $2\pi \times 10$ Hz axial will give a value of 6.5 for equation (7.2) if $K_3 = 5 \times 10^{-37} \text{ m}^6/\text{s}$ and $\sigma = 1.0 \times 10^{-14} \text{ m}^2$. Typically, a ratio of 100 is necessary for good evaporation performance with three-dimensional evaporation, even without considering any heating of the cloud, and so a tight trap is not favorable for ^{85}Rb .

If two-body losses dominate instead, the elastic/inelastic collision ratio is

$$\frac{4\mathbf{s}}{K_2} \sqrt{\frac{k_B T}{\mathbf{p} m}} \quad (7.3)$$

Once again, as the sample gets colder, this ratio gets worse. Now, adiabatically increasing the trap frequencies should help the evaporation performance, but only very slowly since equation (7.3) scales as $\omega_{ho}^{1/2}$. For a typical set of parameters (10^7 atoms a factor of 100 away from the BEC transition, $2\pi \times 17$ Hz radial and $2\pi \times 7$ Hz axial trap frequencies, $\sigma = 0.25 \times 10^{-14} \text{ m}^2$, $K_2 = 7 \times 10^{-20} \text{ m}^3/\text{s}$) this ratio is ~ 275 .

Equation (7.2) indicates that strong trap potentials are to be avoided if the three-body recombination rate can become the dominant loss mechanism during evaporation. Equation (7.3) suggests that stronger trap frequencies would help the elastic/inelastic ratio if two-body inelastic collisions are the dominant loss mechanism, but as the trap frequencies are increased the ‘‘thickness’’ of the atomic cloud begins to limit the evaporative cooling efficiency. As long as the cloud of atoms

in the trap is dilute, a high-energy atom produced anywhere in the cloud by an elastic collision will make it to the evaporation surface. As the cloud gets denser (thicker), however, it becomes more likely that an atom produced in the center of the cloud will collide on the way out of the cloud, reducing the evaporation efficiency. With respect to an adiabatic change in the trap frequencies, the elastic collision rate scales as ω_{ho}^2 while the radial width of the cloud scales as $1/\omega_r$ so the atomic thickness then scales as $\omega_{ho}^2/\omega_r = (\omega_r\omega_z^2)^{1/3}$, and the clouds become thicker with increasing trap frequency. Not only does this increase in atomic thickness reduce the evaporative cooling rate through reducing the efficiency of the elastic collisions, the thicker clouds increase the probability that a hot atom produced through an inelastic collision will scatter on the way out of the cloud, depositing some of its kinetic energy as heat [133].

In addition to the problem of the atomic thickness of the cloud, there is a related effect due to the finite amount of time that an atom takes to move in the harmonic trapping potential. A high-energy atom formed at the center of the cloud that is moving toward the evaporation surface will take one-quarter of the trap period to get there. Depending on the initial location and velocity of a high-energy atom, this time could range from nearly zero to nearly the full radial trap period. Therefore, when the elastic scattering rate exceeds roughly twice the radial trap period, the rate of atoms reaching the evaporation surface will be proportional to roughly twice the trap frequency rather than the elastic collision rate. The inelastic collision rate is unaffected by these considerations, however, and so this effect also decreases the evaporative cooling efficiency.

Again, increasing the trap frequency makes this problem with the finite time an atom takes to move worse. Once the elastic collision rate is limited by the radial trap frequency, the rate will scale as ω_r instead of ω_{ho}^2 during an adiabatic change in the trap frequencies. However, the two-body inelastic rate still scales as $\omega_{ho}^{3/2}$. Instead of the elastic/inelastic ratio improving with an increase in trap frequency, it will begin to decrease once the elastic scattering rate is comparable to twice the radial trap period.

One other effect that needs to be considered is the unitarity limit (equation (5.20)) that limits the elastic collision cross section. Once the collisions are unitarity limited, the elastic cross section will scale roughly as $1/T$, making the ratio in equation (7.3) worse with increasing

temperature. The inelastic collision rates are also smaller at finite temperature as well, mitigating the loss in evaporative cooling efficiency at higher temperatures. However, the inelastic collisions are causing loss throughout the cloud while only relatively high-energy collisions produce an atom that can be removed by the rf-knife, so the unitarity limit would still be expected to degrade the evaporation efficiency.

The optimum trap frequencies are therefore a trade-off between the stronger trap favored by equation (7.3) and the three limiting effects discussed above. For the typical set of parameters presented after equation (7.3), the collision rate is ~ 22 Hz, near the point where the thickness effects are expected to play a role. The unitarity limit is not expected to play a significant role for these conditions because the s-wave scattering length is small enough. Therefore, it seemed that we did not want to increase the trap frequencies significantly, nor decrease them too much from these typical parameters, and therefore these are the trapping frequencies used for the final portion of our evaporation.

The limitations in evaporative efficiency due to atomically thick clouds also need to be considered when selecting the bias magnetic field at which the evaporation is to be performed. The ratio of the elastic scattering cross section (σ) to K_3 and K_2 is determined by the value of the magnetic bias field, and it is in general desirable to choose a field that maximizes that ratio in favor of the elastic collisions. Increasing the elastic collision cross section too much makes the cloud atomically thick, however, and so just maximizing the σ/K_2 or σ/K_3 ratio does not necessarily lead to the most efficient evaporative cooling (see subsection 7.1.5).

In addition to a two- or three-body rate dominating the loss from the cloud, it is also possible that the background loss rate τ will be the dominant loss process if the clouds have a low enough density. After we have loaded the atoms into the magnetic trap initially, the background losses are indeed the limiting loss rate. In this case, there is no strong dependence of the loss rate on the trap frequencies and the optimum elastic/ “inelastic” ratio is just the optimum of the elastic collision rate.

7.1.3 Simple Model of Evaporative Cooling. The precise calculation of the evaporative cooling rate is complicated, and so we have developed a simple model to compare with our observed

evaporation performance. The model is not expected to produce quantitative comparisons with our observations, but rather it serves as a way to test how various factors are expected to affect the evaporation performance. The model is similar to the one presented in Ref. [134], except that it has been modified to account for our one-dimensional evaporative cooling and the effects of a thick atomic cloud.

The calculation of the evaporative cooling rate is performed in the following way. The frequency of the rf-knife is assumed to change so that the evaporation plane removes all atoms whose energy in only one direction (which for concreteness will be referred to as the y direction) is greater than the average energy of the cloud in that direction by a fixed ratio η . In other words, if η is set to be equal to six, all atoms with energy in the y direction greater than $6k_B T$ will be removed from the cloud. The average energy of the atoms removed is set to be $\eta + \kappa$, where κ is a factor that reflects the fact that not all the atoms that are removed by the rf-knife have an energy exactly equal to $\eta k_B T$. κ has not been rigorously calculated and is typically set between 0 and 2. The rate of atoms created by elastic collisions with energies greater than η is calculated by using the idea of detailed balance. The collision rate between atoms with energy greater than η in the y direction and atoms with energy less than that is calculated by summing over all such collisions in a cloud in thermal equilibrium, and this rate will be referred to as ζ . It is then assumed that each one of these collisions results in neither atom in the pair having energy greater than η in the y direction, which is usually true. Further, the collisions between atoms both with y direction energy greater than η are ignored. The principle of detailed balance implies that in thermal equilibrium, the number of atoms with a given energy is constant. Therefore, with the assumptions that have been made the rate at which high-energy atoms collide with low-energy atoms is equal to the rate at which two low-energy atoms collide to produce a high-energy atom in order to maintain the detailed balance. The rate ζ is then the rate at which atoms with energy greater than η in the y direction are produced as a function of time. For values of $\eta > 3$, it is found that ζ is approximately equal to $\Gamma \exp(-\eta)$. If evaporation occurs in three dimensions, this quantity is $\Gamma \eta \exp(-\eta)$, and the decrease in evaporation rate due to the reduction in evaporation dimensions can be readily ascertained.

The effects of the thickness of the atomic cloud are first included with a calculation of the quantity referred to as the evaporation atomic depth. Similar to the optical depth, the evaporation atomic depth is an “absorption length” for atoms, and determines the probability that an atom with enough energy to be evaporated will collide elastically on the way to the evaporation surface. Rather than sum over all of the possible positions of origination and initial velocities, the evaporative atomic depth is calculated for an atom that starts at $y=0$ and has an energy equal to η , all in the y direction, which is considerably easier. Two cases are considered, one where the atom is initially moving toward the evaporation plane and one where it is moving away. The evaporation atomic depth calculated in this way is then averaged over all the positions (x,z) in the cloud to calculate a net depth. Assuming that each elastic collision prevents the atom from getting to the evaporation surface, the total rate of atoms to the evaporation surface after averaging over all (x,z) is then

$$\mathbf{z}' = \frac{\Gamma}{2} \left(\frac{1 - \exp(-AD)}{AD} + \frac{1 - \exp(-3AD)}{3AD} \right) \exp(-\mathbf{h}) \quad (7.4)$$

$$AD = \frac{2\Gamma}{\mathbf{w}_r}$$

The reduction of evaporation efficiency due to the finite time an atom takes to move to the evaporation plane is included in an even less rigorous manner. It is assumed that the net rate of atoms to the evaporation surface was equal to ζ' when Γ is much less than ω_r/π and equal to ω_r/π when Γ was much more than ω_r/π . To interpolate between these extremes, a simple function is used

$$\mathbf{z}'' = \frac{1}{\sqrt{\frac{1}{\mathbf{z}'^2} + \frac{\mathbf{p}^2}{\mathbf{w}_r^2}}} \quad (7.5)$$

In addition to the evaporation rate, the loss rates also need to be included in the model. This can be accomplished quite straightforwardly using equation (6.3) and equation (6.6). The evaporative cooling and losses are then put together in a system of differential equations for the total energy (E) and total number (N) of atoms in the cloud shown in Figure 7.2. The dependence on the trap frequencies is not shown explicitly but rather is contained in the calculation of the density of the cloud.

$$\frac{dN}{dt} = (R_{evap} + R_{BG} + R_{2BDY} + R_{3BDY})$$

$$\frac{dE}{dt} = (\bar{E}_{evap} R_{evap} + \bar{E}_{BG} R_{BG} + \bar{E}_{2BDY} R_{2BDY} + \bar{E}_{3BDY} R_{3BDY})$$

$$T = \frac{E}{3Nk_B} \text{ at all times (i.e. the cloud is always in thermal equilibrium)}$$

$$R_{evap} = -Nz'' \text{ from equation (7.5)}$$

$$\bar{E}_{evap} = (\mathbf{h} + \mathbf{k})k_B T$$

$$R_{BG} = -\frac{N}{\mathbf{t}}$$

$$\bar{E}_{BG} = 3Nk_B T$$

$$R_{2BDY} = -K_2 \langle n \rangle N, \text{ where } \langle n \rangle \text{ is defined by equation (4.12)}$$

$$\bar{E}_{2BDY} = \frac{9}{4} Nk_B T$$

$$R_{3BDY} = -\frac{8}{\sqrt{27}} K_3 \langle n \rangle^2 N$$

$$\bar{E}_{3BDY} = \frac{2}{3} Nk_B T$$

Figure 7.2. The set of differential equations used in the simple model of evaporative cooling.

It is also possible to add additional terms to the simple model to account for the heating of the cloud due to other sources besides the density-dependence of the inelastic collisions. These other sources include such things as the products of the exothermic inelastic collisions depositing energy in the cloud, parametric heating of the cloud due to trap strength variations at twice the trapping frequencies (subsection 4.10.2), or glancing collisions with atoms in the background gas [133].

7.1.4 General Evaporation Strategy for ^{85}Rb . In order to form a BEC in ^{85}Rb , the evaporative cooling must be sufficient to cool the cloud to the transition temperature. The evaporation performance is characterized by an evaporation trajectory. This evaporation trajectory is simply the elastic collision rate plotted versus the temperature of the cloud as the cloud is cooled. If the elastic scattering cross section is temperature-independent, the elastic collision rate and the peak optical depth measured in absorption imaging scale in the same way with number and temperature in a harmonic trap. Therefore, the evaporation trajectory is often plotted as the peak optical depth as a function of the temperature of the cloud. See Figure 7.3 for an example of an evaporation trajectory. When evaporative cooling is working well, “runaway” evaporation occurs [98]. As the cloud of atoms is cooled, the density of the cloud increases fast enough that the elastic collision rate climbs as the temperature gets colder and colder. The increasing collision rate means that the evaporative cooling is working more and more efficiently the colder the cloud gets and therefore an arbitrarily low temperature can be achieved unless some change in the collision physics occurs that stops runaway cooling. Often, the change in physics is the creation of a BEC at extremely low temperatures.

The formation of ^{85}Rb BEC is more “stumble across the finish line” evaporative cooling rather than runaway evaporation. Initially, the elastic collision rate does increase as the cloud is cooled, and the evaporation works fine. However, as the cloud cools further, the losses in ^{85}Rb become dominated by two- or three-body losses at temperatures far from the BEC transition. Since the ratio of elastic/inelastic collisions scales with temperature as $T^{1/2}$ or T^2 , as the sample is cooled the evaporation efficiency must necessarily decrease. Rather than hoping that the collision rate increases as the sample cools, the best that can be done in ^{85}Rb is to maintain sufficient evaporative cooling until the BEC transition is reached.




Figure 7.3. Optimum evaporation trajectory. The evaporation trajectory at 162.5 G is shown as the filled circles in this plot. The open circle represents the optical depth of the cloud of atoms when it is initially loaded. The cooling from ~ 45 to ~ 2 μK was performed at 251 G. The prediction of the optical depth as a function of temperature by the simple model is also shown as the solid line. To generate this curve, the following values were used in the model: $\eta=5$, $\kappa=1.6$, $K_2=5 \times 10^{-20} \text{ m}^3/\text{s}$, $a=167a_0$, $\tau=300 \text{ s}$, $\omega_r = 108.5$, $\omega_z = 42.4$. The value of κ was adjusted to fit the measured trajectory.

Our original attempts to create ^{85}Rb BECs were not successful—the elastic collision rates fell dramatically before the BEC transition temperature was reached. One breakthrough was taking advantage of the steep scaling of the phase-space density with temperature in a magnetic trap (phase-space density is proportional to N/T^3). Unless the loss of atoms from the cloud is very high, almost any decrease in temperature will increase the phase-space density. A crucial step in forming ^{85}Rb BECs was to continue to evaporatively cool even after the evaporation started to flag. Even though we are very close (within 10%) of reaching the transition temperature with a nearly million atoms in the trap, we are only able to form BECs of up to 18,000 atoms. I would say that the resolution to keep evaporating until there were no atoms left or a BEC was created (the “BEC or was one of the more important steps in our creation of ^{85}Rb BECs. This strategy only led to the formation of a BEC at some values of the bias magnetic field, however. At other fields the losses were indeed so high that the phase space density did not increase, steep temperature dependence notwithstanding.

To extend the initial evaporation in which the collision rate increased as the cloud was cooled, we used the following general strategy. We evaporatively cooled the cloud at a magnetic field only until the evaporative performance started to drop, as signified by the elastic collision rate beginning to decrease rather than increase as the cloud was cooled. At this point, we knew that the inelastic collisions were beginning to overwhelm the elastic collisions. This could just be the temperature dependence of equation (7.2) or (7.3) alone, or the elastic collision rate may have increased to nearly the radial trap frequency causing the evaporative cooling efficiency to suffer. If the elastic/inelastic collision rate ratio itself has become too small, then the field was changed in order to improve the ratio, if possible. If the atomic thickness of the cloud is the limitation, then the field is changed in a way that reduces both the inelastic losses and the elastic collision rate. The reduction in the elastic collision rate alleviates the drop in efficiency caused by the thickness of the cloud of atoms. It is necessary to reduce the inelastic collision rate at the same time to maintain a favorable elastic/inelastic collision ratio.

Eventually, changing the value of the magnetic field can no longer extend the temperature over which the evaporation works well. The elastic collision rate then decreases as the cloud is

cooled and the only thing that can be done is to continue evaporating and hope that the BEC transition temperature is reached before all of the atoms go away. Fortunately, for some values of the magnetic field a BEC is formed before all of the atoms are lost.

7.1.5 Optimization of the Evaporation Trajectory. There are two aspects to the optimization of the evaporative cooling of ^{85}Rb . Once a magnetic field is selected, the parameters associated with the rf-knife had to be optimized by trial and error. The other aspect of the optimization of the evaporative cooling was the selection of the magnetic fields at which to evaporate. Trial and error as guided by the above considerations was used in the selection of the magnetic fields as well, but for each change in field the rf parameters had to be revisited, making the test of the evaporation performance time-consuming.

The parameters associated with the rf-knife that were optimized were the rf power and the frequency both as a function of time. The evaporation sequence was divided into a series of stages. Each stage corresponded to a factor of two decrease in the temperature of the cloud. The rf power and frequency as a function of time were optimized to maximize the number of atoms at the end of each evaporation stage (i.e. the elastic collision rate was maximized at a fixed temperature). The rf power was typically varied in step sizes of 5 dBm (a factor of 3). The evaporation performance was not particularly sensitive to the power even with these large changes in magnitude. For instance, an order of magnitude increase in the rf power would decrease the number of atoms left at the end of a later stage by less than a few percent. In general, as the cloud of atoms became colder the optimum power in the rf-knife became smaller. This behavior is expected since the width of the rf-knife depends on the power in the knife. As the cloud of atoms is cooled, the spread in magnetic field across the cloud decreases and therefore a narrower rf-knife is needed to differentiate between the low and high energy atoms in the cloud so that it can selectively remove the highest energy atoms for evaporative cooling.

The frequency of the rf-knife as a function of time ($f(t)$) was set to have the form of an exponential decay:

$$f(t) = (f_{start} - f_{base}) \exp\left(\frac{-t}{T}\right) + f_{base}, \quad T = \frac{-t_{stage}}{\ln\left(\frac{f_{stop} - f_{base}}{f_{start} - f_{base}}\right)} \quad (7.6)$$

The four parameters used to specify the exponential curve shown in equation (7.6) are the start frequency (f_{start}), the stop frequency (f_{stop}), the baseline frequency (f_{base}), and the time between the start and stop frequency (t_{stage}). The start and stop frequency are the initial and final frequencies of the rf-knife for the stage. If the baseline is set to be the frequency that corresponds to the magnetic field at the center of the trapping potential, then the rf-knife should ideally maintain a fixed value of the parameter η (see subsection 7.1.3) as it cuts into the cloud since $f(t)$ has an exponential shape (in fact, this is why $f(t)$ is set to be an exponential function).

The parameters in equation (7.6) were varied to optimize the rf frequency as a function of time. The start frequency was varied to optimize the performance only if the stage was the first one at a new magnetic field, otherwise the start frequency was set to be the stop frequency of the previous stage. The stop frequency was set so that the evaporation stage cooled the cloud by roughly a factor of two. The baseline frequency and t_{stage} were then adjusted to maximize the number of atoms at the end of the stage. Adjusting the baseline frequency is equivalent to changing the rate of change of $f(t)$ during the evaporation stage to cut more or less quickly as the stage progresses.

The performance of the evaporation is robust in the sense that large changes in the rf-knife parameters are required in order to change the performance significantly. The quantity t_{stage} was typically varied by 50% or more during the optimization, and the performance was not particularly sensitive to the value of the baseline frequency, either. The set of the optimized evaporation parameters is shown in Table 7.1, and $f(t)$ for stages 2-8 is shown in Figure 7.4.

The behavior of the relevant collisions during the initial part of the evaporative cooling mean that the evaporation performance is not very sensitive to the value of the magnetic field at which the evaporation occurs. When the atoms are initially loaded into the magnetic trap, the loss rate of atoms from the cloud is not dominated by two- or three-body inelastic collisions but rather by the background loss rate τ . This loss rate is independent of field, and at the relatively high initial temperatures of the sample, the elastic scattering rate will be unitarity limited (subsection 5.2.4)

Stage	f_{start} (MHz)	f_{stop} (MHz)	f_{base} (MHz)	t_{stage} (seconds)	Rf amplitude (dBm before amplifier)
1	139.08	131.23	127.08	T1=35	-12
2	82.6	81.6	80.6	T2=10	-12
3	81.6	81.55	81.28	T3=8	-17
4	81.55	81.43	81.08	T4=8.2	-17
5	81.43	81.34	81.21	T5=4.3	-17
6	81.34	81.29	81.235	T6=5	-25
7	81.29	81.27	81.255	T7=2	-30
8	81.27	81.225	81.222	T8=15	-35

Table 7.1. Summary of evaporation parameters used in our optimum evaporation trajectory. Stage 1 occurred at 251 G while the remaining stages are at 162 G.




Figure 7.4. Log plot of the frequency of the rf-knife as a function of time during the evaporative cooling at 162 G. Only stages 2-8 are shown. 81.209 MHz corresponds to the frequency that drives transitions at the bottom of the trap potential.

except near the $a=0$ minimum of the Feshbach resonance. Therefore, the evaporation performance should initially be similar at a wide range of magnetic fields. The initial magnetic field that is used in our optimum trajectory is 251 G. This field was chosen for several reasons. First, the s-wave scattering length at that field is relatively large, approximately $-340 a_0$, so as the cloud cools the elastic scattering rate will be large. Secondly, the final evaporation must occur at a field with positive scattering length, and from 251 G it is not necessary to ramp the bias field through the high-loss Feshbach resonance peak to get to the positive scattering length region. Also, the Feshbach resonance does exhibit energy dependence and 251 G is far enough away from the resonance that it will not experience significantly enhanced inelastic rates. Finally, the frequency needed for the rf-knife couples well through the rf coil.

The initial evaporative cooling is slow because of the relatively weak trap used in the initial evaporation at 251 G (radial and axial trap frequencies approximately 15 Hz by 7 Hz) and correspondingly low density and temperature of the cloud of atoms plus the one-dimensional nature of the evaporation. For optimum evaporation efficiency, the time for the first stage of evaporation (T1) was over 120 seconds. We rarely used T1 times this long out of a consideration for a reasonable duty cycle and so the initial evaporation stage was usually not optimum. A typical time that was used for T1 was 35 seconds, although occasionally times as short as 20 seconds were used when the initial magnetic trap loading from the MCMO sequence (section 4.7) was very good. When the loading after the MCMO sequence was not as good times as long as 65 seconds were used. The reduction in the initial stage evaporation time resulted in a decrease of 10% in the number of atoms remaining at the end of stage 1.

As the evaporation progresses, the density of the cloud increases as it cools and eventually the three-body inelastic collisions at 251 G begin to overwhelm the elastic collisions and the evaporation starts to flag. If the evaporation is continued at this field, the elastic collision rate begins to decrease as the cloud cools and the loss in evaporation efficiency prevents a BEC from being formed. Therefore, after the cloud has been cooled to $\sim 2 \mu\text{K}$ (from $45 \mu\text{K}$), the field is changed to 162.5 G (scattering length $\sim 170 a_0$). This improves the elastic/inelastic collision ratio and the evaporation efficiency picks up again. Also, the dominating inelastic collisions are two-body rather

than three-body at 162.5 G, which is in general more favorable for evaporative cooling. The elastic collision rate again begins to improve as the cloud is cooled, at least for a time. At ~ 300 nK, the evaporation trajectory peaks and then turns over. Even though the elastic collision decreases as the cloud cools further, it is still possible to form a BEC by continuing the evaporation to very low (~ 10 nK) temperatures. We have indications that a condensate exists at higher temperatures than the ~ 10 nK of the final stage, but the fraction of atoms in the condensate at higher temperatures is small. The indication that there is a condensate at these higher temperatures is the observation of an explosion of atoms formed by ramping the bias magnetic field to large scattering lengths (see Chapter X).

Since the elastic collision rate peaks at 300 nK at 162 G, we tried changing the field at that point to improve the subsequent evaporation. Both increasing the size of the scattering length by a factor of two (ramping to 160 G) and decreasing it by a factor of two (ramping to 164 G) failed to improve the evaporation, however, as described below.

At the point where the evaporation trajectory turns over at 162 G, the elastic collision rate is greater than the radial oscillation frequency (elastic rate 22 Hz vs. 17 Hz radial trapping frequency), and so decreasing the elastic collision rate would alleviate any problems associated with the thickness of the cloud. However, at 164 G the elastic/inelastic collision ratio is worse than at 162 G since the elastic scattering rate falls faster than the inelastic collision rate. It was still possible to create BECs, but the number of atoms in the BEC was smaller and the evaporation took more time (~ 20 seconds) because of the slower elastic scattering rate.

Ramping the field to 160 G increased the elastic and inelastic collision rates in a way that improved the elastic/inelastic collision ratio, but the large scattering length caused the elastic collision rate to greatly exceed the radial trap frequency. This meant that the evaporation efficiency suffered from the difficulties associated with thick clouds of atoms and the increase in elastic scattering rate could not be taken advantage of. The performance at 160 G was worse than the other two fields, and it was not possible to create a BEC.

From these studies, we concluded that there was no way to significantly improve upon the evaporation performance at 162 G by changing the bias magnetic field in the trap and using the

Feshbach resonance. The evaporation trajectories for the three different fields measured is shown in Figure 7.5.

7.1.6 Comparison with Simple Model. In addition to the measured evaporation trajectory shown in Figure 7.3, the predicted evaporation trajectory using the simple model described in subsection 7.1.3 is also shown. The model certainly does not match the data in a quantitative fashion, but the general trend of the observed evaporation trajectory and the model is the same. The inelastic loss rate is simply too high for runaway evaporation to occur. The evaporation is still good enough to increase the phase-space density, however, and so the BEC transition can still be reached.

The simple model also reproduces some other effects as well. The evaporation performance is relatively insensitive to changes in η and κ , which is consistent with the observation that the rf-knife parameters have to be adjusted by relatively large amounts to significantly affect the evaporation performance. Also, if a and K_2 are adjusted to match the conditions at 164 G, the model predicts longer cooling times and less efficient evaporation. If a and K_2 are adjusted to match the conditions at 160 G, however, the model does not predict the large loss in evaporation efficiency observed at that field. Most likely, the effects having to do with an elastic collision rate faster than the radial trap frequency are not taken into account sufficiently in the simple model, such as the deposition of energy in the cloud by inelastic collision products that scatter on the way out of the atomically thick cloud [133]. Finally, the model also predicts that the evaporation at 251 should also perform poorly at low temperatures as was observed.

Since the simple model matches most of the qualitative features of the observed evaporation trajectory, there are no indications that any other exotic effects are responsible for the poor performance in ^{85}Rb as compared to ^{23}Na or ^{87}Rb .

The model parameters can be varied in order to see how much the evaporation trajectory could be improved for changes in the trap frequencies and magnetic field. Already, there is no indication that any change in the field alone can significantly improve the evaporation as discussed above. Simply halving the trap frequencies at 162 G reduces the number of atoms remaining at the BEC transition by a factor of 10. A 10% reduction in all of the trap frequencies leads to a 14% reduction of the number of atoms remaining near the BEC transition temperature. Doubling all of




Figure 7.5. Comparison of Evaporation Trajectories at Different Magnetic Fields. At the peak of the 162 G evaporation trajectory, the bias field was ramped to a different value and the evaporation was then optimized at that new field. The 162 G curve was taken at a different time than the 162.5 G curve shown in Figure 7.3.

the trap frequencies results in a 30% increase in the number of atoms remaining at the transition temperature in the model, although this result is probably an overestimate of the improvement since the drop in evaporation efficiency due to atomic thickness effects is underestimated by the model. Due to the inaccurate treatment of the atomic thickness effects, it is unclear whether an increase in the trap frequencies would improve the actual evaporation trajectory at 162 G.

When the field and trap frequencies are both varied simultaneously in the model, there is no significant improvement that can be found over 162 G and $2\pi \times 17.3$ by $2\pi \times 6.8$ Hz. A reduction in the trap frequencies is found not to be compensated by an increase in the scattering length since the inelastic losses increase as well. Increasing the trap frequencies gives an increase in performance in the model at 162 G as described above. With the doubled trap frequencies, neither increasing nor decreasing the field from the 162 G improves the evaporation performance significantly, however. Therefore, there is no obvious improvement in the evaporation that can be found in the simple model, especially considering that the reduction in evaporative cooling efficiency due to the atomic thickness of the cloud appears to be underestimated in the model.

7.1.7 Achieving Reproducible Evaporation Performance. Since our experiments with ^{85}Rb BECs use destructive methods to image the sample, it is necessary to produce BECs that have the same number of atoms and temperature both from one shot to the next and over longer periods of time. When the variation in the measured condensate number was greater than the variation expected from the uncertainty in the absorption imaging ($\sim 5\%$, see subsection 4.12.3), the primary source of shot-to-shot irreproducibility was determined to be fluctuations in the bias field of the magnetic trap. A change in the bias field will shift the energy at which the rf-knife cuts into the cloud during the evaporation. Longer drifts in the evaporation were also caused by drifts in the bias field of the magnetic trap, as well as by changes in the magnetic trap loading due to changes in the MOT-CMOT-Molasses-Optical pumping (MCMO) stages (section 4.7). The variation in the number of atoms in the BEC was much greater than the variation in the number of atoms at the end of any of the previous evaporation stages, suggesting that the evaporation itself was quite uniform (2% or better number reproducibility at around 300 nK) from shot to shot.

In order to compensate for small drifts in the bias magnetic field, there is an offset that is uniformly added to all of the frequencies used in the rf evaporation, called the "rf shift". As the bias field of the magnetic trap changes, the rf shift can be adjusted to maintain the same number of atoms in the BEC. From measuring the number of atoms in the BEC as a function of this offset, we found the slope was ~ 600 atoms/kHz or 300 atoms/mG. In order to achieve 10% reproducibility in number with 5000 atom condensates, a magnetic field stability of 2 mG is then required. This means that the DC stability of the currents in the baseball and bias coils must be better than 6 ppm (since the baseball coil produces a ~ 300 G field and the bias a ~ 150 G field).

Not only is it important that the field produced by the current through the trap coils be stable, it is also necessary for the atoms not to slosh in the vertical direction in the magnetic trap. If the center-of-mass of the cloud of atoms moves in the vertical direction, then the rf-knife will cut more or less deeply into the cloud. If this sloshing motion varies from shot-to-shot, then that will lead to a variation in the number of atoms in the condensate. Normally, any initial sloshing in the vertical direction is removed during the evaporative cooling since high-energy atoms are removed irrespective of whether their energy comes from being in the high-energy tail of a Maxwell-Boltzmann energy distribution or a center-of-mass motion. So there should ideally be no number variation resulting from this effect. However, if the sloshing is imparted during the evaporative cooling, then it will be impossible to remove all of the sloshing motion.

External vibrations near the radial trap frequency (17.3 Hz) can excite the vertical sloshing motion of atoms in the cloud, and we found that these vibrations were a significant source of shot-to-shot variability in the number of atoms in the condensate. We initially noticed this through the discovery of a correlation between the variability in the position of the center of mass of the cloud and the shot-to-shot variation of the number of atoms in the condensate.

Once this effect was observed, we proceeded to measure and try to eliminate all the sources of vibrations at this frequency. A laser beam was directed in such a way that it was clipped by the baseball coil so that the vibrations of the coil were turned into intensity variations of this laser beam. We then tapped on and pulled the various optics and cables on the table to determine which excited vibrations in the baseball coil. The camera on its long post was found to be a significant source of

17 Hz vibrations and so a counterweight was added to change the camera post's vibration resonances. Also, the trap coils were reeoped to improve their mechanical stability and all the wires leading to the coils were firmly attached to the optical table as close to the coils as possible. This reduced the observed variation in the position of the center of mass of the cloud from as much as 10 μm rms to a typical 1.5 μm rms and the variation in the number of atoms in the condensate from as high as 25% or more to 10% or less.

While it was possible to compensate for the longer time-scale drift in the number of atoms in the condensate through the use of the rf shift, a drift complicates the data analysis and operation of the experiment. It also inevitably results in a loss of duty cycle as some data points need to be taken to monitor the drift. Therefore, we tried to eliminate the drift. There were several sources of this drift that were identified. One source was the variation in stray magnetic fields in the trap region. Moving ferromagnetic materials such as the transfer tube guide magnets away from the magnetic trap reduced the drift. A change in the temperature of the magnetic trap coils can also change the bias field in the magnetic trap though changing the spatial dimensions of the coils. The water cooling helps minimize any temperature changes, but we also found that it was helpful to use a small fan to cool the science MOT anti-Helmholtz coils to prevent them from getting hot and in turn heating the base and bias coils or their support structures.

Another source of a drift in the number of atoms in a condensate at the end of the evaporation sequence is a change in the initial loading of the magnetic trap from some change in the MCMO stages (MOT-CMOT-Molasses-Optical pumping see section 4.7), such as a decrease in the optical pumping efficiency. Since the rf shift is used to maintain the number of atoms in the condensate, a reduction in the evaporation efficiency (the usual direction of the drift) causes us to adjust the final rf frequency away from the bottom of the trap and the drift in initial loading appears to be a change in the bias field in the magnetic trap. However, as this happens the temperature of the condensates with a fixed number of atoms begins to increase. For instance, if originally there is no detectable thermal component in the condensate of a certain size, as the initial loading drifts a thermal cloud will appear around the condensate even as the rf shift is adjusted. This drift in initial loading was reduced in several ways. First, the MCMO conditions are optimized as well as possible.

When these conditions are well optimized, a small change in the parameters does not significantly decrease the loading. The science MOT does drift slowly over months, and so it is necessary to periodically improve the alignment of the MOT, clean the optics, etc. Since we do not usually run with the optimal time for the first evaporation segment, it is possible to increase that time to reduce the sensitivity of the evaporation to the initial conditions, at the loss of duty cycle. Finally, the various laser detunings (such as the laser detuning during the optical pumping) and magnetic trap loading efficiency are checked periodically over the course of the day and any changes are corrected.

7.1.7.1 Use of Inelastic Losses to Reduce Number Variation. Since the two- and three-body loss rates are density-dependent, a condensate with a larger number of atoms than average will tend to lose them faster than the average rate, and a condensate with a smaller number will lose atoms at a slower rate. Therefore, the inelastic collisions tend to smooth out number variations. These inelastic losses are used as an additional way to reduce the fluctuation of the number of atoms in the condensate.

The first way this is done is to lengthen the final stage used in the evaporation sequence. While the time for the seventh stage is 2 seconds, the time for the last stage is 15 seconds. The inelastic losses during this time will smooth out the condensate number. In addition, if the cloud is for any reason a bit hotter at the end of stage 7, it will evaporate more efficiently in the slowed stage 8 and that will also reduce variability. We have observed that this longer time allows the creation of colder condensates as compared to a shorter stage 8 as well.

In addition to a slower final stage of evaporation, the condensate is often ramped near $a=0$ and allowed to evolve there from 2-10 seconds. The reduction in mean-field energy causes the density of the condensate to increase since its spatial size is reduced. The inelastic loss rates therefore increase and even for the relatively small-number condensates remaining after stage 8 the number of atoms can be reduced significantly in this time. No increase in the number of thermal atoms due to the inelastic losses is observed, although the apparent temperature of the condensate does increase because any thermal atoms that remained at the end of the evaporation are at a lower density and are not lost as fast as the condensate atoms. To remove some of these thermal atoms, a weak rf-knife is used at a fixed frequency during the evolution at $a=0$. This rf-knife does not cause

any true evaporation and so it is quite inefficient at cooling the cloud further. It does help maintain the temperature of the cloud by selectively removing some thermal atoms as atoms are lost from the condensate.

7.1.7.2 Reproducibility Performance. With all of the improvements made to the system described above, the measured reproducibility of the number of atoms in large condensates ($\sim 15,000$) is 7%, slightly larger than the apparent variation expected from noise in the absorption imaging, 5% (subsection 4.12.4). This indicates the variation in the actual number of atoms in the condensate is comparable to the variation due to the noise in the imaging. With the institution of the extra stage at $a=0$ using inelastic losses to smooth the number variations, similar performance is observed for condensates down to ~ 2000 atoms. Without using the extra inelastic loss stage, the fractional variation in the number of atoms in the condensate increases as the total number decreases, as would be expected from a fixed amount of variation in the magnetic trap bias field.

7.2 Loss of Atoms From the Condensate When Ramping through the Feshbach Peak.

7.2.1 Absorption Imaging and Ramping through the Feshbach Peak. In order to perform absorption imaging with a high signal-to-noise ratio and be able to interpret the results quantitatively, it is necessary to turn off the magnetic fields that form the magnetic trap. If the atoms were imaged while still trapped, there would be two difficulties associated with the relatively large bias field (~ 160 G). First, because the trap field is oriented perpendicular to the direction of the probe laser and the optics in the imaging system are polarization-sensitive, the amount of light absorbed by the atoms is difficult to interpret (see section 6.4). Second, there is no cycling transition at ~ 160 G so only a few photons can be scattered from an atom before it decays into a state that is out of resonance with the probe laser light, as opposed to the ~ 100 photons that are scattered in the usual absorption imaging at $B \sim 0$.

There is a difficulty associated with turning off the magnetic trap fields, however. Since the condensates are formed at fields higher than the peak of the Feshbach resonance, as the magnetic field is turned off the value of the field will pass through that of the Feshbach resonance. The s -wave scattering length diverges at the peak and since the mean-field energy of the condensates is

proportional to the scattering length, it is possible that the condensate will receive an impulse that causes it to expand outward. Also, the inelastic collision rates also increase markedly near the Feshbach peak (Chapter VI), and so a loss of atoms due to these large inelastic rates also occurs during the ramp. Finally, large loss rates were observed in ^{23}Na BECs near Feshbach resonances [90], and exotic loss mechanisms [94-97] may then also cause atoms to be lost from an ^{85}Rb condensate as the magnetic field is ramped through the Feshbach peak. The distortion of the spatial size and loss of atoms from the condensate are both discussed in the following two subsections.

7.2.2 Outward Impulse from the Magnetic Trap Turn-off. In fact, the atoms in the condensate are observed to receive an outward impulse from the magnetic trap turn-off, and so the measured spatial size is not the same as the spatial size of the atoms when they were actually in the magnetic trap. If the atoms could be imaged immediately after the trap is turned off, then this would not be a concern since the condensate would not have any time in which to expand. However, due to the presence of a magnetic field that persists after the coils are turned off (subsection 4.11.4), we must wait for about 1.6 ms before imaging in order to remove errors in the determination of the number of atoms due to this persistent field. While the Feshbach resonance peak outward impulse does distort the spatial size of the condensates from what they were in the magnetic trap, it is useful in that it stabilizes condensates against collapsing due to the large negative scattering length at $B \sim 0$ (this effect is discussed in more detail subsection 10.3.7).

The fact that the impulse is outward is reasonable given the proportionality of the mean-field energy to the density of the condensate. As the magnetic field is tuned through the Feshbach resonance peak from higher field, the s-wave scattering length diverges first to $+\infty$, then changes discontinuously to $-\infty$, and then increases (decreases in magnitude) towards its value at $B \sim 0$ (see Figure 5.1). Since a has both positive and negative values, it is not immediately clear whether the atoms in the condensate will receive an inward (contracting) or outward (expanding) impulse from the trap turn-off. However, since the scattering length diverges first toward $+\infty$, the atoms in the condensate will initially be accelerated outward and the density of the condensate will be reduced. The following time that the condensate spends at $a < 0$ will then be at a lower density, and so the resulting inward acceleration will be less than the outward acceleration imparted.

Figure 7.6 contains data that show the presence of this outward impulse. As expected, condensates with higher initial density experience more of an impulse from the trap turn-off than those with lower initial density. This outward impulse did not normally distort the data of interest to us, and so this effect has not been characterized any better than shown in Figure 7.6. If this effect becomes important in interpreting any future measurements, a faster field turnoff or a more careful measurement of condensate expansion as a function of initial condensate density and initial scattering length will be necessary.

7.2.3 Loss of Condensate Atoms while Ramping through the Feshbach Peak. The loss of atoms caused by ramping through the Feshbach peak is a greater concern than the impulse imparted to the atoms. This is for two reasons. First, the results with ^{23}Na condensates in Ref. [90] sparked theoretical interest in possible exotic coherent loss mechanisms associated with condensates and the Feshbach resonance (see subsection 7.2.4 below). Second, any loss of atoms during the turn-off of the magnetic trap would mean that the atoms measured in the absorption image would not be the same as the number of atoms originally in the magnetic trap. Since several of our measurements required the knowledge of the number of atoms in the magnetic trap, this loss had to be either eliminated or accounted for in order to accurately determine the number of atoms in the condensate.

The more time the condensate spends at a field near the Feshbach peak, the larger the loss of atoms from the condensate. Figure 7.7 shows the fractional loss as a function of the speed at which the magnetic field is ramped. This data was acquired by increasing the current in the baseball coil with a linear ramp to sweep the magnetic field experienced by the atoms from 162.0 G to ~ 132 G and then immediately turning off the trap and imaging the cloud using absorption imaging. As expected, the loss is greater for slower ramps, although the increase in loss with ramp speed is quite slow. This is reasonable for a loss that is density-dependent, since for slower ramps the condensates have some time to expand due to the mean-field energy, resulting in a decrease in the density. The fraction of atoms lost for the fastest ramp was determined to be less than 9%. This limit was measured in a separate experiment where the atoms were imaged in the trap and the number of atoms before and after the ramp through the Feshbach peak were compared. Although the in trap imaging was difficult to interpret, it was possible to use it for a relative measurement. Included for

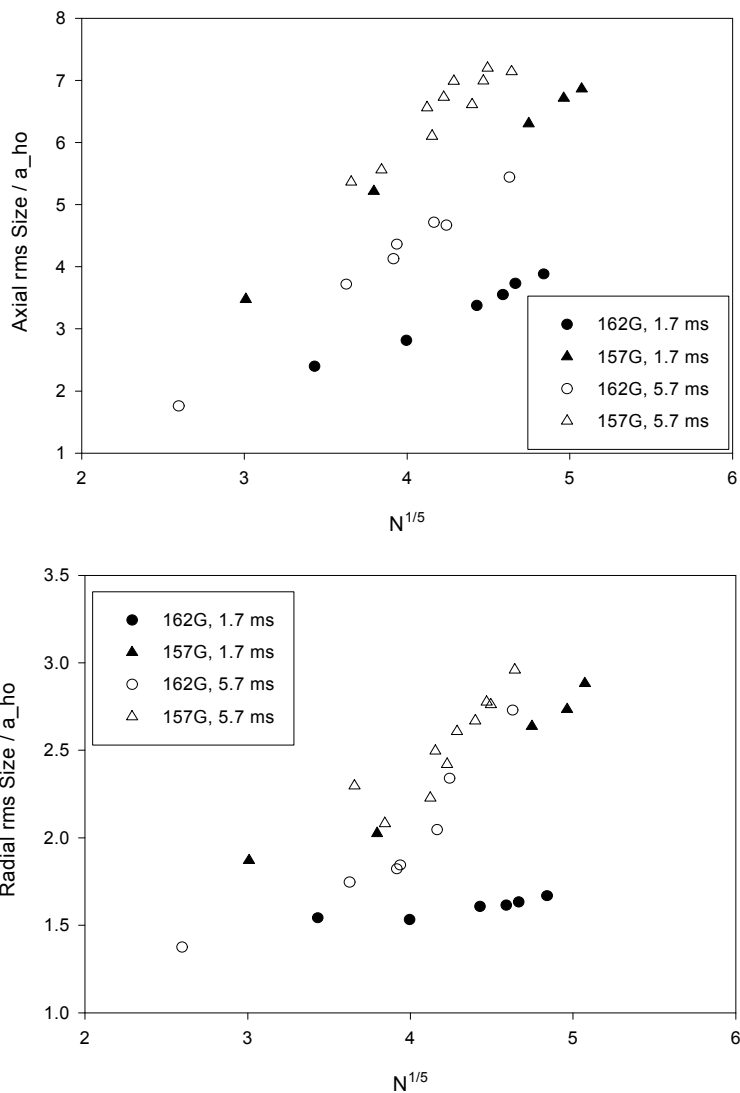


Figure 7.6. Evidence for an impulse imparted during trap turn-off. The filled triangles and circles represent the measured spatial size for condensates originally at 157 and 162 G respectively imaged 1.7 ms after the magnetic trap was turned off. The open symbols represent the measured widths when the condensates were imaged 5.7 ms after the trap turn-off. The observed size is larger for the longer expansion time data than for the shorter expansion time data, indicating that the condensates are expanding outward outside the trap. The relative size increase is larger for the condensates that are originally smaller in spatial extent and are therefore denser. These data were taken soon after ^{85}Rb BECs were formed, before we were able to improve the reproducibility of the system or create condensates with more than a few thousand atoms, and so the number varied by a factor of three during the measurement.




Figure 7.7 Fractional of atoms lost following a rapid sweep of the magnetic field through the peak of the Feshbach resonance as a function of the inverse speed of the ramp. Data are shown for a condensate (●) with a peak density of $n_{\text{pk}} \sim 3 \times 10^{12} \text{ cm}^{-3}$ and for a thermal cloud (○) with a temperature $T = 430 \text{ nK}$ and a peak density of $n_{\text{pk}} = 4.5 \times 10^{11} \text{ cm}^{-3}$.

comparison in Figure 7.7 is a plot showing the fractional loss for a cloud of thermal atoms at a temperature significantly hotter than the BEC transition temperature.

The density of the condensates quoted in Ref. [16] is likely not correct. The expansion of the spatial size due to the impulse imparted during the trap turnoff was not taken into account, and the clouds were small enough that the resolution limit could have been distorting the measurement of the radial direction. A simple estimate of the density using the measured number and the scattering length at 162 G would predict a peak density three times as large as what was quoted in Ref. [16], and that is the density listed in Figure 7.7.

As illustrated by Figure 7.8, the same experiment was repeated with condensates that had a much lower density. This lower density was achieved by changing the mean-field interaction energy to create a collective excitation that expanded the spatial size of the condensate (see section 8.3), and so this lower density was transitory rather than an equilibrium property of the condensate. With the reduction in density, the fraction of the atoms lost was reduced as well. For instance, while at 10 $\mu\text{s}/\text{G}$ the higher density case showed a 25% loss, the low-density case is consistent with no loss within the 2% uncertainty in the measurement. The small amount of loss observed in the low-density case implies that for the fastest trap turn-off speeds, the number measured by the absorption imaging is very close (within a few percent or less) to the number of condensate atoms originally in the trap.

7.2.4 Comparison of our Observed Loss Rate to Coherent Loss Mechanisms. After the observation of the large and unexpected loss rates in ^{23}Na condensates near a Feshbach resonance [90], several proposals for coherent loss mechanisms were put forward to explain the loss [94-97]. All of these mechanisms were based on a formalism originally developed by E. Timmermans and coworkers [94]. In this theory, changing the magnetic field in the vicinity of a Feshbach resonance does not simply change the scattering length between two atoms. Instead, the field operator describing the atoms in many-body theory ($\hat{\Psi}(\mathbf{x})$ in section 3.3.2) is split into two parts, one that describes the atoms in the condensates and one that describes molecules in the bound state that it is resonantly tuned as part of the Feshbach resonance (see Chapter II). The time evolution of the atoms and molecular fields is then calculated in a mean-field treatment. It is found that the atom and




Figure 7.8. Fraction of atoms lost from a condensate following a rapid sweep over the Feshbach peak as a function of inverse ramp speed. The data shown here are similar to those displayed in Figure 7.7, only the peak density of these condensates is $n_{pk} \sim 0.2 \times 10^{12} \text{ cm}^{-3}$.

molecular fields are coupled to each other, and this coupling results in a change in the effective atom-atom scattering length as the energy of the molecular bound state is tuned.

There are then two types of coherent loss mechanisms that are predicted to occur as the magnetic field is swept through the Feshbach resonance. In two cases [94,96], the rapid change in magnetic field results in the conversion of a significant number of atoms to molecules via a Josephson-like oscillation between the atom and molecular states. These molecules can then collide with atoms in the condensate that in turn results in loss from the magnetic trap. It is important to note that the collision rate for these loss-inducing collisions was not determined *a priori*, but rather was selected to fit the ^{23}Na data. In another treatment [95], the sweep through the Feshbach resonance both creates molecules and tunes the bound state so that it becomes a quasi-bound state. This quasi-bound state decays, and the atoms resulting from this decay gain kinetic energy and leave the magnetic trap.

There are several difficulties associated with these theory treatments. One of these difficulties is associated with the fact that it is possible in the theoretical treatment for the whole condensate to be converted from atoms into molecules without any consideration for the fact that two atoms must be close to each other in order to form a molecule. Therefore, there should be a time scale in the theory that reflects the time that the atoms take to move so that they are close enough to another atom to form a molecule. This problem is addressed in a so-called pairing theory developed by M. Holland and coworkers [97]. Pairing theory includes this time for atoms to move close to one another by considering the evolution of not only the mean-field atom and molecular fields, but also beyond mean-field components such as the normal and anomalous densities. The inclusion of these additional terms complicates the numerical evaluation of the theory predictions for the case of ^{85}Rb , and so there has been no calculation of the expected losses as the condensate is swept through the Feshbach resonance.

Another difficulty with the treatments in [94-96] is that they are all mean-field treatments. Near the Feshbach resonance, the mean-field treatment is expected to break down, casting doubt on the validity of the predicted loss rates near the Feshbach peak, which is where most of the loss

occurs. This violation of the mean-field theory is more strongly pronounced in ^{85}Rb due to the relatively low value of the condensate densities. The characteristic energy of the atom-molecule coupling term is $\alpha n^{1/2}$. In order to expect the prediction of the mean-field theories to be valid, it is reasonable to require that the quantity $na^3 \ll 1$ at a detuning from the peak, ϵ , that is large compared to $\alpha n^{1/2}$. For the data in Figure 7.7, when ϵ is equal to $10\alpha n^{1/2}$, $na^3 \sim 9$.

Finally, there is no way to separate any loss due to coherent processes from the ordinary dipole relaxation and three-body recombination inelastic collision rates that are also enhanced near the Feshbach resonance peak (see Chapter VI). The strong and poorly characterized temperature dependence of these known loss rates near the Feshbach peak makes it difficult to determine what fraction of the observed condensate loss can be attributed to the usual inelastic loss processes. We cannot therefore confirm or rule out the existence of any coherent loss processes induced by the ramp of the magnetic field through the peak of the Feshbach resonance. Again, even if we could determine some part of the loss to be coherent, the predictions of Refs. [94-96] are not expected to apply to our ^{85}Rb condensates for the reasons discussed in the above paragraphs.

7.3 ^{85}Rb Condensates with Positive Scattering Lengths.

7.3.1 Stability of the Condensates with Positive Scattering Length. As part of our initial studies of ^{85}Rb condensates, we examined the properties of the condensates at positive scattering lengths. The fact that we were able to create stable ^{85}Rb condensates with up to 18,000 atoms with lifetimes consistent with those expected from the two- and three-body inelastic collision rates was already a consequence of being able to tune the scattering length to various positive values. Away from the Feshbach resonance, the maximum number of atoms in a stable condensate is predicted to be ~ 90 [135] due the large negative scattering length of $\sim -400 a_0$ (see Chapter IX for more details on condensate stability at negative scattering length). Without the ability to operate in a region of positive scattering length, we would therefore never expect to observe stable condensates of much greater than ~ 90 atoms.

7.3.2 Spatial Size of the Condensates as a Function of Magnetic Field. When the sign of the scattering length is positive, the mean-field energy is also a positive quantity that is proportional to the magnitude of a and the density of the atoms in the condensate. Since the condensate can

reduce its mean-field energy by reducing its density, the interaction between atoms with a positive scattering length acts as repulsive force, tending to expand the spatial size of the condensate. As the value of a is increased, the spatial extent of the condensate in its ground state is therefore predicted to increase as well.

In order to study this variation of spatial size with field, a condensate was created at 162 G and then the bias magnetic field in the trap was linearly ramped to a new value to change the value of a . Variations in the value of a that spanned three orders of magnitude were studied. The duration of this ramp was sufficiently long (500 ms) to ensure that the condensate responded adiabatically (see Chapter VIII for more details). Figure 7.9 shows a series of condensate images for various magnetic fields. Moving towards the Feshbach peak from 162 G, the condensate size increases due to the increased strength of the mean-field interaction. The density distribution of the condensate approaches the distribution expected in the Thomas-Fermi limit (see subsection 3.3.4) with the aspect ratio given by the trap frequencies as ω_r/ω_z . Moving in the opposite direction from 162 G, the cloud size becomes smaller than the 7 μm resolution limit shortly before we reach the non-interacting limit where the condensate density distribution is a Gaussian whose dimensions are set by the harmonic oscillator lengths ($l_i = (\hbar/m\omega_i)^{1/2}$ where $i=r,z$). The aspect ratio for these non-interacting condensates is $(\omega_r/\omega_z)^{1/2}$. We observed the aspect ratio decrease as the mean-field interaction energy was reduced, although the resolution limit obscured the true aspect ratio as the radial extent of the condensate became small. Also, we took condensate images similar to those in Figure 7.9 at many field values between 156 and 166 G in order to see if the measured width was scaling with magnetic field in the manner that was expected. Before discussing the result of these measurements, the density distribution of condensates in the Thomas-Fermi limit will be outlined in the following paragraphs.

The ratio of the mean-field energy to the kinetic energy is given by the dimensionless parameter Na/a_{ho} only in the limit where the mean-field interaction energy is small compared to the kinetic energy. As the mean-field energy becomes appreciable, Na/a_{ho} begins to underestimate the ratio of the mean-field to kinetic energy. At $Na/a_{\text{ho}} \sim 0.3$ the two energies are comparable and at $Na/a_{\text{ho}} = 1$, the ratio of the total mean-field energy to total kinetic energy is equal ~ 5 . Therefore, it is




Figure 7.9 (color). False color images and horizontal cross sections of the condensate column density for various magnetic fields. The condensate number was varied to maintain a peak optical depth (OD) of ~ 1.5 . The magnetic field values are (a) $B = 164.9$ G, (b) $B = 162.0$ G, (c) $B = 158.1$ G, (d) $B = 156.9$ G, and (3) $B = 156.1$ G.

not necessary that Na/a_{ho} be much greater than one for the Thomas-Fermi limit to apply, only that Na/a_{ho} be somewhat larger than one.

One advantage of the Thomas-Fermi limit is that the size of the condensate can be calculated analytically. In the Thomas-Fermi limit, the condensate density has the shape of a parabola

$$\begin{aligned}
 n(\vec{x}) &= k \left(1 - \frac{(x^2 + y^2)}{R_r^2} - \frac{z^2}{R_z^2} \right) \\
 k &= m \frac{m}{4p\hbar^2 a}, \quad R_{r,z}^2 = \frac{2m}{m\omega_{r,z}^2} \\
 m &= \frac{\hbar\omega_{ho}}{2} \left(\frac{15Na}{a_{ho}} \right)^{\frac{2}{5}}
 \end{aligned} \tag{7.7}$$

where n is defined only where it is positive (i.e. the density is zero where the equation for n would be negative), μ is the chemical potential defined in equation (3.23), and $\omega_{r,z}$ is the radial/axial trap frequency of the magnetic trap. From equation (7.7), it can be seen that the spatial extent of the condensate is proportional to R_r in the radial direction and R_z in the axial, and scales as the product of the number and the scattering length (Na) to the 1/5 power. Integrating through the y direction gives the column density of the condensate, to which the observed optical depth is proportional

$$\text{Column density} = \frac{4}{3} k R_r \left(1 - \frac{x^2}{R_r^2} - \frac{z^2}{R_z^2} \right)^{\frac{3}{2}} \tag{7.8}$$

Figure 7.10 compares the density and column density profiles in one direction. From equation (7.8) and equation (7.7), the relationship between the FWHM of the condensate column density profile in the z direction (along the line where $x=0$) and the product Na is then

$$\text{FWHM}_z = 1.22 \sqrt{\frac{\hbar\omega_{ho}}{m\omega_r^2}} \left(\frac{15}{a_{ho}} \right)^{\frac{1}{5}} (Na)^{\frac{1}{5}} \tag{7.9}$$

Figure 7.11 shows the measured axial FWHM versus $(Na)^{1/5}$ of condensates with an average number of ~ 5000 atoms at fields ranging from 156 G to 166 G. Figure 7.12 shows a subset of this data plotted against $N^{1/5}$ and $a^{1/5}$ separately. The value of the scattering length at a particular




Figure 7.10. Comparison of the shape of the condensate density and column density in the Thomas-Fermi limit along one spatial direction.




Figure 7.11. Measured condensate width versus the product of the scattering length and the number of atoms in the condensate, with that product raised to the 1/5 power. From equation (7.9), the axial FWHM of a condensate is expected to scale as $Na^{1/5}$, and indeed the data do fall on a straight line. The slope of the line matches the prefactor in equation (7.9) to the precision of the measurement. The deviation from the fit line at small values of Na is discussed in the main text.




Figure 7.12. This figure shows a randomly selected subset of the data shown in Figure 7.11, only here the dependence of the width on the number and scattering length are plotted separately. It is difficult to make out details on this graph, but is possible to observe the spatial size increasing with N at a fixed a and the spatial size increasing with a at a fixed N .

magnetic field was not taken directly from equation (5.1) using the central values of the parameters determined in the study of the elastic collision rate (Chapter V). Instead, the locations of B_{peak} and B_{zero} and the value of a_{bg} were allowed to vary within their uncertainties in order to best match the condensate widths as a function of Na . The error on the measurement of the condensate widths was estimated by the average deviation of the measured widths from their fit values, assuming that all of the deviations were due purely to statistical errors. This last assumption is suspect due to the reasons that are presented below, preventing this measurement of the condensate widths from being used to improve the measured location of B_{peak} . Instead, this measurement demonstrates that the condensate widths have the general scaling expected from mean-field theory, i.e. the width of the condensates increases with increasing product of the scattering length and number of atoms in the condensate to a relatively weak power.

The reason that the FWHM of the condensates was used for this measurement rather than fitting the condensates to the expected density profile given in equation (7.8) is that the FWHM is less sensitive to several systematic effects present in the imaging of the condensates. The finite resolution of the optics in the imaging system causes both diffraction rings around the edge of the condensate and rounds the sharp edges of the Thomas-Fermi density distribution, distorting the observed density profile. Since the edge of the condensate occupies a relatively large area in the image, these distortions can cause significant problems in the fitting. Also, it was assumed in the fitting that no thermal atoms were present, and so if there were thermal atoms in addition to the condensate atoms, the measured density profile would have been distorted. At large values of the scattering length, heating can occur during the ramp from 162 G, and any thermal atoms present will be pushed out of the condensate due to the mean-field interaction energy, since the second-order coherence function results in a factor of two large mean-field energy for thermal atoms as compared to condensate atoms [100]. While the FWHM is not insensitive to the presence of thermal atoms, it is less sensitive than performing a fit to the entire density profile.

For small values of Na , the observed widths are greater than the widths expected from equation (7.9). There are several effects that explain this deviation. The width of the condensate

will not go to zero as Na goes to zero, but will rather approach the size set by the simple harmonic oscillator length. The FWHM of the condensate in the axial direction in the absence of any mean-field interaction is $7 \mu\text{m}$. The resolution limit of the imaging system may also be contributing to the observation of a larger than expected width. Finally, the small size of these condensates means that the density is relatively high and the impulse imparted during the magnetic trap turnoff combined with the 1.6 ms delay time between the trap turn-off and the probe pulse will cause the observed width to be slightly larger than the actual width in the magnetic trap.

Even with these caveats, the condensate width does indeed appear to scale as expected in the region of positive scattering length. This gives us confidence that we understand the condensate ground state in this region of the Feshbach resonance. Condensates in this region of the Feshbach resonance can then serve as the initial condensates for experiments in regions of the Feshbach resonance that are not expected to be as well understood. This includes increasing the positive scattering length to search for a breakdown in the mean-field treatment of the condensate and exploring the stability and dynamics of condensates in the region where the scattering length is negative.

Chapter VIII

Collective Excitations and an Initial Search for Beyond Mean-field Effects

8.1 Overview.

The ability to change the atom-atom interaction strength in ^{85}Rb from very small to very large values provides a unique ability to test the range of validity of mean-field theory in the description of the behavior of Bose-Einstein condensates. While the superfluid properties of liquid helium are attributed to the presence of a Bose-Einstein condensate in the liquid, it is not possible to quantitatively describe the behavior of liquid helium within a mean-field theory because the interactions between helium atoms are too strong. Conversely, Bose-Einstein condensates in dilute trapped atom samples (with fixed s-wave scattering lengths) have been described well by a mean-field treatment [12], at least in the limit of zero temperature. In fact, as I will discuss below, it is prohibitively difficult to create dilute trapped atom condensates with strong enough atom-atom interaction energy (i.e. high enough density) to test the limits of the mean-field description without changing the strength of the atom-atom interactions themselves. The adjustable scattering length of ^{85}Rb can bridge the gap between the dilute and strongly interacting limits, and hence can provide a test of the limits of the applicability of mean-field theory as the strongly interacting limit is approached in a controlled fashion.

The behavior of any system of particles is described in quantum theory with respect to a Hamiltonian, which describes the energy of the system. The energy of a Bose-Einstein condensate in mean-field theory has three components: a kinetic energy, a potential energy, and a mean-field interaction energy proportional to the density of the atoms in the condensate and the value of the s-wave scattering length. As the strength of the atom-atom interaction energy increases, the mean-field term will no longer accurately describe the interaction energy in the condensate (as evidenced by liquid He). In particular, it is predicted that beyond the mean-field description the interaction energy will not be directly proportional to the density of the atoms in the condensate or directly

proportional to the s -wave scattering length [136]. It is also predicted that the first order beyond-mean-field correction to the energy of the condensate will be positive so that the mean-field value of the interaction energy will be lower than the actual interaction energy. None of these predictions concerning the first-order correction to the mean-field energy have been tested in a controlled fashion in condensates.

The deviation of the condensate interaction energy from its mean-field predicted value can be measured by determining the frequency of the lowest energy collective excitations of the condensate. The collective excitations spectrum describes the motion of a condensate that is perturbed from its ground state, and the lowest energy collective excitations describe the motion of the widths of the condensate as it oscillates in response to having its spatial size perturbed from its ground state value. As the condensate gets spatially larger and smaller as it oscillates, the density changes and therefore the interaction energy plays a role in determining the condensate motion and oscillation frequency. The mean-field interaction energy results in a significant shift of the condensate oscillation frequencies from their values in the absence of any atom-atom interactions. Similarly, the presence of corrections to the interaction energy beyond the mean-field value will shift the frequencies of the oscillations as well.

Since it is in general possible to measure frequencies precisely, the detection and quantification of beyond-mean-field corrections by determining condensate oscillation frequencies was very attractive experimentally. The ability to tune the atom-atom interactions in ^{85}Rb condensates also suggested a natural way to observe how an increase in atom-atom interaction strength leads to deviations from the mean-field predicted oscillation frequencies: the oscillation frequencies of condensates could be measured as a function of magnetic field, and therefore s -wave scattering length, over a range of values from where the mean-field was expected to be an accurate description of the condensate motion to where deviations from the mean-field description were observed due to strong atom-atom interactions beginning to limit the validity of the mean-field description. Also, the ability to tune the atom-atom interactions was useful in inducing the oscillations. Rapid changes in the scattering length resulted in the condensate no longer having the ground state spatial extent, and that in turn caused the condensates to oscillate.

After we completed the measurement of the oscillation frequencies of ^{85}Rb condensates at a variety of magnetic fields, we found the measured frequencies scaled roughly (but not precisely) as expected from the mean-field predicted values as the atom-atom interactions were turned up from the non-interacting limit. As the strength of the atom-atom interaction limit was increased further, statistically significant deviations from the mean-field predictions were observed, indicating the presence of beyond mean-field corrections to the atom-atom interaction energy, as expected.

There is a catch, however. At the magnetic fields with strong atom-atom interactions a relatively large heating rate was observed. As the condensate heated up, the oscillation frequency was observed to change as well. These temperature-dependent shifts were the same magnitude but in the opposite direction of the shift due to the beyond-mean-field corrections, and so they obscured the measurement of the magnitude of the beyond-mean-field frequency shifts. Our current measurements therefore demonstrate the presence of beyond-mean-field corrections in the condensate oscillation frequencies but are limited in the precision to which the shifts can be determined. Additional measurements conducted in the future may be able to significantly improve the determination of the magnitude of the beyond-mean-field frequency shifts.

This chapter is organized in the following manner. Section 8.2 describes the collective excitations of condensates in more detail and section 8.3 discusses the predicted size of the leading order beyond-mean-field corrections to the condensate excitation frequencies. Section 8.4 describes how changes in the s-wave scattering length were used to induce oscillations in the condensate. Section 8.5 describes the experimental techniques used to measure the condensate oscillation frequencies. Section 8.6 discusses the initial results of the frequency measurements as a function of magnetic field and the observation of statistically significant beyond-mean-field effects. Section 8.7 and 8.8 detail the difficulties in the measurement of the oscillation frequencies. There were two types of difficulties: for relatively small values of the scattering length the measured oscillation frequencies did not agree with the mean-field predicted frequencies even though the beyond-mean-field corrections were expected to be small; and at magnetic fields with strong interactions the temperature shifts described briefly in the previous paragraph obscured the determination of the magnitude of the beyond-mean-field corrections. Section 8.9 addresses the limits that can be placed

on the magnitude of the beyond-mean-field corrections in spite of the difficulties presented in section 8.8 with the current observations and contains suggestions for how additional data could improve upon the current measurement.

8.2 Introduction to Collective Excitations

8.2.1 Condensate Excitations and the Gross-Pitaevskii Equation. In addition to determining the ground state of a Bose-Einstein Condensate with mean-field interactions, the Gross-Pitaevskii (GP) equation (equation (3.19)) also describes the dynamics of such a condensate. These dynamics can be described through the determination of the excited state spectrum of the GP equation. One way that this can be done is to consider small time-dependent variations of the condensate order parameter around the ground state ($\phi(\mathbf{r})$) and look for solutions of the GP equation of the form [12]

$$\Phi(\vec{r}, t) = e^{\left(\frac{-i\mu t}{\hbar}\right)} \left[\mathbf{f}(\vec{r}) + u(\vec{r})e^{-i\omega t} + v^*(\vec{r})e^{i\omega t} \right] \quad (8.1)$$

where as usual μ refers to the chemical potential of the condensate. By plugging equation (8.1) into the GP equation and keeping terms linear in the functions $u(\mathbf{r})$ and $v(\mathbf{r})$, two coupled equations for $u(\mathbf{r})$ and $v(\mathbf{r})$ are obtained

$$\begin{aligned} \hbar \mathbf{w} u(\vec{r}) &= [H_0 - \mathbf{m} + 2g\mathbf{f}^2(\vec{r})]u(\vec{r}) + g\mathbf{f}^2(\vec{r})v(\vec{r}) \\ -\hbar \mathbf{w} v(\vec{r}) &= [H_0 - \mathbf{m} + 2g\mathbf{f}^2(\vec{r})]v(\vec{r}) + g\mathbf{f}^2(\vec{r})u(\vec{r}) \end{aligned} \quad (8.2)$$

$$H_0 = \frac{-\hbar^2 \nabla^2}{2m} + V_{ext}(\vec{r}), \quad g = \frac{4p\hbar^2 a}{m}$$

These coupled equations can then be solved in order to determine the eigenfrequencies ω and hence the energies $\epsilon = \hbar\omega$ of the excitations of the condensate. These excitations will form a spectrum of states, and by projecting the initial (i.e. $t=0$) condensate order parameter onto this set of states, the time evolution of the condensate can then be determined.

This is not the only procedure that can be used to determine the excitation frequencies of the condensate. It is possible to diagonalize the many-body Hamiltonian given by equation (3.16) in the so-called Bogoliubov approximation [137], in which the fluctuation term \hat{c} is expressed in terms of quasiparticle operators with Bose commutation relations. The resulting equations describing the

excited state eigenfrequencies are equivalent to equation (8.2). Another method that has been employed is to calculate the excited states in the Random-Phase Approximation (RPA) using more traditional atomic physics methods like those used to determine the electronic excited states of an atom (see Chapter III and Ref. [41]). The eigenfrequencies calculated in the RPA again agree with those calculated in the other methods. Finally, it is possible to directly numerically integrate the GP equation in order to determine any condensate dynamics. This can be necessary for extremely large amplitude excitations where the assumption of a small deviation of the order parameter from the ground state no longer holds.

8.2.2 Types of Excitations: Single-particle, Sound Wave, and Collective Excitations. In a uniform gas (homogenous BEC), the excited state amplitudes $u(\mathbf{r})$ and $v(\mathbf{r})$ are plane waves, and equation (8.2) can be used to calculate the dispersion relation between the momentum of an excitation and its energy [138]

$$(\hbar\omega)^2 = \left(\frac{\hbar^2 q^2}{2m} \right) \left(\frac{\hbar^2 q^2}{2m} + 2gn \right) \quad (8.3)$$

where $\hbar q$ is the momentum of the excitation (which is proportional to the inverse of the excitation's wavelength) and as usual n is the density of the condensate. Equation (8.3) has different behaviors for low- and high-momentum excitations. For high-momentum excitations, the dispersion relation is exactly the same as the free-particle case. For low-momentum excitations, a phonon dispersion relation is exhibited such that $\omega=cq$, where c is the sound velocity equal to

$$c = \sqrt{\frac{gn}{m}}. \quad (8.4)$$

In the case of a condensate in a harmonic potential, the finite spatial extent of the condensate introduces an additional length scale that must also be considered in describing the nature of the excitations. Instead of two classes, the excitations are better described by three cases: the case in which the energy of the excitation is higher than the chemical potential, the case in which the energy of the excitation is less than the chemical potential but the wavelength of the excitation is smaller than the spatial extent of the condensate, and the case when the excitation has both low energy and a wavelength comparable to the size of the condensate.

Excitations with energies much greater than the chemical potential will behave approximately like a single particle moving in an effective Hamiltonian that includes the kinetic energy of the single particle, the potential energy due to the external potential, and a term involving the mean-field interaction of the condensate. In this way these high-energy single-particle excitations are analogous to the free-particle excitations at high momentum in the homogenous case --- the finite extent of the condensate and the presence of the external potential merely complicate the single particle motion in the inhomogeneous case as compared to the free-particle-like excitations in the homogenous case.

In the second case, the excitation has a wavelength smaller than the spatial extent of the condensate but an energy smaller than the chemical potential. These excitations propagate as sound waves with a velocity approximately given by equation (8.4) where the value for n is the local density at the position of the excitation. These sound waves have been observed in experiments with ^{23}Na condensates at MIT, where a laser was focused at the center of the trap to create a wavepacket that was observed to propagate outward at the local speed of sound [139]. Also, the nonlinear nature of the GP equation has been used to describe soliton wave propagation in condensates as well [140].

Finally, when the wavelength of the excitation is comparable to the spatial extent of the condensate, then the excitation involves the motion of the condensate as a whole, as opposed to a single particle or some localized wavepacket moving through it. This type of excitation is referred to as a collective excitation, and it is this type of excitation that was studied in the measurement described in this chapter. Collective excitations correspond to the low energy part of the excited state spectrum determined from the solutions to equation (8.2). Often, collective excitations in condensates are considered in two limits, the non-interacting case and within the Thomas-Fermi approximation. The non-interacting limit is discussed in the subsection below, and the Thomas-Fermi limit is then treated in the subsection following that.

8.2.3 The Non-interacting Limit. In the absence of any mean-field interactions, there is no coupling between the two equations in (8.2), and so the $v(\mathbf{r})$ functions can therefore be ignored. The excitation spectrum is then the same as for a simple harmonic oscillator. Once this spectrum is known, the motion of a condensate can be described. Formally, this can be accomplished by

projecting the initial condensate order parameter onto the complete set of eigenstates of the simple harmonic oscillator by expressing it as a sum of the eigenstates

$$\Phi(t=0) = c_0|0\rangle + c_1|1\rangle + c_2|2\rangle + \dots + c_n|n\rangle + \dots \quad (8.5)$$

where $|n\rangle$ is the n^{th} eigenstate of the harmonic oscillator [26] and c_n is a complex number. If the condensate is not in the ground state, then its density profile will not remain constant in time and will oscillate. Each eigenstate will evolve in time with a phase factor proportional to the energy of the eigenstate $\exp(iEt/\hbar)$

$$\Phi(t) = c_0|0\rangle + c_1 \exp(-i\omega t)|0\rangle + c_2 \exp(-i2\omega t)|0\rangle + \dots + c_m \exp(-in\omega t)|0\rangle + \dots \quad (8.6)$$

and thus the condensate order parameter can then be computed for any time.

As a concrete example, consider one dimension of the simple harmonic oscillator trap potential. Since in the absence of any mean-field interaction term the GP equation is separable, the motion in each direction can be described individually. The ground state in this situation is a gaussian with an rms width equal to the simple harmonic oscillator length l , determined by the mass of the atom and the harmonic potential oscillation frequency as $l=(\hbar/m\omega)^{1/2}$. If a condensate is initially prepared so that its order parameter is a gaussian with a width smaller or larger than the harmonic oscillator length, then it will not be in the ground state and will oscillate in the trap. The frequency of this oscillation is determined by the excited state spectrum of the simple harmonic oscillator.

If the formal procedure presented above is followed, when the initial condensate order parameter is projected onto the spectrum of the eigenstates, the largest projection will be into the ground state, followed by the first excited state with even parity. Since the initial order parameter has even parity, it can only be decomposed into eigenstates of even parity as well. Other even-parity states will also be present in the projection of the initial order parameter. The larger the difference between the initial width and the ground state width, the greater the projection of the order parameter into higher excited states, both in an absolute sense and relative to the projection into the first excited state with even parity.

As an example of the oscillatory motion of a condensate in the non-interacting limit, consider a condensate whose width is initially larger than the simple harmonic oscillator length l . By summing over all the eigenstates to calculate the condensate order parameter as a function of time, it can be seen that the gaussian width contracts from its initial width to a size smaller than the simple harmonic oscillator length, where it then turns around and starts expanding in size until it reaches its initial width and then repeats the cycle. The energy of the even-parity excited states is equal to $2\hbar\omega$, $4\hbar\omega$, $6\hbar\omega$, ... and so for any pair of times separated by $2\pi/2\omega$ the phase factors associated with the time evolution of the eigenstates will all be equal. This means that the condensate oscillates at a frequency equal to 2ω . Figure 8.1 shows such an oscillation of the condensate density (the density of the condensate is equal to the square of the absolute value of the order parameter).

The preceding discussion has been a formal way of obtaining an intuitive result. A condensate that is initially larger than the ground state has more potential than kinetic energy and so it will contract inward. There is nothing to stop this contraction when the condensate reaches its ground state size, so it will continue to contract until the increase in kinetic energy slows the contraction to a halt, and then the condensate will begin to expand again until it reaches its initial size. The symmetry of the initial state means that the resulting oscillation will happen at twice the trap frequency. What the formal treatment does highlight is the important fact that the excitation state with the largest projection of the initial state when the width of the initial condensate does not match the width of the ground state is the lowest even-parity excited state, and that the oscillation occurs at the frequency corresponding to that state.

In general, it is the case that when the condensate order parameter has an rms size that is larger or smaller than the ground state rms size, the first excited state of even parity will have the largest projection of the initial order parameter into it of all the excited states. In the limit of a very small deviation of the initial state from the ground state width, the frequency that corresponds to this “lowest” excited state (actually the lowest energy eigenstate with even parity --- hereafter this state will be referred to as the lowest energy eigenstate) will therefore be the frequency of the oscillation of the order parameter around its ground state width. Measuring the oscillation frequency will then




Figure 8.1. Oscillation of a condensate density in one dimension in a harmonic trapping potential in the non-interacting limit. The initial order parameter at $t=0$ was projected into the first 20 simple harmonic oscillator basis states, and from the evolution of the phase factor associated with each basis state the time evolution of the condensate density could be determined. Since the initial condensate width is too large, it contracts as it evolves from $t=0$.

measure the energy of this excited state. As the deviation between the initial condensate width and the ground state width increases (i.e. as the oscillation amplitude gets larger), the majority of the initial order parameter will usually still be projected into the “lowest” energy excited state but other states will begin to become important. The oscillation frequency of the condensate will often no longer be exactly equal to the lowest excited state and will shift away from it as the size of the condensate oscillation increases. This change in frequency with oscillation amplitude is sometimes referred to as the quadratic shift, since its leading term scales with the amplitude of the oscillation squared. Since the eigenstates of a purely harmonic trap are all equally spaced in energy, there is no quadratic shift present for a purely harmonic potential (i.e. in the non-interacting limit and in the absence of any anharmonic terms in the potential).

What is less obvious from the more formal treatment of this problem is that if the initial condensate order parameter is a gaussian, it will remain a gaussian during the oscillation and only its width, and not its shape, will change. This is only true if the initial shape of the order parameter is a gaussian, however. Any other shape whose width is smaller or larger than the ground state value will still oscillate at twice the trap frequency, but while it is oscillating its shape will change. Figure 8.2 shows the oscillation of parabolic initial condensate density. The shape changes in an intuitive way. Since the wings of the parabola are smaller than a gaussian’s, there is less of a “force” acting on them since they are closer to the ground state of the trap than the rest of the order parameter. Therefore, the wings move less during the oscillation than the rest of the order parameter. This also means that the central part of the order parameter is narrower than a gaussian during parts of the oscillation.

8.2.4 Collective Excitations with Mean-Field Interactions. When mean-field interactions are present in the condensate, it is more difficult to calculate the spectrum of the excitations. The mean-field interactions now couple motion in orthogonal directions since the mean-field energy is proportional to the density of the condensate. In the Thomas-Fermi limit (see subsections 3.3.4 and Chapter VII), many of the collective excitations frequencies can be evaluated analytically and for simplicity of discussion this limit will be assumed for the rest of this subsection.




Figure 8.2. Oscillation of a condensate with a parabola-shaped initial density in the non-interacting limit. The time evolution of the condensate density was calculated in the same way as for Figure 8.1. The inset shows the presence of a spatial oscillation in the condensate density that exists at $t=\pi/2\omega$ (one half of the oscillation cycle). These ripples in the density were also present when 30 basis states were used to calculate the condensate evolution. While the size of the ripple is small, it can contain a significant fraction of the condensate atoms. For instance, if the condensate density shown was the same in all three dimensions, then 15% of the atoms would be contained in the ripples rather than the central feature. Notice that the central part of the order parameter is narrower at $t=\pi/4\omega$ than $t=\pi/2\omega$.

In a harmonic potential with spherical symmetry, the collective excitations can be labeled by the number of radial nodes (n_r), the angular momentum (L) and the projection of angular momentum onto the z -axis (m) of the eigenfunction associated with the collective excitation. The eigenfunctions in the spherical case are then [141]

$$\delta n(\vec{r}) = P_L^{(2n_r)}(r/R) r^L Y_{Lm}(\mathbf{q}, \mathbf{f}) \quad (8.7)$$

where δn is the difference between the ground state density of the condensate and the density of the condensate undergoing the collective excitation. R is the usual radius of the Thomas-Fermi order parameter given in equation (7.7). Several types of collective excitations are referred to by specific names due to the frequent interest of theorists and experimentalists in particular modes. For instance, modes with no nodes are referred to as surface excitations. A mode with $n_r=1$ and $L=0$ is referred to as a breathing or monopole mode and in a spherically symmetric potential represents a uniform oscillation of the spatial size of the condensate.

In a trap with axial (as opposed to radial) symmetry, m is still a good quantum number. Among the lowest energy collective excitations are quadrupole excitations of the form $\delta n=r^2 Y_{2m}(\theta, \phi)$ for $m = \pm 2, \pm 1$. The frequencies of these oscillations, which by the way are surface excitations, are given by the simple relations [141]

$$\begin{aligned} \omega^2(L=2, m=\pm 2) &= 2\omega_r^2 \\ \omega^2(L=2, m=\pm 1) &= \omega_r^2 + \omega_z^2 \end{aligned} \quad (8.8)$$

In the case of the $m=0$ oscillations, the $L=0$ and $L=2$ modes are coupled in a trap with axial symmetry. As is the case in a classical system, the coupling shifts the frequencies of the modes from their original values. The frequency of the two decoupled (i.e. normal) modes are given by [141]

$$\omega^2(m=0) = 2\omega_r^2 + \frac{3}{2}\omega_z^2 \mp \frac{1}{2}\sqrt{9\omega_z^4 - 16\omega_z^2\omega_r^2 + 16\omega_r^4} \quad (8.9)$$

It is these two $m=0$ modes that we studied as part of this work. The manner in which the condensate oscillates is shown in Figure 8.3. These two modes are the lowest energy collective excitations that can be excited through changing the mean-field energy of the condensate, and so the widths of the cloud will oscillate at these frequencies in response to a change in the scattering length.




Figure 8.3. Normal modes of the $m=0$ collective excitations. The length of the arrows refers to the relative motion of the axial (horizontal) and radial (vertical) during the normal mode oscillation. Mode (a) is the higher-frequency $m=0$ mode, while (b) shows the motion of the other (lower-frequency) $m=0$ mode.

One odd thing about the frequencies of the collective excitations is that once in the Thomas-Fermi limit the frequencies do not depend on the ratio Na/a_{ho} . In other words, the oscillation frequency is independent of the mean-field interaction strength. This behavior can be seen as a consequence of the relationship between the spatial size of the condensate and the phonon speed of sound at the center of the trap (equation (8.3)) [12]. As the mean-field energy increases, the spatial extent of the condensate increases and the wavelength associated with an excitation involving the whole condensate increases as well. The speed of sound at the center of the cloud increases by the same amount, however, due to the larger mean-field term. Since the dispersion relation is $\omega=cq$ in the phonon-like part of the excitation spectrum there is no change in the excitation frequencies with a change in the Na/a_{ho} ratio.

8.2.5 Experimental Observations of Collective Excitations. Shortly after the formation of dilute atomic condensates in magnetic traps, the frequencies of collective excitations of those condensates were studied. Here at JILA [142], the frequencies of the lowest energy collective excitations, the $m=2$ and one of the $m=0$ modes, were measured in an axially symmetric TOP magnetic trap with an aspect ratio of $1/8^{1/2}$. The frequency of these modes was observed to be a function of the ratio Na/a_{ho} , and tended to the Thomas-Fermi predicted frequencies in the limit of a large value of that ratio. The collective excitations were excited by driving changes in the trap potential with the proper symmetry near the excitation frequencies. The collective excitation frequency of the lowest $m=0$ mode was also measured in an experiment at MIT [143] with a condensate in a more anisotropic trap (aspect ratio of 230/17). In this case, the ratio Na/a_{ho} was much greater than one and the measured frequency was in reasonable agreement with the Thomas-Fermi prediction for the mode. Again, the collective excitation in the MIT experiment was created by introducing an oscillation into the trap potential near the excitation frequency.

In addition to measuring the excitation frequencies for condensates that were as cold as could be created, both the JILA and MIT groups measured the excitation frequencies as a function of the temperature of the condensate as well [144,145]. As the temperature was increased, large shifts in the excitation frequencies (on the order of 20% or more) occurred. Also, in both of these cases

the excitations were observed to damp out over time, and this decay rate was a strong function of temperature, increasing with increasing temperature of the condensate.

As per the discussion in subsection 8.2.3, any change in the width of the condensate will result in the lowest collective excitation frequency of the appropriate symmetry having the largest projection of the order parameter as compared to all of the other excited states. This means that deforming the trap will tend to only excite the lowest energy collective excitation frequencies, and so exciting collective excitations in higher energy states alone cannot be accomplished by deforming the trap potential. By increasing the amplitude of the excitation, the quadratic shift can be measured and that does give information about energies of the higher excited states, but does not measure any of those energies directly. Both the JILA and MIT groups measured the size of the quadratic shift in the collective excitations that were observed. Work at Oxford also examined the dependence of the size of the quadratic shift on the parameters of the magnetic trap [146].

8.2.6 Modes Frequencies between the Non-interacting and Thomas-Fermi Limit.

While the condensate oscillation frequencies are just twice the trap frequency in each direction in the non-interacting limit, and the frequencies in the Thomas-Fermi Limit of the two $m=0$ modes are given in equation (8.9), for the parameters of our experiments we often work in between these two regimes. Therefore, it is necessary for us to know the two $m=0$ collective excitation frequencies as a function of the ratio Na/a_{ho} to compare our observations to the mean-field theory predictions. The excitation frequencies can be calculated directly from the solutions of equation (8.2), as was done in Ref. [147]. Or, these frequencies can be calculated within the RPA [41], which yields equivalent results. These calculation methods are reasonably sophisticated and time consuming, however, and so a simpler method of accurately calculating the lowest excitation frequencies is desirable.

Such a model (the ‘‘PG model’’) has been developed by Perez-García and co-workers [148]. By fixing the shape of the condensate to be a gaussian, a set of three coupled equations for the widths of the condensate in each direction can be obtained:

$$\frac{d^2}{dt^2} v_i + \mathbf{I}_i^2 v_i = \frac{1}{v_i^3} + \frac{P}{v_i^2 v_j v_k}, \quad i = x, y, z, \quad i \neq j \neq k \quad (8.10)$$

with the following definitions

$$\begin{aligned}
V_{ext}(\vec{r}) &= \frac{1}{2} m \omega^2 (\mathbf{I}_x^2 v_x^2 + \mathbf{I}_y^2 v_y^2 + \mathbf{I}_z^2 v_z^2) \\
\mathbf{t} &= \mathbf{n} t \\
a_{PG} &= \sqrt{\hbar / m \omega} \\
\mathbf{s}_i &= a_{PG} v_i, \quad i = x, y, z \\
P &= \sqrt{\frac{2}{\mathbf{p}}} \frac{Na}{a_{PG}}
\end{aligned} \tag{8.11}$$

where σ_i is the rms width of the gaussian condensate order parameter in each direction. Note that the definition of the oscillator length a_{PG} is similar to, but not equal to, the definition often used for the simple harmonic oscillator length (see equation (3.21)): a_{PG} is defined with respect to the radial trap frequency alone while a_{ho} is defined using the geometric mean of all of the trap frequencies. This set of coupled equations in (8.10) is precisely the same as for a classical particle moving in a potential V_{eff}

$$V_{eff} = \frac{1}{2} (\mathbf{I}_x^2 v_x^2 + \mathbf{I}_y^2 v_y^2 + \mathbf{I}_z^2 v_z^2) + \frac{1}{2v_x^2} + \frac{1}{2v_y^2} + \frac{1}{2v_z^2} + \frac{P}{v_x v_y v_z} \tag{8.12}$$

The coordinates of the particle would then correspond to the scaled rms widths of the condensate order parameter. The frequencies of the lowest trap modes can therefore be determined from this effective potential by using standard techniques of classical mechanics. The position of the potential minimum (v_{x0} , v_{y0} , v_{z0}) can be calculated for each of the directions as a function of the parameter P

$$0 = \mathbf{I}_i^2 v_{i0} - \frac{1}{v_{i0}^3} - \frac{P}{v_{i0}^2 v_{j0} v_{k0}}, \quad i = x, y, z, \quad i \neq j \neq k \tag{8.13}$$

The effective potential is then expanded to second order around the position of the potential minimum. From the matrix of these terms evaluated at the minimum of the effective potential

$$\begin{pmatrix} \frac{\partial^2 V_{eff}}{\partial x^2} & \frac{\partial^2 V_{eff}}{\partial x \partial y} & \frac{\partial^2 V_{eff}}{\partial x \partial z} \\ \frac{\partial^2 V_{eff}}{\partial y \partial x} & \frac{\partial^2 V_{eff}}{\partial y^2} & \frac{\partial^2 V_{eff}}{\partial y \partial z} \\ \frac{\partial^2 V_{eff}}{\partial z \partial x} & \frac{\partial^2 V_{eff}}{\partial z \partial y} & \frac{\partial^2 V_{eff}}{\partial z^2} \end{pmatrix} \tag{8.14}$$

the squares of the lowest-energy collective excitations' frequencies is given by the eigenvalues of the matrix, and the eigenvectors describe the normal modes associated with each excitation. Of course the resulting frequencies will be for the scaled time τ , and so to obtain the true frequencies of the oscillations the square root of the eigenvalues must be multiplied by the quantity ω from equation (8.11).

In the limit where the interaction term P is much larger than one, the $1/(2v^2)$ terms associated with the kinetic energy of the condensate in equations (8.10) and (8.12) can be ignored. The resulting equations for the widths then match the frequencies given by equation (8.9) and equation (3.24) for the evolution of the widths in the Thomas-Fermi limit. Therefore, in the limit of large interactions, the PG model will reproduce the predictions for the excitation frequencies in the Thomas-Fermi limit. It also reproduces the frequencies of the oscillations in the non-interacting limit.

To see how well the PG model interpolates between these two extremes, the PG model predictions are compared to RPA predictions for an axially symmetric trap similar to the one that we used in our studies. The comparison is shown in Figure 8.4. There is good agreement (within 0.2%) between both the approximate PG method and the full RPA calculation. Therefore, we have used the PG model to generate the predictions for expected excitations frequencies.

One note of caution in using the PG model is appropriate here. The rms widths of the condensate calculated in the PG model do not agree with the rms width of the condensates calculated in the Thomas-Fermi limit, since the PG model fixes the shape of the condensate to be a gaussian. The rms width of the condensate density from the PG model is $\sim 20\%$ smaller than the rms width calculated from the Thomas-Fermi rms width corresponding to the same conditions. Note that the rms width calculated in the PG model is for the order parameter, which is $2^{1/2}$ larger than the rms width of the condensate density.

8.3 Beyond The Mean-Field Theory.

8.3.1 Pushing the Validity of Mean-field Theory. The range of validity for mean-field theory is determined by the condition that the dimensionless parameter na^3 be small, where n is the peak density of the condensate. The reasons for this are discussed in subsection 3.3.3. In order to




Figure 8.4. Comparison between the $m=0$ lower-frequency mode oscillation frequency calculated in an approximate model [148] developed by Perez-Garcia *et al.* and a full RPA calculation [41]. The approximate model is shown as the filled circles, while the RPA calculation is shown as a solid line.

make this parameter larger, either the density of the condensate needs to be increased, or the scattering length needs to be made larger. Because of the effect of mean-field repulsion, it is difficult to increase the density of the condensate, however. Creating a condensate with a larger number of atoms will increase its spatial size, causing the density to increase only slowly as a function of the number of atoms in the condensate. Because of the strong scaling of na^3 on a , increasing the scattering length allows the $na^3 \ll 1$ limit to be reached much more easily. This advantage of increasing the scattering length rather than the density can be demonstrated explicitly by calculating na^3 in the Thomas-Fermi limit

$$na^3 = \frac{m\omega_{ho}}{8\hbar\mu} \left(\frac{15}{a_{ho}} \right)^{\frac{2}{5}} (Na^6)^{\frac{2}{5}} \quad (8.15)$$

Because the Feshbach resonance can be used to increase the value of a , ^{85}Rb condensates provide a good candidate for tests of beyond-mean-field effects, even though the number of atoms that can be cooled into a condensate is relatively small. For comparison, the value of na^3 for an ^{87}Rb condensate of 10^6 atoms in a trap with $\omega_{ho}=500$ and a scattering length of $\sim 100 a_0$ is $na^3 \sim 6 \times 10^{-4}$. For a ^{23}Na condensate with 2×10^7 atoms, $\omega_{ho}=600$, and a scattering length $\sim 50 a_0$ has $na^3 \sim 10^{-4}$. In ^{85}Rb , values of $na^3 \sim 10^{-2}$ have been achieved. In these calculations, the peak density of the condensate has been used for n , and throughout the rest of this chapter whenever na^3 is computed the peak density will be used unless otherwise noted.

8.3.2 Calculations of the Energy of a Condensate Beyond the Mean-field. The purely mean-field treatment of the evolution of the condensate order parameter (i.e. the GP equation) ignores contributions to the energy due to the normal and anomalous quantum fluctuation terms, $\langle c^\dagger c \rangle$ and $\langle cc \rangle$ respectively, described in Chapter III. In this subsection, $\langle c^\dagger c \rangle$ and $\langle cc \rangle$ will be referred to by n' and m' respectively. Equation (3.18) shows how these quantities contribute to the evolution of the condensate order parameter Φ . When the parameter na^3 is small, then n' and m' are also small at temperatures close to $T=0$ and the GP equation should accurately describe the evolution of a condensate.

In principle, it would be possible to calculate the expectation values leading to the normal and anomalous densities precisely. Also, since the condensate and the normal and anomalous densities can interact with each other, the time-dependence of all of these quantities should be computed as well. In practice, such a calculation is quite complex [45]. Instead of solving the complete coupled problem several groups have calculated the correction to the energy of the condensate in a semi-classical approximation. While still complicated, it is considerably easier to determine n' and m' in the case of a homogenous condensate as a function of the condensate density to first order in a perturbation expansion. By using a local density approximation, the values of n' and m' on in the inhomogeneous case by the condensate density.

The result of the calculation of the condensate energy in this approximation will be a beyond-mean-field contribution to the energy of the condensate in addition to the kinetic, potential, and mean-field terms. This additional term is applied as a perturbation with respect to a given collective excitation mode calculated in mean field theory, and the resulting energy shift then gives the shift in the frequency of the collective excitation due to first-order corrections from beyond-mean-field effects.

Several groups have used this method to calculate the first order corrections to the two $m=0$ modes in an axially symmetric trap. S. Stringari and L. Pitaevskii [149] incorporated the first-order beyond-mean-field correction to the condensate energy determined by Lee, Huang, and Yang [39,136] into the GP equation and calculated the eigenfrequencies to the first order in this correction term. E. Braaten and J. Pearson [150] and separately S. Giorgini [151] calculated the expectation values for n' and m' in a semi-classical approximation as a function of the condensate density in different ways from each other and from Ref. [39,136] that eventually resulted in the same equation for determining the eigenfrequencies of the condensate $m=0$ oscillations with the beyond-mean-field effects included to first order. All of the predictions agreed with each other. The calculation of the first-order correction to the mean-field energy is made complicated by the presence of an ultraviolet divergence in that correction energy when calculated in the thermodynamic limit. This ultraviolet divergence comes from a process by which two atoms each scatter out of the condensate, collide with each other and then reenter the condensate. To remove this divergence, the atom-atom interactions, which are still treated as a contact potential, must be renormalized.

The prediction for the shift in the frequency of the two $m=0$ modes in an axially symmetric trap given by all of these theories in the limit that the temperature (T) of the condensate is equal to 0 is

$$\frac{d\omega}{\omega} = \frac{63\sqrt{\mathbf{p}}}{128} \sqrt{a^3 n_{pk}} \left(\frac{1}{2} \pm \frac{8 + \mathbf{I}^2}{6\sqrt{9\mathbf{I}^4 - 16\mathbf{I}^2 + 16}} \right) \quad (8.16)$$

where n_{pk} is the peak density of the condensate, $\lambda = \omega_z/\omega_r$ is the inverse of the aspect ratio, ω is the frequency determined for either the lower- or higher-frequency $m=0$ mode calculated in equation (8.9). The upper sign in front of the last term should be used with the higher-frequency $m=0$ mode and the lower sign used with the lower-frequency $m=0$ mode. The goal of our measurement was to determine if the frequency of the $m=0$ excitations shifted in the manner predicted by equation (8.16).

The inclusion of this beyond-mean-field energy will not only shift the frequencies of the collective excitation from their mean-field value, it will also affect the shape of the ground state condensate density. Given the precision with which it is normally possible to measure a frequency, the change in the shape of the condensate would be expected to be more difficult to measure than the frequency shifts. The density of all of the atoms in the ground state in the Thomas-Fermi limit, ρ , is given by [150]

$$\mathbf{r} + \frac{32}{3\sqrt{\mathbf{p}}} a^{3/2} \mathbf{r}^{3/2} = \frac{m}{4\mathbf{p}\hbar^2 a} (\mathbf{m} - V_{ext}(\vec{r})) \quad (8.17)$$

Figure 8.5 shows a comparison between the shape of the condensate density with and without the beyond-mean-field term proportional to $a^{1.5}$. It is important to note that equation (8.17) is for the total density of the atoms in the ground state in the trap, not just the density corresponding to the condensate (i.e. there is a normal component present even at $T=0$ and the total density (ρ) is the sum of the normal component density and the condensate density as discussed in subsection 3.3.3). The density of atoms in the condensate alone including beyond-mean-field effects to first order, given in the limit where na^3 is still small ($\sim 10^{-4}$ or so), is [149]

$$n(\vec{r}) \cong n_{TF}(r) - \frac{5}{4} \left(\frac{32a^{3/2}}{3\sqrt{\mathbf{p}}} \right) n_{TF}(\vec{r})^{3/2}. \quad (8.18)$$




Figure 8.5. Comparison of the total density with the beyond-mean-field shift correction and without. For simplicity, a spherical potential was assumed. The chemical potential between the two cases was set so that the number of atoms was the same in each case. The total density with the beyond-mean-field correction was calculated with equation (8.17) with $n_{pk}a^3$ set equal to 10^{-2} .

The quantity $n_{\text{TF}}(\mathbf{r})$ is the density calculated in the Thomas-Fermi limit without the consideration of any beyond-mean-field effects. The difference in the total number of atoms computed by integrating equation (8.17) over all space where it is positive and the number of atoms in the condensate itself determined by integrating equation (8.18) in the same way as (8.17) yields the total quantum depletion

$$\frac{N_{\text{out}}}{N} = \frac{5\sqrt{\mathbf{p}}}{8} \sqrt{n_{pk} a^3} . \quad (8.19)$$

where again this expression will only be valid for small values of na^3 .

There are several assumptions that have gone into the calculation of the beyond-mean-field frequency and energy shifts that may not be valid in the case of our trapped atoms. First, the dynamic coupling between the condensate and the normal and anomalous components has been ignored in the use of the local density approximation. These various components form a coupled system, and oscillations in the condensate could drive oscillations in the normal component, for instance. The coupling can also result in a shift of the oscillation eigenfrequencies. Also, the local density approximation itself will not be exact in the inhomogeneous case. This assumption shouldn't be too bad since the important size scale for the fluctuation terms is given by the healing length $\xi=(8\pi na)^{1/2}$ which is usually much smaller than the spatial extent of the condensate [44]. However, near the edge of the condensate where the density is less, this assumption will no longer be valid and the calculation of the beyond-mean-field energy shifts for that part of the condensate may not be correct. The effect of the spatial variation in the condensate density as a function of position on the size of the beyond-mean-field correction term can be estimated, and for conditions of 1000 atoms with a scattering length of 4000 a_0 in our trap it is less than 1%. Another assumption is that the physics of the Feshbach resonance can be entirely described by only an effective s-wave scattering length that has been taken to arise from a contact potential. For large values of na^3 , this contact potential description may no longer be adequate [152]. Three-body collisions, which are not included in these beyond-mean-field calculations, may become important as well, leading to deviations from the prediction since the interaction energy due to these three-body collisions will not have been taken into account.

Lastly, while the leading beyond-mean-field corrections scales as $(na^3)^{1/2}$, the next-order term has not been calculated as accurately yet and so it is unclear when this first-order correction will no longer be correct. A Monte-carlo simulation of homogenous condensates [152] indicates that the next-order beyond-mean field term is opposite in sign to the first order correction, so that the first-order correction will overestimate the correction to the mean-field as the parameter na^3 is increased. Also, there is an argument [153] that eventually the energy per atom of the condensate must scale as $\hbar^2 n^{2/3}/(m)$ in the limit of very large interactions. This energy is calculated by assuming that all of the atoms in the condensate are confined to boxes with a side equal to $1/n^{1/3}$. In this case, the wavefunctions of the individual particles will not overlap and so there will be no contribution to the total energy of the condensate from the mean-field energy. From a variational perspective, this trial wavefunction for the ground state must therefore set an upper limit on the energy of the ground state of the condensate. Since this limit on the ground state scales with the density less steeply than the energy from the mean-field interaction, it is expected that eventually the beyond-mean-field effects will reduce the energy of the condensate from the value predicted in the mean-field approximation. This effect of the beyond-mean-field interactions is opposite in sign to the leading order correction, and so eventually the higher-order terms should overwhelm the leading correction according to this argument as well.

8.3.3 Finite Temperature Effects and Beyond-Mean-Field Theory. One way of looking at the problem of calculating beyond-mean-field effects is that it is the inclusion of the effects on the condensate dynamics due to atoms that are not in the condensate. While a large value for na^3 leads to atoms being depleted from the condensate (equation (8.19)), a finite temperature of the trapped cloud also means that not all of the atoms are in the condensate. The same considerations and limitations of the theory calculations outlined in the previous subsection apply to the calculation of the effect of finite temperatures on the energy of the condensate as well. The previous measurements at JILA and MIT [142,143] showed that a finite temperature could lead to large shifts in the measured oscillation frequencies. We too observed large shifts in the measured frequency as a function of the temperature of the atoms in the cloud, and ultimately these temperature shifts

obscured the zero temperature beyond-mean-field shifts (equation (8.16)) that were the goal of this measurement.

8.3.4 Predicted Magnitude of the Beyond-Mean-Field Shifts. It is possible to calculate the predicted beyond-mean-field frequency shift for both $m=0$ modes for our typical trap parameters ($2\pi \times 17.413$ Hz by $2\pi \times 6.791$ Hz) from equation (8.16) assuming that the peak density of the condensate is equal to the peak density calculated in the Thomas-Fermi limit

$$\begin{aligned} \frac{dw}{w} &= 0.500 \times 10^{-6} (Na^6)^{\frac{1}{5}} \quad \text{radial oscillation} \\ \frac{dw}{w} &= 0.077 \times 10^{-6} (Na^6)^{\frac{1}{5}} \quad \text{axial oscillation} \end{aligned} \quad (8.20)$$

where a is expressed in units of the Bohr radius (a_0). Clearly, the radial oscillation is more sensitive to the Beyond-Mean-Field frequency shifts for our trap parameters. For 1000 atoms at a scattering length of $3000 a_0$, the fractional shifts in the oscillation frequencies are 3.0% and 0.5% for the radial and axial modes, respectively. The usual precision inherent in frequency measurements is high enough to easily detect a 3% frequency shift, indicating to us that the measurement of the beyond-mean-field shift was feasible.

8.4 Inducing Collective Excitations in ^{85}Rb .

8.4.1 Changing the Scattering Length to Create Collective Excitations. Using the Feshbach resonance, it is easy to change the value of the s-wave scattering length, and therefore the mean-field energy of an ^{85}Rb condensate, by altering the bias magnetic field in the magnetic trap. In fact, the s-wave scattering length changes fast enough with bias magnetic field near the Feshbach resonance that it is difficult to change the radial and axial trap frequencies without changing the scattering length. Therefore, it would be very difficult to excite the collective excitations in ^{85}Rb condensates through the modulation of the trapping currents as was done in Refs. [142-146]. If the bias magnetic field is changed slowly enough, then the condensate remains in the ground state. However, if the magnetic field is changed to alter the mean-field energy at a rate faster than the lowest collective excitation frequency, then the condensate will begin to oscillate.

The basic physics of the induction of the oscillations is easy to understand. For example, if the magnetic field is changed rapidly to increase the value of a , a condensate originally in the ground state will not have time to adjust its spatial size and it will be smaller than the ground state at the new value of a . The increased repulsion between the atoms caused by the increase in a then causes the spatial size of the condensate to expand outward. With no significant damping in the problem, the condensate continues to expand outward past the ground state equilibrium size until the external trapping potential slows the atoms down and causes the condensate to contract again. This motion repeats itself at the oscillation frequency characteristic of the lowest collective excitation states.

In our experiments, the value of a was usually changed on the order of 1 ms, much faster than the two collective mode excitation oscillation periods of ~ 28 and ~ 94 ms. The density of the condensate does not have time to change significantly during the change in the bias magnetic field, and these changes in a can be modeled as sudden (i.e. instantaneous) jumps to a reasonable approximation. The deviation of the widths of the condensate from the ground state at the new field can be expressed as a sum of the normal modes associated with the two $m=0$ collective excitations. Actually, there are three modes that should be considered instead of two, but the cylindrical symmetry of the trap makes any coupling into the other ($m=2$) mode small. The time evolution of the condensate widths can be determined from the oscillation frequency of each normal mode and the projection of the initial condensate into the two normal modes.

For our trap frequencies (roughly $2\pi \times 17.4$ Hz radial by $2\pi \times 6.8$ Hz axial) the high-frequency $m=0$ mode is an oscillation mostly in the radial direction and the low-frequency $m=0$ mode is an oscillation mostly in the axial direction. The two normal modes can be calculated using the PG model in the Thomas-Fermi limit. The resulting eigenvectors of the matrix in equation (8.14) in this limit are

$$\begin{aligned} \text{high-frequency mode} &= \begin{pmatrix} 0.699 \hat{x} \\ 0.699 \hat{y} \\ 0.151 \hat{z} \end{pmatrix} \\ \text{low-frequency mode} &= \begin{pmatrix} -0.107 \hat{x} \\ -0.107 \hat{y} \\ 0.989 \hat{z} \end{pmatrix} \end{aligned} \tag{8.21}$$

Therefore, if the condensate is oscillating only in the low-frequency mode, the axial direction will oscillate with an amplitude that is $0.99/0.11 = 9$ times larger than in the radial direction. The two directions will also be oscillating out of phase. In the high-frequency mode the two directions oscillate in phase with the radial amplitude a factor of ~ 5 times larger than the axial. In general, a sudden change in the scattering length will excite motion in both normal modes. For instance, if the condensate begins in a ground state in the Thomas-Fermi limit, a sudden increase in the scattering length will result in both $m=0$ modes being excited almost equally. The ratio of the excitation amplitude of the higher-frequency mode to the lower-frequency mode would be ~ 0.8 in this case.

Because the two $m=0$ modes each have a stronger oscillation amplitude in one particular direction, a different terminology will be adopted for the rest of the chapter. The high-frequency $m=0$ mode will be referred to as the radial oscillation, and the low-frequency $m=0$ mode will be referred to as the axial oscillation. Even though these names have been given to these two modes, it is still the case that neither is a purely axial or purely radial oscillation.

In order to calculate the expected mean-field oscillation frequencies, the trap frequencies had to be known in all three directions. The frequencies were determined using the slosh method described in subsection 4.10.2 to a precision of better than 0.2%.

8.4.2 Control of the Excitation Amplitude of Each Individual Mode. In the course of our measurements, we were usually interested in measuring one mode or the other. We also desired to control the amplitude of the oscillation in whichever mode that we had chosen to measure. A sudden change from one value of the scattering length to another can be used to control the amplitude of the oscillations by increasing or decreasing the magnitude of the change of the scattering length. However, there is no way to excite one mode selectively using only a single sudden change of the scattering length, and so a more complicated time dependence of scattering length is necessary in order to excite one mode significantly more than the other.

As described in Chapter IV (subsection 4.5.2), the value of the bias magnetic field in the trap is proportional to the current in the bias coils and the baseball coil. During these experiments, the bias coil current was held fixed while the baseball coil current was changed. This had the benefit of keeping the trap frequencies approximately constant while the changing the bias magnetic field

(subsection 4.10.1). The SRS programmable function generator was used to control the baseball current, and so very complicated waveforms could be applied to change the bias field in the magnetic trap. This allowed us to use two strategies to selectively excite individual oscillations in the condensate: a parametric drive and a sequence of timed sudden changes in the scattering length to produce the desired oscillation.

The parametric drive consisted of varying the bias magnetic field of the trap sinusoidally at the frequency of the desired excitation. The resulting time-dependent variation in the scattering length then resonantly drove the desired oscillation mode. In principle, the other mode would not be excited with any significant amplitude since its frequency did not match the frequency of the resonant drive. One large drawback to the parametric drive is that the drive must be on for at least several oscillation periods in order to excite the desired oscillation. With the presence of heating and loss in the condensates, we often could not afford to wait this amount of time.

A timed sequence of sudden jumps allowed us to selectively excite a desired condensate oscillation in a much shorter amount of time. The basic idea behind this method was to use a series of carefully-timed rapid jumps in the scattering length to produce an oscillation of a desired amplitude in one mode while suppressing the oscillation amplitude in the other. The first step was to slowly ramp the bias field in the trap to achieve an initial scattering length to set the initial spatial size of the condensate to a desired value. The bias field for this condition will be referred to as the launch field. Next, the bias field was jumped in ~ 1 ms to an intermediate value between the launch field and the desired final bias field where the oscillation was to be measured. At this intermediate field, a delay was added to let the condensate oscillate. For clarity, it will be assumed that the oscillation in the axial direction is to be reduced, although this same method can be adapted to suppress the oscillation in the radial direction as well. After the jump to the intermediate field, the axial width of the condensate was allowed to expand or contract due to the sudden change in scattering length until it reached its turning point. The intermediate field was chosen so that the spatial extent of the condensate in the axial direction at the turning point was nearly equal to the final field ground state axial spatial extent. By again jumping the field to the final field value when the condensate had reached the axial turning point, the axial direction of the condensate would be the

proper size and it would not be contracting or expanding. In this way, the size of the axial oscillation was suppressed.

While this method of time-sequenced jumps of the fields is an improvement over simply changing the field from some launch field directly to the final field, it is still not possible to control the radial and axial oscillation amplitudes completely independently. If the parameters are chosen to minimize the axial oscillation, then for those parameters the radial oscillation will have a set value. If that radial oscillation amplitude is smaller or larger than desired then a compromise must be reached between suppressing the axial oscillation and getting the desired radial oscillation. Choosing the launch field carefully can improve the suppression for a given desired oscillation size somewhat, but the relative insensitivity of the radial and axial oscillation frequencies to the ratio Na/a_{ho} in the Thomas-Fermi limit limits the gains that can be made by changing the launch field. By waiting longer, the fact that the radial frequency was not a multiple of the axial frequency meant that a more favorable radial oscillation amplitude could be found for a given axial amplitude. However, we often were not able to wait long due to the loss and heating of the condensate and so usually we had to accept some compromise between the desired radial and axial oscillation sizes. Typically, we desired a 20% amplitude oscillation in one direction, and the residual oscillation in the other direction was 7% or less.

In order to estimate the intermediate field and the delay time that would be necessary, we developed a very simple model based loosely on the PG model. The condensate widths were treated like particles in a harmonic potential whose minimum was located at the width corresponding to the ground state width, which depends on the value of the scattering length. The frequencies of this model potential were chosen to be the mean-field oscillation frequencies calculated from the PG model. It was assumed that the changes in bias field were instantaneous and that the effect of a change in field was a change in the location of the minimum in the potential. A sudden change in the field therefore resulted in the widths of the condensate oscillating in response to the sudden change in the potential minimum. The condensate oscillation amplitudes resulting from a given sequence of timed jumps can be calculated in this model and then the parameters corresponding to the most desirable oscillation amplitudes chosen. Figure 8.6 shows the predictions from this simple




Figure 8.6. Comparison of the observed radial and axial oscillation amplitudes to those predicted using the simple model described in the text. For these data the launch, intermediate, and final fields were 162.0, 158.3, and 156.5 G respectively.

model for a typical set of parameters and the comparison of the model prediction to the measured oscillation amplitudes.

This simple model had the advantage of producing estimates that were reasonably close to the observed oscillation amplitudes without any difficult calculations. It did suffer from some problems due to the simplifying assumptions that were made, and that is why the agreement between the observed and predicted oscillation amplitudes in Figure 8.6 is not better than shown. The effective potential used to describe the “motion” of the condensate widths should not be simply harmonic, but should include anharmonic terms as well. Also, the assumption that the condensate has a fixed shape is not valid if any jumps are made from the values of the scattering length where the Thomas-Fermi limit holds into a region of scattering length where the non-interacting limit is a better description. The changes in shape that can be induced by rapid changes in a are not accounted for in the simple model or the PG model and can lead to errors in the calculation of the oscillation amplitudes and the spatial extent of the condensate. Even with these difficulties, the simple model predictions were good enough to set the scattering length jumps to produce oscillation amplitudes that were close to their desired values. Some slight tweaking of the delay or intermediate field was then used to fine-tune the oscillation amplitudes, if desired. Some improvement in the simple model could probably be made by putting the jumps and delays into the PG model directly and calculating the resulting oscillation amplitudes, but this adds complexity to the calculation. The PG model will also do nothing to alleviate the errors due to the change in the shape of the condensate between the Thomas-Fermi and the non-interacting limits.

One limitation to this method of time-sequenced jumps is that it is difficult to induce a radial-only oscillation for small mean-field energies in our trapping potential. In the non-interaction limit, the condensate ground state aspect ratio is about 1.6, while it is about 2.6 in the Thomas-Fermi limit. Therefore, the axial size changes much more in comparison to the radial size between the Thomas-Fermi and non-interacting limits. It then becomes difficult to induce a large radial oscillation without a large axial oscillation also being present.

8.5 Measuring the Oscillation Frequencies.

8.5.1 Experimental Method. The measurement of the condensate oscillation frequencies was performed by measuring the condensate radial and axial width as a function of time at a selected magnetic field. The condensates were first created at 162 G in the usual manner described in Chapter IV. After that, either a parametric drive or time-sequenced jumps in the scattering length were used to excite an oscillation in the condensate at the magnetic field where the oscillation frequencies were to be measured. After a delay time at that field, the magnetic trap was turned off and the atoms were imaged using absorption imaging to determine the number of atoms in the cloud and the axial and radial widths. Since the absorption imaging was destructive, this experimental sequence was repeated many times. The radial and axial frequencies were thus measured as a function of the delay time to map out the oscillations. An example of this measurement of an oscillation frequency is shown in Figure 8.7.

8.5.1.1 Transverse Magnetic Field Imaging. Since the widths of the condensates were changing during the oscillation, occasionally the peak optical depth of the condensate would be greater than 2 or more. This high optical depth will not only cause a distortion in the measurement of the number due to the optical depth ceiling (see subsection 4.12.5), the measured widths can also be distorted since the true peak optical depth of the condensate will be underestimated in the absorption image. When this high optical depth was a problem, we rotated the magnetic field applied during the imaging by 90° to reduce the Clebsch-Gordan coefficients of the transition that the probe laser drove to do the imaging. This had the effect of reducing the optical depth of the cloud by about a factor of two for our experimental conditions. The absorption images with the bias field rotated were compared to images with the bias field in the normal direction in order to calibrate this reduction factor. The calculation of the reduction that is expected would be difficult due to the fact that the optical pumping of the sample in this situation is complex. The polarization of the atoms will eventually reach a steady state as the laser attempts to align the atoms along the probe axis and the transverse magnetic field causes those spins to precess.

8.5.2 Oscillation Frequency Signal-to-Noise Considerations. For a fixed size statistical uncertainty on each measurement of the widths of the condensate (typically 5% in our measurement), the precision of a frequency measurement with a fixed number of data points is




Figure 8.7. This figure shows a typical measurement of radial oscillation frequency. There appears to be a small amount of damping, and this is reflected in the fit to the data. However, this damping may in fact have other causes, as described in section 8.5. These data were taken at a field of 158.7.G

proportional to the amplitude of the oscillation squared. The precision will improve linearly with the delay time spanned by the data (i.e. measuring the oscillation out to twice as long a time with the same number of data points will reduce the statistical uncertainty in the measured frequency by a factor of two), and of course the statistical uncertainty will decrease with the square root of the total number of points taken. Therefore, it was desirable to measure a large amplitude oscillation over a long period of time.

One difficulty with a large amplitude, however, is the fact that the frequency of the oscillation will in general shift quadratically with amplitude of the oscillation (see subsection 8.1.3). To minimize any quadratic shifts, the amplitude was usually kept between 15-20%, the smallest size that still provided a reasonable signal-to-noise ratio. Also, the quadratic shift was measured at a few fields to make sure that the frequency was not being distorted significantly.

Measuring the oscillation over a long period of time also introduced complications into the measurement. In contrast to other collective excitation measurements [142-146], we observed what appeared to be only a slight damping of the collective excitations. This is perhaps not surprising given the very cold temperatures of our condensates [154]. However, there were temperature-dependent frequency shifts and since the temperature of the condensates increased over time, this limited the amount of time that could be spanned measuring the oscillation frequency.

Our usual method of taking data was to take data at 8-10 different delay times that mapped out the first oscillation period. We would then take a similar set of data at a later time, typically after 7-9 oscillation periods. The relative phase of the two oscillations would then fix the frequency of the oscillation. In addition to these two sets of data, two additional sets were taken at intermediate times. These additional sets did little to improve the statistical uncertainty of the measurement, but they were very important for identifying a variety of systematic effects, including shifts in the frequency over time and distortions in the measured widths due to an unexpected systematic distortion, which are both described below.

8.5.3 Condensate Heating and Loss Rates. At magnetic fields with large scattering lengths, the condensate number loss and heating rates are also high. Condensates with a few thousand atoms were observed to have number decay lifetimes as short as 100 ms at fields with

scattering lengths of several thousand a_0 . In addition to the loss of atoms from the trap, these condensates were also observed to “melt.” That is, a thermal component in the trapped cloud was observed to grow with time as the condensate remained at the high-scattering length field. In order to minimize the heating and loss rates, condensates with small numbers of atoms were used to perform the oscillation frequency measurements.

The increase in the loss rate at high-scattering length magnetic fields is expected from the magnetic field dependence of the inelastic collision rates near the Feshbach resonance. In the regions of large positive scattering length in which we wanted to measure the beyond-mean-field shifts, the dominant loss mechanism was three-body recombination (see Chapter VI). The heating rate had the same general dependence on field and atom number as did the inelastic loss rate.

In order to estimate the optimum number of atoms in the condensate for the measurement of beyond-mean-field shifts, it was assumed that the three-body rates did indeed scale as a^4 (see subsection 6.1.2.3). Then, with the density of the condensate again calculated in the Thomas-Fermi limit, the three-body loss rate would be proportional to $N^{4/5}a^{14/5}$. Fixing the predicted beyond-mean-field frequency shift to a given value (and therefore fixing the quantity Na^6), the loss rate would then scale as $N^{1/3}$. Therefore, for a given beyond-mean-field frequency shift, it was advantageous to use condensates with smaller rather than larger number in order to reduce the loss and heating rates. Most of our data near the Feshbach peak were taken with condensates with less than 2000 atoms, and for measurements at the highest scattering lengths condensates with ~ 800 atoms were used.

8.5.4 Accounting for Number Variation. The oscillation measurements described here were performed before the shot-to-shot variations and long term drifts in the number of atoms in the condensates were improved (see subsection 7.1.7). The rf-shift (see subsection 7.1.7) offset could be used to correct for the drift, but doing so usually required taking a few destructive measurements with the resulting loss of duty cycle. Instead we let the number of atoms drift over the course of the oscillation measurement.

This variability in the number posed a potential problem in the measurement of the oscillation frequencies. The beyond-mean-field shift and the mean-field-only frequency do not change quickly with N , and so the variation should not have produced significant errors in that

respect. Rather, the ground state size of the condensate was a function of number (although again weakly) and so therefore the amplitude of the oscillation that was induced was a function of the number. Also, the ground state size around which the condensate oscillates at the final field was a function of number. Since the number of atoms could be determined from each absorption image, however, it was possible to correct for the amplitude and ground state width changes. Usually, this was accomplished by introducing additional parameters to the fitting function used to determine the frequency from the oscillation data: a linear dependence of the oscillation amplitude on the number of atoms and a linear dependence on number of the width around which the oscillations occurred. Also, the oscillation data taken under the same conditions were sometimes binned into number ranges and the frequencies of binned data sets measured to search for any number dependence. Of course, with a smaller number of data points in each binned set the precision of the individual frequency determinations was smaller. Figure 8.8 illustrates the accommodation of the number variation in the fit to determine the frequency.

8.6 Results of the Oscillation Frequency Measurements.

Figure 8.9 shows the measured axial and radial oscillation frequencies as a function of field. The data shown in Figure 8.9 were most often taken in the manner described in the last paragraph of subsection 8.5.2 and shown in Figure 8.7. We were primarily interested in the change in the measured radial frequency as a function of field since it is more sensitive (factor of 6.5) to the leading order beyond-mean-field corrections than the axial mode. However, because the mean-field predicted frequency of the axial mode changes by 20% from the non-interacting limit to the Thomas-Fermi limit (as opposed to 1% for the radial) and because of its lesser sensitivity to the beyond-mean-field corrections, the axial mode serves as a useful comparison between the measured frequencies and their mean-field predicted values.

Examining just the axial mode, the measured frequencies follow the general trend with magnetic field predicted from mean-field theory: as the magnetic field is tuned to increase the s-wave scattering length, the axial oscillation frequency decreases from the twice the axial trap frequency to a constant value in the Thomas-Fermi limit. While the general trend is the same as predicted, the measured axial oscillation frequencies do not agree quantitatively with the mean-field

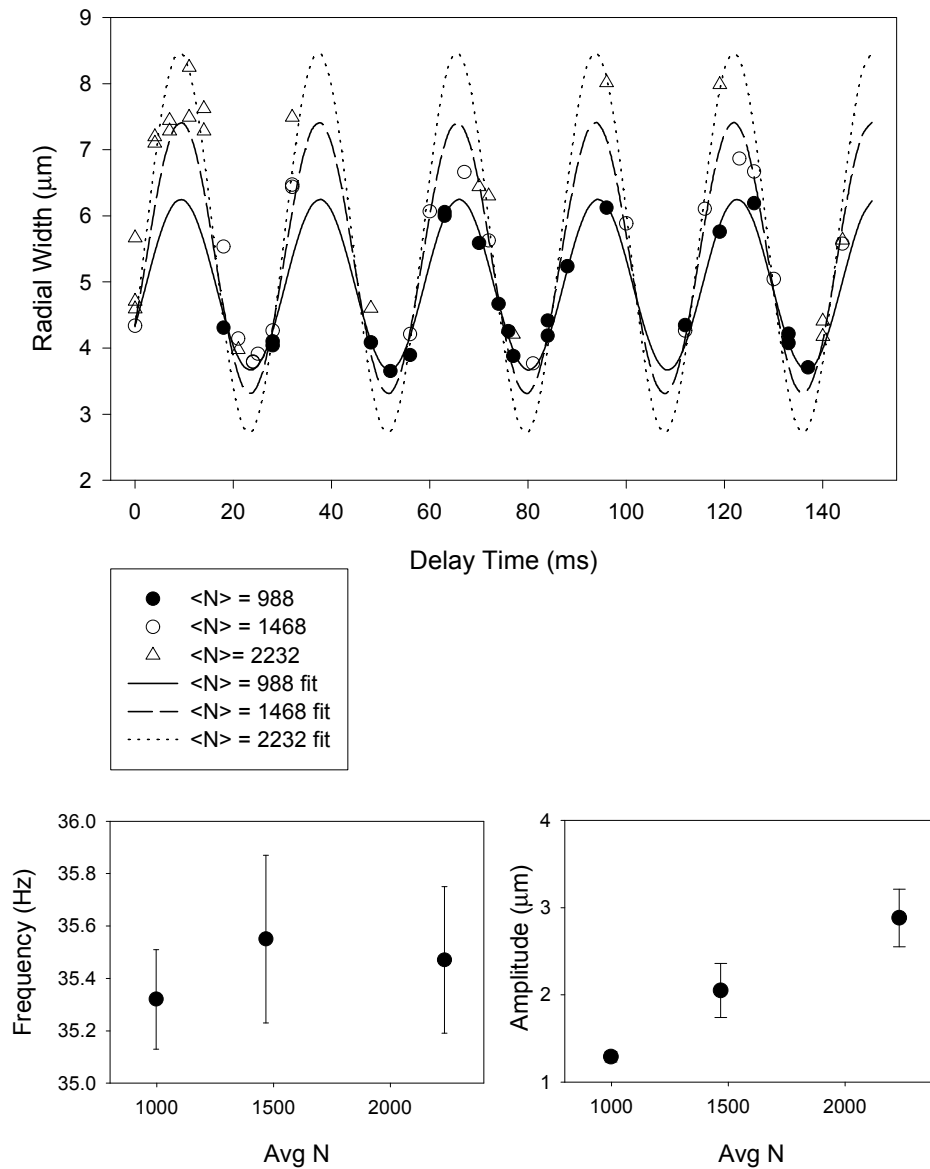


Figure 8.8. Dependence of the oscillation amplitude on condensate number. In this figure, the data have been binned by number into three sets and a sine wave was individually fit to each set. The two lower figures show the lack of any dependence of the measured frequency on the number and the oscillation amplitude of the fit to each set of data. The field at which these data were taken was 156.5 G.




Figure 8.9. Comparison of the measured oscillation frequencies to the mean-field and beyond-mean-field predicted frequencies. Most of the data points were taken over several oscillation periods and were therefore vulnerable to the temperature shift systematic discussed in the text. One data point taken over only two oscillation periods is also shown as the open circle, and the disagreement between that point and the others at the same field is additional evidence for the temperature dependent frequency shifts. The number of atoms used to measure the oscillation frequencies varied from field to field, and that accounts for the structure of the mean-field prediction curves. The dashed line shows the predicted oscillation frequencies from equation (8.16) for $B_{\text{zero}} = 154.9$, while the two dotted lines show the predicted frequencies as B_{zero} is shifted ± 0.4 G (see Chapter V).

predicted frequencies at the higher values of magnetic field that were measured. We believe that it is likely that systematic errors in the frequency measurement are due to both changes in the shape of the condensates between the Thomas-Fermi and non-interacting limits and coupling between the axial and radial directions during the trap turn-off and preparation for imaging of the atom cloud. These systematic effects are described in more detail below in section (8.7). In contrast to the disagreement at higher magnetic fields, at the magnetic fields with large scattering length the axial frequency agreed with the mean-field predictions, suggesting that the systematic errors present at higher fields were not present at lower magnetic fields.

Turning to examine the radial mode, the radial oscillation frequency agrees with the mean-field predictions within the uncertainties in the individual frequency measurements from the 164 G field down to 159 G. The agreement with the mean-field prediction at 159 G is interesting since the beyond-mean-field corrections would predict a shift of 0.22 Hz from the mean-field value, 35.41 Hz, statistically different from the measured radial oscillation frequency of 35.20(7) Hz at that field.

As the magnetic field is tuned closer to the peak of the Feshbach resonance to increase the s-wave scattering length, the measured oscillation frequencies begin to diverge from the mean-field predicted values. This divergence is statistically significant and in the direction predicted for the beyond-mean-field shift, but the magnitude of the observed shifts is smaller than the magnitude predicted in equation (8.16). The divergence originally did not appear to continue as the magnetic field was tuned closer to the Feshbach resonance, however, and the measured oscillation frequency at 155.7 G first agreed with the mean-field prediction.

The loss and heating rates at 155.7 G were large, and we had some indications that the radial oscillation frequency was shifting downward as the condensates were melting. To see if such a temperature shift was obscuring the beyond-mean-field shift at 155.7 G, data were taken over a range of delay times corresponding to only two oscillation periods as opposed to seven in the original data. The data from this measurement are shown in Figure 8.10. The measured radial frequency, 36.09(28) Hz, was different from the mean-field prediction of 35.28 Hz and the original measurement of 35.27(14) Hz. This clearly indicated that the radial oscillation was shifting as the condensates melted. It also clearly indicated the presence of beyond-mean-field frequency shifts.




Figure 8.10. Radial oscillation frequency measured at 155.7 G using a relatively short integration time. The average number of atoms in this set of data was about 800.

Finally, these data also made it apparent that the temperature shifts would severely complicate the determination of the magnitude of the beyond-mean-field shifts. These temperature shifts are discussed in more detail in section 8.8 below and the estimate of the magnitude of the beyond-mean-field shifts is discussed in section 8.9.

8.7 Distortions in the Measurement of Oscillation Frequencies.

This section concerns itself with the deviation of the measured axial oscillation frequencies from the mean-field predictions at magnetic fields with relatively small scattering lengths ($200 a_0$ and less). At the time that the data in this chapter were taken, the source or sources of these discrepancies were not known. During subsequent experiments measuring the collapse of condensates with attractive interactions (see Chapters IX and X), both what we learned about the collapse of the condensates and our increased familiarity with the effects of the manipulation of the scattering length revealed several likely sources of error in the axial oscillation frequency. These likely sources of error are discussed in the two subsections below. It should be emphasized that the mechanisms described below have not been confirmed as the source of the discrepancy.

8.7.1 Measurement of the Axial Frequency at $a=0$. One of the checks that we performed was the measurement of the axial oscillation frequency in the non-interacting limit. Figure 8.11 shows data measuring the axial oscillation frequency at $a=0$. In order to set the final field so that $a=0$, a collapse measurement (see Chapter IX) was performed to determine where a 5000 atom condensate would collapse. From this measurement and the measured value of the stability coefficient k , the magnetic field could be set so that the scattering length was within $0.7 a_0$ or better of $a=0$. The axial trap frequency at this field was $2\pi \times 6.797$ Hz, meaning that the axial oscillation frequency should have been 13.594 Hz. The time-sequenced jumps were used to induce the oscillation for this measurement. We measured an axial oscillation frequency of 13.81(7) Hz, which disagreed with the predicted value by three standard deviations.

Because of the small size of the condensate at $a=0$, the widths of the condensate could not be measured by turning off the magnetic trap directly from the $a=0$ field to image the atoms. The size of the cloud would then have been below the resolution limit, making it difficult to observe changes in the cloud width. So instead, we jumped the field again after letting the condensate




Figure 8.11. Measurement of the axial oscillation frequency near $a=0$ (165.85 G). Twice the axial frequency was 13.594 Hz, so the measured frequency does not agree with the expected value. The problem points, which are discussed in the text, are indicated by the arrows.

oscillate for a delay time at $a=0$. The field to which we jumped was 164.2 G ($a=67 a_0$). After this jump, we then waited ~ 40 ms to let the axial size of the cloud expand to its turning point at this new field. The trap was then turned off to image the cloud. The idea behind this additional jump was to convert the oscillation at $a=0$ into a variation in size in an expanded condensate after the field jump to positive scattering length. If the condensate had a relatively small spatial extent at $a=0$, then the sudden jump would impart more mean-field energy than if the condensate had a larger spatial extent. This means that condensates that were originally small at $a=0$ would expand to a larger size after the field jump to positive scattering length than condensates with a size that was originally larger at $a=0$. By mapping out the expanded size of the condensate as a function of delay time at $a=0$, the oscillation frequency could then be determined.

We worried that this extra jump to positive a would result in the variation of the expanded size not being strictly sinusoidal, but since we were looking for a periodic function, we assumed at the time that any deviation from the sinusoidal shape would occur in the same way in each cycle and so a frequency could still be accurately determined. Basically, since whatever the axial direction was doing should repeat every 73.6 ms we could still accurately determine the frequency even in the presence of any systematic deviations from a sinusoidal pattern that the expanded condensates exhibited.

In light of the disagreement between the measured frequency and the predicted frequency, however, we now believe that we did not properly account for the oscillation in the radial direction coupling into the axial width measurement. Our assumption that the observed axial width should repeat precisely every 73.6 ms is only true if there was just one frequency in the problem --- once another frequency can couple to the motion there is no guarantee that the observed axial width will repeat precisely at its oscillation frequency. The coupling between the radial and axial oscillations could have occurred in several ways, and three speculations of how this could have happened are discussed in this and the following paragraphs. One obvious candidate is the jump back to 164.2 G to expand the condensate size above the resolution limit. The amount of mean-field expansion

induced by this field jump will depend on the density of the condensate, and that in turn depends on both the axial and radial widths, not just the radial width alone.

Another possibility of how the radial oscillation can influence the axial comes about because of the change in the shape of the condensate ground state. Since the condensate was formed in the limit where the Thomas-Fermi approximation is valid, the initial shape of the condensate was more parabolic than gaussian. During the time-sequenced jumps, not only will an oscillation in the condensate width be introduced but the shape of the condensate will oscillate as well. This shape oscillation affects the central density of the condensate and therefore the coupling between the radial and axial directions in a way not taken into account into the calculation of the oscillation frequencies. The shape oscillations may also lead to errors in the determination of the widths of the condensate during the absorption imaging since a gaussian shape is used to fit the condensate widths.

Another way the axial and radial directions could have been coupled in the measurement of the oscillations at $a=0$ was through the evolution of the condensate after the magnetic trap was turned off and before the atoms were imaged. The scattering length is negative at $B \sim 0$ and so it is possible for the condensate to contract inward after the trap has been turned off. Normally, the impulse imparted during the trap turn-off prevents this contraction (subsection 7.2.2), but for samples with a high enough initial density, the attractive mean-field interaction, which scales strongly with the condensate density, is enough to overwhelm this outward impulse and cause the cloud to contract inward, and even collapse as described in subsection 10.3.7. One of the results of this collapse can be a burst of atoms, and this burst could distort the measurement of the condensate widths. Since the condensate collapse and the outward impulse imparted both would have depended on the density of the cloud, and therefore again depended on the radial and axial widths, this represents another way that the two directions can be coupled together.

Although there is no hard evidence, there are some suggestions in the data that these types of effects were the source of the discrepancy in the measurement. There are several data points indicated in Figure 8.11 at later times that show a significant deviation from the fit to the sine wave. These abnormally large axial widths were not only observed at $a=0$ but also at other fields with a relatively small scattering length. They occurred when the radial width was predicted to be close to

its minimum, suggesting that the radial width was affecting the measurement of the axial width. The large deviation of the data points from the expected fit tends to pull the later peak of the axial oscillation closer to $t=0$, resulting in a frequency determination that would be higher than its true value.

8.7.2 Measurement of the Oscillation Frequencies at Small a . The measured oscillation frequencies at 162 and 164 G were also not observed to match the mean-field predictions. The same distortions described in the previous paragraphs were observed to occur in the measurement of the axial oscillation frequency at these fields. There was no extra jump to a positive scattering length field in order to expand the size of the condensate above the resolution limit in this case, but the other considerations about the condensate evolution at $B \sim 0$ and the condensate shape oscillations still apply. Occasionally, the axial width of the condensate would be much larger than expected at the times when the radial direction width was predicted to be small and the axial width was predicted to be large. The distortion was often severe, and the condensates were observed to be flat rather than rounded at the peak. The fits to these images had little meaning and often these points were not included in the fit. However, we could not be certain that this distortion was not occurring to a lesser extent in the data points that were included in the fit. Figure 8.9 shows a comparison of the measured oscillation frequencies against their predicted values, and there is a clear deviation at values where the scattering length is relatively small.

In contrast, the radial oscillations were much closer to the mean-field predicted frequencies at these fields with smaller scattering length. However, the relatively weaker dependence of the radial oscillation frequency on mean-field interaction energy means that the measurement of the radial oscillation frequencies is not as sensitive a test of the mean-field predictions as measuring the axial oscillation frequencies (see section 8.6). Further, since the radial oscillations are faster than the axial oscillations, the problems associated with any unexpected coupling between the two will not distort the radial direction as much. This is because a complete radial oscillation can occur while the axial direction does not change much.

8.8 Temperature-dependent Shifts in the Radial Oscillation Frequency.

8.8.1 Frequency Shift as a Function of Time at High Loss and Heating Fields. As discussed in section 8.6, we found that as the condensates heated up and melted, the measured radial oscillation frequency changed. These temperature-dependent frequency shifts were significant. For instance, the change in radial frequency due to the heating at 155.7 G was enough to completely obscure the beyond-mean-field shift in our initial measurements. Another way that these temperature-dependent shifts can be characterized is shown in Figure 8.12. The data in this figure represents the average of several different data sets. The observed radial oscillations were naively fit with a sine wave whose frequency did not vary with the time during the oscillation. Only the early data could be fit, or the first half of the data, or all of the data. If the frequency is not shifting as a function of time then these different ways of fitting the data should agree with each other. However, we found that as successively later data sets were included in the fit, the radial frequency determined from the fit decreased. This indicates that the frequency of the radial oscillation itself was decreasing over time.

8.8.2 Measurement of the Condensate Heating Rate. Figure 8.13 shows the measured condensate fraction (i.e. the number of atoms in the condensate divided by the total number of atoms in the condensate and in the thermal component) c as a function of time at the same field at which the data in Figure 8.12 were taken. The fact that the thermal component was increasing over the same time scale as the radial oscillation frequency was shifting led us to associate the observed frequency shifts with the condensate heating. The condensate fraction was measured by performing the usual ramp to excite the oscillation, waiting at the final field for some selected time, and then jumping quickly to a small scattering length. This causes the condensate to shrink significantly in size while any thermal component should not shrink as much. The spatial extent of the condensate and thermal components are then different and by then turning off the trap and imaging the sample, it is possible to measure the ratio of the thermal to condensate fraction by fitting the absorption image. A fitting function consisting of the sum of two two-dimensional gaussians, each with its own axial and radial widths and peak height, was fit to the absorption image. The wider component then corresponded to the thermal fraction while the narrower component corresponded to the condensate. Using this




Figure 8.12. This figure shows the radial oscillation frequency measured at 156.1 G as a function of the time over which the oscillation was measured. For instance, if the radial oscillations were measured over times spanning 450 ms, the measured frequency was 35.27 Hz, while if the oscillations were only measured over times spanning 225 ms the frequency was measured to be 35.77 Hz. The decrease in measured frequency with increasing integration time indicates that the frequency of the oscillation is decreasing as the condensate oscillates at 156.1 G.

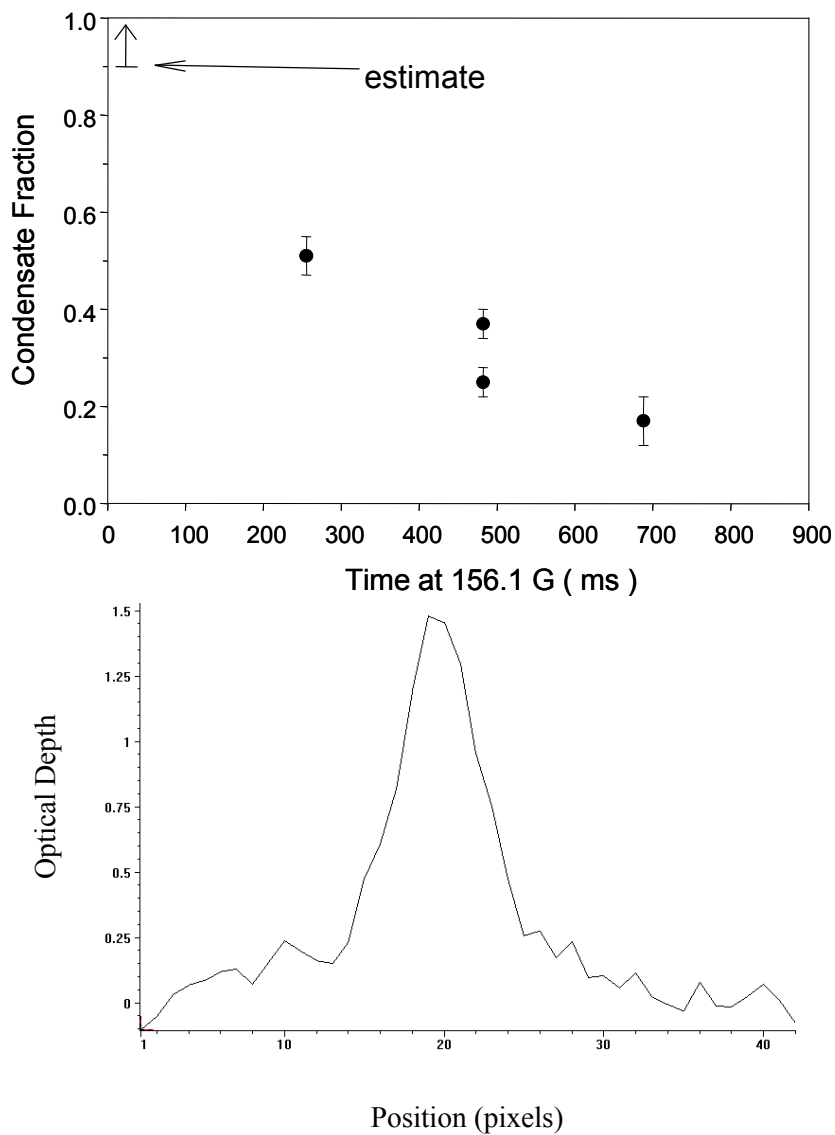


Figure 8.13. Measurement of the melting of the condensate at 156.1 G. The upper plot shows the measured condensate fraction as a function of time that the condensate spent at 156.1 G. No thermal component was originally visible, and so we estimate that the condensate fraction must therefore be greater than 90%. The lower plot shows the cross section of an absorption image used to measure the condensate fraction. The narrow component corresponds to the condensate while the broader base is the thermal component.

method to measure the condensate fraction as a function of delay time at the final field, we found that the condensate was melting since the thermal component grew with increasing delay time.

At the time that these data were taken, we were unaware of the explosion of atoms that could occur from a condensate collapse when the magnetic trap was turned off and the field went to approximately zero (see subsection 10.3.7). Also, the condensates in the images were small enough that the OD ceiling and resolution limit (see section 4.12) could have been distorting the measured condensate number as well. Therefore, the data in figure 8.13 likely possess some errors due to these effects. These errors should not significantly affect any of the conclusions made in this chapter.

The heating rate was found to increase both with increasing scattering length and increasing number of atoms that were originally in the condensate. The origin of this heating rate is somewhat mysterious since it can cause most of the atoms to leave the condensate and go into the thermal component without a very large loss of atoms from the atom cloud. For instance, at 156.1 G, in the time that it takes for the sample to heat up so that only 20% of the atoms are still in the condensate, less than 30% of the atoms have been lost from the sample. The atomic thickness of the condensates with large scattering length may contribute to this heating rate [133].

8.8.3 Evidence for Coupling between the Condensate and Thermal Cloud Motions.

From the data in Figure 8.13, it can be seen that by 500 ms delay time at 156.1 G a substantial fraction of atoms are in the thermal cloud. At the same time in Figure 8.12, the radial direction of the cloud is still observed to be oscillating. This means that it is likely that both the thermal component and condensate component are oscillating in the trap. The thermal atoms feel a factor of two higher mean-field repulsion than condensate atoms due to the difference in the correlations between the position of two thermal atoms and two condensate atoms. This not only tends to repel the thermal atoms out of the center of the condensate, it also means that the oscillation of the thermal cloud and the condensate are coupled to each other.

We found evidence of this coupling in a measurement of the radial oscillation frequency at 157.1 G, shown in Figure 8.14. In this measurement, a shift in the frequency as a function of time at the final field was observed, but it did not uniformly decrease. At later times, the frequency started




Figure 8.14. Radial oscillation measurement at 157.1 G. The measured widths are suggestive of the presence of a beat note in the oscillation, which would indicate that two frequencies were present in the oscillation.

to increase with time. Looking at the oscillation itself, we observed that the amplitude of the oscillation decreased and then increased again. This behavior is indicative of the existence of a beat note between two frequencies, indicating that there were two components oscillating in the radial direction. We postulate that these two components are the condensate and the thermal component.

At the large positive scattering length field, there was no clear separation in sizes between the thermal and condensate component observed. It is still possible to determine the overall width of the cloud precisely, but attempting to measure the individual widths of each component in the cloud is far less precise. Therefore, it was not possible to trace the oscillations of the condensate and the thermal component individually. Since the thermal component fraction is not constant, the radial oscillation shown in Figure 8.14 cannot be fit to determine the frequency of the individual components in any simple way.

It is probably important to point out at this point that we really have two separate sets of observations. The frequency of the radial oscillation is certainly decreasing as a function of time at the large a fields. The fraction of atoms in the thermal component was increasing as a function of the time at these same fields as well. It seems likely that the shifts in the frequency that were observed were due to the increase in the thermal component of the cloud, but no explicit tests were made to link the two, such as deliberately creating a condensate with a thermal component at the end of the evaporative cooling and measuring the radial oscillation frequency.

8.9 Estimation of the Magnitude of the Beyond-Mean-Field Frequency Shifts and Outlook for Future Measurements.

8.9.1 Difficulties in Estimating the Beyond-Mean-Field Shift. Since the frequency shifts due to the melting of the condensates shift in the direction opposite that of the beyond-mean-field shifts, the presence of the temperature-dependent shifts does not change our conclusion that we have observed beyond-mean-field shifts. In other words, the observed deviations from the mean-field prediction in the radial direction are towards higher frequency while the temperature-dependent shifts are observed to be in the direction of lower frequency. Therefore, the temperature-dependent shifts are not responsible for the observed deviation of the measured radial frequency from the mean-field prediction.

As amply demonstrated by the original data point at 155.7 G, however, the temperature-dependent shifts can be the same size as the beyond-mean-field shifts themselves, as so the temperature-dependent shifts make the determination of the magnitude of the beyond-mean-field shifts difficult. In order to determine the beyond-mean-field shift value at $T=0$ to test equation (8.16), it is necessary to extrapolate the radial oscillation frequencies measured at finite temperatures to $T=0$. This extrapolation has two important difficulties associated with it. First, the functional form of the frequency shift as a function of temperature is not known. This of course complicates an extrapolation to $T=0$. Second, it was not possible for us to observe a thermal component of less than 10% of the total atoms in the cloud with the signal-to-noise that we had at the time. Since no thermal component was observed initially, we could not measure the initial temperature other than to say it was below a certain value (the value that corresponded to a thermal fraction of 10%). Because of the steep scaling of the thermal fraction with temperature (see equation (3.6)), a thermal fraction of 10% corresponds to a temperature of $T/T_c \sim 0.5$ in the non-interacting limit. Since we did not know the initial temperature, it was also unclear how close to $T=0$ the samples were and therefore how much of an extrapolation was required.

In principle, the functional form of the radial oscillation frequency as a function of temperature could be calculated or measured experimentally. In Ref. [151], a method to calculate the excitation frequencies at finite temperatures was developed. However, the predictions do not compare well with our data. For instance, for our trap parameters and the scattering length at 156.1 G ($\sim 3000 a_0$) the predicted shift is on the order of 20% for a thermal fraction of 70% [155]. By comparing the measured radial oscillation frequencies of the data shown in Figure 8.12 at short times to the data taken at times between 360 and 800 ms delay times (thermal fraction $\sim 70\%$ as determined from the data in Figure 8.13), a shift of only 3(1)% was observed. There could be several reasons for this discrepancy: the dynamic coupling between the thermal and condensate components is not considered in Ref. [151]; the prediction is made in the thermodynamic limit and so effects due to the finite number of atoms in the condensate are not considered; and the assumptions made in the calculation may not be applicable to our system either due to physics associated with the Feshbach resonance or the inadequacy of using only the leading order correction to calculate the thermal shifts.

In any case, this calculation is not a reliable way of extrapolating to $T=0$. It should be noted that Ref. [151] does match the observed temperature dependence of the axial oscillation mode in ^{23}Na [145].

In an absence of a theoretical determination of the scaling of the temperature shift with temperature, the scaling must be measured experimentally. There are several complications in such a determination. First, we could not measure thermal fractions smaller than 10% and so it was not possible to measure the magnitude of any shifts from $T=0$ to the initial temperature. In addition, the initial temperature could only be measured to be smaller than the limit implied by a 10% thermal fraction. The precise measurement of an oscillation frequency is easiest if the oscillation is measured over several oscillation periods, but during these oscillations the sample heats up making it difficult to measure the oscillation frequency in a narrow range of temperature without significantly more data. The observed coupling between oscillation in the thermal and condensate component also leads to several complications itself. The oscillation in the widths measured at longer delay times is not a measurement of just the condensate widths but some combination of the condensate and thermal widths. A frequency shift may be due to an actual change in the condensate oscillation frequency or it may only be an apparent shift caused by a beat note between the condensate and thermal components whose relative number of atoms changes as a function of time. Finally, the frequency shift may not be only a function of the temperature, as defined by the thermal fraction, but the shift may depend on both the temperature and the state of oscillation of the thermal component.

Even if the functional form of the temperature-dependent frequency shift could be calculated or measured at relatively large temperatures ($T/T_c > 0.5$), it is unclear how far an extrapolation back to $T=0$ should be carried. Partly this is due to the fact that the initial thermal fraction cannot be measured more accurately than establishing that it is below 10%. It is also due to the fact that the temperature should probably not be extrapolated all the way back to $T=0$. If the chemical potential (μ) is much greater than $k_B T$ where T is temperature of the cloud, then the thermal effects should be unimportant and the oscillation frequencies should not depend on the temperature of the sample as long as this condition holds. It is then necessary only to extrapolate back to the temperature where the oscillation frequency no longer scales with frequency, but it is

unclear precisely what this temperature is. It can be determined from the calculations in Ref. [151] to be the point where $k_B T = 0.3 \mu$, but it is unclear how well this calculation applies to our system for the reasons listed above.

8.9.2 Estimate of the Magnitude of the Beyond-Mean-Field Shift. With the above limitations in mind, it is possible to estimate the limits on the beyond-mean-field shift given the data that have already been taken.

Before detailing how such an estimation is performed, there is the question of what data will be most useful in performing the estimation. The measurement at 155.7 G has the virtue of possessing a large beyond-mean-field shift. However, the observed thermal fraction is $\sim 20\%$ after just two radial oscillations meaning that the temperature of the sample is changing relatively rapidly within a single radial oscillation, making the quantitative determination of the oscillation frequency as a function of temperature at this field difficult. Of more interest is the data at 156.1. In this case, there is clearly a beyond-mean-field shift present if only data from relatively short delay times are considered in the radial oscillation frequency determination. There are also data that show the decrease in measured frequency at longer delay times.

The first step in performing the estimation is to determine the change in the thermal fraction of the clouds as a function of delay time. The data in Figure 8.13 are fit using a second-order polynomial with two different intercepts at zero delay time. One intercept is 0% thermal fraction and the other is 10% of thermal fraction. These two values span the possible range of the initial thermal fraction of the cloud.

The change in the observed radial oscillation frequency as a function of delay time must be determined for both possible values of the intercept at zero delay time. As discussed in the previous section, the functional form of the frequency shift with thermal fraction that really should be used is unknown. However, the radial oscillation data exhibited a sinusoidal oscillation at all the delay times that were measured within the signal-to-noise of the individual width measurements without large changes in the measured amplitude or frequency of an individual oscillation. Also, the total frequency shift was only a few percent. This suggests that it is reasonable to assume that the frequency shift can be treated as a perturbation to a sinusoidal oscillation with a fixed frequency.

Therefore, the estimation of the frequency shift is probably not very sensitive to the choice of the fitting function as long as the fitting function is reasonable. The fitting function used is

$$d\omega_{thermal} = \alpha f_{thrm}^n \quad (8.22)$$

where α and n are the parameters varied in the fit and f_{thrm} is the thermal fraction. From equation (3.6), the temperature of the cloud scales as f_{thrm} to the 1/3 power. This relation between f_{thrm} and the temperature is strictly correct only for the non-interacting limit, but it should still approximately describe the temperature of the cloud as a function of the thermal component even in the presence of strong atom-atom interactions. The fitting function given by equation has the virtue of being simple and allowing the rate of change of the frequency shift to change as a function of temperature in the cloud.

Figure 8.15 shows the result of fitting equation (8.22) to the data from Figure 8.12. From these fits the oscillation frequency at low temperatures can be determined by extrapolating back towards $T=0$. There are two temperatures that can be considered for the point to which the extrapolation should be taken, $T=0$ and the temperature at which the radial oscillation frequency is flat as a function of temperature (see the discussion in the previous subsection). Estimating the condition for the frequency to be flat as a function of temperature to be $\mu/3$, the thermal fraction that corresponds to this temperature would only be 2%, so there is little difference between these two points for this field. The important uncertainty in the extrapolation then becomes the initial thermal fraction in the cloud. If the initial thermal fraction was truly zero, then no extrapolation is required. Because the temperature-dependent shift is only in one direction, there is an asymmetry in the uncertainty of the intercept. The central value of the zero delay time frequency is $36.0^{+0.7}_{-0.2}$ Hz in the case that the initial temperature is zero. For the case where the initial thermal fraction was 10%, the fit yields $36.2^{+2.7}_{-0.4}$ Hz. Extrapolating back not to $T=0$ but to the temperature equivalent to $\mu/3$ ($T/T_c = 0.26$ in the non-interacting limit, thermal fraction 2%) yields $36.2^{+2.3}_{-0.4}$ Hz. It should be noted again that the plots in Figure 8.12 and 8.15 are not really frequency versus time, but apparent frequency determined from a single sine wave fit as a function of the range of the delay time used in the measurement. This does not affect the determination of the intercept point, which is the quantity of interest here.




Figure 8.15. Estimation of the beyond-mean-field shift in the presence of temperature-dependent fits. The data shown here are the data also displayed in Figure 8.12. The initial thermal fraction was assumed to be 10% for the fits shown here. The fitting function given in equation (8.22) is shown as the solid line, which is the best fit to the data, and the dotted line which represents the fit corresponding to the value of the intercept that is one standard deviation away from the best fit value. The parameter n is ~ 0.33 for the dotted line and so it corresponds to a temperature shift that changes linearly with the temperature of the cloud.

These estimates of the beyond-mean-field shift in the absence of any temperature-dependent shift can be compared to the predicted value from equation (8.16). The predicted value is $36.25^{+0.62}_{-0.25}$ Hz. The uncertainty in this predicted value is due to the uncertainty in the precise location of the Feshbach resonance peak (B_{peak} , see Chapter V). The measured beyond-mean-shift-value and the predicted value are then in agreement at this field, although the uncertainty in the measured value would also be consistent with shifts a factor of ~ 4 larger than equation (8.16). The uncertainty is only in the direction of higher oscillation frequencies, however, so the measured radial oscillation frequency is clearly not consistent with the mean-field predicted value.

The measured oscillation frequency at 158.7 G, 35.20(7) Hz, shows no shift from the mean-field predicted value of 35.19 Hz, even though the beyond-mean-field predicted frequency is 35.41 Hz. The heating rate at this field is much less, such that the thermal fraction was measured to be 50% only after 3 seconds of condensate evolution at 158.7 G. Unlike the data at 156.1 G, the data at 158.7 G did not show any clear shift as with increasing time of a range of delay times spanning up to 500 ms. Does this suggest that there is no beyond-mean-field shift at this field even though one is expected? Not necessarily. The data taken were not extended to delay times long enough to detect a clear shift in the oscillation frequency and so the analysis applied at 156.1 G can not be directly applied here. The lack of an observed frequency shift as a function of time may only be due to the fact that the thermal fraction changes much more slowly over the course of the radial oscillation frequency measurement than was the case at 156.1 G. More data would therefore be necessary in order to determine whether a beyond-mean-field shift was present at this field or not. From the magnitude of the temperature shifts determined from the fits to the 156.1 G data, a 0.2 Hz temperature shift for thermal fractions of 15% and less is not unreasonable.

8.9.3 Improving the Determination of the Magnitude of Beyond-Mean-Field Shifts with Additional Measurements. The large uncertainty in the determination of the magnitude of the beyond-mean-field shift is unsatisfactory for a precision test of the predicted frequency shifts due to beyond-mean-field effects. Also, there are still some questions concerning the validity of the data given the disagreement with the predicted oscillation frequencies as low values of a . This last subsection contains suggestions for future data that could improve this measurement.

In order to have more confidence in the experimental determinations of the oscillation frequencies, it would be good to confirm the source of the deviations from the mean-field predictions at low value of a . Removing the expansion time after the magnetic trap has been turned off will introduce an error in the measured number, but it can be used to test whether a collapse induced by the trap turnoff is causing any of the systematic shifts. The fact that the lifetimes of the condensates at fields with low a are relatively long means that low-amplitude parametric drives could be used to induce the oscillations in the condensate. Inducing oscillations in this way instead of jumping rapidly between fields could partially or completely remove systematic shifts associated with changes in the shape of the condensate during the oscillation since the ground state shape of the condensates would be much more constant during the creation of the oscillation.

Improving the measurement of small thermal components would help reduce the uncertainty due to the extrapolation back towards colder temperatures significantly. There have been a number of improvements in the apparatus since the data taken in this subsection were recorded, and it may be possible to extend the measurement sensitivity to smaller numbers of atoms in the thermal component.

More data to reduce the measurement uncertainties at 156.1 G and other fields near it could improve the measurement markedly. A reduction of a factor of two in the uncertainties in the data shown in Figure 8.12 could potentially reduce the upper limit on the determination of the beyond-mean-field shift from 2.7 Hz to 0.5 Hz. Also, the data shown in Figure 8.12 were not taken in a manner optimally designed to measure the temperature-dependent shift and extrapolate back to low temperatures, so more efficient data collection is most likely possible. The analysis of the temperature-dependent shift presented here is simple, and more sophisticated methods of fitting to determine the temperature-dependent shift could also be employed to reduce the uncertainty in the measurement and also reduce any systematic errors. In particular, rather than fitting to the apparent frequency shift determined from the oscillation data by simple sine waves, a fitting function that included the frequency shifts could be fit directly to the data themselves. The number variability and the fact that several data sets were included in the data shown in figure 8.12 make such a more sophisticated analysis difficult to apply to the data taken so far.

The coupling between the thermal and condensate components has the potential to complicate the determination of the temperature-dependent shifts. When measuring any oscillation frequencies in the future, some attention should be paid to the appearance of any beat notes or other indication of the presence of two frequencies. As long as large distortions from a sine wave are not observed, it is probably adequate to treat the temperature-dependent shifts as a perturbation as discussed in the previous subsection.

In order to reduce the uncertainty in the predicted values of the beyond-mean-field shifts, a more precise determination of the location of the Feshbach resonance peak would be desirable. This could potentially be accomplished through a more precise measurement like that outlined in Chapter V, but the high loss rates and the large scattering lengths (and hence complications due to the unitarity limit) will complicate more precise thermal equilibration measurements. Like the improved determination of B_{zero} , there may be some properties of condensates near the Feshbach peak that allows them to be used to determine its location more precisely and so additional experiments near the Feshbach resonance peak may remove this difficulty.

Finally, the ability to intentionally create hotter clouds after evaporation was not used in the analysis of the data taken so far. By measuring the oscillation frequencies with hotter clouds, the importance of the dynamics of the condensate and thermal component coupling could be tested. Also, it would provide an alternative way to determine the temperature-dependent shifts at fields such as 158.7 G where the heating rate is low but the initial temperature may be high enough to distort any measurement of beyond-mean-field shifts.

Chapter IX

Controlled Collapse of Bose-Einstein Condensates with Attractive Interactions

9.1 Attractive Interactions and Bose-Einstein Condensate Stability.

9.1.1 Bose-Einstein Condensate Mechanical Stability. In the past two Chapters, the mean-field energy of the condensates that were studied was positive since the s-wave scattering length between atoms in the condensate was itself positive. Since the mean-field energy is proportional to the density of atoms in the condensate (see equation (3.20)), with positive scattering length the mean-field interaction acts to repel the atoms from one another because lowering the density of the condensate lowers its energy as well.

Conversely, when the s-wave scattering length is negative, the atoms attract rather than repel each other. Therefore, an increase in density of the condensate lowers its energy. This has a destabilizing influence on the condensate itself, which is predicted to prevent the formation of stable condensates in the homogeneous case when any attractive interactions are present. Long-wavelength fluctuations in the condensate that result in an increase in the density of the condensate lower its energy (short-wavelength excitations have more kinetic energy and so the reduction in energy due to the increase in density is offset by the kinetic energy increase). Consequently, these long-wavelength fluctuations can grow with time and the condensate cannot then exist in a stable state.

In the inhomogeneous case, the spatial extent of the confining potential adds another spatial and energy scale that needs to be considered in the stability of the condensate. The confining potential cuts off long-wavelength excitations and can therefore stabilize a condensate with attractive interactions, as long as the magnitude of the attractive interactions is not too large. Considering the energy of the condensate as a function of its spatial extent illustrates this stabilization. There are three contributions to the energy of a condensate in the inhomogeneous case: the kinetic energy, the potential energy due to the confining potential, and the mean-field energy. For a fixed number of

atoms, the potential and mean-field energy of the condensate will decrease as the spatial extent of the condensate decreases, while the kinetic energy increases as the spatial size decreases since the condensate wavefunction must then possess more curvature. In the absence of any mean-field interactions, a balance between the kinetic and potential energy determines the spatial extent of the condensate in the ground state. If the spatial extent of the condensate is increased or decreased from the ground state in this case, then the energy of the condensate increases. When attractive mean-field interactions are included, then the spatial extent of the ground state will decrease from the non-interacting value. If the mean-field interactions are not too strong, this decrease can still be balanced by the kinetic energy of the condensate and the condensate will still be stable. However, because of the energy dependence on the spatial extent of the condensate is steeper for the mean-field energy than the kinetic energy, for strong enough mean-field interactions the kinetic energy can be overwhelmed so that it can no longer stabilize the condensate.

These arguments are illustrated in a specific example. If the shape of the condensate order parameter in a spherically symmetric harmonic potential is fixed to be a gaussian with an rms width R , then the kinetic, potential, and mean-field energy can be computed as a function of R using equation (3.15). For simplicity, the length scale used in equation (3.15) for the total energy can be rescaled so that one unit of length is equal to the simple harmonic oscillator length $a_{ho}=(\hbar/m\omega)^{1/2}$ so that

$$E = \hbar\omega \left\{ \int d^3r \Phi^*(r) \left(\frac{-\nabla^2}{2} + \frac{r^2}{2} + \frac{2\mathbf{p}Na}{a_{ho}} |\Phi(r)|^2 \right) \Phi(r) \right\} \quad (9.1)$$

where spherical symmetry has been assumed both in the equation itself and the definition of a_{ho} .

Also, the normalization of Φ is set so that $\int d^3r |\Phi(r)|^2 = 1$. With $\Phi(r)=\exp(-r^2/2R^2)/(\pi R^2)^{3/4}$, the energy as a function of R when a is negative is

$$E = \frac{3}{2} \hbar\omega \left(\frac{1}{2R^2} + \frac{R^2}{2} - \sqrt{\frac{2}{\mathbf{p}}} \frac{N|a|}{a_{ho}} \frac{1}{3R^3} \right) \quad (9.2)$$

where the first term is the kinetic energy, the second the potential energy, and the third the mean-field energy. Several plots of the energy as a function of R for different values of $N|a|/a_{ho}$ are shown

in Figure 9.1. As long as the ratio $N|a|/a_{ho}$ is smaller than 0.671, there is a minimum in the energy as a function of the spatial extent of the condensate. For values of the ratio $N|a|/a_{ho}$ larger than that, there is no minimum. The mean-field interactions will be strong enough to overcome the kinetic energy, and so any condensate with this fixed shape is predicted not to be stable for any value of the rms width R .

9.1.2 Theory Predictions for Condensate Stability. Several predictions have been made concerning the point at which condensates in harmonic potentials are no longer stable. All of these predictions are computed in mean-field theory, and can be expressed as a stability criterion

$$\frac{N_{crit} |a|}{a_{ho}} = k \quad (9.3)$$

where k is a dimensionless number. For any condensate whose product of the number of atoms in the condensate and the magnitude of the s-wave scattering length ($N|a|$) is larger than ka_{ho} there is no stable ground state solution of the mean-field (i.e. Gross-Pitaevskii) equation describing the condensate.

For condensates whose density is fixed to have a gaussian shape, k is calculated to be 0.671 as per the discussion in the previous subsection. The shape of condensates with attractive mean-field interactions is not described exactly by a gaussian, however. In order to predict the stability criterion more precisely it is therefore necessary to solve for the ground state of the condensate without fixing the shape. Ruprecht *et al.* [135] computed the ground state as a function of the ratio $N|a|/a_{ho}$ by integrating the time-dependent GP equation, equation (3.19). In this treatment, the ratio $N|a|/a_{ho}$ was set to be zero at time $t=0$, and the condensate was set to be initially in the ground state. Since there were no interactions at $t=0$ it was trivial to set the condensate ground state. After these initial conditions had been established, the ratio $N|a|/a_{ho}$ was increased adiabatically slowly as a function of time in the direction of more attractive interactions. The condensate would therefore remain in the ground state and by integrating the time-dependent GP equation the ground state as a function of ratio $N|a|/a_{ho}$ could then be computed. It was found that the ground state energy had an asymptotic cutoff at $C_{nl}^{3D} \approx -1.62$, where the nonlinear interaction parameter is defined as




Figure 9.1. This figure shows the energy of a condensate whose density profile is fixed to be a gaussian as a function of the width of the condensate R (x-axis) and the value of the ratio $N|a|/a_{ho}$ (number label on each curve). Note that for the curve with $N|a|/a_{ho}$ equal to 0.7 there is no minimum in the energy as a function of R . The energy was calculated from equation (9.2).

$C_{nl}^{3D} = 2Na(2m\mathbf{w}/\hbar)^{1/2}$ in Ref. [135]. Therefore, the asymptotic cutoff in ground state energy corresponds to a value of k in equation (9.3) of 0.573. This calculation was performed using a spherically symmetric potential.

Several other groups have confirmed the prediction of the stability coefficient obtained in Ref. [135]. All of these other methods used the GP equation as the starting point, but different methods were used in order to calculate the stability coefficient k . In Ref. [156], stable solutions of the GP equation ground state could not be found once the number of atoms in a condensate with a fixed negative scattering length exceeded a value consistent with the prediction in Ref. [135]. Ref. [157] calculated the ground state by minimizing the energy using the steepest descent numerical method. The calculated stability coefficient was higher than Ref. [135] ($k = 0.638$) but the authors claimed that the two predictions were consistent. A variational method was used to find the ground state in Ref. [158], and a numerical integration of equation (3.20) to find the ground state energy in Ref. [159]. Both of these calculations produced the stability coefficient found in Ref. [135]. A different method of solving the time-dependent GP equation [160] found a result consistent with Ref. [135], as did an evaluation of the condensate ground state stability through the calculating the effect of the introduction of small variations around the numerically calculated ground state solution of the GP equation [161]. This list should not be considered to be exhaustive.

All of the above predictions were performed in a spherically symmetric harmonic trapping potential. The steepest descent method was used to minimize the energy of and thus find the ground state of condensates with attractive interactions in Ref. [162], both for the spherical case ($k = 0.545$, close to Ref. [135]), and in the case where the axial trap frequency is much less than the radial frequency. For our radial trap frequency, the stability coefficient in this limit of large aspect ratio was calculated to be $k = 0.58$. In our case the axial trap frequency is only a factor of 2.5 less than the radial and so this limit is not truly satisfied, but the lack of a significant dependence on the ratio of the axial to radial trap frequencies indicate that the predictions for the spherical traps should apply to our axial trap as well at the few percent level. This can be understood qualitatively from considering the condensate behavior as the stability limit is approached. As the mean-field attraction strength increases so that the kinetic energy can barely stabilize the condensate, the condensate tends to be

pulled into a spherical shape. Therefore, it is reasonable to expect that the difference in the trap frequencies should not significantly affect the stability criterion.

In addition to Ref. [162], we have asked Brett Esry to calculate the lowest energy collective excitation frequency as a function of Na/a_{ho} for our trap conditions. One of the predicted characteristics of the condensate instability is the frequency of the lowest energy collective excitation frequency going to zero. The disappearance of a stable ground state is considered to be equivalent to the disappearance in the minimum of the energy of the condensate as a function of its spatial extent. As this minimum disappears, the curvature of the total condensate energy vs. width in at least one direction must therefore also disappear. As that curvature decreases, the collective excitation frequency associated with that direction must therefore go to zero as well. Using the techniques for calculating the collective excitation frequencies in the RPA [41], it was possible to determine the ratio of Na/a_{ho} where the lowest collective excitation frequency was zero. That point corresponds to a stability coefficient $k = 0.548$ [163], again indicating that the predictions for a spherical trap apply to our trap as well at the few percent level. Another prediction [164] for our trap parameters gives a value of 0.552 for k .

One note of caution is in order when comparing our measurement to theory predictions. We define the value of a_{ho} in equation (9.3) as given in equation (3.21) and repeated here

$$a_{ho} = \sqrt{\frac{\hbar}{m(\omega_x \omega_y \omega_z)^{1/3}}} \quad (9.4)$$

where the trap frequency of each direction is given by $\omega_{x,y,z}$. In the literature, the harmonic oscillator length is sometimes defined with respect to just the radial oscillator frequency as $(\hbar/m\omega_r)^{1/2}$, and so some caution is necessary in comparing different predictions to our results and to each other.

Most of the above calculations concern themselves only with the stability of the ground state as a function of the attractive interaction strength and do not consider the dynamics that could be introduced due to a non-infinitely slow change in the attractive interaction strength. The calculation in Ref. [160] suggests that a sudden change from a scattering length of zero to an attractive scattering length should reduce the stability coefficient by 10%. This prediction is not very applicable to the work discussed in this chapter since the scattering length was slowly and not

suddenly changed. The stability of the condensates ramped to negative scattering length fields has also been considered in the PG model (see subsection 8.2.6 for a description of the PG model and Refs. [165,166] concerning the stability calculation). For the typical rate of change of the scattering length used in this work, these calculations using the PG model do not predict any significant change in the stability coefficient (less than 2% reduction).

9.1.3 Dynamics of the Condensate Instability. The question of what happens to the condensate once the stability criterion given in equation (9.3) is violated is a natural one to ask, but the discussion of the dynamics of unstable condensates will be postponed until the next chapter. In general, the mechanical instability of the condensate would suggest that it would start to shrink in size, and it was postulated that the density of the condensate would increase until two- and three-body loss mechanisms reduced the number of atoms in the condensate [156,167]. This has been referred to as the “collapse” of the condensate. Because of these predictions, we initially looked for a loss of atoms from the condensate as a signature of the collapse, and indeed we did observe a loss. Further experiments, however, revealed the collapse dynamics to be much more complicated than such a simple picture. In spite of this, we will still refer to the onset of the instability as the condensate collapse in this and the following chapter.

9.1.4 Previous Work in ^7Li and Comparison to our Work. The first atom with a negative scattering length that was cooled to Bose-Einstein transition temperatures was ^7Li [168], and the maximum number of atoms in the condensate has now been observed to be consistent with the prediction in Ref. [135]. There are several differences between these ^7Li experiments and our work, however. In ^7Li , condensates are formed from a rapidly cooled thermal cloud at a negative scattering length whose value remains fixed, and condensate formation kinetics has been a central study [169]. The condensates grow in number until they reach a critical value and then their number decreases. Following the decrease, the number begins to grow again and the growth and collapse cycle is repeated as atoms are fed into the condensate from the thermal cloud that is present. These cycles are observed to be variable and so a statistical treatment of post-collapse ensembles has been used to examine the ^7Li growth and collapse process [170].

Here, we used the Feshbach resonance in ^{85}Rb to tune the scattering length by changing the magnetic field. After forming stable and nearly thermal-atom-free condensates with positive scattering length, we tuned the scattering length to a selected negative value in the vicinity of the expected critical value to study the instability. This allows us to study the onset of instability and the collapse dynamics in a very controlled manner at a variety of different scattering lengths. In addition to the ability to tune the scattering length and cold initial condensates, there were other advantages in our ^{85}Rb work that do not exist in the ^7Li system. Our ability to form stable condensates allowed us to characterize the statistical uncertainties in our system independent of any variability brought about through the collapse process. Also, the ability to manipulate the mean-field interactions in the condensates allowed us to expand the size of our condensates above our imaging resolution limit to measure the number and shape of the atoms in the cloud, while in ^7Li the imaging had to be done below the resolution limit of the imaging system.

9.1.5 Goals of this Measurement. The measurement presented in this Chapter concerned itself with four questions: how sharp is the transition from stable to unstable condensates as a function of number and scattering length, how many atoms remain after the collapse has occurred, does the onset of collapse depend solely on the $N|a|$ product as in equation (9.3), and if so, what is the value of the stability coefficient k ? We found the transition from stability to instability very sharp, with no width greater than that expected from the statistical uncertainties in our system. The number of atoms remaining after the collapse occurred was typically around 60% of the initial condensate number. The onset of instability did scale with the $N|a|$ product, and the measured value of k was close to, but slightly lower than, the theory predictions. The details of the answers to each of these questions are presented in the following sections in this Chapter, after the procedure used in the experiment is outlined.

9.2 Measuring the Onset of Instability.

9.2.1 Creation of the Condensates to be Studied. Initially, 3.5×10^8 atoms were loaded into the baseball magnetic trap at 45 μK in the $F=2$, $m_f = -2$ state as described in Chapter IV. The condensates used in this measurement were then formed using the evaporative cooling described in Chapter VII at a field of 162.3 G. After the condensates were formed, the field was ramped slowly

(800 ms) to near $a=0$ ($B \sim 165.85$ G). The condensates were then kept at that field for 2-10 seconds with an rf-knife present to let the inelastic losses reduce any number variations as described in subsection 7.1.7.1.

9.2.2 Ramping to $a < 0$ and the Measurement of the Number of Condensate Atoms.

After the condensates had been prepared as described in the previous subsection, the field was ramped to a selected value of $a < 0$. The region of the Feshbach resonance in which we worked was that near $a = 0$, not the region that involved going over the peak of the Feshbach resonance. The magnetic field ramp was performed in 200 ms, slowly enough to be adiabatic for most of the ramp (see subsection 9.6.2), but fast enough that a significant number of atoms will not be lost to inelastic collision loss mechanisms during the ramp. After that ramp, the condensates were allowed to evolve at the $a < 0$ field for 50 ms and would either collapse or not, depending on the condensate number and the value of a . This 50 ms time was increased to 200 ms and no significant differences were observed due to the increase in time at the $a < 0$ field.

While the condensate was at the fields corresponding to $a \sim 0$ and $a < 0$, the spatial size of the condensates was smaller than the $7 \mu\text{m}$ resolution limit of our imaging system (subsection 4.12.7). Therefore, after the 50 ms evolution time, the condensates were expanded in spatial size using the mean-field interaction to measure the number of atoms in the condensate. This mean-field expansion was not only convenient; it was necessary for ^{85}Rb because the usual expansion of a condensate by turning off the magnetic trap [6,171] is not effective for ^{85}Rb . In ^{87}Rb and ^{23}Na , turning off the magnetic trap causes the condensates to expand outward rapidly due to their kinetic and (repulsive) mean-field interaction. In ^{85}Rb , a is $\sim -400 a_0$ at $B \sim 0$ and so with the trap off the mean-field interactions do not drive the condensate outward. The relatively weak magnetic traps and small atom numbers used with ^{85}Rb also mean that the kinetic energy in the condensate is small as well. For instance, the radial rms width of a condensate originally at $a = 0$ in our traps will only increase from $1.9 \mu\text{m}$ to $2.74 \mu\text{m}$ in 20 ms of expansion after the magnetic trap has been turned off. Because the atoms are accelerating due to gravity, they quickly fall away before they expand sufficiently to be imaged above the resolution limit of the imaging system.

The mean-field expansion of the condensate spatial size was performed by ramping the magnetic field from the $a < 0$ field to 157 G ($a \sim +1500 a_0$) in 5 ms and then waiting for 7 ms to let the condensate expand under the influence of the large positive scattering length. After this expansion, condensates that did not collapse typically had rms widths of $\sim 12 \mu\text{m}$, larger than the resolution limit. The magnetic trap was then turned off and the atoms were imaged using absorption imaging (section 4.11) in order to determine the number of atoms in and the shape of the condensate. The number of atoms in the condensate was then calculated using all of the considerations listed in section 4.12.

The selection of the parameters of this expansion ramp was made using a few different considerations. In order to avoid the difficulties associated with a large optical depth and still have a reasonable signal-to-noise, the ramp was used to keep the peak optical depth of the cloud around one. We have found that ramps from $a \sim 0$ to large positive scattering lengths with duration 1 ms or less can introduce loss in the number of atoms for condensates whose density is initially high. Therefore, the ramp duration was set to 5 ms to let the condensate spatial size expand so that this loss was avoided. If the field is ramped too close to the peak, then there is a significant loss in the number of atoms from the condensate due to the losses at fields with large scattering lengths. If the field is not ramped far enough, the clouds do not expand as much and the raw observed number of atoms in the condensate will drop as the optical depth of the clouds increases and the imaging difficulties associated with high optical depth clouds (subsection 4.12.5) becomes significant. The parameters chosen represent a compromise between high loss and high optical depths.

9.2.3 Measuring the Temperature of the Condensates. The temperature of the condensates (i.e. the number of atoms in a thermal component) was measured for a two reasons. First, the stability condition in equation (9.3) is expected to apply to the number of condensate atoms, not the total number of atoms in the cloud (see subsection 9.6.1 below and the references contained therein). Since the total number of atoms is the quantity that is measured by the absorption imaging, the thermal fraction needed to be known in order to know the number of atoms in the condensate. Secondly, we wished to test to make sure that the stability coefficient k was not very sensitive to the temperature of the cloud. This test is described in subsection 9.6.1.

To measure the number of thermal atoms present, the condensates were created, ramped to $a \sim 0$ and left to evolve there just as described in subsection 9.2.1. The magnetic trap was then shut off without using any expansion ramp to increase the size of the condensate. This meant that the condensate was below the resolution limit of the imaging system and that there should be a separation between the spatial extent of the thermal cloud and the spatial extent of the condensate. A gaussian fitting function was then fit only to the wings of the atom cloud (i.e. the central, below-resolution limit condensate was excluded from the fit). From the widths and peak optical depth of the gaussian that had been fit to the wings, the thermal fraction was determined.

There was a complication present when trying to determine the number of thermal atoms that were present. The dense condensates would collapse due to the negative scattering length at $B \sim 0$, ejecting a burst of atoms from the condensate (see subsection 10.3.7). Due to the 1.6 ms time outside of the trap (subsection 4.11.4), these burst atoms would expand to a size greater than the condensate ($\sim 4 \mu\text{m}$ rms) and would appear to be an additional component. It was difficult to separate the central component, the burst component, and any thermal atoms with a high degree of accuracy. After this measurement was completed, an improved method for measuring the number of thermal atoms was developed. Rather than turning off the trap immediately, the condensate was briefly ramped to a positive scattering length field in order to impart enough outward expansion that the cloud did not initially collapse at $B \sim 0$.

For most of the data used to determine the onset of instability for condensates at $a < 0$, we cooled the cloud as best we could and usually we could not observe a thermal component. To establish a limit on the possible number of thermal atoms, we examined the precision of our imaging techniques and estimated that the limit of our sensitivity was 10% of the total atoms in the sample.

9.2.4 Monitoring the Number Variation and Drift. In the absence of any non-destructive imaging, it was not possible to precisely know the number of atoms in the condensate before the field was ramped to the fields with $a < 0$. The initial number of atoms could therefore only be known to within the measured number reproducibility ($\sim 5\%$ at best, 7-10% was common). In order to monitor both the variation of the number of atoms from shot-to-shot and any drift in the initial number of condensate atoms, data points were occasionally taken without ramping to $a < 0$ from a

~ 0 . In every other respect these monitor points were the same as the points with the ramp to $a < 0$. These points were taken often, typical every other or every third data point, in order to tightly monitor any drift in number or increase or decrease in the shot-to-shot variability of the number of atoms in the condensate.

9.3 Signature of the Collapse and Sharpness of the Transition from Stability to Instability.

9.3.1 Signature of the Collapse Event. As long as the product $N|a|$ of a condensate was below a critical value, the condensate showed no effect of an excursion into the region where $a < 0$. Once this critical value was exceeded, the condensates were observed to collapse. There were two signatures of this collapse: the number of atoms was observed to decrease, and the condensates were observed to be oscillating in highly excited states. The number loss was greater than the uncertainty from the variations in initial condensate atom number and so was clearly observable. The shape of the condensates was also often very different from its non-collapsed shape. Occasionally the aspect ratio of a condensate that had collapsed would be much higher than a stable condensate's aspect ratio; sometimes it would be much less. The relatively slow ramp down to $a < 0$ meant that the precise time at which the collapse occurred varied from shot to shot since the number of atoms initially in the condensate was not exactly reproducible. For instance, if the number of atoms initially in the condensate fluctuated upward, the condensate would cross the onset of instability sooner in the ramp and therefore collapse sooner. As the condensates oscillate in the trap after the collapse, the phase of the oscillation will therefore vary from shot to shot and that explains the observed variations in the condensate shapes. Figure 9.2 compares absorption images of non-collapsed and post-collapsed condensate images.

9.3.2 Sharpness of the Transition from Stability to Instability. There are two possibilities for the nature of the probability of a condensate to collapse as its $N|a|$ product nears the stability limit defined in equation (9.3). The collapse could be a binary event such that if the $N|a|$ product is less than the stability limit, the condensate is stable and there is no significant shape change or large loss of atoms. Once the $N|a|$ product exceeds this limit, then the condensate will




Figure 9.2 (Color). Comparison of Stable and Collapsed Condensates at $a < 0$. All of these images were taken with the mean-field expansion ramp described in the main text. Frame (a) corresponds to a stable condensate while frames (b) and (c) show collapsed condensates. All three images were taken under similar conditions, with the only difference being that in (b) and (c) $|a|$ was greater than in (a). The false color scale is not the same from image to image.

collapse. The other possibility is that as the $N|a|$ product approaches the stability criterion, fluctuations of some kind [172] can cause the condensate to occasionally collapse. For instance, the width of the condensate could tunnel through the potential energy barrier shown in Figure 9.1 [173,174], causing the condensate to collapse even though its $N|a|$ product would have suggested that it would be stable according to equation (9.3). In the case where fluctuations can cause a condensate to collapse, the collapse would not appear to be a binary event, but rather the probability of the collapse would increase as the stability criterion was approached. We refer to the change in probability of collapse as a function of magnetic field (i.e. s-wave scattering length) as the sharpness of the transition from stability to instability. If the collapse were a binary event, the probability for collapse would change rapidly as a function of field and the transition from stability to instability would be very sharp.

In order to determine the sharpness of the transition from stability to instability, data were taken maintaining the number of the atoms in the condensate as close to 6400 as possible. Monitor data points described in subsection 9.2.4 were alternated between data points in which the magnetic field was ramped to several selected values around the field where the condensates became unstable and collapsed. The number of atoms in the condensate that remained after a ramp to $a < 0$ was then compared to the number of atoms measured by the neighboring monitor points as a function of field. The results of this measurement are shown in Figure 9.3.

Even if the transition from stability to instability as a function of the scattering length is in actuality infinitely sharp, there would still be an observed width to the transition due to the variation in the initial number of atoms in the condensate. For instance, if the field was ramped to precisely the value corresponding to the onset of instability for the average number of atoms in the condensates in the data set, half of the time the condensates would collapse since the initial number would be higher than average and half the time the condensates would be stable and not collapse since the initial number would be lower than the average. At a field slightly different from the field where the collapse probability is 50%, there will still be a significant probability that the condensates will collapse since the number can fluctuate above the critical number at that field, and also a significant probability that the condensates will be stable since the number could also fluctuate

below the critical number. Therefore there will be a range of magnetic fields where the condensates sometimes collapse and sometimes are stable, and the transition from stable to instability will not appear infinitely sharp, even if it truly is. The larger the variation in the initial number, the larger the apparent width of the transition. The width due to the number variation is quoted in terms of the rms variation in number (σ_{num}) as

$$\text{width} = \sigma_{\text{num}} \cdot \left. \frac{\partial}{\partial B} \left(\frac{ka_{ho}}{a(B)} \right) \right|_{B=B_0}, \quad a(B) = -380 \left(1 - \frac{10.95}{B-154.9} \right) a_0 \quad (9.5)$$

where B_0 is the field where the average number of atoms in the data set would become unstable. For instance, in Figure 9.3 B_0 is 165.96 G. This width corresponds to the spread in magnetic field between the fields where the condensates have a 32% chance of collapsing to the field where there is a 68% chance of collapse.

The onset of the collapse is observed to be as sharp as we can measure it within the variation of the initial number of atoms in the condensate. This means that either the collapse truly is a binary event, or that the probability for the condensate to collapse increases extremely rapidly as a function of scattering length. In either case, the collapse can be treated as a binary event within the precision of our measurement.

The distribution of data at 165.96 G is also indicative of the collapse being like a binary event. These data, and other data similar to it, fall into two separate groups, with one group consistent with the stable condensates at fields with lower values of $|a|$ and the other group consistent with the collapsed condensates at fields with higher values of $|a|$. This is what would be expected if the collapse were a binary event. The variation in the number of atoms initially in the condensate would cause the condensates with an atom number that was higher than the average to collapse, and the condensates whose number was lower than the average to be stable.

9.4 Fraction of Atoms Remaining in the Condensate After Collapse.

9.4.1 Fraction Remaining after Collapse. From the data in Figure 9.3 and other data sets like it, the fraction of atoms remaining after a condensate with an initial number of ~ 6400 collapses can be measured. That fraction is 0.58(3). When the starting number of atoms in the condensate was changed to ~ 2600 atoms, the fraction remaining was 0.76(3). These fractions are measured only




Figure 9.3. Transition from stable to unstable condensates. This figure shows the fraction of atoms remaining as the magnetic field was ramped to higher magnetic fields (i.e. stronger attractive interactions). The hatched region indicates the expected 4 mG width due to the initial number variation. The initial number of atoms in the condensate was kept near 6400. At the field in the hatched region, the data fell into two groups and the individual points (\square) are displayed rather than the average of those points.

when the magnetic field is tuned just past the onset of instability. As the value of a is made to be more negative, the fraction of atoms remaining decreases.

9.4.2 Loss of Condensate Atoms at $a < 0$ in the Absence of Collapse and Unusual Statistics. The all-or-nothing nature of the collapse led to complications in the determination of the fraction of atoms that were remaining, both after the collapse and after an excursion into the $a < 0$ region in the absence of a collapse. Near the onset of the instability, the condensates whose initial number was greater than the average due to an upward variation in the initial number will collapse. Therefore, the average number of atoms in the condensates that are stable at a field near the onset of instability will be lower than the average number of atoms initially in the condensate before the ramp to the $a < 0$ region.

In order to determine if any atoms are being lost due to the ramp to fields with $a < 0$ even in the absence of the collapse, the probability distribution of the number of atoms in stable condensates at $a < 0$ is modeled as a truncated gaussian. This idea is illustrated in Figure 9.4. The critical number used to truncate the gaussian is determined in the same way as the stability coefficient is measured as described below in section 9.5. The ratio (ξ) of the average number of condensate atoms in the truncated probability distribution to the average number of condensate atoms in the full probability distribution can be calculated as a function of the critical number

$$\mathbf{x} = \frac{\int_{-\infty}^{N_{crit}} N \exp\left(\frac{-(N - N_{avg})^2}{2\sigma_{num}^2}\right) dN}{N_{avg} \int_{-\infty}^{N_{crit}} \exp\left(\frac{-(N - N_{avg})^2}{2\sigma_{num}^2}\right) dN} \quad (9.6)$$

where N_{avg} and σ_{num} are the average number and the standard deviation of the full probability distribution, respectively, and N_{crit} is the critical number ($N_{crit} = ka_{ho}/a$). Setting the lower end of the integration to $-\infty$ is a matter of mathematical convenience and does not affect the calculation of ξ .

The loss of atoms due to a ramp to $a < 0$ in the absence of a collapse is then calculated in the following way. Monitor points (subsection 9.2.4) were used to determine the standard deviation of the full (initial) probability distribution of the number of atoms in the condensate (i.e. N_{avg} and σ_{num} in equation (9.6)). From this measured standard deviation and the critical number, the ratio ξ




Figure 9.4. This figure shows the truncated gaussian probability distribution used in analyzing the fraction of atoms lost after a ramp to fields with $a < 0$. The average of the initial number of atoms in the condensate before the ramp to $a < 0$ and the standard deviation of those points determine the location of the peak of the width of the overall gaussian distribution. The probability distribution of the number of atoms in stable condensates is shown by the solid curve, while the probability distribution of the pre-collapse number of atoms in condensates that did collapse is shown as the dashed line. The critical number N_{crit} depends on the value of $|a|$ and therefore varies as a function of magnetic field. Both the stable and collapsed condensate distributions are normalized appropriately before being used in a calculation.

can be calculated for each field in the measurement. The number of atoms in a stable condensate is then divided by the value ξ appropriate for its field. These adjusted numbers are then averaged together, and their average is compared to the average of the number of initial atoms in the condensates (i.e. the average of the number of atoms in the monitor points). This procedure corrects for the tendency of stable condensates to have a smaller number than the average number of atoms. When this procedure is followed, the fraction of atoms that remain after a ramp into $a < 0$ without a collapse was found to be 0.98(1). The small amount of loss that may be present (2%) is consistent with the loss of atoms from the condensate due to inelastic collisions during the 200 ms ramp to the $a < 0$ field. This ratio is the same both for ~ 6400 atoms and also for smaller (~ 2000 and less) condensates.

The same considerations were used in the computation of the fraction of atoms remaining after the collapse. In this case, the initial number of atoms in condensates that collapsed tends to be higher than the average number of initial atoms near the onset of instability. In order to compute the average number of atoms in the condensates before the collapse, the same truncated gaussian used for the stable condensates was used. The average initial number was then computed as

$$\bar{N}_{collapse} = \frac{\int_{N_{crit}}^{\infty} N \exp\left(\frac{-(N - N_{avg})^2}{2s_{num}^2}\right) dN}{\int_{N_{crit}}^{\infty} \exp\left(\frac{-(N - N_{avg})^2}{2s_{num}^2}\right) dN} \quad (9.7)$$

where the quantities in equation (9.7) are defined in the same way as in equation (9.6).

9.4.3 Increased Variation in the Collapsed Number. The fractional variation of the number of atoms remaining after the condensate collapse is greater than the variation in the initial number of condensate atoms measured with the monitor points. Using the data shown in Figure 9.3 and several other data sets similar to it with an average number of ~ 6400 , the standard deviation of the initial condensate number was 7% of the average condensate number, while the comparable fractional standard deviation of the number of atoms measured in the collapsed condensates was 20%. Both of these variations include the $\sim 5\%$ imaging uncertainty (subsection 4.12.3) as well as the variation of the actual number of atoms in the condensates. For the 2600 atom case the

variations were 9% for the initial number of condensate atoms and 20% for the number remaining in the collapsed condensates.

There is a possibility that this observed increase in the number variation for collapsed condensates is an artifact of the absorption images, but we do not believe that this is likely the case. The collapsed condensates shapes were different than the pre-collapse condensates, but the collapsed condensate shapes were still larger than the resolution limit. Also, the number was calculated both by summing over the optical depth measured in each pixel in the absorption image and fitting the condensate images using a 2D gaussian fitting function. The two methods agreed on average, so the different shapes should not have distorted the measurement of the number. Therefore, it does not seem likely that the number variability was due to systematic effects in the calculation of the number from the absorption images of the collapsed condensates. The increased variability then indicates that the collapse process is not the same every time, but varies from collapse event to collapse event.

9.5 Determination of the Stability Coefficient.

9.5.1 Measurement of the Onset of Instability as a Function of the Number of Atoms in the Condensate. To examine the validity of the theory prediction in equation (9.3), we varied the number of atoms that were initially in the condensate and determined the minimum magnetic field at which that number of atoms would collapse. Since the scattering length varies approximately linearly with magnetic field for the negative scattering lengths of interest, equation (9.3) would predict that a plot of $1/N_{\text{crit}}$ vs. the collapse field would be a straight line. The slope of this line would then determine the stability coefficient k .

The drift and random variation in the initial number complicated this measurement since it was impossible to know the exact number of atoms that were in the condensate before the ramp to $a < 0$. The drift in the initial number of atoms in the condensate and the random variation were determined by frequently interspersing monitor points described in subsection 9.2.4 among the data points with ramps to fields with $a < 0$. After a ramp to $a < 0$, the condensate was classified as either collapsed or non-collapsed based on its shape and the number of atoms in that condensate compared to the initial number estimated from both the non-collapsed data points and the monitor points. If either the shape was different than usual for the non-collapsed data points or there was significant

loss of atoms from the condensate, the condensate was classified as collapsed. Two days' negative a data are shown in Figure 9.5 (a) and 9.5 (b). Figure 9.6 shows a summary of these two days' data and two additional sets that were taken in a similar fashion. From the data shown in Figures 9.6 and 9.7, it is evident that the onset of instability scales well with the product $N|a|$.

For the data in figs. 9.5 (a) and (b), the field was measured using the baseball and bias currents through the calibration described in subsection 4.10.1. The assignment of the number to each data point varied depending on whether the data point had collapsed or not. If the data point had not collapsed, then the pre-collapse number of atoms in that condensate was simply equal to the number of atoms measured by the absorption imaging for that data point. If the data point had collapsed, then the initial number could not be best determined from the remaining number of atoms since the fractional variation of the number of atoms remaining after a collapse is greater than the variation of the pre-collapse number of atoms in the condensate. Instead, the number of atoms present in the condensate before the collapse was estimated by averaging neighboring monitor or non-collapsed data points. If a non-collapsed data point was involved in the initial number estimation, a correction was applied due to the propensity of the non-collapsed points to have an initial number lower than the average, as discussed in subsection 9.4.2. This correction was division by the ratio ξ (equation (9.6)).

The all-or-nothing nature of the collapse also led to the use of an unusual approach in the determination of the stability coefficient k in the presence of the variations in the initial atom number. Instead of fitting a line to a set of random data, we had to determine a boundary between collapse and non-collapse events. Data were taken at high number (~ 6500) and a selected low number that was varied. Many points were taken at the critical magnetic fields at which the initial number variation resulted in some condensates collapsing while other condensates remained stable. This ensures that some of the stable condensate points are just on the verge of collapsing. The two stable condensate data points closest to the boundary between stability and instability, one from the low-number set and one from the high-number set, were then used. The magnetic field drifted by tens of mG from day to day due to variations in the ambient magnetic field. To allow day-to-day comparisons, the value of the high-number collapse field was used to measure the field drift. The

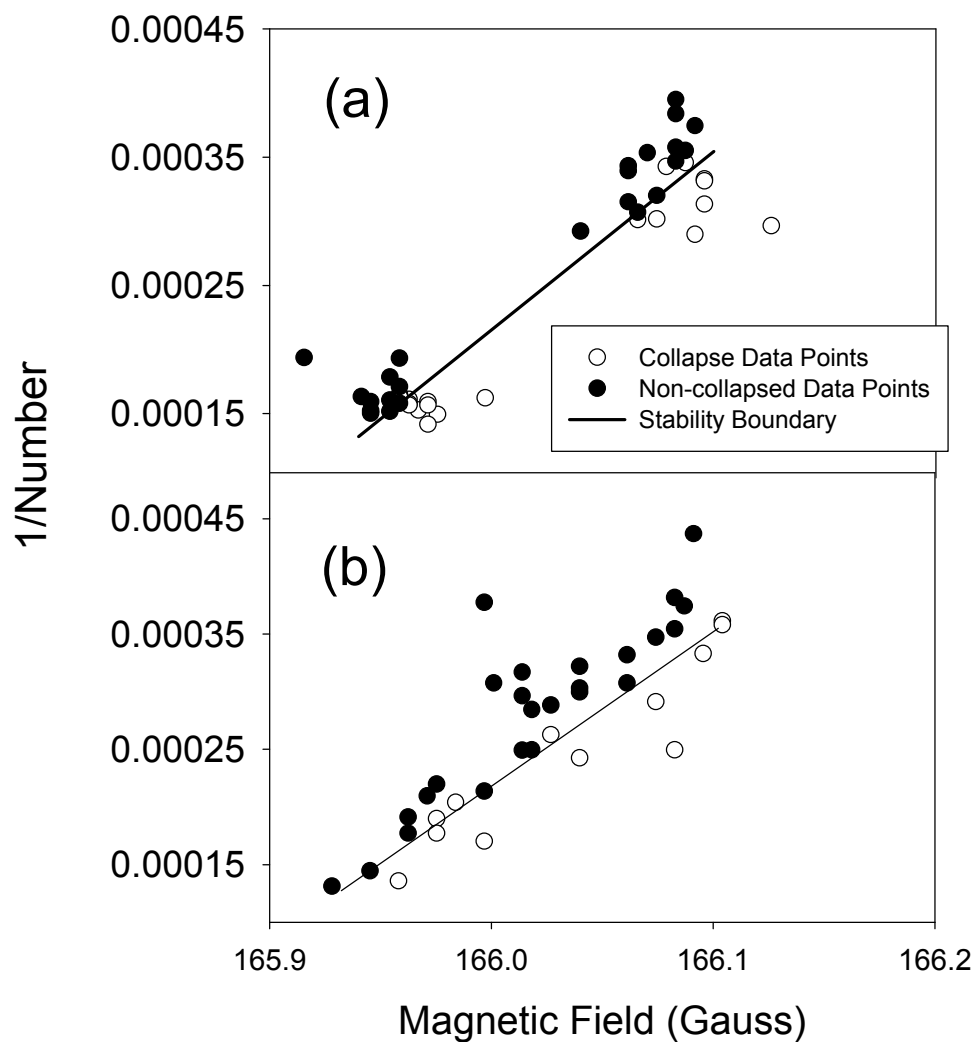


Figure 9.5. Determining the onset of collapse as a function of the condensate number and the magnetic field. Figures 9.5 (a) and 9.5 (b) show data sets used to determine the stability coefficient via the boundary that divides the data between collapse (\square) and non-collapse (\bullet) events. In Figure 2(a), the data were concentrated in two field regions to accurately measure the slope of the boundary line. In Figure 2(b), the initial number was varied over a larger range to illustrate the function form of the onset of instability. Due to initial number variations, some collapse points appear to be on the wrong side of the stability boundary.




Figure 9.6. This figure shows stable condensate numbers known to be on the verge of instability (see text) as determined from several data sets like those in figs. 9.6(a) and 9.6 (b). The fit to these points provides the value of the stability coefficient k .

remaining closest-to-boundary points for various numbers are the points shown in Figure 9.7 and are fit to determine a slope.

While this is not the most efficient way to determine the slope, the drift in initial number and the noise in the absorption imaging itself complicate more sophisticated analyses, such as the one described in subsection 9.5.3 below. Also, the closest-to-boundary-points method already gives a statistical error for k that is much smaller than the systematic errors. In order to estimate the uncertainty in the initial slope determinations, a standard Monte-Carlo simulation [175] was used.

To perform the Monte-Carlo simulation, a set of points corresponding in field location and average number of atoms to the data points in the actual sets of data was created. These points were randomly varied around their assigned central positions by the amount equal to the estimated number variation in the actual data. Once this was done, the random points with numbers greater than a critical value were caused to collapse, and were labeled as such. Noise was then again added to the remaining non-collapsed random points to simulate the noise inherent in the absorption imaging. The two points closest to the boundary were then taken in the same way as for the actual data sets and a slope was computed. This procedure was repeated ~ 100 times and the variation of the slopes provided an estimate of the uncertainty in the determination of the stability coefficient in a single set of data. This uncertainty agreed with the observed day-to-day variation of the measured slope of $1/N_{\text{crit}}$ vs. magnetic field. Also, this analysis showed that there were no significant ($> 2\%$) systematic errors that were introduced by using this method of analysis. One distortion that did appear was that occasionally the imaging noise would cause an anomalously large number of atoms to appear to be stable at higher values of $|a|$. The anomalously large number of atoms appears at large $|a|$ rather than small $|a|$ because the slope dN_{crit}/da is much shallower at large $|a|$ than small $|a|$. Therefore, the same-sized fractional noise in the absorption imaging tends to cause condensates at large $|a|$ to appear farther from the true onset of instability. This typically occurred once per Monte-Carlo set (and therefore likely once per actual data set), and so instead of using the point closest to the collapse boundary for the low-number sets, that point was eliminated and the second-closest point was used for the analysis.

The average of all of the data gives a slope of $0.00126(3)$ (atoms gauss) $^{-1}$ for $1/N_{\text{crit}}$ vs. magnetic field. In order to compare this result to the theory predictions, it is necessary to convert the measured slope from $1/N_{\text{crit}}$ vs. magnetic field to $1/N_{\text{crit}}$ vs. a . This is done using the determination of the scattering length presented in Chapter V. From this determination, the change in scattering length vs. magnetic field near 166.05 G is -33.47 ± 1.72 a₀/G (or -1.771 nm/G).

In order to estimate the uncertainty on the stability coefficient k determination, several factors need to be taken into account. There is of course the statistical error on the determination of the slope of $1/N_{\text{crit}}$ vs. magnetic field. The uncertainty of the slope of a vs. magnetic field must be taken into account. Also, the systematic uncertainty associated with the determination of the number of atoms from the absorption imaging has to be included as well (see section 4.12). There are two parts to this systematic uncertainty. One part of the systematic uncertainty would result in an equal error in the determination of the number of atoms in a condensate and the number of atoms in a thermal cloud, which has a larger spatial extent. The error due to this part of the systematic uncertainty in the number determination from the absorption imaging would cause an error in the determination of the slope of the scattering length vs. magnetic field as well since the measurement of the number of atoms in the thermal equilibration measurement described in Chapter V would have been affected. This error in the slope of a vs. B determination would partially offset the error in the determination of N_{crit} in the stability coefficient measurement. Since the thermal equilibration measurement scales as a^2 , and the stability coefficient measurement depends on a , a 10% systematic uncertainty in the number determination in both would result in a 5% uncertainty in k . The second part of the uncertainty in the number determination is due to errors associated with the small spatial size of the condensates. This uncertainty will only apply to the determination of N_{crit} and no partial cancellation will occur. We estimate the size of these second uncertainties to be 10%.

Another factor needs to be included in the uncertainty of the determination of N_{crit} is the uncertainty in the fraction of atoms in the thermal component of the condensate. We cannot accurately determine the thermal fraction once it is below 10%. It is unlikely that all of the thermal atoms have been cut away, and so it is assumed that 5(3) % of the atoms in the cloud are in the thermal component. This assigned uncertainty is also included in the systematic uncertainty of k .

Finally, the small errors (0.5%) associated with the calibration of the magnetic field are also included, although they do not significantly affect the size of the uncertainty in k .

Including all of these sources of uncertainty, the value measured for k in equation (9.3) is

$$k = 0.460 \pm 0.012 \text{ (statistical)} \pm 0.058 \text{ (systematic)}. \quad (9.8)$$

9.5.2 Comparison of the Measured Stability Coefficient to Theory Predictions. The functional form for the stability criterion in equation (9.3) seems to be correct – the onset of instability is observed to scale very well with the product of $N|a|$ as clearly evidenced by the data in figs. 9.6(a), 9.6(b), and 9.7. The measured stability coefficient is slightly lower than the predictions presented in Refs. [135,156-162], being 0.80 ± 0.10 of the predicted value of $k = 0.574$, or two standard deviations away. Using the slightly lower value of $k = 0.548$ from the predictions in an axially symmetric trap, the measured value is 0.84 ± 0.10 of the predicted value, still a 1.6 standard deviation disagreement. This means that the condensates appear to be collapsing at a lower value of the $N|a|$ product than the theory predicts [176]. Possible reasons for this deviation probably can only be accurately discussed in light of the dynamics of the collapse, which are currently unexplained by theory. See Chapter X for more details on the collapse dynamics.

Before discussing the examination of two potential sources of systematic error in the measurement of k , the promised aside on more sophisticated analysis to determine the stability coefficient from the collapse data will be included here. Neither of the two potential sources of systematic error was found to significantly affect the measurement of the value of k .

9.5.3 Aside on more Sophisticated Analysis of the Collapse Data to Determine k . Rather than using just the two closest-to-boundary points from each data set, it would be more efficient to employ a data analysis that used the information provided by all of the data points. For instance, the truncated gaussians used in the analysis of the fraction of atoms lost during the collapse (subsection 9.4.2) could form the basis for such an analysis. Three probabilities could be calculated using these truncated gaussians for a given value of the stability coefficient k and therefore a certain value of N_{crit} for each magnetic field. One is the probability that the condensates would collapse at a particular field given the average initial number and variation in the initial number of atoms. The second probability is the probability that condensates would not collapse at a given field given the

initial number and variation in the initial number of atoms. The last probability is the probability that none of the numbers of atoms in the non-collapsed condensate would be closer to the critical number than the number of atoms in the non-collapsed condensate that is closest to the critical number. For example, if the critical number at a particular field was predicted to be 2000 atoms, and the average initial number of atoms in the condensates was 1500 atoms with a standard deviation of 150 atoms, then it would be very unlikely that the maximum number of atoms observed in a stable condensate out of several points would be only 1000 atoms. For this example, the third probability would be very low. By varying the stability coefficient k to maximize the product of all three probabilities, the value of k that best fits the data could be determined.

There are two difficulties that make such an analysis almost hopelessly complicated. The first difficulty is that in addition to the actual variation of the initial number, there is a variation in the measured number due to the $\sim 5\%$ noise in the absorption imaging. This variation would have to be incorporated into the probabilities described in the previous paragraph, which is quite complicated. Even when that was accomplished, it would be very difficult to determine the average initial number and the variation in that initial number for each data point due to the drift in the initial number. It would be necessary to fit the drift in the value of the initial number while simultaneously fitting the stability coefficient. Since the rf-shift was often adjusted to keep the number of atoms in the condensate close to the desired value, the experiment conditions changed frequently further complicating the analysis. So instead of using any such analysis technique, the two points closest to the boundary were used as described above. In that case, it was only important to take enough data points straddling the onset of instability in order to ensure that two of the points were close to the collapse boundary.

9.6 Search for Systematic Shifts in k due to Finite Temperature and Non-adiabatic Field Ramps.

9.6.1 Finite Temperature and the Onset of Instability. The presence of a thermal component has been predicted to cause condensates to become unstable at smaller values of N_{crit} than at $T=0$ in both ${}^7\text{Li}$ [158,177,178] and ${}^{85}\text{Rb}$ [178]. However, the effect of finite temperature on the critical number is predicted to be small for the low temperatures studied in this work. Ref. [178]

predicts that the shift in the stable condensate number between zero temperature and the temperatures at which we took our data should be less than 1%.

In order to test whether the finite temperature of our clouds was affecting our measurement of the stability coefficient k , we deliberately created hot clouds that had roughly equal numbers of condensate and thermal atoms (~ 6800 atoms total) and then performed a collapse measurement. This was roughly a factor of ten more thermal atoms than were present in the previous cold-temperature measurements. The collapse of the condensates in these clouds was more difficult to determine since the thermal atoms obscured the change in shape of the condensates that collapsed. There was still a number loss that was detectable, however. The onset of instability was determined using these clouds by taking several data points that straddled the field corresponding to the onset of instability. Immediately following that, the onset of instability for condensates without a detectable thermal component and roughly the same number of condensate atoms was also measured in the same way. We did not observe any decrease in the value of N_{crit} for the hotter clouds.

In fact, the presence of the thermal component seemed to stabilize the condensates, if anything. The measured N_{crit} with the hotter clouds exceeded the N_{crit} at colder temperatures by 14%. This apparent increase in the stability of the condensates is not really greater than the uncertainties in the measurement of the thermal fraction, however. The mean-field expansion ramp (subsection 9.2.2) increased the size of the cloud beyond the resolution limit so that the total number of atoms could be determined accurately. However, given the probable coupling between the thermal component and condensate during the expansion ramp (see Chapter VIII), it is unclear how the condensate and thermal atom shapes will change, making determination of the individual number in each component difficult. Therefore, the measurement of the number of atoms in just the thermal component and the total number of atoms were used to determine the condensate atom number.

The measurement of the number of atoms in the thermal component was performed in two ways. First, the magnetic trap was turned off immediately after the condensates had been ramped to $a \sim 0$ and imaged immediately after the magnetic trap turn-off (no 1.6 ms expansion time as in subsection 4.11.4). At the time, a correction was applied to the number to account for the stray magnetic field present immediately after the trap turn-off, but we did not realize at the time that this

correction could vary by 20% of the initial condensate number from day to day. The other method used to determine the thermal atom number was to measure the spatial size of the thermal component. From this measurement a temperature could be determined. From the relation between the temperature and condensate fraction, equation (3.6), and the total number of atoms measured in the cloud, the thermal fraction could then also be computed. There is some question, however, how the thermal equilibrium in the cloud will be maintained during the ramp to $a \sim 0$ due to the reduction in the rate of elastic collisions due to the decrease in the elastic scattering cross section, and if any distortions in the spatial extent of the thermal cloud result. These two measurements agreed with each other at the 10 % level. Subsequent measurements have shown that the size of the uncertainties that were present in both methods of measuring the thermal fraction were as large as 14 %, making it unclear whether a shift in N_{crit} with temperature exists.

Given the fact that the apparent change in N_{crit} was no greater than 14% even though the thermal fraction was increased by an order of magnitude or so, we concluded that there were no significant shifts due to finite temperature in the determination of k at the cold temperatures used in our measurement.

9.6.2 Condensate Dynamics and the Onset of Instability. In this measurement we wanted to determine the onset of instability for condensates in the ground state. If the magnetic field ramp to $a < 0$ is too fast, then the condensate will no longer be in the ground state but will exhibit collective excitations (see Chapter VIII). To search for a dependence of the onset of instability on condensate dynamics, we increased the spatial size of the initial condensates by starting with an initial scattering length of $\sim +200 a_0$ instead of zero and used ramp times of both 200 ms and 1 ms to magnify the possible effect of any condensate dynamics. The previous data were taken using a magnetic field ramp duration (200 ms) that was more than twice the period of the lowest collective excitation of the condensate, to avoid exciting the condensate prior to collapse. This condition is true except for very near the point of instability where the lowest collective excitation frequency goes to zero (see subsection 9.1.2). In hindsight, since the condensates appear to be collapsing before the predicted critical point, the condensates may collapse before the region where the lowest collective excitation frequency diverges toward zero is reached.

For the 200 ms ramp there was no detectable change in the collapse point between ramps starting at $a = +200 a_0$ and $a=0$. The 1 ms, initial $a = 200 a_0$ ramp imparted a significant amount of pre-collapse excitation to the condensates, but the apparent value of k only decreased by 8%. Therefore, we concluded that there was no significant error introduced in the measurement of k due to the fact that the ramp speeds to $a < 0$ were not infinitely slow.

Using the PG model (subsection 8.2.6 and Ref. [165,166]), the effect of these ramps on the stability coefficient could be calculated for our trap frequencies. For the purposes of this calculation, the criterion for the collapse in the PG model was the occurrence of a point in time when the axial and radial spatial sizes went asymptotically to zero. The 200 ms ramp beginning at $a= 200 a_0$ was predicted to have a 10% lower value of k than the 200 ms ramp beginning at $a=0$. The 1 ms ramp beginning at $a= 200 a_0$ was actually found to stabilize the condensate against collapse, increasing k by 10%. Neither of these predictions matched the observed behavior, suggesting the PG model is inadequate for either describing the criterion for collapse or describing the oscillations of the condensate since the PG model uses a fixed shape.

9.7 Determination of B_{zero} .

9.7.1 Improved Measurement of B_{zero} . Since the onset of instability scaled well with the $N|a|$ product, the measurement of the onset of instability as a function of magnetic field provided for an improved determination of the magnetic field where the scattering length is equal to zero, B_{zero} (see Chapter V). Once the boundary between stable and unstable condensates were measured in data sets like those in Figure 9.5, the stability boundary could be extrapolated to the field where $1/N_{\text{crit}}$ (i.e. $N_{\text{crit}} = \square$). At this field, equation (9.3) implies that $a=0$.

Chapter X

Dynamics of Collapsing Condensates and Future Studies of ^{85}Rb Bose-Einstein Condensates

10.1 Overview.

While the experiments discussed in the previous Chapter were useful in determining the onset of instability, they did not examine the dynamics of collapsing condensates. The dynamics of the condensates while they collapsed (i.e. became unstable due to attractive interactions) are both puzzling and surprising. Perhaps the most surprising observation was that the collapse resulted in an explosion of “hot” atoms from the condensate. We also observed the formation of sharp density peaks during the collapse, a rapid loss of atoms, and a condensate that remained after the collapse that was in a highly excited state of motion. This Chapter describes the qualitative features of the collapse dynamics, but does not discuss quantitative details. Measurements of those are currently underway and will be carried out largely by my successors. The interested reader is referred to Ref. [18] for these details.

After describing the methods used to observe the collapse dynamics and their basic features, some brief attention will be given to possible future experiments in ^{85}Rb . The region of large scattering length will almost certainly be revisited. The nature of the losses induced by rapid ramps to large values of $a > 0$, the time that it takes for the condensates to respond to a rapid change in the scattering length, and perhaps the expansion rate at large scattering lengths will be investigated. It will also be interesting to study the behavior of condensates ramped quickly over the Feshbach peak into the region of large negative scattering length.

10.2 Measuring the Collapse Dynamics.

10.2.1 Creating the Initial Condensates. The condensates used in these experiments were created in the same manner as in subsection 9.2.1. A wider range of initial condensate number was used in these studies, and we began with anywhere from 3,000 to 15,000 atoms in condensates with

a scattering length near $a=0$ (165.85 G). The value of the magnetic field each day was calibrated by determining the point of the onset of instability in the manner described in Chapter IX.

10.2.2 Ramping the Magnetic Bias Field to Change a . Once the condensate had been prepared, a wide variety of different ramps of the magnetic field (and hence ramps of a) were used to investigate the dynamics of the collapse. Figure 10.1 shows the types of ramps used in these studies. The first part of any magnetic field ramp was a rapid (~ 0.1 ms) ramp to a selected value of $a < 0$ (a_{collapse}). The condensate was then allowed to evolve at that field for a variable amount of time (τ_{evolve}). The condensate collapse process was then halted with another rapid ramp to a positive or zero scattering length field (corresponding to a scattering length value a_{quench}). At this point, the condensate was often smaller than the resolution limit of our imaging system. To measure the number of atoms without distortion due to their small size, we used the mean-field repulsion at positive scattering lengths in the manner shown in Figure 10.1(a) and described in subsection 9.2.2 to expand the condensates above the resolution limit to measure their number. To measure features of the collapsing condensates that would be obscured by the mean-field expansion ramp, we used the ramp in Figure 10.1(b) that halted the collapse with a_{quench} , but did not expand the condensates further. Another ramp that was used was one in which the magnetic field was turned off to image the atoms without any ramp back from the $a < 0$ field (Figure 10.1(c)). This ramp was used to observe the motion of the condensate during its evolution at $a < 0$ and to measure the timing of the explosion of atoms from the condensate during the collapse.

10.2.3 Imaging the Condensates. After the magnetic trap fields were turned off, the shape, spatial size, and number of atoms in the cloud were measured using absorption imaging (see section 4.11). After the collapse, there were often several components to the cloud, a condensate that remained after the loss due to the collapse, a thermal component left over from the evaporative cooling, and atoms ejected in the explosion during the collapse. Fortunately, these various components had different spatial sizes allowing us to separate one component from another.

10.3 Dynamics of the Collapse.

10.3.1 Number vs. τ_{evolve} . Figure 10.2 shows the number of atoms in the condensate as a function of the time τ_{evolve} after the magnetic field had been ramped to a_{collapse} (see Figure 10.1 for a

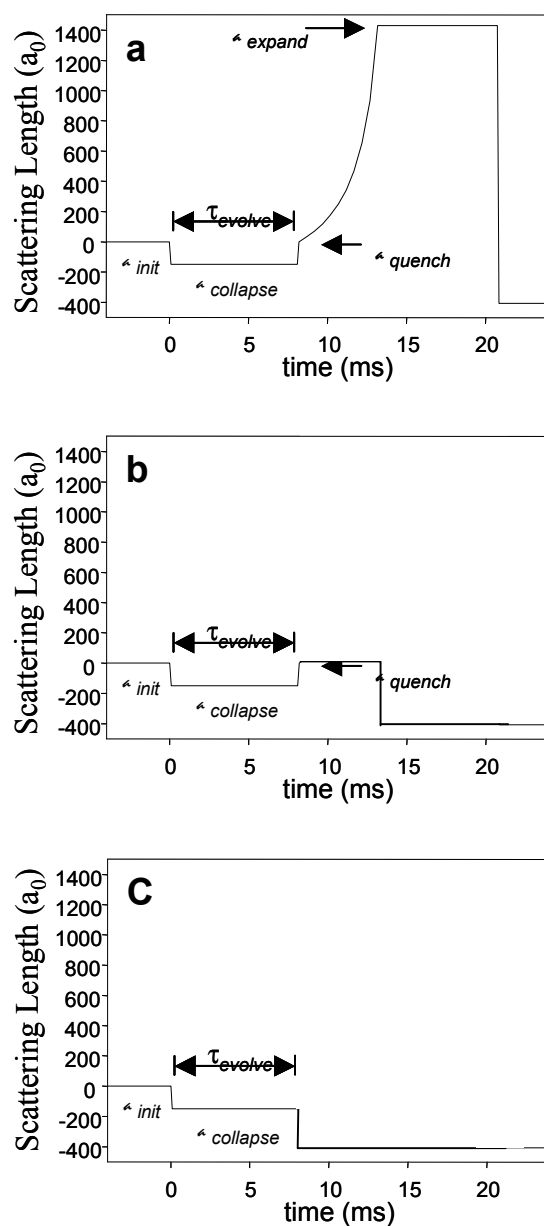


Figure 10.1. This figure illustrates several magnetic field ramps used to measure the condensate collapse dynamics. Rather than plotting the magnetic field as a function of time, the corresponding value of the scattering length is shown. The last change in the scattering length to ~ -400 a₀ represents the turnoff of the magnetic trap. The condensates were imaged 1.6 ms after this trap turnoff.




Figure 10.2. Number of condensate atoms as a function of the evolution time at a magnetic field with attractive interactions (τ_{evolve}). These data were taken using the ramp shown in Figure 10.1(a) at a field with a value of $a = \square 30 a_0$. The initial number of atoms in the condensate is shown as N_0 and the number of atoms remaining after the completion of the collapse is given as $N_{remnant}$. The solid line shows the exponential fit used to determine the decay constant during the rapid loss of atoms during the collapse.

definition of these quantities). There are three prominent features in this curve: there is a time t_{collapse} where there is little or no number loss; a rapid loss of atoms during a time t_{decay} , and a number of atoms that remain in the condensate (the “remnant”) after the rapid loss is completed. The evolution of the condensate during each of these three time periods will be discussed in the following subsections.

10.3.2 Contraction and Formation of Density “Spikes.” During the time t_{collapse} , the number of atoms in the condensate remains stable. The spatial size of the condensate contracts, and so the density of the condensate increases. This contraction cannot be observed directly since the condensates are below the resolution limit of the imaging system. However, the contraction can be inferred from the measured spatial size of the condensates after the mean-field expansion part of the ramp (see Figure 10.1(a)).

When the value of a is increased more rapidly than the condensate can adjust its ground state, a condensate that is denser will feel a greater outward force than one that is less dense since the mean-field interaction energy will be larger in the denser case. This means that condensates whose spatial size is small before the mean-field expansion tend to be made large, and originally large spatial size condensates tend not increase in size as much. During t_{collapse} , the spatial size of the condensates after the expansion ramp is observed to get larger, indicating that the condensates are contracting. The amount of contraction can be estimated by using the PG model described in subsection 8.2.6. Integrating the PG model equations while including the time dependence of a during the expansion ramp indicates that the change in the condensate density during the contraction is not very large --- the density increases only by $\sim 50\%$.

When the condensates were allowed to evolve at $a_{\text{quench}} \sim 0$ for a few milliseconds before the trap turnoff (e.g. Figure 10.1(b)) or a mean-field expansion ramp, we saw small highly anisotropic features form and grow during t_{collapse} . These features, which we refer to as “jets,” are shown in Figure 10.3. They are indicative of narrow regions of high density (density “spikes”) forming in the condensates as they collapse. The jets anisotropy indicates that these spikes are spatially non-uniform, and from the outward velocity of the atoms in the jet we are able to determine the spatial extent of the density spikes by using the uncertainty principle.




Figure 10.3. Jet images for a series of τ_{evolve} values for the conditions of Figure 10.2. The evolution times were 2, 3, 4, 6, 8, and 10 ms (from a to f). Each image is $150 \times 255 \mu\text{m}$. The bar indicates the optical depth scale. The jets were longest (i.e. most energetic) and contained the most atoms at values of τ_{evolve} for which the slope of the loss curve (Figure 10.2) was greatest. A tiny jet is barely visible for $\tau_{\text{evolve}} = 2$ ms (image a), which is 1.7 ms before the onset of the rapid loss of atoms from the condensate. The images also show how the number of condensate atoms decrease with time.

The length of the time period t_{collapse} where these density contractions were occurring varied in an intuitive way with the value of a_{collapse} , becoming shorter as the attractive interactions became stronger (i.e. t_{collapse} decreased as $|a_{\text{collapse}}|$ increased). The length of t_{collapse} also depends on the initial scattering length before the ramp to $a < 0$ (a_{init}), becoming longer as the value of a_{init} is increased. This implies that condensates with initially lower density take longer to contract. This trend is also seen in the dependence of t_{collapse} on the initial condensate number, with condensates with larger number having shorter t_{collapse} times.

10.3.3 Onset of Rapid Loss. After the condensates have spent some time contracting and forming density spikes, a loss of atoms suddenly begins and atoms are rapidly (\sim few ms) lost from the condensate. The insensitivity of τ_{decay} to the collapse field, sudden increase in the loss rate, and rapid rate of atom loss are puzzling. Simple estimates of the three-body recombination coefficient (see Chapter VI) and the measurement of the condensate density would give loss rate time scales orders of magnitude longer than what is observed.

10.3.4 Explosion of “Hot” Atoms. At the same time that the rapid loss rate begins, we also observe an explosion of atoms ejected from the condensate. The number of atoms in the explosion increases on the same time scale as the atoms are rapidly lost from the condensate. The average energy of these atoms in the explosion is typically on the order of a few hundred nK and can be highly anisotropic. This is much hotter than the temperature scale associated with the condensate ($T_c \sim 12$ nK) or any mean-field interaction energy (\sim tens of nK), but much smaller than the energy associated with the decay products of inelastic collisions (\sim tens of μK to mK). The fraction of atoms ejected in this explosion is always about 20% of the initial condensate number, regardless of the initial number of atoms or the value of a_{collapse} .

Figure 10.4 shows the absorption image in which the burst is measured. The burst appears as a high-aspect-ratio atom cloud that extends far past the central (condensate) component of the cloud. At first glance, this may suggest that the burst of atoms is highly anisotropic and consists of atoms flying away in the axial direction, but this is not the case. The image shown in Figure 10.4 is taken one-half of a radial trap oscillation period after the onset of the rapid loss of atoms during the condensate collapse, and so it is not an image of the explosion as it is occurring. Rather, we image a

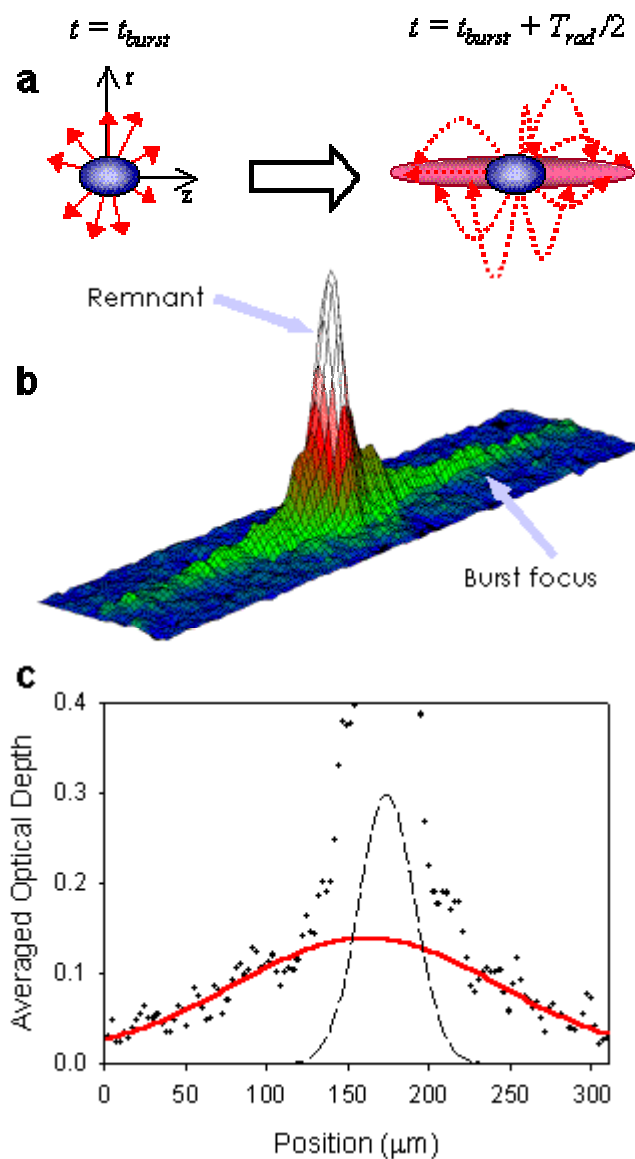


Figure 10.4 (Color). Plot (a) shows a conceptual illustration of a focus of the atoms in the explosion. Plot (b) shows an image of a refocused explosion. The temperature in the axial direction (along the direction of the long axis of the refocused cloud) determined from this data is 62 nK. Plot (c) Shows the radially averaged cross section of plot (b) with a Gaussian fit (red line) to the burst energy distribution. The central 100 μm were excluded from the fit to avoid distortion in the fit due to the remnant and the thermal cloud. The latter is present in the pre-collapse sample and appears to be unaffected by the collapse. The dashed line indicates the fit to this initial thermal component. Ramps like the one shown in Figure 10.1(c) were used to collect these data.

similar atom density distribution at one radial trap oscillation period, 1.5 radial oscillation periods, and so on. At one half of an axial oscillation period, we observe the same high-aspect-ratio atom cloud, only this time it is oriented along the radial, rather than the axial, direction.

These measurements indicate that what we are observing is the oscillation in the magnetic trap of the atoms ejected during the explosion. Essentially, the ejected atom cloud is exhibiting breathing motion as described in subsection 4.6.6. Due to the harmonic trap potential, the radial/axial coordinate of every particle is the precisely the same at time intervals separated by one radial/axial oscillation period. Since the atoms in the explosion were originally ejected from the collapsing condensate, at the appropriate times after the explosion the magnetic trap will refocus their coordinates to their initial values possessed at the time of the explosion. If the trap frequencies were the same in the radial and axial directions, then the initial density distribution of the ejected atoms would be recovered. However, the radial and axial frequencies are not the same, and so the radial coordinate is refocused at a different time than the axial coordinate, producing the observed density distributions like that in Figure 10.4. By measuring the spatial extent of the ejected atom cloud at a focus position, the energy in the radial or axial direction can be determined. The fact that the atoms ejected in the explosion oscillate in the magnetic trap in the way that they do means that the explosion did not alter the magnetic sublevel state of these atoms.

10.3.5 Remnant Condensate. Not all of the atoms are lost from the condensate during the collapse process. Eventually, the rapid loss and explosion cease and what remains is a condensate, which we refer to as the “remnant.” At least, we presume it is a condensate. If it were in thermal equilibrium (which it is not), then the atoms are dense and cold enough that their phase space density would suggest that the cloud was a condensate. The atoms also respond to a change in the scattering length, expanding due to increases in the value of a , but this behavior is not exclusive to condensates. More investigations may definitively determine the coherence properties of the remnant. The number of the atoms in the remnant was observed to be stable after the rapid loss was completed. In other words, there has been no evidence of a second collapse after the first one.

The remnants are observed to be oscillating in the magnetic trap with large amplitudes. Figure 10.5 shows the measured widths of a remnant as a function of time. The small size of the





Figure 10.5(Color). This figure demonstrates the oscillations of the remnant by plotting the radial width as a function of time at the $a < 0$ field. The solid line is a fit to the widths composed of two sine waves. The two frequencies of the sine waves correspond to twice the radial and axial frequencies within the measurement error. The two false color images show the peaked structures observed in the condensates as the remnant is oscillating.

remnants makes a definitive determination of the widths difficult, but the curve in Figure 10.5 does show that periodic motion is occurring. Also, the remnants are observed to have structure in their density distribution as they are oscillating, as shown in the two inset figures of Figure 10.5. The source of this structure is not known at this time.

10.3.6 “Missing” Atoms. One can sum the number of atoms observed in the burst and the number of atoms in the remnant and the sum of these two is not equal to the initial number of atoms in the condensate. Therefore, some atoms are no longer observed, and we refer to these as “missing” atoms. The burst fraction remains constant as $|a_{\text{collapse}}|$ is increased while the remnant number decreases, so the fraction of missing atoms increases as $|a_{\text{collapse}}|$ increases.

What has happened to the missing atoms to make them unobservable is still unknown. One possibility is that these atoms have somehow had enough energy imparted to rapidly leave the region around the condensate. The optical depth of the missing atoms would then drop quickly and they would not be observed in the absorption images. If we could improve the sensitivity of our recapture fluorescence measurement (subsections 4.11.1 and 4.11.2), then it may be possible to detect such energetic atoms. Another possibility is that the missing atoms have been converted to molecules so that they would not scatter any of the light from the probe laser during the absorption imaging. It is also possible that both of these hypotheses are true, i.e. energetic molecules are created. Determination of the reasons that the missing atoms are unobservable would be useful in illuminating the nature of their creation.

10.3.7 Condensate Collapse After the Magnetic Trap Turnoff. At low magnetic fields, the value of a is large and negative ($\approx 400 a_0$). Therefore, it would be expected that the condensates would collapse after the magnetic trap was turned off. Often, however, the impulse imparted to the condensates during the magnetic trap turnoff (subsection 7.2.2) prevents the condensates from collapsing before the atom cloud is imaged. When the initial density is high enough, though, the outward impulse is not enough to overcome the mean-field attraction at $B \sim 0$, and the condensates do collapse after the magnetic trap is turned off.

The basic features observed in the condensate collapses due to the trap turnoff are the same as those observed for collapses in the magnetic trap. There is a rapid loss of atoms, and an explosion

of hot atoms occurs. Since we usually let the condensates expand for ~ 2 ms before the cloud is imaged, the atoms in the explosion can look like a thermal component and so the turnoff collapse complicates the determination of the number of the thermal atoms. By using a magnetic field ramp to impart some outward velocity to dense condensates, this turnoff explosion can be eliminated.

10.4 Possible Future Experiments with ^{85}Rb BEC.

In addition to the detailed quantitative measurements of the collapse dynamics that are nearly completed, there are several other lines of investigation currently being considered and I will close this thesis with some brief speculation about possible future experiments in ^{85}Rb . As I mentioned in the introduction to this Chapter, I believe that ^{85}Rb will produce interesting studies of a variety of physics in the future. Given our past experience, it is difficult to predict what the interesting physics will be. Future observations will likely reveal surprising behavior, such as the features observed in the measurements of the collapse dynamics, presenting new avenues of investigation that cannot be discerned now.

10.4.1 Revisiting the Region of Large Positive a . The next set of experiments will likely investigate condensate dynamics at large values of $a > 0$, like the work discussed in Chapter VIII. Before attempting to measure any corrections to mean-field theory predictions using collective excitation frequencies again, several other experiments are likely. The loss and heating rates and mechanisms at large positive scattering length can be measured, both so that the rates are known for the purpose of planning and interpreting any measurements performed at those fields, and to investigate if the losses and heating rates are consistent with the ordinary two-body dipole relaxation and three-body recombination mechanisms.

By quickly ramping to large values of $a > 0$, it will also be possible to investigate the time that it takes for condensates to respond to a rapid change in the scattering length. Such an experiment is likely examining a beyond-mean-field effect, since the probability of finding one atom near another will likely be important for establishing the new value of the scattering length in the condensate. There are several possible time scales that could be relevant: the time scale associated with the bound state energy $\sim \hbar/2ma^2$ (which is independent of the density of the cloud); the time scale of the fluctuations associated with the beyond-mean-field corrections (i.e. fluctuations on the

spatial scale of the condensate healing length [12]) $\sim 4\pi\hbar na/m$; the time scale associated with molecular oscillations in pairing theory [97] $\sim gn^{1/2}/\hbar$ (where $g \sim 2.8 \times 10^{-38} \text{ Jm}^{3/2}$); or perhaps some other time scale or some combination of these. To give an idea of what these various time scales are, for a density of $1 \times 10^{12} \text{ cm}^{-3}$ and a scattering length of $a = 10,000 a_0$, the three time scales listed above are $\sim 0.7 \text{ ms}$, $\sim 0.2 \text{ ms}$, and $\sim 40 \mu\text{s}$. A possible experiment to measure the time scale for condensates to respond to a rapid change in the scattering length would be to change the magnetic field very rapidly to a large positive scattering length field, wait a variable time, and then ramp back to the original field. If the condensate does not respond to this ramp (i.e. it does not expand as much as expected), then the wait time can be used to calculate the time that it takes for a condensate to respond to a rapid change in the scattering length.

The measurement of beyond-mean-field corrections to the GP equation (see Chapter VIII) could be performed by using rapid ramps to fields with large scattering length. Instead of measuring the shift in the condensate collective excitation frequencies, the expansion rate of condensates in response to a change in the value of a could be measured. Since the expansion rate will depend on the energy per particle in the condensate, deviations from the mean-field predicted expansion rate would be observed as the beyond-mean-field corrections became significant. An expansion rate cannot in general be measured as precisely as a frequency shift. However, the avoidance of the temperature-dependent effects described in Chapter VIII may well offset any loss in sensitivity, since it will be possible to conduct experiments at fields with larger values of na^3 .

10.4.2 Thermodynamic Functions as a Function of Scattering Length. The value of a is expected to affect both the BEC phase transition critical temperature and the relationship between the fraction of atoms in the thermal component and the temperature of the cloud [12]. A positive value of a is expected to lower the critical temperature since the mean-field repulsion tends to lower the peak density of the cloud. A study of the critical temperature as a function of scattering length could confirm this expected effect. Some preliminary work was done studying this effect with limited success, and the increased knowledge of how ramping the magnetic field can elucidate condensate properties should allow a better characterization of this effect.

10.4.3 Going Over the Feshbach Resonance Peak. The studies detailed in subsection 7.2.3 show that it is possible to ramp over the Feshbach resonance peak without a significant amount of atom loss as long as the density of the condensates is relatively low. This means that it will be possible to study condensate evolution on the lower-magnetic-field side of the Feshbach resonance. Since the scattering length is negative at these fields, it is expected that the condensates will collapse. The first study will be to see if the condensates collapse in the same way at these fields as at the fields with negative scattering length at higher magnetic field. The ability to tune the scattering length towards negative infinity means that the collapse can be studied in a region that is beyond the range of the applicability of mean-field theory. Along different lines, predicted molecule-atom oscillations [94-97] could be searched for, although they are predicted to occur on a very fast time scale (\sim tens of microseconds). Of course, the high loss rates in the vicinity of the Feshbach peak will almost certainly complicate the experiments in this region of magnetic field, but by ramping past the peak, these loss rates should be able to be minimized and useful investigations performed.

Bibliography

- [1] Chris J. Myatt, Ph. D. Thesis, University of Colorado, 1997; C. J. Myatt *et al.*, Opt. Lett. **21**, 290 (1996).
- [2] S. Bose, Z. Phys. **26**, 178 (1924).
- [3] A. Einstein, Sitzungsber. Kgl. Preuss. Akad. Wiss. **1924**, 261 (1924); A. Einstein, Sitzungsber. Kgl. Preuss. Akad. Wiss. **1925**, 3 (1925).
- [4] For a review of superfluid Helium properties see J. Wilkes, *Properties of Liquid and Solid Helium* (Clarendon, Oxford, 1967).
- [5] J. L. Lin and J. P. Wolfe, Phys. Rev. Lett. **71**, 1222 (1993).
- [6] M. H. Anderson *et al.*, Science **269**, 198 (1995).
- [7] K. B. Davis *et al.*, Phys. Rev. Lett. **75**, 1969 (1995).
- [8] C. C. Bradley *et al.*, Phys. Rev. Lett. **75**, 1687 (1995).
- [9] For a review see C. E. Wieman *et al.*, Rev. Mod. Phys. **71**, S253 (1999).
- [10] H. F. Hess, Phys. Rev. B **34**, 3476 (1986).
- [11] D. G. Fried, Phys. Rev. Lett. **81**, 3811 (1998).
- [12] F. Dalfovo *et al.*, Rev. Mod. Phys. **71**, 463 (1999)
- [13] J. L. Roberts *et al.*, Phys. Rev. Lett. **81**, 5109 (1998).
- [14] J. L. Roberts *et al.*, Phys. Rev. A, 2001, in press.
- [15] J. L. Roberts *et al.*, Phys. Rev. Lett. **85**, 728 (2000).
- [16] S. L. Cornish *et al.*, Phys. Rev. Lett. **85**, 1975 (2000).
- [17] J. L. Roberts, *et al.*, Phys. Rev. Lett. **86**, 4211 (2001).
- [18] E. A. Donley *et al.*, submitted to Nature (London), 2001.
- [19] H. Feshbach, Ann. Phys. (N.Y.) **19**, 287 (1962).
- [20] W. C. Stwalley, Phys. Rev. Lett. **37**, 1628 (1976).
- [21] E. Tiesinga *et al.*, Phys. Rev. A **46**, R1167 (1992); E. Tiesinga, B. J. Verhaar, and H. Stoof, Phys. Rev. A **47**, 4114 (1993).
- [22] See, for instance, L. D. Landau and E. M. Lifshitz, *Mechanics*, 3rd Ed. (Pergamon Press, Oxford, 1976), pg. 159.

-
- [23] For instance, see J. R. Taylor, *Scattering Theory: the Quantum Theory of Nonrelativistic Collisions* (Kreiger, Malabar, Florida, 1987).
- [24] For example, see the treatment in R. Shankar, *Principles of Quantum Mechanics* (Plenum Press, New York, 1994).
- [25] See N. F. Mott and H. S. W. Massey, *The theory of Atomic Collisions* (Clarendon Press, Oxford, UK, 1949) for details.
- [26] J. J. Sakurai, *Modern Quantum Mechanics* (Addison-Wesley, New York, 1994) pg. 395-400.
- [27] See J. P. Burke, Jr., Ph.D. Thesis, University of Colorado, 1999 for a more complete discussion and set of references concerning cold collisions between alkali atoms.
- [28] M. Born and J. Oppenheimer, *Ann. Phys. (Leipzig)* **84**, 457 (1927).
- [29] C. Cohen-Tannoudji, B. Diu, and F. Laloe, *Quantum Mechanics* (John Wiley and Sons, New York, 1977).
- [30] J. M. Vogels, B. J. Verhaar, and R. H. Blok, *Phys. Rev. A* **57**, 4049.
- [31] J. P. Burke, Jr., C. H. Greene, and J. L. Bohn, *Phys. Rev. Lett.* **81**, 3355 (1998).
- [32] M. Gajda and K. Rzazewski, *Phys. Rev. Lett.* **78**, 2686 (1997); P. Navez *et al.*, *Phys. Rev. Lett.* **79**, 1789 (1997); M. Holthaus, E. Kalinowski, and K. Kirsten, *Annals of Physics* **270**, 198 (1999).
- [33] For a concise treatment see R. Kubo, *Statistical Mechanics* (North Holland, Amsterdam, 1965), pp. 29-31.
- [34] J. R. Ensher *et al.*, *Phys. Rev. Lett.* **77**, 4984 (1996).
- [35] For example, see S. M. Barnett, K. Burnett, J. A. Vaccaro, *Journal of the Results of the National Institute of Standards and Technology* **101**, 593 (1997).
- [36] M. R. Andrews *et al.*, *Science* **275**, 637 (1997); J. E. Samisarian *et al.*, *Phys. Rev. Lett.* **85**, 2040 (2000).
- [37] E. A. Burt *et al.*, *Phys. Rev. Lett.* **79**, 337 (1997).
- [38] For vortices in trapped atomic gas BEC see M. R. Matthews *et al.*, *Phys. Rev. Lett.* **83**, 2498 (2000); K. W. Madison *et al.*, *Phys. Rev. Lett.* **84**, 806 (2000).
- [39] K. Huang, *Statistical Mechanics* (Wiley, New York, 1963), Chapters 13 and 19.
- [40] B. D. Esry and C. H. Greene, *Phys. Rev. A* **60**, 1451 (1999)
- [41] B. D. Esry, *Phys. Rev. A* **55**, 1147 (1997).
- [42] See, for example, A. L. Fetter and J. D. Walecka, *Quantum Theory of Many-particle Systems* (McGraw-Hill, San Francisco 1971).
- [43] E. P. Gross, *J. Math. Phys.* **4**, 195 (1963); L. Pitaevskii, *Sov. Phys. JETP* **13**, 451 (1961).
- [44] J. O. Andersen and E. Braaten, *Phys. Rev. A* **60**, 2330 (1999).

-
- [45] N. P. Proukakis, K. Burnett, and H. T. C. Stoof, *Phys. Rev. A* **57**, 1230 (1998).
- [46] D. A. W. Hutchinson, E. Zaremba, and A. Griffin, *Phys. Rev. Lett.* **78**, 2511 (1997); R. J. Dodd, M. Edwards, C. W. Clark, and K. Burnett, *Phys. Rev. A* **57**, R32 (1998); S. Giorgini, *Phys. Rev. A* **61**, 063615 (2000).
- [47] R. Walser *et al.*, *Phys. Rev. A* **59**, 3878 (1999); R. Walser *et al.*, *Phys. Rev. A* **63**, 013607 (2001).
- [48] U. Al Khawaja and H. T. C. Stoof, *Phys. Rev. A* **62**, 053602.
- [49] M. R. Matthews *et al.*, *Phys. Rev. Lett.* **81**, 243 (1998).
- [50] A. Corney, *Atomic and Laser Spectroscopy* (Clarendon Press, Oxford, 1977).
- [51] U. Volz and H. Schmoranzer, *Physica Scripta* **T65**, 48 (1996).
- [52] G. P. Barwood, P. Gill, and W. R. C. Rowley, *Appl. Phys. B* **53**, 142 (1991).
- [53] E. Arimondo, M. Inguscio, and P. Violino, *Rev. Mod. Phys.* **49**, 31 (1977).
- [54] C. Cohen-Tannoudji, J. Dupont-Roc, and G. Grynberg, *Atom-Photon Interactions* (John Wiley and Sons Inc., New York, 1992), pp. 368-370.
- [55] P. D. Lett *et al.*, *J. Opt. Soc. Am B* **6**, 2084 (1989).
- [56] J. Dalibard and C. Cohen-Tannoudji, *J. Opt. Soc. Am.* **6**, 2023 (1989).
- [57] E. Rabb *et al.*, *Phys. Rev. Lett.* **59**, 2631 (1987).
- [58] C. Monroe *et al.*, *Phys. Rev. Lett.* **65**, 1571 (1990).
- [59] C. Wieman, G. Flowers, and S. Gilbert, *Am. Jour. Phys.* **63**, 317 (1995).
- [60] T. Walker, D. Sesko, and C. Wieman, *Phys. Rev. Lett.* **64**, 408 (1990).
- [61] For instance, A. Gallagher and D. E. Pritchard, *Phys. Rev. Lett.* **63**, 957 (1989); M. H. Anderson *et al.*, *J. Opt. Soc. Am. B* **11**, 1332 (1994);
- [62] C. G. Townsend *et al.*, *Phys. Rev. A* **52**, 1423 (1995).
- [63] W. D. Phillips and H. Metcalf, *Phys. Rev. Lett.* **48**, 596 (1982).
- [64] Jason Ensher, Ph. D. Thesis, University of Colorado, pp. 79-81 (1998) or B. Demarco, H. Rohner, and D. S. Jin, *Rev. Sci. Instru.* **70**, 1967 (1999).
- [65] J.H. Moore, C. C. Davis, and M. A. Copken, *Building Scientific Apparatus, 2nd ed.*, (Addison-Wesley, Reading, Massachusettes, 1989), pg. 79.
- [66] M. Stephans, R. Rhodes, and C. Wieman, *Journal of Applied Physics* **76**, 3479 (1994).
- [67] K. B. MacAdam, A. Steinbach, and C. Wieman, *Am. J. Phys.* **60**, 109 (1992).

-
- [68] W. Demtröder, *Laser Spectroscopy Basic Concepts and Instrumentation, 2nd Enlarged Ed.* (Springer, Berlin, 1996), pg. 444.
- [69] K. L. Corwin *et al.*, *Appl. Opt.* **37**, 3295 (1998).
- [70] This liquid crystal shutter had the advantage of being relatively cheap and easy to install. However, the liquid crystal absorbs about 50% of the laser light and will be likely be replaced in the near future to improve the collection MOT loading time.
- [71] W. Petrich *et al.*, *J. Opt. Soc. Am. B* **11**, 1332 (1994).
- [72]
- [73] See, for instance, V. D. Barger and M. G. Olsson, *Classical Electricity and Magnetism a Contemporary Perspective* (Allyn and Bacon, Inc., Boston 1987), pg. 68.
- [74] W. Petrich *et al.*, *Phys. Rev. Lett.* **74**, 3352 (1995); K. B. Davis *et al.*, *ibid.* **75**, 3969 (1995).
- [75] T. Bergeman, E. Erez, and H. Metcalf, *Phys. Rev. A* **35**, 1535 (1987).
- [76] F. Reif, *Fundamentals of Statistical and Thermal Physics* (McGraw-Hill, New York, 1965).
- [77] C. Cohen-Tannoudji, J. Dupont-Roc, and G. Grynberg, *Atom-Photon Interactions* (John Wiley and Sons Inc., New York, 1992), pp. 460-489.
- [78] P. Horwitz, *Appl. Phys. Lett.* **26**, 306 (1975); J. C. Camparo and R. D. Frueholz, *Phys. Rev. A* **30**, 803 (1984).
- [79] L. D. Landau and E. M. Lifshitz, *Mechanics, 3rd Ed.* (Pergamon Press, Oxford, 1976), pg. 80.
- [80] The resolution limit listed here is the usual one given by the Rayleigh criterion (see R. A. Serway, *Physics for Scientists and Engineers with Modern Physics, 3rd Ed.* (Saunders College Publishing, Philadelphia, 1990), pg. 1080, or any other general physics or introductory optics textbook). This criterion does not apply as obviously to our system since we are not trying to resolve two separate objects, but the FWHM of a sub-resolution limit object is roughly equal to the resolution limit given in the Rayleigh criterion.
- [81] P. Meystre and M. Sargent III, *Elements of Quantum Optics, 2nd Ed.* (Springer-Verlag, Berlin, 1991).
- [82] For instance, see V. Sanchez-Villicana, S. D. Gensemer, and P. L. Gould, *Phys. Rev. A* **54**, R3730 (1996).
- [83] M. Marinescu and L. You, *Phys. Rev. Lett.* **81**, 4596 (1998).
- [84] P. O. Fedichev, *et al.*, *Phys. Rev. Lett.* **77**, 2913 (1996).
- [85] J. L. Bohn and P. S. Julienne, *Phys. Rev. A* **56**, 1486.
- [86] A. J. Moerdijk *et al.*, *Phys. Rev. A* **51**, 4852; J. M. Vogels *et al.*, *Phys. Rev. A* **56**, R1067 (1997).
- [87] C. R. Monroe *et al.*, *Phys. Rev. Lett.* **70**, 414 (1993).
- [88] N. R. Newbury, C. J. Myatt, and C. E. Wieman, *Phys. Rev. A* **51**, R2680 (1995).

-
- [89] For a review, see P. S. Julienne, *J. Res. Natl. Inst. Stand. Technol.* **101**, 487 (1996).
- [90] S. Inouye *et al.*, *Nature (London)* **392**, 151 (1998); J. Stenger *et al.*, *Phys. Rev. Lett.* **82**, 2422 (1999).
- [91] Ph. Courteille *et al.*, *Phys. Rev. Lett.* **81**, 69 (1998).
- [92] J. L. Roberts *et al.*, *Phys. Rev. Lett.* **81**, 5109 (1998).
- [93] V. Vuletic *et al.*, *Phys. Rev. Lett.* **82**, 1406 (1999); V. Vuletic *et al.*, *Phys. Rev. Lett.* **83**, 943 (1999).
- [94] E. Timmermans *et al.*, *Phys. Rep.* **315**, 199 (1999); E. Timmermans *et al.*, *Phys. Rev. Lett.* **83**, 2691 (1999).
- [95] F. A. van Abeelen and B. J. Verhaar, *Phys. Rev. Lett.* **83**, 1550 (1999).
- [96] V. A. Yurovsky, *et al.*, *Phys. Rev. A* **60**, R765 (1999).
- [97] M. Holland, J. Park, and R. Walser, *Phys. Rev. Lett.* **86**, 1915 (2001).
- [98] C. R. Monroe *et al.*, *Phys. Rev. Lett.* **70**, 414 (1993).
- [99] A. Griffin, W. C. Wu, and S. Stingari, *Phys. Rev. Lett.* **78**, 1838 (1997).
- [100] E. A. Burt *et al.*, *Phys. Rev. Lett.* **79**, 337 (1997).
- [101] This value for κ agrees to within 4% of a recent Monte-carlo calculation in B. DeMarco *et al.*, *Phys. Rev. Lett.* **82**, 4208 (1999).
- [102] J. Burke, private communication
- [103] Before the completion of our studies, two other groups observed Feshbach resonances in low temperature collisions between alkali atoms (see Refs. [90] and [91]). We took data while our system was suffering from very short magnetic trap lifetimes, on the order of 14 seconds (see subsection 4.4.3), and saw indications of a change in scattering length near 160 G. Systematic effects present at the time prevented us from being confident that we really were seeing a change in the scattering length, however, until the vacuum problem was fixed so that we could better prepare our samples. In the meantime, we alerted the group at the University of Texas of our preliminary results, and this explains the acknowledgement in Ref. [91].
- [104] F. H. Mies *et al.*, *J. Res. Natl. Inst. Stand. Technol.* **101**, 521 (1996).
- [105] G. F. Gribakin and V. V. Flambaum, *Phys. Rev. A* **48**, 546 (1993).
- [106] The values for C_8 and C_{10} were taken from M. Marinescu, H. R. Sadeghpour, and A. Dalgarno, *Phys. Rev. A* **49**, 982 (1994).
- [107] H. M. J. M. Boesten *et al.*, *Phys. Rev. Lett.* **77**, 5194 (1996).
- [108] J. P. Burke, Jr. *et al.*, *Phys. Rev. Lett.* **80**, 2097 (1998).
- [109] C. C. Tsai *et al.*, *Phys. Rev. Lett.* **79**, 1245 (1997).

-
- [110] P. S. Julienne *et al.*, Phys. Rev. Lett. **78**, 1880 (1997); J. P. Burke, Jr. *et al.*, Phys. Rev. A **55**, R2511 (1997).
- [111] H. M. J. M. Boesten *et al.*, Phys. Rev. A **56**, 636 (1997).
- [112] M. R. Matthews *et al.*, Phys. Rev. Lett. **81**, 243 (1998).
- [113] C. J. Myatt *et al.*, Phys. Rev. Lett. **78**, 586 (1997).
- [114] J. Burke, Jr. and Carl Williams have both communicated to me that the position of the Feshbach resonance should change by several gauss from the zero temperature location when the collision energies are on the order of 100 μ K. Communications with D. Heinzen indicated that they did not take this shift in the position of the Feshbach resonance into account in his and his coworkers' original analysis.
- [115] C. J. Williams, private communication. This theoretical analysis was preliminary at the time.
- [116] A. Derevianko *et al.*, Phys. Rev. Lett. **82**, 3589 (1999). In this work, C_6 was calculated to be 4691 ± 50 .
- [117] J. M. Vogels *et al.*, Phys. Rev. A **61**, 043407 (2000). Here, C_6 was measured to be 4650 ± 50 .
- [118] J. P. Burke, Jr. *et al.*, Phys. Rev. Lett. **80**, 2097 (1998).
- [119] F.H. Mies *et al.*, J. Res. Natl. Inst. Stand. Technol. **101**, 521 (1996).
- [120] P. O. Fedichev, M. W. Reynolds, and G. V. Shlyapnikov, Phys. Rev. Lett. **77**, 2921 (1996).
- [121] A. J. Moerdijk, H. M. J. M. Boesten, and B. J. Verhaar, Phys. Rev. A **53**, 916 (1996).
- [122] B. D. Esry, C. H. Greene, and J. P. Burke, Jr., Phys. Rev. Lett. **83**, 1751 (1999). The authors of this paper have communicated to us that the predictions contained in it are a factor of six too high due to a mistake in the calculations.
- [123] E. Nielsen and J. H. Macek, Phys. Rev. Lett. **83**, 1566 (1999).
- [124] P. F. Bedaque, E. Braaten, and H. W. Hemmer, Phys. Rev. Lett. **85**, 908 (2000).
- [125] P. Leo, C. Willimas, and P. Julienne, Phys. Rev. Lett. **85**, 2721(2000); P. L. Leo *et al.*, Phys. Rev. Lett. **81**, 1389 (1998).
- [126] S. Kotochigova, E. Tiesinga, and P. S. Julienne, Phys. Rev. A **63**, 012517 (2000).
- [127] P. R. Bevington, *Data Reduction and Error Analysis for the Physical Sciences* (McGraw-Hill, New York, 1969).
- [128] C. J. Williams (NIST, Gaithersburg) private communication. The methods described in Ref. [119] were used with an updated Rb-Rb interaction potential. The dipolar rates at several collision energies were given for each field.
- [129] For a description of the polarization rotation imaging used in the ^7Li work, see C. C. Bradley, C. A. Sackett, and R. G. Hulet, Phys. Rev. Lett. **78**, 985 (1997). This treatment does not solve Maxwell's equations in the slowly varying envelope approximation, however, and so their results are not guaranteed to be correct. Most polarization rotation imaging treatments do not consider a spin-polarized sample. See for example, B. P. Ablitt, *et al.*, Waves in Random Media **9**, 561 (1999).

-
- [130] See B. P. Anderson and M. A. Kasevich, *Phys. Rev. A* **59**, R938 (1999) and the references therein.
- [131] M. R. Andrews *et al.*, *Science* **273**, 84 (1996).
- [132] <http://fermion.colorado.edu/~chg/Collisions/>
- [133] H. Beijerinck, preprint.
- [134] K. B. Davis, M. O. Mewes, and W. Ketterle, *Appl. Phys. B* **60**, 155 (1995).
- [135] P. A. Ruprecht *et al.*, *Phys. Rev. A* **51**, 4704 (1995).
- [136] T. D. Lee, K. Huang, and C. N. Yang, *Phys. Rev.* **106**, 1135 (1957).
- [137] See, for instance, A. L. Fetter, *Phys. Rev. A* **53**, 4245 (1996).
- [138] N. Bogoliubov, *J. Phys. (Moscow)* **11**, 23 (1947).
- [139] M. R. Andrews *et al.*, *Phys. Rev. Lett.* **79**, 553 (1997).
- [140] J. Denschlag *et al.*, *Science* **287**, 971 (2000); S. Burger *et al.*, *cond-mat /9910487* at xxx.lanl.gov e-print website; B. P. Anderson *et al.*, *Phys. Rev. Lett.* **86**, 2926 (2001).
- [141] S. Stringari, *Phys. Rev. Lett.* **77**, 2360 (1996).
- [142] D. S. Jin *et al.*, *Phys. Rev. Lett.* **77**, 420 (1996).
- [143] M.-O. Mewes *et al.*, *Phys. Rev. Lett.* **77**, 988 (1996).
- [144] D. S. Jin *et al.*, *Phys. Rev. Lett.* **78**, 764 (1997).
- [145] D. M. Stamper-Kurn *et al.*, *Phys. Rev. Lett.* **81**, 500 (1998).
- [146] G. Hechenblaikner *et al.*, *Phys. Rev. Lett.* **85**, 692 (2000).
- [147] M. Edwards *et al.*, *Phys. Rev. Lett.* **77**, 1671 (1996).
- [148] V. M. Perez-García *et al.*, *Phys. Rev. Lett.* **77**, 5320 (1996).
- [149] L. Pitaevskii and S. Stringari, *Phys. Rev. Lett.* **81**, 4541 (1998).
- [150] E. Braaten and J. Pearson, *Phys. Rev. Lett.* **82**, 255 (1999).
- [151] S. Giorgini, *Phys. Rev. A* **61**, 063615 (2000).
- [152] S. Giorgini, J. Boronat, and J. Casulleras, *Phys. Rev. A* **60**, 5129 (1999).
- [153] This argument is attributed to A. Leggett.
- [154] See Ref. [12] and the discussion and references therein concerning the damping of collective excitations.
- [155] S. Giorgini, private communication.

-
- [156] R. J. Dodd *et al.*, Phys. Rev. A **54**, 661 (1996).
- [157] F. Dalfovo and S. Stingari, Phys. Rev. A **53**, 2477 (1996).
- [158] M. Houbiers and H. T. C. Stoof, Phys. Rev. A **54**, 5055 (1996).
- [159] A. Gammal, T. Frederico, and L. Tomio, Phys. Rev. E **60**, 2421 (1999).
- [160] A. Eleftheriou and K. Huang, Phys. Rev. A **61**, 043601 (2000).
- [161] L. Berge, T. J. Alexander, and Y. S. Kivshar, Phys. Rev. A **62**, 023607 (2000).
- [162] V. M. Perez-Garcia, H. Michinel, and H. Herrero, Phys. Rev. A **57**, 3837 (1998).
- [163] B. Esry, private communication.
- [164] A. Gammal, T. Frederico, and L. Tomio, cond-mat/0104231 at <http://xxx.lanl.gov> e-print server.
- [165] A. L. Pattenayak *et al.*, Phys. Rev. A **63**, 033604 (2001).
- [166] Eric Cornell and Susannah Kreim, private communication.
- [167] Y. Kagan, A. Muryshev, and G. V. Shlyapnikov, Phys. Rev. Lett. **81**, 933 (1998).
- [168] C.C. Bradley, C. A. Sackett, and R. G. Hulet, Phys. Rev. Lett. **78**, 985 (1997).
- [169] J. M. Gerton *et al.*, Nature (London) **408**, 6813 (2000).
- [170] C. A. Sackett *et al.*, Phys. Rev. Lett. **82**, 876 (1999).
- [171] Ernst *et al.*, Appl. Phys. B **67**, 719 (1998).
- [172] C. A. Sackett, H. T. C. Stoff, and R. G. Hulet, Phys. Rev. Lett. **80**, 2031 (1998).
- [173] H. T. C. Stoof, J. Stat. Phys. **87**, 1653 (1997).
- [174] In contrast to Ref. [173] M. Ueda and A. J. Leggett, Phys. Rev. Lett. **80**, 1576 (1998) and J. L. Bohn, B. D. Esry, and C. H. Greene, Phys. Rev. A **58**, 584 (1998) both find that the quantum tunneling probability is extremely small unless the number of atoms in the condensate is within a few atoms of the critical number.
- [175] See, for instance, W. H. Press *et al.*, *Numerical Recipes: The Art of Scientific Computing* (Cambridge University Press, Cambridge, United Kingdom 1989), pp. 529-538.
- [176] It has been suggested to us that this lower value of the stability coefficient may be indicative of a local collapse at the peak density of the condensate. The theory underlying this claim has not been fully developed, however, and we are by no means convinced that this is the explanation for the low value of k that we measured. See M. Ueda and K. Huang, Phys. Rev. A **60**, 3317 (1999).
- [177] M. J. Davis, D. A. W. Hutchinson, and E. Zaremba, J. Phys. B **32**, 3993 (1999).
- [178] E. J. Mueller and G. Baym, Phys. Rev. A **62**, 0.53605 (2000).

