

Copyright © 2006 by Joseph Michael Kinast
Some rights reserved

This work is licensed under the Creative Commons
Attribution-NonCommercial-ShareAlike License.

To view a copy of this license, visit
<http://creativecommons.org/licenses/by-nc-sa/2.5/>

THERMODYNAMICS AND SUPERFLUIDITY OF A
STRONGLY INTERACTING FERMI GAS

by

Joseph Michael Kinast

Department of Physics
Duke University

Date: _____

Approved:

Dr. John Thomas, Supervisor

Dr. Steffen Bass

Dr. Calvin Howell

Dr. Henry Everitt

Dr. Stephen Teitsworth

Dissertation submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy
in the Department of Physics
in the Graduate School of
Duke University

2006

ABSTRACT

(Physics)

THERMODYNAMICS AND SUPERFLUIDITY OF A
STRONGLY INTERACTING FERMI GAS

by

Joseph Michael Kinast

Department of Physics
Duke University

Date: _____

Approved:

Dr. John Thomas, Supervisor

Dr. Steffen Bass

Dr. Calvin Howell

Dr. Henry Everitt

Dr. Stephen Teitsworth

An abstract of a dissertation submitted in partial fulfillment of
the requirements for the degree of Doctor of Philosophy
in the Department of Physics
in the Graduate School of
Duke University

2006

Abstract

Strongly interacting two-component Fermi gases are prototypes for other exotic systems in nature, including high temperature superconductors, quark-gluon plasmas and neutron stars. Interactions between spin components can be varied widely, permitting exploration of the crossover from Bose to Fermi statistics. In the strongly interacting limit, the behavior of these systems becomes independent of the microscopic details of their interactions. Hence, strongly interacting Fermi gases provide a testing ground for many-body nonperturbative calculations of strongly interacting matter.

In this dissertation, I describe the first thermodynamic study of a strongly interacting degenerate Fermi gas. This study of the heat capacity shows a transition in behavior at $T/T_F = 0.27(0.02)$, which is interpreted as a superfluid phase transition by Kathy Levin's pseudogap theory. This is the first direct measurement of the superfluid transition temperature in the strongly interacting regime. I also describe the first measurement of the temperature dependence of the radial breathing mode in a strongly interacting Fermi gas. As the temperature of the gas is lowered, we find an increase in the lifetime of the breathing mode oscillation at the hydrodynamic frequency. This is inconsistent with expectations for a collisional system, and provides evidence for a superfluid state. As the temperature increases, we observe a transition in behavior at $T/T_F = 0.35$. In the high temperature regime, an abrupt increase in the damping rate is interpreted as the

breaking of noncondensed atom pairs. Further, I describe studies of the magnetic field dependence of the radial breathing mode in a low temperature system, which test the best current many-body predictions for the equation of state of the gas. An increase in the damping rate above the center of a broad Feshbach resonance is interpreted as the breaking of atom pairs.

I describe our techniques for producing degenerate, stable, strongly interacting two-component mixtures of ${}^6\text{Li}$ confined in an optical trap. Evaporative cooling yields $2.0(0.2) \times 10^5$ atoms at temperatures as low as $T/T_F = 0.06$. Starting from these cold atom samples, I describe the techniques used to study the heat capacity and radial breathing mode in a strongly interacting Fermi gas.

Acknowledgements

It seems appropriate that the first few words in this space should be addressed to my parents, who succeeded in the gargantuan task of raising me without killing me, tempting though it may have been. It seems I have inherited my father's stubbornness, perfectionism, and sense of humor, a combination that was tempered by my fortune to inherit some of my mother's quiet thoughtfulness. In my toddlerhood and childhood, I clashed often and vociferously with my parents, who wanted nothing more than to impress upon me the futility of trying to be perfect while giving me just enough room to commit and then learn from my mistakes. As I have already half-jokingly suggested, the fact that I survived my childhood is a testament to their tolerance and parenting skills. That I now view my parents as friends is a testament to their inherent goodness. Without their patience, compassion, and willingness to listen, the last quarter century would have been substantially more difficult and my accomplishments less rewarding.

If my parents were charged with trying to get me to see the big picture as I grew, my siblings were tasked with teaching me more pragmatic lessons. I can thank my older brother, Matt, for teaching me that it is unwise to pick a fight with someone older, larger and stronger than you, a lesson that required many repetitions before it took hold. My brother has taught me much about accepting loss without feeling like a loser and cherishing victory without gloating. A friend and one of the nicest guys I know, my brother has influenced me more than he

probably knows.

Though I claim with confidence that I have learned many good lessons from my siblings, I suspect that my younger sister Lori matured and learned in spite of my influence. As I could not seem to beat Matt in much of anything, Lori was often pressed into being an unwilling participant in competitions whose point she did not see and whose rules were stacked against her. But bless her boundless heart, she participated all the same and still talks to me even today. I thank Lori for showing me that not everything needs to be a competition, and for teaching me my toughest lesson: sometimes taking a break from work is necessary.

Finally, to round out my immediate family, I thank Liz, my sister-in-law, who is sufficiently good natured to pretend that my explanations about my research are both clear and interesting. As anyone present for such a discussion can attest, feigning interest is not for the fainthearted, and yet Liz pulls it off beautifully. Even more remarkable, Liz tolerates my brother and me in spite of the fact that our collective IQ never rises above 75 when we are in the same room.

There are, of course, many in the Duke physics community whose guidance, help, and companionship I must recognize. First, there is John Thomas, my advisor. Sharp, insightful, jovial, and quick-witted, John is a ruthlessly logical beast in discussions of physics, and there is no question that I am a better researcher for having been in his company over the past several years. On more than one occasion, he has taken the entire research group on his back and willed us through a difficult stretch. Of course, no acknowledgement of John would be complete without sharing a few of his more memorable statements. Without further ado, John's insights:

- On universal behavior: “We measure stuff that doesn’t depend on anything. Therefore, we know nothing.”
- On his teenage years: “It’s amazing what a pinch of high explosive will do.”
- On a former colleague: “He was a nice guy, but a total tyrant.”

Readers who enjoyed these statements are encouraged to speak to other members of the research group for more entertaining, but less printable, musings from John on life, death, and everything in between.

Other members of the research group here have played instrumental roles in my development. I spent little time with Ken O’Hara and Stephen Granade, but the work in this dissertation would not have been possible without their substantial efforts. I thank Mike Gehm and Staci Hemmer for guiding me in my early days in the laboratory. Their willingness to teach at a time when they were very busy themselves speaks volumes about their altruistic natures. While Staci has returned to her western Montana roots, I consider myself fortunate to have remained in contact with Mike while he served as a postdoctoral researcher in another Duke laboratory. I thank Andrey Turlapov, a postdoctoral researcher in our lab, for his companionship and many insights during our hours in the lab together. Andrey’s relentless optimism was a force not to be trifled with, even in our darkest moments. I wish Andrey all the best as he moves toward establishing his own research group.

My time in the lab was also made more enjoyable by the presence of more recent inductees into the research group: Bason Clancy, Le Luo, and James Joseph. In them, the future of the lab is in good, capable hands. In particular, it seems

appropriate to single out Bason, who has served the dual role of co-worker and friend during my time at Duke. Though I kept no records of such things, I would guess that at least 80% of the beer I consumed in Durham occurred in Bason's presence, and I state without hesitation that I might not have finished my degree had Bason not joined the group. To round out my labmates, I thank Ingrid Kaldre, an undergraduate researcher in our laboratory, for managing to make me laugh no matter how bad a mood I was in. I cannot think of any other person whose company I enjoy as much in spite of the fact that we agree on almost nothing.

Outside of our lab, I express gratitude to Andy Dawes, Michael Stenner, and other members of the Gauthier lab for their willingness to share optical elements and expertise when we were dreaming up and implementing solutions on a short time scale. Assistance from the front office staff, notably Donna Ruger and Maxine Stern, was invaluable as I navigated Duke's bureaucratic landscape. John Wambaugh, Amanda Sabourov, Brad Marts, and Josh(topher) Tuttle made the survival of my first year here palatable if not pleasant. Laughing at a huge pile of work in the company of others who are similarly condemned to suffer is a far richer experience than weeping openly in front of the pile by yourself. All the best to my "grads01" comrades as they finish their degrees and scatter about the country. Thanks also to Brian Bunton, the aforementioned Brad, and Brian Tighe for their company at Durham Bulls games. Now that's a group of people who truly appreciate the majesty of beating a bad joke into the ground.

I thank the members of my advisory committee, including Steffen Bass, Henry Everitt, Calvin Howell, and Stephen Teitsworth for taking time from their busy

schedules to review my dissertation and offer ways in which it could be improved.

There have been many good teachers along the way who, whether they realized it or not, were integral in getting me to this point. The faculty at Drew University, including Drs. Fenstermacher, Supplee, Carter, McGee, and Boeshaar deserve my gratitude for preparing me for graduate school. Going even further back, I would like to single out the efforts of my 4th grade science and math teacher, Carl Shipley. Mr. Shipley had a unique ability to challenge each student in his classroom, regardless of his or her individual capabilities. I cannot think of a finer measure of an instructor.

In the end, for all of the help I received from members of the academic community, it may well have been trumped by the support I received from my friends and family. Brian Joos, Phil Hessler, Maddy Hartman, Bason Clancy, and my family (including my parents' saccharine sweet pooch, Casey) tolerated my whining, bolstered me when necessary, and patiently waited when I disappeared into the lab, library, or office for weeks or months at a time. Even when I was not the finest son, brother, or friend, they supported me without question, and I thank them all.

Contents

Abstract	iv
Acknowledgements	vi
List of Tables	xviii
List of Figures	xix
1 Introduction	1
1.1 Motivation for studies	3
1.1.1 High temperature superconductors	3
1.1.2 Quark-gluon plasma	6
1.1.3 Neutron stars	7
1.2 Significance of current work	8
1.2.1 Original studies of strongly interacting Fermi gases	9
1.2.2 Contributions to laboratory operation	10
1.3 Dissertation organization	12
2 Crossover physics & Feshbach resonances	16
2.1 The BEC-BCS Crossover	17
2.2 Low energy scattering	19
2.2.1 Unitary Fermi gases	22

2.3	Collisional Feshbach resonances	26
2.4	The electronic ground state of ${}^6\text{Li}$	28
2.5	The quest for superfluidity	36
3	General Experimental Methods	43
3.1	Basic cooling and trapping techniques	44
3.1.1	Preliminary cooling stages	45
3.1.2	The Magneto-Optical Trap (MOT)	45
	Basic physics of a MOT	45
	The ${}^6\text{Li}$ MOT	51
3.1.3	The far off-resonance dipole trap (FORT)	53
	Electric Dipole force	54
	FORT geometry	56
	Loading the FORT	57
3.1.4	Evaporative cooling	58
3.2	Experimental Apparatus	62
3.2.1	Double-pass AO arrangement	62
3.2.2	The vacuum chamber	63
3.2.3	The oven	64
3.2.4	The Zeeman slower	67
3.2.5	Optical beam generation	69
3.2.6	Frequency standard – the “locking” region	75
3.2.7	FORT beam generation and conditioning	78
3.2.8	Radio-frequency antenna	82

3.2.9	High field magnets	83
3.2.10	MOT bias coils	84
3.2.11	CCD camera and imaging system	84
3.2.12	Timing system	86
3.3	Typical experimental sequence	86
3.3.1	Production of a degenerate noninteracting Fermi gas	90
3.3.2	Production of a degenerate strongly interacting Fermi gas	94
4	Expansion Dynamics	96
4.1	The Relaxation Approximation	97
4.2	Ballistic expansion	99
4.3	Ballistic expansion with magnetic field	101
4.4	Hydrodynamic expansion with magnetic field	103
4.5	Expansion from a slightly anharmonic trap	105
4.6	Measuring trap oscillation frequencies	106
4.6.1	Measuring optical trap oscillation frequencies	107
4.6.2	Measuring magnetic field curvature	112
5	Image Processing	119
5.1	Image acquisition procedure	120
5.2	Image Analysis	124
5.3	Column density for an ideal two-level system	127
5.4	Optical Pumping	131
5.5	Correcting for probe beam depolarization	141
5.6	Measuring ϕ	143

5.7	Determining the saturation intensity I_{sat}	144
5.8	Final expression for the column density	147
5.9	Measuring the imaging system magnification	149
6	Heat Capacity of Fermi Gases	151
6.1	Energy input calculation	152
6.2	Experimental sequence	160
6.2.1	Preparation of the strongly interacting Fermi gas	161
6.2.2	Preparation of the noninteracting Fermi gas	162
6.3	Temperature measurement	164
6.3.1	Noninteracting Gases	167
6.3.2	Strongly interacting gases	172
	Temperature calibration	178
	Applicability of calibration technique away from resonance	181
	Other temperature measurement schemes in the strongly interacting region	182
6.3.3	Influence of trap anharmonicity on temperature measurement	183
	Generating anharmonic spatial densities	184
6.4	Results and discussion	189
6.4.1	Heat capacity of a noninteracting Fermi gas	190
6.4.2	Heat capacity of a strongly interacting Fermi gas	194
	Measurement of β	202
6.4.3	Virial theorem for a unitary Fermi gas	203
6.5	Conclusion	207

7	Breathing Mode Measurements	208
7.1	Exciting the breathing mode	209
7.2	Experimental sequence	212
7.3	The relaxation approximation and the radial breathing mode . . .	215
7.4	Hydrodynamic gases obeying a polytropic equation of state	223
7.4.1	Anharmonic corrections for a unitary Fermi gas	231
7.4.2	Anharmonic corrections for a noninteracting Fermi gas . . .	232
7.5	Magnetic field dependence of the breathing mode	233
7.6	Temperature dependence of the breathing mode at unitarity . . .	240
7.7	Conclusion	252
8	Conclusion	255
8.1	Chapter summary	256
8.2	Anticipated upgrades to the experimental apparatus	259
8.3	Outlook	260
A	Harmonically trapped Fermi gases	262
A.1	Sommerfeld expansion	263
A.2	Derivation of basic quantities	266
A.2.1	Density of states	266
A.2.2	Chemical potential	268
A.2.3	Average energy per particle	271
A.3	Density profiles for harmonically trapped noninteracting Fermi gases	273
A.3.1	Zero temperature profiles	274
A.3.2	Sommerfeld profiles	276

A.3.3	Finite temperature profiles	277
A.3.4	High temperature limit	279
B	Imaging at high magnetic field	281
B.1	Zeeman tuning of the ground and excited states	282
B.2	Determining the frequency-to-voltage conversion factor	286
C	Custom Software Upgrades	290
C.1	New usage of Stanford Pulse Generators	290
C.1.1	Error checking from within <i>Alter Channels.vi</i>	292
C.1.2	Use of stored configurations	294
C.2	<i>AutomateAcquisition.vi</i>	296
C.2.1	Intended usage of <i>AutomateAcquisition.vi</i>	297
C.2.2	Using <i>AutomateAcquisition.vi</i>	298
	To automate running of the system only	298
	To automate running the system and saving the images	302
C.2.3	Operation of <i>AutomateAcquisition.vi</i>	306
C.3	<i>Arbitrary Waveforms.vi</i>	309
C.3.1	Using <i>Arbitrary Waveforms.vi</i>	311
	Step-by-step instructions for using <i>Arbitrary Waveforms.vi</i>	318
	Adding lowering curves to the drop-down list	319
	Adding new timing files to the list in <i>Arbitrary Waveforms.vi</i>	320
C.3.2	Operation of <i>Arbitrary Waveforms.vi</i>	321
	Bibliography	327

List of Tables

3.1	Temperature profiles for atomic source oven	68
3.2	Summary of AO frequency shifts during different experimental phases	90
C.1	Correspondence between dummy variables in lowering curve Perl script and numeric values in timing file or front panel of <i>Arbitrary Waveforms.vi</i>	324

List of Figures

2.1	Schematic of Feshbach resonance phenomenon	27
2.2	Hyperfine energy levels of ${}^6\text{Li}$ ground state	32
2.3	Feshbach resonances in ${}^6\text{Li}$	34
3.1	Mechanism behind optical molasses	46
3.2	Spatial restoring force in the MOT	50
3.3	Double-pass AO arrangement	63
3.4	Simplified diagram of the vacuum chamber	65
3.5	Schematic of oven	66
3.6	Horizontal MOT beams entering vacuum chamber	73
3.7	Vertical MOT beam and probe beam entering vacuum chamber	74
3.8	Legend for optical layout diagrams	75
3.9	Optical layout for generating MOT beams	76
3.10	Optical layout for generating camera beam	77
3.11	CO_2 laser and related optics	79
3.12	CCD camera and imaging optics	85
4.1	Parametric resonance for radial trap oscillation frequencies	108
4.2	Parametric resonance for axial trap oscillation frequency	109
4.3	Atom cloud moving in magnetic potential	115
4.4	Center of mass position of atom cloud as function of time	116

5.1	Imaging procedure: signal and reference shots	121
5.2	Sample absorption image	125
5.3	Absorption image with balance and cloud regions of interest . . .	126
5.4	Sample image showing column density	129
5.5	Transition probabilities for on-resonance absorption imaging . . .	136
5.6	Estimate of error in detected atom number	140
5.7	Spatial profile showing amount of probe beam depolarization . . .	144
5.8	Camera magnification measurement	150
6.1	One-dimensional spatial distribution of an atom cloud	166
6.2	Thomas-Fermi fits to strongly interacting density profiles	175
6.3	Strongly interacting temperature calibration	180
6.4	Temperature measurement error arising from anharmonicity . . .	187
6.5	Energy vs. temperature for a noninteracting Fermi gas	193
6.6	Energy vs. temperature for strongly interacting Fermi gas	195
6.7	Energy vs. temperature for strongly interacting Fermi gas: low temperature	196
6.8	Energy vs. temperature for strongly interacting Fermi gas: ln-ln plot	198
6.9	Energy vs. temperature for strongly interacting and ideal Fermi gases	200
6.10	Heat capacity for strongly interacting Fermi gas	201
6.11	Mean square size vs. energy for unitary gas	206
7.1	Radial breathing mode absorption images	211
7.2	Radial breathing mode: width vs. time	212
7.3	Schematic of breathing mode excitation	214

7.4	Radial breathing mode locus plot	222
7.5	$\langle \tilde{x}^4 \rangle / \langle \tilde{x}^2 \rangle^2$ versus T/T_F for harmonically trapped Fermi gases . . .	230
7.6	Radial breathing mode frequency vs. magnetic field	235
7.7	Frequency and damping vs. magnetic field	236
7.8	Scheme for studying temperature dependence of radial breathing mode	244
7.9	Breathing mode frequency versus \tilde{T}	245
7.10	Damping time versus \tilde{T}	247
7.11	Damping time versus \tilde{T} for alternate conditions	251
7.12	Locus plot with unitary damping data	252
A.1	Three-dimensional harmonic oscillator density of states	268
A.2	Chemical potential for Fermi gas	271
A.3	Energy versus temperature for Fermi gas	272
A.4	One-dimensional density profiles	279
B.1	Fluorescence signal from the ${}^6\text{Li } D_2$ line	284
B.2	Frequency shift of imaging transitions versus magnetic field	287
C.1	<i>Alter Channels.vi</i> diagram panel showing SPG error check	293
C.2	<i>Arbitrary Waveforms.vi</i> front panel	313
C.3	Sample lowering curve for forced evaporation	315
C.4	Dual lowering curve sequence for forced evaporation	316

Chapter 1

Introduction

I suspect one could mark the passage of years by noting the annual scramble among university biology, chemistry, and physics departments to enlist majors from the limited pool of scientifically-minded freshman. Inevitably, in the course of lobbying, physics professors boast, “Physics is the fundamental science.” Biology and chemistry, the argument proceeds, could be rendered obsolete if we had adequate knowledge of the state of every constituent particle in a chemical reaction or a living organism. It is an extreme view, to be certain, as I never expect to explain why my family dog howled at firetrucks starting from Schrödinger’s equation. That said, the notion of studying “the fundamental science” is an appealing one. The argument can be taken one step further, however, for those who study fermions, the fundamental building blocks of matter. Unfortunately, basic studies of fundamental physical systems rarely garner much attention, as they are so far removed from potential applications in the “real world.” As a result, it might seem surprising that the field of strongly interacting degenerate Fermi gases has enjoyed so much exposure in popular science publications in recent years, and it is natural to question why. The answer, or some portion of it, requires a digression into quantum statistics.

All particles in the universe fall into two classes of matter which obey dif-

ferent quantum statistics: fermions and bosons. Typically, quantum mechanics textbooks run amok with personification when describing the differences between the two types of particles. Bosons are described as jovial, gregarious, and sociable, while fermions are characterized as misfits who absolutely, positively need their personal space. The analogies, while painful at times, have their merit. In a quantum world of discretized energy levels, multiple bosons are allowed to occupy the same quantum state, a privilege not extended to fermionic particles thanks to the Pauli exclusion principle. While this establishes the primary difference between bosons and fermions, it suggests no reason why the study of one would be more fundamental than the other. For this, we must draw on the concept of spin, which can be thought of as a classification scheme for quantum particles. Bosons have integer spin (0, 1, and so on) while fermions have half-integer spin ($1/2$, $3/2$, and so on). Perhaps most importantly for the purposes of this discussion, even numbers of fermions can be assembled to produce a boson.

The chest-thumping fermion researcher might argue that many bosonic systems are little more than a cluster of fermions in one place. Surely this is an unfair characterization of the many interesting studies of bosonic systems, but its essence explains the widespread interest in the field of degenerate Fermi gases. Fermions are the fundamental building blocks of matter, and so fermionic systems are ubiquitous in the universe. Not all of these fermionic systems are easily accessible or subject to human manipulation, however, and this is where the ultracold vapor of fermionic ${}^6\text{Li}$ atoms confined in our optical trap becomes relevant. In studying the comparatively simple system of a gas of fermions, potentially we can learn about many more exotic systems in nature, ranging from strongly interacting nuclear matter to high temperature superconductors. And, unlike some

of these other systems, we can exercise good control over many of the physical parameters of our system.

In Section 1.1, we consider in greater detail some of the motivation for studying strongly interacting Fermi gases. This is followed in Section 1.2 by a discussion of the significance of the research presented in this dissertation. Finally, Section 1.3 provides an outline of this dissertation.

1.1 Motivation for studies

Justifying the need for basic research is often a struggle, particularly in a society which focuses heavily on results and practical applications. As for the latter, one would be hard pressed to argue that studies of strongly interacting Fermi gases will yield substantial improvement in the average person's quality of life in the near future. While a hardcore pragmatist might have difficulty justifying this type of research, those of a more flexible nature will be pleased to learn that studies of strongly interacting Fermi gases are of interest to researchers beyond the atomic, molecular, and optical physics community (see for example [1–3]). Below, we discuss some of the areas of physical inquiry that have taken an interest in the research we conduct in our lab.

1.1.1 High temperature superconductors

In ordinary metals, the motion of electrons is impeded by frequent collisions with other matter. This gives rise to ohmic heating, a process that reduces the efficiency of the transmission of electricity. There is a class of materials, however, that undergoes a very unusual transition at low temperature to a state in which

the electrons flow without resistance. These metals are known as superconductors, and the economic consequences of developing a superconductor which operates above room temperature (≈ 295 K) could be staggering. Naturally, then, there is great interest in developing such a material.

The term “high temperature superconductor” might be a bit misleading to those unfamiliar with the history of the field, as the current high temperature superconductors do not make the transition to a superconducting state until they are cooled to temperatures roughly 160 °C below room temperature. Theoretical understanding of the operation of these materials remains incomplete, unfortunately. There are those who hope that the development of a dependable theory of high temperature superconductors might yield insights into methods for developing better superconductors (that is, superconductors with a higher critical transition temperature). It should come as no surprise, then, that there is also keen interest in the progress of the theory of superconductors.

Perhaps the most widely recognized theory of superconducting behavior is BCS theory, developed by Bardeen, Cooper, and Schrieffer in the late 1950’s ([4] provides a fine summary of the theory). In the theory, a collection of fermions, often referred to as the “Fermi sea,” allows a pair of spin-up and spin-down fermions that are weakly attracted to each other to form a delicate bound state. The paired fermions, known as Cooper pairs, are correlated primarily in momentum space (they have equal and opposite momenta), and their separation is often larger than the average interparticle spacing of fermions in the Fermi sea. If many of the fermions in the sea undergo pairing, then the aggregate collection of fermions will undergo a transition to a superconducting state, in which they flow without resistance. As the pairs are delicate objects, the superconducting

state can be a fairly fragile being as well, and both are easily broken by thermal fluctuations. Consequently, superconducting transitions tend to occur at rather low temperatures, where thermal fluctuations are suppressed.

This is all quite nice, but how does the physics of a *high temperature* superconductor relate to the physics governing an *ultracold* Fermi gas? Indeed, it might seem odd to claim the existence of parallels between materials that appear to operate at very different temperatures. In fact, there is only an apparent discrepancy. The physics of both systems are governed not by their absolute temperature T , but the ratio of their temperature to a characteristic temperature for that system. This characteristic temperature, known as the Fermi temperature T_F , is a property of the physical system which depends on the density of the fermions involved, among other things. For some of the best currently known high temperature superconductors, where a typical Fermi temperature is on the order of 10^4 kelvin, the critical transition temperature T_c to the superconducting state occurs at $T_c/T_F \approx 0.01$ [5]. In ultracold Fermi gases, we have a very similar scenario to what occurs in a superconductor. As the Fermi gas is cooled, the presence of attractive interactions between opposite spin fermions can lead to pairing and a phase transition to a superfluid state. Where electrons in a superconductor flow without resistance, fermions in a superfluid Fermi gas can flow without viscosity. Furthermore, in the presence of strong attractive interactions, the transition temperature to the superfluid state can be a substantial fraction of the Fermi temperature (see [6–10], for example), upwards of 0.25. Consequently, the behavior of a superfluid Fermi gas at reduced temperature $T/T_F = 0.25$ might behave like a metal superconductor operating at a temperature of several thousand degrees kelvin, well above room temperature. Studying the behavior of such a superfluid

might yield insights into the development of the highly sought-after above room temperature superconductor. More immediately accessible, however, is the use of data from studies of degenerate Fermi gases to test the accuracy of theories developed for superconductors. In metals and other condensed matter systems, researchers often exercise limited control over physical properties such as the interparticle spacing between fermions or the strength of their interactions. Ultracold Fermi gases are far more flexible, however, and the ability to produce such gases under different physical conditions provides an extensive battery of tests for theories of superconductivity. In Chapter 6, we report on the comparison between the measured heat capacity of a strongly interacting Fermi gas and the predictions of a theory originally developed for high temperature superconductors.

1.1.2 Quark-gluon plasma

Once believed to be elementary particles, baryons such as protons and neutrons are now known to be composed of smaller particles called quarks. Quarks belong to the class of fermionic particles, and under normal conditions, are bound to other quarks via the exchange of gluons. However, for very high temperature and density conditions, quarks and gluons can become unbound, creating a quark-gluon plasma (QGP). Such a condition is believed to have existed some tens of microseconds following the Big Bang at the beginning of the universe (see [11], for example).

Recently, researchers at the Relativistic Heavy Ion Collider (RHIC) succeeded in producing a QGP under controlled laboratory conditions (for a discussion of this somewhat controversial result, see [12]). A remarkable experimental achievement requiring 100 GeV/nucleon [13, 14], the resulting plasma at 2×10^{12} Kelvin

surprisingly shares some properties with our ultracold atomic gases. Early theoretical efforts predicting the existence of the QGP hinted that it would be a weakly interacting system [15]. However, data from RHIC suggested instead that the behavior of the system was more readily described by nearly perfect hydrodynamics, in which frequent collisions drive the dynamics of the system. In fact, the QGP has been touted as a nearly perfect fluid [15], given its hydrodynamic properties. The observation of extremely good hydrodynamics in a strongly interacting ultracold gas of ${}^6\text{Li}$ [16] was one of the primary reasons dialogue between ultracold Fermi gas and QGP researchers began. The anisotropic expansion of a strongly interacting ${}^6\text{Li}$ gas, first observed in 2002, resembles the “elliptic flow” of particles that occurs when heavy nuclei collide slightly off-axis [15]. Given the wide disparity in energy and temperature scales between ultracold Fermi gases and a QGP, it is surprising to note that researchers in both fields have undertaken similar calculations in an effort to explain the behavior of their respective systems. For instance, calculations of the viscosity and hydrodynamic behavior arising from binary collision processes have proceeded in much the same fashion in both fields [17].

1.1.3 Neutron stars

As a final example of the parallels between strongly interacting Fermi gases and other physical systems in the universe, we consider the case of neutron stars. Such stars are composed of very densely packed neutrons. The tremendous density of the stars and the associated gravitational attraction would be great enough to cause collapse of the star were it not for the Fermi pressure provided by the Pauli exclusion principle.

We have observed similar behavior in a strongly attractive Fermi gas [18], where the inward attractions are induced by a Feshbach resonance (see Section 2.3) and are opposed by the outward Fermi pressure. Comparison between the strongly attractive, degenerate Fermi gases produced in our laboratory and the strongly interacting matter in a neutron star is made possible by the universal nature of these systems. For very strongly interacting matter, the behavior of the system becomes independent of the details of how strong interactions are produced. In Section 2.2.1, we discuss a universal many-body parameter β which characterizes the strength of interparticle interactions in very strongly interacting matter. While the nuclear theory community has been performing calculations of β for several decades, the first measurements of β were conducted by atom cooling and trapping groups studying strongly interacting Fermi gases. This is discussed in greater depth in Section 2.2.1.

There is substantial excitement that the similarities between systems such as high temperature superconductors, the QGP, and neutron stars and strongly interacting degenerate Fermi gases will allow strongly interacting Fermi gases to serve as proving grounds for theories developed for other systems. Moreover, one could argue that the existence of dialogue between historically disparate subfields of physics is a worthwhile end in itself.

1.2 Significance of current work

In Section 1.2.1, I assess the importance of the original research on strongly interacting Fermi gases discussed in this dissertation. This research builds heavily on the efforts of many others. In Section 1.2.2, I outline my contributions to the

improved operation of the experimental apparatus used in our laboratory.

1.2.1 Original studies of strongly interacting Fermi gases

The two primary experimental efforts outlined in this dissertation are the measurement of the heat capacity in the strongly interacting regime (Chapter 6) and studies of the radial breathing mode throughout the crossover region (Chapter 7). Our study of the heat capacity [10] is the first thermodynamic measurement of a strongly interacting Fermi gas, and the first direct measurement of the superfluid transition temperature in the strongly interacting regime. This measurement required the development of new experimental techniques, including an approximate scheme for measuring the temperature of the gas in the strongly interacting regime. Temperature measurement in the presence of strong interactions is a fairly controversial topic, and while we do not suggest that our scheme is the ultimate solution, we do believe it is a practical alternative to other temperature measurement schemes, as discussed in Section 6.3.

Our exhaustive studies of the radial breathing mode in a strongly interacting Fermi gas, discussed in Chapter 7, also yielded important insights. As the breathing mode effectively tests how the gas holds itself together while being stretched and compressed, it provides a convenient test of the equation of state of the gas. As researchers continue to develop better theories to describe the behavior of strongly interacting matter, simple tests of these theories are of great importance. Measurements of the breathing mode constitute one such test.

Our initial study of the temperature dependence of the radial breathing mode lifetime [19] provided the first evidence for superfluid hydrodynamics in a strongly interacting Fermi gas. This was followed by a more definitive study of the temper-

ature dependence of the radial breathing mode [20], which verified the conclusions of our initial study and provided several new insights. This second study revealed unexpected features in the damping rate of the breathing mode as a function of temperature which have been interpreted as possible indicators of a superfluid phase transition and the breaking of noncondensed atom pairs. We also conducted exhaustive measurements of the magnetic field dependence of the radial breathing mode for a low temperature gas. The results of this study were in general agreement with predictions based on superfluid hydrodynamics. Unexpected increases in the damping rate and frequency above resonance, first observed in [21], have been interpreted in terms of pair-breaking.

1.2.2 Contributions to laboratory operation

While a good deal of effort was expended in acquiring the data for this dissertation, none of this would have been possible without the substantial efforts of my predecessors. Of particular note are the contributions of Ken O'Hara, Stephen Granade, Michael Gehm, and Staci Hemmer, who are responsible for the design and construction or purchase and installation of nearly every element presently in use on the system. All of the laser systems, basic optical design choices, vacuum system elements, and other myriad components were in place when I began working in the lab. Moreover, all of the necessary computer and software equipment had been developed and debugged. Very simply, the credit for the construction of the system should be directed to the aforementioned people. Without them, the work presented in this dissertation would not have been possible.

The majority of my time in the lab was consumed by making particular processes more efficient or precise, and increasing the reliability of the system. In

my early days in the lab, I was guided in these efforts by Staci Hemmer. Shortly thereafter, we were joined by postdoctoral researcher Andrey Turlapov. Together, we simplified a number of the optical designs on the system, and where possible, replaced unreliable electronic circuits and power supplies with more dependable pieces of equipment. Improving the reliability of the system also required upgrades in the custom software used to control the experimental apparatus. Of the contributions I made to the laboratory from an operational standpoint, I consider the last of these to be the most substantial.

Some of the studies presented in this dissertation required that particular events occur with microsecond precision. While all of the physical components which comprise the experimental apparatus were capable of functioning with microsecond precision, not all of the various microsecond time delays were accounted for in the software which controls the system. These delays have now been documented in the “timing files” (scripts which control the activity of the experimental apparatus). Another important upgrade to the software was the development of an error checking scheme for the DG535 Stanford pulse generators which we use to control some of our most time-sensitive tasks. The general programming of the Stanford pulse generators also underwent an overhaul, as the use of configurations stored in nonvolatile memory in the pulse generators now prevents some errors which could have occurred when the user switched between different timing files. These upgrades are discussed in Appendix C.1.

Perhaps my single biggest contribution to the laboratory’s custom software library was the development of a LabVIEW program which automates the process of taking and saving data. As the reliability of the experimental system improved, it became apparent that the system could be run for hours at a time with minimal

attention from the user. To capitalize on this new degree of reliability, a program was developed which allows the user to vary a single experimental parameter over a user-defined range of values, saving data at each value of the parameter. Further, the system was designed to randomize the order in which the data is acquired, thereby lessening the effects of systematic drifts over the course of a day. The software that handles this is discussed in Appendix C.2.

The final software upgrade of note, described in Appendix C.3, involves the development of a LabVIEW file which enables the user to command the behavior of particular GPIB programmable devices involved in the final stage of the atom cooling process. This upgrade did not provide any further functionality to the laboratory, but it did automate a number of tedious processes which formerly were handled manually.

The common thread among all of the upgrades summarized above is that they enabled the experimental system to operate more reliably and with less attention from the user. And while the data in this dissertation could conceivably have been acquired without these upgrades, the process would have been substantially more painful.

1.3 Dissertation organization

Chapter 2 covers most of the basic theoretical background required for an understanding of the relevance of the experimental results discussed elsewhere in this dissertation. Where possible, I have attempted to provide simple physical descriptions, and where simple descriptions are unavailable, I have tried to direct the reader to more demanding treatments of the subject matter. In Chapter 2,

I also present a brief summary of the major experimental results from the sub-field of degenerate, strongly interacting Fermi gases over the past several years. It is hoped that such a summary will help to place the studies described in this dissertation in a larger context.

Before delving into the novel physical studies described in Chapters 6 and 7, we must first cover some basic atom cooling and trapping methods. Chapter 3 summarizes the basic experimental techniques currently employed in our laboratory, and provides a general description of the experimental apparatus as well. Much of this chapter is similar to content contained in previous theses from this group [22–24], and is included here for completeness. Furthermore, some of the experimental techniques in use have evolved in complexity from previous theses, and these areas are discussed in greater detail.

All of the data presented in this dissertation was acquired in the form of absorption images of the atom clouds. Due to the small radial size of the trapped atom cloud, reliable imaging requires that we release the cloud from the trap and allow it to expand to a size much larger than the resolution of our imaging system prior to imaging the cloud. In order to infer information about the trapped cloud, we need to understand how the cloud expands upon release from our optical trap. In Chapter 4, we consider the expansion dynamics of gases in the noninteracting and hydrodynamic regimes. This is followed in Chapter 5 by a thorough discussion of the techniques used for probing the atom cloud and processing the resulting images.

In Chapter 6, we report our studies of the heat capacity of a strongly interacting Fermi gas, the first thermodynamic study of such a system. In the chapter, I describe the novel energy input technique and temperature measurement meth-

ods used to study the heat capacity. Our understanding of our results was aided greatly by a collaboration with Kathy Levin's theory group at the University of Chicago. Included in this chapter is a very basic description of their theory along with a comparison of their theory and our data.

Using experimental techniques very similar to the ones that make a study of the heat capacity possible, we can excite the radial breathing mode of our strongly interacting Fermi gas, leading to the studies discussed in Chapter 7. We have conducted fairly exhaustive studies of the breathing mode, considering both the dependence on interaction strength and temperature. In this chapter, I will describe the techniques used to excite and then monitor the evolution of the breathing mode. Interpretation of the sometimes complex results of these studies is provided.

Chapter 8 offers a conclusion to this dissertation and provides an outlook for future studies of strongly interacting degenerate Fermi gases. In spite of the tremendous experimental progress provided by a number of research groups worldwide, there is still much to learn about this very physically rich system.

Finally, there are three appendices included in this dissertation. Although many of our studies focus on the behavior of *strongly interacting* harmonically trapped Fermi gases, our understanding of these results often relies on concepts drawn from theoretical treatments of *noninteracting* harmonically trapped Fermi gases. As a result, Appendix A presents some basic results from the area of noninteracting Fermi gases. This appendix is written at a basic level and is intended as a handy reference for future researchers in our laboratory. Appendix B covers some of the issues relevant to the generation of an on-resonance probe pulse needed for high field imaging, a technique we use to take pictures of our

gas. Finally, Appendix C provides a description of the custom software upgrades I implemented during my time in the group. This section will be of limited interest to anyone outside our laboratory, save those who enjoy tedious technical descriptions of software.

Chapter 2

The BEC-BCS crossover and Feshbach resonances

Much of the appeal of studying ultracold gases resides in the excellent degree of control we exercise when manipulating them. While we can readily control the density and temperature of the gases, perhaps the most useful “knob” to turn is the strength of the interactions between the particles which comprise the gas. In this chapter, we explore the basic physics of the so-called “BEC-BCS crossover,” which explores the transition between Bose-Einstein and Fermi statistics as the nature and strength of interparticle interactions is varied. We will find that this realm of very rich physical behavior is made accessible by Feshbach resonances which can greatly enhance the rate at which collisions occur in the gas. Some basic scattering theory will be presented, and we will consider the special case of a unitary Fermi gas. Much of the discussion will focus on broad physical concepts, but toward the end of the chapter we will focus on issues specific to the ${}^6\text{Li}$ atom. In Section 1.2.1, I provided a brief description of the significance of the original research presented in this dissertation. In an attempt to position these results in a larger hierarchy of progress by the field as a whole, I include a discussion of the major experimental results from the past several years in Section 2.5. We begin, however, with a basic discussion of the BEC-BCS crossover.

2.1 The BEC-BCS Crossover

A qualitative description of the BEC-BCS crossover is perhaps most easily given by examining the two extremes in the crossover spectrum. BEC stands for “Bose-Einstein condensation,” a phenomenon first predicted by Einstein early in the twentieth century [25].¹ In Bose-Einstein condensation, as a group of bosonic particles is cooled below a critical temperature, many of the bosons collapse into the ground state of the system. As BECs amount to macroscopic occupation of the lowest quantum state available, the phenomenon obviously cannot be observed with fermionic particles, as the Pauli exclusion principle prevents fermions with the same quantum numbers from occupying the same quantum state. However, recall that even numbers of fermions can bind together to form a composite boson. In that case, paired fermions with small interparticle separation will behave like bosons and can undergo Bose-Einstein condensation. Such molecules can be formed when weak repulsive interactions exist between two spin components of a Fermi gas.

Suppose, however, that we have a two-component mixture of fermions which is governed by weak attractive interactions between the two components. This is the BCS limit. BCS theory was introduced in Section 1.1.1, but it is worth summarizing again here. Under the right conditions, two fermions can become weakly paired via correlations in momentum space. In order for a shallow bound state to appear, the presence of other fermions (often referred to as the “Fermi sea”) is necessary. It is also critical that the interactions between the fermions be

¹At the time of Einstein’s prediction, the technology did not exist to produce such a state in a controlled laboratory setting. In fact, Bose-Einstein condensation was not directly observed in bosonic gases until 1995 [26–28], an experimental achievement which garnered the 2001 Nobel prize in physics.

attractive rather than repulsive. Finally, due to the weak nature of the coupling between the components of the pair, the paired state can be quite delicate. Pairs can be broken easily by thermal fluctuations, and so the phenomenon of pairing tends to occur at low temperatures. If a substantial fraction of the fermions in the gas form pairs, then the gas can undergo a transition to a superfluid state, analogous to the transition to a superconducting state in a superconductor.

In summary, BECs can form in fermionic systems provided pairs of fermions with small interparticle separation form, and act as composite bosons. As with atomic BECs, BECs of molecules composed of fermions can form in a sufficiently cold system governed by weak repulsive interactions. In the other extreme, a BCS-type superfluid can form in a sufficiently cold system in which there exist weak attractive interactions between the fermions. Suppose, however, we have a physical system in which the nature (attractive versus repulsive) and strength of interparticle interactions can be continuously tuned between the BCS and BEC regimes. In this case, we can investigate the “BEC-BCS crossover.” This is precisely the type of system we can generate in our laboratory. We can explore the crossover from Fermi to Bose statistics by manipulating the interactions between two spin components of a Fermi gas. All of this is made possible by the use of Feshbach resonances, to be discussed in Section 2.3. Before delving into the many useful features of a Feshbach resonance, it will be helpful to review some basic scattering theory.

2.2 Low energy scattering

Scattering processes involving two quantum particles are covered in nearly every elementary quantum mechanics textbook (examples include [29–32]) and in previous theses from this group (see [22, 24] for some of the more comprehensive treatments). Consequently, only the basics of two body scattering will be reviewed here. In the discussion that follows, we will use the notation conventions in [31], though the conventions are much the same in any standard quantum mechanics textbook.

Consider a stream of particles impinging on a target. If we suppose that the incident stream of particles is, more or less, a plane wave with wave vector \mathbf{k} , and if we suppose that the scattered particles look roughly like a spherical wave with outgoing spherical wave amplitude $f(\mathbf{k}', \mathbf{k})$, where \mathbf{k}' represents the wave vector of the outgoing wave, then the outgoing spherical wave amplitude is related to the differential scattering cross section $d\sigma/d\Omega$ by

$$\frac{d\sigma}{d\Omega} d\Omega = |f(\mathbf{k}', \mathbf{k})|^2 d\Omega. \quad (2.1)$$

The quantity on the left hand side of (2.1) represents the ratio of particles scattered into a detector subtending solid angle $d\Omega$ per second to the number of particles per unit area per second impinging on the detector.

If the scattering potential provided by the target is spherically symmetric, and if we let θ be the angle between the incident and outgoing wave vectors \mathbf{k} and \mathbf{k}' , then we can expand the outgoing wave in partial waves as

$$f(\mathbf{k}', \mathbf{k}) = f(\theta) = \sum_{l=0}^{\infty} (2l+1) f_l(k) P_l(\cos \theta), \quad (2.2)$$

where l indicates the order of the angular momentum of the outgoing wave and $P_l(x)$ are Legendre polynomials. The partial wave amplitudes $f_l(k)$ are related to the scattering phase shifts δ_l by [31]

$$f_l(k) = \frac{\exp(i \delta_l) \sin \delta_l}{k}. \quad (2.3)$$

The total scattering cross section can be calculated using (2.1), (2.2) and (2.3):

$$\sigma_{tot} = \frac{4 \pi}{k^2} \sum_{l=0}^{\infty} (2l + 1) \sin^2 \delta_l. \quad (2.4)$$

Rewriting the scattering cross section in this manner might seem like a step in the wrong direction, as it suggests we must determine the value of an infinite number of partial wave phase shifts δ_l . Matters simplify substantially for ultracold gases, however, where p-wave and higher order angular momentum processes ($l > 0$) are largely suppressed. Before presenting a simple argument as to why this is the case, however, suppose that instead of considering the scattering of many particles from a target, we focus on the process of one particle scattering from another. The above formalism still holds if we suppose that the individual particles are represented by wavepackets.

We now consider a simple argument for the dominance of s-wave scattering processes in our laboratory. Under typical experimental conditions, the temperature of the particle is often 1 μ K or lower, which corresponds to a de Broglie wavelength λ_{dB} of roughly 700 nm. This corresponds to a linear momentum of $p_{max} = h/\lambda_{dB} = 9.5 \times 10^{-28}$ m/s. Further, if we consider an interparticle potential with an effective range $r_0 = 20$ bohr, a reasonable approximation for the electronic triplet potential which dominates the interaction of the particles in our

gas, then the maximum quantized angular momentum l_{max} associated with an interaction between two ${}^6\text{Li}$ atoms at $1 \mu\text{K}$ is much less than one:

$$\hbar l_{max} = r_0 p_{max} \Rightarrow l_{max} = \frac{2\pi r_0}{\lambda_{dB}} \approx 0.001. \quad (2.5)$$

Since l must be a nonnegative integer, the primary interactions between ultracold atoms at the temperatures we have considered must involve $l = 0$, or s-wave, processes. In this case, the infinite sum in (2.4) reduces to a single term,

$$\sigma_{tot} = \frac{4\pi \sin^2 \delta_0}{k^2}. \quad (2.6)$$

That is, if s-wave collisions dominate, we need to calculate only the lowest order scattering phase shift δ_0 .

Trigonometric identities yield

$$\sin^2 \delta_0 = \frac{\tan^2 \delta_0}{1 + \tan^2 \delta_0}, \quad (2.7)$$

and (2.7) can be rewritten using the definition of the low energy s-wave scattering length [30]

$$a_s \equiv -\lim_{k \rightarrow 0} \frac{\tan \delta_0}{k}, \quad (2.8)$$

which in turn permits us to express the scattering cross section in terms of the incident wave vector k and the s-wave scattering length a_s as

$$\sigma_{tot} = \frac{4\pi a_s^2}{1 + k^2 a_s^2}. \quad (2.9)$$

For very low energy collisions, $k \rightarrow 0$ and the scattering cross section becomes

$$\sigma_{tot} \approx 4\pi a_s^2. \quad (2.10)$$

Equation (2.10) is an important result, as it suggests that low energy scattering processes can be effectively parameterized by a single number, the s-wave scattering length a_s . If the magnitude of a_s is large, then the interaction between the scattering particles is strong. Conversely, small values of a_s indicate weak interparticle interactions. Furthermore, the sign of the s-wave scattering length indicates whether the effective potential associated with the scattering process is repulsive ($a_s > 0$) or attractive ($a_s < 0$).

When we discuss the operation of Feshbach resonances in Section 2.3, we will rely on the s-wave scattering length as a means of characterizing the nature of interparticle interactions throughout the BEC-BCS crossover region. Before doing so, however, we consider a special case in which we reach the quantum mechanical limit of strong interactions. That is, we consider a unitary Fermi gas.

2.2.1 Unitary Fermi gases

We saw in Section 2.2 that the strength of a low energy collisional process could be expressed in terms of the s-wave scattering length. What happens if we let $a_s \rightarrow \pm\infty$? Regardless of the sign of the interaction, (2.9) indicates that

$$\lim_{a_s \rightarrow \pm\infty} \sigma_{tot} = \frac{4\pi}{k^2}. \quad (2.11)$$

That is, the collisional behavior of the system effectively becomes independent of the sign and strength of the interparticle interactions. When these conditions exist in our gas of fermions, we have created a unitary Fermi gas.

To make further progress with the concept of unitarity, we must step back for a moment and consider a very different case, that of the noninteracting Fermi gas. For a gas confined in a potential $U(r)$, we can write the equation of state as [33]

$$\frac{\hbar^2 k_F(r)^2}{2m} + U(r) = \mu, \quad (2.12)$$

where $k_F(r)$ represents the position dependent local Fermi wave vector and μ is the global chemical potential of the system. If we were to introduce weak interparticle interactions into the system, then (2.12) could be modified [34, 35] to include an interaction energy term U_{int} proportional to the product of the scattering length a_s and the particle density $n(r)$. The latter of these quantities is related to the local Fermi wave vector via $n(r) \propto k_F(r)^3$. As the strength of the interparticle interactions was increased, however, the mean field formalism would cease to be valid, and the s-wave scattering length would need to be replaced by an effective scattering length a_{eff} . Eventually, the scattering cross section would reach the limit provided by (2.11), where the lone remaining length scale in the system is the local Fermi wave vector, $k_F(r)$. By dimensional analysis, then, we require that $a_{eff} \propto 1/k_F(r)$. But, if $U_{int} \propto a_{eff} n(r)$, then the respective dependence of a_{eff} and $n(r)$ on $k_F(r)$ yields $U_{int} \propto k_F(r)^2$. Since the local Fermi energy $\epsilon_F(r) = \hbar^2 k_F(r)^2/(2m)$, we can rewrite the interaction energy term as [16, 36]

$$U_{int} = \beta \epsilon_F(r), \quad (2.13)$$

where β is a dimensionless constant. Then, in the unitarity limit, the zero temperature equation of state is

$$(1 + \beta) \frac{\hbar^2 k_F(r)^2}{2m} + U(r) = \mu^*. \quad (2.14)$$

Note that we have placed an asterisk on the global chemical potential in (2.14) to distinguish it from the global chemical potential in the noninteracting case. The parameter β is a universal many-body parameter which gives the ratio of the interaction energy to the local Fermi energy [36].

There is a simple interpretation of the equation of state given by (2.14). Suppose we define an effective mass m^*

$$m^* = \frac{m}{1 + \beta}, \quad (2.15)$$

where m is the bare mass appearing in (2.14). Then we can rewrite (2.14) as

$$\frac{\hbar^2 k_F(r)^2}{2m^*} + U(r) = \mu^*. \quad (2.16)$$

Very simply, the equation of state for a unitary gas is the same as the equation of state for a noninteracting gas with a scaled mass. When we discuss the heat capacity of a unitary Fermi gas in Chapter 6, we will be interested in the specific case of a harmonically trapped unitary gas. For a zero temperature harmonically trapped *noninteracting* Fermi gas with $N/2$ particles per spin state, the global chemical potential μ is equal to the Fermi energy at the center of the trap (see Appendix A.2.2)

$$\mu = \varepsilon_F = \hbar \bar{\omega} (3N)^{1/3}, \quad (2.17)$$

where $\bar{\omega} = (\omega_x \omega_y \omega_z)^{1/3}$ is the geometric mean of the trap oscillation frequencies. Since $\bar{\omega} \propto \sqrt{1/m}$, we can write the scaled geometric mean of the trap oscillation frequencies as $\bar{\omega}^* = \sqrt{1 + \beta} \bar{\omega}$. Consequently, the chemical potential for the zero temperature unitary case is

$$\mu^* = \sqrt{1 + \beta} \mu = \sqrt{1 + \beta} \varepsilon_F, \quad (2.18)$$

where again ε_F is the global Fermi energy for a noninteracting gas.

There are interesting consequences to having interaction energy of the form given by (2.13). Since the scattering length is effectively infinite, it cannot appear in any physically meaningful terms. Likewise, (2.14) is independent of the details of the interparticle interaction potential (recall that $U(r)$ arises from an external potential which confines all of the particles). Consequently, we have written an equation of state which is independent of the details of the interparticle interactions. In other words, unitarity limited systems obeying (2.14) should behave in similar fashion regardless of the way in which strong interactions form in the system. Indeed, the nuclear theory community has been performing calculations of β , relevant to strongly interacting nuclear matter, for some time (see [37], for example). However, experimental verification of the value of β , which recent theoretical investigations [38–40] suggest is roughly $\beta = -0.55$, waited until researchers of strongly interacting Fermi gases began measuring the quantity [10, 16, 18, 41–43]. This is one example of universal behavior, or universality [44, 45], in which the details of interparticle interactions become unimportant in the unitarity limit, permitting the comparison of strongly interacting matter of many different ilk.

In closing, we offer one caution about the preceding formalism. When dis-

cussing collisional processes, we focused on the binary collision model, in which two particles are involved in an interaction process. As a result of the Pauli exclusion principle, we would expect binary collision processes to play a dominant role in two-component Fermi mixtures, where three particle processes would be unlikely. However, in the strongly interacting regime, some of the experimental data are difficult to explain in terms of binary collision models. As a result, we caution against relying too heavily on the two-body collision picture, particularly in the regime of strong interactions.

2.3 Collisional Feshbach resonances

Having explored the basics of the BEC-BCS crossover in Section 2.1 and basic scattering theory in Section 2.2, we turn now to the physical mechanism which makes the study of weak and strong interactions in our gas possible, the collisional Feshbach resonance.

Consider colliding particles interacting via the potentials depicted in Figure 2.1. In the Figure, the combined energy of the colliding particles is represented by the solid horizontal line. The particles have sufficient energy to interact via the open channel represented by the solid curve, but do not have sufficient energy to interact via the potential represented by the dashed curve. (Note that the energy of the colliding particles is below the mouth of the dashed curve.) This latter potential represents an energetically forbidden, closed collision channel. Under normal circumstances, the particles will interact via the open collision channel, and the closed channel will have little or no impact on the process. However, if the total energy of the colliding particles could be raised or lowered such that it

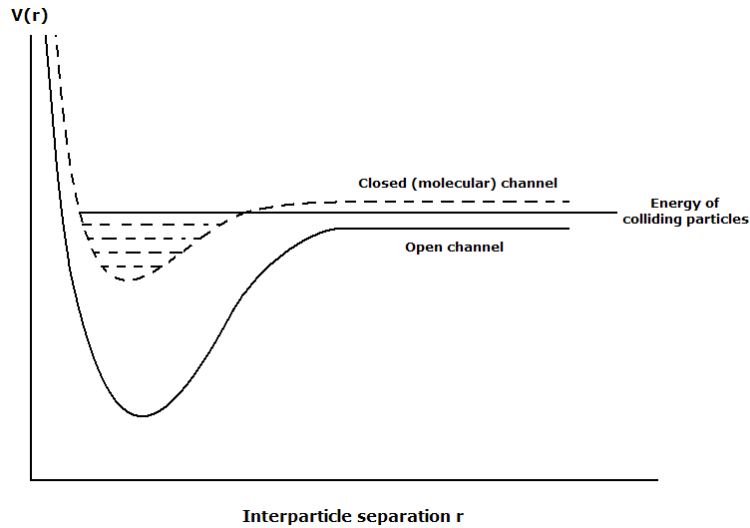


Figure 2.1: A collisional Feshbach resonance arises when the energy of two colliding particles in an open collisional channel coincides with the energy of a bound state in a closed molecular channel. In the above diagram of potential energy $V(r)$ versus interparticle separation r , the solid curve represents the potential energy in the open channel, the solid horizontal line represents the energy of the colliding particles, the dashed curve represents the potential in the closed (molecular) channel, and the horizontal dashed lines represent bound molecular states. When the energy of the colliding particles coincides with a bound state in the closed channel, there is a resonant enhancement of the interparticle collision rate and the scattering length diverges.

matched the energy level of a bound state in the closed channel, then the scattering process in the open channel is greatly enhanced. This type of resonance is often referred to as a Feshbach resonance [46]. The scenario depicted in Figure 2.1 is a generic representation of the Feshbach resonance phenomenon. For the states of ${}^6\text{Li}$ which we study, the triplet potential (“open channel”) is very shallow compared to the singlet potential (“closed channel”).

The strong enhancement of the scattering rate near a Feshbach resonance is manifested in the form of a diverging scattering length. Often, control over the

energy of the colliding particles is exercised by the application of a magnetic field which Zeeman tunes the energy of the colliding particles in the open triplet channel into resonance with a bound state in the closed singlet channel.² In the vicinity of a Feshbach resonance, the general form for the scattering length as a function of magnetic field B is given by [47]

$$a_s(B) = a_b \left(1 + \frac{\Delta}{B - B_0} \right), \quad (2.19)$$

where a_b represents the background scattering length, B_0 is the magnetic field which marks the center of the Feshbach resonance, and Δ is the width of the resonance. Plotting (2.19) as a function of magnetic field reveals that the Feshbach resonance allows us to continuously explore the behavior of the colliding particles as they are tuned from weak to strong attractive interactions, before changing abruptly to strong repulsive and then weak repulsive interactions. In Section 2.4, we will discuss two important Feshbach resonances in ${}^6\text{Li}$ which allow us to study the BEC-BCS crossover. A plot of the scattering length versus magnetic field for these resonances is given in Figure 2.3.

2.4 The electronic ground state of ${}^6\text{Li}$

While the majority of this chapter has focussed generally on two-component Fermi systems, we now consider specifically the two-component Fermi system we use in our laboratory. As mentioned previously, we work with an isotope of lithium, ${}^6\text{Li}$. The ${}^6\text{Li}$ atom is composed of 3 protons, 3 neutrons, and 3 electrons, yielding an

²In the singlet potential, the electronic spins of the two atoms are antiparallel. In the triplet potential, the electronic spins are parallel. Consequently, the energy of the triplet potential can be Zeeman tuned while the singlet potential remains stationary.

overall charge neutral atom. The nuclear ground state has nuclear spin $I = 1$, while the electronic ground state consists of 2 electrons in the $1s$ orbital and the final, unpaired electron in the $2s$ orbital. Throughout the remainder of this thesis, “the electron” refers to the unpaired electron in the valence band. In the electronic ground state, the orbital angular momentum quantum number $L = 0$, while the intrinsic spin of the electron is $S = 1/2$. As we will need to consider the hyperfine structure of the ${}^6\text{Li}$ atom, it is worthwhile to establish now the total angular momentum quantum number F arising from the nuclear, orbital, and electron spins. In the electronic ground state ($I = 1$, $L = 0$, and $S = 1/2$), angular momentum addition gives two possible values of the total angular momentum, $F = 3/2$ and $F = 1/2$.

In the absence of a magnetic field, the $F = 1/2$ manifold of the ground state is two-fold degenerate, with spin projections $m_F = \pm 1/2$, while the $F = 3/2$ level has four-fold degeneracy, where $m_F = \pm 3/2, \pm 1/2$. The application of an external magnetic field lifts this degeneracy, and six distinct eigenstates emerge. We wish to solve for the eigenstates of the internal Hamiltonian

$$H_{int} = \frac{a_{hf}}{\hbar^2} \mathbf{S} \cdot \mathbf{I} + \frac{\mu_B}{\hbar} \left(g_J^{gnd} \mathbf{S} + g_I \mathbf{I} \right) \cdot \mathbf{B}, \quad (2.20)$$

where $a_{hf}/\hbar = 152.137$ MHz is the magnetic dipole constant and $g_J^{gnd} = 2.002$ is the total electronic g-factor for the ${}^6\text{Li}$ ground state, $g_I = -0.000448$ is the total nuclear g-factor, μ_B is the bohr magneton, and B is the external magnetic field.³

³The reader is encouraged to peruse Appendix A in [24] and references therein for a more thorough discussion of the interaction of the ${}^6\text{Li}$ atom with electromagnetic fields.

We find that the eigenstates of (2.20) are [48]:

$$|1\rangle = \sin \Theta_+ |1/2 0\rangle - \cos \Theta_+ |-1/2 1\rangle \quad (2.21)$$

$$|2\rangle = \sin \Theta_- |1/2 -1\rangle - \cos \Theta_- |-1/2 0\rangle \quad (2.22)$$

$$|3\rangle = |-1/2 -1\rangle \quad (2.23)$$

$$|4\rangle = \cos \Theta_- |1/2 -1\rangle + \sin \Theta_- |-1/2 0\rangle \quad (2.24)$$

$$|5\rangle = \cos \Theta_+ |1/2 0\rangle + \sin \Theta_+ |-1/2 1\rangle \quad (2.25)$$

$$|6\rangle = |1/2 1\rangle, \quad (2.26)$$

where the basis kets $|m_S m_I\rangle$ give the electronic spin projection m_S and the nuclear spin projection m_I . The hyperfine states are labelled according to increasing energy, with $|1\rangle$ being the eigenstate with the lowest energy eigenvalue, and $|6\rangle$ possessing the largest energy eigenvalue. The coefficients in (2.21) through (2.26) are magnetic field dependent and obey the following relations

$$\sin \Theta_{\pm} = \frac{1}{\sqrt{1 + (Z^{\pm} + R^{\pm})^2 / 2}} \quad (2.27)$$

$$\cos \Theta_{\pm} = \sqrt{1 - \sin^2 \Theta_{\pm}} \quad (2.28)$$

$$Z^{\pm} = \frac{\mu_B B}{a_{hf}} (g_J^{gnd} - g_I) \pm \frac{1}{2} \quad (2.29)$$

$$R^{\pm} = \sqrt{(Z^{\pm})^2 + 2}. \quad (2.30)$$

The energy eigenvalues E_n associated with the eigenstates $|n\rangle$ are

$$E_1 = -\frac{1}{4} (a_{hf} - 2 g_I \mu_B B + 2 a_{hf} R^+) \quad (2.31)$$

$$E_2 = -\frac{1}{4} (a_{hf} + 2 g_I \mu_B B + 2 a_{hf} R^-) \quad (2.32)$$

$$E_3 = \frac{a_{hf}}{2} - \frac{\mu_B B}{2} (2 g_I + g_J^{gnd}) \quad (2.33)$$

$$E_4 = \frac{1}{4} (-a_{hf} - 2 g_I \mu_B B + 2 a_{hf} R^-) \quad (2.34)$$

$$E_5 = \frac{1}{4} (-a_{hf} + 2 g_I \mu_B B + 2 a_{hf} R^+) \quad (2.35)$$

$$E_6 = \frac{a_{hf}}{2} + \frac{\mu_B B}{2} (2 g_I + g_J^{gnd}). \quad (2.36)$$

We can readily plot the energy eigenvalues as a function of magnetic field. Due to the smallness of the energy scale, it is useful to plot the energy levels in frequency units, using $\mu_B/h = 1.3996$ MHz/gauss. This is done in Figure 2.2, where the energy eigenvalues have been labelled by their associated eigenstate $|n\rangle$.

In our laboratory, the initial stages of our experimental sequence are designed to produce a 50-50 mixture of the $|1\rangle$ and $|2\rangle$ hyperfine ground states. Why are these states preferable to other two-component combinations of hyperfine states? There are several reasons. First, there are two Feshbach resonances in the $|1\rangle$ - $|2\rangle$ mixture which occur at experimentally accessible magnetic fields. These resonances are depicted in Figure 2.3. The broad Feshbach resonance centered at 834 gauss [49] is the resonance we primarily use. It is quite broad (several hundred gauss) and is the resonance of choice for many researchers working with ${}^6\text{Li}$. Above the center of the Feshbach resonance, attractive interactions of arbitrary strength can be generated. Below the center of the Feshbach resonance, the scattering length is positive, enabling us to create repulsive interactions of arbitrary strength.




Figure 2.2: The hyperfine energy eigenvalues of the ${}^6\text{Li}$ ground state are plotted in frequency units versus applied magnetic field in gauss. There are six energy levels at nonzero magnetic field, and convention dictates that the eigenstates are labelled $|1\rangle$, $|2\rangle$, and so on, in order of increasing energy. For reasons discussed in the main text, we use a mixture of the $|1\rangle$ - $|2\rangle$ states in our lab. At the left, we note the total angular momentum quantum number ($F = 1/2$ or $F = 3/2$) for each of the states at zero magnetic field.

Although the true BEC and BCS limits correspond to weakly interacting systems, the high side of the resonance ($B > 834$ gauss) is often referred to as the “BCS side” of the resonance and the low side of the resonance ($B < 834$ gauss) is referred to as the “BEC side” of the resonance, even when discussing the strongly interacting regime. Depicted in the inset of Figure 2.3 is the position of the narrow Feshbach resonance near 544 gauss [50]. This resonance is extremely narrow, and has not been used for the crossover studies discussed in this dissertation. Below the narrow Feshbach resonance, however, is a feature known as the “zero crossing” [51, 52], where the scattering length $a_s = 0$. This feature is important for the production of noninteracting Fermi gases, as discussed in Section 3.3.1.

A second reason for studying the $|1\rangle$ - $|2\rangle$ mixture is that it is an energetically stable combination of states. In the ground state of ${}^6\text{Li}$, the total angular momentum spin projection $m_F = m_S + m_I$ is conserved in s-wave collisions. If we ignore energy considerations for a moment, a collision of atoms in the $|4\rangle$ ($m_F = -1/2$) and $|1\rangle$ ($m_F = 1/2$) states could yield atoms in the $|1\rangle$ ($m_F = 1/2$) and $|2\rangle$ ($m_F = -1/2$) states following the collision. A quick glance at Figure 2.2 reveals, however, that such a collision would result in the conversion of internal energy to kinetic energy (since the combined energy of the $|4\rangle$ and $|1\rangle$ states is greater than that of the $|1\rangle$ and $|2\rangle$ states). But, how much energy is released? Close to zero magnetic field, the total energy released would be roughly $h \times 228$ MHz, or more than 10 mK in temperature units. Even larger energies would be released at higher magnetic field. Since we wish to work at microkelvin or sub-microkelvin temperatures, we obviously cannot have collisions which release such large amounts of energy. This problem is avoided if we start with a $|1\rangle$ - $|2\rangle$ mixture at a temperature of $150 \mu\text{K}$, a condition we can readily produce, as discussed in

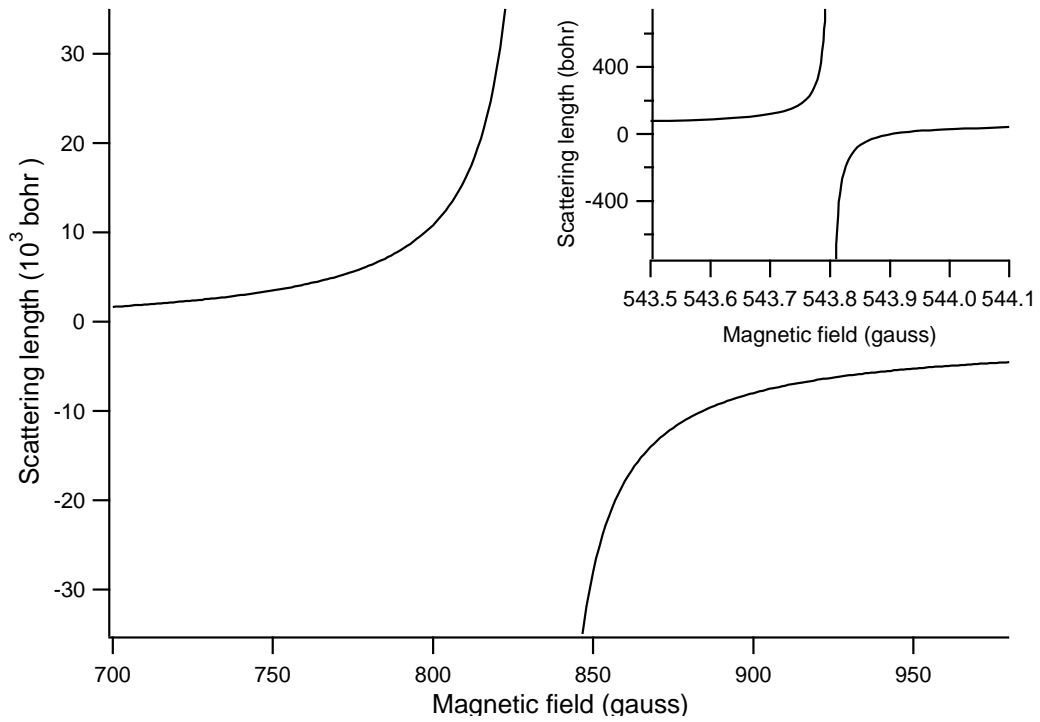


Figure 2.3: S-wave scattering length is plotted versus magnetic field for a $|1\rangle$ - $|2\rangle$ mixture of the ${}^6\text{Li}$ ground state. The main plot displays the scattering length for the broad Feshbach resonance centered at 834 G. The scattering length formula used to generate the plot is from [49]. In the inset, the scattering length versus magnetic field for the narrow Feshbach resonance near 544 gauss is displayed. The data for this plot was extracted from [50].

Section 3.1.2. While angular momentum considerations would allow two atoms beginning in the $|1\rangle$ and $|2\rangle$ states to end up in the $|1\rangle - |4\rangle$, $|2\rangle - |5\rangle$, $|4\rangle - |5\rangle$ or $|3\rangle - |6\rangle$ states, energy considerations prohibit all of these collision exit channels provided that the starting temperature of the $|1\rangle$ - $|2\rangle$ mixture is much less than 10 mK. Since the $|1\rangle$ - $|2\rangle$ mixture is stable against inelastic collisions, and possesses Feshbach resonances at experimentally accessible magnetic fields, it is an excellent choice for researchers studying two-component strongly interacting Fermi gases.

Recall that our ultimate goal is the study of two-component atomic Fermi gases. Further, consider a very simple two-component Fermi system: two spin $1/2$ particles, one with “spin-up” and the second with “spin-down.” This is precisely the type of system we create and study in our lab, where the $|1\rangle$ state has total angular momentum spin projection $m_F = 1/2$ and the $|2\rangle$ state has $m_F = -1/2$. If we consider (2.21) and (2.22), however, we see that $\cos \Theta_{\pm}$ becomes very close to unity at high magnetic field, in which case the electronic spin projection in the $|1\rangle$ and $|2\rangle$ states is nearly identical. Consequently, when we discuss our “spin-up”-“spin-down” mixture at high magnetic field, we are actually considering a mixture of atoms which is very nearly electron spin polarized (and interacts via a triplet electronic potential).

Now that we have outlined the properties of the ${}^6\text{Li}$ atom which allow us to explore the BEC-BCS crossover in a controlled laboratory setting, we are ready to consider the general experimental techniques employed in our laboratory in Chapter 3. Before doing so, however, it will be helpful to view the novel studies described in Chapters 6 and 7 as part of a worldwide effort to study strongly interacting Fermi gases. To this end, I have included an overview of the major experimental results in our field in Section 2.5.

2.5 The quest for superfluidity

In recent years, studies of degenerate Fermi gases have yielded major experimental progress and a mind-boggling number of theoretical predictions and insights. Even simple descriptions of the major theoretical developments could easily consume the remainder of this dissertation, so I will spare both myself and the reader such a summary, leaving the task instead for a theorist who is wiser and more quixotic than I am. Likewise, a description of every experimental result is beyond the scope of this document. However, a general timeline of the major experimental achievements in the field is tractable, and will be important in helping to place the studies described in this dissertation in a larger context.

A reasonable starting point for the brief history of degenerate Fermi gases is the first production of such a gas in a controlled laboratory environment. While the first degenerate Fermi gases were produced in magnetic traps [53–55], the ultimate goal of realizing a strongly interacting degenerate Fermi gas required the development of ultrastable optical traps capable of trapping stable mixtures of multiple spin states. In the mid-1990’s, the longest reported lifetimes of optical traps was on the order of ten seconds [56, 57]. After developing a theory for understanding the impact of noise sources on optical trap lifetimes [58, 59], in 1999 the Duke group produced a stable optical trap with a lifetime of 300 seconds [60], comparable to the lifetimes of magnetic traps. In 2001, the Duke group produced the first degenerate Fermi gas using all optical techniques [51].

Having produced degenerate Fermi gases, most groups then turned their sights to what would become a holy grail for the field: the creation and observation of a superfluid Fermi gas. The detection of a superfluid state requires a demonstration of macroscopic coherence of the atom sample. Macroscopic coherence can be seen

in the form of interference phenomena, superfluid hydrodynamics or quantized vortices, for example.

An important first step toward producing and detecting a superfluid state was determining the location of Feshbach resonances in the spin state mixtures of interest [52, 61, 62]. The use of Feshbach resonances was essential, as the transition temperature to a superfluid state was predicted to be experimentally accessible only in the strongly interacting regime [6, 7]. With the position of the broad Feshbach resonance in ${}^6\text{Li}$ somewhat localized, the Duke group produced the first strongly interacting degenerate Fermi gas in 2002 [16]. The signature of strong interactions was the observation of anisotropic expansion following release of the gas from the optical confining potential. The hydrodynamic expansion of the gas was suggestive of a superfluid state, but the researchers noted that collisional evolution of the cloud could not be completely ruled out as a cause of the observed expansion dynamics. In response to this study, a model based on zero temperature collisions was developed which suggested that the hydrodynamic expansion observed by the Duke group could be accounted for by collisional processes [63]. The validity of this conclusion, however, was called into question by a more rigorous model [64] which suggested that collisional processes would lead to substantial entropy and temperature increases in the gas as it expanded. Such temperature increases were not detected by the Duke group, however, and superfluid hydrodynamics remains the most likely explanation for the observed expansion dynamics of the cloud. The first observation of strongly interacting degenerate Fermi gases by the Duke group also marked the first measurement of β in a strongly interacting Fermi gas. This was followed by other measurements of the interaction energy and studies of expansion dynamics in the strongly interacting regime [65, 66]. Im-

portant experiments conducted around this time also include the observation of the absence of clock shifts in a Fermi gas in the unitarity limited regime [67,68].

In early 2003, several groups began to focus on the production of molecules on the side of the Feshbach resonance where it is energetically favorable for two fermions to form a molecule. Many groups succeeded in producing molecules from two-component Fermi gases [50,69–72], and found that the lifetime of those molecules was quite long in the vicinity of the Feshbach resonance. Having created molecules, the next logical step was to induce those molecules to Bose-condense. Reports of the production and study of a BEC of molecules formed by paired fermions came near the conclusion of 2003 and early 2004 [41,42,73–75]. While the production of a molecular BEC was a noteworthy accomplishment, the primary experimental goal remained the observation of superfluid behavior on the attractive, “BCS side” of the Feshbach resonance.

Claims of condensation in an attractive two-component mixture of fermions began to be voiced in early 2004. The JILA group made the first claim of fermionic condensation [76], based on evidence which relied on the projection of a fermionic condensate onto a molecular condensate. The interpretation of their data was questioned by the MIT group, who conducted a similar experiment and suggested that the supposed fermionic condensate may have been molecular in nature [77]. Furthermore, there was some concern regarding the validity of the JILA group’s projection technique, as it was not clear if the molecular condensate which was used as an indicator of a fermionic condensate could have formed during the projection process. These concerns were ultimately laid to rest in early 2005, when the MIT group demonstrated that the formation time of the condensate was sufficiently long that the projection technique was a valid indicator of the

existence of a fermionic condensate [78]. While these studies demonstrated the presence of a condensate, they were unable to show macroscopic coherence of the condensate, a key indicator of a superfluid state.

Several attempts to observe macroscopic coherence came in the form of studies of the breathing mode of a strongly interacting Fermi gas [19–21, 79]. Studies conducted at Duke included investigations of the magnetic field dependence of the radial breathing mode in the low temperature regime [19, 79] and the temperature dependence of the mode in a unitarity limited system [19, 20]. These results are discussed in greater detail in Chapter 7. A study by a group at the University of Innsbruck examined the magnetic field dependence of the axial and radial breathing modes in the low temperature regime [21].

Evidence for a superfluid state arose from the observation of increased breathing mode oscillation lifetimes in the low temperature regime [19, 20]. The measured oscillation frequency was in excellent agreement with predictions for a strongly interacting, hydrodynamic Fermi gas. However, frequency information alone is insufficient to establish whether hydrodynamic behavior arises from collisional processes or from a macroscopic wavefunction associated with a superfluid state. At very low temperatures, Pauli blocking suppresses collisions, which should result in decreased oscillation lifetimes at the hydrodynamic frequency. The data taken at Duke, however, showed increased oscillation lifetimes at the hydrodynamic frequency in the low temperature regime, which we regard as evidence for a superfluid state. Studies of the radial breathing mode as a function of magnetic field displayed an unexpected increase in the oscillation frequency and damping rate above the Feshbach resonance. This behavior was observed by both the Innsbruck group [21] and the Duke group [79], and was first interpreted by

the Innsbruck group as a possible indication of pair breaking.

An expected precursor to the transition to a superfluid state is the formation of pairs of fermions. Consequently, a number of techniques were devised to record the presence of fermion pairs. Radio-frequency spectroscopy [80], rapid oscillation of the interaction strength between spin components [81], detection of spatial correlations in the gas [82], and detection of shot noise in the atom cloud [83] demonstrated the presence of atom pairs throughout the BEC-BCS crossover region. However, the demonstration of pairing alone is insufficient to establish the existence of a superfluid state. Pseudogap theory (see [5] for a summary of pseudogap theory) suggests that fermions undergo pairing and a phase transition at different temperatures, unlike the case in traditional BCS theory. Instead, as the collection of fermions is cooled, pairs begin to form at one temperature, often denoted T^* , before undergoing a phase transition to a superfluid/superconducting state at some lower temperature, traditionally labelled T_c . While pseudogap theory remains somewhat controversial, several variants of pseudogap theory have achieved qualitative, and in some cases, quantitative agreement with data from the field of strongly interacting Fermi gases. As the body of available experimental results has failed to rule out pseudogap theory as a possible explanation for the behavior of strongly interacting Fermi gases, many in the field have accepted that a demonstration of superfluid behavior must include more than an observation of pairing between fermions.

It should be noted that around this time, two groups conducted refined spectroscopy studies to better localize the Feshbach resonances in ^6Li [49, 84]. While less glamorous than some other studies seeking demonstrations of superfluidity, these efforts were critical to continued progress in the field, and deserve mention.

The first study of the thermodynamic properties of the unitary Fermi gas was conducted at Duke in 2004, and was reported in [10]. Specifically, we studied the heat capacity, and with theoretical support from a group at the University of Chicago, made the first direct measurement of the superfluid phase transition temperature in a unitarity limited Fermi gas. This result is covered in greater detail in Chapter 6. The study of the heat capacity also yielded interesting insights into universality and a virial theorem for a unitary gas [85].

Another relatively recent study that deserves mention is a molecular probe experiment which determined the population of the closed-channel molecular state [86]. Here, researchers found that the population of the closed-channel singlet state is vanishingly small at and above the center of the broad Feshbach resonance in ${}^6\text{Li}$.

In May 2005, the MIT group provided further evidence for superfluidity in a strongly interacting Fermi gas: observation of quantized vortices. For the conditions of their experiment, vortices appear to be the manifestation of macroscopic coherence in their gas.

Having reached the major milestone of producing and detecting a superfluid state in the low temperature regime, one might question if research into strongly interacting Fermi gases will grow stagnant. For all that has been made of the importance of observing a superfluid state, the answer is no. There is still much to be understood about the superfluid state, and there is much to be learned about the physics of strongly interacting Fermi systems. Recent experiments have probed the momentum distributions in a Fermi gas in the crossover region [87] as well as modifications to the physics when the two spin components of the Fermi gas are not equally populated [43, 88]. Furthermore, briefly explored lines of study such

as the investigation of p-wave processes [84, 89, 90] and interspecies Feshbach resonances and heteronuclear molecules [91, 92] are candidates for additional study. Finally, while the studies mentioned above involved a single atom trap, there is a great deal of interest in modelling condensed matter systems using Fermi gases confined in optical lattices.

Chapter 3

General Experimental Methods

A typical experimental cycle in our laboratory, from the initial trapping and cooling sequence to the acquisition of an absorption image of the cooled gas of atoms, takes roughly 25 seconds. In Chapters 6 and 7, I will describe the specific experimental manipulations used to conduct studies of the heat capacity and breathing mode of strongly interacting Fermi gases. However, the bulk of the experimental cycle is the same for these experiments, as well as many others conducted in our laboratory. In fact, for a typical experimental cycle time of 25 seconds, more than 24 seconds of that time will feature identical experimental manipulations for the heat capacity and breathing mode experiments. The goals of this chapter are to outline the behavior of the experimental system during these 24 seconds, as well as present some basic theory for understanding this behavior. The experimental apparatus has been described extensively in previous theses from this group [22–24, 93], and the reader is encouraged to review those sources for greater detail regarding the experimental apparatus. The description presented in this dissertation is included for completeness and to note recently implemented upgrades which have not been described elsewhere.

We begin with a discussion of basic atom cooling and trapping theory in Section 3.1. This is followed by a description of the experimental apparatus in

Section 3.2. With all of the basic theory and experimental elements exposed, we describe a typical experimental sequence in Section 3.3. In our studies, it is important that we are able to produce both noninteracting and strongly interacting degenerate Fermi gases. Consequently, Sections 3.3.1 and 3.3.2 discuss some of the issues specific to the production of noninteracting and strongly interacting gases, respectively.

3.1 Basic cooling and trapping techniques

The technology for cooling and trapping atoms has changed little in recent years. While upgrades to our experimental apparatus have enabled us to utilize some of these techniques more efficiently, the underlying principles have remained the same. In this section, I will cover the basic cooling and trapping strategies used in our laboratory. The production of a vapor of atoms, and the subsequent slowing of these atoms in the Zeeman slower is summarized in Section 3.1.1. This initial stage of cooling provides atoms which are slow enough to be captured in a magneto-optical trap. The operation of this trap is discussed in Section 3.1.2. Some portion of the atoms confined in the magneto-optical trap are loaded directly into a second optical trap, a far off-resonance dipole trap. The operation of the dipole trap is considered in Section 3.1.3. Finally, we consider the last stage of cooling in Section 3.1.4 which relies on a method for selectively removing the hottest atoms from the atom cloud.

3.1.1 Preliminary cooling stages

If we wish to cool a gas of atoms, it follows that we must first produce a gas of atoms. As ${}^6\text{Li}$ is a solid at room temperature, we must heat the gas until a vapor forms, a topic covered in Section 3.2.3. This vapor of atoms then travels through a Zeeman slower, where a near-resonant beam opposes the motion of the atoms, a process discussed in greater detail in Section 3.2.4. Upon exiting the Zeeman slower, the atoms are travelling at speeds of roughly 30 meters per second. While this vapor of atoms is substantially cooler than the vapor that exists in the oven, we are still far from the temperatures required for the onset of quantum degeneracy. Further cooling is provided by the magneto-optical trap.

3.1.2 The Magneto-Optical Trap (MOT)

The first cooling step after the Zeeman slower relies on a Magneto-Optical Trap (MOT), a combination of laser beams and magnetic fields which reduce the temperature of the ${}^6\text{Li}$ atoms while simultaneously providing a linear spatial restoring force. We will begin with a discussion of the basic physics behind the MOT before discussing specific issues that arise when confining ${}^6\text{Li}$ in a MOT.

Basic physics of a MOT

The first topic we will consider in the operation of the MOT is the phenomenon of optical molasses. Depicted in Figure 3.1 is a one-dimensional scenario. Consider an atom (represented by the grey circles in Figure 3.1) that has a simple two-level electronic transition of frequency ω_{atom} and is constrained to move in one dimension (left and right in Figure 3.1). The atom is illuminated by coun-




Figure 3.1: Mechanism behind optical molasses. Here, we consider an atom (grey circle) confined to move in one direction (left or right). The small black arrows indicate the velocity vector of the atom. The large grey and white arrows represent a stream of photons whose frequency ω_{phot} is red detuned from the “true” (un-Doppler-shifted) transition frequency of the atom ω_{atom} . (a) Here the atom is moving to the left, which results in the photon beam from the left (grey arrow) being Doppler shifted toward resonance with the atom, while the photon beam from the right (white arrow) is Doppler-shifted even further from the atomic resonance. (b) Here the atom is moving to the right, which results in the photon beam from the right (grey arrow) being Doppler shifted toward resonance with the atomic transition. The photon beam from the left is Doppler-shifted away from the atomic transition. In both (a) and (b) the atom has an increased probability of absorbing a photon from the beam whose direction opposes the atom’s velocity vector. This results in the atom’s momentum being restricted to a small range of values close to zero momentum.

terpropagating laser beams composed of photons whose frequency ω_{phot} is red detuned from the atomic transition frequency ($\omega_{phot} < \omega_{atom}$). In scenario (a) in Figure 3.1, the motion of the atom (represented by the small black arrow) will cause the laser beam from the left to be Doppler-shifted toward resonance with the atomic transition. Meanwhile, the laser beam propagating from right to left is further red-detuned from the atomic transition. In scenario (b), the situation is reversed, with the photon beam propagating from right to left being Doppler-shifted toward resonance with the atomic transition. In both (a) and (b), the atom is more likely to absorb a photon from the laser beam which is propagating in a direction opposite to the atom's velocity vector. This preferential absorption of atoms which oppose the atomic motion will result in a reduction of the atom's momentum in that direction. Since the reduced momentum of the atom makes it seem as if the atom is moving in a viscous liquid, this phenomenon is referred to as optical molasses. While we have considered the one-dimensional case, the argument easily generalizes to a three-dimensional scenario involving three sets of red-detuned, orthogonal, counterpropagating beams.

Optical molasses does not tell the whole story of the MOT, nor does it provide sufficiently low temperatures to reach quantum degeneracy. Every cycle of absorption from one of the laser beams is accompanied by re-emission of a photon. The momentum recoil from the photon emission cycles competes with the momentum-reducing influence of the laser beams, and the net result is a lower bound on the temperatures which can be reached by relying on optical molasses. The limit is known as the Doppler cooling limit and is roughly $140 \mu\text{K}$ for ${}^6\text{Li}$. While 140 millionths of a degree above absolute zero seems like an impressively small temperature, it is still too high for the types of studies we wish to conduct.

In Section 3.1.3, we will discuss the next steps in cooling the gas of ${}^6\text{Li}$, but before doing so, we must address a second limitation of optical molasses.

In the discussion of optical molasses, it should be noted that the primary physical result was confinement of the atoms to a small region of momentum space close to zero momentum. There was no spatial confinement of the atoms, however, so by a random walk process atoms in the optical molasses could slowly wander away from the intersection of the laser beams providing the optical molasses effect. Furthermore, in Section 3.1.3, we will see that the next stage of cooling will require that the atoms confined in the MOT be loaded into a second optical trap, and the success of this procedure depends on maintaining a high atomic density in the MOT. The solution to the random walk problem and the need for high densities rests in the application of spatial confinement to the atoms.

Thus far, we have considered the action of the *optical* part of the magneto-*optical* trap, so it should not be terribly surprising that the next portion of the discussion will focus on the *magnetic* elements of the trapping mechanism. As with the discussion of optical molasses, we will consider a simple case of atomic motion in one dimension. Furthermore, while the angular momentum and associated electronic structure of the ${}^6\text{Li}$ atom is more complicated than what will be discussed in the simplified model below, the effect of the applied magnetic field will be more transparent if we consider a simple atomic structure.

Consider an atom whose ground state total angular momentum is $F = 0$ and whose excited state total angular momentum is $F' = 1$. In the ground state manifold, the total angular momentum projection m_F is obviously restricted to $m_F = 0$. In the excited state, however, we have $m'_F = 0, \pm 1$, and these sublevels should Zeeman tune differently upon application of an external mag-

netic field. Furthermore, angular momentum considerations require that the $|F = 0 m_F = 0\rangle$ to $|F' = 1 m'_F = 1\rangle$ transition occurs only upon absorption of σ^+ circularly-polarized light, while the $|F = 0 m_F = 0\rangle$ to $|F' = 1 m'_F = -1\rangle$ transition relies upon absorption of σ^- circularly-polarized light. This combination of preferential absorption of light with a given polarization along with Zeeman tuning of the m'_F sublevels in the excited state can be exploited to provide a spatial restoring force to the atoms.

Suppose we have two wire coils arranged in the anti-Helmholtz configuration. In this scenario, the coils are coaxial, and current flow through the upper coil is opposite in direction to current flow through the lower coil. The result is a spherical quadrupole magnetic field which features a null magnetic field on the axis of the coils midway between the two coils. As atoms move away from this midpoint in any direction, the magnetic field gradient varies linearly. The presence of this field gradient will result in Zeeman tuning of the m'_F sublevels in the excited state. In Figure 3.2, we depict the energy tuning of the magnetic field sublevels in the ground and excited states of our fictional $F = 0, F' = 1$ atom as a function of one spatial dimension x . Suppose that we arrange our optical beams such that photons travelling from right to left have σ^- circular polarization, while photons travelling from left to right have σ^+ circular polarization. We assume that the photons are slightly red-detuned with respect to the $F = 0 \rightarrow F' = 1$ transition at zero magnetic field (the dashed horizontal line and the thick vertical arrow indicate the energy of the incident photons in Figure 3.2). The spatially dependent force arises from the preferential selection of photons with a given polarization as a function of distance from the zero magnetic field point. An atom that wanders into the spatial region $x > 0$ will be more likely to absorb a




Figure 3.2: The spatial restoring force of the MOT arises from the preferential absorption of a photon with a particular polarization based on the atom's position. The vertical axis represents energy and the horizontal represents spatial position, with $x = 0$ at the intersection of the axes. A magnetic field exists near $x = 0$ which is well approximated by $B(x) = B_0 x$, where B_0 is a constant. We consider a hypothetical two-level system in which the total angular momentum in the ground and excited states is $F = 0$ and $F' = 1$, respectively. The narrow horizontal and slanted lines depict the Zeeman tuning of the angular momentum sublevels as a function of position. The dashed horizontal line and thick vertical line with an arrow indicate the energy of a photon which is interacting with the atom. If the atom wanders toward a position $x > 0$, it is more likely to absorb a σ^- photon, as this transition is Zeeman shifted toward resonance. Conversely, if an atom moves to a position $x < 0$, it is more likely to absorb a σ^+ photon. This action gives rise to a spatial restoring force.

σ^- photon travelling from the right, as this quantum transition is Zeeman shifted toward resonance. Likewise, atoms in the $x < 0$ region will be more likely to absorb a σ^+ photon, as this transition is Zeeman shifted into resonance. For small displacements about $x = 0$, the restoring force is linear in space. Though we have discussed this spatial restoring force in one dimension, it generalizes to three dimensions. This three-dimensional spatial confinement, when combined with the optical molasses effect, can yield a large number of cold atoms.

Throughout this discussion, I have adopted simplified models to illustrate the basic physical principles which govern the operation of MOTs. Often, the assumptions involved simplifications in the angular momentum and energy level structure of the atom which was interacting with the optical beams in the MOT. As the real angular momentum and energy level structure of the ${}^6\text{Li}$ atom is more complex than the model atoms discussed above, it is worthwhile to consider some practical issues that arise when designing a MOT to capture ${}^6\text{Li}$ atoms.

The ${}^6\text{Li}$ MOT

Ideally, the ${}^6\text{Li}$ quantum energy level structure would allow us to work on a closed, two-level transition, where the atom would shuttle back and forth between the ground and excited states. Unfortunately, the situation is not this simple. As discussed in Section 2.4, the electronic ground state of ${}^6\text{Li}$ has two possible values of the total angular momentum, $F = 1/2, 3/2$. In the lowest electronic excited state, the orbital angular momentum $L = 1$, and the nuclear and electronic spins remain the same as in the ground state, $I = 1$ and $S = 1/2$. Angular momentum addition allows for three possible values of the total angular momentum of the excited state, $F' = 1/2, 3/2, 5/2$. From an energy standpoint, these three excited

state hyperfine levels are unresolved, as the natural linewidth of the ground to excited state transition (≈ 5.9 MHz) is larger than the hyperfine splitting between the $F' = 1/2$ and $F' = 5/2$ excited states (≈ 4.4 MHz). The ground state, meanwhile, is well-resolved, with a hyperfine splitting of roughly 228 MHz between the $F = 1/2$ and $F = 3/2$ levels. Consequently, light that is resonant with the $F = 1/2$ to excited state transition will not be resonant with the $F = 3/2$ to excited state transition, and vice versa.

The ultimate consequence of this large ground state hyperfine splitting is the need for two different frequencies of light in the optical beams used in the MOT. As both ground state hyperfine levels will be occupied when atoms emerge from the Zeeman slower, we need light resonant with both levels to maximize our atom capture. Once captured, we must continue to use bichromatic MOT beams, as failure to do so would lead to optical pumping into dark quantum states which cannot be trapped by the MOT. We refer to light tuned to the $F = 3/2$ to excited state transition as the “MOT beam,” while light tuned to the $F = 1/2$ to excited state transition is referred to as the “repumper beam.” Empirically, we find that the optimized ratio of the power in the MOT beam to the repumper beam is roughly 3:1. When loading the MOT, we typically red-detune the MOT beams $\approx 30 - 35$ MHz from resonance, while detuning the repumper beams roughly 20 MHz from resonance. When working properly, our ${}^6\text{Li}$ MOT is estimated to contain roughly 500 million atoms.

Earlier, it was noted that the role of the MOT is to act as a pre-cooler for yet another optical trap, to be discussed in Section 3.1.3. We can optimize the loading of this second optical trap by conducting two additional phases at the end of the MOT loading cycle. When working properly, we can load our MOT

to a reasonable level within 5 seconds. Following this 5 second loading time, we implement the MOT “cooling phase”, where both the MOT and repumper beams are brought to within 3 MHz of resonance, and the intensity of both the MOT and repumper beams are substantially reduced (the intensity of the beams during the cooling phase must be adjusted on a daily basis to optimize the cooling phase). The cooling phase lasts roughly 5 ms and is designed to cool the atoms to the Doppler limit of $140 \mu\text{K}$ while increasing the density of the MOT. Following the cooling phase, we conduct an “optical pumping” phase, during which the repumper beams are extinguished while the MOT beams remain on. Here, the frequency of the MOT beams are tuned directly into resonance while the intensity of the beams remains close to the intensities used during the cooling phase. The optical pumping phase has a duration of $200 \mu\text{s}$ and is designed to pump all of the atoms in the MOT into the $F = 1/2$ ground state manifold. Recall from the discussion in Section 2.4 that the desired two component mixture of hyperfine states, the $|1\rangle$ - $|2\rangle$ mixture, can be generated by populating the $F = 1/2$ ground state level and applying a bias magnetic field. Hence, the optical pumping phase places the atoms in the appropriate ground state hyperfine level to produce our desired two component mixture. At the same time, removing atoms from the $F = 3/2$ hyperfine level prevents inelastic collisions which could lead to substantial heating of the atomic sample.

3.1.3 The far off-resonance dipole trap (FORT)

Whereas the operation of the MOT relies on optical beams whose frequency is nearly resonant with the ${}^6\text{Li}$ atom, the operation of the far off-resonance dipole trap (FORT) relies on a high power beam whose frequency is quite far from the

atomic resonance. The FORT is produced by a Coherent DEOS LC100-NV CO₂ laser with peak output power of 140 watts. The wavelength of the laser is 10.6 μm , which is nearly 16 times longer than the wavelength of the light used in the MOT beams. As the CO₂ laser beam is so far from resonance, the probability of absorption of a photon is quite small. Instead, the CO₂ laser beam polarizes the ⁶Li atoms, and the atoms are then attracted to regions of high laser beam intensity. A simple spatial variation in the beam intensity can be produced by focusing the CO₂ laser beam, with the position of maximum intensity being at the focus of the beam. When the FORT is operating properly, roughly two million atoms can be confined in a small region near the focus of the CO₂ laser beam. The remainder of this section will elaborate on this process.

Electric Dipole force

If we position the CO₂ laser beam and associated optics such that the focus of the laser beam coincides with the location of the ⁶Li atoms confined by the MOT, the electric field from the CO₂ laser beam will influence the ⁶Li atoms. While the atoms are overall charge neutral, spatial variations in the electric field provided by the CO₂ laser beam will induce an electric dipole moment in the atoms. The interaction energy associated with an atom of polarizability α in the presence of an electric field \mathbf{E} is given by [23, 24]

$$U_{dip} = -\frac{1}{2} \alpha \overline{\mathbf{E}^2}, \quad (3.1)$$

where the bar indicates that the electric field is averaged over many optical cycles. The electric field associated with a photon propagating through space assumes a

sinusoidal form, and if we take \mathcal{E} to be the slowly varying amplitude of the electric field, then (3.1) can be rewritten as $U_{dip} = -\alpha \mathcal{E}^2/4$. Of course, the square of the electric field can be readily related to the field's intensity I , and the potential associated with the electric dipole can be expressed in the MKS units system as

$$U_{dip} = -\frac{1}{2\epsilon_0 c} \alpha I, \quad (3.2)$$

where ϵ_0 is the permittivity of free space and c is the speed of light. As ϵ_0 , c and I are all positive quantities, we see the attractive or repulsive nature of the dipole potential depends on the sign of the polarizability α . Simple models of atomic polarizability (see [23,24], for example) indicate that if the optical frequency of the electric field is lower than the transition frequency of the atom, the polarizability α will be positive, and the dipole potential will be attractive. Strictly speaking, atoms in different quantum states which experience the same electric field will experience different forces. However, if the electric field frequency is very far detuned from any of the atomic transition frequencies considered, then atoms in different quantum states will experience essentially the same force. This is important, as we intend to confine two different hyperfine ground states of the ${}^6\text{Li}$ atom in our trap. Since the frequency of the CO_2 laser beam is so far detuned from all of the transition frequencies for the hyperfine ground states of ${}^6\text{Li}$, the potential formed by this laser beam provides a state-independent trapping mechanism. In this case, the polarizability α is effectively equal to the static polarizability α_0 associated with the DC stark effect [24].

FORT geometry

As noted earlier, the emergence of a dipole force requires spatial variations in the electric field. This is accomplished by focusing the collimated CO₂ laser beam such that the focus of this beam overlaps with the location of the atoms confined in the MOT. For a cylindrically symmetric system, the intensity of a focused Gaussian beam is [94]

$$I(r, z) = \frac{I_0}{1 + (z/z_0)^2} \exp\left(-\frac{2r^2}{r_0^2}\right), \quad (3.3)$$

where I_0 is the maximum beam intensity, $z_0 = \pi r_0^2/\lambda$ is the Rayleigh range, and r_0 is the $1/e^2$ intensity radius of the beam at $z = 0$. λ is the wavelength of the electric field associated with the beam. This expression for the spatial variation of the intensity can be combined with (3.2) to express the dipole potential energy as

$$U_{dip}(r, z) = -\frac{U_0}{1 + (z/z_0)^2} \exp\left(-\frac{2r^2}{r_0^2}\right), \quad (3.4)$$

where the maximum trap depth in MKS units is given by

$$U_0 = \frac{\alpha_0 I_0}{2 \epsilon_0 c}. \quad (3.5)$$

Note that in (3.5), we have used the static polarizability α_0 .

For many applications, we find that the atoms confined in the dipole potential occupy only the deepest portion of the trap. In this case, the combined Lorentzian confinement in the axial direction and the Gaussian confinement in the radial dimension can be approximated by harmonic confinement in all dimensions. Taylor

expanding (3.4) yields

$$U_{dip}(r, z) \approx -U_0 + \frac{2U_0}{r_0^2} r^2 + \frac{U_0}{z_0^2} z^2. \quad (3.6)$$

If we consider a harmonic trap containing particles of mass m where the trap oscillation frequencies ω_i are defined in the usual manner

$$U(r, z) = -U_0 + \frac{m\omega_r^2}{2} r^2 + \frac{m\omega_z^2}{2} z^2, \quad (3.7)$$

then we can make the identifications

$$\omega_r^2 = \frac{4U_0}{m r_0^2} \quad (3.8)$$

$$\omega_z^2 = \frac{2U_0}{m z_0^2}. \quad (3.9)$$

Loading the FORT

In the discussion of the operation of the MOT, I alluded to two phases, known as the “cooling” and “optical pumping” phases, which are conducted at the end of each MOT loading cycle. These two phases are necessary to optimize the loading of the FORT from the MOT. In fact, while the MOT is loading, we are loading the FORT as well. Provided that the FORT and MOT are well-overlapped, and provided that the cooling and optical pumping phases are working properly, we can load approximately 2 million atoms into the FORT. With the FORT loaded, the MOT beams have served their purpose and can be extinguished until the beginning of the next experimental cycle. The atoms loaded into the FORT are at approximately 140 μK , the Doppler limited temperature of the MOT. This is still

far from the temperatures we seek, however. To achieve the desired temperatures, we must conduct evaporative cooling, a process we now discuss.

3.1.4 Evaporative cooling

With the MOT beams extinguished and the FORT loaded, we still have a non-interacting classical gas. Classical because the temperature is still too high, and noninteracting because the magnetic field at the location of the FORT is zero, where the scattering length between the $|1\rangle$ and $|2\rangle$ states is zero. (Recall that the MOT magnetic field is a spherical quadrupole field with a magnetic field zero equidistant between the two anti-Helmholtz coils).

The process of evaporative cooling is relatively straightforward. Suppose we have a trap containing a 50-50 mixture of two spin states and we produce conditions such that the particles begin to collide with each other. If the trap depth is many orders of magnitude deeper than the energy of an average atom inside the trap, then the probability of atoms leaving the trap is quite small. However, suppose that the trap depth is only a few times greater than the average energy of a particle in the trap. As the particles confined in the trap undergo collisions, some atoms will cede energy to others and fall into deeper portions of the trap while others will manage to acquire enough energy to escape the trapping potential. If only the highest energy atoms leave the trap, then the remaining atoms will be left with a lower average energy. That is, the remaining atoms will be cooler. This process occurs continuously: the hottest atoms leave the trap, while the remaining atoms rethermalize via collisions. This process is known as evaporative cooling, a procedure in which the experimenter sacrifices hot atoms to produce a cooler atomic sample. In an ideal world, evaporative cooling would continue

ad infinitum or until the experimenter intervened. In practice, the process loses its effectiveness when the average energy of a trapped particle is substantially lower than the trap depth, at which point it becomes statistically unfavorable for a given atom to acquire enough energy from collisions with other particles to escape the trap.

All is not lost, however, if the trap depth is lowered to the point where evaporative cooling can begin again. We use the terminology “free evaporation” to describe the process of evaporative cooling while the trap depth is held constant at its maximum value. We use the term “forced evaporation” to describe the process of evaporative cooling while lowering the depth of the optical trap. Forced evaporation proceeds by time-dependent amplitude reduction of the confining potential. A simple exponential lowering curve can be applied to conduct forced evaporation, but a more efficient method can be derived based on simple scaling laws [23, 24, 95]. If we wish to keep the trap depth an order of magnitude greater than the thermal energy $k_B T$ of an average particle in the trap, then the proper time dependence of the trap depth is given by

$$U(t) = U_0 \left(\frac{1}{1 + t/\tau_{lc}} \right)^{1.45}, \quad (3.10)$$

where τ_{lc} is the lowering curve time constant. Depending on the duration of time over which the trap is lowered and the selected value of τ_{lc} , we can lower the trap to any desired depth. In practice, the trap depth is reduced not by reducing the power of the CO₂ laser, but by reducing the efficiency of an AO placed at the laser’s output (see Section 3.2.7). We use an IntraAction Corp AGM-4010BJ1 AO designed to operate at maximum efficiency when 50 watts of 40 MHz RF power is

applied to the AO's crystal. Before reaching the RF amplifier for the AO crystal, a 40 MHz sine wave of fixed amplitude generated by an Agilent E4423B signal generator is multiplied by the output of an Agilent 33250A arbitrary waveform generator. Prior to being multiplied by the signal from the Agilent E4423B signal generator, the output of the Agilent 33250A waveform generator is directed through a low-pass filter to minimize electronic noise which contributes to atom heating in the optical trap. The arbitrary waveform generator is loaded with a voltage versus time curve which will produce the desired trap depth versus time given by (3.10). When the user wishes to begin forced evaporation, the arbitrary waveform generator is triggered, and the CO₂ laser power reaching the trapping region is reduced in a controlled fashion. After the trap depth has been lowered to the desired value, the user can leave the trap depth at its minimum point for a period of time before adiabatically recompressing the trap to any fraction of its original depth. Details of this procedure are addressed in Section C.3 of Appendix C.

We note that the appearance of τ_{lc} in (3.10) raises another important issue: the trap depth cannot be lowered at an arbitrarily fast rate. The time constant involved in the lowering process must be substantially larger than the average time between collisions of the trapped particles. If this condition is not met, the trapped atoms will not have sufficient time to rethermalize after ejecting hot atoms, and the efficiency of evaporative cooling will suffer. Empirically determined values of the lowering curve time constant for common experimental conditions are given in Sections 3.3.1 and 3.3.2.

To this point, it should be clear that the efficiency of the evaporative cooling process depends critically on the presence of collisions between the two trapped

spin states. The efficiency of cooling is impacted also by the degree to which the two spin states are equally populated. Consider an extreme case, where 90% of the atoms are in state $|1\rangle$ and the remaining 10% are in state $|2\rangle$. Here, many atoms in state $|1\rangle$ will be unable to undergo collisions, and will simply be spilled from the trap as the trap depth is reduced. The optimum scenario involves a 50-50 mixture of both spin states. To ensure that this is the case before evaporative cooling commences, we first ramp the magnetic field from 0 gauss to a small bias field of roughly 8 gauss. A noisy radio-frequency (RF) pulse is then applied to the atoms confined in the FORT via a small antenna placed in our vacuum chamber. The center frequency of this noisy RF pulse matches the frequency splitting between the $|1\rangle$ and $|2\rangle$ states at 8 gauss, and drives the atoms into an incoherent equally-populated mixture. Once the spin state populations are properly balanced, the bias magnetic field can be ramped to higher values to initiate the processes of free and forced evaporation.

In closing, I offer a brief note about terminology. Throughout this chapter, I have described a number of cooling processes, including cooling provided by the Zeeman slower, the MOT and evaporative cooling. As we now take for granted the proper operation of the cooling mechanisms of the Zeeman slower and the MOT, the only remaining cooling process of substantial interest is evaporative cooling. For this reason, I will frequently use the term “cooling” to mean “evaporative cooling” throughout the remainder of this dissertation.

3.2 Experimental Apparatus

When compared to some other atom cooling and trapping laboratories, our experimental apparatus appears to be relatively simple. That said, the length of this chapter should indicate that our system still features a reasonable level of complexity. In this section, I will attempt to provide a basic description of the major components in our experimental apparatus. Much of the system has been described in detail in previous theses from this group [22–24, 93]. Consequently, for more detailed information than what is provided in the following descriptions, the reader is encouraged to peruse earlier theses.

3.2.1 Double-pass AO arrangement

Acousto-optic modulators (AOs) are used extensively throughout the experiment to generate optical beams of the desired frequency. All of the AOs presently in use on the experimental system operate in the double-pass configuration. A schematic is shown in Figure 3.3. A polarizing beam-splitting cube directs a linearly polarized optical beam into the AO crystal. Lenses placed before and after the crystal focus and recollimate the beam. An iris placed behind the crystal blocks all orders but the first order upshifted beam, which then travels through a quarter waveplate before contacting a retroreflecting mirror. The mirror sends the beam back through the quarter waveplate, iris and AO crystal. The first order beam from this backgoing pass will now be propagating along the same line as the incident beam, and will have experienced two equal frequency shifts from the AO. This backgoing beam also has had its polarization flipped, such that it will exit the polarizing beam splitting cube on a face orthogonal to the face where the

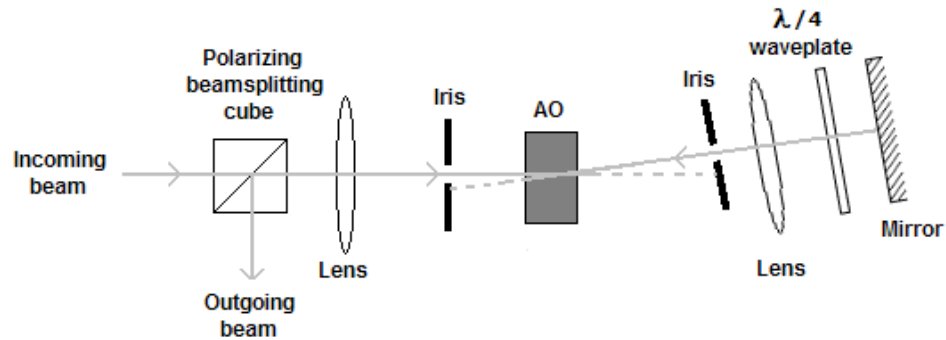


Figure 3.3: Schematic of an acousto-optic modulator operating in the double-pass configuration, as described in Section 3.2.1.

beam first entered. This arrangement has a number of advantages over the single-pass arrangement. First, when the AO is in the “off” position, rejection of leakage light is better than for an AO operating in the single-pass configuration. Further, as the frequency shift provided by the AO is changed, the angular misalignment in the outgoing beam is reduced for the double-pass arrangement compared to the single-pass scenario. Finally, the double-pass configuration allows for larger frequency shifts than those possible with a single pass through the AO. Unless otherwise noted, throughout the remainder of this dissertation, reported values of frequency shifts provided by AOs refer to the *total* frequency shift after the optical beam has passed through the AO twice.

3.2.2 The vacuum chamber

The cooling and trapping techniques discussed in Section 3.1 are effective provided that the atom cloud is in a low pressure environment. High background gas pressure will result in collisions between background gas particles and trapped atoms, which will lead to heating of the trapped atom cloud as well as loss of atoms.

Consequently, the business of cooling and trapping atoms must be conducted in an ultrahigh vacuum environment.

The primary vacuum chamber comprises the oven, the Zeeman slower, and the so-called “trapping region” where the MOT and FORT are located. Detailed diagrams as well as descriptions of the entire vacuum chamber are available in [22–24, 93]. A side view of the vacuum system is shown in Figure 3.4. At the right of the diagram is the oven region. Atomic flux from the oven region travels through the Zeeman slower before being trapped by the MOT and FORT in the main trapping region. In the main trapping region, the pressure is less than 2×10^{-11} torr, which is roughly two orders of magnitude lower than the pressure maintained in the oven region when the oven is at operating temperatures. The pressure differential is made possible by a small aperture copper nipple between the oven region and the Zeeman slower. The ultrahigh vacuum is maintained by continuously operating ion pumps along with occasional operation of titanium sublimation pumps. The operation and specifications of these pumps are discussed in greater detail in [22–24].

3.2.3 The oven

Lithium is a soft, silver-gray solid at room temperature, and while our ultimate goal is to produce very cold ${}^6\text{Li}$ atoms, doing so requires that we manipulate a vapor of ${}^6\text{Li}$. The melting point of lithium is around 180°C [96], but we need to keep our lithium oven source closer to 400°C to provide adequate atom flux into the main trapping region. For more information regarding the vapor pressure of ${}^6\text{Li}$ as a function of temperature (which determines the atom flux to the main trapping region), the reader is encouraged to consult Appendix A in [24].




Figure 3.4: Simplified diagram of the vacuum chamber. The oven region is shown at the right. Atomic flux from the oven passes through the Zeeman slower before being confined in the main trapping region by the MOT and FORT. Alphanumeric labels on the various vacuum ports match the convention used in [24]. Figure contents are modified versions of similar diagrams appearing in [24]. Drawing not to scale.




Figure 3.5: Schematic of the ${}^6\text{Li}$ oven. At the left is the “can,” where most of the ${}^6\text{Li}$ resides. The nozzle of the oven sprouts from the can and collimates the atomic vapor while conducting it to the Zeeman slower. The oven assembly mates with the vacuum chamber using a rotatable 2 $\frac{3}{4}$ inch ConFlat (CF) flange, at the right of the schematic. Dimensions are given in inches. The boldface numbers on the diagram mark the 5 regions of the oven. Wrapped around the outside of each oven region is a nichrome wire, which is used as the heating element for the oven. As the nichrome wires in different oven regions are wired in parallel, we have independent control of the temperature in each region of the oven. Thermocouples monitor the temperature in each of the regions.

The oven itself resembles a corn-cob pipe (see Figure 3.5), and is mummified by 5 independent sets of nichrome wire windings (Omega PN NI80-020-50) and thermal cement (Omega CC High Temperature Cement). As the nichrome wires in the 5 oven regions are wired in parallel, the current flowing through each wire can be customized to produce the desired temperature for that region of the oven. Thermocouples are placed in each of the 5 regions to monitor the temperature profile. Both the thermocouples and the nichrome wire are encased in thermal cement, and the oven is enclosed by firebrick to insulate the assembly from its room temperature surroundings.

The interior of the oven is lined with stainless steel mesh, which is designed to wick liquid lithium which strikes the inner bore of the nozzle back toward the can region. This design is often referred to as a recirculating oven, and the basic goal is to maximize the life of the oven by recycling any lithium which fails to successfully exit the oven on its initial attempts. The success of the wicking action, the atomic flux from the oven, and the lifetime of the oven all depend on the chosen temperature profile. Choosing an appropriate temperature profile remains something of a black art, and the “correct” temperature profile varies from oven to oven. As a general guideline, Table 3.1 lists temperature profiles that worked with two different ovens.

3.2.4 The Zeeman slower

The first stage of the atom cooling process occurs in the Zeeman slower (see Figure 3.4). Here, the motion of the atoms is impeded by a slowing beam which enters the vacuum chamber system from the left via port S1 in Figure 3.4. The frequency of the slowing beam is red detuned from the D_2 transition in ${}^6\text{Li}$ by

Region No.	1	2	3	4	5
Oven A	370	380	425	350	250
Oven B	370	380	415	410	350

Table 3.1: Temperature profiles for two different ovens. The region numbers correspond to the numbers shown in Figure 3.5. Temperatures are given in °C. Determining the correct temperature profile for a given oven is a trial-and-error process, but the above examples provide a reasonable starting point when breaking in a new oven.

roughly 200 MHz. However, as the atomic beam from the oven and the slowing beam are counterpropagating, the finite velocity of the atoms causes a Doppler shift which brings some portion of the atoms into resonance with the slowing beam. Atoms that are resonant with the slowing beam will absorb photons. The excited atoms will then emit photons in random directions. Since the atoms are always absorbing photons which are travelling in a specific direction (toward the oven), but emitting photons in a random direction, on average the atoms' speed along the slowing beam axis will decrease. However, this slowing will occur only for a small portion of the atomic flux for which the atoms are Doppler-shifted into resonance with the slowing beam. Furthermore, as this small subset of atoms is slowed, the atoms will eventually reach sufficiently small velocities that they will no longer be Doppler-shifted into resonance with the slowing beam. A more efficient method of harvesting fast atoms from the oven is needed.

The solution to this problem is the Zeeman slower. Here, several wire coils coaxial with the slowing beam direction apply a spatially varying bias magnetic field to the atoms. In Figure 3.4, the 10 wire coils are separated by larger diameter air fins to dissipate heat. The coils are wired in parallel, so the current flowing through each coil can be controlled independently. The bias magnetic

fields Zeeman shift the energy level spacing of the atoms such that the atoms remain resonant with the Doppler-shifted slowing beam as they travel the length of the Zeeman slower. Some fraction of the atoms from the oven (those with very high velocities) remain untrappable, as the maximum bias field applied by the slower is incapable of Zeeman shifting the energy level spacing of the atom to be resonant with the slowing beam. Nevertheless, with prudent choices of the magnetic fields, a reasonably large number of atoms can be harvested from the “fast” atomic beam coming from the oven. The amount of atom slowing (i.e., cooling) which occurs in the Zeeman slower is quite impressive. Atoms travelling around 2 kilometers per second upon entering the Zeeman slower can be slowed to roughly 30 meters per second upon exiting the slower. Most importantly, this latter speed is sufficiently slow for the atoms to be captured by the MOT discussed in Section 3.1.2. For more information regarding the Zeeman slower, the reader is encouraged to examine [93, 97].

3.2.5 Optical beam generation

The discussion of the operation of the MOT and the Zeeman slower makes obvious the need to generate several optical beams near the D_2 resonance in ${}^6\text{Li}$. This section addresses the production of such beams.

The beam generation workhorse is a Coherent 699-21 dye laser, which is pumped by a recently acquired Coherent Verdi V-10 diode-pumped solid state laser. The latter of these can output up to 10 W at 532 nm, and has excellent pointing stability. A typical output pump power for the Verdi is around 6.5 W, though only a little more than 5 W reaches the dye laser after a beam splitter inserted between the Verdi and the dye laser redirects a portion of the pump beam.

We utilize 1.17 grams of LD688 dye dissolved in 1.1 liters of 2-phenoxyethanol in our dye laser, and we find that the dye laser can output over 850 mW in single mode operation near 671 nm when the dye is fresh.

An important consequence of the apparatus used to establish the frequency standard of the dye laser (see Section 3.2.6) is that the dye laser output frequency is roughly 200 MHz below the desired atomic resonance in ${}^6\text{Li}$. This is by design, as the operation of the Zeeman slower requires that the slowing beam be roughly 200 MHz below resonance. A quarter waveplate¹ and polarizing beam splitting cube after the dye laser output directs 80-130 mW of laser power to the slowing beam port (see Figure 3.9). A glan prism and quarter waveplate are placed in the slowing beam path prior to the vacuum chamber to ensure that the slowing beam has good circular polarization to maximize its interaction with the atomic flux from the oven. In addition, a telescope placed in the slowing beam path expands the beam diameter to roughly 1 cm before entering the vacuum chamber. The slowing beam then converges to a focus in the oven region. The high intensity of the slowing beam near the oven creates a dipole force which attracts atoms to the slowing beam and improves the rate of atom loading.

Laser power which is not diverted to the slowing beam path is directed to the MOT AO, where a double-pass arrangement upshifts the laser light by 165-170 MHz during the MOT loading phase (or more, for the cooling and optical pumping phases). A half waveplate and polarizing beam splitting cube direct a portion of the MOT light to the repumper AO, where the light is upshifted even further to produce the repumper light needed to excite atoms from the

¹Ideally, we would use a half waveplate to determine the portion of the dye laser output that forms the slowing beam. However, given the desired power distribution after the waveplate and polarizing beam splitting cube, a quarter waveplate works fine.

$F = 1/2$ ground state. The MOT and repumper beams are then recombined on a polarizing beam splitting cube and directed toward the vacuum system, as shown in Figure 3.9. Before reaching the vacuum system, however, the combined MOT and repumper beams are split into 3 separate beams, each with its own expanding telescope. These 3 beams are directed to different ports on the main vacuum chamber, where they make up the 3 mutually orthogonal MOT beams needed to confine the atoms in three-dimensional space. Note that this scheme differs from the “bowtie” MOT beam generation scheme described in previous theses [22–24]. The present scheme, which uses three independent optical beams, allows for easier beam alignment and optimization of trap loading.

The MOT beams enter the vacuum chamber as shown in Figure 3.6 and panel (a) of Figure 3.7. In Figure 3.6, the horizontal MOT beams enter the chamber after passing through a quarter waveplate to produce circularly polarized light. Upon exiting the vacuum chamber, each beam passes through an additional quarter waveplate before being retroreflected.

In Chapter 5, we will discuss the procedure for extracting data from the system by taking absorption images of the atom clouds. These absorption images require the production of a resonant probe beam and the introduction of this probe beam to the vacuum chamber. In previous theses, the camera probe ports and the vertical MOT beam ports were independent. However, to observe anisotropic expansion of the atom cloud in the strongly interacting regime [16], the imaging system was modified, and the vertical MOT beam and probe beam now share the ports formerly used exclusively by the vertical MOT beam. In panel (a) of Figure 3.7, we show the operation of the vertical MOT beam. Here, the beam passes through a polarizing beam splitting cube before encountering a quarter

waveplate. Upon exiting the bottom of the vacuum chamber, the beam passes through a second quarter waveplate and a second polarizing beam splitting cube before being retroreflected. The polarization of the vertical MOT beam is chosen such that the beam passes through the polarizing beam splitting cubes unchanged. The cubes are necessary, however, to introduce the probe beam into the vacuum system, a scenario depicted in panel (b) of Figure 3.7. Here, the linear polarization of the probe beam is selected such that it will enter the polarizing beam splitting cube at the bottom of the vacuum chamber and be directed upwards. Before entering the vacuum system, the probe beam passes through a quarter waveplate, which produces the desired circular polarization for interactions with the trapped atoms. At the top of the vacuum system, the beam passes through the quarter waveplate, producing linear polarization which is deflected to the left when it contacts the upper polarizing beam splitting cube. This deflected beam passes through a simple imaging setup (not shown in Figure 3.7) before being imaged by the camera.

Having settled the geometry of the MOT and probe beams, we turn now to the production of the probe beam. We do not require that the MOT and repumper beams be available during all phases of an experimental cycle. Consequently, when the MOT and repumper beams are not needed, the MOT AO is turned off, and the zeroth-order beam which passes through the MOT AO strikes a pick-off mirror and is directed through consecutive double-passed AOs. This scenario is depicted in Figure 3.10. The camera AOs are used to generate the probe beam used in imaging, and hence are given the rather unexciting but accurate names of Camera AO #1 and Camera AO #2. After navigating through both camera AOs, the camera beam is directed to a single-mode optical fiber which guides the beam




Figure 3.6: Top view of the horizontal MOT beams entering the vacuum chamber. Prior to entering the chamber, each linearly polarized beam passes through a quarter waveplate (gray rectangle) to produce circularly-polarized light. Upon exiting the chamber, the beam passes through a quarter waveplate before being retroreflected by a mirror (black rectangle). The vertical MOT beam (not shown) enters the vacuum chamber via port M1. This figure is not to scale, and is a modified version of a similar figure appearing in [24]. Alphanumeric labels on the various vacuum ports match the convention used in [24].

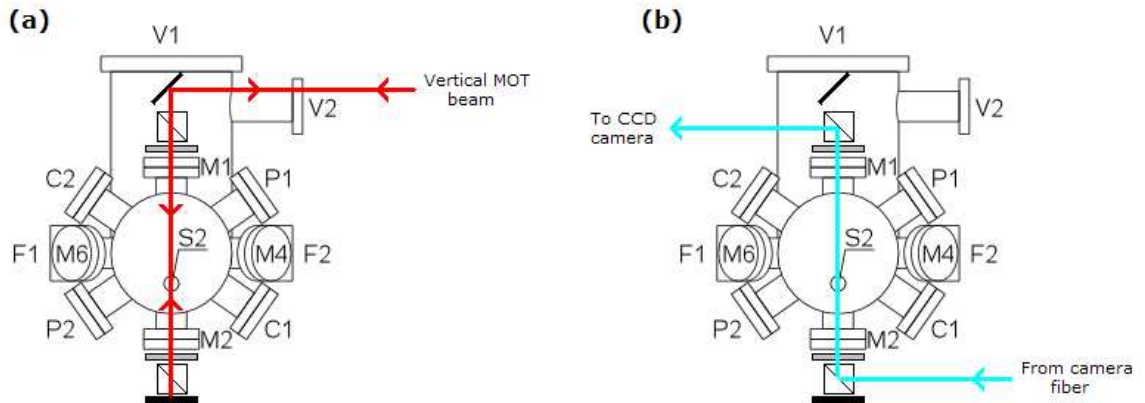


Figure 3.7: View from the oven side of the vacuum chamber of the geometry of the vertical MOT and probe beams. Gray rectangles represent quarter waveplates, solid rectangles are mirrors, and squares with a single diagonal line represent polarizing beam splitting cubes. (a). The vertical MOT beam configuration. (b). The probe beam configuration. These figures are not to scale, and are modified versions of similar figures appearing in [24]. Alphanumeric labels on the various vacuum ports match the convention used in [24].

closer to the vacuum system and removes many beam imperfections. Upon exiting the camera fiber, the beam is collimated after expanding to a size of roughly 2.5 cm in diameter, and directed to the vacuum system as displayed in panel (b) of Figure 3.7. An important limitation of this experimental arrangement is that we cannot simultaneously generate the MOT and probe beams. In practice, however, this limitation rarely presents problems.

In closing, we offer a few comments about the simplified diagrams of the current optical layout included in Figures 3.9 and 3.10. The legend for these figures is shown in Figure 3.8. Please note that the diagrams are not to scale, and that there are a number of elements on the optical table that are not represented in the diagrams. For instance, various telescopes, irises, and mirrors have been omitted to avoid cluttering the diagrams. Only the essential elements for understanding




Figure 3.8: The legend for the optical layouts shown in Figures 3.9 and 3.10. Waveplates in the diagrams are labelled according to their phase shift ($\lambda/2$ or $\lambda/4$). The mechanical shutter shown above is in the closed position. Mechanical shutters which are open are shown with a white area at the center of the narrow section.

the generation of optical beams with different frequency shifts have been included.

3.2.6 Frequency standard – the “locking” region

As some processes during an experimental cycle require that we produce optical beams within a MHz of the desired atomic transition, we require that our dye laser be locked to an absolute frequency standard. This frequency standard is supplied by an independent vacuum chamber and ^6Li oven arrangement, which we refer to as “the locking region.” A beam splitter takes a small portion of the dye laser’s output and directs it through a double-passed AO, where the beam is upshifted by roughly 200 MHz. This probe beam (with a power of approximately 1 mW) is then directed to the locking region vacuum system, where it intersects perpendicularly with an atomic beam of ^6Li . The fluorescence from the beam is monitored by a PMT while a small amount of frequency modulation introduced by the AO causes the frequency of the probe beam to scan over the D_2 transition at a frequency of 14.2 kHz. A lock-in amplifier and electronic servo [98] process the frequency-modulated fluorescence signal from the PMT, and send an error signal

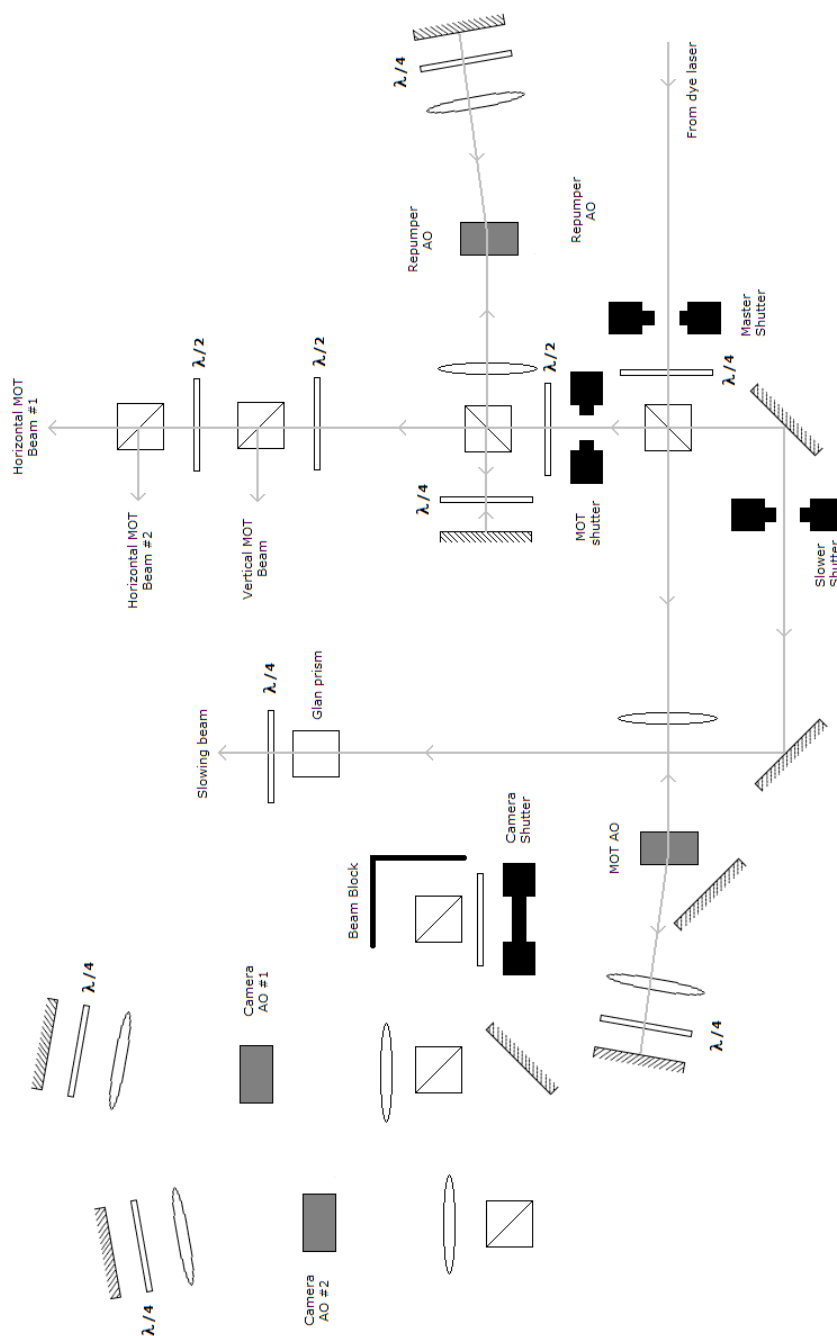


Figure 3.9: Optical layout for generating the MOT and slowing beams. The legend is shown in Figure 3.8. The schematic is not to scale and some optical elements have been omitted.




Figure 3.10: Optical layout for generating the camera beam. The MOT AO is in the “off” position, allowing a pickoff mirror to direct the zeroth order beam to the camera AOs. The legend is shown in Figure 3.8. The schematic is not to scale and some optical elements have been omitted.

to the dye laser control box to keep the dye laser's reference cavity locked to the desired atomic resonance. More information about the locking region design can be found in [99].

3.2.7 FORT beam generation and conditioning

In some sense, generation and manipulation of the laser beam which produces the FORT is simpler than the generation of beams associated with the MOT. In the case of the FORT, we do not care what the precise frequency of the trapping beam is, only that it is sufficiently far from resonance. Furthermore, we need only one beam to produce the FORT, compared to the numerous beams required for a MOT. However, the generation of the FORT does pose several problems.

The laser source for the FORT is a Coherent DEOS LC100-NV CO₂ laser, which has a peak output power of 140 watts at a wavelength of 10.6 μm . This wavelength is well into the infrared portion of the electromagnetic spectrum, where it is invisible. This combination of high power and invisibility can make beam alignment procedures complicated and/or dangerous. Fortunately, for sufficiently low powers or for beams with large cross sectional areas, the position of the beam can be tracked using fluorescent viewing cards. Beam alignment procedures are also aided by the use of an AO placed near the CO₂ laser's output, as depicted in Figure 3.11. The AO's efficiency can be reduced to a level such that the first order transmitted beam which produces the FORT is cut to a safe power level.

The location of the AO must be carefully chosen, as the high output power of the CO₂ laser taxes the ability of the water-cooled² AO to dissipate energy

²In fact, the AO is cooled by a mixture of distilled water (75%) and Dowfrost (25%), an




Figure 3.11: CO₂ laser and placement of beam conditioning optics relevant to generation of the FORT trapping beam. Drawing not to scale. Drawing of the vacuum chamber taken from [24]. Alphanumeric labels on the various vacuum ports match the convention used in [24].

absorbed from the beam. For this reason, the distance between the CO₂ laser output and the AO is increased by adding a “zig-zag” mirror arrangement between the elements to increase the path length (not depicted in Figure 3.11). This increased path length allows the cross sectional area of the CO₂ laser beam to expand, which reduces the intensity of the beam when it enters the AO, and improves the far-field quality of the beam. When first installed, the AO is rotated to produce maximum output power in the first order beam, represented by the solid gray line in Figure 3.11. Some portion of the beam incident on the AO will remain in the zeroth order, however, and this beam (represented by the thick dashed gray line) is directed to a water-cooled beam dump. In spite of the efforts to reduce thermal effects in the AO, we still observe thermal lensing of the first-order beam. This thermal lensing leads to astigmatism in the beam, which is corrected by a cylindrical telescope formed by lenses L1 and L2. The beam then passes through a diffractive sampling optic, where a small portion of the beam is extracted and directed to a detector (the thin gray line in Figure 3.11). This detector arrangement can be used to measure noise characteristics of the laser as well as make precision measurements of the laser power which are necessary for effective evaporative cooling. The vast majority of the beam, however, continues to a polarizer, which is positioned to allow this forward-going beam to pass unimpeded. The beam then passes through an expanding telescope formed by aspheric lenses L3 (focal length = 1.2 inches) and L4 (focal length = 11.25 inches). The telescope expands the beam by a factor of 10, approximately. After being redirected by two mirrors, the beam is incident on aspheric lens L5 (focal length = 7.5 inches), which we refer to as the “focusing lens” or the “final lens.” Recall

anti-corrosion agent. This same mixture cools the CO₂ laser as well.

from the discussion of FORT operation that the depth of the trap is proportional to the intensity of the beam. Lens L5 produces a point of very high intensity in the middle of the vacuum chamber, where the FORT is formed. At the focus, the beam radius is between 40 and 50 μm and the Rayleigh length in the axial direction is roughly an order of magnitude larger. Note that the 10x telescope is essential for producing a tight focus, as basic gaussian optics predicts a tighter focus for a beam with a larger waist entering the focusing lens. The aspheric lens L6 (focal length = 7.5 inches) recollimates the beam on the other side of the vacuum chamber. In between, the beam passes through two special windows, to be discussed in a moment.

Once the beam has been recollimated by lens L6, there are two options for its subsequent path. The “chopper” is an element after lens L6 that is essentially a mirror which can be raised or lowered pneumatically [23, 24]. In the scenario depicted in Figure 3.11, the “chopper” is in the down position, and the beam is directed to a power meter. The chopper remains in this position throughout much of the free and forced evaporation processes. However, when loading the FORT, we can increase the efficiency of atom capture from the MOT by increasing the power of the beam which forms the FORT. This is accomplished by placing the chopper in the “up” position, in which case the beam strikes the rooftop mirror. The rooftop mirror retroreflects the beam while flipping its linear polarization 90 degrees. The retroreflected beam retraces the path of the input beam all the way to the polarizing optic, at which point its orthogonal polarization causes the beam to be directed into a beam dump, as demonstrated by the thin dashed gray line in Figure 3.11.

One might question: why not leave the retroreflected beam on at all times?

While this would provide us with more trapping power, we find that the pointing instability of the backgoing beam heats the trapped atoms, a scenario that is largely absent in the forward-going beam. In fact, low pointing noise and other stability factors are critical in being able to perform atom cooling and trapping experiments on a long time scale using all-optical techniques [58]. For this reason, the laser providing the trapping beam must be extremely quiet, and all optics along the CO₂ laser beam path must be well mounted on sturdy bases.

In closing, we noted earlier that the CO₂ laser beam passes through special windows when entering and exiting the vacuum chamber. In fact, all of the optics along the CO₂ laser path, including the windows, are made of zinc-selenide, a yellowish, but largely transparent material which is non-absorptive in the infrared portion of the electromagnetic spectrum. Glass, unfortunately, absorbs at 10.6 μm , and is therefore unsuitable for use with the CO₂ laser. In spite of the efforts to reduce absorption of the CO₂ laser beam, the $\approx 70\%$ maximum efficiency of the AO and the number of surfaces between the laser output and the power meter noticeably reduces the beam power. While peak powers near 140 W are possible at the CO₂ laser output, 60 W is a normal power level detected at the position of the power meter.

3.2.8 Radio-frequency antenna

In 2002, a vacuum system upgrade was conducted in which a radio-frequency (RF) antenna was installed in the vacuum chamber. Transitions between the hyperfine ground states of the ⁶Li atom occur in the radio-frequency range, so the RF antenna allows us to alter the populations of the various hyperfine ground states.

In spite of the wide array of potential experimental manipulations provided by the antenna, to date we have used the antenna exclusively to balance spin state populations prior to initiating evaporative cooling. This is accomplished by applying an 8 gauss bias magnetic field to the atoms and sending a noisy RF pulse centered at ≈ 7.4 MHz to the antenna. The bandwidth of the pulse is 2 MHz, and the pulse duration is 100 milliseconds. The center frequency of the pulse is chosen to match the frequency splitting of the $|1\rangle$ and $|2\rangle$ states at 8 gauss. Via rate equation pumping, the populations of the $|1\rangle$ and $|2\rangle$ spin states are equalized.

The source of the RF signal is a Hewlett Packard 33120A function generator whose output is directed to a Mini-Circuits 15542 ZAD-1 frequency mixer. Here, the signal is multiplied by a DC signal provided by a buffered output from the multiplexer, which allows us to easily control the RF circuit from the computer. Prior to reaching the RF antenna, the RF signal is amplified by an Amplifier Research Model 30L amplifier.

3.2.9 High field magnets

Designed and built largely by Ken O'Hara, the high field magnets presently installed on the apparatus enable studies of the behavior of strongly interacting gases near the broad Feshbach resonance in ${}^6\text{Li}$ centered at 834 gauss. The two magnets are each powered by two Agilent 6682A power supplies wired in series. At peak output, the current flow is 240 amps, producing bias fields of up to 1114 gauss. Each magnet dissipates close to 5 kW, and therefore must be water-cooled. Interlocks monitoring the power consumption of the magnets as well as the flow of cooling water protect against overheating and destruction of the magnets.

The high field magnets are involved in almost every phase of the experimental

cycle. During the MOT loading phase, the magnets run in the anti-Helmholtz configuration, in which the current flow through the upper magnet is opposite to the current flow in the lower magnet, producing a spherical quadrupole magnetic field. Following conclusion of the MOT loading sequence, the so-called “magnet flipper” box, also designed by Ken O’Hara, reverses the polarity of one of the magnets. When the magnets are turned on again, they operate in the Helmholtz configuration, producing a nearly uniform bias magnetic field at the location of the trapped atoms.

3.2.10 MOT bias coils

When optimizing the loading of atoms from the MOT to the FORT, it is helpful to be able to shift the MOT position slightly. This is accomplished using two pairs of electromagnets. The small coils coaxial with the direction of the CO₂ laser beam [23,24] allow for shifting of the MOT along the CO₂ laser beam axis. Larger coils coaxial with the slowing beam axis shift the MOT in that direction. The current flow through both sets of bias coils is controlled by the “Hummer,” a custom-built current regulation device.

Finally, the large bias coils coaxial with the slower do more than provide a bias field during the MOT loading phase. They also provide the 8 gauss magnetic field needed for equalizing spin state populations as discussed in Section 3.2.8.

3.2.11 CCD camera and imaging system

The generation of the probe beam used for imaging was addressed in Section 3.2.5, and the path of the probe beam through the vacuum chamber is depicted in panel (b) of Figure 3.7. After exiting the polarizing beam splitting cube on top of the




Figure 3.12: Figure showing the geometry of the CCD camera and imaging beam when imaging in Fast Kinetics mode. The black arrows indicate the direction of probe beam propagation. Figure not to scale.

vacuum system, the beam is directed horizontally toward the imaging optics and the camera. The remaining optical elements in the imaging beam path are shown in Figure 3.12.

After exiting the polarizing beam splitting cube, the imaging beam passes through a converging achromat lens. Immediately after the achromat is an iris with an aperture roughly 0.5 inches in diameter. Light collected from the outermost portions of the imaging lens tends to degrade image quality and resolution. The iris ensures that only light striking the center portion of the lens, where aberration is less prevalent, is directed to the camera. For all of the data discussed in this dissertation, the camera was operating in Fast Kinetics acquisition mode, a topic covered in greater detail in Section 5.1. When operating in Fast Kinetics mode, a razor blade is placed in the image plane of the large achromat lens. In the normal imaging mode, the razor blade is removed. Finally, after a portion of the imaging beam is blocked by the razor blade, the remaining portion is collected by a microscope objective fixed to the end of a bayonet mount on the front of the CCD camera. Images are acquired using an Andor Technology DV434-BV CCD

camera, featuring a 1024×1024 array of pixels (pixel size = $13 \mu\text{m}$ on a side). The acquisition and processing of images is discussed extensively in Chapter 5.

3.2.12 Timing system

Throughout Section 3.2, I have described the operation of many of the major elements in our experimental system. The task of making this collection of instruments play nicely with each other falls to the timing system. The timing system consists primarily of a computer-controlled 32-channel I/O box and GPIB output managed by LabVIEW software. These elements in turn control the actions of a number of GPIB devices on the system, as well as elements which respond to digital inputs. Of course, given the complexity of the system, the operation of some physical instruments requires switching between different values of analog signals. For this reason, a large custom built multiplexer translates the digital signals from the 32-channel control box into the appropriate analog signals. Much of the software design and development for the system was completed by Stephen Granade, and the operation of the timing system is described extensively in his dissertation [23]. Many of the software upgrades implemented in recent years are outlined in Appendix C of this dissertation.

3.3 Typical experimental sequence

In Sections 3.1 and 3.2, I introduced and described the operation of the basic tools used to cool and trap atoms in our laboratory. In this section, I will present a chronological account of a typical experimental sequence. While the sequence of events which occur during a typical experimental cycle are the same for nearly

every experiment we conduct, some details of the process differ based on whether we are producing a noninteracting or a strongly interacting Fermi gas. In Sections 3.3.1 and 3.3.2, I will address the production of noninteracting and strongly interacting Fermi gases, respectively.

The experimental cycle begins with the loading of the MOT. A vapor of ${}^6\text{Li}$ atoms is produced in the oven, which collimates the vapor beam and directs it to the Zeeman slower. Inside the slower, the motion of the atomic beam is opposed by a slowing beam, and the velocity of the atoms is substantially reduced. At the end of the Zeeman slower, the atoms have slowed to a point where they can be captured by a magneto-optical trap (MOT). During the MOT loading process, the MOT bias coils shift the MOT position slightly to maximize overlap between the MOT and the FORT. Once the MOT loading sequence is complete, the optical slowing beam is turned off via a mechanical shutter.

Throughout the MOT loading sequence, the FORT is also present. We can increase the capture of atoms from the MOT by maximizing the power in the CO_2 laser beam which forms the FORT. To create maximum FORT depth, the chopper is in the “up” position during the loading process, allowing the retroreflected CO_2 beam to overlap the forward-going beam. To further improve the efficiency of FORT loading, we conduct two phases at the conclusion of the MOT loading sequence. First, we conduct a “cooling phase,” during which the MOT beams are ramped closer to the atomic resonance while simultaneously being reduced in intensity. The cooling phase lasts 5 ms, and brings the atoms in the MOT to a temperature of roughly $140 \mu\text{K}$. At the conclusion of the cooling phase, a brief ($200 \mu\text{s}$) “optical pumping” phase is conducted, during which atoms are pumped into the lowest hyperfine ground state of the ${}^6\text{Li}$ atom. Once the optical pumping

phase has been completed, the MOT beams are extinguished. We are left with approximately 2 million atoms in the FORT, at a temperature of $140 \mu\text{K}$.

After the MOT beams have been extinguished, and the FORT loaded, the atoms are still sitting at the center of the spherical quadrupole magnetic field, where the magnetic field is zero. The high field magnets providing the spherical quadrupole magnetic field are turned off, and the direction of current flow through one of the magnets is reversed using the magnet “flipper” box. An 8 gauss bias magnetic field produced by one pair of the MOT bias coils is applied to the atoms confined in the FORT. The application of this magnetic field lifts the two-fold degeneracy which exists in the lowest hyperfine ground state of ${}^6\text{Li}$ at zero magnetic field. Two new eigenstates emerge, labelled the $|1\rangle$ and $|2\rangle$ states. A noisy RF pulse is then applied to the atoms. The center frequency of the RF pulse matches the frequency splitting between the $|1\rangle$ and $|2\rangle$ at 8 gauss, and results in equal spin populations of the $|1\rangle$ and $|2\rangle$ states.

The high field magnets are then turned on after which the MOT bias coils providing the 8 gauss field are turned off. The field supplied by the high field magnets is brought to 300 gauss, where 10 seconds of free evaporation occurs. Several seconds after the start of free evaporation, the retroreflected portion of the CO_2 laser beam which forms the FORT is adiabatically blocked as the chopper is slowly lowered, leaving only the forward-going beam to provide the confining potential. This results in a weaker trapping potential, but no position noise from the backgoing beam. At the conclusion of the free evaporation sequence, we have roughly 1 million atoms at a temperature of around $50 \mu\text{K}$.

Once free evaporation is complete, forced evaporation can begin. The choice of magnetic field for forced evaporation depends on whether the user wishes to

produce a noninteracting (Section 3.3.1) or strongly interacting (Section 3.3.2) gas. The trap depth of the FORT is controllably lowered to allow hot atoms to be ejected from the trap. The remaining atoms rethermalize and result in an overall cooler atom cloud with approximately 2×10^5 atoms. The depth of the FORT is then adiabatically recompressed to a user-defined value. For more information about the lowering curves which guide the process of forced evaporation, see Appendix C.3.

Following the desired experimental manipulation (see Chapters 6 and 7 for techniques used to study the heat capacity and breathing mode of the confined gas), information is extracted from the system in the form of an absorption image. Unfortunately, the size of the trapped atom cloud is too small to be imaged reliably, so we must first release the atom cloud from the trap. Following release, the cloud begins to expand, and the manner in which this expansion occurs depends on the nature of the interactions between the two spin states which comprise the cloud (see Chapter 4). Once the cloud has expanded to a sufficiently large size, a resonant pulse of light strikes the cloud, and the resulting image is recorded on a CCD camera. This process is destructive in nature, and results in the atoms in the atom cloud being scattered. After these atoms have had time to disperse, a second pulse of light strikes the CCD array. By comparing the photon counts in this second pulse to the photon counts in the image of the cloud, the spatial distribution of the atoms can be reconstructed (see Chapter 5).

At this point, the experimental cycle is effectively complete. The experimental apparatus returns to its MOT loading configuration in preparation for the next cycle. In closing, we include a table (Table 3.2) summarizing the frequency shifts provided by AOs during different phases of the experimental sequence. These

AO Name	MOT Loading	MOT Cooling	MOT Optical pumping	Camera beam
Locking	+200	+200	+200	+200
MOT	+170	+196	+198	Off
Repumper	+247	+228	Off	Off
Camera AO #1	Off	Off	Off	+214
Camera AO #2	Off	Off	Off	+214

Table 3.2: Summary of the frequency shifts provided by the AOs during different experimental phases. All AOs operate in a double-pass configuration and the reported frequency shifts are given in MHz after completing both passes through the AO. The values reported in the table provide ballpark figures, and may need to be modified to maximize system performance.

values are approximate, and will depend on other factors such as the overall laser power and the selected beam power during different experimental phases. While these frequency shifts are not to be taken as the definitive experimental settings, they do provide ballpark figures that have been associated with good performance from the experimental apparatus.

3.3.1 Production of a degenerate noninteracting Fermi gas

When producing a noninteracting Fermi gas, the loading of the MOT and the FORT proceed as described in Section 3.3. Following the process of equalizing the spin state populations, the magnetic field is ramped to 300 gauss, where free evaporation takes place. Following 10 seconds of free evaporation, forced evaporation is initiated while the magnetic field remains at 300 gauss. The depth of the optical trap is lowered according to (3.10), where we use a lowering time constant $\tau_{lc} = 0.25$ seconds for a duration of roughly 12 seconds. The optical trap is held at this minimum depth for roughly 1 second before being recompressed to some final value in 1 or 2 seconds. We typically wait 0.5 to 1 second before

initiating any further experimental manipulation, to allow the atom sample to come to equilibrium following the optical trap recompression sequence. When the experimental system is operating properly, the above procedure will net about 10^5 atoms per spin state at a reduced temperature of $T/T_F \approx 0.25$. Of course, as long as we maintain the magnetic field at 300 gauss where the gas is weakly interacting, we do not have a truly noninteracting sample. To reach the noninteracting limit, we have two options:

1. Ramp the magnetic field to 0 gauss, where the $|1\rangle$ - $|2\rangle$ scattering length is 0.
2. Ramp the magnetic field to the so called “zero crossing” near 530 gauss [62], where the scattering length is also 0.

In practice, the latter of these options is superior on two counts. First, following the experimental manipulation of the atomic sample, we will need to image the atom cloud. At zero magnetic field, there is substantial optical pumping to dark states during the on-resonance probe pulse [24], resulting in reduced signal-to-noise. This effect is mitigated at the zero crossing, where the imaging transitions are well-approximated by a two-level system. This effect is discussed in greater detail in Section 5.4. As the imaging transition is nearly a two-level system, and therefore undergoes minimal optical pumping to dark states during the probe pulse, the signal-to-noise which results from imaging at 530 gauss is superior to that at zero magnetic field.

There is a second, and in some ways, more physically interesting reason we elect to produce a noninteracting gas by working at the zero crossing rather than at zero magnetic field. Early attempts on our part to study a noninteracting gas by ramping the applied magnetic field to zero gauss resulted in an apparent

“explosion” of the gas. Comparison of images acquired at 300 gauss prior to the magnetic field sweep and those images taken after the sweep to zero magnetic field revealed a substantial loss of atoms as well as significant heating of the remaining atoms. At the time, we were unable to explain the cause of this behavior, but more recent studies [84,90] reveal that there are three p-wave Feshbach resonances for a $|1\rangle - |2\rangle$ mixture between 159 and 215 gauss. As the magnets which apply the bias magnetic field are quite large, the time needed to conduct magnetic field sweeps is on the order of 0.35 seconds. At such a sweep rate, the atom cloud appears to spend a sufficiently large amount of time in the vicinity of these p-wave resonances for there to be significant atom loss and heating. Fortunately, there are no such problems in conducting a magnetic field ramp from 300 gauss upwards to the zero crossing, and as a result, we produce noninteracting gases by working near the zero crossing.

Those familiar with the properties of $|1\rangle$ - $|2\rangle$ mixtures of the ${}^6\text{Li}$ ground state might be tempted to ask: why not produce a noninteracting gas by performing free and forced evaporation near the broad Feshbach resonance centered at 834 gauss before sweeping the magnetic field to the zero crossing? After all, evaporative cooling could be conducted more quickly and efficiently at 834 gauss, where we have a unitarity limited scattering length, compared to the relatively weak scattering length at 300 gauss. There are two reasons we do not produce noninteracting samples in this way. First, we cannot leave the large bias magnets at high magnetic field (say, 800 gauss, or larger) for more than about 10 seconds, before protective interlocks engage to prevent the magnets from overheating. Efforts to conduct free and forced evaporation near the broad Feshbach resonance would likely tax these safety limits. The second reason we do not produce non-

interacting gases by cooling near the broad Feshbach resonance before ramping the magnetic field to the zero crossing is due to heating between 650 and 680 gauss [61, 62].

Finally, one might ask: why not generate a noninteracting gas by utilizing the large scattering lengths available near the narrow Feshbach resonance at 544 gauss to conduct evaporative cooling prior to ramping the magnetic field to the (nearby) zero crossing value? This is not experimentally viable, as the so-called narrow Feshbach resonance is quite aptly named. The width of the resonance is less than 1 gauss [50, 84], making it too narrow to study reliably.³ The required magnetic field precision exceeds the precision of our magnetic field controls. Furthermore, as the large magnets do not produce a completely uniform bias magnetic field, it is possible for atoms contained in different portions of the optical trap to exist on different sides of the narrow Feshbach resonance. (That is, an atom at the very center of the magnetic field axis might occupy a region in space where the magnetic field is above the Feshbach resonance value, while an atom far from the magnetic field axis might experience a bias magnetic field whose magnitude is less than the narrow Feshbach resonance value). While the narrow Feshbach resonance centered at 544 gauss might eventually prove to be an interesting object of study, we did not attempt to examine it for the data in this dissertation. Unless otherwise stated, references to “the Feshbach resonance” indicate we are speaking about the broad Feshbach resonance centered at 834 gauss.

³One of the few studies using the narrow resonance is presented in [50].

3.3.2 Production of a degenerate strongly interacting Fermi gas

Compared to the issues considered in Section 3.3.1 regarding the production of a noninteracting Fermi gas, the production of a strongly interacting Fermi gas is comparatively simple. Following free evaporation for 10 seconds at a bias magnetic field of 300 gauss, the magnetic field is ramped to a destination value near the broad Feshbach resonance centered at 834 gauss. Forced evaporation starts with the lowering of the optical trap depth for 3-4 seconds with a typical time constant of 80 ms. (The specific lowering durations and time constants depend on the magnetic field and the efficiency of FORT loading from the MOT, and are therefore tweaked often.) Note that this is shorter than the lowering times and time constants used for the noninteracting gas; the very large scattering length which exists near the Feshbach resonance allows evaporative cooling to proceed with great efficiency, which results in increased speed. After the lowering sequence is complete, the optical trap depth is held at its minimum value for roughly 0.5 seconds before being recompressed to the desired final trap depth in approximately 1 second. After allowing 0.5 seconds to ensure that the gas is in equilibrium at the final optical trap depth, the atom cloud is ready for further experimental manipulation. When the experimental apparatus is functioning properly, the above sequence yields strongly interacting Fermi gases containing 2×10^5 atoms at estimated reduced temperatures $T/T_F \leq 0.10$.⁴

One of the advantages of producing a strongly interacting gas near the Fesh-

⁴Why “estimated” reduced temperatures? In the strongly interacting regime, temperature measurement remains controversial. See Section 6.3.2 for our pragmatic solution to the problem of measuring temperatures in the strongly interacting regime.

bach resonance is that we can employ on-resonance imaging techniques without having to change the magnetic field. As discussed in Section 5.4, the optical transitions used in imaging are very nearly closed two-level cycling transitions, which maximize the signal-to-noise ratio. The only scenario for which optical imaging becomes problematic is well below the center of the Feshbach resonance, where it is energetically favorable for ${}^6\text{Li}$ atoms to form ${}^6\text{Li}_2$ dimers. Trying to image these dimers by tuning the probe beam frequency to the free atom optical transition will be successful provided that the dimers are weakly bound. If the binding is sufficiently weak, an atom in the dimer will absorb a photon and the associated momentum recoil will break the pair, yielding two atoms which can then be imaged readily. In practice, we find that the apparent atom number and signal-to-noise ratio is decreased when imaging at magnetic fields below the Feshbach resonance, where ${}^6\text{Li}_2$ dimers are formed.

With the completion of the descriptions of the production of interacting and noninteracting degenerate Fermi gases, we are in a position to discuss extracting data from the system. We do so by acquiring absorption images of the clouds, but we must release the clouds from the trapping potential before performing absorption imaging. We cover the topic of expansion dynamics in the next chapter.

Chapter 4

Expansion Dynamics

All of the images of atom clouds analyzed and presented in this dissertation were acquired following release of the atoms from the optical trap. Ideally, we would be able to image the trapped clouds directly, as that would eliminate the need to understand the expansion dynamics of the cloud. However, imaging of expanded clouds is necessary for two reasons. First, the transverse size of the trapped cloud is comparable to the resolution of our imaging system. Extracting information about the spatial density of the cloud requires that the overall dimensions of the cloud be much larger than the resolution of our imaging system. Second, the cloud is released from the optical trap because we want to reduce the optical thickness of the cloud. The trapped cloud is so optically thick that attempts to calculate the spatial profile of the cloud based on on-resonance absorption imaging become unreliable. Releasing the cloud from the optical trap overcomes both of these problems: the expanded cloud is much larger than the resolution of our imaging system, and the expanded cloud is less optically thick.

The expansion dynamics of quantum gases is a rich topic, particularly in the strongly interacting regime, where high collision rates and superfluidity strongly influence the evolution of the expanding atom cloud. In this section, I will discuss the expansion dynamics of a noninteracting atom cloud (ballistic expansion) and

the evolution of an expanding hydrodynamic cloud. The latter case is applicable in the strongly interacting regime, where the gas appears to obey the hydrodynamic equations when trapped [19,21,79] and following release from the optical trap [16].

In the sections that follow, we will make use of the relaxation approximation to obtain formulas governing the time-dependent scale factors $b_i(t)$ that describe the evolution of the overall dimensions of the atom cloud. For a cloud expanding by a scale transformation, the shape of the cloud will remain the same following release from the trap, and the size of the cloud in the i -direction at time t is $b_i(t)$ times its trapped size. We will find that the expansion dynamics are governed by the optical trap oscillation frequencies and the slight magnetic field curvature provided by our magnets. Procedures for measuring these oscillation frequencies are covered at the end of this chapter.

4.1 The Relaxation Approximation

Ideally, one would predict the behavior of a gas by measuring the initial conditions for each particle and deterministically evolving the state of the system, thereby allowing one to make predictions for the macroscopic observables. Of course, for systems involving a large number of particles, this approach is unreasonable. While predicting the behavior of individual particles in the atom cloud is unreasonable, predicting the evolution of the phase space density is not. Here, the behavior of single atoms is disregarded, and the statistical behavior of the aggregate sample followed instead. This technique is a powerful one, as we can explore the kinds of behavior a system might exhibit without possessing detailed knowledge of the interparticle interactions. In the following, we will consider the

evolution of the *classical* phase space density.

We begin by taking the total time derivative of the phase space density $f(t, x_i, v_i) = f(t, x, y, z, v_x, v_y, v_z)$, yielding

$$\frac{\partial f}{\partial t} + \mathbf{v} \cdot \frac{\partial f}{\partial \mathbf{r}} - \frac{1}{m} \frac{\partial U}{\partial \mathbf{r}} \cdot \frac{\partial f}{\partial \mathbf{v}} = I_{coll}[f], \quad (4.1)$$

where $I_{coll}[f]$ is an integral related to collisional processes in the gas. While one could proceed by writing down and solving the expression for the collisional integral (which would require knowledge of interparticle interactions), we will instead use the relaxation approximation (also referred to as the relaxation time approximation, [100]). This approximation involves the following simplification [101]

$$I_{coll}[f] \approx -\frac{f - f_{le}}{\tau_R}, \quad (4.2)$$

where f_{le} is the phase space density for local equilibrium conditions and the relaxation time τ_R is related to the average time between collisions in the gas. Next, we use the phase space density scaling ansatz [101]

$$f(t, x_i, v_i) = \frac{1}{\prod_j (b_j \theta_j^{1/2})} f_0 \left[\tilde{x}_i, \frac{1}{\theta_i^{1/2}} (v_i - \dot{b}_i \tilde{x}_i) \right], \quad (4.3)$$

where $\tilde{x}_i = x_i/b_i$, the $b_i = b_i(t)$ factors track the overall size of the gas in the i -direction, dots represent derivatives with respect to time, and the $\theta_i = \theta_i(t)$ factors track the effective temperature in the i -direction. Further, f_0 denotes the equilibrium phase space density. Combining (4.1), (4.2), and (4.3), we arrive at

the following system of equations,

$$\ddot{b}_i + \frac{1}{b_i m \langle \tilde{x}_i^2 \rangle} \left[\left\langle \tilde{x}_i \frac{\partial U_{EV}(b_x \tilde{x}, b_y \tilde{y}, b_z \tilde{z})}{\partial \tilde{x}_i} \right\rangle - \theta_i \left\langle \tilde{x}_i \frac{\partial U_{EQ}(\tilde{x}, \tilde{y}, \tilde{z})}{\partial \tilde{x}_i} \right\rangle \right] = 0 \quad (4.4)$$

$$\dot{\theta}_i + 2 \theta_i \frac{\dot{b}_i}{b_i} = -\frac{1}{\tau_R} (\theta_i - \bar{\theta}), \quad (4.5)$$

where $\bar{\theta} = (\theta_x + \theta_y + \theta_z)/3$ in the latter equation. In (4.4), I have used two different forms of the potential energy. U_{EQ} represents the potential energy terms present under equilibrium conditions prior to perturbations. U_{EV} represents potential energy terms relevant as the system evolves in time. For scenarios in which the potential energy contributions that establish the initial equilibrium remain constant as the system evolves in time, the functional forms of U_{EQ} and U_{EV} will be similar. However, in general, the two potential energy terms will not be equal, as we will see in the coming sections.

4.2 Ballistic expansion

We begin with the simple case of a noninteracting gas. At first, we will consider expansion from a harmonic optical potential with trap oscillation frequencies ω_i . Ballistic expansion is observed in degenerate mixtures of the $|1\rangle$ and $|2\rangle$ hyperfine states of ${}^6\text{Li}$ at zero magnetic field [48] and at the “zero crossing” [52, 62] near 530 gauss, where the s -wave scattering length between the two hyperfine states is zero. In this section, we consider ballistic expansion at zero magnetic field. Ballistic expansion in the presence of a magnetic field, a scenario that occurs at the zero crossing, will be addressed in Section 4.3.

We start with the equations derived in Section 4.1 using the relaxation ap-

proximation

$$\ddot{b}_i + \frac{1}{b_i m \langle \tilde{x}_i^2 \rangle} \left[\left\langle \tilde{x}_i \frac{\partial U_{EV}(b_x \tilde{x}, b_y \tilde{y}, b_z \tilde{z})}{\partial \tilde{x}_i} \right\rangle - \theta_i \left\langle \tilde{x}_i \frac{\partial U_{EQ}(\tilde{x}, \tilde{y}, \tilde{z})}{\partial \tilde{x}_i} \right\rangle \right] = 0 \quad (4.6)$$

$$\dot{\theta}_i + 2 \theta_i \frac{\dot{b}_i}{b_i} = -\frac{1}{\tau_R} (\theta_i - \bar{\theta}). \quad (4.7)$$

Recall that τ_R is related to the characteristic time between collisions of two particles in the gas. In the noninteracting limit, there are no collisions, and $\tau_R \rightarrow \infty$. In this case, (4.7) yields $\theta_i = 1/b_i^2$. Prior to release from the optical trap, the only potential energy present is that provided by the optical trap. If we assume the optical trap provides harmonic confinement, then

$$U_{EQ}(\tilde{x}, \tilde{y}, \tilde{z}) = \frac{m}{2} (\omega_x^2 \tilde{x}^2 + \omega_y^2 \tilde{y}^2 + \omega_z^2 \tilde{z}^2). \quad (4.8)$$

After the optical trap is extinguished, the atom cloud evolves free of external forces.¹ In this case, $U_{EV} = 0$. The above system of six equations then simplifies to a three equation system,

$$\ddot{b}_i - \frac{1}{m b_i^3 \langle \tilde{x}_i^2 \rangle} \left\langle \tilde{x}_i \frac{\partial U_{EQ}(\tilde{x}, \tilde{y}, \tilde{z})}{\partial \tilde{x}_i} \right\rangle = 0. \quad (4.9)$$

Combining (4.8) and (4.9), we arrive at a simple differential equation for the $b_i(t)$ factors,

$$\ddot{b}_i - \frac{\omega_i^2}{b_i^3} = 0. \quad (4.10)$$

¹We assume the effects of gravity are negligible for typical experimental conditions. Further, we are not concerned with the center of mass motion of the atom cloud, so much as we are interested in the temporal evolution of the cloud's spatial profile.

To (4.10) we add the initial conditions

$$b_i(0) = 1 \quad (4.11)$$

$$\dot{b}_i(0) = 0. \quad (4.12)$$

Solving (4.10) for these initial conditions yields

$$b_i(t) = \sqrt{1 + (\omega_i t)^2}. \quad (4.13)$$

The physical meaning of the $b_i(t)$ factors is straightforward. An atom cloud with width σ_i in the i -direction at time $t = 0$ will have a width $b_i(t) \times \sigma_i$ at time t .

4.3 Ballistic expansion with magnetic field

In the preceding section, we considered the expansion dynamics of a noninteracting gas. The results of Section 4.2 are valid for a noninteracting system where there are no additional forces present during time of flight. This would be the case for the two lowest hyperfine ground states of ${}^6\text{Li}$ at zero magnetic field. However, for reasons addressed in Section 3.3.1, we prefer to produce noninteracting gases at the zero crossing near 530 gauss.

For ballistic expansion near the zero crossing, the presence of the magnetic field can alter the expansion dynamics. If the magnetic field generated by the large magnets were perfectly uniform, the gradient of the magnetic field would be zero, and no magnetic force would be exerted on the atoms. In reality, however, the finite size of the magnets results in finite curvature of the magnetic field. The presence of this magnetic field curvature exerts a force on the atoms as they

expand. Although the presence of the magnetic field gradient represents a small perturbation for typical experimental conditions, the effect can be significant if the atom cloud is released from a shallow optical trap, as detailed in Section 4.6.2.

We return to equations (4.4) and (4.5). Once again, we let $\tau_R \rightarrow \infty$ to obtain

$$\ddot{b}_i + \frac{1}{m b_i \langle \tilde{x}_i^2 \rangle} \left[\left\langle \tilde{x}_i \frac{\partial U_{EV}(b_x \tilde{x}, b_y \tilde{y}, b_z \tilde{z})}{\partial \tilde{x}_i} \right\rangle - \frac{1}{b_i^2} \left\langle \tilde{x}_i \frac{\partial U_{EQ}(\tilde{x}, \tilde{y}, \tilde{z})}{\partial \tilde{x}_i} \right\rangle \right] = 0. \quad (4.14)$$

In contrast to the analysis of Section 4.2, magnetic forces are present at all times. Consequently, if we approximate both the dipole optical potential and magnetic potential as three-dimensional harmonic potentials, U_{EV} and U_{EQ} are

$$U_{EV}(b_x \tilde{x}, b_y \tilde{y}, b_z \tilde{z}) = \frac{m}{2} (\omega_{mx}^2 b_x^2 \tilde{x}^2 - \omega_{my}^2 b_y^2 \tilde{y}^2 + \omega_{mz}^2 b_z^2 \tilde{z}^2) \quad (4.15)$$

$$U_{EQ}(\tilde{x}, \tilde{y}, \tilde{z}) = \frac{m}{2} [(\omega_x^2 + \omega_{mx}^2) \tilde{x}^2 + (\omega_y^2 - \omega_{my}^2) \tilde{y}^2 + (\omega_z^2 + \omega_{mz}^2) \tilde{z}^2], \quad (4.16)$$

where ω_i represents the optical trap oscillation frequency in the i -direction, and ω_{mi} represents the magnetic potential oscillation frequency in the i -direction. Before continuing, note that the magnetic potential is attractive in the x - and z -directions, but repulsive in the y -direction (along the magnet axis), a consequence of the symmetry of the magnet's design and magnetostatics which is discussed in Section 4.6.2.

If we plug (4.15) and (4.16) into (4.14), the equations governing the time

evolution of the b_i factors are as follows,

$$\ddot{b}_x + \omega_{mx}^2 b_x - \frac{\omega_x^2 + \omega_{mx}^2}{b_x^3} = 0 \quad (4.17)$$

$$\ddot{b}_y - \omega_{my}^2 b_y - \frac{\omega_y^2 - \omega_{my}^2}{b_y^3} = 0 \quad (4.18)$$

$$\ddot{b}_z + \omega_{mz}^2 b_z - \frac{\omega_z^2 + \omega_{mz}^2}{b_z^3} = 0. \quad (4.19)$$

In the absence of the magnetic field curvature, we were able to obtain a simple analytic solution describing ballistic expansion of the noninteracting gas. Unfortunately, the presence of the magnetic field potential prevents us from doing the same here. Equations (4.17) through (4.19) must be solved numerically, making use of the initial conditions (4.11) and (4.12).

4.4 Hydrodynamic expansion with magnetic field

As the case of ballistic expansion is well understood, and as the noninteracting regime can be readily accessed given the two hyperfine states of ${}^6\text{Li}$ we trap, it is natural to question why one would conduct time of flight expansion under any other circumstances. Indeed, if one is using time of flight expansion as a tool to probe the gas, rather than a probe of the expansion dynamics themselves, it would be more convenient to conduct all time of flight measurements in the noninteracting regime. Unfortunately, technical limitations prevent us from doing so. Much of the interesting physics of this system occurs near the broad Feshbach resonance located near 834 G. If one can very quickly sweep from the Feshbach resonance region to the zero crossing near 530 G, then all time of flight measurements can be conducted in the noninteracting regime. In our laboratory, a magnetic field sweep

of around 300 gauss can take several tenths of a second, much longer than the microsecond time scales we would require to ensure that the sweep rate was much faster than all physically relevant time scales. Consequently, we are often forced to conduct our time of flight measurements in the strongly interacting regime, where we have observed very good hydrodynamic behavior.

The underlying cause of hydrodynamic expansion can be a high collision rate in the gas, or the presence of superfluidity in the gas. While there is still some debate about the cause of hydrodynamic behavior in degenerate Fermi gases under various conditions, the behavior appears to be well described by the hydrodynamic equations of motion. Unlike the case of ballistic expansion, we will not consider hydrodynamic expansion in the absence of a magnetic field. Observations of hydrodynamic behavior occur at magnetic fields in the vicinity of the broad Feshbach resonance centered at 834 gauss, where the gas is strongly interacting. Once these interactions are removed, hydrodynamic behavior is not observed. Hence, we will consider the most experimentally relevant scenario in which a hydrodynamic gas expands from a combined optical and magnetic potential into a purely magnetic potential.

Once again, we begin with the equations based on the relaxation approximation. For the noninteracting gas, we let the characteristic time between collisions go to infinity. Here, we consider the case of an infinite collision rate, in which the characteristic time between collisions goes to zero. In that case, $\tau_R \rightarrow 0$. Then the equations governing the b_i are given by

$$\ddot{b}_i + \frac{1}{m b_i \langle \tilde{x}_i^2 \rangle} \left[\left\langle \tilde{x}_i \frac{\partial U_{EV}(b_x \tilde{x}, b_y \tilde{y}, b_z \tilde{z})}{\partial \tilde{x}_i} \right\rangle - \frac{1}{\Gamma^{2/3}} \left\langle \tilde{x}_i \frac{\partial U_{EQ}(\tilde{x}, \tilde{y}, \tilde{z})}{\partial \tilde{x}_i} \right\rangle \right] = 0, \quad (4.20)$$

where $\Gamma \equiv b_x b_y b_z$. Once again, we approximate both the magnetic and optical potentials as three-dimensional harmonic potentials. Substituting (4.15) and (4.16) into (4.20), we obtain the following system of equations,

$$\ddot{b}_x + \omega_{mx}^2 b_x - \frac{\omega_x^2 + \omega_{mx}^2}{b_x \Gamma^{2/3}} = 0 \quad (4.21)$$

$$\ddot{b}_y - \omega_{my}^2 b_y - \frac{\omega_y^2 - \omega_{my}^2}{b_y \Gamma^{2/3}} = 0 \quad (4.22)$$

$$\ddot{b}_z + \omega_{mz}^2 b_z - \frac{\omega_z^2 + \omega_{mz}^2}{b_z \Gamma^{2/3}} = 0. \quad (4.23)$$

In contrast to the case for the noninteracting gas, where the expansion in each direction was independent, we see here that for the hydrodynamic gas, the equations governing the b_i must be solved simultaneously using numerical techniques. When numerically solving equations (4.21) through (4.19), we again use the initial conditions given by equations (4.11) and (4.12).

4.5 Expansion from a slightly anharmonic trap

In our studies of degenerate Fermi gases, we attempt to produce experimental conditions in which the atomic cloud occupies only the deepest portion of the optical potential, where treating the optical trap as a harmonic well should be approximately correct. Even with these efforts, we have found that slight anharmonicity can be significant for some of our measurements. In preceding sections, we have considered the expansion dynamics of clouds released from harmonic optical potentials with trap oscillation frequencies ω_i . How are these frequencies altered if the cloud expands from an anharmonic confining potential?

For a perfectly harmonic trap, the expansion dynamics are precisely deter-

mined by the trap oscillation frequencies. For slightly anharmonic traps, we can characterize the trap using an effective oscillation frequency. We consider the noninteracting limit given by (4.14), and set $U_{EV} = 0$. Then, we are left with

$$\ddot{b}_i - \frac{1}{m b_i^3 \langle \tilde{x}_i^2 \rangle} \left\langle \tilde{x}_i \frac{\partial U_{EQ}(\tilde{x}, \tilde{y}, \tilde{z})}{\partial \tilde{x}_i} \right\rangle = 0. \quad (4.24)$$

We can rewrite (4.24) as

$$\ddot{b}_i - \frac{(\omega_i^2)_{eff}}{b_i^3} = 0, \quad (4.25)$$

where the effective trap oscillation frequency is defined as

$$(\omega_i^2)_{eff} = \frac{1}{m \langle \tilde{x}_i^2 \rangle} \left\langle \tilde{x}_i \frac{\partial U_{EQ}(\tilde{x}, \tilde{y}, \tilde{z})}{\partial \tilde{x}_i} \right\rangle. \quad (4.26)$$

Note that for harmonic forms of U_{EQ} , the effective trap oscillation frequency will be the standard harmonic trap oscillation frequency. For “soft” potentials, such as the one produced by our gaussian optical dipole trap, the effective trap oscillation frequency will be lower than its harmonic counterpart.

4.6 Measuring trap oscillation frequencies

We have seen that the expansion dynamics in the noninteracting and hydrodynamic limits are governed by the optical trap oscillation frequencies ω_i and the magnetic potential oscillation frequencies, ω_{mi} , for $i = x, y, z$. Consequently, we need methods for measuring these quantities. Accurate measurement of the optical trap frequencies relies on a technique known as parametric resonance. The

results from parametric resonance measurements can be verified by measuring the breathing mode frequency of the noninteracting gas as well. Measurement of the effect of magnetic field curvature can be accomplished by monitoring the center of mass motion of the atomic cloud following release into the magnetic potential. We begin, however, with the optical trap.

4.6.1 Measuring optical trap oscillation frequencies

The importance of measuring the optical trap oscillation frequencies accurately cannot be overstated. The physical behavior of the trapped atom cloud is determined primarily by two factors: the interparticle interactions, and the confining potential. If we can characterize accurately and precisely the confining potential of the optical trap, this will leave us free to study the more physically interesting interparticle interactions in the trapped atom cloud. We have seen already that the optical trap oscillation frequencies influence the dynamics of an expanding atom cloud. Perhaps more importantly, however, the optical trap parameters help determine the natural temperature and energy scales of the trapped atom cloud. In this section, I will discuss the use of parametric resonance techniques and measurements of the breathing mode as methods to characterize the optical trap oscillation frequencies.

The parametric resonance technique is relatively straightforward. First, we evaporatively cool a 50-50 mixture of the two lowest hyperfine spin states of ${}^6\text{Li}$ at 300 gauss to the lowest temperatures we can achieve. Then, the optical trap is recompressed to the desired depth. At this point, we initiate the process of parametric resonance, in which the optical trap experiences small amplitude modulation at a particular drive frequency. After the amplitude modulation is

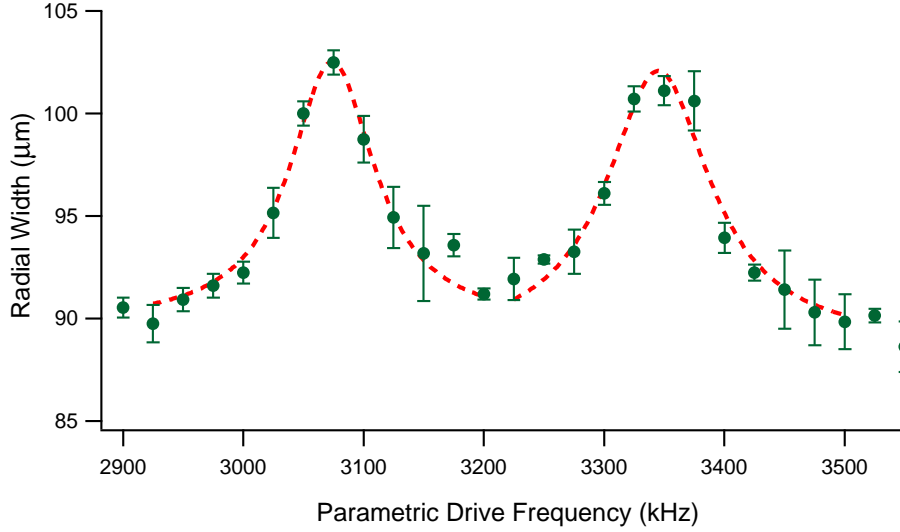


Figure 4.1: The results of a parametric resonance measurement of the radial trap oscillation frequencies for our optical potential. We monitor the radial width of the released atom cloud as a function of the frequency of the amplitude modulation (the parametric drive frequency). Spikes in the width are observed at frequencies which are twice the natural trap oscillation frequencies. Dots with error bars are the measured radial widths of the atom cloud (measured approximately using a zero temperature Thomas-Fermi spatial profile) and the dashed lines are Lorentzian fits to the peaks. The above data was taken for an optical trap depth $U_0 = 35 \mu\text{K}/k_B$, roughly 5% of the maximum trap depth.

complete, the magnetic field is ramped to the zero crossing at 530 gauss, where the atom cloud is released from the trap and imaged. The width of the released atom cloud can be tracked as a function of the amplitude modulation frequency to produce plots like Figures 4.1 and 4.2.

In Figures 4.1 and 4.2, we see the results of performing the parametric resonance procedure in a frequency range close to twice the natural trap oscillation frequencies. While it is tempting to assume that the resonant excitation of the atom cloud will be strongest when driving at the natural trap frequencies, in fact the phenomenon of parametric resonance is observed when exciting the system

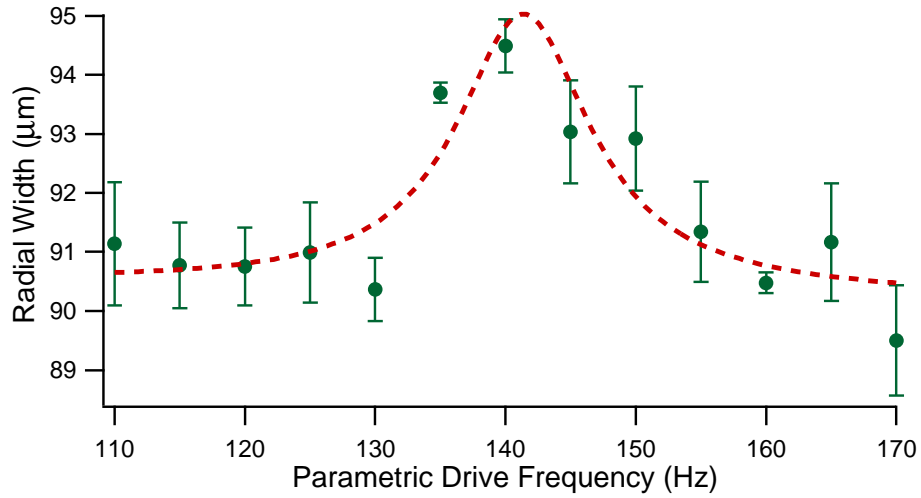


Figure 4.2: The results of a parametric resonance measurement of the axial trap oscillation frequency for our optical potential. We monitor the radial width of the released atom cloud as a function of the frequency of the amplitude modulation (the parametric drive frequency). A spike in the width is observed at twice the natural axial trap oscillation frequency. Dots with error bars are the measured radial widths of the atom cloud (measured approximately using a zero temperature Thomas-Fermi spatial profile) and the dashed line is a Lorentzian fit to the peak. The above data was taken for an optical trap depth $U_0 = 35 \mu\text{K}/k_B$, roughly 5% of the maximum trap depth.

at twice its natural frequencies [102]. The spikes in the measured cloud width indicate that the atom cloud has become hotter at those drive frequencies, which in turn indicates that the drive frequency is close to twice one of the natural frequencies of the optical trap. To obtain accurate parametric resonance results, we typically start with an atom cloud whose reduced temperature $T/T_F < 0.3$. Starting from degenerate samples has two advantages. First, small increases in temperature are most easily observed at low temperature, where the increase in temperature, while small, can still constitute a noticeable fractional change in temperature. Second, working at low temperature where the atom cloud occupies the lowest portion of the trap minimizes the impact of trap anharmonicity on the measurement.

The results displayed in Figures 4.1 and 4.2 were acquired for an optical trap depth of $U_0 = 35 \mu K/k_B$, roughly 5% of the maximum trap depth. The amplitude modulation signal for the radial trap oscillation frequencies was a 2 mV peak-to-peak sine wave with a duration of 1 second. For the axial measurement, the amplitude modulation was a 5 mV peak-to-peak signal which lasted for 4 seconds. In both cases, these sine waves were passed through a 15 dB attenuator before being summed with the DC signal which provides the optical trap depth, yielding less than 1% amplitude modulation of the trap.

While parametric resonance gives us a reasonably accurate measurement of the optical trap oscillation frequencies, trap anharmonicity can result in an underestimate of the harmonic trap oscillation frequencies. Fortunately, we can measure the radial breathing mode of the trapped atom cloud, which will allow for a consistency check with the parametric resonance measurement. Further, we have simple analytic formulas for estimating the breathing mode oscillation

frequency shift resulting from trap anharmonicity.

The process of measuring the breathing mode of the trapped atom cloud is discussed at length in Chapter 7, and will be summarized here. After conducting evaporative cooling at 300 gauss, the magnetic field is ramped to the zero crossing at 530 gauss, and the optical trap is briefly extinguished. After the optical trap is restored, the atom cloud is trapped once more and begins to oscillate. The measured oscillation frequency should be in close agreement with one of the measured radial parametric resonance frequencies. Using the results in Section 7.4.2, we can calculate the breathing mode oscillation frequency shift arising from trap anharmonicity. This allows us to fix the “true” optical trap oscillation frequency for this dimension. Having established the most accurate measurement of one of the radial trap oscillation frequencies, we can use the ratio of the two radial trap oscillation frequencies measured using parametric resonance to fix the “true” trap oscillation frequency for the remaining radial trap dimension. Unfortunately, for technical reasons, it is inconvenient to excite the axial breathing mode of the trapped atom cloud to provide a similar consistency check and anharmonicity estimate for the measured optical trap oscillation frequency in the axial dimension. Consequently, we take the axial trap oscillation frequency measured using parametric resonance as the “true” oscillation frequency.

In closing, the astute reader might be concerned that the parametric resonance excitation process was conducted at 300 gauss rather than the zero crossing at 530 gauss. At the latter magnetic field, the gas is truly a noninteracting system, and we would expect to obtain the most accurate measurements of the natural trap oscillation frequencies. In contrast, at 300 gauss, the gas is weakly interacting, and the presence of interactions could lead to a systematic shift in the measured

trap oscillation frequencies. However, as the system is weakly interacting at 300 gauss, this should contribute only a small shift in the measured trap frequencies. This contention is supported by measurements in which parametric resonance is conducted at 300 gauss and at 530 gauss. We measure the same parametric excitation frequency for one of the radial dimensions in both cases. Moreover, conducting parametric resonance at 300 gauss is necessary to measure all of the trap oscillation frequencies, as we use the fractional change in size of the one radial width visible on the camera images as an indicator of heating. If we had a truly noninteracting system, pumping energy into the system axially would serve to increase the axial width of the cloud without impacting the radial width. Moreover, we would be completely unable to measure the second radial trap resonance, as there is no way to observe this dimension on the two-dimensional camera image. Conducting parametric resonance at 300 gauss allows for cross-relaxation between the primary axes of the optical trap. Adding energy to the atom cloud by exciting the axial mode will appear as additional radial energy if the atoms in the cloud are colliding. This process is possible at 300 gauss, but not at 530 gauss. Hence, we conduct our parametric resonance excitations at 300 gauss.

4.6.2 Measuring magnetic field curvature

The magnetic potential oscillation frequencies arising from the magnetic field curvature are measured by monitoring the center of mass motion of the atom cloud following release from the optical trap. In an ideal world, we would be able to measure all of the magnetic potential oscillation frequencies ω_{mi} independently. However, given the current geometry of our imaging system, we can measure only

two of the magnetic potential oscillation frequencies directly. In fact, we need to measure only one of these, as symmetry considerations and magnetostatics fix the values of the other two potential oscillation frequencies.

First, there are some practical issues to discuss. The magnetic potential oscillation frequencies will change as the magnetic field is varied. As a substantial number of our measurements were conducted at 840 gauss, slightly above the center of the broad Feshbach resonance at 834 gauss, I will present a measurement of the magnetic potential oscillation frequencies at 840 gauss. Note that we are interested only in the center of mass motion of the atom cloud. This might lead one to believe that the optical trap depth has no impact on the magnetic field curvature measurement. While this is true, the signal-to-noise ratio of the measurement can be improved by releasing the atom cloud from a very shallow optical trap. As we have seen in the preceding chapter regarding expansion dynamics, atom clouds confined in optical traps expand more rapidly as the optical trap oscillation frequencies are increased. Rapid expansion of the atom cloud results in lower column densities, and reduced signal-to-noise ratios. Hence, when measuring the magnetic potential oscillation frequencies, we release the atom cloud from the shallowest possible optical trap.

In Figure 4.3, we see false color absorption images of an atom cloud released from an optical trap into the magnetic potential at a magnetic field of 840 gauss. The top panel shows the density profile after 2 ms time of flight. The lower two images were taken after 6 and 8 ms time of flight. While the aspect ratio of the spatial density of the atom cloud is changing, we are concerned instead with the center of mass motion of the cloud.² This is shown in Figure 4.4, where

²In this particular panel of images, we see that the cloud moves only to the left of the image,

the transverse position of the cloud is plotted as a function of time elapsed from the removal of the optical trap. The dots represent the measured position of the atom cloud, while the dashed line shows a sinusoidal fit to the data, from which we can extract the magnetic potential oscillation frequencies. Due to the cylindrical symmetry of our magnet design, the two orthogonal dimensions (x and z) visible in the camera images have the same magnetic potential oscillation frequencies, $\omega_{mx} = \omega_{mz}$. As only a fraction of a single period of motion can be obtained, repeated measurements of the magnetic potential oscillation frequencies yield a range of values. At a magnetic field of 840 gauss, for the two directions visible in the camera images, we measure $2\pi \times 19 < \omega_{mx} < 2\pi \times 24$.

While we can determine ω_{mx} and ω_{mz} experimentally with little effort, we cannot directly measure the final magnetic potential oscillation frequency, ω_{my} . We are unable to track the motion of the atom cloud in the y -direction, as it coincides with the direction of propagation of our imaging beam. However, basic magnetostatics allows us to calculate the magnetic potential oscillation frequency in this direction.

In the discussion that follows, we will assume that the magnetic dipole moment of the atom adiabatically follows the local magnetic field direction. This will be the case if the atom moves sufficiently slowly. For the low temperatures at which most of our experiments are conducted, this is a reasonable approximation. Furthermore, the $|1\rangle$ and $|2\rangle$ hyperfine states of ${}^6\text{Li}$, are called “high field seeking” states (see [22–24] for example), as their magnetic energy is minimized as an

indicating that the atom cloud is located at the equilibrium position of the magnetic potential in the other visible dimension. If the atom cloud were not in the center of the magnetic potential in either dimension visible on the camera, the center of mass of the cloud would move roughly diagonally in the image plane following release from the optical trap.

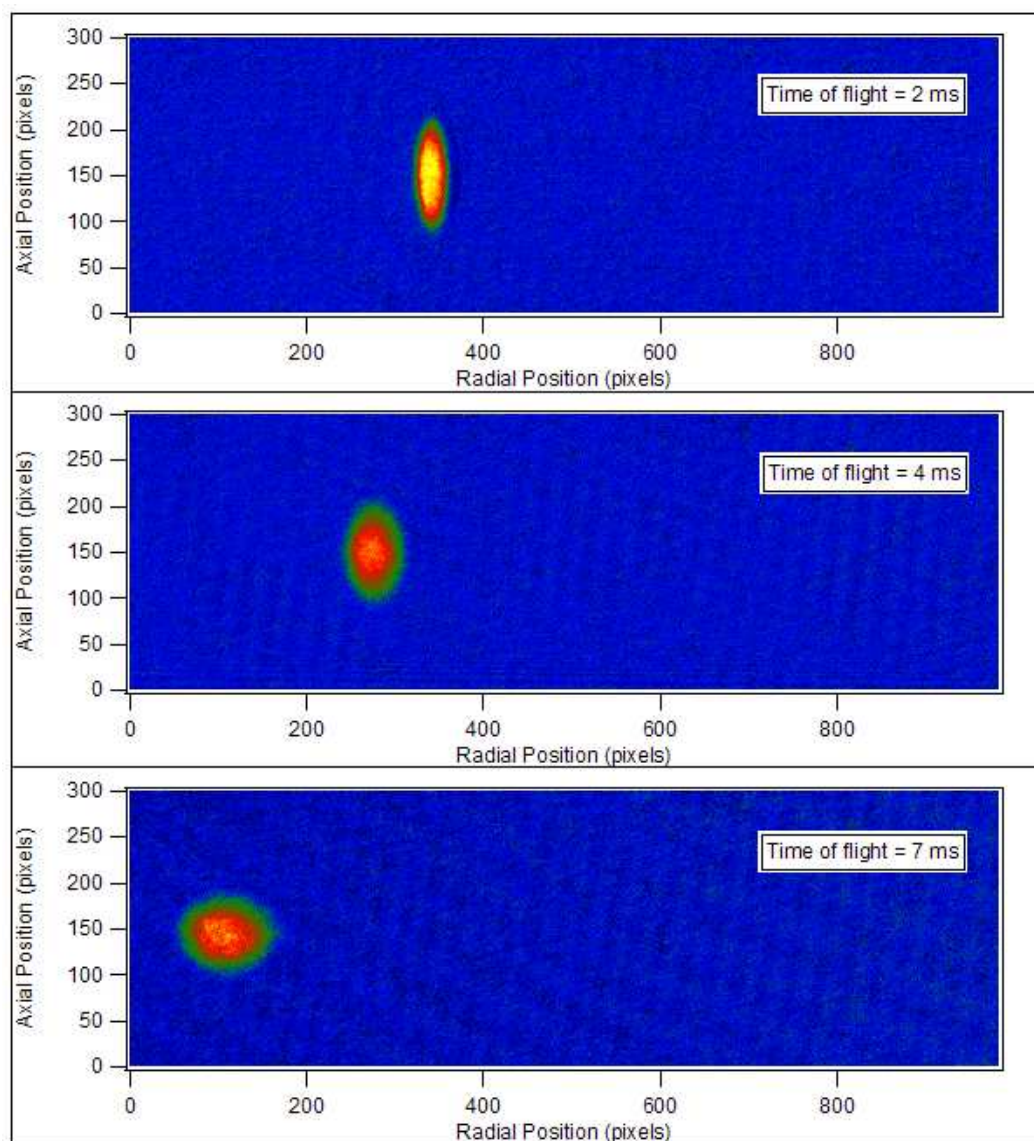


Figure 4.3: False color absorption images of the atom cloud after release from a shallow optical trap into the magnetic potential at 840 gauss. The time elapsed from release from the optical trap is indicated in the upper right hand corner of each image. The center of mass motion of the atom cloud can be analyzed to determine the magnetic potential oscillation frequency in the plane of the camera image.




Figure 4.4: The center of mass position of the atom cloud plotted versus time of flight. The dots represent the measured position of the cloud, while the dashed line is a sinusoidal fit to the data. As only a fraction of one period of the oscillation is available, there is some error in the measured magnetic potential oscillation frequency. At 840 gauss, measurements indicate a magnetic potential oscillation frequency of 19 to 24 Hz in the plane imaged by the camera.

external magnetic field is increased.

Our magnet arrangement roughly resembles two tires on an imaginary axle, where each tire represents a coil of wires and the imaginary axle coincides with the y -axis.³ The trapped atom cloud is located on the y -axis, midway between the magnets. As an atom moves in the x - z plane away from the y -axis, the magnetic field drops. Meanwhile, an atom which begins on the y -axis midway between the magnets will experience an increasing magnetic field as it moves toward one of the magnets. Since the spin states we trap are drawn to regions of high magnetic field, we conclude that the magnetic potential is attractive in the x and z directions, but repulsive in the y direction. Consequently, if we approximate the magnetic

³The common convention is to label the symmetry axis of the magnets the z -axis. However, to be consistent with our published papers, we label the symmetry axis the y -axis.

potential as harmonic, we can write

$$U_{mag} = \frac{m}{2} (\omega_{mx}^2 x^2 - \omega_{my}^2 y^2 + \omega_{mz}^2 z^2). \quad (4.27)$$

The magnetic potential energy can be written in a more general fashion using basic magnetostatic theory. If the magnetic dipole moment of the atom is aligned with the local magnetic field direction, then

$$U_{mag} = -\boldsymbol{\mu} \cdot \mathbf{B} = -\mu |\mathbf{B}|. \quad (4.28)$$

Furthermore, Maxwell's equations state that in free space $\nabla \cdot \mathbf{B} = 0$ and $\nabla \times \mathbf{B} = 0$. Consequently, $\nabla^2 \mathbf{B} = 0$. Using (4.27) and (4.28), we can write

$$\nabla^2 \mathbf{B} = \nabla^2 \left(\frac{-U_{mag}}{\mu} \right) = \omega_{mx}^2 - \omega_{my}^2 + \omega_{mz}^2 = 0. \quad (4.29)$$

Due to cylindrical symmetry, however, $\omega_{mx} = \omega_{mz}$, which implies

$$\omega_{my}^2 = 2\omega_{mx}^2. \quad (4.30)$$

In other words, we need to measure only ω_{mx} . We can then use cylindrical symmetry to fix ω_{mz} , and (4.30) provides ω_{my} .

In summary, by measuring one of the magnetic potential oscillation frequencies, we can determine all of the magnetic potential oscillation frequencies. Of course, it would be inconvenient to have to measure the magnetic potential oscillation frequencies each time we choose to work at a new magnetic field. Fortunately, by measuring the magnetic potential oscillation frequencies at one magnetic field,

we can readily calculate the trap oscillation frequencies at an arbitrary magnetic field. Consider that $\mathbf{F}_{mag} \propto |\mathbf{B}|$, and $F_{mi} \propto \omega_{mi}^2$. Knowing that the magnetic force is null at zero magnetic field, and supposing we measure the trap oscillation frequency $\omega_{mi}(B_0)$ at a magnetic field B_0 , we can write

$$\omega_{mi}^2(B) = \omega_{mi}^2(B_0) \frac{B}{B_0}. \quad (4.31)$$

Chapter 5

Image Processing

After creating a degenerate Fermi gas and subjecting it to some form of experimental manipulation, we require a method for monitoring the response of the gas. By far, the most widely used technique for extracting information from confined quantum gases involves an imaging technique. While several schemes exist for acquiring images of trapped gases, all of the data in this dissertation was acquired using on-resonance absorption imaging. In absorption imaging, a probe beam pulse which is resonant with an optical transition for one of the trapped hyperfine states passes through the atom cloud. The resulting combination of incident and scattered light is collected by a simple imaging system and directed to a CCD (charge coupled device) camera, where the spatial distribution of the light is recorded. With the information recorded by the CCD camera, we can reconstruct the spatial distribution of the atom cloud. The spatial distribution of the atom cloud, in turn, allows us to calculate such quantities as the number of trapped atoms, as well as the temperature and size of the atom cloud.

There are two basic challenges to be met when producing absorption images. The first involves generating a probe beam whose frequency is resonant with the desired optical transition. This topic is addressed in Appendix B. The second challenge involves reproducing the spatial distribution of atoms in the cloud by

examining the CCD camera’s recorded array of photon counts, a topic we cover in this chapter.

5.1 Image acquisition procedure

Acquiring an image of the atom cloud is a two-step process. Following release from the optical potential, the cloud is allowed to expand for a well defined duration, known as the time of flight. A probe beam illuminates the atom cloud, and the incident and scattered fields from the probe beam are imaged by the CCD camera. This array of photon counts recorded on the CCD camera is referred to as the “signal” shot, and is displayed in the upper panel of Figure 5.1. Here, the position of the atom cloud is revealed by the dark spot in an otherwise relatively flat background.

On-resonance absorption imaging is a destructive process. After waiting a short time for the atoms to leave the imaging area, a second pulse of light is generated, and the resulting spatial distribution of photons is recorded by the CCD camera. This process is referred to as taking the “reference” shot. The reference shot should be identical to the signal shot in every way, with the exception that there should be no atoms present while taking the reference shot. A typical reference shot is displayed in the lower panel of Figure 5.1. For a perfect imaging system that always produces identical probe pulses composed of an ideal plane wave, acquisition of a reference shot would be unnecessary. However, a glance at the lower panel of Figure 5.1 reveals a mottled spatial distribution of recorded photon counts even when the atom cloud is not present. Having information about the spatial inhomogeneity of the probe beam is necessary to reconstruct

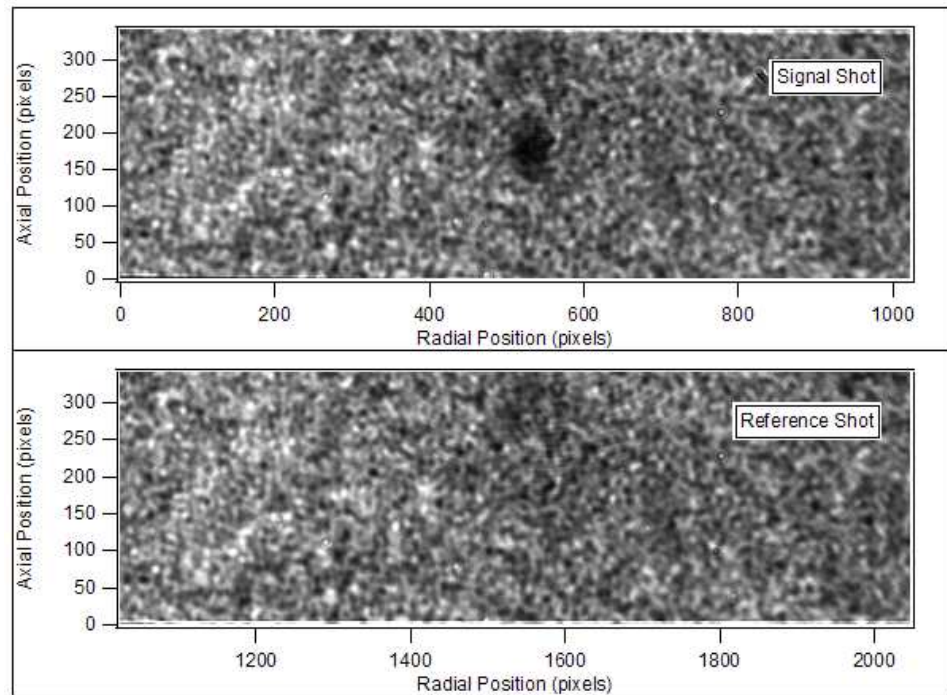


Figure 5.1: False color images showing the spatial distribution of a probe pulse of light. The upper panel, labelled the signal shot, shows the recorded spatial distribution of photons when the resonant probe pulse illuminates an atom cloud. The atom cloud casts a shadow, which appears as the dark spot in an otherwise relatively flat background. The lower panel, labelled the reference shot, shows the image of the probe beam when no atom cloud is present. Calculating an absorption image or the column density of the atom cloud requires both a signal and a reference shot.

accurately the spatial distribution of atoms in the cloud.

There are two issues that I have swept under the rug thus far. First, the quality of our absorption images is related to the time delay between acquisition of the signal and reference shots. The shorter the duration, the better the resulting images. Of course, we must wait for the atoms to disperse before taking the reference shot, so we cannot make the delay time arbitrarily short. However, we can reduce the delay time by making use of a Fast Kinetics feature on our CCD camera. When working in Fast Kinetics mode, a portion of the CCD array is shielded by a razor blade placed at the image plane in our imaging system. Roughly one-third of the CCD array is exposed. In the process of taking the signal shot, the first third of the CCD array is exposed. After the signal shot is complete, the camera is instructed to shift the camera pixels to expose a fresh third of the CCD array. The array of pixels containing the signal shot is now protected by the razor blade, and the reference shot is captured on the newly exposed pixels. With both the signal and reference shots acquired, the camera can then begin the process of reading out the photon counts in the CCD array. The time savings that results from using the Fast Kinetics mode is substantial. Our camera requires $16 \mu\text{s}$ to shift each row of pixels. As we have roughly 340 rows of pixels exposed when working in Fast Kinetics mode, we must wait roughly 6 ms between taking the signal and reference shots.¹ In comparison, if we were to use the entire CCD array when acquiring both the signal and reference shots, we would have to wait for the camera to shift each row *and read out the photon counts in each pixel*. The fastest time to process a pixel is $1 \mu\text{s}$ for our camera

¹In practice, we wait longer than 6 ms to accommodate other necessary delays in the timing system.

model, and our camera has a megapixel array. As a result, the minimum delay time between the signal and reference shots would be in excess of 1 second. As the quality of our images degrades with increasing delays between acquisition of the signal and reference shots, Fast Kinetics has become the default imaging mode for our lab.

The second issue that I have glossed over is the need to take a “background” image. As with any other photon counting device, our CCD array is susceptible to dark counts, in which charge buildup in the CCD array results in apparent photon counts even with the camera shutter closed. This effect can be mitigated by reducing the temperature of the array using thermoelectric coolers. Typically, we cool our CCD array to $-40\text{ }^{\circ}\text{C}$, resulting in a substantial reduction in dark counts compared to room temperature conditions. Even with this reduction, the number of dark counts (≈ 500 per camera pixel) remains too large to be considered negligible. To account for this, we acquire a “background” image in which no atoms or probe pulse is present and the shutter to the camera is closed. Essentially, the background image is a spatial distribution of apparent “photon” counts when no light is reaching the CCD array. This spatial distribution tends to be very flat and varies minimally throughout the course of an experimental run. As a result, only one background is acquired each day, and all of the succeeding signal and reference shots have this array of background counts subtracted from them.

5.2 Image Analysis

Even with all the effort that goes into producing optimum conditions for acquisition of the signal and reference shots, the resulting images still suffer from a number of imperfections. Much of the image analysis procedure is designed to reduce the impact of these imperfections. To assess the performance of our experimental system, it is convenient to view the images in real time. We use Andor MCD software to view the images acquired by the camera, and the preferred method of displaying the data in real time is as an absorption image. Image processing involves the manipulation of three arrays of information: the signal shot, the reference shot, and the background. We let the number of background counts in the pixel located at (x, z) be represented by $I_B(x, z)$. Then the total counts in the signal (I_{sig}) and reference (I_{ref}) shots in a pixel at position (x, z) are given by

$$I_{sig}(x, z) = I_s(x, z) + I_B(x, z) \quad (5.1)$$

$$I_{ref}(x, z) = I_r(x, z) + I_B(x, z), \quad (5.2)$$

where $I_s(x, z)$ and $I_r(x, z)$ represent the counts not due to background dark counts. $I_{sig}(x, z)$, $I_{ref}(x, z)$ and $I_B(x, z)$ are measured quantities from which we calculate $I_s(x, z)$ and $I_r(x, z)$. The absorption image can be calculated using

$$abs(x, z) = 100 \left[1 - \frac{I_s(x, z)}{I_r(x, z)} \right], \quad (5.3)$$

where $abs(x, z)$ gives the percentage absorption in a pixel located at position (x, z) . If we apply (5.3) to the signal and reference shots displayed in Figure 5.1,




Figure 5.2: False color image showing the calculated absorption based on the signal and reference shots displayed in Figure 5.1. The light pill-shaped object is the atom cloud.

we produce the false-color absorption image shown in Figure 5.2.

Even when working in Fast Kinetics mode, where the time elapsed between acquisition of the signal and reference shots is intentionally small, we still observe fluctuations in probe beam power between the two shots. Of course, the entire purpose of the reference shot is to obtain information about what the array of photon counts looks like in the absence of an atom cloud. Any other differences between the signal and reference shots will introduce errors when calculating the absorption image or column density. As a crude first step toward overcoming this problem, we make use of a “balance” region in the signal and reference shots. Figure 5.3 shows the absorption image displayed in Figure 5.2 along with two rectangles. The rectangle at the left of the image shows the outline of the balance region of interest (ROI). The rectangle that frames the atom cloud shows the “cloud” region of interest, which is necessary for more sophisticated background subtraction algorithms which will be discussed in a moment. Note that the balance ROI is far from the atom cloud, in a region where the absorption should be zero. In each of the signal and reference shots, the background corrected number

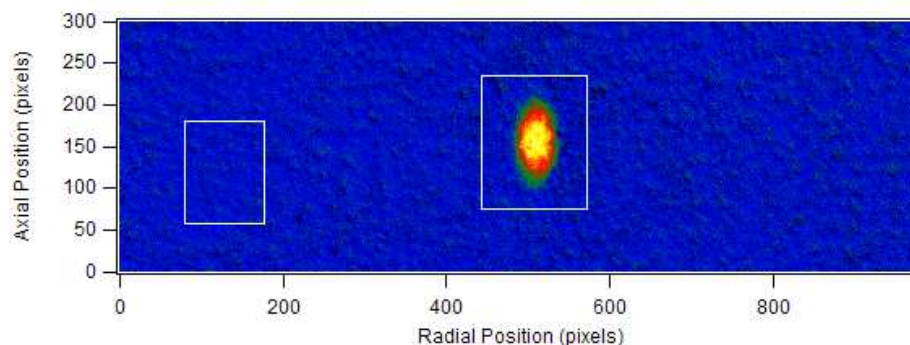


Figure 5.3: False color absorption image showing two regions of interest (ROI). The rectangle at the left shows the outline of the balance ROI, a portion of the image used to correct for probe beam power fluctuations between the acquisition of the signal and reference shots. The rectangle that frames the atom cloud shows the cloud ROI, which is used for more sophisticated background subtraction techniques discussed in the text.

of counts in the balance region is tallied. The ratio of the counts in the signal to reference shots should be 1, but if it differs, the number of counts in *every pixel* of the reference shot is scaled until there is an equal number of counts in the balance ROI in the signal and reference shots.

While the use of the balance region does a reasonable job of subtracting any residual background from the image, the analysis program we use for processing the signal and reference shots (Igor Pro v4.0) has a more sophisticated method for subtracting residual backgrounds. To make use of Igor Pro’s background subtraction algorithms, we need to specify the location of the atom cloud, which is accomplished using the previously mentioned cloud ROI. Igor then fits a plane to the absorption image *outside the cloud region of interest*, and subtracts this residual background.²

²In reality, we perform this background subtraction procedure on the calculated column density (see Section 5.3) rather than the absorption image. The principle is precisely the same in both cases, however, and the discussion of this step was included here to avoid a digression

While the use of absorption images is convenient for tracking the performance of the system during an experimental run, the extraction of meaningful physical quantities from our images often requires the conversion of these absorption images to column densities. That is, we want to convert the calculated absorption for a particular pixel into the number of atoms represented by that level of absorption. The calculation of column densities is discussed in Section 5.3.

5.3 Column density for an ideal two-level system

In this section, we consider the absorption of a probe beam whose frequency is close to resonance in a two-level closed atomic system. This type of calculation arises in many applications in quantum optics, and has been presented in sources too numerous to reference here. Examples of recent treatments can be found in [24, 103, 104].

We consider a probe beam propagating along the y -axis and illuminating an atom cloud. As the beam passes through the cloud, the intensity of the beam will diminish. We can express the attenuation of the beam as

$$\frac{dI}{dy} = -I \alpha(I) n(x, y, z), \quad (5.4)$$

where I is the intensity of the beam, $n(x, y, z)$ is the spatial density of the atom cloud, and

$$\alpha(I) = \frac{\sigma_R}{1 + I/I_{sat} + \delta^2}. \quad (5.5)$$

when discussing the column density calculation.

In (5.5), σ_R is the resonant optical cross section for a two-level quantum transition of wavelength λ

$$\sigma_R = \frac{3 \lambda^2}{2 \pi}, \quad (5.6)$$

and δ is the detuning in half-linewidths of the probe beam from resonance

$$\delta = \frac{\omega - \omega_0}{\Gamma_s/2}. \quad (5.7)$$

In the preceding expressions, Γ_s (I_{sat}) is the linewidth (saturation intensity) of the transition, ω is the frequency of the incident beam, and ω_0 is the transition frequency of the two-level system.

If we plug (5.5) into (5.4), separate variables and integrate, we obtain

$$\int_{I_0}^{I_{y'}} dI \frac{1 + I/I_{sat} + \delta^2}{I} = -\sigma_R \int_{-\infty}^{y'} n(x, y, z) dy, \quad (5.8)$$

where $I_{y'}$ indicates the beam intensity at position y' . While this formalism allows us to calculate the intensity of the beam at some arbitrary position inside the atom cloud, in practice we are concerned with the intensity of the beam after it passes through the atom cloud. If we let $y' \rightarrow \infty$ in the integral on the right hand side of (5.8), we obtain the definition of the column density

$$n(x, z) = \int_{-\infty}^{\infty} n(x, y, z) dy. \quad (5.9)$$

Using this definition and performing the integral on the left hand side of (5.8), the column density can be written

$$n(x, z) = -\frac{1}{\sigma_R} \left\{ (1 + \delta^2) \ln \left[\frac{I(x, z)}{I_0(x, z)} \right] + \frac{I(x, z) - I_0(x, z)}{I_{sat}} \right\}. \quad (5.10)$$




Figure 5.4: False color image showing the column density calculated from the data in the signal and reference shots displayed in Figure 5.1. The light region shows the position of the atom cloud.

Note that in (5.10), $I_0(x, z)$ is the intensity of the probe beam at position (x, z) before the probe beam reaches the atom cloud. $I(x, z)$ is the intensity of the probe beam after it has passed through the atom cloud. A typical two-dimensional column density is displayed in Figure 5.4.

In closing we note that (5.10) can be simplified if the incident probe beam is weak, $I_0(x, z), I(x, z) \ll I_{sat}$, in which case the second term in braces is approximately zero. Then the expression for the column density simplifies to

$$n(x, z) = -\frac{1 + \delta^2}{\sigma_R} \ln \left[\frac{I(x, z)}{I_0(x, z)} \right]. \quad (5.11)$$

This formula for the column density has its merits, as it eliminates the need to obtain an absolute calibration of the imaging system. Consider that I_{sat} is a calculated number, whereas $I_0(x, z)$ and $I(x, z)$ are measured quantities. To properly compare the ratio of these quantities to calculate the second term in braces in (5.10), one needs to consider the following issues:

- What is the efficiency of the CCD camera?

- Does the camera employ some sort of gain mechanism? (The CCD camera in use in our lab has gain characteristics which are dependent on the readout time per pixel on the CCD array.)
- Do optical elements along the probe beam path significantly attenuate the beam after it has illuminated the atom cloud? Or, does the probe beam intensity detected at the CCD array provide a good estimate of probe beam intensity when it illuminates the atom cloud?

Essentially, we must ask the question: if a probe beam has intensity I_{sat} at the location of the atoms, what intensity will be detected by the camera?

Obviously, it would be convenient to disregard these issues and work with a weak probe beam pulse [$I_0(x, z), I(x, z) \ll I_{sat}$]. However, we also desire a short probe beam pulse, as long probe times would result in “smearing” of the atom cloud. Recall that for time of flight imaging, the atom cloud is expanding as it is being illuminated. We want a very short imaging pulse to take a true snapshot of the atoms, rather than a long imaging pulse which provides a time-averaged image of the expanding atom cloud. Unfortunately, a short, weak imaging pulse can result in a low signal-to-noise ratio, so we use a shorter, stronger imaging pulse along with (5.10) to calculate our column densities. In doing so, we can obtain good signal to noise ratios with short probe beam pulses. Furthermore, a brief imaging pulse should minimize optical pumping to dark states, a phenomenon discussed in Section 5.4. The cost of using this formula is that we must consider all of the questions in the bulleted list above. That is, rather than using the calculated value of I_{sat} for the optical transition we use for imaging, we use an effective value of I_{sat} which takes into account the imperfect performance of our

imaging system. These issues are considered in Section 5.7.

Before considering the calculation of I_{sat} for our imaging system, there are other concerns we must address. At the beginning of this section, we assumed that we had an ideal two-level quantum system. As we will see, for typical imaging conditions used in our laboratory, this is not precisely true. The number of interactions between atoms in the cloud and incident photons will be reduced if the photons have the “wrong” polarization or if the atoms are in the “wrong” quantum state. These scenarios are considered in Sections 5.4 and 5.5.

5.4 Optical Pumping

Early in the discussion of measuring the column density, we assumed that we had a perfect, two-level quantum system. Although lithium is a relatively simple, hydrogenic atom, there are many energy levels to consider, and depending on the choice of optical transition, the assumption of an ideal two-level system might be invalid. Hence, during a probe pulse, an atom might end up occupying a quantum state that is not resonant with the probe beam. This atom is said to be in a “dark” state. As this atom is present in the atom cloud but not detectable by the probe beam, the phenomenon of optical pumping into dark states will lead to an underestimate of the number of atoms present in the cloud. Since the number of atoms in the cloud sets the characteristic energy and temperature scales of the physical system, it is worthwhile to estimate the error in the detected number which results from assuming a perfect two-level system.

In Section 2.4 we found that the two lowest hyperfine ground states of the ${}^6\text{Li}$

atom are given by

$$|1\rangle = \sin \Theta_+ |1/2 0\rangle - \cos \Theta_+ |-1/2 1\rangle \quad (5.12)$$

$$|2\rangle = \sin \Theta_- |1/2 -1\rangle - \cos \Theta_- |-1/2 0\rangle, \quad (5.13)$$

where the kets $|m_S m_I\rangle$ are given in the basis of the electron and nuclear magnetic spin projections, and the coefficients in front of the kets are magnetic field dependent. In the ground state, the electronic orbital angular momentum is zero, so if we let $J = L + S$ be the total electronic angular momentum, where L is the orbital angular momentum, and $S = 1/2$ is the intrinsic spin of the electron, we find that the hyperfine ground states of ${}^6\text{Li}$ are in the $J = 1/2$ manifold. Further, as the magnetic field is increased, we find that

$$|1\rangle \rightarrow -|-1/2 1\rangle \quad (5.14)$$

$$|2\rangle \rightarrow -|-1/2 0\rangle \quad (5.15)$$

with increasing magnetic field. That is, at “high” magnetic field ($B > 500$ gauss), the $|1\rangle$ and $|2\rangle$ states are essentially in the same electronic state and differ only by nuclear magnetic spin projection. This is fortuitous, as the electronic dipole transition we use to image the atoms depends on the electronic state of the atom only. Since the two hyperfine states of interest are essentially in the same electronic state, we should be able to image both states equally well. In this high field regime, we see that the electronic spin projection is “down” with respect to the magnetic quantization axis, which is coaxial with the probe beam propagation axis. We represent the electric dipole operator by $T(1, q)$, where $q = -1, 0, +1$

for left-circular, linear, and right-circular polarized light, respectively. The first argument of $T(1, q)$ reminds us that the only allowed transitions involve a change of one unit in the orbital angular momentum. That is, if we consider transitions between states enumerated in the total product basis $|L m_L\rangle |S m_S\rangle |I m_I\rangle$, we find

$$\begin{aligned} \langle I' m'_I | \langle S' m'_S | \langle L' m'_L | T(1, q) | L m_L \rangle | S m_S \rangle | I m_I \rangle \\ = \delta_{I', I} \delta_{m'_I, m_I} \delta_{S', S} \delta_{m'_S, m_S} \delta_{L', L+1} \delta_{m'_L, m_L+q}. \end{aligned} \quad (5.16)$$

Equation (5.16) suggests a way to minimize optical pumping to dark states during the imaging process. If we are starting from the $|L = 0 m_L = 0\rangle |S = 1/2 m_S = -1/2\rangle$ state, and if we image using left-circularly polarized light [corresponding to $T(1, -1)$], the only electronic excited states which can be accessed are the $|L = 1 m_L = -1\rangle |S = 1/2 m_S = -1/2\rangle |I = 1 m_I = 0, \pm 1\rangle$ states.

There are now two separate issues to consider. First, as the energy levels tune differently as a function of magnetic field, will the selection of one imaging frequency cause some of these transitions to be far off resonance? Second, for the transitions which are resonant, how strong is the transition probability between the ground and excited states? First, we consider the issue of detuning. At high magnetic field, the $|1\rangle$ and $|2\rangle$ states are separated by nearly 75 MHz, which is much larger than the natural linewidth (FWHM = 5.9 MHz) of the ${}^6\text{Li } D_2$ line [arising from the $2^2S_{1/2} \rightarrow 2^2P_{3/2}$ transition (see [24], for example)]. As such the $|1\rangle$ and $|2\rangle$ states are clearly resolved and we need not worry about exciting the $|2\rangle$ to excited state transition when exciting the $|1\rangle$ to excited state transition, and

vice versa. What about the excited states? Due to the small hyperfine splitting in the $J = 3/2$ manifold of the $2P$ level in ${}^6\text{Li}$, we can safely ignore hyperfine structure for all but the lowest magnetic fields. At the magnetic fields of interest ($B > 500$ gauss), we assume that the magnetic field dependence of the energy levels is well-described by the Zeeman shift alone. Here the energy shift is (see [24] for a more thorough discussion)

$$\Delta E = \mu_B (g_J^{ex} m_J + g_I m_I) B, \quad (5.17)$$

where μ_B is the Bohr magneton. Note that the desired excited electronic state

$$|L = 1 \ m_L = -1\rangle |S = 1/2 \ m_S = -1/2\rangle = |J = 3/2 \ m_J = -3/2\rangle, \quad (5.18)$$

so the magnetic field tuning of the desired excited state is given by

$$E_{ex}(B) = \mu_B \left(-\frac{3}{2} g_J^{ex} + g_I m_I \right) B, \quad (5.19)$$

where $g_I = -0.0004476540$ and $g_J^{ex} = 1.335$ for the states of interest [105]. We can see now that the separation of the energy levels due to the nuclear magnetic spin projection will be quite small. For instance, at $B = 840$ G, the separation between the $m_I = -1$ and the $m_I = +1$ states will be roughly 1 MHz, which is smaller than the natural linewidth of the D_2 transition. As a result, the three excited states associated with the $|J = 3/2 \ m_J = -3/2\rangle$ level are unresolvable. However, since $g_J^{ex} \gg g_I$, the $|J = 3/2 \ m_J = -3/2\rangle$ and $|J = 3/2 \ m_J = -1/2\rangle$ are well-separated, and there is negligible probability of exciting atoms to the $|J = 3/2 \ m_J = -1/2\rangle$ level when the dye laser frequency is tuned to promote

atoms to the $|J = 3/2 \ m_J = -3/2\rangle$ level. To quickly summarize, the two hyperfine ground states ($|1\rangle$ and $|2\rangle$) are well-resolved, but the three nuclear states in the $|J = 3/2 \ m_J = -3/2\rangle$ excited state manifold are not resolvable if we consider only the frequency of the probe beam.

To this point, we have established what quantum energy levels are “in play” based on issues related to the probe beam frequency. We now consider which electric dipole transitions can be excited by the probe beam. For convenience, we will label the excited states of interest in the total product basis as

$$|A\rangle = |L = 1 \ m_L = -1\rangle |S = 1/2 \ m_S = -1/2\rangle |I = 1 \ m_I = -1\rangle \quad (5.20)$$

$$|B\rangle = |L = 1 \ m_L = -1\rangle |S = 1/2 \ m_S = -1/2\rangle |I = 1 \ m_I = 0\rangle \quad (5.21)$$

$$|C\rangle = |L = 1 \ m_L = -1\rangle |S = 1/2 \ m_S = -1/2\rangle |I = 1 \ m_I = +1\rangle. \quad (5.22)$$

If we image using left-circularly polarized light, then $T(1, q) \rightarrow T(1, -1)$, and (5.16) gives the following transition probabilities from the ground to excited states:

$$|\langle C | T(1, -1) | 1 \rangle|^2 = \cos^2 \Theta_+ \quad (5.23)$$

$$|\langle B | T(1, -1) | 2 \rangle|^2 = \cos^2 \Theta_-, \quad (5.24)$$

with all other electric dipole transition probabilities between the $|1\rangle$, $|2\rangle$, $|A\rangle$, $|B\rangle$, and $|C\rangle$ states being zero. For a perfect, two-level system, the coefficients in (5.23) and (5.24) would be 1, so our task now is to consider how the quantities $\cos^2 \Theta_+$ and $\cos^2 \Theta_-$ vary with magnetic field.

In Section 2.4, we provided the analytic forms for $\cos \Theta_-$ and $\cos \Theta_+$. Squaring




Figure 5.5: The above displays the transition probabilities for on-resonance absorption imaging. $\cos^2 \Theta_-$ (solid line) and $\cos^2 \Theta_+$ (dashed line) give the transition probabilities for atoms in the $|2\rangle$ and $|1\rangle$ states, respectively, to the unresolved excited state. For magnetic fields above 450 gauss, these probabilities are in excess of 99%, and the imaging transition will act approximately like a closed, two-level system.

these quantities, we have the desired transition probabilities expressed in (5.24) and (5.23), respectively. These quantities are plotted as a function of magnetic field in Figure 5.5. We find that the transition probabilities between the ground and excited states used for imaging are in excess of 99% for magnetic fields above 450 gauss, where most of our experimental work is conducted. Since these transition probabilities are close to unity, the imaging transition will act approximately like a closed, two-level quantum system.

Although our imaging transition will behave approximately like an ideal two-level system, it is worthwhile to obtain an estimate of the error in the detected number of atoms resulting from optical pumping to dark states. We consider imaging at 840 gauss with a probe beam duration of $5 \mu\text{s}$ and intensity $0.3 I_{\text{sat}}$. The photon scattering rate for a two-level system interacting with a resonant

probe beam is

$$\Gamma_{sc} = \frac{1}{2\tau_{sp}} \frac{I/I_{sat}}{1 + I/I_{sat}}, \quad (5.25)$$

where $\tau_{sp} = 27.1$ ns is the excited state lifetime for the transition of interest in ${}^6\text{Li}$. As the probe beam propagates through the atom cloud, the intensity of the beam will drop. Hence, plugging $I/I_{sat} = 0.3$ into (5.25) will provide an upper limit for the average photon scattering rate for atoms in the cloud. For $I/I_{sat} = 0.3$, (5.25) reveals that the photon scattering rate is 4.26 million photons per second, or roughly 21 photons during a $5 \mu\text{s}$ pulse. Suppose we focus on a single atom as it absorbs and emits photons during the $5 \mu\text{s}$ probe pulse. For a perfect two-level system, on average the atom will absorb and emit 21 photons during the probe pulse. For a system with optical pumping to a dark state, the number of photons absorbed and emitted will be smaller. We can obtain an estimate of this reduction in the number of scattered photons if we consider the departure of our imaging transition probabilities from unity. At 840 gauss the optical transition probabilities between the $|1\rangle$ and $|2\rangle$ ground states and the excited state are nearly equal,

$$\cos^2 \Theta_+ \approx \cos^2 \Theta_- \approx 0.998. \quad (5.26)$$

That is, an atom in the $|1\rangle$ state has 99.8% probability of absorbing a photon when the probe pulse is first turned on. Likewise, it has a $99.8\% \times 99.8\%$ probability of returning to the $|1\rangle$ state after re-emitting the photon. Consequently, only $99.8\% \times 99.8\%$ of the atoms will be in a quantum state which is resonant with the probe beam after the first absorption-emission cycle. The probability of absorbing a second photon is $99.8\% \times 99.8\% \times 99.8\% = 0.998^3$, and so on. If we follow this

process over 21 optical cycles, we find that, on average, the atom will absorb

$$\sum_{n=1}^{21} 0.998^{2n-1} = \frac{0.998(1 - 0.998^{42})}{1 - 0.998^2} = 20.14 \quad (5.27)$$

photons during the 5 μ s pulse. Of course, it is unphysical to discuss fractional photons, but we are considering average quantities here. Equation (5.27) reveals that an atom in our system will scatter, on average, 20.14 photons for every 21 photons that would be scattered in a perfect two-level system. This means that $\kappa = 20.14/21 = 0.959$ times as many photons will be absorbed by the atom cloud, which will result in a larger number of photons reaching the CCD camera.

This increase in the number of photons reaching the camera will result in an apparent reduction in the number of atoms in the cloud. Consider that I_0 in (5.10) represents the probe beam intensity prior to reaching the atom cloud, while I is the probe beam intensity after exiting the atom cloud. The ratio I/I_0 provides the percentage of unabsorbed light, while $1 - I/I_0$ gives the percentage of absorbed light. For a system with optical pumping to a dark state, the amount of unabsorbed light will increase by a factor proportional to $1 - \kappa$. If we let

$$\frac{I}{I_0} \rightarrow \frac{I}{I_0} + (1 - \kappa) \left(1 - \frac{I}{I_0}\right) \quad (5.28)$$

in (5.10), and compare the resulting expression for the column density to what we would calculate for an ideal system, we obtain

$$N_{err}(I, \kappa) = \frac{\ln\{[I + (1 - \kappa)(I_0 - I)]/I_0\} + [I + (1 - \kappa)(I_0 - I) - I_0]/I_{sat}}{\ln(I/I_0) + (I - I_0)/I_{sat}}. \quad (5.29)$$

Equation (5.29) gives the fraction of atoms detected in our nearly two-level system

compared to a perfect two-level system. A plot of $N_{err}(I, \kappa)$ for $\kappa = 0.959$ and $I_0/I_{sat} = 0.3$ is given in Figure 5.6. We note that the percentage error becomes quite large as $I \rightarrow 0$. However, this limit corresponds to the case of complete absorption of the probe beam, a scenario we avoid when imaging. For typical imaging conditions, we do not exceed 80% absorption (or, $I/I_0 = 0.20$). If we consider an unweighted average of $N_{err}(I, \kappa)$ for absorption levels from 0% to 80%, we obtain

$$\frac{1}{I_0 - 0.20 I_0} \int_{0.20 I_0}^{I_0} N_{err}(I, 0.959) dI = 0.95, \quad (5.30)$$

for $I_0 = 0.3 I_{sat}$. Since the value of $N_{err}(I, \kappa)$ is relatively flat for absorptions between 0% and 80%, we expect (5.30) to provide a reasonable estimate of the error in the measured atom number.³ Consequently, at 840 gauss, if we image using a probe pulse with intensity $I_0 = 0.3 I_{sat}$ and duration $5 \mu\text{s}$, we anticipate that our measured atom number will be $\approx 5\%$ lower than the true value. Since the Fermi energy $\varepsilon_F \propto N^{1/3}$, a 5% underestimate of the atom number will lead to a 2% underestimate of the Fermi energy. Likewise, since the Fermi radii $\sigma_i \propto N^{1/6}$, a 5% underestimate of the atom number yields a 1% underestimate of the Fermi radii.

The preceding formalism should provide an upper limit on the magnitude of the error in the detected number of atoms. Recall that at the beginning of the

³We can calculate a weighted average of the number error by considering theoretically generated density profiles for harmonically trapped Fermi gases. The expression for the spatial density of the clouds is given by (A.47). This atomic density can then be related to the ratio I/I_0 using (5.10) for a chosen value of I_{sat} . Next, we use (5.29) to estimate the error in the detected atom number for that value of I/I_0 . Integrating the errors for all density levels provides a weighted average of the error in the atom number. We find that for typical imaging conditions, the weighted average of the number error indicates that we underestimate the atom number by 5% or 6%, in agreement with (5.30).




Figure 5.6: Estimate of the error in detected atom number resulting from optical pumping to dark states during probe pulse. The percentage of atoms detected N_{err} given by (5.29) (solid line) varies as a function of the ratio of I/I_0 . The dashed line is $N_{err} = 1$. The above plot was generated for $I_0 = 0.3 I_{sat}$ and $\kappa = 0.959$.

estimate, we assumed that all of the atoms were interacting with a probe beam with intensity $0.3 I_{sat}$. This is not true, as the probe beam intensity weakens as it propagates through the atom cloud, resulting in a lower photon scattering rate for many of the atoms in the cloud. A lower photon scattering rate results in fewer chances to transition to a dark state, and a smaller error in the number of detected atoms. Consequently, for typical imaging conditions (a $5 \mu s$ probe pulse with intensity $0.3 I_{sat}$ at 840 gauss), the 5% underestimate of the atom number is an upper limit on the magnitude of the error in the detected atom number.

For fields above 840 gauss, we expect the error in the detected atom number to diminish further, as the imaging transition asymptotically tends toward a closed, two-level transition. Many cycles of the experimental apparatus indicate that the statistical fluctuations in the number of atoms is roughly 10%. Since the error in

atom number arising from optical pumping to dark states is smaller than the 10% statistical fluctuations in the atom number, we neglect optical pumping when reporting the number of atoms in our cloud. Moreover, the error bars on the reported atom number indicate statistical fluctuations only.

In this section, we assumed that the incident photons were in the correct state (all have left-circular polarization), and considered what portion of atoms were in the “wrong” quantum state. However, interactions between atoms and photons will be reduced also if the atoms are in the correct quantum state, but the incident photons have the incorrect polarization. This latter scenario is considered in Section 5.5.

5.5 Correcting for probe beam depolarization

Thus far, we have assumed that the probe beam is composed entirely of photons capable of interacting with the atom cloud. In this case, complete absorption of the probe beam by the atom cloud at position (x_0, z_0) yields $I_s(x_0, z_0) = 0$. However, suppose that some small portion of the probe beam cannot interact with the atoms in the cloud because the photons have the incorrect polarization. In this case, some portion of the probe beam will still propagate to and be detected by the CCD array despite the perfect absorption of the portion of the probe pulse that can interact with the atom cloud. To obtain the most accurate reconstruction of the column density, we need to account for this effect.

In (5.10), we stated that $I_0(x, z)$ is the intensity of the probe beam before the atom cloud and $I(x, z)$ is the intensity after passing through the cloud. For a perfectly polarized probe beam, we note that $I_0(x, z)$ would correspond to the

“reference shot” quantity $I_r(x, z)$ in (5.2). Similarly, $I(x, z)$ would be equivalent to the “signal shot” quantity $I_s(x, z)$ in (5.1).

We now modify these definitions slightly to account for imperfect probe beam polarization. Suppose $I_0(x, z)$ and $I(x, z)$ represent the *portion of the probe beam which is absorbable by the atom cloud*. In this case, we can relate $I_0(x, z)$ and $I(x, z)$ to the quantities $I_s(x, z)$ and $I_r(x, z)$ in (5.2) and (5.1) by the following:

$$I_0(x, z) = \phi I_r(x, z) \quad (5.31)$$

$$I(x, z) = I_s(x, z) - (1 - \phi) I_r(x, z). \quad (5.32)$$

In (5.31) and (5.32), ϕ represents the percentage of the input probe beam which can be absorbed by the atoms. For $\phi < 1$, (5.31) states that not all of the incident beam can be absorbed. Similarly, (5.32) indicates that for perfect absorption of the absorbable part of the probe beam, the probe beam will not be completely attenuated, as the portion of the beam with the wrong polarization will pass through the cloud. If we substitute (5.31) and (5.32) into (5.10), we obtain

$$n(x, z) = -\frac{1}{\sigma_R} \left\{ (1 + \delta^2) \ln \left[\frac{I_s(x, z) + (\phi - 1) I_r(x, z)}{\phi I_r(x, z)} \right] + \frac{I_s(x, z) - I_r(x, z)}{I_{sat}} \right\}. \quad (5.33)$$

We now have an expression for the column density expressed in terms we can readily extract from the signal and reference shots. Before applying (5.33), however, we need to determine the value of ϕ , a topic we address in Section 5.6.

5.6 Measuring ϕ

The column density for clouds confined in the optical trap will be highest at the center of the trap, and the spatial variation of the trapping potential should provide a spatially varying column density. However, we find that for very cold atom clouds that are imaged after short expansion times following release from the optical potential, absorption images of the cloud reveal a “flat-top” region near the center of the cloud. Here, while the column density of the atom cloud still varies in space, the apparent absorption is saturated, as all of the photons which can be absorbed, are absorbed. For a probe beam composed entirely of photons with the correct polarization to interact with a perfect two-level atom, the absorption should be 100% (or, $\phi = 1$) in this “flat-top” region. However, in practice, this number is slightly less than 1, and we measure $\phi = 0.97$ for our system.

We measure ϕ by producing a highly degenerate atom cloud ($T/T_F < 0.10$) in the strongly interacting regime, where the process of evaporative cooling is most efficient. We allow the atom cloud to expand from the optical trap for a short time before illuminating it with a resonant probe pulse. After acquiring a reference shot (see Section 5.3), we can produce an absorption image. The measured value of ϕ is the average percentage absorption for each pixel in the “flat-top” region of the absorption image. A sample absorption image with a “flat-top” region is displayed in Figure 5.7.

With the value of ϕ in hand, we are almost ready to apply (5.33). Before doing so, however, we need to consider how imaging system imperfections influence the value of the saturation intensity we should use.

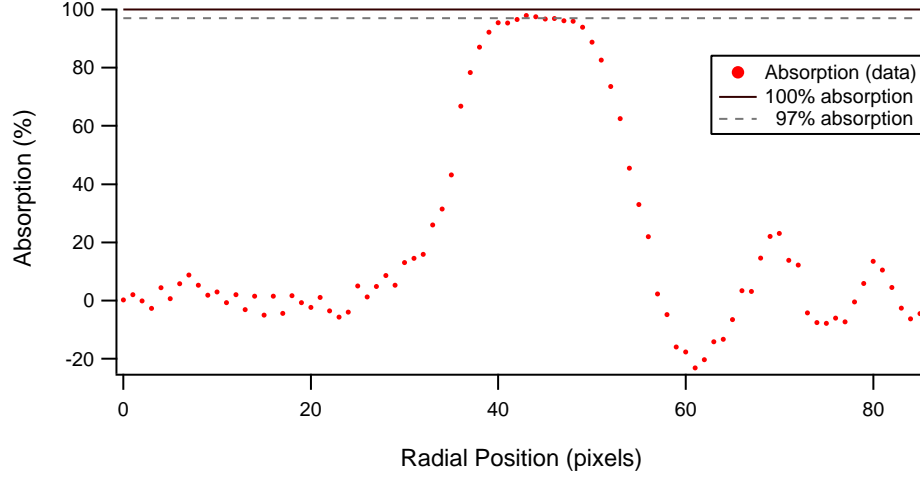


Figure 5.7: At short expansion times for cold atom samples, the spatial variation of the absorption profile (shown as dots) acquires a “flat-top.” In the distribution above, the maximum absorption is roughly 97%, indicated by the dashed line. This would correspond to $\phi = 0.97$, as described in the text. The solid line shows 100% absorption. The ripples to the right of the spatial profile are an imaging artifact arising from diffraction and fluorescence from the atom cloud.

5.7 Determining the saturation intensity I_{sat}

The saturation intensity for the D_2 line of ${}^6\text{Li}$, the transition we use for imaging, is $I_{sat} = 2.54 \text{ mW/cm}^2$ [24]. We wish to convert this intensity into the number of photons per pixel for a probe pulse of duration t_{probe} . The amount of energy E_{probe} striking a single pixel of area A_{pix} in a probe pulse with duration t_{probe} is

$$E_{probe} = I_{sat} A_{pix} t_{probe}. \quad (5.34)$$

The energy associated with a photon of wavelength λ is

$$E_{photon} = \frac{hc}{\lambda}, \quad (5.35)$$

where h is Planck's constant and c is the speed of light. Then the total number of photons striking a single pixel during the probe pulse is given by E_{probe}/E_{photon} . Measurements of the imaging system magnification (discussed in Section 5.9) indicate that $A_{pix} = 6.2 \times 10^{-12} m^2$ for our system. A typical probe pulse for the data presented in this dissertation is $t_{probe} = 5 \mu s$. For these values, an imaging pulse at the saturation intensity of the D_2 line, where $\lambda = 670.977 nm$, would yield 2660 photons per pixel.

To this point, we have assumed that our imaging system and camera are perfectly efficient. That is, we have assumed that there is no further attenuation of the probe beam after it passes through the atom cloud. In addition, we have assumed that the camera detects every photon that strikes the CCD array. In practice, of course, this is not true. It is difficult to estimate the attenuation of the probe beam between the location of the atom cloud inside the vacuum chamber and the probe beam's exit port, as we cannot place a power meter inside the vacuum chamber to measure the probe beam power at the location of the atom cloud. However, we know that the probe beam's exit port contains an anti-reflection coated window, which should minimize attenuation of the probe beam. Consequently, we make the assumption that a measurement of the probe power at the probe beam's exit port provides a reasonable estimate of the beam power when it strikes the atom cloud.

While it is difficult to estimate the decrease in probe beam power between the location of the atom cloud and the probe beam's exit port on the vacuum chamber, it is straightforward to measure the CCD camera's efficiency. That is, we want to know what percentage of photons that strike the camera are detected by the camera. To accomplish this task, we place a mask in the image plane of

the first lens in the probe beam path. The mask should be small enough that all of its edges can be seen on the CCD camera. The proper distance between the camera and the mask is determined by the sharpness of the mask on the camera. If the mask is constructed from razor blades and folded pieces of aluminum foil, its distance can be set precisely by moving the mask until no diffraction ripples appear at the edges of the mask. The probe beam can then be turned on, and a power meter placed between the mask and the camera. This measured power, which we take to be the “true” beam power, can be converted to the number of photons per second using (5.35). Next, atom-free camera images can be acquired using probe pulses of various durations (such as 5 μs , 10 μs , and 20 μs , for example), and the number of counts inside the masked border can be summed. Plotting the number of counts versus the probe pulse duration should yield linear data in which the slope of the line provides the number of counts per second detected by the CCD camera. This can then be compared to the measurement made using the power meter to determine the efficiency of the camera. Using this technique, we measure our camera efficiency to be $\eta_{cam} = 0.50$, after we correct for gain issues, a topic we now address.

When using any detector, one must be wary of gain issues. Our CCD camera has gain characteristics which depend on the user-specified readout time per pixel. Further, for users of the Andor DV434-BV camera, we issue one further caution. Andor’s definition of gain is well-described by any number of adjectives, none of which are printable in this dissertation. With great restraint, I will generously call their definition of gain *counterintuitive*. Consider, for example, you elect to use a 16 μs readout time per pixel, which the Andor literature indicates will provide a gain of 1.4. It would be tempting to assume that this value of the gain means

that for every 10 photons that strike the CCD array, 14 detection events occur. However, Andor defines gain as the number of photons it takes to register a single count (detection event) on the array. In other words, a gain of 1.4 according to Andor’s definition means that 14 photons must strike the CCD array to produce 10 detection events. Sadly, Andor’s definition of gain is what the rest of the planet defines as the multiplicative inverse of gain. In the discussion that follows, I use g_{cam} to represent the “true” gain of the camera (or the multiplicative inverse of the gain specified in the Andor literature).

5.8 Final expression for the column density

The preceding sections have considered modifications to (5.10) to account for depolarization in the probe beam, optical pumping to dark quantum states, imperfect camera efficiency and camera gain issues. We have already modified the expression for the column density to account for probe beam depolarization, yielding (5.33). We now modify this expression to account for imperfect camera efficiency and camera gain. Recall that (5.33) takes the form

$$n(x, z) = -\frac{1}{\sigma_R} \left\{ (1 + \delta^2) \ln \left[\frac{I_s(x, z) + (\phi - 1) I_r(x, z)}{\phi I_r(x, z)} \right] + \frac{I_s(x, z) - I_r(x, z)}{I_{sat}} \right\}. \quad (5.36)$$

In our earlier discussion, we noted that $I_s(x, z)$ and $I_r(x, z)$ are interpreted as the (background-corrected) number of photon counts in a pixel located at position (x, z) in the signal and reference shots, respectively. We now need to determine how the above expression must be modified to account for imperfect camera efficiency and non-unity gain. First, we note that the argument of the natural log

in (5.36) will be unaltered, as multiplicative factors correcting for the camera efficiency and gain will cancel. The second term in braces in (5.36), however, will require some modification. Both $I_s(x, z)$ and $I_r(x, z)$ are measured quantities which are subject to the camera's imperfect detection characteristics, whereas I_{sat} is a calculated quantity denoting the number of photons per pixel in a probe pulse of duration t_{probe} at the saturation intensity. To compare these measured and calculated quantities on equal footing, we must consider a modification to this term to account for the camera efficiency η_{cam} and camera gain g_{cam} . We submit the following as the final expression for the column density,

$$n(x, z) = -\frac{1}{\sigma_R} \left\{ (1 + \delta^2) \ln \left[\frac{I_s(x, z) + (\phi - 1) I_r(x, z)}{\phi I_r(x, z)} \right] + \frac{I_s(x, z) - I_r(x, z)}{g_{cam} \eta_{cam} I_{sat}} \right\}. \quad (5.37)$$

In closing, note that the correct application of (5.37) requires that camera gain issues be accounted for when measuring the camera efficiency η_{cam} .⁴

Now that we have our final expression for the column density, we should be able to analyze images and extract physically meaningful quantities from them. Doing so, however, will require that we know any magnification factors introduced by our imaging system. Measuring the magnification is discussed in Section 5.9.

⁴That is, if your power meter indicates that there should be 10 photons striking the camera, and you detect 20 photons on your camera, and your camera has a true gain of 2, then your camera efficiency $\eta_{cam} = (20/2)/10 = 1$, rather than $\eta_{cam} = 20/10 = 2$.

5.9 Measuring the imaging system magnification

The imaging system presently in use relies on a fairly simple optical setup. A single achromat lens and a microscope objective deliver the image of the atom cloud to the CCD array in the imaging camera. These optics magnify the image of the atom cloud, and measuring this magnification factor is a relatively simple process. The final lens in the CO₂ laser beam path that creates the focus where the atoms are confined is mounted on translation stages which allow us to controllably translate the position of the trapped atoms. By moving the final lens transverse to the direction of propagation of the CO₂ laser beam, the position of the trapped atom cloud can be shifted on the CCD camera. As the size of individual camera pixels is known (13 μm on a side), and as the focusing lens position can be translated by a known amount, the magnification can be determined. Figure 5.8 shows the results of such a measurement. The apparent position of the atom cloud on the camera is plotted against the known translation of the CO₂ focusing lens. For our imaging system, the measured magnification is 5.22.

Accurate measurement of the imaging system magnification is critical. Failure to measure the magnification correctly will result in errors when calculating the atom number and the size of the atom cloud, among other physically relevant quantities.

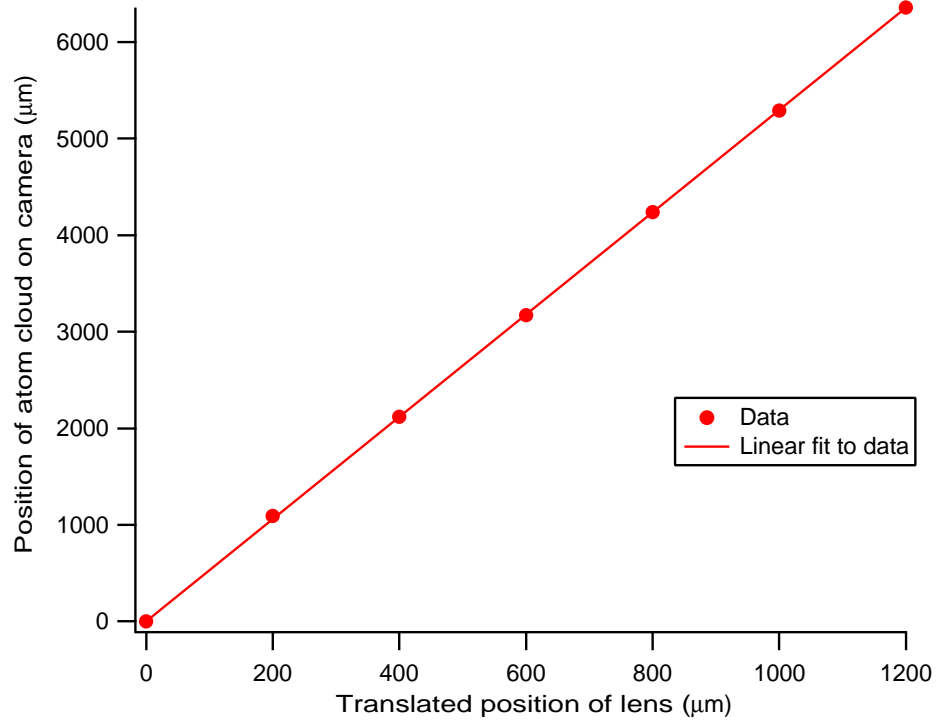


Figure 5.8: Measurement of the magnification of the imaging system. The atom cloud is translated a known amount by shifting the focussing lens which determines the position of the optical trap. The apparent motion of the optical trap on the camera is monitored, using the known size of a camera pixel. The ratio of the apparent motion on the camera to the true motion of the lens yields the magnification, which is provided by the slope of the line through the data.

Chapter 6

Heat Capacity of a Strongly Interacting Fermi Gas

Measurements of the heat capacity reveal information about one of the most basic thermodynamic properties of a material, and have been used to identify phase transitions in a number of classic thermodynamic or condensed matter systems (see, [106], for example). The heat capacity C is the change in energy associated with a change in temperature,

$$C = \frac{dE}{dT}. \quad (6.1)$$

Often, the heat capacity is measured while one thermodynamic quantity is held constant, such as the volume or the pressure. In our studies of degenerate Fermi gases there is no simple way of producing constant pressure or constant volume conditions. Instead, we measure the heat capacity of our degenerate Fermi gases in the presence of constant trap depth of our confining potential.

In this chapter, we report our measurement of the heat capacity of a strongly interacting Fermi gas, the first measurement of a thermodynamic quantity in a strongly interacting Fermi gas. In order to conduct a measurement of the heat capacity, we developed a technique for adding energy to the gas, as well as a method for measuring the temperature of the gas. The novel energy input

method allows for a precise estimate of the energy added to the gas, as discussed in Section 6.1. This is followed by a brief summary of the experimental sequence for studying the heat capacity in Section 6.2. Our temperature measurement scheme is an approximate method for measuring the temperature of a strongly interacting gas which is based on spatial profiles for noninteracting Fermi gases. Temperature measurement in the strongly interacting regime remains a somewhat controversial topic, and we compare our scheme to other methods presently in use in Section 6.3.

Interpretation of the results of our studies was aided greatly by interaction with a theory group at the University of Chicago. Frequent dialogue with Kathy Levin and Qijin Chen of the Chicago group resulted in a collaboration in which we obtained good quantitative agreement between their theory and our data. An apparent transition in our data is interpreted by their theory as a superfluid phase transition, which marks one of the first direct measurements of the critical temperature for the superfluid phase transition in a strongly interacting Fermi gas. This and other results are discussed in Section 6.4.

6.1 Energy input calculation

In principle, one can produce gases at different temperatures by stopping the evaporative cooling process prior to reaching the lowest temperatures. However, in practice, this technique suffers from less than ideal repeatability. Further, since evaporative cooling relies on ejecting hot atoms from the trapping potential, stopping the evaporative cooling process before reaching low temperature results in atom clouds with a larger number of atoms. As a result, if one wishes to study

a temperature dependent effect, the number varies along with the temperature. This is obviously not ideal. A second possible method of varying the temperature involves parametric heating, in which the trap depth is modulated at twice its natural frequency. In this case, the amplitude modulation of the trap results in resonant heating of the confined gas. Unfortunately, determining the amount of heating from parametric excitation can be difficult. In place of these two techniques, we developed an energy input method based on changing the potential energy of the atom cloud. Unlike the first method discussed (premature stoppage of evaporative cooling), this method permits variation of the temperature of the gas while holding the atom number constant. Further, unlike parametric heating, this method allows one to write a simple analytic expression for the energy input.

The starting point for our energy input method is a gas that has already been cooled to the lowest possible temperature. For our purposes, we can approximate the spatial density of our starting gas as a zero temperature Thomas-Fermi spatial distribution. We consider a two component Fermi gas with $N/2$ particles per spin state. The three-dimensional density profile for a noninteracting Fermi gas in the local density approximation is given by (see Section A.3.1)

$$n(x, y, z) = \frac{4 N}{\sigma_x \sigma_y \sigma_z \pi^2} \left[1 - \left(\frac{x}{\sigma_x} \right)^2 - \left(\frac{y}{\sigma_y} \right)^2 - \left(\frac{z}{\sigma_z} \right)^2 \right]^{3/2}, \quad (6.2)$$

where the Fermi radii $\sigma_i = (2 \varepsilon_F / m)^{1/2} / \omega_i$, ε_F is the Fermi energy, ω_i is the trap oscillation frequency in the i -direction, and m is the mass of a single fermion. To obtain the corresponding density in the zero temperature, unitary regime (see Section 2.2.1), let $\sigma_i \rightarrow \sigma_i^* = (1 + \beta)^{1/4} \sigma_i$ in (6.2). Further, in the unitarity limit $\varepsilon_F \rightarrow \varepsilon_F^* = \sqrt{1 + \beta} \varepsilon_F$.

The total potential energy *associated with a single spin state* is given by

$$U_{total} = \int n(x, y, z; t) U(x, y, z) dx dy dz, \quad (6.3)$$

where $U(x, y, z)$ is the potential energy provided by the confining potential. As written, equation (6.3) allows for a time-dependent density distribution. Naturally, for this calculation to be self-consistent, we should require that the trapping potential assume a harmonic form, as we have already assumed a density profile associated with a harmonic trapping potential. However, as we will be calculating small corrections to the harmonic assumption, we will leave the analytic expression for the trapping potential in its generic form for now.

The method of adding energy to the gas is clever, but relatively simple, and was first suggested by John Thomas. Starting with a very low temperature cloud, we abruptly turn off the trapping potential at time t_0 and allow the gas to expand for a time t_{heat} . We then abruptly restore the trapping potential, at which point the expanded cloud occupies a larger portion of the trap. It is this increased size that determines the amount of energy added to the gas. Immediately after recapture of the cloud, it will oscillate, as this method will excite the breathing mode of the gas (see Chapter 7). However, after the oscillations damp and the cloud has reached equilibrium, the increased potential energy has been redistributed throughout the atom cloud. Mathematically, we can express the added energy ΔE as

$$\Delta E = \int [n(x, y, z; t_0 + t_{heat}) - n(x, y, z; t_0)] U(x, y, z) dx dy dz. \quad (6.4)$$

Assuming that we have a well-characterized trapping potential, the task becomes determining the density profiles before and after the release and recapture

sequence. Prior to release, we assume that we are starting with a zero temperature profile given by equation (6.2). To obtain a form for the spatial profile after t_{heat} , we assume a scaling ansatz [107, 108] of the form

$$n(x, y, z; t) = \frac{n_0(x/b_x, y/b_y, z/b_z)}{b_x b_y b_z}, \quad (6.5)$$

where the $b_i = b_i(t)$ carry all of the time dependence and n_0 represents the initial density distribution at time t_0 for which $b_i(t_0) = 1$. Note that the $b_i(t)$ represent the same scaling factors as those discussed in our treatment of expansion dynamics in Chapter 4. The scaling ansatz expressed in (6.5) and the density (6.2) are valid provided that the initial density distribution is harmonically confined and obeys a polytropic equation of state. In our case, since we are starting from a very low temperature gas, the condition of harmonic confinement is well met prior to release from the optical trap. However, for long heating durations, t_{heat} , the gas will expand to sizes for which trap anharmonicity becomes important. Does the scaling ansatz (6.5) still hold? The answer is yes, because the scaling ansatz is valid if the *original* state is a harmonically trapped gas. If the gas does not heat during expansion (and we have not seen evidence for heating during expansion when starting with a low temperature cloud), then the scaling ansatz is valid. The gas can become aware of anharmonicity in the trap only upon restoration of the trapping potential, at which point the scaling ansatz has served its purpose. Finally, since most of our experiments occur in the strongly interacting or noninteracting regime, we can use hydrodynamic or ballistic expansion factors, respectively, to determine the values of the b_i .

We can rewrite equation (6.4) using equation (6.5), making the following sub-

stitutions: $\tilde{x} = x/b_x$, $\tilde{y} = y/b_y$, and $\tilde{z} = z/b_z$. We then obtain

$$\begin{aligned} \Delta E &= \int n(\tilde{x}, \tilde{y}, \tilde{z}) U(b_x \tilde{x}, b_y \tilde{y}, b_z \tilde{z}) d\tilde{x} d\tilde{y} d\tilde{z} \\ &\quad - \int n(x, y, z) U(x, y, z) dx dy dz. \end{aligned} \quad (6.6)$$

Letting $\tilde{x} \rightarrow x$, $\tilde{y} \rightarrow y$, and $\tilde{z} \rightarrow z$ in the first integral in equation (6.6) allows us to write the added energy in the more convenient form

$$\Delta E = \int n(x, y, z) [U(b_x x, b_y y, b_z z) - U(x, y, z)] dx dy dz. \quad (6.7)$$

Ultimately, it will be convenient to write the fractional change in the energy of the gas. To this end, we will normalize the added energy to the starting energy. For a zero temperature harmonically trapped Fermi gas, the average energy per particle is $3\varepsilon_F/4$. Consequently, for a zero temperature harmonically confined Fermi gas with $N/2$ particles per spin state, the total energy *of one of the spin states* (prior to the release and recapture heating sequence) is $E_{tot} = 3\varepsilon_F/4 \times N/2 = 3N\varepsilon_F/8$ (see Appendix A.2.3). Later, we will introduce a correction factor to account for the fact that we actually start from a small but finite temperature gas. If we consider a zero temperature system, however, the fractional change in energy produced by the release and recapture sequence is

$$\frac{\Delta E}{E_{tot}} = \frac{8}{3\varepsilon_F N} \int n(x, y, z) [U(b_x x, b_y y, b_z z) - U(x, y, z)] dx dy dz. \quad (6.8)$$

All that remains is to select a form for the trapping potential $U(x, y, z)$ and calculate the integral in equation (6.8).

We begin by assuming that we have a harmonic trapping potential given by

$$U_{HO}(x, y, z) = -U_0 + \frac{m}{2} (\omega_x^2 x^2 + \omega_y^2 y^2 + \omega_z^2 z^2). \quad (6.9)$$

Using the relationship between the trap oscillation frequencies ω_i and the Fermi radii σ_i , we can write the trapping potential in the form

$$U_{HO}(x, y, z) = -U_0 + \varepsilon_F \left[\left(\frac{x}{\sigma_x} \right)^2 + \left(\frac{y}{\sigma_y} \right)^2 + \left(\frac{z}{\sigma_z} \right)^2 \right]. \quad (6.10)$$

Plugging equations (6.10) and (6.2) into equation (6.8), and performing the integral yields

$$\Delta E_{HO} \equiv \frac{\Delta E}{E_{tot}} = \frac{1}{6} (b_x^2 + b_y^2 + b_z^2 - 3). \quad (6.11)$$

Consequently, if we start from a zero temperature gas and allow time t_{heat} to elapse between release and recapture of the gas, we can write the resulting energy of *both spin states* as

$$E(t_{heat}) = E_0 \left[1 + \frac{1}{6} (b_x^2 + b_y^2 + b_z^2 - 3) \right], \quad (6.12)$$

where $E_0 = 3 N \varepsilon_F / 4$ for a noninteracting gas and $E_0 = 3 N \varepsilon_F^* / 4$ for a unitary Fermi gas, where $\varepsilon_F^* = \sqrt{1 + \beta} \varepsilon_F$. Note that the b_i factors in (6.12) are dependent on t_{heat} . The simplicity of equation (6.12) is a large part of the appeal of this heating technique. The magnitude of the expansion factors $b_i(t_{heat})$ determine the energy of the gas after it has had time to reach equilibrium.

Thus far, we have assumed that we begin with a zero temperature gas prior to the energy input sequence. What if we do not start from a gas at zero temperature, but from a finite temperature system? If the temperature is very high, the above

formalism will fail, as the assumed density profile was for a zero temperature system. Naturally, one could simply repeat the above analysis using a numerically generated finite temperature spatial profile, but if the temperature of the sample is sufficiently low (that is, low enough that its density profile still looks very similar to a zero temperature density profile), we can make a single tweak to equation (6.12) to correct for the finite temperature of the starting atom cloud. We introduce a parameter η ,

$$E(t_{heat}) = \eta E_0 \left[1 + \frac{1}{6} (b_x^2 + b_y^2 + b_z^2 - 3) \right], \quad (6.13)$$

where $\eta = E(T_i)/E(T = 0)$. Here, T_i is the temperature of the gas prior to energy input, and η is the ratio of the energy in a cloud with temperature T_i to the energy in a zero temperature cloud. In the limit that we start with a zero temperature cloud, η obviously is equal to one. The value of $E(T)$ for a noninteracting Fermi gas is given by (A.28) in Section A.2.3. We note that (6.13) is valid for noninteracting as well as strongly interacting Fermi gases, provided that η and E_0 are calculated properly in each case.

The preceding analysis is valid for a harmonically confined Fermi gas. However, for high precision measurements, anharmonic corrections can be important. Analyzing the contribution of these anharmonic corrections is our next step. Before diving into the gory details, however, please note that we assume that the anharmonic corrections are small. We will still assume that we are starting from a low temperature cloud whose spatial profile can be approximated by a zero temperature harmonically trapped Fermi gas. For large degrees of anharmonicity, this assumption will break down and the following analysis will be invalid.

We begin with a more realistic potential for a focused dipole trap,¹

$$U(x, y, z) = \frac{-U_0}{1 + \frac{\varepsilon_F}{U_0 \sigma_z^2} z^2} \exp \left[-\frac{\varepsilon_F}{U_0} \left(\frac{x^2}{\sigma_x^2} + \frac{y^2}{\sigma_y^2} \right) \right]. \quad (6.14)$$

Let $\varepsilon \rightarrow \varepsilon_F^* = \sqrt{1 + \beta} \varepsilon_F$ and $\sigma_i \rightarrow \sigma_i^* = (1 + \beta)^{1/4} \sigma_i$ in (6.14) to obtain the corresponding expression for the potential in the unitarity limit. Note that a Taylor expansion of equation (6.14) up through quadratic terms will reproduce (6.10). If we keep quartic terms, the potential is approximately

$$U(x, y, z) \simeq -U_0 + \varepsilon_F \left(\frac{x^2}{\sigma_x^2} + \frac{y^2}{\sigma_y^2} + \frac{z^2}{\sigma_z^2} \right) - \frac{\varepsilon_F^2}{U_0} \left(\frac{x^4}{2\sigma_x^4} + \frac{y^4}{2\sigma_y^4} + \frac{z^4}{\sigma_z^4} + \frac{x^2 y^2}{\sigma_x^2 \sigma_y^2} + \frac{x^2 z^2}{\sigma_x^2 \sigma_z^2} + \frac{y^2 z^2}{\sigma_y^2 \sigma_z^2} \right). \quad (6.15)$$

Plugging (6.15) into (6.8), we arrive at the following expression for the energy of the gas following energy input and equilibration:

$$E(t_{heat}) = \eta E_0 (1 + \Delta E_{HO} + \Delta E_{AN1}), \quad (6.16)$$

where ΔE_{HO} is given by equation (6.11), $E_0 = 3 N \varepsilon_F / 4$ for a noninteracting gas and $E_0 = 3 N \varepsilon_F^* / 4$ for a unitary Fermi gas, and the anharmonic correction is

$$\Delta E_{AN1} = -\frac{\varepsilon_F}{40 U_0} \left[b_x^4 + b_y^4 + 2 b_z^4 + \frac{2}{3} (b_x^2 b_y^2 + b_x^2 b_z^2 + b_y^2 b_z^2) - 6 \right]. \quad (6.17)$$

If we keep sixth order polynomial terms in the Taylor expansion of (6.14), the

¹Note that equation (6.14) is not completely correct. If cylindrical symmetry ($\sigma_x = \sigma_y$) does not hold, there should be two Rayleigh lengths in the z -direction. However, provided that the system is nearly cylindrically symmetric, equation (6.14) is a good approximation.

calculated energy input becomes

$$E(t_{heat}) = \eta E_0 (1 + \Delta E_{HO} + \Delta E_{AN1} + \Delta E_{AN2}), \quad (6.18)$$

where ΔE_{AN2} is given by

$$\begin{aligned} \Delta E_{AN2} = & \frac{\varepsilon_F^2}{1440 U_0^2} [5 b_x^6 + 5 b_y^6 + 30 b_z^6 + 3 b_x^4 (b_y^2 + b_z^2) \\ & + b_x^2 (3 b_y^4 + 2 b_y^2 b_z^2 + 6 b_z^4) + 3 b_y^4 b_z^2 + 6 b_y^2 b_z^4 - 66]. \end{aligned} \quad (6.19)$$

In closing, we emphasize once more the broad applicability of (6.18). It has been observed that low temperature one-dimensional profiles in the strongly interacting regime closely resemble zero temperature Thomas-Fermi profiles [16]. Consequently, one of the powerful features of (6.18) is that it is equally applicable in the noninteracting and strongly interacting regimes, provided that we have a way of estimating the Fermi energy and the small temperature correction parameter η .

6.2 Experimental sequence

Sections 3.3.1 and 3.3.2 describe many of the general steps taken when producing degenerate noninteracting and strongly interacting Fermi gases. In this section, we provide greater detail on the final stages of cooling and energy input in our measurements of the heat capacity of a noninteracting and strongly interacting Fermi gas. A measurement of the heat capacity of the noninteracting Fermi gas, while of little interest from a theoretical standpoint, is an important test of our energy input technique. We discuss the preparation and manipulation of the

noninteracting gas in Section 6.2.2, but first we consider the strongly interacting gas in Section 6.2.1.

Before considering issues specific to the noninteracting and strongly interacting cases, there are some physical parameters of interest that are the same in both scenarios. Following forced evaporative cooling, the optical trap depth is recompressed to roughly 4.6% of the maximum trap depth, which gives a trap depth $U_0/k_B = 35 \mu\text{K}$. At this depth, the trap oscillation frequencies in the nearly cylindrically symmetric radial dimensions are $\omega_x = 2\pi \times 1785$ and $\omega_y = 2\pi \times 1612$, yielding a geometric mean of $\omega_\perp = \sqrt{\omega_x \omega_y} = 2\pi \times 1696(10)$. The oscillation frequency in the axial dimension is approximately $\omega_z = 2\pi \times 72(5)$. The values of the radial trap oscillation frequencies have been corrected for anharmonicity as described in Sections 4.6.1 and 7.4.2, while the axial trap oscillation frequency is the measured (uncorrected) value. Following forced evaporation and recompression, we have $N = 2(0.2) \times 10^5$ atoms, yielding a global Fermi energy, given by (A.22), of $\varepsilon_F/k_B \approx 2.4 \mu\text{K}$, for a noninteracting gas. Note that this Fermi energy is small compared to the trap depth, $U_0/k_B = 35 \mu\text{K}$.

6.2.1 Preparation of the strongly interacting Fermi gas

Following the initial stages of the experimental cycle discussed in Section 3.3.2, the applied bias magnetic field is ramped to 840 gauss, just above the center of the broad Feshbach resonance, and forced evaporation begins. The depth of the optical trap is reduced as described in Section C.3.1 over a period of roughly 4 seconds with a lowering time constant of 80 ms (see equation (C.1)). After remaining at low well depth for 0.5 seconds, the trap is recompressed to 4.6% of its maximum trap depth over a period of 1 second. After reaching this final trap

depth, no further experimental manipulation occurs for 0.5 seconds, which allows the atoms to reach equilibrium. Next, the energy input scheme outlined in Section 6.1 is applied, wherein the optical trap is turned off for a brief period. For the experiments, t_{heat} is varied from 0 to 460 μs , corresponding to $0 \leq \omega_{\perp} t_{heat} \leq 4.9$. For all of the heating times under consideration, it is a reasonable approximation to let $b_z(t_{heat}) = 1$, as the axial trap oscillation frequency is much less than the radial trap oscillation frequencies ($\omega_z/\omega_{\perp} = 0.042$). After restoration of the trap, the gas is allowed to equilibrate for 0.1 seconds. The gas is then released from the trap once again, at which point it expands and is imaged as outlined in Chapter 5. The time of flight used prior to imaging varied depending on the duration of t_{heat} . For $t_{heat} < 300 \mu\text{s}$, the time of flight was 1 ms or 800 μs , while for $t_{heat} > 300 \mu\text{s}$, the time of flight was 500 μs . The shorter time of flight for longer values of t_{heat} was necessary to maintain a reasonable signal-to-noise ratio when imaging the hotter atom clouds. When calculating the expansion factors during both the heating sequence and for time of flight imaging, the equations of hydrodynamics in the presence of a magnetic potential, outlined in Section 4.4, are applicable.

The experimental cycle was repeated many times to acquire all of the data presented in Section 6.4.2. The data for different values of t_{heat} was acquired in random order and ten measurements were conducted for each value of t_{heat} .

6.2.2 Preparation of the noninteracting Fermi gas

In preparing the noninteracting gas, the initial stages of cooling proceeded as discussed in Section 3.3.1. Forced evaporation was conducted at 300 gauss, where the optical trap depth was lowered for 12 seconds with a time constant of 0.25 seconds (see equation (C.1)). The trap depth remained at its minimum value for 1

second before being recompressed to 4.6% of the maximum trap depth in 1 second. Upon reaching the final trap depth, no experimental manipulations occurred for 0.5 seconds to allow the atoms to reach equilibrium. Energy was then added to the gas using the technique outlined in Section 6.1. Values of the heating time t_{heat} ranged from 0 to 300 μs in steps of 10 μs . As with the strongly interacting gas, we make the approximation $b_z(t_{heat}) = 1$, as the axial dimension remains nearly stationary for the values of t_{heat} under consideration. The magnetic field was maintained at 300 gauss for several tenths of a second to allow the cloud to equilibrate before being ramped to the zero crossing value near 530 gauss, where the scattering length $a_s = 0$. The gas was then released from the optical trap and allowed to expand for 1 ms before being imaged as described in Chapter 5. When calculating the expansion factors for the heating sequence and time of flight imaging, the equations for ballistic expansion in the presence of a magnetic field, given in Section 4.3, are applicable. However, the force supplied by the slight magnetic field curvature alters the ballistic expansion factors $b_i(t)$ by less than 0.2% from their zero magnetic field values for $t \leq 1$ ms. Consequently, when calculating the expansion factors for heating and for time of flight, the results for ballistic expansion in the absence of a magnetic field, given in Section 4.2, are a good approximation. The experimental cycle was repeated many times to acquire all of the data presented in Section 6.4.1. The data for different values of t_{heat} was acquired in random order and a total of ten measurements were conducted for each value of t_{heat} .

A close reading of the previous paragraph reveals that the energy input technique was applied to a weakly interacting gas at 300 gauss rather than a truly noninteracting gas, such as we have when the applied magnetic field matches the

zero crossing value. Recall from the discussion in Section 4.6.1, where we considered measurements of the parametric resonance frequencies at both 300 and 530 gauss, that the results were indistinguishable. That is, the weakly interacting system behaved essentially as a noninteracting system, though we elected to conduct parametric resonance excitations at 300 gauss, as that allowed for cross-relaxation between the primary axes of the trap. We employ the same reasoning here. In measuring the heat capacity, we first wish to add energy to the system and then allow the atoms to redistribute that energy and come to equilibrium. This process can occur in a weakly interacting system, but not in a noninteracting system.

Now that we have outlined the experimental manipulations which occurred during the process of adding energy to the noninteracting and strongly interacting Fermi gases, we must discuss the techniques used for measuring the temperatures of the resulting atom clouds. Combined knowledge of the energy and temperature will allow us to measure the heat capacity in the noninteracting and strongly interacting regimes.

6.3 Temperature measurement

Measuring the temperature of harmonically trapped noninteracting Fermi gases is a relatively straightforward process. The spatial density of the cloud can be compared to a theoretical prediction for a cloud of a particular temperature. A curve fitting algorithm can be employed to find the temperature of the theoretical curve that most closely matches the data. This technique is widely used to determine the temperature of noninteracting harmonically trapped Fermi gases (see [51, 109, 110], for example). There are some subtleties, however, which we

cover in Section 6.3.1.

Temperature measurement in the strongly interacting regime remains a rather contentious topic. Whereas the spatial profiles of noninteracting gases are well-known, at present there is no consensus on the appropriate spatial profiles in the strongly interacting regime. This condition is part of a larger problem, in which there is no single theory which is regarded as the “correct” theory for describing strongly interacting matter. The lone exception to this statement is the case of a zero temperature unitary gas, where the equation of state provided in Section 2.2.1 has gained widespread acceptance. In general, however, there is no agreed-upon technique for measuring temperature in the strongly interacting regime. Given the importance of measuring temperature in the presence of strong interactions, several techniques have been proposed. In Section 6.3.2, we outline our temperature measurement scheme for strongly-interacting Fermi systems, and compare it to some other temperature measurement methods presently in use.

Before discussing issues specific to the measurement of temperature in the noninteracting or strongly interacting regimes, we provide a few comments applicable to both cases. In Chapter 5, we discussed how images of the atom clouds are acquired and processed to yield two-dimensional column densities such as the one displayed in Figure 5.4. In principle, we could measure the temperature by fitting a two-dimensional theoretical density profile to the two-dimensional column density. However, we are concerned about the sensitivity of a two-dimensional fit routine to trap anharmonicity in the axial dimension of the atom cloud. Consequently, we prefer to analyze one-dimensional radial density distributions formed by binning the two-dimensional column density along the axial direction. An example of a one-dimensional radial density profile is displayed in Figure 6.1. Be-




Figure 6.1: A sample one-dimensional spatial distribution of an atom cloud, formed by binning the two-dimensional column density along the axial direction to produce a radial profile. The radial coordinate is shown along the horizontal axis while the vertical axis displays the atomic density in arbitrary units.

fore using a curve fitting routine to determine the temperature of the cloud, the density profile $n(x)$ is normalized such that

$$\int_{-\infty}^{\infty} n(x) dx = 1. \quad (6.20)$$

This normalized one-dimensional spatial profile can then be compared to theoretical profiles to determine the temperature of the cloud.

Finally, it is worth mentioning once more that all of the data discussed in this dissertation was released from the optical trap prior to imaging. Consequently, if we desire to know the temperature of the trapped cloud, we must be concerned about possible temperature and shape changes as the cloud expands. In Chapter 4, we covered the equations which describe the expansion of the atom clouds under ballistic and hydrodynamic conditions. In both cases, the shape of the

cloud in any one direction did not change, though its size did. Consequently, since we are trying to detect temperature by monitoring the shape of the cloud, we suggest that measuring the shape of the expanded cloud will reveal the shape of the trapped cloud, which in turn provides us with information about the temperature of the trapped cloud. In the noninteracting regime, this is not much of a concern, as the lack of collisions between the particles in the gas should prevent changes in the energy distribution, and therefore, the temperature of the cloud. In the presence of interactions, however, there are concerns that collisions could lead to heating of the atom cloud as it expands [64]. Our efforts to observe heating as an ultracold strongly interacting gas expands showed no discernable temperature increase when starting from a low temperature gas. Consequently, we believe that measuring the temperature of the expanded atom clouds in the low temperature regime provides a good estimate of the temperature of the cloud prior to release from the optical trap.

6.3.1 Noninteracting Gases

In the introduction to this section, we noted that the temperature of a harmonically trapped noninteracting Fermi gas influences the shape of the atom cloud. In practice, one can compare the shape of a real piece of data to a theoretically generated spatial profile with a known temperature. Using a curve fitting routine, the best fit to the spatial profile associated with the real piece of data will yield the temperature that best describes that piece of data.

Unfortunately, the process is not that simple, as there are several parameters to consider. In Appendix A.3.3, we present analytic expressions for the spatial profiles for harmonically trapped noninteracting Fermi gases. There, we find that

the normalized one-dimensional profile assumes the form

$$n(x; T) = -\frac{3N}{\sqrt{\pi}\sigma_x} \left(\frac{T}{T_F}\right)^{5/2} Li_{5/2} \left[\exp \left(\frac{\frac{\mu}{\varepsilon_F} - \frac{x^2}{\sigma_x^2}}{T/T_F} \right) \right], \quad (6.21)$$

where Li_n is the polylogarithm function, μ is the chemical potential and σ_x is the Thomas-Fermi radius in the x -direction. We see that the shape of the profile depends on several parameters. As noted already, the temperature of the atom cloud will influence its shape. Furthermore, the Fermi radius σ_x plays an important role in determining the spatial profile. If we suppose that we can measure the atom number to sufficiently high accuracy,² then a curve fitting routine based on equation (6.21) will need only σ_x and T/T_F as free fit parameters.

In addition to considering the spatial profiles for finite temperature harmonically trapped noninteracting Fermi gases in Appendix A, we also consider the zero temperature and high temperature limits. We call on those results now to illustrate the challenge given to a curve fitting algorithm that is handed a piece of real data and equation (6.21). We see that in the zero temperature limit (A.40), the one-dimensional profile becomes

$$n_Z(x) = \frac{8N}{5\pi\sigma_x} \left(1 - \frac{x^2}{\sigma_x^2}\right)^{5/2}. \quad (6.22)$$

In this case, if we know the number of atoms N , then only the Fermi radius of the cloud σ_x needs to be determined. In the high temperature regime given by

²The measured atom number generally agrees with the normalization of the atom cloud to within 5%. However, we do use the amplitude of the spatial profile as an additional fit parameter to obtain the best possible fit. The use of this additional fit parameter is largely independent of the curve fitting issues discussed in Section 6.3, and as such, to simplify the discussion, I consider the case where the atom number is known to high accuracy.

(A.51), the profile assumes a classical, gaussian form,

$$n_c(x) = \frac{N}{2\pi^{1/2}\sigma_x} \left(\frac{T}{T_F}\right)^{-1/2} \exp\left[-\frac{T_F}{T} \left(\frac{x^2}{\sigma_x^2}\right)\right]. \quad (6.23)$$

Here, we see that both the Fermi radius σ_x and the temperature T/T_F appear, and they enter the expression for the density profile symmetrically. That is, at high temperature, the curve fit routine can determine only the product $\sigma_x^2 \times T/T_F$, but not the individual values of the temperature and the Fermi radius. This is in stark contrast to the situation at zero temperature, where the Fermi radius and the temperature are wholly uncorrelated. In between these two limits, the Fermi radius and temperature gradually evolve from being completely independent fit parameters at low temperature to perfectly correlated fit parameters at high temperature. That is, the fit routine has a relatively easy task determining the Fermi radius and temperature when the atom clouds are quite cold, and it has an impossible task at high temperature.

Determining the temperature and the Fermi radius in the classical (high temperature) regime is only impossible if the fit routine considers both the Fermi radius and the temperature to be free fit parameters. If we provide the fit routine information about the Fermi radius of the cloud, then we can measure its temperature. There are several techniques for overcoming this indeterminacy. First, one can calculate the Fermi radius of the cloud using knowledge of the trap oscillation frequencies and the number of atoms (see (A.37)). This calculated Fermi radius can then be held constant while the curve fit routine determines the temperature. As always, the devil is in the details. Recall that we acquire images of clouds following release and expansion from the optical trap. The cost

of performing time of flight measurements is that one needs to understand the expansion dynamics, which requires knowledge of the trap oscillation frequencies, as discussed in Chapter 4. This is not a problem if we consistently produce the same trap conditions in each of our experiments. However, if we wanted to study a particular effect as a function of trap depth, it would be a very tedious task to measure the trap oscillation frequencies at each new trap depth.

There is a more convenient way of overcoming the correlation between the atom cloud's Fermi radius and its temperature in the classical regime. The solution is to measure the Fermi radius σ_x at low temperature (where the Fermi radius and the temperature are nearly independent), and hold the Fermi radius constant at the measured value for all higher temperatures.³ This technique has the advantage that one need not go through the time consuming process of thoroughly characterizing the trap conditions and then determining the expansion dynamics to measure the temperature.

Typically, the power of the laser which generates the optical trap is stable on the order of hours. This means that the trap oscillation frequencies ω_i are also stable on this time scale. The number of atoms which we load into the trap, however, will vary from shot to shot. As the number of atoms will affect the apparent Fermi radius of the cloud, we must take these fluctuations into account. Using the definition of the Fermi radius (A.37) and the Fermi energy (A.22), we can write

$$\sigma_i = \sqrt{\frac{2\varepsilon_F}{m}} \frac{1}{\omega_i} = \sqrt{\frac{2\hbar\bar{\omega}6^{1/3}}{m\omega_i^2}} \left(\frac{N}{2}\right)^{1/6} = c_i \left(\frac{N}{2}\right)^{1/6}, \quad (6.24)$$

³Note that the Fermi radius σ_x is a zero temperature parameter. As it does not have any temperature dependence, the value of σ_x measured at low temperature is applicable to all higher temperature clouds generated under the same trap conditions.

where $c_i = [2 \hbar \bar{\omega} 6^{1/3} / (m \omega_i^2)]^{1/2}$ is the “number independent Fermi radius.” (Recall that N is the total number of particles in the gas, so $N/2$ is the number of particles per spin state in a 50-50 mixture). One can accurately measure the number independent Fermi radius by measuring the Fermi radius σ_i of the cloud at low temperature and correcting for the number of atoms in that spin state by dividing σ_x by $(N/2)^{1/6}$. Averaging the measured values of c_i from 10 to 20 atom samples generated under the same conditions is sufficient to fix the value of c_i . The temperature of hotter atom clouds can then be determined by using (6.24), using the measured value of c_i along with the measured atom number for that particular cloud, to fix the value of σ_i . Once σ_i is fixed, the fit routine, based on (6.21) can then independently determine the temperature T/T_F .

The primary advantages of this technique are its speed and simplicity. Acquiring twenty pieces of data to fix the value of c_i takes no more than 10 minutes, and provided that the expansion dynamics of the cloud during time of flight are reasonably well understood, one can readily measure the temperature of the cloud. The primary drawback of this technique is that the initial determination of c_i requires very low temperatures. The fit routine begins to have difficulty distinguishing between Fermi radius and temperature for $T/T_F \geq 0.20$. Hence, to obtain a reliable measurement of the temperature, one should measure c_i using atom clouds at temperatures $T/T_F \leq 0.15$. If such low temperatures are unattainable, one must resort to measuring each of the trap frequencies and calculating the value of c_i .

While the discussion in this subsection has focussed on temperature measurement of noninteracting clouds, we will see that many of the same issues must be considered when attempting to extract temperature information from strongly

interacting clouds. Many of the solutions to these issues are the same, but as we will see in Section 6.3.2, some of the consequences of these solutions will have a profound impact on the natural energy and temperature scales for the strongly interacting system.

6.3.2 Strongly interacting gases

In the introduction to Section 6.3, we alluded to the difficulties of measuring temperatures of atom clouds in the strongly interacting regime. In the noninteracting regime, we have seen that the spatial density profiles are well known. This is not the case in the strongly interacting regime, however, so another method must be employed. In this section, we discuss an empirical thermometry method based on the similarities between the one-dimensional densities of noninteracting and strongly interacting gases. It is important to note that we consider this temperature measurement scheme to be applicable in the strongly interacting regime at or above the center of the broad Feshbach resonance in ${}^6\text{Li}$ at 834 gauss. Spatial profiles below the resonance, where molecules can form and Bose condense, reveal bimodal structures composed of a condensate surrounded by a thermal cloud [41, 42, 73–75]. Our empirical thermometry scheme is not applicable in this regime, the “BEC side” of the resonance. Unless otherwise stated, the discussion that follows pertains to temperature measurement in a strongly *attractive* Fermi gas.

Very simply, in the strongly interacting regime where some portion of the atom cloud might be superfluid while other regions are not, there is little reason to believe that the density profiles should look like noninteracting gas profiles. Yet this is what was first observed in [16], where one-dimensional profiles of ultracold,

strongly interacting Fermi gases were well-fit by zero temperature Thomas-Fermi profiles of the form given by (A.40). Even as the temperature of the gas is increased, the profile still appears to be well-fit by a Thomas-Fermi profile. Eventually, as the temperature rises to a level where the gas begins to acquire a more or less classical, gaussian profile, the Thomas-Fermi shape continues to be a “good” shape for the gas, as the Thomas-Fermi shape naturally tends to a gaussian profile for large temperatures. In Figure 6.2, we show three different one-dimensional density profiles of strongly interacting gases acquired at 840 gauss. Panels (a), (b), and (c) show density profiles of atom clouds of increasing temperature. The dots represent the data, while the solid curves show best fit Thomas-Fermi profiles for noninteracting, harmonically trapped Fermi gases. In spite of the strong interactions, the cloud profiles are well-fit by the noninteracting profiles [10, 111]. We utilize this surprising similarity to implement an approximate temperature measurement scheme.

Before continuing, it is important to note that the success of this temperature measurement scheme is not necessarily an indication that the spatial distribution of atoms in a strongly interacting, harmonically confined gas is nearly indistinguishable from a similarly confined noninteracting system. It does seem to be the case, however, for binned *one-dimensional* profiles. The three-dimensional clouds and the two-dimensional column densities associated with strongly interacting gases might have more complex features (kinks or bends) that get “washed out” when the distributions are integrated to produce one-dimensional profiles. Given our current imaging techniques, it is impossible to reconstruct the shape of the three-dimensional cloud. However, viewing of the two-dimensional column densities does not reveal any obvious nonmonotonic features in the spatial distribution

of atoms. A more substantial study of these two-dimensional profiles involving the averaging of many atom clouds may reveal the existence of unusual features in the atomic density. Such a study was not undertaken for this dissertation.

Recall that when measuring the temperature of noninteracting Fermi gases, we hold the Fermi radius σ_i , which in turn is based on c_i as given in (6.24), constant at a measured value while using the reduced temperature T/T_F as a floating fit parameter. We will use essentially the same procedure here. Using noninteracting Fermi gas profiles, we will measure the Fermi radius σ_i^* of a cold atom cloud. As with the noninteracting gas, we can measure a “number independent Fermi radius” c_i^* using,

$$\sigma_i^* = c_i^* \left(\frac{N}{2} \right)^{1/6}, \quad (6.25)$$

where asterisks are being used to distinguish quantities from their noninteracting gas counterparts. After measuring c_i^* for ultracold gases, we can then hold this value of c_i^* constant for hotter clouds, using (6.25) to calculate the appropriate value of σ_i^* for a cloud with $N/2$ atoms per spin state. With the value of σ_i^* fixed, we use a modified version of (6.21),

$$n(x; \tilde{T}) \propto Li_{5/2} \left[\exp \left(\frac{\frac{\mu}{\varepsilon_F} - \frac{x^2}{(\sigma_x^*)^2}}{\tilde{T}} \right) \right], \quad (6.26)$$

to measure the empirical reduced temperature \tilde{T} for that atom cloud.

At this point, we have a way of labelling the temperature of the gas. The above formalism cannot *determine* the true temperature T/T_F in the strongly interacting regime, but it does allow us to assign an empirical temperature \tilde{T} to strongly interacting clouds. The next logical step, then, is to consider the physical significance of measuring temperature in this manner.

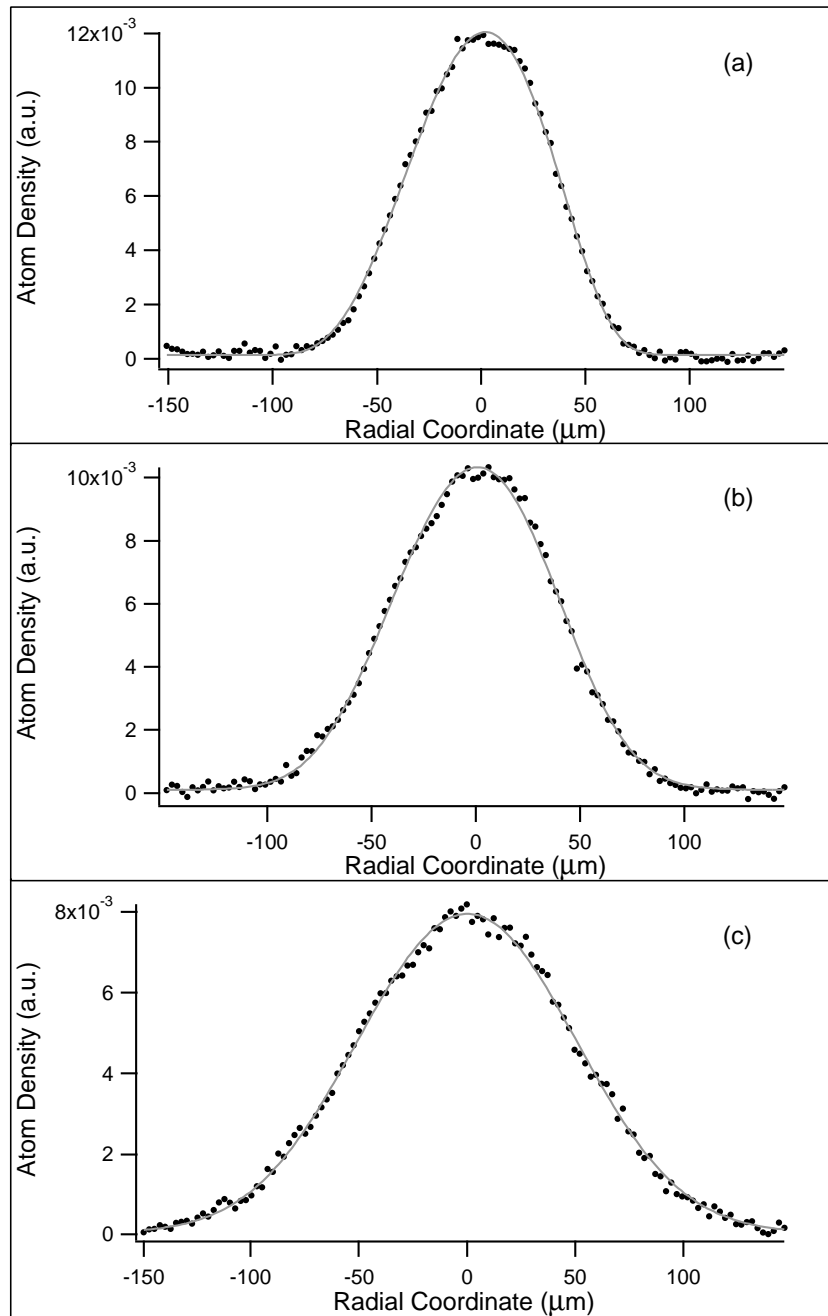


Figure 6.2: Three normalized one-dimensional density profiles (dots) of clouds imaged at 840 gauss in the strongly interacting regime. The solid gray lines represent best fit Thomas-Fermi density profiles for noninteracting Fermi gases. Panel (a) shows a cloud at empirical temperature $\tilde{T} = 0.11$, panel (b) shows a cloud at $\tilde{T} = 0.32$, while panel (c) shows a cloud at $\tilde{T} = 0.62$. The meaning of the empirical temperature parameter \tilde{T} is addressed in Section 6.3.2.

Our first attempt at understanding the meaning of \tilde{T} will draw on our discussion of unitary Fermi gases in Section 2.2.1. There, we saw that dimensional analysis allows us to write a modified version of the equation of state for a noninteracting gas to express the equation of state for a unitary system. Furthermore, recall that the modification involved the introduction of the parameter β , which is a universal constant representing the ratio of the interaction energy to the local Fermi energy. Simple algebraic manipulation of (2.14) revealed that the unitary gas behaves in many regards like a noninteracting gas composed of particles with an effective mass $m^* = m/(1 + \beta)$. Since $\beta < 0$, the effective mass exceeds the bare mass, $m^* > m$. For a harmonically trapped zero temperature Fermi system, the Fermi radii obey (see Appendix A.3.1)

$$\sigma_i = \left(\frac{2 \varepsilon_F}{m} \right)^{1/2} \frac{1}{\omega_i}, \quad i = x, y, z. \quad (6.27)$$

If we substitute the expression for the effective mass in (6.27), along with the appropriate modifications to the expressions for the Fermi energy and the trap oscillation frequencies, $\omega_i^* = \sqrt{1 + \beta} \omega$ and $\varepsilon_F^* = \sqrt{1 + \beta} \varepsilon_F$, we find that the Fermi radii in the unitary case σ_i^* are related to the noninteracting Fermi radii σ_i via

$$\sigma_i^* = (1 + \beta)^{1/4} \sigma_i. \quad (6.28)$$

Before continuing, it is worth noting that (6.28) has been used to measure the value of β based on cloud sizes in the strongly interacting regime [10, 16, 18, 41, 43].

This discussion of β is intended as more than an illuminating aside. The preceding discussion should make it apparent that the value of β sets the characteristic length scale (via the Fermi radii) along with the characteristic energy

scale (via the Fermi energy). Further, since the Fermi energy in the unitarity limit $\varepsilon_F^* = \sqrt{1 + \beta} \varepsilon_F$ is proportional to the Fermi temperature $T_F^* = \sqrt{1 + \beta} T_F$, we propose a natural reduced temperature scale for the unitary regime,

$$\tilde{T}_{nat} = \frac{T}{T_F \sqrt{1 + \beta}}. \quad (6.29)$$

We are now in a position to consider whether the natural temperature scale given by (6.29) is equivalent to the empirical reduced temperature \tilde{T} which appears in the expression for the density profiles given by (6.26). At zero temperature, we require that $T/T_F = \tilde{T} = 0$. The natural temperature scale \tilde{T}_{nat} obviously meets this requirement. Next, we consider the high temperature regime. Recall that in the classical (high temperature) regime, the Thomas-Fermi profiles become gaussian (see (6.23)). Furthermore, recall that the Fermi radius and reduced temperature enter gaussian profiles as the product, $\sigma_x^2 \times T/T_F$. Even in the presence of interactions, we expect the gaussian shape to be valid in the high temperature regime, where the gas becomes more classical in nature. If we let $\sigma_x \rightarrow \sigma_x^*$ in the product $\sigma_x^2 \times T/T_F$, we require that $T/T_F \rightarrow \tilde{T}_{nat}$ to maintain appropriate normalization of the gaussian cloud. That is,

$$\sigma_x^2 \frac{T}{T_F} = \frac{(\sigma_x^*)^2}{\sqrt{1 + \beta}} \frac{T}{T_F} = (\sigma_x^*)^2 \tilde{T}_{nat}. \quad (6.30)$$

Consequently, if we let $\sigma_x \rightarrow \sigma_x^*$ and $T/T_F \rightarrow \tilde{T}_{nat}$ in (6.23), we can measure the temperature of a high temperature, interacting atom cloud by holding the Fermi radius constant at σ_x^* . The reported temperature will be \tilde{T}_{nat} , which we can then relate to the reduced temperature T/T_F using (6.29). Since the finite temperature Thomas-Fermi shape given in (6.26) becomes gaussian in the classical limit, we

conclude that the empirical temperature $\tilde{T} = \tilde{T}_{nat}$ in the high temperature regime.

To this point, we have established that $\tilde{T} = \tilde{T}_{nat}$ in the zero temperature and high temperature limits. However, we do not have any simple means of relating the empirical reduced temperature \tilde{T} to the “true” temperature T/T_F in the intermediate temperature regime. In order to do so, we require a more sophisticated approach, one that allows us to calibrate \tilde{T} over all temperature ranges.

Temperature calibration

The temperature calibration technique is relatively simple, though it does result in a model-dependent temperature scale. The essential thrust involves generating theoretical one-dimensional profiles of strongly interacting clouds at a number of different temperatures, and then subjecting them to the temperature measurement scheme outlined in Section 6.3.2. In doing so, we can establish a one-to-one correspondence between the empirical temperature \tilde{T} and the “true” temperature T/T_F . As noted earlier, there is no universally-accepted theory for strongly interacting gases. Consequently, each unique theory of strongly interacting matter will predict different spatial profiles, which will yield a different temperature calibration. Furthermore, each temperature calibration will be valid for a particular strength of interparticle interactions. Each time one wishes to measure the temperature at a new magnetic field (i.e., new interaction strength), a new temperature calibration is needed.

The above calibration procedure was implemented with the help of Kathy Levin and Qijin Chen from the University of Chicago. Their theory is based on a pseudogap formalism in which particles in the trap can fall into three cate-

gories: condensed fermion pairs, noncondensed fermion pairs, and unpaired fermions [5, 8, 10, 111]. Their theory accounts for fermion-fermion and fermion-boson correlations, but not boson-boson correlations. Consequently, their theory is expected to have inappropriate behavior when the gas is deep in the BEC regime. This is not a concern to us, however, as we are interested in a temperature calibration in the strongly interacting limit. For more information on Levin’s pseudogap theory, the reader is encouraged to consult [5, 8, 10, 111].

In Figure 6.3, we show the results of the temperature calibration using Levin’s density profiles. In the Figure, the vertical axis gives the “true” temperature of the theoretical density profiles, while the horizontal axis shows the measured empirical temperature \tilde{T} scaled by a factor of $\sqrt{1+\beta}$. For this calibration, we used $\beta = -0.49$, as measured in [10]. The dots represent the measured temperature calibration, while the gray dotted line is the diagonal $T/T_F = \sqrt{1+\beta}\tilde{T}$. If \tilde{T} is equal to the natural temperature scale \tilde{T}_{nat} (6.29), then we would expect all of the dots to overlap the diagonal gray line. We see that close to zero temperature and for temperatures $T/T_F > 0.30$, the quantity $\sqrt{1+\beta}\tilde{T}$ provides a reasonably accurate estimate of the “true” temperature T/T_F . However, for finite reduced temperatures less than 0.30, $\sqrt{1+\beta}\tilde{T}$ tends to underestimate the true temperature. It is believed that the discrepancy between the measured empirical temperature \tilde{T} and the temperature T/T_F arises from the appearance of condensed fermion pairs. In Levin’s theory, condensation first appears at a critical temperature $T_c/T_F = 0.29$, which leads to a modification in the density profile of the gas. This is evidently reflected in the temperature calibration shown in Figure 6.3.




Figure 6.3: Temperature calibration in the strongly interacting regime based on theoretical density profiles [5, 8, 10, 111]. The vertical axis gives the “true” temperature of the theoretical density curves T/T_F , while the horizontal axis gives $\sqrt{1+\beta}\tilde{T}$. The circles represent the measured temperature correspondence between T/T_F and $\sqrt{1+\beta}\tilde{T}$, while the dotted line shows $T/T_F = \sqrt{1+\beta}\tilde{T}$. Figure taken from [10].

Based on the results shown in Figure 6.3, we conclude that

$$\tilde{T} \approx \tilde{T}_{nat} = \frac{T}{T_F \sqrt{1 + \beta}}, \text{ for } \tilde{T} \geq 0.42. \quad (6.31)$$

When attempting to determine the true temperature for $\tilde{T} < 0.42$, we will need to rely on the correspondence between the temperature scales shown in Figure 6.3. A power law fit to the measured temperatures which is constrained to pass through the point $\sqrt{1 + \beta} \tilde{T} = T/T_F = 0$ yields the following correspondence,

$$\sqrt{1 + \beta} \tilde{T} = 1.8 \left(\frac{T}{T_F} \right)^{1.49}, \text{ for } \tilde{T} < 0.42. \quad (6.32)$$

To summarize, to convert the empirical temperature \tilde{T} to the “true” temperature T/T_F , we can use (6.32) for $\tilde{T} < 0.42$ and (6.31) for all higher temperatures. In closing, we caution once again that these conversion factors will be dependent upon the theory used to generate the theoretical density profiles which serve as the basis for the temperature calibration.

The primary advantage of measuring the temperatures of strongly interacting profiles in this fashion is that we can use essentially the same temperature measurement techniques we use for noninteracting Fermi gases. However, extra effort is needed to determine the reduced temperature T/T_F associated with the measured value of \tilde{T} .

Applicability of this temperature measurement method outside of the strongly interacting regime

Thus far, we have focussed on temperature measurement in the unitary regime, as we intend to study the heat capacity in the unitarity limit. However, the

temperature calibration technique based on curve fitting of theoretically generated density profiles should be applicable outside the unitarity limit as well, provided that the one-dimensional density profiles continue to resemble noninteracting Thomas-Fermi shapes. While this precludes using this technique with bimodal density distributions in the BEC regime, it poses no restrictions for regions featuring strong attractive interactions. Of course, when we first introduced β , it was presented as a parameter valid only at unitarity. However, in an approximate fashion, we could parameterize the strength of the interparticle interactions by measuring a value of β using (6.28), for example, as a function of magnetic field. That is, we could measure $\beta = \beta(k_F a_s)$, where the dimensionless parameter $k_F a_s$ denotes the interaction strength. Here, k_F is the Fermi wave vector and a_s is the s-wave scattering length. With each new interaction strength, a new value of β could be measured, and a new calibration curve relating $\sqrt{1 + \beta \tilde{T}}$ and T/T_F could be developed.

Before considering the influence of trap anharmonicity on the accuracy of our temperature measurement scheme in Section 6.3.3, we briefly mention other temperature assessment techniques in the strongly interacting regime.

Other temperature measurement schemes in the strongly interacting region

Other schemes have been considered for estimating the temperature of gases in the strongly interacting regime. Some have used measurements of the condensate fraction on the BEC side of the resonance as a means of assessing the initial temperature of the gas. Here, estimates of the temperature can be made by assessing the fraction of the atom cloud which has Bose-condensed, compared

to the portion of the cloud composed of thermal, unpaired atoms. This is a common temperature measurement technique for weakly interacting atomic Bose gases. However, due to the width of the broad Feshbach resonance in ${}^6\text{Li}$, it is difficult to reduce the magnetic field to the point where the gas is truly weakly interacting. As a result, when estimating the temperature of the gas, care must be taken to account for the presence of interactions. For a discussion of some of the issues involved in measuring temperature in this way, see [112]. Once the temperature is estimated deep in the BEC regime, researchers conduct isentropic field sweeps into the strongly interacting regime (see [41, 78], for example). For ${}^{40}\text{K}$, isentropic field sweeps from the noninteracting regime above resonance into the strongly interacting regime have been used to assess the initial temperature of the gas [76]. However, isentropic field sweeps can fix the temperature of the gas in the strongly interacting regime only if the relationship between entropy and temperature is known for a strongly interacting system. Otherwise, there is no means of connecting the temperature in the weakly or noninteracting regimes, where the entropy versus temperature dependence is known, to the corresponding temperature when strong interactions are involved. As there has not been a measurement of the entropy versus temperature in the strongly interacting regime, estimates of the temperatures following isentropic sweeps into the presence of strong interactions rely on theoretical input [113, 114].

6.3.3 Influence of trap anharmonicity on temperature measurement

When discussing the energy input technique in Section 6.1, we considered the impact of slight anharmonicity in the trapping potential. As a measurement of

the heat capacity will require knowledge of both the energy and temperature of the cloud, we need a way of determining the influence of trap anharmonicity on temperature measurement as well. In Section 6.3, the temperature measurement technique we outlined relied on spatial profiles for a noninteracting Fermi gas confined in a perfectly harmonic trap. In this section, I will examine the errors arising from measuring the temperature of *anharmonically* confined atom clouds using a curve fitting routine based on *harmonically* trapped profiles. Throughout this section, the formalism we use will be drawn from noninteracting Fermi gas theory. However, we once again draw on the similarity of one-dimensional noninteracting and strongly interacting Fermi gas density profiles to justify the use of these results for both interaction regimes. To determine the impact of trap anharmonicity on the measured temperature, we first need to generate some anharmonic density profiles.

Generating anharmonic spatial densities

We begin with (A.33) from Appendix A:

$$n(x, y, z) = \frac{2\pi (2m)^{3/2}}{(2\pi\hbar)^3} \int_0^\infty d\epsilon f(\epsilon) (\epsilon - U)^{1/2} \Theta[\epsilon - U], \quad (6.33)$$

where $f(\epsilon)$ is the usual Fermi occupation function

$$f(\epsilon) = \frac{1}{\exp[(\epsilon - \mu)/(k_B T)] + 1}. \quad (6.34)$$

Further, μ is the chemical potential and U is the trap potential energy. We can write (6.33) in slightly more compact fashion by normalizing the energy terms to the Fermi energy and temperature. Making the substitutions $\epsilon \rightarrow \varepsilon_F \epsilon'$, $\mu \rightarrow \varepsilon_F \mu'$,

$U \rightarrow \varepsilon_F U'$, and $T \rightarrow T_F T'$, (6.33) becomes

$$n(x, y, z) = \frac{2\pi (2m)^{3/2} \varepsilon_F^{3/2}}{(2\pi\hbar)^3} \int_0^\infty d\epsilon' \frac{(\epsilon' - U')^{1/2} \Theta[\epsilon' - U']}{\exp[(\epsilon' - \mu')/T'] + 1}. \quad (6.35)$$

If we absorb all of the lead coefficients outside the integral into a normalization factor A , we can write the expression for the density as

$$n(x, y, z) = A \int_0^\infty d\epsilon' \frac{(\epsilon' - U')^{1/2} \Theta[\epsilon' - U']}{\exp[(\epsilon' - \mu')/T'] + 1}. \quad (6.36)$$

Equation (6.36) is valid provided that the trapping potential is infinitely deep. However, any realistic trap will have finite depth, which means that the integration over energy should be terminated when the energy of a particle exceeds the maximum depth of the trap (if the energy of the particle exceeds the maximum trap depth, then the particle is not trapped). In unscaled units, for a trap potential with maximum depth U_0 , the upper limit on the energy integral should be $\epsilon_{max} = U_0$. Since the energy integral in (6.36) is expressed in terms of scaled units, the upper limit on the integral for a trap with maximum depth U_0 is $\epsilon'_{max} = U_0/\varepsilon_F$. That is, for a trap potential with maximum depth U_0 , the spatial profile can be obtained by using the following:

$$n(x, y, z) = A \int_0^{U_0/\varepsilon_F} d\epsilon' \frac{(\epsilon' - U')^{1/2} \Theta[\epsilon' - U']}{\exp[(\epsilon' - \mu')/T'] + 1}. \quad (6.37)$$

As an example, when assessing the impact of trap anharmonicity on our measurements, we often consider a three-dimensional gaussian trapping potential

$$U'(r) = \frac{U_0}{\varepsilon_F} \left[1 - \exp\left(-\frac{\varepsilon_F r^2}{U_0 \sigma_r^2}\right) \right], \quad (6.38)$$

where $\sigma_r = \sqrt{2\varepsilon_F/m/\bar{\omega}}$. Plugging (6.38) into (6.37) and performing the integration numerically, one can obtain a three-dimensional density profile. Lower-dimensional spatial profiles can be obtained by integrating over the appropriate coordinates. The choice of a three-dimensional gaussian to approximate the trapping potential of the optical trap is motivated by the ability to calculate the chemical potential for the three-dimensional gaussian using a simple correction factor to the related result for a three-dimensional harmonic oscillator [60].

We now use the temperature measurement scheme outlined in Section 6.3.1 to examine the error arising from measuring the temperature of an anharmonically confined noninteracting Fermi gas using a curve fitting routine which assumes the gas is perfectly harmonically trapped. We assume that the gas is probed while it is confined in a three-dimensional gaussian potential as given by (6.38). Figure 6.4 depicts the “measured” temperature versus the “actual” temperature for two values of the ratio of the Fermi energy to the trap depth: $\varepsilon_F/U_0 = 0.049$ (squares and dots), a typical value for our experimental conditions, and $\varepsilon_F/U_0 = 0$ (gray dashed line), the case of perfect harmonic confinement. For $\varepsilon_F/U_0 = 0$, we obviously have perfect agreement between the measured and “actual” temperatures. However, for the anharmonically confined gas (squares), we find that the fit routine based on in situ temperature measurement systematically overestimates the temperature of the cloud. This is to be expected, as the “soft” gaussian potential increases the width of the atom cloud, resulting in an apparent larger temperature. The error becomes more pronounced as the temperature is increased: for temperatures below $0.6 T_F$, the measured temperature is within 5% of the true temperature. The temperature overestimate rises to 10% above $1.1 T_F$ and 20% above $1.8 T_F$. We now consider the significance of the dots shown in Figure 6.4.




Figure 6.4: Error introduced by measuring the temperature of a noninteracting Fermi gas confined in a three-dimensional gaussian potential using a curve fitting routine based on density profiles for a harmonically confined Fermi gas. The squares are the results of an in situ temperature measurement, while the dots represent temperature measurement after expansion according to an effective trapping frequency. The in situ temperature measurement shows monotonically increasing error as the temperature rises, while the temperature measurement after time of flight remains closer to the true value. The ratio of the Fermi energy to the trap depth considered here is $\varepsilon_F/U_0 = 0.049$ (squares and dots) and $\varepsilon_F/U_0 = 0$ (gray dashed line).

Thus far, we have found that for a “soft” potential, in situ temperature measurements based on harmonically trapped profiles will overestimate the actual temperature of the cloud. We now show that the temperature error is reduced if the measurement is made on a cloud released from the trap. For simplicity, we will consider the case of ballistic expansion from the trapping potential, as discussed in Section 4.2.

When a cloud of atoms is released from a trap, the expansion dynamics is heavily influenced by the trap conditions prior to release. In the case of perfect harmonic confinement, the expansion of the gas is parameterized by the trap oscillation frequencies ω_i . However, if the initial confinement is not perfectly harmonic, we can assign an effective frequency $(\omega_i)_{eff}$ which differs from the harmonic value, as outlined in Section 4.5.

We note that for “soft” potentials, the effective frequency is reduced from its harmonic value. This lower effective confinement frequency will result in slower expansion during time of flight. While anharmonic effects arising from a “soft” potential tend to make the cloud appear hotter when imaged in situ, the reduced effective frequency results in slower expansion during time of flight. Hence, the cloud starts larger, but expands more slowly than expected.

We return again to Figure 6.4, and consider the temperature measurement error introduced by fitting Thomas-Fermi spatial densities for a harmonically confined gas to a cloud trapped in a three-dimensional gaussian potential. The dots in the figure represent the measured temperature after ballistic expansion according to an effective trapping frequency. In this case, the x -dimension of the cloud expands by a factor of $b_{tof} = \sqrt{1 + [(\omega_x)_{eff} t]^2}$ after time of flight t . We employ the temperature measurement procedure outlined in Section 6.3.1, holding

the Fermi radius constant at the value measured for the lowest temperature data point. Furthermore, we assume that $\omega_x t > 5$. The temperature error in this case never exceeds 11% for reduced temperatures up to $T/T_F = 1.8$. In contrast, the in situ temperature measurement error can be twice as large in this temperature regime.

In the preceding estimates of temperature errors resulting from the measurement of anharmonically confined atom clouds using theoretical profiles of harmonically confined clouds, we considered a noninteracting gas undergoing ballistic expansion. If we repeat the analysis with atom clouds expanding hydrodynamically, we obtain essentially the same estimates of the temperature measurement error. Furthermore, as the spatial profiles of one-dimensional noninteracting and strongly interacting atom clouds appear to be quite similar, we believe that the temperature measurement errors depicted in Figure 6.4 are good estimates of temperature errors for a noninteracting gas following ballistic expansion as well as a strongly interacting gas following hydrodynamic expansion. Since we find that the error in the measurement is a few percent or less for reduced temperatures up to $1.0 T/T_F$, when considering the analysis of the heat capacity data in Section 6.4, we do not attempt to correct the measured temperatures for errors arising from anharmonicity.

6.4 Results and discussion

In Section 6.2, we outlined the experimental sequence for creating noninteracting and strongly interacting Fermi gases. Following energy input as described in Section 6.1, we measure the temperature of the gas as outlined in Section 6.3.

The result is a series of ordered pairs consisting of energies calculated using (6.18) and measured temperatures.

We begin in Section 6.4.1 with analysis of the heat capacity of a noninteracting Fermi gas. As noninteracting Fermi gases are well understood, this is not a particularly interesting case. It does, however, provide a check of our energy input and temperature measurement techniques for a noninteracting system. The more physically interesting investigation of the heat capacity of a strongly interacting Fermi gas is discussed in Section 6.4.2. Finally, we can use the heat capacity data to verify the virial theorem for a strongly interacting Fermi gas, as discussed in Section 6.4.3.

6.4.1 Heat capacity of a noninteracting Fermi gas

Before considering our heat capacity data for a noninteracting Fermi gas, we will consider theoretical predictions for the energy of the gas as a function of temperature, normalized to the ground state (zero temperature) energy. In Appendix A.2.3, we calculate the average energy per particle for a noninteracting, harmonically trapped Fermi gas as a function of temperature. The result is given by (A.28). Obtaining the heat capacity is then a simple matter of applying (6.1) to the energy versus temperature relation.

In earlier sections, we considered the influence of trap anharmonicity on our measurements, and we found that trap anharmonicity will provide a measurable effect on the energy input, but a negligible effect on the measured temperature of the gas for most of the temperature range of interest. Consequently, when developing the theoretical prediction of the energy versus temperature curve for the noninteracting gas, we must take trap anharmonicity into account. This can

be done with a simple modification to (A.28). For a three-dimensional gaussian potential such as (6.38), we can introduce a numeric factor $g(x)$ which provides the ratio of the density of states in a three-dimensional gaussian well to the density of states in a three-dimensional harmonic potential. The formula for $g(x)$ is given by [60]

$$g(x) = \frac{[-\ln(1-x)]^{3/2} (1-x)^{1/2}}{x^2} \frac{16}{\pi} \int_0^1 du u^2 \sqrt{\exp[(u^2-1)\ln(1-x)] - 1}, \quad (6.39)$$

where the argument x is the energy of the atom in units of the well depth such that $x = 0$ at the bottom of the well and $x = 1$ at the top. Using (6.39), we can write the temperature dependence of the energy of a Fermi gas trapped in a three-dimensional gaussian well of depth U_0 , normalized to the ground state energy, as

$$\frac{E_G(T')}{E_G(T=0)} = \frac{\int_0^{U_0/(k_B T_F)} d\epsilon' \epsilon'^3 g(k_B T_F \epsilon'/U_0) / \{\exp[(\epsilon' - \mu')/T'] + 1\}}{\int_0^1 d\epsilon' \epsilon'^3 g(k_B T_F \epsilon'/U_0)}. \quad (6.40)$$

As in Appendix A, μ' and ϵ' are the chemical potential and energy in units of the Fermi energy, respectively, and T' is the temperature in units of the Fermi temperature. We will make use of (6.40) when analyzing the measured energy versus temperature relation for a noninteracting Fermi gas.

With all of the theoretical framework in place, we are ready to examine the results of our energy versus temperature measurement in the noninteracting regime. Recall that prior to adding energy to the gas, we cool the gas to the lowest possible value. For the noninteracting gas, the lowest attainable temperature was $T/T_F = 0.24$, which gives $\eta = 1.34$ for the temperature correction in (6.18).

In Figure 6.5, we plot the measured energy versus temperature for a noninteracting Fermi system. Recall that the vertical axis showing the normalized energy is a calculated quantity based on the duration of t_{heat} , while the horizontal axis indicates the measured temperature. This breaks with the usual convention of displaying the dependent variable along the vertical axis. However, this method of displaying the data was chosen to facilitate a discussion of the heat capacity, which can be determined readily by taking the derivative of the energy versus temperature curve. We note that the data points have error bars for the temperature but not the energy, as the energy input is very precisely determined by t_{heat} . The temperature error bars indicate the standard deviation of the mean of the measured temperatures from multiple repetitions of the same value of t_{heat} . Recall from Section 6.2.2 that data for different values of t_{heat} was acquired in random order, and each data point displayed in Figure 6.5 is the average of ten repetitions for a given value of t_{heat} .

There are no surprises in the data. The agreement between the theory given by (6.40) (with $k_B T_F/U_0 = 0.068$) and the data in Figure 6.5 is quite good, validating both the energy input and temperature measurement techniques. Of course, Figure 6.5 is not truly a presentation of the heat capacity. For that, we would need to take the derivative of the energy versus temperature data points. However, due to statistical fluctuations in the data, calculating the slopes between adjacent data points would not yield an especially meaningful presentation of the data. What is important is that the measured energy versus temperature curve is in good agreement with theory, which implies that the measured heat capacity is in good agreement with the theory as well.

Having tested our methods with the well-known case of a noninteracting Fermi




Figure 6.5: The normalized energy versus reduced temperature for a noninteracting Fermi gas. The dots represent data points where the energy, based on the duration of t_{heat} , is calculated using (6.18) and the temperature is measured using the procedure discussed in Section 6.3.1. The solid line shows the theoretical prediction based on (6.40). E_0 represents the ground state energy. Figure first published in [10].

gas, we can now examine the results for the more physically interesting case of a strongly interacting Fermi gas.

6.4.2 Heat capacity of a strongly interacting Fermi gas

Prior to adding energy to the strongly interacting Fermi gas, we were able to lower the temperature of the gas to $\tilde{T} = 0.04$. In order to use (6.18), we need to estimate η , the correction factor which accounts for the finite temperature of the gas prior to the heating sequence. Due to the very low temperatures considered, where the energy varies little with increasing temperature, we suggest that ideal gas scaling can be used to estimate the temperature correction factor. Using a Sommerfeld scaling factor for the energy with the empirical temperature \tilde{T} , we obtain,

$$\eta = 1 + \frac{2\pi^2}{3} \tilde{T}^2 \approx 1.01, \quad (6.41)$$

which hardly affects the energy scale.

In the analysis that follows, comparison between data and theory will require a measured value of β . We measure $\beta = -0.49(0.04)$, a topic that is addressed in more detail at the end of this section.

Following energy input and equilibration, we obtain the data shown in Figure 6.6. As with the noninteracting gas data, we show the energy, a calculated quantity, along the vertical axis, and the measured empirical temperature \tilde{T} along the horizontal axis. In the Figure, the diamonds represent the average of ten data points for a given value of t_{heat} , while the solid line shows the theoretical prediction for ideal gas scaling. For the theory, we have used (6.40) with rescaled temperatures ($T' \rightarrow \tilde{T}$), and the ratio of the Fermi energy to the trap depth is




Figure 6.6: Energy versus temperature for a strongly interacting Fermi gas. The diamonds represent data points, while the solid line represents ideal gas scaling given by (6.40). In general, ideal gas scaling is in reasonable agreement with the measured energy versus temperature data points for a strongly interacting Fermi gas. The lowest temperature data point (square) is constrained to lie on the solid line. Figure first published in [10].

$\sqrt{1 + \beta} k_B T_F / U_0 = 0.049$. Since we have used ideal gas theory to estimate η , the lowest temperature data point, represented by a square in Figure 6.6, is constrained to lie on the ideal gas scaling curve. We see that over a wide range of temperatures, the measured energy versus temperature data for a strongly interacting Fermi gas is in surprisingly good agreement with that predicted by ideal gas scaling.

If we look closely at the coldest data, however, we find that the data does depart from ideal gas scaling. This is shown in Figure 6.7. As with Figure 6.6, the data are represented by diamonds, while the ideal gas scaling predicted by theory is the solid line. We see that in this low temperature region, the data are better fit by a power law of the form $E(t_{heat})/E_0 - 1 = 9.8 \tilde{T}^{2.53}$, shown in Figure 6.7 as the dotted line.




Figure 6.7: Energy versus temperature for a strongly interacting Fermi gas at low temperature. The data (diamonds) deviate from ideal gas scaling (solid line), but are in agreement with a best-fit power law $E(t_{heat})/E_0 - 1 = 9.8 \tilde{T}^{2.53}$ (dotted line) in the low temperature regime. The lowest temperature data point (square) is constrained to lie on the solid line. Figure first published in [10].

This departure from ideal gas scaling is even more evident if we plot the energy versus temperature data on a $\ln - \ln$ scale as shown in Figure 6.8. As with the other figures, diamonds represent the data, while the solid line represents ideal gas scaling. The dotted line is associated with the best-fit power law given in the preceding paragraph. There is a change in behavior at $\tilde{T} \approx 0.33$, where the value is chosen by the intersection of the best-fit power law and the noninteracting gas theory prediction. While this change in behavior is suggestive of a phase transition, we need theoretical support to provide such an interpretation.

Such theoretical support was provided by Kathy Levin's group at the University of Chicago. Frequent correspondence between Kathy Levin and Qijin Chen, both of Chicago, and members of the Duke group resulted in a joint publication [10]. The primary results of that collaboration are displayed in Figure 6.9. In the figure, both noninteracting and strongly interacting gas data are compared to theoretical predictions provided by the Chicago group. As with all other figures in this section, the theoretical predictions were made using the three-dimensional gaussian potential given by (6.38) to account for slight anharmonicity in our trapping potential.⁴ The agreement between the noninteracting data (dots) and the noninteracting theory (upper solid line) is quite good, as expected. The agreement between the strongly interacting data (diamonds) and the Chicago group's strongly interacting theory is also reasonably good over a wide temperature range. In all preceding figures, we have plotted the strongly interacting data using the empirical temperature scale \tilde{T} . In Figure 6.9, however, we use the temperature

⁴Note that the ratio of the Fermi energy to the trap depth differs between the noninteracting and strongly interacting cases. Since the strong interactions lower the Fermi energy by a factor of $\sqrt{1 + \beta}$ for the strongly interacting case, we have a lower ratio $\sqrt{1 + \beta} k_B T_F / U_0 = 0.049$ than we do for the noninteracting case $k_B T_F / U_0 = 0.068$.




Figure 6.8: Energy versus temperature for a strongly interacting Fermi gas plotted on a $\ln - \ln$ scale. The data (diamonds) agree well with ideal gas scaling (solid line) above $\tilde{T} \approx 0.33$. For $\tilde{T} < 0.33$, the data departs from ideal gas scaling. The dotted line gives a best fit power law in the low temperature regime. The lowest temperature data point (square) is constrained to lie on the solid line. The change in behavior at $\tilde{T} \approx 0.33$ is interpreted as a superfluid phase transition. Figure first published in [10].

calibration scheme outlined in Section 6.3.2 to replot the data in the more familiar reduced temperature scale T/T_F . The inset in Figure 6.9 shows the temperature calibration which was determined using the procedure considered in Section 6.3.2.⁵ Despite the good overall agreement between the strongly interacting data and the theory, we find that the low temperature strongly interacting data is better fit by a power law of the form $E(t_{heat})/E_0 - 1 = 97.3(T/T_F)^{3.73}$ represented by the dashed line in the figure. If we fit a power law to the strongly interacting data in the high temperature regime, we obtain $E(t_{heat})/E_0 - 1 = 4.98(T/T_F)^{1.43}$. The intersection of these two curves is $T/T_F = 0.27(0.02)$, where the error arises from statistical uncertainty only. The change in behavior at this temperature is interpreted by the theory as a superfluid phase transition, and the measured transition temperature $T/T_F = 0.27(0.02)$ is in good agreement with the theoretical prediction $T_c/T_F = 0.29$.

Using the power law fits to the energy versus temperature in the high and low temperature regimes, we can estimate the jump in the heat capacity at the apparent superfluid phase transition. Taking the derivative of the power laws provided in the preceding paragraph, we can obtain the heat capacity, as shown in Figure 6.10. If we calculate the “jump” in the heat capacity, we obtain

$$\frac{C_{<} - C_{>}}{C_{>}} \approx 1.51(0.05), \quad (6.42)$$

where $C_{<}$ ($C_{>}$) represents the heat capacity below (above) the critical temperature. The error bar in (6.42) takes into account statistical errors only. This mea-

⁵The comparison between theory and data occurred after a slight detuning of the theoretical model away from resonance to $1/(k_F a_s) = 0.11$, so that the model has the same value of β as what was measured experimentally. Given the width of the resonance, this amounts to a small adjustment of the theory.




Figure 6.9: Energy versus temperature for strongly interacting and ideal Fermi gases plotted on a $\ln - \ln$ scale. Noninteracting gas data (circles) agree well with ideal gas theory (upper solid line) over the entire range of temperatures considered. Strongly interacting Fermi gas data (diamonds) agrees reasonably well with strongly interacting gas theory (lower solid line) over a wide range of temperatures. The dashed line shows a best-fit power law to the low temperature data. The separation between the noninteracting and strongly interacting gas data is interpreted by theory as a superfluid phase transition. The arrow and marker $T/T_F = 0.27$ indicate our estimate of the superfluid phase transition temperature. The inset shows the temperature calibration used to convert the empirical temperature \tilde{T} to theoretical temperatures T/T_F for the strongly interacting gas data. In the inset, the arrow and T_c label indicate the critical temperature for a superfluid phase transition in the theory. Figure first published in [10].




Figure 6.10: Heat capacity and energy curve for the strongly interacting Fermi gas data. The dashed curve shows the piecewise function of power law fits to the energy versus temperature data for the strongly interacting Fermi gas above and below the change in behavior observed at $T/T_F = 0.27$. The solid curve shows the heat capacity derived from the energy versus temperature curve. The heat capacity jump near $T/T_F = 0.27$ is interpreted as an indicator of a superfluid phase transition.

sured value of the jump in the heat capacity is slightly larger than that expected for an s-wave BCS superconductor in a homogeneous case, where $(C_{<} - C_{>})/C_{>} = 1.43$. In closing, we caution that this estimate of the heat capacity jump may prove to be model-dependent, as the jump is estimated using data which has undergone the model-dependent temperature calibration procedure outlined in Section 6.3.2.

While we consider the measurement of a superfluid phase transition temperature of a strongly interacting Fermi gas to be the primary result of our study of the heat capacity, there is another observation worth mentioning. In Figure 6.9, we

note that the theory curves for the strongly interacting and noninteracting data do not overlap until a temperature well above the critical temperature for the superfluid phase transition. This is consistent with, although it does not prove, the existence of noncondensed atom pairs in the strongly interacting case prior to the onset of a superfluid phase transition.

Measurement of β

In our discussion of unitary Fermi gases in Section 2.2.1, we introduced the dimensionless parameter β , which parameterizes the interaction strength in a unitary Fermi system. In Section 6.3.2, we saw that β plays an important role in establishing the natural temperature and energy scales in the strongly interacting regime. Consequently, β plays an important role in the analysis of the heat capacity data.

To measure β , we produced ten ultracold strongly interacting Fermi gases at 840 gauss. Following preparation of the gas as outlined in Section 6.2.1, the strongly interacting atom cloud was released from the optical potential and allowed to expand for 1 ms before the cloud was imaged. The two-dimensional column densities were integrated along the axial dimension to produce radial density distributions. The Fermi radius of the one-dimensional distribution, σ_x^* was then measured for each atomic sample. Using a calculated hydrodynamic expansion factor ($b_x(1\text{ms}) = 13.3$), the size of the measured cloud was rescaled to determine the cloud size prior to release from the optical trap. Using (6.28), along with the value of the noninteracting Fermi radius σ_x calculated using (6.24) and verified by measurements of noninteracting Fermi gases, the value of β for each cloud was measured. Averaging the measured values of β , we obtain $\beta = -0.49(0.04)$, where the reported error is statistical error only. This measured value is in reasonable

agreement with predictions [38–40].

6.4.3 Virial theorem for a unitary Fermi gas

Before leaving our discussion of the heat capacity experiment, we consider an additional physical insight that can be obtained from further examination of the data. Using a simple theoretical argument, along with the energy versus temperature data from our study of the heat capacity, we verify the existence of a virial theorem for unitary Fermi gases [85]. We begin with the theoretical argument.

Recall from our discussion in Section 2.2.1 that in unitary Fermi gases, the details which govern the short-range interparticle interactions cease to be relevant, and universal behavior is expected. Consider such a gas to be confined in a harmonic potential $U(\mathbf{r})$. If we look at a small volume ΔV of the gas at position \mathbf{r} , we find that the enclosed number of particles ΔN have some energy ΔE . Further, suppose that the number of particles enclosed in this small volume remains constant, such that the spatial density $n = \Delta N/\Delta V$. Since the gas is in the unitary regime, the energy in this small volume can be related to the local Fermi energy $\epsilon_F(n)$ and the temperature. Consequently, we can write

$$\Delta E = \Delta N \epsilon_F(n) f_E \left[\frac{T}{T_F(n)} \right], \quad (6.43)$$

where $T_F(n)$ indicates that the local (density dependent) Fermi temperature. For a zero temperature noninteracting gas, $f_E = 3/5$, while for a zero temperature unitary Fermi gas, $f_E = 3(1 + \beta)/5$. Similarly, we can write an expression for the local entropy

$$\Delta S = \Delta N k_B f_S \left[\frac{T}{T_F(n)} \right], \quad (6.44)$$

where $k_B f_S$ is the average entropy per particle. We can use the thermodynamic relation

$$P = - \left[\frac{\partial \Delta E}{\partial \Delta V} \right]_{\Delta N, \Delta S} \quad (6.45)$$

to determine the local pressure P . When taking the partial derivative in (6.45), we need to hold ΔS constant, which implies that we must hold $T_F(n)$ constant according to (6.44). Consequently, equations (6.43) and (6.45) yield

$$P = \frac{2}{3} n \epsilon_F(n) f_E \left[\frac{T}{T_F(n)} \right]. \quad (6.46)$$

Equation (6.46) reveals that in the unitary gas, the pressure and the local energy density are related in the same manner as for an ideal, homogeneous gas. Equation (6.46) was first derived in [44], and rederived in [85].

In mechanical equilibrium the local pressure and the force of the trapping potential cancel, yielding

$$\nabla P(\mathbf{r}) + n(\mathbf{r}) \nabla U(\mathbf{r}) = 0. \quad (6.47)$$

If we take the inner product of (6.47) with \mathbf{r} and assume a harmonic oscillator potential, we obtain

$$N \langle U \rangle = \frac{3}{2} \int d^3 \mathbf{r} P(\mathbf{r}), \quad (6.48)$$

where $\langle U \rangle$ is the average potential energy per particle. If we then combine

$$\int d^3 \mathbf{r} P(\mathbf{r}) = \int d^3 \mathbf{r} \frac{2}{3} n \epsilon_F(n) f_E \left[\frac{T}{T_F(n)} \right] = E - N \langle U \rangle, \quad (6.49)$$

with (6.48), we obtain

$$N \langle U \rangle = \frac{E}{2}. \quad (6.50)$$

Equation (6.50) is a statement of the virial theorem for a harmonically trapped unitary Fermi gas. Finally, since $\langle U \rangle$ is proportional to the mean square size of a harmonically confined gas, (6.50) is equivalent to

$$\frac{E}{E_0} = \frac{\langle x^2(E) \rangle}{\langle x^2(E_0) \rangle}, \quad (6.51)$$

where E_0 represents the ground state energy of the gas.

Equation (6.51) can be verified easily using the data acquired for the heat capacity measurement. We already have all of the energy information for each value of t_{heat} . To measure the mean square size of the cloud, we fit each one-dimensional spatial profile with a finite temperature Thomas-Fermi shape. The mean square size of the best fit Thomas-Fermi shape is taken as the mean square size of the actual atom cloud. If we plot the energy versus mean square size, we obtain Figure 6.11. Here, we plot the mean square size versus the energy of the cloud, where both the size and the energies have been normalized to their ground state values. The data are represented by dots and the line of best fit is given by the dashed line. The best fit yields $\langle x^2(E) \rangle / \langle x^2(E_0) \rangle = 1.03(0.02)E/E_0$, in close agreement with (6.51).

This result indicates that in spite of the strong interactions in the gas, the potential energy still constitutes half of the total energy, just as it does for a noninteracting gas. In addition to verifying the virial theorem for a unitary Fermi gas, the linearity of the energy versus the size of the cloud serves as further confirmation of the accuracy of the energy input method outlined in Section 6.1.




Figure 6.11: Mean square size versus energy for a unitary Fermi gas. Both the size and energy have been normalized to their ground state values. The data is represented by dots, while the dashed line is a best fit to the data. The slope of the line is nearly one, in good agreement with (6.51). Figure first published in [85].

6.5 Conclusion

In this chapter, we considered a measurement of the heat capacity of a strongly interacting Fermi gas. The measurement was made possible by the development of a novel technique for adding energy to the gas. Further, an empirical temperature measurement technique yielding a reduced temperature \tilde{T} was developed for measuring the approximate temperature in the strongly interacting regime. A measurement of the heat capacity in a noninteracting gas confirmed the reliability of the energy input technique. For the strongly interacting gas, the energy versus temperature data departs from ideal gas scaling near $\tilde{T} = 0.33$. With the aid of Kathy Levin's theory group at the University of Chicago, we were able to calibrate our approximate temperature scale using their theory. We found good quantitative agreement with their BEC-BCS crossover theory, and our rescaled energy versus temperature data for a strongly interacting Fermi gas shows a transition at $T/T_F = 0.27(0.02)$, which is interpreted as a superfluid phase transition. This estimated critical transition temperature is close to the predicted critical temperature from their theory, $T_c/T_F = 0.29$. This constitutes the first measurement of a thermodynamic quantity of a degenerate Fermi gas in the strongly interacting regime, and the first attempt to directly measure the superfluid transition temperature at unitarity.

In the process of measuring the heat capacity of a strongly interacting gas, we also measured the universal many-body interaction parameter $\beta = -0.49(0.04)$. Further, the mean square widths of the clouds used in the heat capacity study were analyzed to verify the virial theorem for a harmonically confined unitary Fermi gas.

Chapter 7

Breathing Mode Measurements

Studying the collective modes in quantum gases is a useful technique, as it offers the possibility of learning about microscopic interactions by monitoring macroscopic behavior. In our recent studies of strongly interacting Fermi gases, we have spent a great deal of time monitoring the radial breathing mode of the gas, in which the trapped gas rhythmically expands and contracts much like a balloon being alternately inflated and deflated. As the response of the trapped atom cloud to this compression process reveals information about the pressure forces governing the particles, investigations of the radial breathing mode yield information about the equation of state of the gas. In light of the intense interest in developing a successful theory of strongly interacting fermionic matter, collective mode studies are a useful tool for testing the equations of state proposed by competing theories.

Our early studies of the radial breathing mode in a strongly interacting Fermi gas provided evidence of a superfluid phase transition [19]. While not a proof of the existence of a superfluid state, our data is difficult to explain in terms of non-superfluid scenarios. Subsequent and more exhaustive studies of the radial breathing mode [20, 79] conducted in our lab revealed features which seem to indicate the existence of atom pairs.

In Section 7.1, we discuss the technique for exciting the radial breathing mode of the gas. In fact, it is essentially the same procedure used in our studies of the heat capacity discussed in Chapter 6, though the time scales involved are markedly different. In Section 7.2, we discuss the preparation of the gas for excitation of the breathing mode. This is followed in Section 7.3 by an examination of predicted frequencies and damping times for the radial breathing mode based on the relaxation approximation formalism. In Section 7.4, we make predictions for the radial breathing mode by starting from the equations of hydrodynamics. We will recover some of the results of Section 7.3, in addition to determining formulas for correcting the measured breathing mode frequencies for trap anharmonicity. Finally, in Sections 7.5 and 7.6 we will discuss the results of our experimental investigations. In Section 7.5, we will consider the magnetic field dependence of the radial breathing mode for low temperature strongly interacting Fermi gases. In Section 7.6, we will examine the temperature dependence of the breathing mode frequency and damping rate in the unitarity limit. A summary of these results is provided in Section 7.7.

7.1 Exciting the breathing mode

Before discussing the technique we use to excite the breathing mode, it is instructive to consider what the breathing mode, once excited, looks like. For the conditions produced by our optical dipole trap, the confined atom cloud has a cigar shape. It is nearly cylindrically symmetric, with a longer length in the axial direction than in the radial dimensions. When the breathing mode of the atom cloud is excited, it looks like the cigar-shaped cloud is being rhythmically

inflated and deflated. As the trap oscillation frequencies in the radial and axial dimensions differ substantially ($\omega_z \ll \omega_x, \omega_y$), the breathing mode can be preferentially excited in either the radial or axial dimension. For technical reasons, it is more convenient for us to excite the radial breathing mode rather than the axial breathing mode, and so we deal exclusively with the former here. For studies of the axial breathing mode, the reader is encouraged to review the research carried out by Rudolf Grimm's group at the University of Innsbruck [21].

As the depth of our optical trap is controlled via an acousto-optic modulator (see Section 3.2.7), we can very quickly extinguish and restore the optical trap. It is this precise time control that allows us to excite the breathing mode. Suppose at time t_0 , we turn off the dipole trap. The cloud of atoms, no longer confined in the optical trap, begins to expand. After a short expansion time t_{excite} , the optical trap is turned on once again. When the optical trap is restored, the expanded atom cloud finds itself no longer in equilibrium with the trapping potential. As a result, the size of the cloud oscillates (see Figures 7.1 and 7.2). Eventually the oscillations damp as the cloud returns to nearly its original size (the size at time t_0). We assume here that the time t_{excite} is sufficiently small that very little energy is added to the gas in the process of exciting the breathing mode. Of course, some energy must be added to coax the atom cloud out of its equilibrium state, but for typical values of t_{excite} used in our measurements, the energy added and corresponding increase in temperature are negligible.

Once the breathing mode is excited, there are two major physical quantities of interest: the lifetime of the oscillations and the frequency of the oscillations. These quantities can be studied as the nature and strength of the interparticle interactions are varied, and they can be studied as a function of temperature.




Figure 7.1: Two dimensional false color absorption images showing the time evolution of the radial breathing mode of a strongly interacting cloud of ${}^6\text{Li}$. Images show the spatial density of the cloud, with lighter colors indicating higher atom densities. The times listed to the right of each panel show the time elapsed since the restoration of the optical dipole trap following the excitation time t_{excite} . The gas grows to a maximum size around 250 μs before shrinking back to a smaller size around 450 μs . The panels show roughly one complete period of oscillation.




Figure 7.2: The Fermi radius of the atom cloud as a function of hold time following the restoration of the trapping potential. Dots with error bars represent measured cloud sizes, while the solid line displays a least-squares fit of a damped sinusoid to the data. Both the frequency and the lifetime of the oscillations contain information of interest.

Before considering basic theoretical predictions for the oscillation frequency and damping time in Section 7.3, we will describe the preparation of the gas and excitation of the radial breathing mode in Section 7.2.

7.2 Experimental sequence

Our studies of the radial breathing mode included investigations throughout the BEC-BCS crossover region, as well as noninteracting Fermi gases. As with the study of the heat capacity, the noninteracting gas was studied as a check on our experimental techniques.

Each cycle of the experiment began with cooling the gas to the lowest temperature possible. This was accomplished by evaporative cooling using similar lowering curve times and time constants to those outlined in Section 6.2. For most of the experiments, the trap was recompressed to 4.6% of its maximum depth, and contained $N \approx 2 \times 10^5$ atoms. As with our study of the heat capacity, this yields a Fermi energy of around $\varepsilon_F/k_B \approx 2.4 \mu\text{K}$ for a noninteracting gas, small compared to the trap depth $U_0/k_B = 35 \mu\text{K}$. Following recompression to

4.6% of the maximum trap depth, the gas is given 0.5 seconds to reach equilibrium before exciting the breathing mode. As noted in Section 7.1, the radial breathing mode is excited by turning off the trapping potential for a brief period, t_{excite} . For most of the data discussed in this chapter, $t_{excite} = 25 \mu\text{s}$. Upon restoration of the trapping potential, the cloud begins to oscillate, and those oscillations are then monitored. Of course, the process of exciting the breathing mode adds energy to the gas, but this additional energy (given by (6.18), letting $t_{heat} = t_{excite}$) results in a temperature increase of $\Delta T/T_F < 0.015$ after the gas thermalizes, using non-interacting gas scaling. As this temperature increase is smaller than typical error bars on our temperature measurements, we consider it to be negligible.

Following the excitation of the breathing mode, the atom cloud is held in the optical trap for a time t_{hold} . The cloud is then released from the trap for time of flight imaging. For much of the data presented in this chapter, 1 ms time of flight was used prior to on-resonance absorption imaging. A schematic of the experimental sequence is displayed in Figure 7.3. Since on-resonance imaging techniques are destructive, many repetitions of the experiment were needed to study the breathing mode oscillations. Each data run typically consisted of 60-90 values of t_{hold} acquired in random order. Three complete data runs of this type were taken, and the resulting data was averaged to produce a curve such as the one shown in Figure 7.2. We were concerned with the width of the cloud as a function of time after excitation of the breathing mode. As with the heat capacity study, the two-dimensional absorption images were integrated along the axial direction to produce one-dimensional radial distributions. These one-dimensional profiles were then fit with a zero temperature Thomas-Fermi shape (as given by (A.40)) to extract an estimate of the Fermi radius σ_x . For much of the data discussed in this




Figure 7.3: Schematic of the sequence for exciting the radial breathing mode. The vertical axis indicates whether the trap is on (dark regions) or off (blank spaces), while the horizontal axis indicates time. The radial breathing mode is excited by turning off the trap for a brief period t_{excite} . When the trap is restored, the cloud begins to oscillate. The trap is kept on for a duration t_{hold} before being extinguished once more for time of flight imaging. The vertical line at the right indicates the time of the imaging probe pulse. Note that the figure is not to scale.

chapter, where $T/T_F < 0.20$, this technique is a reasonable way to parameterize the width of the atom cloud. For higher temperature conditions, such as those considered in Section 7.6, estimates of the cloud width based on zero temperature profiles were corrected using calculations of the second moment of the density distribution for finite temperature Thomas-Fermi profiles.

Although we presently lack a theory of the damping rate of the breathing mode of a strongly interacting Fermi gas, we found that the averaged widths of the cloud as a function of time following the breathing mode excitation were well fit by a damped sinusoid

$$\sigma_x(t) = \sigma_{x0} + A \exp(-t/\tau) \sin(\omega t + \phi). \quad (7.1)$$

Of particular interest to us is the damping time τ and the oscillation frequency ω .

In closing, we note that for much of the data, the gas was cooled directly at the magnetic field at which the breathing mode was excited. There are a few exceptions, however. For the noninteracting gas, evaporative cooling occurred at 300 gauss, before the magnetic field was ramped to the zero crossing near 530 gauss for excitation of the breathing mode. For much of the data in the BEC-BCS crossover region (750 gauss or above), the gas was cooled and studied at the same magnetic field. However, for magnetic fields in excess of 930 gauss, the gas was cooled closer to the Feshbach resonance before being ramped to the desired field at which the breathing mode was excited. This additional magnetic field sweep is believed to be adiabatic, and was used to prevent overheating of the magnets.

Before considering the results of our studies, it will help to have some expectations about the type of behavior we expect to observe under different experimental conditions. For this, we return to the machinery developed for the relaxation approximation in Section 4.1.

7.3 The relaxation approximation and the radial breathing mode

In Section 4.1, we used the relaxation approximation to calculate expansion factors following release of the atom cloud from the optical trap. We can use the same formalism to predict breathing mode oscillation frequencies and damping rates for a collisional gas. Recall that consideration of the time dependence of the phase space density of the gas in the relaxation approximation yields the following

equations:

$$\ddot{b}_i + \frac{1}{b_i m \langle \tilde{x}_i^2 \rangle} \left[\left\langle \tilde{x}_i \frac{\partial U_{EV}(b_x \tilde{x}, b_y \tilde{y}, b_z \tilde{z})}{\partial \tilde{x}_i} \right\rangle - \theta_i \left\langle \tilde{x}_i \frac{\partial U_{EQ}(\tilde{x}, \tilde{y}, \tilde{z})}{\partial \tilde{x}_i} \right\rangle \right] = 0 \quad (7.2)$$

$$\dot{\theta}_i + 2\theta_i \frac{\dot{b}_i}{b_i} = -\frac{1}{\tau_R} (\theta_i - \bar{\theta}). \quad (7.3)$$

In the preceding equations, $b_i(t)$ represents the scale factor in the i -direction, $\theta_i(t)$ is the effective temperature in the i -direction, and

$$\bar{\theta} = \frac{1}{3} \sum_i \theta_i. \quad (7.4)$$

U_{EQ} includes potential energy terms for equilibrium conditions, while U_{EV} includes potential energy terms which may be time dependent.¹ In the case of the breathing mode, the perturbation is weak (the gas expands very little during t_{excite}), and the optical trap is present as the gas vibrates. Further, the weakness of the perturbation means that we can consider the limit of small oscillations, in which b_i and θ_i remain close to their equilibrium values of unity. If we let $b_i(t) \rightarrow 1 + e_i(t)$ and $\theta_i(t) \rightarrow 1 + \phi_i(t)$, where $e_i, \phi_i \ll 1$, and linearize the resulting equations, (7.2) and (7.3) become

$$\ddot{e}_i + \omega_i^2 (2e_i - \phi_i) = 0 \quad (7.5)$$

$$\dot{\phi}_i + 2\dot{e}_i = -\frac{1}{\tau_R} (\phi_i - \bar{\phi}), \quad (7.6)$$

¹When considering the expansion dynamics of the gas following release from the trapping potential, we considered the influence of the magnetic field curvature on the expanding gas. The behavior of the breathing mode is dominated by the optical trap for the conditions under which the breathing mode data discussed in this chapter was taken. Consequently, we consider the effect of the optical trap only.

for the case of harmonic confinement. In (7.6),

$$\bar{\phi} = \frac{1}{3} \sum_i \phi_i. \quad (7.7)$$

We now study the radial breathing mode. Suppose the radial breathing mode is excited and the axial dimension remains stationary. In that case, $e_z = \dot{e}_z = 0$, and the system of equations governing the radial oscillations is

$$\ddot{e}_x + \omega_x^2 (2e_x - \phi_x) = 0 \quad (7.8)$$

$$\ddot{e}_y + \omega_y^2 (2e_y - \phi_y) = 0 \quad (7.9)$$

$$\ddot{\phi}_x + 2\ddot{e}_x + \frac{1}{3\tau_R} (3\dot{\phi}_x + 2\dot{e}_x + 2\dot{e}_y) = 0 \quad (7.10)$$

$$\ddot{\phi}_y + 2\ddot{e}_y + \frac{1}{3\tau_R} (3\dot{\phi}_y + 2\dot{e}_x + 2\dot{e}_y) = 0. \quad (7.11)$$

In a moment, we will numerically solve the above equations to extract both the damping rate and the oscillation frequency of the radial breathing mode. First, however, we will consider the noninteracting limit ($\tau_R \rightarrow \infty$) and the hydrodynamic limit ($\tau_R \rightarrow 0$).

If we let $\tau_R \rightarrow \infty$ in equations (7.8) through (7.11), the system of four equations decouple into two independent two-equation systems,

$$\ddot{e}_i + \omega_i^2 (2e_i - \phi_i) = 0 \quad (7.12)$$

$$\ddot{\phi}_i + 2\ddot{e}_i = 0. \quad (7.13)$$

Solving (7.13) for $\ddot{\phi}_i$ and substituting it into the second time derivative of (7.12)

yields

$$e_i^{(4)} + 4\omega_i^2 \ddot{e}_i = 0, \quad (7.14)$$

where $e_i^{(n)}$ denotes the n^{th} time derivative with respect to e_i , for $n > 2$. As the above equation obviously admits an oscillatory solution at a frequency $\omega = 2\omega_i$, we conclude that exciting the breathing mode of a noninteracting harmonically confined gas will yield oscillations at a frequency $2\omega_i$ in the i -direction.

We now consider the hydrodynamic limit, in which the characteristic time between collisions goes to zero. In this case, (7.5) and (7.6) become

$$\ddot{e}_x + \omega_x^2 (2e_x - \phi_x) = 0 \quad (7.15)$$

$$\ddot{e}_y + \omega_y^2 (2e_y - \phi_y) = 0 \quad (7.16)$$

$$3\dot{\phi}_x + 2(\dot{e}_x + \dot{e}_y) = 0 \quad (7.17)$$

$$3\dot{\phi}_y + 2(\dot{e}_x + \dot{e}_y) = 0. \quad (7.18)$$

Taking the time derivative of equations (7.15) and (7.16), and substituting expressions for the $\dot{\phi}_i$ from equations (7.17) and (7.18), we can reduce the above four equation system to a two equation system,

$$e_x^{(3)} + \omega_x^2 \left(\frac{8}{3} \dot{e}_x + \frac{2}{3} \dot{e}_y \right) = 0 \quad (7.19)$$

$$e_y^{(3)} + \omega_y^2 \left(\frac{8}{3} \dot{e}_y + \frac{2}{3} \dot{e}_x \right) = 0. \quad (7.20)$$

In the limit of perfect cylindrical symmetry ($\omega_x = \omega_y$), both the x - and y -dimensions are governed by the following equation:

$$e_i^{(3)} + \frac{10}{3} \omega_i^2 \dot{e}_i = 0. \quad (7.21)$$

The above obviously admits oscillatory solutions with a frequency $\omega = \sqrt{10/3}\omega_i \approx 1.826\omega_i$. If we do not have perfect cylindrical symmetry, we can still derive an analytic expression for the radial breathing mode in the hydrodynamic limit. If we plug solutions of the form $e_x(t) = F \exp(i\omega t)$ and $e_y(t) = G \exp(i\omega t)$ into (7.19) and (7.20) and eliminate F and G , we obtain a fourth order polynomial equation,

$$9\omega^4 - 24(\omega_x^2 + \omega_y^2)\omega^2 + 60\omega_x^2\omega_y^2 = 0. \quad (7.22)$$

Of the four solutions to (7.22), only two are linearly independent. The lower frequency solution corresponds to the quadrupole mode² while the higher frequency solution is associated with the radial breathing mode [101, 115], and is given by

$$\omega = \left[\frac{4}{3}(\omega_x^2 + \omega_y^2) + \frac{2}{3}\sqrt{4(\omega_x^4 + \omega_y^4) - 7\omega_x^2\omega_y^2} \right]^{1/2}. \quad (7.23)$$

Equation (7.23) can be simplified substantially if ω_x and ω_y are nearly equal. We define the geometric mean of the radial trap oscillation frequencies as

$$\omega_\perp = \sqrt{\omega_x\omega_y}. \quad (7.24)$$

If we let $\omega_x = \omega_\perp + \xi_x$ and $\omega_y = \omega_\perp - \xi_y$, where $\xi_x, \xi_y > 0$ and $|\xi_x - \xi_y| \ll \omega_\perp$, then (7.23) is approximately

$$\omega \simeq \sqrt{\frac{10}{3}\omega_x\omega_y} = \sqrt{\frac{10}{3}}\omega_\perp. \quad (7.25)$$

In the experiments, we monitor the radial breathing mode in the x -direction.

²For a cylindrically symmetric cigar-shaped trap, the radial hydrodynamic quadrupole mode oscillates at $\omega = \sqrt{2}\omega_\perp$, where ω_\perp is given by (7.24).

Equations (7.14) and (7.25) indicate that for a nearly cylindrically symmetric system, the relaxation approximation formalism predicts that the radial breathing mode oscillation frequency in the x -direction will remain in the following range:

$$\sqrt{\frac{10}{3}\omega_x\omega_y} \leq \omega \leq 2\omega_x. \quad (7.26)$$

To this point, we have made no mention of the lifetime of the breathing mode excitation. That is, we have not discussed how long the oscillations persist. In the two limits already examined (the hydrodynamic and noninteracting cases), we found that the solution for ω was real, corresponding to an infinite lifetime. In reality, of course, there will be experimental limitations to the observed lifetime of the mode. However, it is instructive to calculate the predicted lifetime and corresponding oscillation frequency for an ideal system. Our next goal, then, is to calculate the lifetime and frequency of the breathing mode oscillation for an arbitrary choice of τ_R . Recall that when τ_R was first introduced, it was remarked that τ_R was *related to* the characteristic time between collisions between constituent particles in the atom cloud, which we will call τ_0 for convenience. We never discussed the specific relationship between τ_R and τ_0 , as this was unnecessary to examine the two limits, $\tau_R \rightarrow \infty$ and $\tau_R \rightarrow 0$. Further, as τ_R is restricted to be between 0 and ∞ , we can simply assign a numerical value to τ_R between 0 and ∞ and numerically solve for the associated breathing mode frequency and lifetime. This is a considerable simplification, as the relationship between τ_R and τ_0 can in general be a complicated function governed by interparticle interactions.

We return once more to the system of equations (7.8) through (7.11). We consider solutions of the form: $e_x = A \exp(i\omega t)$, $e_y = B \exp(i\omega t)$, $\phi_x = C \exp(i\omega t)$,

and $\phi_y = D \exp(i\omega t)$. Making these substitutions into the system of equations, and eliminating A , B , C and D , we are left with the polynomial equation,

$$3\tau_R^2\omega^6 - 3[1 + 4\tau_R^2(\omega_x^2 + \omega_y^2)]\omega^4 + 8[\omega_y^2 + \omega_x^2(1 + 6\tau_R^2\omega_y^2)]\omega^2 - 20\omega_x^2\omega_y^2 - i\tau_R[6\omega^5 - 20(\omega_x^2 + \omega_y^2)\omega^3 + 64\omega_x^2\omega_y^2\omega] = 0. \quad (7.27)$$

Note that we have *not* assumed cylindrical symmetry ($\omega_x = \omega_y$) in (7.27). In the noninteracting limit ($\tau_R \rightarrow \infty$), (7.27) has four nontrivial ($\omega \neq 0$) solutions: $\omega = \pm 2\omega_x, \pm 2\omega_y$. In the hydrodynamic limit ($\tau_R \rightarrow \infty$), (7.27) reduces to a fourth-order polynomial with the solutions $\omega = \pm\sqrt{2}\omega_\perp, \pm\sqrt{10/3}\omega_\perp$ for a nearly cylindrically symmetric trap. Of the hydrodynamic solutions, the higher frequency solution ($\sqrt{10/3}\omega_\perp$) corresponds to the radial breathing mode in a cigar-shaped trap [115] such as the one in our laboratory. If we solve (7.27) numerically for the case of perfect cylindrical symmetry, we do indeed find that the radial breathing mode oscillation frequency smoothly varies between $1.826\omega_i$ and $2\omega_i$ for $i = x, y$ as τ_R is varied between 0 and ∞ . However, for the case of unequal radial confinement, we find that the limits of the oscillation frequency are slightly altered.

Under typical experimental conditions, we are quite close to having a cylindrically symmetric trap. However, as comparison between theory and experiment for breathing mode measurements requires substantial precision, it is critical to characterize the degree of trap anisotropy. Accurate measurements of ω_x and ω_y can be made using the parametric resonance excitation technique discussed in Section 4.6.1. For typical trap conditions, we find that $\omega_\perp = 2\pi \times 1696$, $\omega_x = 2\pi \times 1785 \approx 1.05\omega_\perp$, and $\omega_y = 2\pi \times 1612 \approx 0.95\omega_\perp$. Plugging these values




Figure 7.4: Locus plot showing the radial breathing mode oscillation frequency in the x -direction versus the lifetime of the mode (solid curve). The dashed line in the lower portion of the graph shows the hydrodynamic breathing mode oscillation frequency, while the dot-dashed line in the upper portion of the graph marks the collisionless (noninteracting) frequency. We assume that the axial dimension remains stationary while the radial directions oscillate. The computation was performed for typical trap conditions where the harmonic confinement frequencies are given by: $\omega_{\perp} = 2\pi \times 1696$, $\omega_x = 2\pi \times 1785 \approx 1.05 \omega_{\perp}$, and $\omega_y = 2\pi \times 1612 \approx 0.95 \omega_{\perp}$. Note that the oscillation frequency is a multi-valued function of the damping rate.

into (7.27) and solving numerically for ω , we can obtain both the predicted oscillation frequency ($Re[\omega]$) and the predicted lifetime of the oscillations ($Im[\omega]$). The results for the predicted behavior of oscillations in the x -direction (which we measure in the experiment) are shown in the locus plot in Figure 7.4.

There are several important features shown in Figure 7.4. First, we note that for a given lifetime of the breathing mode, there are two available oscillation frequencies, which correspond to two different values of the characteristic time

between interparticle collisions. Further, we note that starting on the lower leg of the graph close to the minimum damping rate (maximum oscillation lifetime) and following the solid curve counterclockwise, we are progressing from the hydrodynamic limit (infinite collision rate) to the noninteracting limit (infinite time between collisions). Between these two limits, we find a maximum damping rate at which the oscillation frequency is single-valued.

We note in closing that while the relaxation approximation formalism is often applied to weakly interacting systems, we expect that it can shed some light on the behavior of more complex, strongly interacting systems. At no point in the preceding analysis did we make any assumptions about the explicit nature of the interparticle interactions. We merely made assumptions about the evolution of the phase space density of the system, and then systematically varied the relaxation time τ_R between its two well-known limits (0 and ∞). Based on these assumptions, we were able to establish limits on the radial breathing mode oscillation frequency as well as make a prediction about the maximum possible damping rate for this formalism.

7.4 Hydrodynamic gases obeying a polytropic equation of state

In Section 7.3, we considered the behavior of the radial breathing mode as the relaxation time τ_R was varied from zero to infinity. The relaxation approximation formalism allowed us to explore the collisionless and hydrodynamic limits, as well as all points in between. In this section, we will consider another approach to calculating the oscillation frequency of the radial breathing mode. We will derive

predictions for the radial breathing mode for a gas which obeys a polytropic equation of state, in which the pressure $P \propto n^{\gamma+1}$, where γ is a constant and $n = n(x, y, z, t)$ is the spatial density of the gas. In doing so, we will recover the results from Section 7.3 in the hydrodynamic limit. However, in this section, we will go beyond the harmonic approximation, and consider how anharmonicity impacts the measured breathing mode frequency. While these results will be derived using the equations of irrotational hydrodynamics, we will see that a simple correspondence between the equations for a hydrodynamic gas and those for a noninteracting gas will allow us to simultaneously establish the effect of anharmonicity on a noninteracting system as well.

We begin with Euler's equation for irrotational flow

$$\frac{d\mathbf{u}}{dt} + \nabla \left(\frac{\mathbf{u}^2}{2} \right) = -\frac{1}{m} \nabla U - \frac{1}{m n} \nabla P, \quad (7.28)$$

where \mathbf{u} is the stream velocity obeying $\nabla \times \mathbf{u} = 0$, U is the trap potential, P is the pressure, and m is the mass of a particle. We now assume a scaling ansatz [107, 108] of the form

$$n(x, y, z, t) = \frac{1}{\Gamma} n_0(\tilde{x}, \tilde{y}, \tilde{z}), \quad (7.29)$$

where $\Gamma \equiv b_x b_y b_z$, $\tilde{x} = x/b_x$, $\tilde{y} = y/b_y$, $\tilde{z} = z/b_z$, the factors $b_i = b_i(t)$ carry all of the time dependence [$b_i(0) = 1$, $\dot{b}_i(0) = 0$], and n_0 denotes the density distribution at time $t = 0$. Substituting (7.29) into (7.28) and looking only at the equation for the x -direction, we obtain

$$\frac{n_0(\tilde{x}, \tilde{y}, \tilde{z})}{\Gamma} \tilde{x} \ddot{b}_x + \frac{n_0(\tilde{x}, \tilde{y}, \tilde{z})}{m \Gamma b_x} \frac{dU}{d\tilde{x}} + \frac{1}{m b_x} \frac{dP}{d\tilde{x}} = 0, \quad (7.30)$$

where $U = U(b_x \tilde{x}, b_y \tilde{y}, b_z \tilde{z})$ and $\ddot{b}_x = d^2 b_x / dt^2$. Multiplying (7.30) by \tilde{x} and integrating over the density distribution yields

$$\frac{\ddot{b}_x}{\Gamma} \langle \tilde{x}^2 \rangle + \frac{1}{m \Gamma b_x} \langle \tilde{x} \frac{dU}{d\tilde{x}} \rangle + \frac{1}{m b_x} \int \tilde{x} \frac{dP}{d\tilde{x}} d\tilde{x} d\tilde{y} d\tilde{z} = 0, \quad (7.31)$$

where the density averaged quantities are defined as

$$\langle \xi \rangle \equiv \frac{1}{N} \int \xi(\tilde{x}, \tilde{y}, \tilde{z}) n_0(\tilde{x}, \tilde{y}, \tilde{z}) d\tilde{x} d\tilde{y} d\tilde{z}. \quad (7.32)$$

Note that n_0 obeys the following normalization

$$N = \int n_0(\tilde{x}, \tilde{y}, \tilde{z}) d\tilde{x} d\tilde{y} d\tilde{z}. \quad (7.33)$$

We proceed by assuming a polytropic equation of state $P = c n^{\gamma+1}$, where c and γ are constants. Using this expression for the pressure along with equilibrium force balance considerations ($\nabla U_0 + \nabla P_0 / n_0 = 0$, where “0” subscripts indicate equilibrium conditions), we can write

$$\frac{dP}{d\tilde{x}} = - \frac{n_0(\tilde{x}, \tilde{y}, \tilde{z})}{\Gamma^{\gamma+1}} \frac{dU(\tilde{x}, \tilde{y}, \tilde{z})}{d\tilde{x}}, \quad (7.34)$$

Plugging (7.34) into (7.31), we get

$$\ddot{b}_x + \frac{1}{m b_x \langle \tilde{x}^2 \rangle} \left[\langle \tilde{x} \frac{dU(b_x \tilde{x}, b_y \tilde{y}, b_z \tilde{z})}{d\tilde{x}} \rangle - \frac{1}{\Gamma^\gamma} \langle \tilde{x} \frac{dU(\tilde{x}, \tilde{y}, \tilde{z})}{d\tilde{x}} \rangle \right] = 0, \quad (7.35)$$

with similar equations governing \tilde{y} and \tilde{z} . In principle, one could input an expression for the trap potential and solve equations (7.35) numerically to obtain the breathing mode frequency. However, further progress can be made by employing

the following two assumptions:

1. We have a cigar-shaped trap with cylindrical symmetry: $b_x = b_y$, where the trap oscillation frequencies ω_i obey the relations $\omega_x = \omega_y$ and $\omega_z \ll \omega_x, \omega_y$.
2. The method used to induce small oscillations preferentially excites radial rather than axial vibrations. In this case, the projection of the density profile in the z -direction remains nearly stationary while the radial direction oscillates.

With these two assumptions, we find $b_z = 1$ and the only remaining equation is

$$\ddot{b}_x + \frac{1}{m b_x \langle \tilde{x}^2 \rangle} \left[\langle \tilde{x} \frac{dU(b_x \tilde{x}, b_y \tilde{y}, \tilde{z})}{d\tilde{x}} \rangle - \frac{1}{b_x^{2\gamma}} \langle \tilde{x} \frac{dU(\tilde{x}, \tilde{y}, \tilde{z})}{d\tilde{x}} \rangle \right] = 0. \quad (7.36)$$

We can now study the behavior of the gas for the case of small oscillations, where the $b_i(t)$ factors deviate only slightly from their equilibrium values of 1. The procedure for determining the oscillation frequency is as follows:

1. Select a form for U and Taylor expand the potential to the desired order.
2. Make the substitution $b_x \rightarrow 1 + e_x(t)$, where $e_x(t) \ll 1$. Linearize (7.36) to obtain an equation of the form $\ddot{e}_x + \omega_{meas}^2 e_x = 0$ to extract the predicted radial oscillation frequency ω_{meas} for the chosen form of the potential.
3. Choose an appropriate equation of state. That is, choose a value of γ .

For a perfectly harmonic potential, the expectation values in (7.36) drop out, and the above procedure can be completed without choosing a form of the density distribution. For the harmonic case, (7.36) reduces to

$$\ddot{b}_x + \omega_x^2 \left(b_x - \frac{1}{b_x^{2\gamma+1}} \right) = 0. \quad (7.37)$$

Making the substitution $b_x(t) \rightarrow 1 + e_x(t)$ and linearizing (7.37), we obtain

$$\ddot{e}_x + 2(\gamma + 1)\omega_x^2 e_x = 0. \quad (7.38)$$

Consequently, for a gas confined in a cylindrically symmetric ($\omega_x = \omega_y = \omega_\perp$) harmonic trap, the oscillation frequency is

$$\omega_{meas}^2 = \omega_\gamma^2 \equiv 2(\gamma + 1)\omega_\perp^2. \quad (7.39)$$

However, we now consider a more realistic expression for the potential produced by an optical dipole trap,

$$U(x, y, z) = U_0 - \frac{U_0}{1 + \left(\frac{z}{z_0}\right)^2} \exp \left[- \left(\frac{x}{x_0}\right)^2 - \left(\frac{y}{y_0}\right)^2 \right], \quad (7.40)$$

where $x_0^2 = 2U_0/(m\omega_x^2)$ and similarly for y_0 and z_0 .³

If we Taylor expand (7.40), keep the first two orders of anharmonic corrections, and use the three-step procedure to determine the oscillation frequency, we find

$$\begin{aligned} \omega_{meas}^2 &= \omega_\gamma^2 - \omega_\gamma^2 \left[\frac{4(2 + \gamma)}{3(1 + \gamma)} \frac{\langle \tilde{x}^4 \rangle}{x_0^2 \langle \tilde{x}^2 \rangle} + \frac{\langle \tilde{x}^2 \tilde{z}^2 \rangle}{z_0^2 \langle \tilde{x}^2 \rangle} \right] \\ &\quad + \omega_\gamma^2 \left[\frac{4(3 + \gamma)}{5(1 + \gamma)} \frac{\langle \tilde{x}^6 \rangle}{x_0^4 \langle \tilde{x}^2 \rangle} + \frac{4(2 + \gamma)}{3(1 + \gamma)} \frac{\langle \tilde{x}^4 \tilde{z}^2 \rangle}{x_0^2 z_0^2 \langle \tilde{x}^2 \rangle} + \frac{\langle \tilde{x}^2 \tilde{z}^4 \rangle}{z_0^4 \langle \tilde{x}^2 \rangle} \right], \quad (7.41) \end{aligned}$$

where we have assumed cylindrical symmetry, $x_0^2 = y_0^2$. For typical trap conditions, the anharmonic corrections are small and the term in brackets on the

³The given expression for the dipole potential is only approximately correct, as different confinement in the x - and y -directions would require two Rayleigh lengths in the z -direction. However, if $x_0 \approx y_0$, then the use of one Rayleigh length in the z -direction in the expression for the dipole potential is a reasonable approximation.

second line of (7.41) can be neglected. In this case, the expression relating ω_{meas} and ω_γ reduces to

$$\omega_{meas}^2 = \omega_\gamma^2 \left\{ 1 - \left[\frac{4(2 + \gamma)}{3(1 + \gamma)} \frac{\langle \tilde{x}^4 \rangle}{x_0^2 \langle \tilde{x}^2 \rangle} + \frac{\langle \tilde{x}^2 \tilde{z}^2 \rangle}{z_0^2 \langle \tilde{x}^2 \rangle} \right] \right\}. \quad (7.42)$$

We can simplify (7.42) even further by demonstrating that the first term in brackets in (7.42) is the dominant anharmonic correction. We can estimate the values of the $\langle \xi(\tilde{x}, \tilde{z}) \rangle$ terms by assuming a normalized two-dimensional gaussian profile of the form⁴

$$n_0(\tilde{x}, \tilde{z}) = \frac{N}{\pi \sigma_x \sigma_z} \left(\frac{T}{T_F} \right)^{-1} \exp \left[-\frac{T_F}{T} \left(\frac{\tilde{x}^2}{\sigma_x^2} + \frac{\tilde{z}^2}{\sigma_z^2} \right) \right]. \quad (7.43)$$

The Fermi radii in (7.43) can be related to x_0 and z_0 in (7.42) via

$$\sigma_x^2 = \frac{\varepsilon_F}{U_0} x_0^2 \quad (7.44)$$

$$\sigma_z^2 = \frac{\varepsilon_F}{U_0} z_0^2, \quad (7.45)$$

where we have made use of the definitions of x_0 and z_0 along with (A.37). Using (7.32), (7.43), (7.44) and (7.45), we find that

$$\frac{\langle \tilde{x}^2 \tilde{z}^2 \rangle}{z_0^2 \langle \tilde{x}^2 \rangle} = \frac{1}{3} \frac{\langle \tilde{x}^4 \rangle}{x_0^2 \langle \tilde{x}^2 \rangle}. \quad (7.46)$$

In a moment, we will consider anharmonic corrections for $\gamma = 2/3$ and $\gamma = 1$. For both of these values, the lead coefficient in front of the $\langle \tilde{x}^4 \rangle / (x_0^2 \langle \tilde{x}^2 \rangle)$ term

⁴This formula is given in equation (A.50) in Section A.3.4. All of the density profiles given in Appendix A are normalized such that integration over all coordinates yields $N/2$. The lead factor in (7.43) is twice as large as that given in (A.50) to ensure that (7.43) satisfies the normalization condition given by (7.33).

in (7.42) is approximately 2. Consequently, (7.46) reveals that the second term in brackets in (7.42) is roughly 6 times smaller than the first term in brackets in (7.42). Furthermore, the second term in brackets in (7.42) has the opposite sign as the next order anharmonic correction shown in (7.41). As a result, when calculating anharmonic corrections to the measured breathing mode frequencies for a cylindrically symmetric trap, we neglect the second term in brackets in (7.42) and use

$$\begin{aligned}\omega_{meas}^2 &= \omega_\gamma^2 \left[1 - \frac{4(2+\gamma)}{3(1+\gamma)} \frac{\langle \tilde{x}^4 \rangle}{x_0^2 \langle \tilde{x}^2 \rangle} \right] \\ &= \omega_\gamma^2 \left[1 - \frac{2(2+\gamma)}{3(1+\gamma)} \frac{m \omega_\perp^2 \langle \tilde{x}^4 \rangle}{U_0 \langle \tilde{x}^2 \rangle} \right].\end{aligned}\quad (7.47)$$

To make use of (7.47), one needs to choose the form of the density profiles to determine the values of the $\langle \tilde{x}^k \rangle$ terms. Further progress can be made by choosing an analytic form of the density profile, though of course, these values can also be obtained numerically for an arbitrary choice of density profile. Assuming a zero temperature Thomas-Fermi profile for a harmonically trapped gas as given by (A.40),⁵ we find that $\langle \tilde{x}^4 \rangle / \langle \tilde{x}^2 \rangle^2 = 12/5$. Consequently, for zero temperature spatial profiles, (7.47) becomes

$$\omega_{meas}^2 = \omega_\gamma^2 \left[1 - \frac{8(2+\gamma)}{5(1+\gamma)} \frac{m \omega_\perp^2 \langle \tilde{x}^2 \rangle}{U_0} \right].\quad (7.48)$$

For classical profiles, as given by (A.51), $\langle \tilde{x}^4 \rangle / \langle \tilde{x}^2 \rangle^2 = 3$, in which case (7.47)

⁵Of course, for this calculation to be self-consistent, we should consider spatial densities related to anharmonically confined clouds. However, for small degrees of trap anharmonicity, it is a reasonable approximation to use a spatial density based on harmonic confinement to estimate the first order anharmonic correction.




Figure 7.5: The ratio $\langle \tilde{x}^4 \rangle / \langle \tilde{x}^2 \rangle^2$ (solid line) versus temperature for a harmonically trapped noninteracting Fermi gas. The ratio was calculated using the density profile given by (A.48), and is relevant for anharmonic corrections to the radial breathing mode frequency. The dashed line in the lower portion of the figure is the zero temperature limit (2.4) of the ratio, while the dot-dashed line in the upper portion marks the classical limit (3).

becomes

$$\omega_{meas}^2 = \omega_\gamma^2 \left[1 - \frac{2(2 + \gamma)}{(1 + \gamma)} \frac{m \omega_\perp^2 \langle \tilde{x}^2 \rangle}{U_0} \right]. \quad (7.49)$$

In between the zero and high temperature limits $\langle \tilde{x}^4 \rangle / \langle \tilde{x}^2 \rangle^2$ varies smoothly between 12/5 and 3, as shown in Figure 7.5. The ratio of the moments of the distribution was calculated using normalized finite temperature Thomas-Fermi profiles as given by (A.48).

Applying equations (7.47), (7.48), and (7.49) is relatively straightforward. Both U_0 and ω_\perp can be determined from parametric resonance measurements as described in Section 4.6.1, while m is the mass of a single atom. Using measured values of the oscillation frequency ω_{meas} and the $\langle \tilde{x}^k \rangle$ terms, one can solve (7.47), (7.48), (7.49) for ω_γ , which is the oscillation frequency we would measure

for a harmonic trap. In closing, we note that all of the $\langle \tilde{x}^k \rangle$ terms in (7.47), (7.48), and (7.49) are related to the trapped size of the cloud. If anharmonic corrections are being estimated based on images of expanded atom clouds, the cloud sizes must be rescaled to their trapped dimensions before applying (7.47), (7.48), or (7.49).

7.4.1 Anharmonic corrections for a unitary Fermi gas

As much of our data is taken in the strongly interacting regime, where we have a unitary or nearly-unitary system, we are interested in the value of (7.47) for a unitary gas. Unitary Fermi gases obey a polytropic equation of state in which $\gamma = 2/3$ [115]. In this case, (7.47) becomes

$$\omega_{meas}^2 = \omega_\gamma^2 \left[1 - \frac{16 m \omega_\perp^2 \langle \tilde{x}^4 \rangle}{15 U_0 \langle \tilde{x}^2 \rangle} \right]. \quad (7.50)$$

In the zero temperature limit, setting $\gamma = 2/3$ in (7.48) yields

$$\omega_{meas}^2 = \omega_\gamma^2 \left[1 - \frac{64 m \omega_\perp^2 \langle \tilde{x}^2 \rangle}{25 U_0} \right]. \quad (7.51)$$

The high temperature limit is given by setting $\gamma = 2/3$ in (7.49),

$$\omega_{meas}^2 = \omega_\gamma^2 \left[1 - \frac{16 m \omega_\perp^2 \langle \tilde{x}^2 \rangle}{5 U_0} \right]. \quad (7.52)$$

For known values of U_0 , ω_\perp , and m , measured values of ω_{meas} and $\langle \tilde{x}^k \rangle$ can be used to calculate ω_γ with the aid of (7.50), (7.51), and (7.52).

7.4.2 Anharmonic corrections for a noninteracting Fermi gas

If we return to the relaxation approximation formalism and consider the collisionless limit, we obtain [101, 116]

$$\ddot{b}_x + \frac{1}{m b_x \langle \tilde{x}^2 \rangle} \left[\langle \tilde{x} \frac{dU_{EV}(b_x \tilde{x}, b_y \tilde{y}, \tilde{z})}{d\tilde{x}} \rangle - \frac{1}{b_x^2} \langle \tilde{x} \frac{dU_{EQ}(\tilde{x}, \tilde{y}, \tilde{z})}{d\tilde{x}} \rangle \right] = 0. \quad (7.53)$$

We note that if we let $\gamma = 1$ in (7.36), we obtain (7.53). Hence, to determine the anharmonic corrections for a noninteracting gas, we can set $\gamma = 1$ in (7.47), yielding

$$\omega_{meas}^2 = \omega_{param}^2 \left[1 - \frac{m \omega_{\perp}^2 \langle \tilde{x}^4 \rangle}{U_0 \langle \tilde{x}^2 \rangle} \right]. \quad (7.54)$$

Setting $\gamma = 1$ in (7.48) to obtain the zero temperature limit yields

$$\omega_{meas}^2 = \omega_{param}^2 \left[1 - \frac{12 m \omega_{\perp}^2 \langle \tilde{x}^2 \rangle}{5 U_0} \right]. \quad (7.55)$$

In the high temperature limit, setting $\gamma = 1$ in (7.49) gives

$$\omega_{meas}^2 = \omega_{param}^2 \left[1 - \frac{3 m \omega_{\perp}^2 \langle \tilde{x}^2 \rangle}{U_0} \right]. \quad (7.56)$$

Note that for a perfectly harmonic trap, where $\omega_{meas} = \omega_{param}$, setting $\gamma = 1$ in (7.39) gives $\omega_{param} = 2 \omega_{\perp}$, which agrees with the result derived in Section 7.3 for a noninteracting gas. For anharmonic traps, values of ω_{\perp} and U_0 based on parametric resonance measurements⁶ and measured values of ω_{meas} and $\langle \tilde{x}^k \rangle$ can

⁶Parametric resonance measurements are subject to trap anharmonicity as well. Consequently, when first characterizing the dipole trap, we use values of $(\omega_x)_{anharm}$ and $(\omega_y)_{anharm}$ measured using parametric resonance to estimate the value of $\omega_{\perp} \approx [(\omega_x)_{anharm} (\omega_y)_{anharm}]^{1/2}$

be entered into (7.54), (7.55), and (7.56) to calculate ω_{param} , the oscillation frequency one would measure for a harmonic trap. In closing, we note that although we assumed cylindrical symmetry in deriving (7.47), which we use as the basis for the equations displayed in this subsection, (7.54), (7.55), and (7.56) are valid in the absence of cylindrical symmetry as well.

7.5 Magnetic field dependence of the breathing mode

Our initial measurements of the radial breathing mode frequency [19] were restricted to magnetic fields between 770 gauss and 910 gauss. We found over this range of interaction strengths that the measured radial breathing mode frequencies are in general agreement with predictions based on an equation of state for a superfluid hydrodynamic gas. Moreover, close to the center of the Feshbach resonance, we found that our measured radial breathing mode frequency was within 0.3% of the value expected for a unitarity-limited hydrodynamic Fermi gas ($\sqrt{10/3}\omega_{\perp}$). At roughly the same time, Bartenstein *et al* [21] conducted similar studies of both the radial and axial breathing mode frequencies. While they found that their axial data was in good agreement with hydrodynamic predictions, they observed a systematic shift away from hydrodynamic predictions in the radial mode. Near the center of the broad Feshbach resonance, their measured

which appears in (7.54). A measurement of the radial breathing mode oscillation frequency in the x -direction for a noninteracting gas for the same trap conditions then allows us to use (7.54) to calculate $\omega_{param}/2$, which we regard as the “true” harmonic oscillation frequency of the trap in the x -direction, ω_x . We estimate the “true” harmonic oscillation frequency of the trap in the y -direction using $\omega_y = (\omega_y)_{anharm} \times \omega_x / (\omega_x)_{anharm}$. We then use the “true” harmonic trap oscillation frequencies ω_x and ω_y to calculate the value of $\omega_{\perp} = \sqrt{\omega_x \omega_y}$ used in all subsequent applications of the equations appearing in Section 7.4.

radial breathing mode frequency was roughly 8.7% below the predicted value for a unitarity-limited hydrodynamic Fermi gas [117]. Furthermore, on the BCS side of the resonance, they observed a dramatic increase in the radial breathing mode frequency and damping rate near 910 gauss. As this magnetic field represented the limit of our initial studies of the radial breathing mode, we were motivated to perform an additional study. The results of this more extensive investigation of the radial breathing mode frequency and damping rate were published in [79]. In this latter investigation, we found that measurements conducted below 950 gauss confirmed the results of our initial study [19]. However, for magnetic fields near 1080 gauss, we observed a breakdown in hydrodynamic behavior. This breakdown was signalled by an increase in the measured breathing mode frequency and an even more abrupt increase in the damping rate of the oscillations. This general behavior is in qualitative agreement with the results of Bartenstein *et al*, although there are important quantitative differences which will be explored in a moment.

Following preparation of the gas and excitation of the breathing mode as described in Section 7.2, the lifetime and frequency of the resulting oscillations were measured. The measured frequencies were corrected for anharmonicity, and are plotted as a function of magnetic field and interaction strength in Figure 7.6. The lower horizontal scale provides the interaction strength in terms of the dimensionless parameter $1/(k_F a_s)$, where $k_F = \sqrt{2 m \varepsilon_F / \hbar^2}$ is the Fermi wave vector at the center of the trap, $\varepsilon_F = \hbar \bar{\omega} (3 N)^{1/3}$ is the Fermi energy for a noninteracting gas, and a_s is the s-wave scattering length. The upper horizontal scale gives the magnetic field at which the data was acquired. Note that in converting from magnetic field to $1/(k_F a_s)$, we assume that the Feshbach resonance is centered at 834 gauss. The solid curve in Figure 7.6 provides a theoretical prediction based on superfluid




Figure 7.6: Magnetic field dependence of the frequency ω of the radial breathing mode, normalized to the geometric mean of the trap oscillation frequencies in the two radial dimensions, ω_{\perp} . The solid line is the theory based on superfluid hydrodynamics from Hu *et al* [118]. The dashed horizontal line represents the predicted hydrodynamic frequency for a unitary Fermi gas, $\omega/\omega_{\perp} = \sqrt{10/3}$. The lower horizontal axis parameterizes the interaction strength using the dimensionless parameter $1/(k_F a_s)$, where k_F is the Fermi wave vector at the center of the trap and a_s is the s-wave scattering length. The upper horizontal axis gives the associated magnetic field in gauss. Note that the magnetic field axis is not linear. Figure first published in [79].

hydrodynamics given by Hu *et al* in [118].

In Figure 7.7, we display both the frequency and damping time dependence on interaction strength. Once again, the dimensionless parameter $1/(k_F a_s)$ parameterizes the interaction strength on the lower horizontal axis, while the magnetic field is given on the upper horizontal axis. The frequency data shown in Figure 7.6 is given once more, with the frequency axis labelled at the left of the figure. The right vertical axis gives the damping rate in dimensionless units, $1/(\tau \omega_{\perp})$.

For magnetic fields close to resonance, the measured radial breathing mode

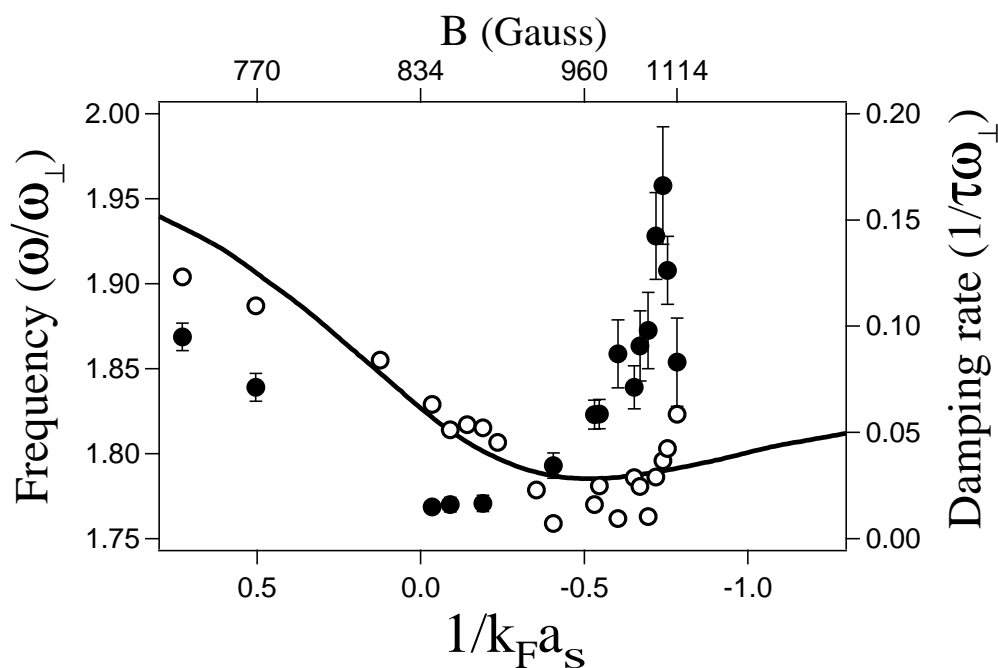


Figure 7.7: Magnetic field dependence of the damping rate $1/(\tau\omega_{\perp})$ and frequency ω/ω_{\perp} of the radial breathing mode. The damping rate data is represented by the solid circles, and the axis is labelled on the right hand side of the graph. The frequency data is represented by the open circles (error bars have been omitted to avoid clutter), with the axis labelled on the left hand side of the graph. The solid line represents the predicted frequency of the radial breathing mode by Hu *et al* [118]. The frequency and damping time are plotted against interaction strength in both the dimensionless parameter $1/(k_F a_s)$ (bottom axis) and the magnetic field (top axis). Note that the top (magnetic field) axis is not linear. Figure first published in [79].

frequencies are in good agreement with the hydrodynamic predictions of Hu *et al*, confirming the measurements first made in [19]. The measured breathing mode frequency just above the center of the Feshbach resonance, $\omega/\omega_{\perp} = 1.829(0.006)$, is in good agreement with the value expected for a unitary, hydrodynamic Fermi gas, $\sqrt{10/3} = 1.826$. In Figure 7.7, we see that the damping rates are quite small close to the center of the resonance, where damping times of up to 7 ms were observed. This corresponds to $1/(\tau\omega_{\perp}) = 0.014$, or roughly 20 periods of oscillation.

As we move away from the center of the broad Feshbach resonance, the data remain in qualitative agreement with the predictions of Hu *et al* [118], but quantitative discrepancies begin to appear. Data taken at 750 and 770 gauss are close to two standard deviations below the predicted frequencies. Above the resonance, the observed frequency dips below the predicted value before rising above the predicted values. The lowest measured frequency occurred at 925 gauss, before rising rapidly around 1080 gauss. The highest magnetic field for which the breathing mode was studied, 1114 gauss, represents the largest magnitude field we can produce with our magnets. At this field, the measured breathing mode frequency is above the predictions of superfluid hydrodynamics but still well below the expected frequency for a noninteracting gas, $2\omega_x$. This is in contrast to the measurements of Bartenstein *et al* [21], where the measured oscillation frequency at high magnetic field was found to exceed predictions for a noninteracting system.

The observed upward shift in the radial breathing mode frequency is indicative of a breakdown in hydrodynamic behavior which is seen more clearly in the damping rate as a function of magnetic field, as shown in Figure 7.7. The abrupt

increase in the damping rate at 1080 gauss is in qualitative agreement with similar behavior reported in [21]. However, there are two ways in which our data differ from Bartenstein *et al*:

1. Near resonance, we observe good agreement between our predictions and hydrodynamic theory, in contrast to [21], where the measured frequencies are systematically lower than predictions [117].
2. We observe a breakdown of hydrodynamics at a magnetic field near 1080 gauss, significantly larger than the magnetic field of 910 gauss at which Bartenstein *et al* observed similar behavior.

As the abrupt increase in frequency and damping rate above resonance was an unexpected observation, there was and is interest in explaining the source of the behavior. Bartenstein *et al* have proposed that the breakdown of hydrodynamic behavior occurs when the binding energy of the atom pairs (the zero temperature BCS energy gap Δ) becomes smaller than the quantized collective mode energy $\hbar\omega$. Using a crude estimate based on BCS theory, we find that the trap-averaged gap $\bar{\Delta} \approx 2\hbar\omega$ at 910 gauss for Bartenstein *et al* and $\bar{\Delta} \approx \hbar\omega$ at 1080 for the conditions in our experiment. We would expect single-particle excitations associated with pair breaking to occur at $2\Delta = \hbar\omega$, so it is unusual that we have $\Delta \simeq \hbar\omega$ for our trap conditions at the magnetic field where hydrodynamics breaks down. However, Heiselberg has noted that the leakage of energy to surface modes could account for the rapid decay of the breathing mode oscillations when $\Delta > \hbar\omega$ [119]. Unfortunately, this argument cannot explain why we observe the breakdown in hydrodynamics at a different value of the quantized collective mode energy from that seen by Bartenstein *et al*. Finally, it is worth noting that the

application of predictions from BCS theory, a theory based on weak interactions, should be used with skepticism in a regime where the gas is strongly interacting.

Another estimate of the breakdown magnetic field comes from Falco and Stoof [120], who suggested that the gas will behave like a BCS system when the energy of the molecular state which gives rise to the Feshbach resonance exceeds the Fermi energy of a two atom pair. In such a case, the small but finite temperature of our gas might be enough to destroy a superfluid state. Setting the approximate energy of a bound atom pair equal to twice the trap-averaged value of the local Fermi energy,

$$\frac{\hbar^2}{m a_s^2} = 2 \langle \epsilon_F(\mathbf{x}) \rangle, \quad (7.57)$$

we can estimate the breakdown magnetic field. The left hand side of (7.57) represents the binding energy of an atom pair relative to the zero of energy in the triplet potential. For our trap conditions, (7.57) gives $1/(k_F a_s) = -0.79$, which is quite close to the value $1/(k_F a_s) = -0.74$ for our trap conditions at 1080 gauss, the magnetic field at which we observe a breakdown in hydrodynamic behavior. While this is an appealing explanation for the behavior we observe in our system, it is unable to explain the results of Bartenstein *et al*, who observed the breakdown of hydrodynamics at the same magnetic field for several different trap depths. Ultimately, our simple models of pair-breaking are unable to explain the discrepancy in the magnetic fields at which the breakdown of hydrodynamics is observed by Bartenstein *et al* and by the Duke group.

7.6 Temperature dependence of the breathing mode at unitarity

We conducted a systematic study of the frequency and damping time of the radial breathing mode of a strongly interacting gas of ${}^6\text{Li}$ just above the center of a broad Feshbach resonance. Here, the zero energy scattering length is effectively infinite and the gas is in the unitary regime. This measurement constitutes the first systematic study of the temperature dependence of the radial breathing mode in a strongly interacting Fermi gas. For reasons outlined in Section 6.3, temperature measurement in the strongly interacting regime remains controversial. Throughout this section, we will present the behavior of the radial breathing mode as a function of the empirical temperature \tilde{T} . Of course, this temperature scale can be calibrated to theoretical temperature scales using the technique discussed in Section 6.3.2, but it is not necessary here.

Our initial investigation of the temperature dependence of a strongly interacting Fermi gas was given in [19], where we found that the lifetime of the radial breathing mode oscillation increased as we lowered the temperature of the gas. In that study, conducted at 870 gauss, different temperatures were achieved by stopping the evaporative cooling process prior to reaching the lowest possible temperature, resulting in different atom numbers for different temperature conditions. Furthermore, our understanding of temperature measurement in the strongly interacting regime was somewhat limited, and so the temperatures reported in [19] show the trend in temperature, but might not be reliable estimates of the absolute temperature. Recall that in Section 7.2, we reported that the duration of the release and recapture sequence to excite the breathing mode was $25 \mu\text{s}$. This is

true for the data reported in [20], but not for our initial study of the temperature dependence of the radial breathing mode reported in [19]. In the initial study, an excitation time of $50 \mu\text{s}$ was used at 4.6% of the maximum trap depth, which limited our maximum oscillation lifetime to $3.85(0.40)$ ms. Even with these limitations, this initial study of the damping time yielded evidence for a superfluid state in the low temperature regime.

Evidence for superfluidity arose in the form of the observation of a hydrodynamic oscillation frequency at 870 gauss. Of course, such behavior can arise from collisional processes as well. However, as the temperature of the gas is lowered, Pauli blocking should begin to suppress collisional processes. The observation of hydrodynamic behavior, in a regime where collisional hydrodynamics is an unlikely explanation, strongly suggests superfluid behavior. Attempts to fit the measured breathing mode lifetimes with predictions based on two-body Pauli blocking were unsuccessful, and the momentum relaxation time required to explain the $3.85(0.40)$ ms lifetime at the lowest temperature data point was very small. As we saw in Section 7.3, in the hydrodynamic limit where the relaxation time $\tau_R \rightarrow 0$, the relaxation approximation model predicts hydrodynamic breathing mode oscillation frequencies ω_{hydro} with an infinite lifetime. If we consider a first order correction to this result, we find that in the hydrodynamic limit, two-body collision models [116, 121, 122] predict that the lifetime of the mode is inversely proportional to the relaxation time,

$$\frac{1}{\tau} \propto \omega_{hydro}^2 \tau_R. \quad (7.58)$$

However, in the low temperature limit, collisional models which incorporate Pauli

blocking [116,123] predict the following relationship between relaxation time τ_R and temperature

$$\frac{1}{\omega_{\perp} \tau_R} \propto \left(\frac{T}{T_F} \right)^2. \quad (7.59)$$

Combining (7.58) and (7.59), we see that two-body collisional models in the hydrodynamic limit predict that $\tau \rightarrow 0$ as $T/T_F \rightarrow 0$. This is precisely the opposite of the observed relationship between τ and T/T_F in the low temperature regime, where the lifetime of hydrodynamic oscillations *increases* as the temperature is lowered. As collisional hydrodynamics was unable to explain the increased lifetime of the breathing mode oscillations as the temperature of the gas was lowered, it was seen as inconsistent with the data.

We also considered the possibility that the observed behavior arose from a collisionless mean-field scenario. Such a scenario is easily rejected based on the observed oscillation frequency. While the data agree almost precisely with the expected hydrodynamic frequency, which is well below the noninteracting frequency $2\omega_x$, a collisionless mean field scenario based on the Vlasov equation [108] leads to a breathing mode frequency above the noninteracting frequency [19].

To summarize, we find that collisional hydrodynamics can explain the observed frequency, but not the observed oscillation lifetimes at low temperature. A collisionless mean field scenario does not agree with the observed oscillation frequency. Having rejected both collisional hydrodynamics and collisionless mean field scenarios as unlikely explanations for the observed behavior, we concluded that the observation of increased breathing mode oscillation lifetimes as the gas was cooled was solid evidence for superfluidity, if not a definitive proof.

While this study presented exciting evidence for superfluid hydrodynamics, it was not without its imperfections. The method used to produce gases at different

temperatures was not ideal, as the stoppage of evaporative cooling meant that gases at different temperatures had a different number of atoms. If the damping rate depended on the Fermi energy of the system (which in turn depends on the number of atoms in the gas), then these fluctuations in atom number could confound the interpretation of the data. A more reliable method of producing gases at different temperatures was required. Furthermore, as already noted, our understanding of temperature was somewhat limited at the time we conducted our first study of the temperature dependence of the radial breathing mode.

Fortunately, the methods discussed in Chapter 6 in our study of the heat capacity allowed us to conduct a more definitive investigation of the temperature dependence of the radial breathing mode. The results of this effort were reported in [20]. In this study, we evaporatively cooled the gas at 840 gauss, just above the center of the broad Feshbach resonance, to the lowest possible temperature, \tilde{T} . Using the energy input techniques discussed in Section 6.1, we were then able to produce a gas at arbitrary temperature using various values of t_{heat} . After allowing 0.1 seconds for the gas to reach equilibrium, the breathing mode was then excited using a 25 μs excitation time. A schematic of this process is shown in Figure 7.8. This technique for producing gases at arbitrary temperatures is superior to the premature stoppage of evaporative cooling, as it allows for fairly constant atom number even as the temperature of the cloud is varied. Furthermore, as we have very good control over t_{heat} , this method permitted us to take higher resolution damping time versus temperature data. Even with these improved techniques, our study of the temperature dependence of the damping time of the radial breathing mode took several months, and more than 6300 repetitions of the experimental cycle. The most important results of these labors are shown in Figures 7.9, 7.10




Figure 7.8: The schematic shows the technique used to study the temperature dependence of the radial breathing mode. Dark regions indicate times during which the optical trap is on, while blank regions indicate that the trap is off. Following an energy input duration of t_{heat} , the gas is allowed to equilibrate. The breathing mode is then excited during t_{excite} , after which the gas oscillates for a time t_{hold} in the trap before being released for time of flight imaging. Note that the schematic is not to scale.

and 7.11.

The data presented in Figure 7.9 was taken at 840 gauss at 4.6% of the maximum trap depth and atom number $N = 2.0(0.2) \times 10^5$. We see that the measured oscillation frequency varies relatively smoothly as a function of empirical temperature \tilde{T} . In the Figure, the open circles represent the measured oscillation frequency prior to applying anharmonic corrections. The black dots show the corresponding frequencies following anharmonic corrections where the ratio of the spatial moments ($\langle x^4 \rangle / \langle x^2 \rangle$, for example) in (7.50) were estimated using finite temperature Thomas-Fermi fits to the data. While the anharmonicity-corrected data points in Figure 7.9 do not have error bars, the frequency error arising from uncertainty in the correction is estimated to be comparable to the statistical error of the uncorrected data points. Over the range of temperatures studied, we find that the measured oscillation frequency remains close to the hydrodynamic value (represented by the dot-dashed line in Figure 7.9) and is substantially less than




Figure 7.9: Frequency of the radial breathing mode for a unitary Fermi gas versus empirical temperature \tilde{T} . The open circles represent the measured frequencies, while the black dots represent the data following correction for anharmonicity using finite temperature Thomas-Fermi profiles. The dot-dashed line is the unitary hydrodynamic frequency $\omega/\omega_{\perp} = \sqrt{10/3}$. The dashed line at the top of the graph is the frequency $2\omega_x$ observed for a noninteracting gas in the low temperature regime. Figure first published in [20].

the value expected for a noninteracting gas, $2\omega_x = 2.10\omega_\perp$, shown in Figure 7.9 as the dashed line. The closeness of the breathing mode frequency to the hydrodynamic frequency could be explainable in terms of universal hydrodynamics for a system governed by isentropic conditions [85].

While the measured oscillation frequency varies smoothly as a function of empirical temperature, we observe much richer behavior from the temperature dependence of the damping rate, shown in Figure 7.10. For $\tilde{T} < 0.50$, the damping rate varies linearly with \tilde{T} . A linear fit for this temperature range yields

$$\frac{1}{\tau\omega_\perp} = 0.146(0.004)\tilde{T} - 0.0015(0.0014). \quad (7.60)$$

We note that the damping rate extrapolates close to zero at zero temperature, consistent with zero temperature expectations for a superfluid system. As with our initial study of the temperature dependence of the radial breathing mode [19], the observation of increased oscillation lifetimes in the very low temperature regime is inconsistent with collisional hydrodynamics and is consistent with superfluid behavior.

Above $\tilde{T} = 0.50$, the damping rate departs strongly from linear scaling. While we do not have a satisfactory theory of the damping rate as a function of temperature, we suggest that the observed departure from linear scaling could signal a superfluid transition. Using the “natural temperature” conversion given by (6.31), we find that $\tilde{T} = 0.50$ is associated with $T/T_F = 0.35$, which is reasonably close to the measured transition temperature ($T/T_F = 0.27$) in the heat capacity study of a unitary Fermi gas. As the measurements of the heat capacity and the breathing mode lifetime are quite different, it is not apparent that critical transi-




Figure 7.10: Temperature dependence of the damping rate for the radial breathing mode of a unitary gas of ${}^6\text{Li}$ at 840 gauss. The solid dots are the main data set taken at 4.6% of the maximum trap depth and $N = 2.0(0.2) \times 10^5$. The dashed line is Equation (7.60), which extrapolates close to zero at zero temperature. Figure first published in [20].

tion signatures should be observed at precisely the same temperature. Finally, it is worth noting once more that the interpretation of the damping rate's departure from linear scaling with \tilde{T} is speculative, and not based on theoretical insights.

Near $\tilde{T} = 0.60$, there appears to be a reproducible notch in the damping rate. The three data points in the notch fall between 2.2 and 4.1 standard deviations below the linear extrapolation at $\tilde{T} = 0.60$. A similar feature was predicted for a two-component fluid composed of a Bose-Einstein condensate and a thermal cloud, where beating of the oscillation frequencies of the different components leads to an apparent faster decay of the oscillation amplitude [124]. We have monitored the decay of the mode for durations exceeding several decay time constants, but find no evidence of the revival of the oscillations. At present, we do not have a satisfactory explanation for the appearance of the notch in the damping at $\tilde{T} = 0.60$.

For temperatures above the notch, $0.65 \leq \tilde{T} \leq 1.0$ (or, $0.45 \leq T/T_F \leq 0.71$, using (6.31)), the damping rate appears to be nearly independent of temperature. This is followed by an increase in the damping rate between $1.0 \leq \tilde{T} \leq 1.2$ (or $0.71 \leq T/T_F \leq 0.86$). Above $\tilde{T} = 1.2$, the damping rate appears to be nearly independent of temperature.

We have considered several possible explanations for the behavior of the damping rate versus temperature for $\tilde{T} > 1.0$. We find that predictions for a normal Fermi gas undergoing binary collisions [122] should exhibit slow variations (on the order of T_F) in the damping for temperatures above $T/T_F = 0.70$, in contrast to the variations we observe which occur over a fraction of the Fermi temperature. Moreover, the binary collision model predicts a decrease in damping in this temperature regime, whereas we observe an increase in damping followed by weak

temperature dependence above $\tilde{T} > 1.2$. Finally, while we find that the oscillation frequency remains close to the hydrodynamic value throughout the range of temperatures studied, the binary collision model predicts a frequency close to the noninteracting gas value for $T/T_F > 0.70$. Interestingly, the maximum observed damping time is relatively consistent with the maximum damping time predicted by the binary collision model for our trap conditions, $1/(\tau \omega_\perp)_{max} = 0.13$ (for a trap with perfect cylindrical symmetry, $1/(\tau \omega_\perp)_{max} = 1/\sqrt{120} \simeq 0.09$). In spite of the apparent agreement on the maximum damping time, we observe too many discrepancies between the predictions of the binary collision model and our data to consider it a likely explanation for the observed high temperature behavior.

Another possible explanation for the increase in damping in the high temperature regime is trap anharmonicity, which increases the bandwidth. However, if trap anharmonicity were the dominant cause of damping, $1/\tau$ would be proportional to the frequency correction. That is,

$$\frac{1}{\tau \omega_\perp} \propto \frac{m \omega_\perp^2 \langle x^2 \rangle}{U_0} \propto \frac{k_B T}{U_0}, \quad (7.61)$$

in which case we would expect the damping rate to rise rapidly and monotonically with increasing temperature, in contrast to our observations.

As a final possible explanation for the increase in damping rate for $1.0 \leq \tilde{T} \leq 1.2$, we consider the breaking of noncondensed pairs. Recall from Section 2.5 that pseudogap theory, in which pairing of particles and the condensation of those particles can occur at different temperatures, might be applicable to strongly interacting Fermi gases. We considered pair breaking as a possible mechanism for the observed increase in damping rates for low temperature gases at magnetic

fields above the center of the Feshbach resonance in Section 7.5. The same basic argument can be applied here. We find that the region of increasing damping of the breathing mode corresponds roughly to the temperature at which noncondensed pairs are expected to vanish [10]. A prediction from pseudogap theory for the trap-averaged gap for a unitary Fermi system indicates that $\bar{\Delta} \leq \hbar\omega$ for $T \geq 0.75 T_F$, or $\tilde{T} \geq 1.06$ [125].

We have also examined the dependence of the damping rate in the low temperature regime ($\tilde{T} \leq 0.50$) on the trap oscillation frequency ω_{\perp} and the number of atoms N . In the unitary regime, dimensional analysis restricts the possible dependence to the following form,

$$\frac{1}{\tau} = \omega_{\perp} f(T/T_F, N, \lambda), \quad (7.62)$$

where $\lambda = \omega_z/\omega_{\perp}$ is the ratio of the axial to the radial trap oscillation frequencies, and f is a dimensionless function. Experimentally, we are unable to determine the dependence of the damping time on λ , but we can consider the effect of different trapping frequencies, atom number and temperature. The results of these studies are shown in Figure 7.11. We find that decreasing the number of atoms in the cloud by a factor of three has no apparent impact on the lifetime of the oscillations. Furthermore, the data in Figure 7.11 indicate that $1/\tau$ versus \tilde{T} scales approximately as ω_{\perp} when the trap is at 0.85% and 19% of the maximum trap depth. Ultimately, we can restrict the damping rate dependence given by (7.62) even further,

$$\frac{1}{\tau} = \omega_{\perp} f(T/T_F, \lambda), \quad (7.63)$$

wherein the damping rate can depend on the number only via the Fermi temper-




Figure 7.11: Damping rate $1/(\tau\omega_{\perp})$ for the system with scaled parameters. The two solid squares represent data acquired at 0.85% of the maximum trap depth. The four open diamonds are associated with data acquired at 19% of the maximum trap depth. The two open circles are associated with data acquired for atom numbers roughly 3 times smaller than that for the main data set. The dashed line is Equation (7.60), the best fit to the low temperature data for the main data set. Figure first published in [20].

ature T_F .

In closing, we show the locus plot of the frequency versus damping rate for the main data set taken at 4.6% of the maximum trap depth in Figure 7.12. Here, we see reasonable agreement between the theory (solid curve) developed using the relaxation approximation, and the data. For high damping rates, however, the agreement between data and theory worsens. In spite of the reasonable agreement for small damping rates, this plot should not be construed as validation of the binary collision model, as we have effectively disregarded temperature information in producing this plot. We note merely that if we consider only the frequency and damping rate of the oscillations, there is reasonable agreement between the data




Figure 7.12: Locus plot of the damping rate and frequency for a unitary Fermi gas. The dots represent data following correction for anharmonic effects, while the solid line is the theoretical prediction provided by the relaxation approximation formalism discussed in Section 7.3. The solid line is the same as that shown in Figure 7.4. The dashed line in the lower portion of the graph shows the hydrodynamic breathing mode oscillation frequency, while the dot-dashed line in the upper portion of the graph marks the collisionless (noninteracting) limit.

and the relaxation approximation prediction.

7.7 Conclusion

Studies of the breathing mode in a strongly interacting Fermi gas provide a convenient method for testing the equation of state of the system. In this chapter, we presented our studies of both the magnetic field and temperature dependence of the radial breathing mode in a strongly interacting Fermi gas of ${}^6\text{Li}$ atoms. Our initial study of the temperature dependence of the lifetime of the oscillations [19]

was inconsistent with both collisionless mean-field and collisional hydrodynamic scenarios. Consequently, we interpreted our results as evidence for superfluid hydrodynamics. This study also featured a limited investigation of the radial breathing mode as a function of magnetic field, and our results were generally in good agreement with predictions based on superfluid hydrodynamics.

In response to similar studies carried out at the University of Innsbruck [21], we conducted a more extensive study of the breathing mode as a function of magnetic field [79]. In this study, we replicated the results of our previous investigations, and observed an unexpected increase in the breathing mode frequency and damping rate at a magnetic field above the center of the Feshbach resonance. While these observations were in qualitative agreement with similar observations first made at Innsbruck, we have been unable to explain the discrepancy in the magnetic fields at which these phenomena were observed. The Innsbruck group has postulated that the mechanism behind the increase in oscillation frequency and damping rate involves the breaking of paired atoms, an idea we also consider plausible.

Our final major study of the radial breathing mode in a strongly interacting Fermi gas focused on a high resolution measurement of the frequency and damping time in the unitarity limit [20]. We found that the measured frequency varied smoothly as a function of temperature, and remained relatively close to the predicted hydrodynamic frequency. In contrast, the damping time displayed very rich behavior as a function of empirical temperature \tilde{T} . The damping rate was found to vary linearly with \tilde{T} in the low temperature regime, before departing abruptly from linear scaling. Following a downward notch in the damping rate at $\tilde{T} = 0.60$, the damping rate was nearly independent of temperature before rising

once more between $1.0 \leq \tilde{T} \leq 1.2$. We suggest that the initial departure from linear scaling near $\tilde{T} = 0.50$ ($T/T_F = 0.35$) could be an indicator of a superfluid phase transition, as the abrupt change in behavior occurs at a temperature close to the transition in the heat capacity (see Chapter 6). Finally, we consider it plausible that the increase in damping rate for $1.0 \leq \tilde{T} \leq 1.2$ is indicative of the breaking of noncondensed pairs.

Chapter 8

Conclusion

This dissertation has described experimental investigations of a strongly interacting gas of fermionic ${}^6\text{Li}$ atoms. Studies of the heat capacity and radial breathing mode of a strongly interacting Fermi gas were made possible by upgrades to an existing cooling and trapping apparatus [22–24,93] and computer control system. The heat capacity study is the first measurement of a thermodynamic quantity in a strongly interacting Fermi gas. We observe a transition in behavior in the energy versus temperature near $T/T_F = 0.27(0.02)$, which is interpreted as the onset of superfluidity. This represents the first direct measurement of the superfluid transition temperature in the strongly interacting regime. The measured energy versus temperature dependence is in good quantitative agreement with predictions provided by a pseudogap theory developed by Kathy Levin’s group at the University of Chicago.

Our studies of the radial breathing mode included efforts to determine the magnetic field dependence and temperature dependence of the breathing mode oscillation frequency and damping rate. Our investigation of the magnetic field dependence revealed an unexpected increase in the damping rate and oscillation frequency above the center of the Feshbach resonance. This effect, first observed in [21], is interpreted as a possible signature of the breaking of fermionic atom

pairs. Our study of the breathing mode in the unitarity limit revealed that while the oscillation frequency depended weakly on the temperature of the gas, the damping rate exhibited several noteworthy features. As the temperature of the gas was lowered, the lifetime of hydrodynamic breathing mode oscillations increased, providing evidence for a superfluid state at low temperature. Abrupt changes in the breathing mode oscillation lifetime have been interpreted as possible indicators of a superfluid phase transition and the breaking of noncondensed atom pairs.

In the remaining sections of this chapter, I will provide a brief summary of the preceding chapters before considering anticipated upgrades to the experimental apparatus and possible future lines of study in the field of strongly interacting Fermi gases.

8.1 Chapter summary

Chapter 1 began by discussing the motivation for studying strongly interacting Fermi gases in a controlled laboratory setting. Parallels between our system and high temperature superconductors, the quark-gluon plasma and neutron stars make our system an interesting object of study, as the other systems are not as easily manipulated and studied as our own. Also in Chapter 1, I outlined the significance of the work presented in this dissertation. This included a discussion of my efforts with regard to upgrading the experimental apparatus as well as the significance of the primary results of our experimental efforts. Chapter 1 concluded with an outline of this dissertation.

Chapter 2 introduced the concepts of the BEC-BCS crossover and Feshbach resonances. In the BEC limit, weakly repulsive interactions can lead to bind-

ing of two fermions to form a molecule. In the BCS limit, weak, attractive interactions can lead to the formation of Cooper pairs. In between, the system evolves smoothly between the two weakly interacting limits to a strongly interacting regime where the nature of the interparticle interactions (attractive versus repulsive) flips sign abruptly. Studies of the BEC-BCS crossover in atomic gases are made possible through the use of Feshbach resonances, in which the energy of two colliding particles can be Zeeman tuned into resonance with a bound state in a closed scattering channel. After discussing Feshbach resonances, Chapter 2 then focused on the particular hyperfine ground states of ^6Li which we trap and cool in our laboratory. Finally, a brief overview of the major experimental results from the past few years in the field of strongly interacting Fermi gases was included.

Chapter 3 covered the basic experimental techniques we employ to produce degenerate, strongly interacting Fermi gases. The operation of the magneto-optical trap and far-off resonance dipole trap were considered. This was followed by a general description of the experimental apparatus.

To extract information from our experimental system, we acquire on-resonance absorption images of the atom clouds following expansion from the optical trap. Chapter 4 considered the dynamics of an atom cloud expanding by ballistic or hydrodynamic scale transformations. As complete understanding of the expansion dynamics of the gas requires knowledge of the optical trap oscillation frequencies and magnetic field curvature, Chapter 4 also considered how these quantities are measured. Once the cloud has expanded to a size much larger than the imaging system resolution, an on-resonance probe pulse illuminates the cloud. Chapter 5 addressed the acquisition and processing of these absorption images.

Chapter 6 presented our study of the heat capacity of a unitary Fermi gas.

Such a study was made possible by the development of novel energy input and temperature measurement techniques. Using these new techniques, we measured the heat capacity of a strongly interacting and noninteracting Fermi gas, and found good agreement between our results and theoretical predictions. A collaboration with a theory group at the University of Chicago led by Kathy Levin helped us interpret our results as a signature of a phase transition in the unitary Fermi gas. Our study of the heat capacity also yielded data which was used to demonstrate the virial theorem in the unitarity limit.

Chapter 7 presented our extensive studies of the radial breathing mode in a strongly interacting Fermi gas. Using techniques similar to those used in the study of the heat capacity, we were able to excite the breathing mode and monitor the lifetime and frequency of the oscillations. The breathing mode was studied as a function of magnetic field as well as temperature. Some of the results of our investigations have been interpreted as evidence for superfluidity and the breaking of condensed and noncondensed pairs in the unitarity limit.

The present chapter provides an overview of the work discussed in this dissertation, while also contemplating some anticipated upgrades to the experimental apparatus and future research directions.

There are three appendices in this dissertation. Appendix A presents a number of basic theoretical results for noninteracting, harmonically trapped Fermi gases. While none of the material contained here is particularly novel, this appendix is intended as a useful resource for other group members. Appendix B considers issues related to the production of a resonant imaging probe pulse at nonzero magnetic field for the quantum states of interest. Finally, Appendix C discusses many of the custom software upgrades which have been implemented during my

time as a member of the research group.

8.2 Anticipated upgrades to the experimental apparatus

The Coherent 699-21 dye laser which generates our slowing, MOT, and probe beams has been in use since 1982. In the intervening decades, it has performed admirably in many important experiments. Of course, as dye lasers are notoriously finicky instruments, our 699 dye laser has also caused graduate students, post doctoral researchers, and tenured professors to launch into epic expletive-laced tirades. Dye laser failures and the attendant swearing episodes have grown more frequent in recent years, as both optical elements in the laser cavity and electronic components in the dye laser's control box have given the last full measure in the name of science. Fortunately, help is on the way in the form of a new Coherent 899-21 dye laser. We anticipate that the new 899 dye laser can be made to perform as well or better than its predecessor, and be more reliable as well.

Further, we expect to expand the radio-frequency capabilities of our experimental system. At present, we use only noisy RF excitations to balance spin state populations after loading the FORT. An upgrade to our RF antenna and electronics arrangement should allow for more complex RF manipulation of the strongly interacting gas, such as transitions from the $|1\rangle$ and $|2\rangle$ states to higher energy hyperfine ground states.

Finally, over the past several years, construction of a new cooling and trapping apparatus has been underway in a laboratory adjacent to the one where the work described in this dissertation was performed. At present, construction is nearing

completion, and experimental investigations will begin shortly.

8.3 Outlook

It might be tempting to think that with the proof of superfluidity now in hand, the field of strongly interacting Fermi gases will become stagnant. I consider this prospect fairly unlikely. Even prior to the demonstration of vortices in a strongly interacting Fermi gas [126], many in the field felt that the existing data provided sufficient evidence for the existence of a superfluid state at low temperature in the strongly interacting regime. The observation of vortices confirmed what many already believed. But for all the focus that has been placed on resonance superfluidity, there is still much to be learned about strongly interacting fermionic matter in general. Clear maps of the phase diagram in the crossover region still remain to be determined. Once the superfluid transition temperature is clearly established, researchers can work at temperatures above the transition, where collisional hydrodynamics is believed to govern the system. Such studies would be of great interest to quark-gluon plasma theorists, for example.

Other possible lines of research have been mentioned in Section 2.5. These include investigations of higher angular momentum scattering [84, 89, 90], mismatched Fermi surfaces [43, 88], and perhaps most prominently, Fermi systems in reduced dimensions and lattices. A number of groups, theoretical as well as experimental, have already begun to study one-, two-, and three-dimensional optical lattices containing fermions. Their findings and predictions are too numerous to mention here, but such systems are expected to exhibit behavior that sometimes differs markedly from the three-dimensional bulk gases investigated in this dis-

sertation. With the proof of superfluidity no longer in question, many research groups are free to pursue divergent lines of inquiry into the behavior of strongly interacting Fermi gases.

Appendix A

Harmonically trapped Fermi gases

Although the title of this dissertation suggests that we deal exclusively with strongly interacting Fermi gases, we rely heavily on concepts drawn from the area of noninteracting Fermi gases. As the noninteracting Fermi gas is well understood, and as we are capable of producing noninteracting Fermi gases in our laboratory, the ideal (noninteracting) Fermi system provides a convenient method for testing our experimental procedures before applying them to the less well-understood strongly interacting system. Furthermore, in Chapter 6, we present an approximate temperature measurement scheme relying on noninteracting Fermi gas results which is applied to strongly interacting systems. Some of the physical quantities of interest for noninteracting Fermi gases are derived in this chapter.

Throughout this Appendix, and the rest of this dissertation as well, I assume that the local density approximation [33] is valid. In this case, we suppose that our trapping potential varies smoothly and slowly and that a small volume of our trapped atom cloud contains enough atoms for that small volume to act as a homogeneous Fermi system. While these assumptions will break down at the edge of our atom cloud, much of our experimental work appears to support the validity of the local density approximation for our trap conditions. For a

more rigorous mathematical statement of the conditions associated with the local density approximation, the reader is encouraged to examine [33].

Before considering the derivation of physical quantities such as the density of states, chemical potential, and spatial profiles of harmonically trapped noninteracting Fermi gases, we consider the Sommerfeld expansion, which will be useful in deriving the low temperature behavior of many of these quantities.

A.1 Sommerfeld expansion

In the remainder of this Appendix, we will often be interested in analytic expressions for physical quantities in the low temperature regime. Such expressions can be obtained using the Sommerfeld expansion [127].

We begin with a generic integral of the form

$$I = \int_0^{\infty} d\epsilon H(\epsilon) f(\epsilon), \quad (\text{A.1})$$

where $H(\epsilon)$ is a general function of the energy and $f(\epsilon)$ is the Fermi occupation number for energy ϵ [127],

$$f(\epsilon) = \frac{1}{\exp\left(\frac{\epsilon - \mu}{k_B T}\right) + 1}. \quad (\text{A.2})$$

In the preceding equation, μ is the chemical potential, k_B is Boltzmann's constant, and T is the temperature. Next, we define

$$K(\epsilon) = \int_0^{\epsilon} d\bar{\epsilon} H(\bar{\epsilon}), \quad (\text{A.3})$$

which allows us to rewrite (A.1) as

$$I = \int_0^\infty d\epsilon \frac{dK(\epsilon)}{d\epsilon} f(\epsilon). \quad (\text{A.4})$$

If we integrate (A.4) by parts and suppose that the function $H(\epsilon)$ has properties such that the surface term vanishes, the integral I can be written as

$$I = - \int_0^\infty d\epsilon K(\epsilon) \frac{df(\epsilon)}{d\epsilon}. \quad (\text{A.5})$$

We now consider a Taylor expansion of the function $K(\epsilon)$ about the point $\epsilon = \mu$, where μ represents the chemical potential. Then, we can approximate $K(\epsilon)$ as

$$K(\epsilon) \simeq K(\mu) + (\epsilon - \mu)K'(\mu) + \frac{(\epsilon - \mu)^2}{2}K''(\mu), \quad (\text{A.6})$$

where primes denote derivatives with respect to ϵ . Plugging (A.6) into (A.5) and making the substitution $y = (\epsilon - \mu)/(k_B T)$, the approximate value of the integral I becomes

$$\begin{aligned} I \simeq & K(\mu) \int_{-\mu/(k_B T)}^\infty \frac{e^y dy}{(e^y + 1)^2} + k_B T K'(\mu) \int_{-\mu/(k_B T)}^\infty \frac{y e^y dy}{(e^y + 1)^2} \\ & + \frac{(k_B T)^2}{2} K''(\mu) \int_{-\mu/(k_B T)}^\infty \frac{y^2 e^y dy}{(e^y + 1)^2}. \end{aligned} \quad (\text{A.7})$$

We now make one further approximation. For very low temperatures, $k_B T \ll \mu$, so we can let $-\mu/(k_B T) \rightarrow -\infty$ in the lower limit of the integrals in (A.7). If we

do this and evaluate the integrals, we are left with

$$\begin{aligned} I &\simeq K(\mu) + \frac{\pi^2}{6} (k_B T)^2 K''(\mu) \\ &= \int_0^\mu d\epsilon H(\epsilon) + \frac{\pi^2}{6} (k_B T)^2 H'(\mu). \end{aligned} \quad (\text{A.8})$$

The integral in (A.8) can be split into two integrals,

$$I = \int_0^{\epsilon_F} d\epsilon H(\epsilon) - \int_\mu^{\epsilon_F} d\epsilon H(\epsilon) + \frac{\pi^2}{6} (k_B T)^2 H'(\mu). \quad (\text{A.9})$$

For a low temperature system, the chemical potential will be nearly equal to the Fermi energy. Consequently, the second integral in (A.9) can be approximated by

$$\int_\mu^{\epsilon_F} d\epsilon H(\epsilon) \simeq (\epsilon_F - \mu) H(\epsilon_F). \quad (\text{A.10})$$

Returning once more to (A.9), we see that the final term is quadratic in the temperature and contains the quantity $H'(\mu)$. Since we are concerned only with the first order correction to the temperature, we can let $H'(\mu) \rightarrow H'(\epsilon_F)$, in which case

$$I = \int_0^{\epsilon_F} d\epsilon H(\epsilon) - (\epsilon_F - \mu) H(\epsilon_F) + \frac{\pi^2}{6} (k_B T)^2 H'(\epsilon_F). \quad (\text{A.11})$$

To make further progress, we need to determine the low temperature expression for the chemical potential. In Section A.2.2, we will derive Equation (A.24), which we use to calculate the chemical potential. If we compare (A.24) to (A.1), we make the identifications: $I = \epsilon_F^3/3$ and $H(\epsilon) = \epsilon^2$. If we plug these values into (A.11) and solve for μ , we obtain the Sommerfeld approximation for the chemical

potential,

$$\mu \simeq \varepsilon_F \left[1 - \frac{\pi^2}{3} \left(\frac{T}{T_F} \right)^2 \right], \text{ for } \frac{T}{T_F} \ll 1. \quad (\text{A.12})$$

We can now use this analytic expression for the low temperature chemical potential to simplify (A.11). Plugging (A.12) into (A.11), we obtain our final expression for the Sommerfeld expansion,

$$I = \int_0^{\varepsilon_F} d\epsilon H(\epsilon) - \varepsilon_F \frac{\pi^2}{3} \left(\frac{T}{T_F} \right)^2 H(\varepsilon_F) + \varepsilon_F^2 \frac{\pi^2}{6} \left(\frac{T}{T_F} \right)^2 H'(\varepsilon_F). \quad (\text{A.13})$$

Consequently, obtaining Sommerfeld expansions requires writing an integral of the form (A.1), identifying $H(\epsilon)$ and using (A.13) to obtain the lowest order temperature dependence.

A.2 Derivation of basic quantities

In this section, I will derive the density of states for a harmonically trapped gas. Using this result, we can calculate an expression for the number of trapped atoms, which we can then use to determine the chemical potential for a noninteracting Fermi gas as a function of temperature.

A.2.1 Density of states

As we are assuming that shell effects can be neglected and the density profile of our trapped atom cloud can be treated as a continuous distribution, we will need to derive an expression for the density of states for a harmonically trapped Fermi gas. The density of states can be derived quickly by pursuing a geometrical argument. We begin with the expression for the energy eigenvalues for a three-dimensional

harmonic oscillator with trap oscillation frequencies ω_i , for $i = x, y, z$ [30],

$$\epsilon = \hbar \left[\omega_x \left(n_x + \frac{1}{2} \right) + \omega_y \left(n_y + \frac{1}{2} \right) + \omega_z \left(n_z + \frac{1}{2} \right) \right]. \quad (\text{A.14})$$

Here, the n_i can take on nonnegative integer values. If we drop the zero point energy factors of $1/2$ in (A.14), we have

$$\epsilon = \hbar (\omega_x n_x + \omega_y n_y + \omega_z n_z), \quad (\text{A.15})$$

which defines a plane in energy space. To determine the density of states, we can use the usual technique of calculating the volume of states in energy space enclosed by a particular value of the energy ϵ . Taking the derivative of this volume with respect to energy then provides the density of states. Consider Figure A.1. The shaded volume enclosed by a particular choice of the n_i is the volume of a pyramid,

$$\begin{aligned} V_{DoS} &= \frac{1}{3} \times \text{height of pyramid} \times \text{area of base} \\ &= \frac{1}{3} \times n_z^{max} \times \left(\frac{1}{2} n_x^{max} \times n_y^{max} \right), \end{aligned} \quad (\text{A.16})$$

where the n_i^{max} factors represent the value of n_i when $n_j = n_k = 0$ (in other words, the n_i^{max} are the coordinates of the intersection of the slanted plane and the coordinate axes in Figure A.1). Using (A.15), we see that

$$n_i^{max} = \frac{\epsilon}{\hbar \omega_i}. \quad (\text{A.17})$$




Figure A.1: Energy space for the three dimensional harmonic oscillator. The choice of energy eigenvalues for the three orthogonal coordinates defines a plane in energy space which can be used to calculate the density of states.

Plugging (A.17) into (A.16), we obtain

$$V_{DoS} = \frac{\epsilon^3}{6 (\hbar\bar{\omega})^3}, \quad (\text{A.18})$$

where the geometric mean of the trap oscillation frequencies is given by $\bar{\omega} = (\omega_x \omega_y \omega_z)^{1/3}$. If we take the derivative of (A.18) with respect to energy, we obtain the density of states for a harmonically trapped three-dimensional gas,

$$g(\epsilon) = \frac{dV_{DoS}}{d\epsilon} = \frac{\epsilon^2}{2 (\hbar\bar{\omega})^3}. \quad (\text{A.19})$$

A.2.2 Chemical potential

Using the density of states (A.19) along with the Fermi occupation number (A.2), we can calculate the chemical potential of the system for a particular temperature

and number of atoms. Throughout this dissertation, we will frequently consider two-component Fermi gases with equal spin populations. I will define N as the total number of particles in the gas, which naturally gives $N/2$ particles for each spin component. Hence, we can write

$$\frac{N}{2} = \int_0^\infty g(\epsilon) f(\epsilon) d\epsilon = \frac{1}{2(\hbar\bar{\omega})^3} \int_0^\infty \frac{\epsilon^2 d\epsilon}{\exp\left(\frac{\epsilon-\mu}{k_B T}\right) + 1}. \quad (\text{A.20})$$

If we consider a zero temperature Fermi gas, the preceding integral simplifies. The Fermi occupation number becomes unity for all energy levels below the Fermi energy ε_F , and is zero for all energy levels above the Fermi energy. In that case,

$$\frac{N}{2} = \frac{1}{2(\hbar\bar{\omega})^3} \int_0^{\varepsilon_F} \epsilon^2 d\epsilon = \frac{\varepsilon_F^3}{6(\hbar\bar{\omega})^3}, \quad (\text{A.21})$$

and a simple manipulation gives the Fermi energy for a noninteracting Fermi gas confined in a harmonic trap with $N/2$ particles per spin state,

$$\varepsilon_F = \hbar\bar{\omega}(3N)^{1/3}. \quad (\text{A.22})$$

The Fermi energy sets a characteristic energy scale which is of great importance when discussing both interacting and noninteracting Fermi gases. Similarly, one can define a related temperature, the Fermi temperature T_F . For a harmonically trapped noninteracting Fermi gas with $N/2$ particles per spin state, the Fermi temperature is

$$T_F = \frac{\varepsilon_F}{k_B} = \frac{\hbar\bar{\omega}(3N)^{1/3}}{k_B}. \quad (\text{A.23})$$

The definitions of the Fermi energy and temperature will allow us to write an

integral equation which can be used to determine the chemical potential μ . Using the definition of the Fermi energy, we can rewrite (A.20) as

$$\frac{\varepsilon_F^3}{3} = \int_0^\infty \frac{\epsilon^2 d\epsilon}{\exp\left(\frac{\epsilon - \mu}{k_B T}\right) + 1}. \quad (\text{A.24})$$

Often, the relevant physics will depend not on the absolute temperature T but instead on the ratio of the temperature to the Fermi temperature. To this end, we define a reduced temperature $T' \equiv T/T_F$, a reduced energy $\epsilon' \equiv \epsilon/\varepsilon_F$, and we normalize the chemical potential to the Fermi energy as well, $\mu' \equiv \mu/\varepsilon_F$. With these substitutions, (A.24) can be written

$$\frac{1}{3} = \int_0^\infty \frac{\epsilon'^2 d\epsilon'}{\exp\left(\frac{\epsilon' - \mu'}{T'}\right) + 1}. \quad (\text{A.25})$$

Equation (A.25) can be used to determine the chemical potential for a given value of the reduced temperature. Simply select a value for the reduced temperature T' and numerically determine the scaled chemical potential μ' .

Finally, we repeat here the Sommerfeld expression for the chemical potential derived in Section A.1,

$$\mu \simeq \varepsilon_F \left[1 - \frac{\pi^2}{3} \left(\frac{T}{T_F} \right)^2 \right], \quad \text{for } \frac{T}{T_F} \ll 1. \quad (\text{A.26})$$

The Sommerfeld approximation and the result of the full numerical calculation of the chemical potential are plotted in Figure A.2.




Figure A.2: Chemical potential in units of the Fermi energy versus temperature. The solid line shows the full numerical calculation using (A.25). The dashed line shows the Sommerfeld approximation (A.26).

A.2.3 Average energy per particle

Determining the average energy per particle for a harmonically confined Fermi gas is a relatively straightforward task. We begin by writing the total energy of *both spin states* of the gas,

$$E(T) = 2 \int_0^\infty \epsilon g(\epsilon) f(\epsilon) d\epsilon = \frac{1}{(\hbar\bar{\omega})^3} \int_0^\infty \frac{\epsilon^3 d\epsilon}{\exp\left(\frac{\epsilon-\mu}{k_B T}\right) + 1}. \quad (\text{A.27})$$

Using (A.22) and the reduced parameters defined in Section A.2.2, we can derive an expression for the average energy per particle,

$$\frac{E(T)}{N} = 3 \varepsilon_F \int_0^\infty \frac{\epsilon'^3 d\epsilon'}{\exp\left(\frac{\epsilon'-\mu'}{T'}\right) + 1}. \quad (\text{A.28})$$




Figure A.3: Average energy per particle in units of the Fermi energy versus temperature. The solid line is generated using (A.28), while the Sommerfeld approximation (A.29) is the dashed line.

Obtaining values of the average energy per particle involves choosing a value of the reduced temperature, retrieving the value of the chemical potential from (A.25), and calculating the integral in (A.28). Note that for a zero temperature, harmonically trapped noninteracting Fermi gas, the average energy per particle is $3\varepsilon_F/4$. Using the Sommerfeld expansion in (A.13), the low temperature energy per particle is

$$\frac{E(T)}{N} \simeq \frac{3\varepsilon_F}{4} \left[1 + \frac{2\pi^2}{3} \left(\frac{T}{T_F} \right)^2 \right], \text{ for } \frac{T}{T_F} \ll 1. \quad (\text{A.29})$$

Both the full numerical result (A.28) and the Sommerfeld approximation (A.29) for the average energy per particle are plotted in Figure A.3.

A.3 Density profiles for harmonically trapped noninteracting Fermi gases

The primary method for extracting data from our experimental arrangement involves taking an absorption image of our atom clouds. These absorption images are then processed to yield spatial density profiles of the gas. As these profiles are the easiest way to extract information from our experimental system, we now discuss the theoretical density profiles for a harmonically trapped noninteracting Fermi gas. We start with zero temperature profiles before considering low temperature Sommerfeld profiles. Next, spatial profiles at arbitrary finite temperature, which we refer to as Thomas-Fermi profiles, will be presented. Finally, spatial profiles in the high temperature, classical limit will be presented.

Before deriving results for specific temperature conditions, however, I will present the general derivation of the spatial density profiles. As noted previously, we will assume that the local density approximation is valid. In this approximation, one supposes that a small volume within the trap contains a large number of fermions which act like its own homogeneous fermionic system. These individual homogeneous systems can then be “glued together” to produce the complete density profile. This technique is valid provided that there are a large number of particles in a small volume in the gas, and that the trapping potential varies slowly on the order of the size of this small volume [33]. If these conditions are met, then the three-dimensional density distribution is given by

$$n(x, y, z) = \frac{1}{(2\pi\hbar)^3} \int d^3p f[H(\mathbf{x}, \mathbf{p})], \quad (\text{A.30})$$

where f is the Fermi occupation number given by (A.2) and the single particle Hamiltonian is

$$H(\mathbf{x}, \mathbf{p}) = \frac{p^2}{2m} + U. \quad (\text{A.31})$$

Here, $U = U(x, y, z)$ is the potential energy. We can rewrite (A.30) as

$$n(x, y, z) = \frac{1}{(2\pi\hbar)^3} \int d\epsilon f(\epsilon) \int d^3p \delta \left[\epsilon - \left(\frac{p^2}{2m} + U \right) \right]. \quad (\text{A.32})$$

Performing the momentum integral over spherical coordinates, (A.32) can be written in a more convenient form as an integral over energy

$$n(x, y, z) = \frac{2\pi(2m)^{3/2}}{(2\pi\hbar)^3} \int_0^\infty d\epsilon f(\epsilon) (\epsilon - U)^{1/2} \Theta[\epsilon - U]. \quad (\text{A.33})$$

In (A.33), $\Theta[\epsilon - U]$ restricts the integration to energy values for which $\epsilon > U$. Finally, note that we have made no assumptions about the temperature in deriving (A.33).

A.3.1 Zero temperature profiles

We first consider a zero temperature noninteracting two-component gas of Fermi atoms with $N/2$ atoms per spin state. Here, the Fermi occupation number is unity below the Fermi energy and zero above. Equation (A.33) then becomes the straightforward integral

$$\begin{aligned} n_Z(x, y, z) &= \frac{2\pi(2m)^{3/2}}{(2\pi\hbar)^3} \int_0^{\epsilon_F} d\epsilon (\epsilon - U)^{1/2} \Theta[\epsilon - U] \\ &= \frac{4\pi(2m)^{3/2}}{3(2\pi\hbar)^{3/2}} \epsilon_F^{3/2} \left(1 - \frac{U}{\epsilon_F} \right)^{3/2} \Theta[\epsilon_F - U]. \end{aligned} \quad (\text{A.34})$$

We consider a three-dimensional harmonic oscillator potential of the form

$$U_{ho}(x, y, z) = \frac{m}{2} (\omega_x^2 x^2 + \omega_y^2 y^2 + \omega_z^2 z^2). \quad (\text{A.35})$$

We can rewrite the potential in the more convenient form

$$U_{ho}(x, y, z) = \varepsilon_F \left(\frac{x^2}{\sigma_x^2} + \frac{y^2}{\sigma_y^2} + \frac{z^2}{\sigma_z^2} \right), \quad (\text{A.36})$$

using the Fermi radii defined as

$$\sigma_i = \left(\frac{2 \varepsilon_F}{m} \right)^{1/2} \frac{1}{\omega_i}, \quad i = x, y, z. \quad (\text{A.37})$$

Using the harmonic potential along with the Fermi energy for a harmonic oscillator (A.22), we can rewrite (A.34) as

$$n_Z(x, y, z) = \frac{4N}{\sigma_x \sigma_y \sigma_z \pi^2} \left(1 - \frac{x^2}{\sigma_x^2} - \frac{y^2}{\sigma_y^2} - \frac{z^2}{\sigma_z^2} \right)^{3/2} \Theta \left[1 - \frac{x^2}{\sigma_x^2} - \frac{y^2}{\sigma_y^2} - \frac{z^2}{\sigma_z^2} \right]. \quad (\text{A.38})$$

Obtaining lower dimensional projections of (A.38) involves integration over the appropriate coordinate(s). The two-dimensional projection of $n_Z(x, y, z)$ is

$$\begin{aligned} n_Z(x, z) &= \int_{-y_{max}}^{y_{max}} n_Z(x, y, z) dy \\ &= \frac{3N}{2\pi \sigma_x \sigma_z} \left(1 - \frac{x^2}{\sigma_x^2} - \frac{z^2}{\sigma_z^2} \right)^2 \Theta \left[1 - \frac{x^2}{\sigma_x^2} - \frac{z^2}{\sigma_z^2} \right], \end{aligned} \quad (\text{A.39})$$

where $y_{max} = \sigma_y(1 - x^2/\sigma_x^2 - z^2/\sigma_z^2)^{1/2}$ in the integral. Likewise, using $z_{max} = \sigma_z(1 - x^2/\sigma_x^2)^{1/2}$, the one-dimensional zero temperature distribution is found to

be

$$\begin{aligned}
 n_Z(x) &= \int_{-z_{max}}^{z_{max}} n_Z(x, z) dz \\
 &= \frac{8 N}{5 \pi \sigma_x} \left(1 - \frac{x^2}{\sigma_x^2}\right)^{5/2} \Theta \left[1 - \frac{x^2}{\sigma_x^2}\right].
 \end{aligned} \tag{A.40}$$

Note that integration of (A.40) over x from $-\sigma_x$ to $+\sigma_x$ yields $N/2$, the number of fermions per spin state. A normalized one-dimensional zero temperature density profile is plotted in Figure A.4.

A.3.2 Sommerfeld profiles

To derive density profiles for low temperature harmonically trapped noninteracting Fermi gases, we return once more to (A.33), which provides the general expression for the three-dimensional density profile using the local density approximation. This integral has the same form as (A.1), and we make the identification

$$\begin{aligned}
 H(\epsilon) &= \frac{2 \pi (2 m)^{3/2}}{(2 \pi \hbar)^3} (\epsilon - U_{ho})^{1/2} \Theta [\epsilon - U_{ho}] \\
 &= \frac{6 N}{\epsilon_F^{3/2} \sigma_x \sigma_y \sigma_z \pi^2} (\epsilon - U_{ho})^{1/2} \Theta [\epsilon - U_{ho}],
 \end{aligned} \tag{A.41}$$

where again we will assume harmonic confinement of the form given by (A.36). Using (A.13), we find that the three-dimensional Sommerfeld density profile for a noninteracting Fermi gas with $N/2$ atoms per spin state is

$$\begin{aligned}
 n_s(x, y, z) &= \frac{4 N \Theta [1 - U_{ho}/\epsilon_F]}{\sigma_x \sigma_y \sigma_z \pi^2} \left\{ \left(1 - \frac{U_{ho}}{\epsilon_F}\right)^{3/2} \right. \\
 &\quad \left. - \frac{\pi^2}{2} \left(\frac{T}{T_F}\right)^2 \left[\left(1 - \frac{U_{ho}}{\epsilon_F}\right)^{1/2} - \frac{1}{4} \left(1 - \frac{U_{ho}}{\epsilon_F}\right)^{-1/2} \right] \right\} \tag{A.42}
 \end{aligned}$$

As done previously for the zero temperature profiles, we can obtain lower-dimensional versions of (A.42) by integrating over the appropriate coordinates. The two-dimensional Sommerfeld profile becomes

$$\begin{aligned} n_s(x, z) &= \int_{-y_{max}}^{y_{max}} n_s(x, y, z) dy \\ &= \frac{3N}{2\pi\sigma_x\sigma_z} \left[\left(1 - \frac{x^2}{\sigma_x^2} - \frac{z^2}{\sigma_z^2}\right)^2 \right. \\ &\quad \left. - \frac{\pi^2}{3} \left(\frac{T}{T_F}\right)^2 \left(1 - \frac{2x^2}{\sigma_x^2} - \frac{2z^2}{\sigma_z^2}\right) \right] \Theta \left[1 - \frac{x^2}{\sigma_x^2} - \frac{z^2}{\sigma_z^2}\right], \quad (\text{A.43}) \end{aligned}$$

where we have used the same limits of integration as in the calculation of (A.39). Likewise, using the same limits of integration employed in (A.40), we can integrate (A.43) over the z -direction to obtain

$$\begin{aligned} n_s(x) &= \frac{8N}{5\pi\sigma_x} \Theta \left[1 - \frac{x^2}{\sigma_x^2}\right] \left\{ \left(1 - \frac{x^2}{\sigma_x^2}\right)^{5/2} \right. \\ &\quad \left. + 5\pi^2 \left(\frac{T}{T_F}\right)^2 \left[\frac{1}{8} \left(1 - \frac{x^2}{\sigma_x^2}\right)^{1/2} - \frac{1}{6} \left(1 - \frac{x^2}{\sigma_x^2}\right)^{3/2} \right] \right\}. \quad (\text{A.44}) \end{aligned}$$

Note that integration of $n_s(x)$ over the x -direction will yield $N/2$, the number of atoms per spin state. Furthermore, as expected, in the zero temperature limit, (A.44) reduces to (A.40). A normalized one-dimensional Sommerfeld density profile for $T/T_F = 0.10$ is plotted in Figure A.4.

A.3.3 Finite temperature profiles

As the Sommerfeld profiles are valid only in the low temperature limit, there is a need to consider density profiles for arbitrary temperatures. Previous theses

[23, 24] from our research group have gone into substantial detail regarding the derivation of these profiles. As such, I will present only the results.

As in preceding subsections, we consider a noninteracting Fermi gas with $N/2$ particles per spin state. Integration of the complete phase space density over all momenta yields the three-dimensional spatial distribution,

$$n(x, y, z; T) = -\frac{3N}{\pi^{3/2} \sigma_x \sigma_y \sigma_z} \left(\frac{T}{T_F}\right)^{3/2} Li_{3/2} \left[\exp \left(\frac{\frac{\mu}{\varepsilon_F} - \frac{x^2}{\sigma_x^2} - \frac{y^2}{\sigma_y^2} - \frac{z^2}{\sigma_z^2}}{T/T_F} \right) \right], \quad (\text{A.45})$$

where the polylog function is defined as,

$$Li_n(x) = \sum_{k=1}^{\infty} \frac{x^k}{k^n}, \quad |x| < 1. \quad (\text{A.46})$$

Using the series expansion of the polylog function, one can obtain lower dimensional spatial distributions [23, 24]. The two- and one-dimensional spatial profiles, respectively, are

$$n(x, z; T) = -\frac{3N}{\pi \sigma_x \sigma_z} \left(\frac{T}{T_F}\right)^2 Li_2 \left[\exp \left(\frac{\frac{\mu}{\varepsilon_F} - \frac{x^2}{\sigma_x^2} - \frac{z^2}{\sigma_z^2}}{T/T_F} \right) \right] \quad (\text{A.47})$$

and

$$n(x; T) = -\frac{3N}{\sqrt{\pi} \sigma_x} \left(\frac{T}{T_F}\right)^{5/2} Li_{5/2} \left[\exp \left(\frac{\frac{\mu}{\varepsilon_F} - \frac{x^2}{\sigma_x^2}}{T/T_F} \right) \right]. \quad (\text{A.48})$$

In closing, we note that the polylogarithm function $Li_n(x)$ has a singularity at $x = 1$. Consequently, care should be exercised when performing numerical manipulations where the argument of the polylogarithm approaches 1. An example of a one-dimensional Thomas-Fermi profile is plotted in Figure A.4.




Figure A.4: Normalized one-dimensional density profiles for zero and finite temperature harmonically trapped noninteracting Fermi gases. The dotted curve is a zero temperature profile, while the dashed curve is a Sommerfeld density profile for $T/T_F = 0.10$. The solid curve is a Thomas-Fermi density profile for $T/T_F = 0.35$.

A.3.4 High temperature limit

At sufficiently high temperature, quantum statistics become unimportant, and the spatial density profiles can be derived by assuming Maxwell-Boltzmann statistics. While the finite temperature Thomas-Fermi shapes derived in Section A.3.3 will naturally tend toward classical shapes at high temperature, it is worthwhile to present the analytic forms in the classical regime to facilitate a discussion of temperature measurement techniques in Section 6.3.

The spatial profiles for harmonically confined classical gases are gaussian and have been derived elsewhere (see, for example, previous theses from this group [23,24]). The three-dimensional spatial profile for a noninteracting gas with $N/2$

particles per spin state is

$$n_c(x, y, z) = \frac{N}{2 \pi^{3/2} \sigma_x \sigma_y \sigma_z} \left(\frac{T}{T_F} \right)^{-3/2} \exp \left[-\frac{T_F}{T} \left(\frac{x^2}{\sigma_x^2} + \frac{y^2}{\sigma_y^2} + \frac{z^2}{\sigma_z^2} \right) \right]. \quad (\text{A.49})$$

The two- and one-dimensional profiles are, respectively,

$$n_c(x, z) = \frac{N}{2 \pi \sigma_x \sigma_z} \left(\frac{T}{T_F} \right)^{-1} \exp \left[-\frac{T_F}{T} \left(\frac{x^2}{\sigma_x^2} + \frac{z^2}{\sigma_z^2} \right) \right]. \quad (\text{A.50})$$

and

$$n_c(x) = \frac{N}{2 \pi^{1/2} \sigma_x} \left(\frac{T}{T_F} \right)^{-1/2} \exp \left[-\frac{T_F}{T} \left(\frac{x^2}{\sigma_x^2} \right) \right]. \quad (\text{A.51})$$

Appendix B

Imaging at high magnetic field

As discussed in Chapter 5, our primary method for extracting information from our system involves absorption imaging of the fermion clouds in the presence of a bias magnetic field. The bias field causes Zeeman tuning of the energy states which we use for imaging. For the magnetic fields of interest, the degree of Zeeman tuning of these states is large enough that the use of AOs to generate the desired frequency shift in the probe beam is inconvenient. While we still use AOs so we can accurately control the duration of the imaging probe pulse, the primary technique for generating the desired frequency involves briefly disengaging the servo system which keeps the dye laser locked to the D_2 transition and shifting the dye laser's frequency using a GPIB programmable power supply. Once the imaging process is complete, the dye laser is returned to its original frequency and the frequency servo is engaged once again.

I will present information regarding the Zeeman tuning of the ground and excited states in the presence of a magnetic field in Section B.1. This will tell us the frequency shift required to conduct imaging at high magnetic field. However, the process of shifting the dye laser's frequency involves feeding a voltage to the control box for the laser. Calibrating the relationship between this shift voltage and the associated frequency shift is covered in Section B.2.

B.1 Zeeman tuning of the ground and excited states

In Section 2.4, we saw that the energy eigenvalues associated with the $|1\rangle$ and $|2\rangle$ hyperfine ground states assume the following form:

$$E_1(B) = -\frac{1}{4} (a_{hf} - 2 g_I \mu_B B + 2 a_{hf} R^+) \quad (\text{B.1})$$

$$E_2(B) = -\frac{1}{4} (a_{hf} + 2 g_I \mu_B B + 2 a_{hf} R^-). \quad (\text{B.2})$$

Recall that in the preceding expressions, $a_{hf}/h = 152.137$ MHz is the hyperfine coupling constant and $g_J^{gnd} = 2.002$ is the total electronic g-factor for the ${}^6\text{Li}$ ground state, $g_I = -0.000448$ is the total nuclear g-factor, μ_B is the bohr magneton, and B is the external magnetic field. The R^\pm factors that appear in (B.1) and (B.2) are

$$R^\pm = \sqrt{(Z^\pm)^2 + 2}. \quad (\text{B.3})$$

$$Z^\pm = \frac{\mu_B B}{a_{hf}} (g_J^{gnd} - g_I) \pm \frac{1}{2}. \quad (\text{B.4})$$

Equations (B.1) through (B.4) provide all the necessary information about the energy level tuning of the $|1\rangle$ and $|2\rangle$ hyperfine ground states as a function of magnetic field.

In Section 5.4, we saw that the Zeeman tuning of the desired excited state is given by

$$E_{ex}(B) = \mu_B \left(-\frac{3}{2} g_J^{ex} + g_I m_I \right) B, \quad (\text{B.5})$$

where $g_I = -0.000448$ and $g_J^{ex} = 1.335$ for the excited state. As the g-factor

for the nucleus is much smaller than the g -factor for the electron, (B.5) can be approximated by

$$E_{ex}(B) \simeq -\frac{3}{2} \mu_B g_J^{ex} B. \quad (\text{B.6})$$

Now that the tuning of the ground and excited states has been established, we are in a position to consider the absolute frequency shifts needed to image on a transition between one of the hyperfine ground states and the desired excited state. To do this properly, we need to outline a few properties of our frequency standard which was first discussed in Section 3.2.6. Recall that the frequency standard is established by monitoring the fluorescence from the perpendicular intersection of a beam from the dye laser and an atomic beam of ${}^6\text{Li}$ inside the locking region vacuum system. The perpendicular intersection of the beams means that the frequency of the D_2 transition can be detected without any Doppler shifts. Furthermore, as this excitation and detection occurs in the absence of a magnetic field, the hyperfine structure is visible in the form of two quantum transitions (from the $F = 1/2$ and $F = 3/2$ ground state manifolds). If the dye laser is scanned over the D_2 line, the resulting fluorescence signal looks like Figure B.1. The lower frequency transition (the tall peak) represents transitions from the $F = 3/2$ ground state manifold to the excited state, while the higher frequency transition (short peak) is associated with transitions between the $F = 1/2$ ground state manifold and the excited state. Under normal operation, the dye laser's reference cavity is locked to the taller fluorescence peak.

Before writing an equation for the total frequency shift between the hyperfine ground states and the excited state, we should consider some frequency offsets which exist in the system. Recall that our locking scheme places the dye laser output frequency roughly 200 MHz below the $F = 3/2$ ground state to excited




Figure B.1: Fluorescence in arbitrary units as a function of laser frequency in GHz (this is not an absolute frequency). We see here the D_2 line in ${}^6\text{Li}$. The taller peak to the left represents the transition from the $F = 3/2$ ground state manifold to the excited state while the shorter peak on the right represents the transition from the $F = 1/2$ ground state manifold to the excited state. The two peaks are separated by roughly 228 MHz. Under normal operation an electronic servo locks the dye laser cavity to the tall fluorescence peak.

state transition. Further, as discussed in the previous paragraph, we are interested in the tuning of the hyperfine states which start in the $F = 1/2$ ground state manifold at zero magnetic field. These states are excited at a frequency roughly 228 MHz larger than the $F = 3/2$ ground state to excited state transition. Hence, to excite the $F = 1/2$ ground state on the D_2 transition at zero magnetic field, we would need to upshift the laser frequency by roughly 428 MHz. This is accomplished by two AOs in series in the camera beam path. Each AO is set up in a double-pass configuration, such that the total frequency shift for each AO is roughly 215 MHz. Using two AOs in series has the added benefit that turning off the amplitude control to the AOs provides better protection against leakage light than one AO alone. After exiting the second AO, the probe beam is directed to a fiber which guides the probe beam toward the vacuum system. What is important for our discussion here, however, is that the probe beam is designed to be resonant with the $F = 1/2$ ground state to excited state transition when the laser is locked and the atoms are at zero magnetic field. If we consider (B.1) and (B.2), we see that for the two $F = 1/2$ states $E_1(0) = E_2(0) = -a_{hf}$. This non-zero value arises because when calculating the energy eigenvalues for the hyperfine ground state, the “zero” is taken to be the energy level if one neglects the coupling between the electronic and nuclear spins. When calculating how the energy levels of the $|1\rangle$ and $|2\rangle$ states shift as function of magnetic field, however, we want to set our energy origin at their zero magnetic field position. This is readily accomplished by adding a_{hf} to (B.1) and (B.2). Then the total frequency

shifts for imaging on the $|1\rangle$ and $|2\rangle$ states at nonzero magnetic field are

$$\Delta E_1(B) = E_{ex}(B) - E_1(B) + a_{hf} \quad (\text{B.7})$$

$$\Delta E_2(B) = E_{ex}(B) - E_2(B) + a_{hf}, \quad (\text{B.8})$$

where $E_{ex}(B)$ is given by (B.6), and $E_1(B)$ and $E_2(B)$ are given by (B.1) and (B.2), respectively. Equations (B.7) and (B.8) are plotted in frequency units in Figure B.2 as a function of magnetic field. As (B.7) and (B.8) yield increasingly negative values with increasing magnetic field, we conclude that the excited and ground state energy levels involved in the imaging process tune toward each other as a function of magnetic field. Consequently, to image on resonance at high magnetic field, we need to shift the dye laser's frequency to a lower value during the imaging process. Some details of this operation are addressed in Section B.2.

B.2 Determining the frequency-to-voltage conversion factor

During normal operation, the dye laser's frequency is locked to the atomic fluorescence from the locking region vacuum system. A servo system keeps the dye laser on the atomic resonance by sending an error voltage to the "EXT SCAN" input on the back of the 699 dye laser control box. This BNC post takes a ± 5 volt input which is capable of driving the dye laser through the complete range of the frequency scan. The range of the frequency scan is specified by the controls on the front of the dye laser control box. When the laser is locked to the




Figure B.2: The frequency shift of the imaging transitions between the hyperfine ground states and the desired excited state is plotted as a function of magnetic field. The solid line represents the frequency shift for the transition between $|1\rangle$ and the excited state, while the dashed line shows the frequency shift for the transition between $|2\rangle$ and the excited state.

external signal provided by the atomic fluorescence from the locking region, the voltage sent to the “EXT SCAN” post remains relatively constant, with small fluctuations to compensate for excursions in the dye laser’s frequency.

When conducting high field imaging, however, the servo system is briefly disengaged, and the voltage sent to the “EXT SCAN” input is shifted by an amount that depends on several factors. As we saw in Section B.1, the hyperfine ground states of interest Zeeman tune as a function of magnetic field. Hence, both the magnetic field and the desired imaging transition must be considered when determining this voltage shift. Furthermore, the frequency scan range specified by the user on the front panel of the dye laser’s control box will be important in determining the voltage which corresponds to the desired frequency shift.

As seen in Figure B.2, to image either spin state $|1\rangle$ or $|2\rangle$ at magnetic fields up to 1000 gauss, we will need to be able to decrease the dye laser’s frequency by about 1.5 GHz. To make this possible, when locking the dye laser, we must place the atomic fluorescence peak near the upper portion of a frequency scan whose width exceeds 1.5 GHz. To provide sufficient headroom, we set the dye laser to have a 2.3 GHz wide scan. As the full 2.3 GHz scan should correspond to ramping the input to the “EXT SCAN” post from -5 to +5 volts, we conclude that the frequency-to-voltage conversion should be 2300 MHz per 10 volts, or 230 MHz/volt. However, this is a somewhat crude estimate of the frequency-to-voltage conversion.

A better conversion factor can be obtained by using a Fabry-Perot interferometer. Using a Coherent 240-2-B Fabry-Perot interferometer with a free spectral range of 1.5 GHz, we can more precisely determine the frequency-to-voltage conversion. With the dye laser scan width set to 2.3 GHz, and with the dye laser

locked to the D_2 fluorescence from the locking region, set the interferometer to show two modes at once. Next, run a timing file which unlocks and shifts the laser. If you fail to specify a frequency shift voltage, you will find that the dye laser frequency remains at its “locked” value during the unlock and shift sequence. However, as you specify larger and larger frequency shift voltages, you will find that the modes on the interferometer will shift by a larger and larger amount. On an oscilloscope, set cursors to mark the locations of the detected peaks from the interferometer when the laser is locked. Then, find the voltage required to make one of the detected peaks jump to the location of the adjacent peak. In doing so, you find the voltage needed to shift the dye laser frequency by 1.5 GHz, the free spectral range of the interferometer. Measurements indicate that it takes roughly 5.7 volts to shift the dye laser frequency 1.5 GHz, or approximately 263 MHz/volt. This is a little more than 10% larger than the conversion factor crudely estimated from knowledge of the dye laser’s frequency scan width.

Appendix C

Custom Software Upgrades

Much of the custom software required to operate the experimental system was developed and installed by Stephen Granade, and has been described in his dissertation [23]. Here, I describe recent software upgrades as well as new tools that have been added to the lab’s custom software library. In Section C.1, I describe two new usages of the DG535 Stanford Pulse Generators which are employed for precise timing control. The remainder of the appendix is devoted to describing the operation of two LabVIEW virtual instruments (VIs). In Section C.2, I discuss the operation of *AutomateAcquisition.vi*, which allows the user to automate the acquisition of data. Finally, to send the so-called “lowering curves” to the experimental system for conducting forced evaporation, we have been using the LabVIEW program named *Arbitrary Waveforms.vi*. In Section C.3, I will describe the operation of this program. We begin, however, with a description of the usage of the Stanford Pulse Generators.

C.1 New usage of Stanford Pulse Generators

The DG535 Stanford pulse generators, hereafter referred to as SPGs, are assigned tasks that require high timing precision. For example, the processes of exciting

the breathing mode, turning off the CO₂ laser and generating the resonant probe pulse are all handled by SPGs. The operation of the generators is relatively simple. Upon receiving an external trigger, an SPG can output up to 4 independent timing events, labelled A, B, C and D. The user can define the delays between each of these timing events. While the output voltages and impedances of the SPG can be configured by the user, they are essentially digital instruments: each output of the SPG can exist in an “on” or “off” position. While there are only 4 independent timing events, there are a greater number of outputs on the SPG. The outputs labelled with a single letter (A, B, C, or D) change state according to the delay times commanded by the user and remain in that state until all timing tasks are completed by the SPG, at which point all channels are returned to their initial states. However, there are also outputs which correspond to square pulses. These outputs are labelled A-up-B-down, A-down-B-up, C-up-D-down, and C-down-D-up (on the front panel of the SPGs, these outputs are identified by the upright or inverted square pulses sandwiched between the letters A and B or C and D). The A-up-B-down channel, for example, will emit a square pulse with a rising edge when event A occurs and a falling edge when event B occurs.

In Section C.1.1, I will outline a software upgrade to *Alter Channels.vi* which allows us to conduct simple error checking of the SPGs prior to running each timing file. In Section C.1.2, I will describe the current scheme for storing configurations in the SPGs to prevent timing errors which can arise when switching between timing files.

C.1.1 Error checking from within *Alter Channels.vi*

Due to the complexity of some of the timing files which control the 32 timing channels and the SPGs, it is easy to send an invalid command to the SPGs. In the past, the control of the timing system was a one-way street. The timing system software would send commands to the system, but was unable to read back information from elements within the timing system.¹ This is no longer the case. A simple error checking scheme has been implemented which will prevent the user from making many common mistakes when commanding the operation of the SPGs.

This software upgrade was deemed necessary because while certain SPG errors can be detected by the LabVIEW software, LabVIEW was not able to detect all such errors, including some of the most common ones. As the LabVIEW software could not detect some SPG errors, the only way to become aware of such errors was to read the front panel of the SPGs, which tended to occur only when the system was obviously misbehaving or during the process of shutting down the system.

The software upgrade amounts to a relatively small change in the diagram panel of *Alter Channels.vi*. A small portion of the Alter Channels diagram panel is shown in Figure C.1. The sub-VI *SPG Error Check.vi* handles all of the error checking and alerts the user to the existence of errors in any of the SPGs. The user must specify the GPIB address of all of the SPGs using a numeric array. *SPG Error Check.vi* works by querying the SPGs for the existence of all classes of errors. If no errors are reported, *Alter Channels.vi* proceeds by running the

¹The only element which was able to record information from the system and funnel it back to the computer was the Andor CCD camera, which is involved in only a small portion of the timing system's operation.

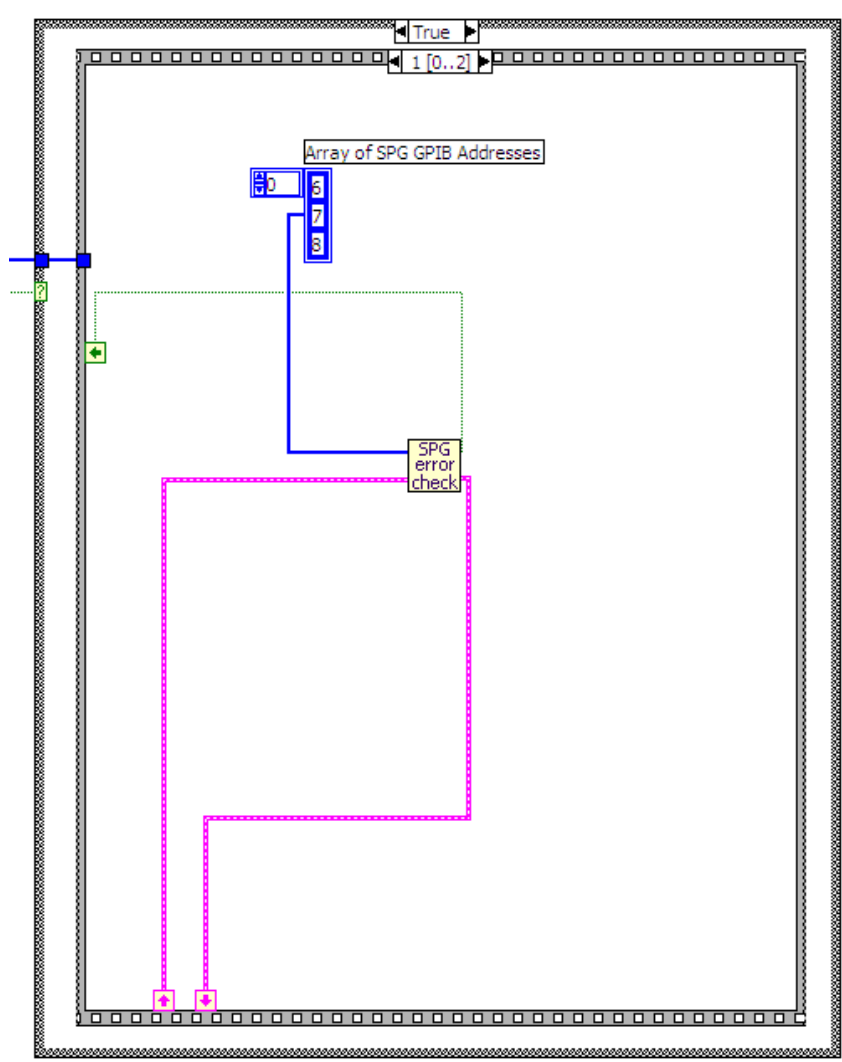


Figure C.1: Near the bottom of the large While loop on the *Alter Channels.vi* diagram panel is a Sequence structure nested inside a Case structure. The first panel of the Sequence structure sends GPIB commands to the SPGs (and any other GPIB programmable instruments needed for that timing file). The second panel of the Sequence structure checks for errors in the SPGs using the sub-VI *SPG Error Check.vi*. The user must tell the program the GPIB addresses of the SPGs using a numeric array.

timing file. If an error is reported, the user is alerted to the class of error that has occurred. For an explanation of the error, the user must consult the SPG owner's manual. In closing, I note that the additional time required to perform this error checking is small compared to the overall cycle time of the experimental apparatus.

C.1.2 Use of stored configurations

As noted earlier, each SPG can output 4 time edges (A, B, C, and D) or pulses (labelled A-up-B-down, C-up-D-down, A-down-B-up, and C-down-D-up). In some experiments, all of the available time edges will be used, while in others, one or more may be unnecessary. This creates a potential danger if you are not careful when writing timing files. Suppose that in timing file #1, you use pulses A-up-B-down and C-up-D-down, and you program the appropriate pulse widths and delay times. Then, suppose you switch to timing file #2, in which you wish to use pulse A-up-B-down only. If you forget to turn off pulse C-up-D-down, this pulse will still be output when you run timing file #2, an obviously undesirable scenario. To overcome this problem, it is safest to make use of configurations which can be stored in nonvolatile RAM in the SPGs.

Perhaps it is easiest to think of the stored configurations like templates. In each timing file that requires the use of an SPG, select the template that most closely matches your goals, then write only a few commands that modify that template. Using stored configurations makes the GPIB programming process cleaner, and reduces the likelihood that unintended pulses will remain active from a previous timing file. The process of writing a configuration and storing it in the SPG's memory is handled by the timing file named *SPG Configurations.txt*. This

file is located in the same directory as the other timing files.

After a few introductory notes at the beginning of the file, there are two major sections. The first contains a list of all the GPIB commands that are used to define the various configurations. The second portion of the timing file contains a list of the 32 channels used in the timing system. This list of channels is necessary to avoid getting an error from *Alter Channels.vi* when you run the *SPG Configurations.txt* file.

For each configuration listed in the *SPG Configurations.txt* timing file, there are several parts. First, there is a brief description of when you should use that particular configuration. Next, there is a list of GPIB commands that define a particular configuration. Finally, in the commented out lines at the end of each configuration are a list of the commands that need to be used in each timing file that makes use of that particular configuration.

If you want to understand all of the commands in the configuration file, you should refer to the SPG manual. However, it is worthwhile at this point to outline the general structure of the command sequence in the configuration files. For each configuration, the first command is **CL**. This restores the SPG to its factory default settings. That is, it sets the trigger mode to single shot, and it configures the external trigger to occur at +1 volt on a positive slope with a high impedance termination. All delay times are set to zero with T_0 as the sync edge. All outputs are configured to drive high impedance loads to TTL levels. All edges are returned to normal polarity (the A, B, C, and D outputs all produce rising edges). The commands that follow the **CL** command modify the factory default settings. The next line usually sets the trigger mode to External and sets the trigger level to 2.5 volts. Subsequent lines often configure the output impedances,

levels, and offsets of various edges or pulses. This is the location where one can set a given pulse to be “off” in a particular configuration. By starting from the factory default settings, you have already ensured that all pulse durations are zero seconds. However, as an added level of security, you can set the output amplitude for channels you wish to remain off to 0.1 volts (the minimum allowed pulse amplitude). Finally, each configuration concludes with the `ST` command, which stores the configuration in the SPG’s nonvolatile RAM.

As mentioned earlier, there are also commented out lines that contain GPIB command sequences that need to be called in the timing files that use that particular configuration. The first command is always `RC`, which recalls a stored configuration. The remaining commands require numeric values of timing variables that must be generated in the individual timing files. Places where you need to input a timing variable are indicated by braces.

Typically, you should not need to run the *SPG Configurations.txt* file unless you notice an error in one of the configurations, or you need to write a new configuration. However, **when you write a new configuration or modify an existing one, you must run the *SPG Configurations.txt* file in *Alter Channels.vi* to send the updated configurations to the SPGs. Also, be very careful when modifying an existing configuration; the behavior of all timing files that use that stored configuration will change.**

C.2 *AutomateAcquisition.vi*

AutomateAcquisition.vi (often referred to as “the VI” in the remainder of Section C.2) is a crude attempt to automate the acquisition of data during lengthy

data runs. It relies heavily on the machinery developed in *Alter Channels.vi*, and as it was written hastily, some of the programming choices were not especially elegant. What it lacks in beauty, however, it compensates for in function. Experiments such as the studies of the breathing mode would not have been possible without it.

Use of the program is fairly straightforward. First, the user selects a particular timing file to be used in the data run. Next, he or she specifies a timing file parameter that is to be varied during the data run. The user provides a start and stop value for this parameter, as well as a step size. The VI then generates a bunch of timing files in which the timing parameter of interest is varied from “start value” to “end value” in units of “step size.” (The user also has the option of providing a comma separated list of parameter values in lieu of the start, stop and step values). The VI then randomizes the order of the timing files (to reduce the effect of systematic drifts in the experiment) and runs the files one at a time. When used in conjunction with the *FK_AR_autolevel_autosave.pgm* BASIC program written for the Andor MCD software, the VI will automatically run the experimental system and save the data.

C.2.1 Intended usage of *AutomateAcquisition.vi*

When choosing between saving data manually using *Alter Channels.vi* and automating the data acquisition process using *AutomateAcquisition.vi*, consider the following questions:

- Will you need to vary a particular experimental parameter during the course of a data acquisition run?

- If the answer to the first question is “yes”, will you need to vary this parameter over a wide range of values?
- Can the experimental parameter of interest be adjusted in the timing file?
- Will the data acquisition process be lengthy? (On the order of hours, for example).
- Are you lazy, tired, or some combination of the two?

If the answer to all of the above questions is “yes”, then *AutomateAcquisition.vi* will likely be useful to you.

C.2.2 Using *AutomateAcquisition.vi*

The following series of steps should guide you through the process of using *AutomateAcquisition.vi*. You can automate only the running of the experiment, or you can automate the processes of running the experiment and saving the data. There are separate subsections for these two cases.

To automate running of the system only

If you want to automate the control of the experimental apparatus while still saving data manually, do the following:

1. If you are going to be saving background corrected images of your data, make sure that you have already taken a background using *Alter Channels.vi*. Unfortunately, *AutomateAcquisition.vi* does not allow you to acquire a background image.

2. Examine the timing file you are going to use during your data acquisition run. With the exception of the timing parameter that will be varied during the data run, make sure that all of the timing parameters in the timing file are set to the desired values.
3. Open *AutomateAcquisition.vi* and make sure you are looking at the contents of the tab named “Load Timing File” (the leftmost tab). Press the little folder icon at the right of the “Timing File” path control box to bring up an “open file” dialog box. Select the desired timing file, and click “Open”. Next press the large “Load Timing File*” button in the upper right portion of the tab. (Make sure that the VI is running at this point). You should now see the contents of the timing file in the “Text of Timing File” window at the bottom of the screen.
4. Scroll through the “Text of Timing File” window until you find the timing parameter that you wish to vary during the experimental run. Highlight and copy this parameter.
5. Click on the “Configure Acquisition Sequence” tab (the middle tab). On this tab, you will define the range of values over which to vary the timing parameter of interest. The box in the upper left corner of the tab defines the paths to various files needed to run the program; these filepaths should not need to be changed often. In the box labelled “ACQUISITION SEQUENCE PARAMETERS”, click in the text line beneath “Name of Timing Parameter to be Varied” and paste the value that you copied in step 4. When pasting this timing parameter, the lead “\$” is optional. (The VI will accept both *\$mot_loading_time* and *mot_loading_time*, for example). You have two

options for defining the range of values over which the timing parameter of interest will be varied. A toggle switch named “List Creation Method” at the right selects between the options. When the toggle switch is in the “up” position, you should enter values in the “Start Value”, “End Value” and “Step Size” boxes. During the course of the experimental run, *AutomateAcquisition.vi* will vary the “Name of Timing Parameter to be Varied” from “Start Value” to “End Value” in units of “Step Size”. Make sure that “Start Value” is less than “End Value”. When the “List Creation Method” toggle switch is in the “down” position, the timing parameter of interest will be varied using a comma-separated list of values that can be entered in the box to the right of the toggle switch. Note that a Mathematica notebook named *List for automated acquisition.nb* (located in the LAB1_DOCS directory on the network) can be used to create a comma-separated list of parameter values which can be copied and pasted into *Automate Acquisition.vi*.

6. In the upper right corner of the “Configure Acquisition Sequence” tab is a button labelled “Prepare Acquisition Sequence”. Press this button. You should now see a sample text file in the window named “PREVIEW SAMPLE MODIFIED TIMING FILE.” Scroll through the timing file until you find the timing parameter that you wish to vary. Verify that the value of this parameter has been set equal to the provided “End Value.”

Note that the “Trial Number” control is meaningless when you do not intend to automate the process of saving images.

7. Click on the rightmost tab named “Take Data.” A description of each of

the controls and indicators is provided below.

- “Start Acquisition Sequence” button: press this button to begin the data run
- “Stop” button: press this button to stop the data run after the present loop of the timing file is completed
- “Queue Array” control: allows you to look at the sequence of timing files that will be employed during the data run
- “Queue Offset” control: when set to zero, this control will instruct the VI to loop through the entire list of timing parameters. When set to a positive integer, this control allows the user to start a data run “midway” through the list of timing parameters. (This feature might be useful, for instance, if the laser falls out of lock midway through a data run and you need to stop the program while you relock the laser. When you restart the program, you can set the “queue offset” control to start at the value of the timing parameter immediately before you lost laser lock).
- “Currently Acquiring Shot Number” indicator: displays how many shots you have taken during the present data run
- “Shots Remaining” indicator: displays how many shots are left in the present data run
- “Present Shot” indicator: displays the value of the experimental parameter for the piece of data presently being acquired
- “Last Shot Acquired” indicator: indicates the value of the experimental parameter for the previous loop of the timing file.

- “Alter Channels Indicators*” box of indicators: displays several of the indicators from the front panel of *Alter Channels.vi*. These are provided so you can track the progress of each acquisition.

To begin the data run, simply set “Queue Offset” to the desired value (usually zero) and press “Start Acquisition Sequence.” When you press the start button, a dialog box will appear. If you are not going to be automating the image saving process, you can ignore this dialog box. Just press “Continue.” The VI will loop through all the desired values of the timing parameter and then stop. If you wish to stop the data acquisition process prematurely, press “STOP.” The VI will complete data acquisition for the present value of the timing parameter and then stop data acquisition. When the end of the data run is reached (either when you press “STOP” or when the VI naturally completes the data loop), another dialog box will appear. If you are not automating the image saving process, you can ignore this dialog box. Simply press “CONTINUE.”

Note: at the beginning of a data run, *AutomateAcquisition.vi* will loop through the first value of the experimental parameter 3 times. This is done because we have noticed that system performance varies depending on the duration of MOT loading. It is recommended that you not save any data until you reach the third shot at the first value of the experimental parameter.

To automate running the system and saving the images

If you wish to automate the processes of running the system and saving the images, do the following:

1. Complete steps 1 through 5 listed in the section titled “**To automate running of the system only.**”
2. In the upper right corner of the “Configure Acquisition Sequence” tab is a button labelled “Prepare Acquisition Sequence.” Press this button. You should now see a sample text file in the window named “PREVIEW SAMPLE MODIFIED TIMING FILE.” Scroll through the timing file until you find the timing parameter that you wish to vary. Verify that the value of this parameter has been set equal to the provided “End Value.”

Finally, note the “Trial Number” control just below the “Prepare Acquisition Sequence” tab. When you automate the image saving process, *AutomateAcquisition.vi* will assign filenames to the recorded data using the following format:

(value of timing parameter)T(trial number).asc

If you want to take multiple data points at a given value of the timing parameter of interest, you will need to take several data runs. (At present, *AutomateAcquisition.vi* does not allow the user to take several data points at each value of the experimental parameter). In this case, you will need to set the “Trial Number” control to 1 for the first data sequence, 2 for the second sequence, and so on. **BE CAREFUL HERE: IF YOU NEGLECT TO UPDATE THE “TRIAL NUMBER” CONTROL AT THE BEGINNING OF A DATA RUN IN WHICH YOU WISH TO ACQUIRE A SECOND (OR THIRD, OR FOURTH...) PIECE OF DATA FOR EACH VALUE OF THE TIMING PARAMETER OF INTEREST, YOU MIGHT OVERWRITE SOME OF YOUR**

PREVIOUS DATA.

3. Click on the rightmost tab named “Take Data.” A description of each of the controls and indicators is provided below.
 - “Start Acquisition Sequence” button: press this button to begin the data run
 - “Stop” button: press this button to stop the data run after the present loop of the timing file is completed
 - “Queue Array” control: allows you to look at the sequence of timing files that will be employed during the data run
 - “Queue Offset” control: when set to zero, this control will instruct the VI to loop through the entire list of timing parameters. When set to a positive number, this control allows the user to start a data run “midway” through the list of timing parameters. (This feature might be useful, for instance, if the laser falls out of lock midway through a data run and you need to stop the program while you relock the laser. When you restart the program, you can set the “queue offset” control to start at the value of the timing parameter immediately before you lost laser lock).
 - “Currently Acquiring Shot Number” indicator: displays how many shots you have taken during the present data run
 - “Shots Remaining” indicator: displays how many shots are left in the present data run
 - “Present Shot” indicator: displays the value of the experimental parameter for the piece of data presently being acquired

- “Last Shot Acquired” indicator: indicates the value of the experimental parameter for the previous loop of the timing file.
- “Alter Channels Indicators*” box of indicators: displays several of the indicators from *Alter Channels.vi*. These are provided so you can conveniently track the progress of each acquisition.

To begin the data run, simply set “Queue Offset” to the desired value (usually zero) and press “Start Acquisition Sequence.” A dialog box will appear. The dialog box reminds you to perform tasks that will be discussed in more detail in the remaining steps.

4. The dialog box first reminds you that you need to set the “Trial Number” control to the right value. Next, you are reminded that you should have already saved a background image in the Andor MCD program (this task should have been completed in step 1). Next, you are told to select *FK_AR_autolevel_autosave.pgm* using “Run program by filename” (available in the File menu) in the Andor MCD program. When you have opened this program in Andor MCD, search for line of code which defines *datadirectory\$*. Set *datadirectory\$* equal to the desired filepath for the data you will save. When you start running *FK_AR_autolevel_autosave.pgm*, the Andor MCD program will display a dialog box reminding you to properly set the target directory for the data. Press any key to dismiss this dialog box. With the Andor MCD program now running, you can return to *AutomateAcquisition.vi* and press “CONTINUE” to dismiss the dialog box reminding you to complete all the tasks discussed in this step.
5. At this point, both the Andor MCD program and *AutomateAcquisition.vi*

should be running. The VI will run the system with a given value of the timing parameter of interest, and the Andor MCD program will automatically save the data with the appropriate filename. The VI will loop through all the desired values of the timing parameter and then stop. If you wish to stop the data acquisition process prematurely, press “STOP.” The VI will complete the data acquisition for the present value of the timing parameter and then stop data acquisition. When the end of the data run is reached (either when you press “STOP” or when the VI naturally completes the data loop), another dialog box will appear. This dialog box urges you to press the “Stop” button in the Andor MCD program once it is finished saving the most recent data shot. If you do not stop the Andor MCD program, it will save files with an incorrect filename when you start acquiring images again.

Note: at the beginning of a data run, *AutomateAcquisition.vi* will loop through the first value of the experimental parameter 3 times. This is done because we have noticed that system performance varies depending on the duration of MOT loading. To be certain that all data is acquired under similar conditions, the Andor MCD program will not save any data until the third shot at the first value of the experimental parameter.

C.2.3 Operation of *AutomateAcquisition.vi*

This section describes in more detail the operation of the VI. The information that follows will be useful if you encounter some bugs in the VI or if you attempt to upgrade the VI at some point in the future.

The initial stages of the program are fairly straightforward. First, you read

in a given timing file. Next, you tell the VI which timing parameter you wish to vary, and then you provide the VI with a start value, an end value, and a step size for this timing parameter. (Alternatively, you can provide the program with a comma-separated list of timing parameter values). When you instruct the VI to prepare the acquisition sequence, it generates a different timing file for each value of the timing parameter. Suppose you want to vary the *\$time_of_flight* variable from 200 μs to 500 μs in steps of 100 μs . When you press “PREPARE ACQUISITION SEQUENCE” in the “Configure Acquisition Sequence” tab, the VI will generate 4 timing files: 200.txt, 300.txt, 400.txt and 500.txt. The timing files are generated in the following manner: the VI searches through the original timing file and locates the timing parameter variable name. It then sets the parameter equal to the “Start Value”. It saves a new timing file with the name (Start Value).txt. It then repeats the process for all other values of the timing parameter. These timing files are placed in the same directory as the original timing file.

Next, the VI randomizes the order of these newly generated timing files. This is done to compensate for systematic drifts in the experimental apparatus. When you press the “START ACQUISITION SEQUENCE” button on the “Take Data” tab, the VI takes the first timing file from the randomized list of timing files, commands *Alter Channels.vi* to load the timing file, and then run it. (Note that a copy of *Alter Channels.vi* appears in the diagram screen for *AutomateAcquisition.vi*). At the beginning of each data run, the first timing file in the queue will be run 3 times. This is done in the hope that the user will save only the third piece of data, thereby minimizing the impact of long MOT loading times on system performance. When the end of an acquisition run is reached, all of the

timing files (200.txt, 300.txt, 400.txt and 500.txt, in the example introduced in the previous paragraph) are deleted. If you press “STOP” before the end of an experimental run, these files will not be deleted.

As *AutomateAcquisition.vi* works its way through the list of timing files (200.txt, 300.txt, 400.txt and 500.txt, in the example scenario) it saves a list of filenames to a text file called *NameHolding.txt* located in the following directory: Jet-project/jetlab/Network Harddrives for Lab Computers/LAB1_DOCS/Computer Control System/Control VIs and Programs/Timing System – Automated Acquisition. The first line of this text file lists the filename to be used by *FK_AR_autolevel_autosave.pgm* in the Andor MCD program when saving the first file. The second piece of data will be saved under the filename listed on the second line of the *NameHolding.txt* file, and so on. Unfortunately, *FK_AR_autolevel_autosave.pgm* is written in Andor BASIC, which has a puny set of file handling capabilities. When the *read()* function is called in Andor BASIC, the program will read only the first line of text in a text file. If you call the *read()* function yet again for the same text file, it will read the second line of that text file, and so on. As long as the name of the text file does not change, Andor BASIC will continue to read successive lines of that text file. **As a result, the user must always remember to press STOP in the Andor MCD program at the end of a data acquisition run. If not, the Andor MCD program and *AutomateAcquisition.vi* will fall out of sync.** Consider: each time you start a new data run in *AutomateAcquisition.vi*, the VI will delete the old *NameHolding.txt* file and begin generating a new one. Consequently, if you failed to stop Andor MCD after the most recent data run (suppose you took 40 shots in the data run), it will look on the 41st line of *NameHolding.txt* file for

a filename. Of course, at this point, *AutomateAcquisition.vi* has already deleted the old *NameHolding.txt* file and generated a new one which is only one line long after the first piece of data is acquired. Andor MCD will get cranky, which will undoubtedly make you cranky as well. The VI will display dialog boxes reminding the user to avoid this problem by stopping the Andor MCD program at the end of each data run.

Finally, it is important to note one more aspect about the way in which pieces of data are automatically saved: *AutomateAcquisition.vi* and *FK_AR_autolevel_autosave.pgm* are written under the assumption that the read-out time for the camera CCD array is long. How long? Long enough that by the time the image is finally ready to be saved, *AutomateAcquisition.vi* has already started running the next timing file. It is important to note that the VI does not update *NameHolding.txt* until it reaches the end of a particular timing file. This should not be a problem under most circumstances since the camera shots are usually taken less than 2 seconds from the end of a timing file, and it typically takes close to 10 or 15 seconds to read out a camera shot. By this time, *NameHolding.txt* file will have been updated, and LabVIEW and Andor will be in sync. If short readout times are used, however, then pieces of data might not be saved with the appropriate name.

C.3 *Arbitrary Waveforms.vi*

In the process of performing forced evaporation, two Agilent 33250A signal generators are used to control the depth of the FORT via amplitude modulation of the RF signal supplied to an IntraAction AO at the CO₂ laser's output. The

mathematical expression used to determine the FORT depth as a function of time is given in [23,24], and when used in conjunction with measurements of the laser power transmitted through the AO as a function of the AO crystal's RF drive amplitude, the user can establish the relationship between the amplitude control voltage directed to the AO and the resulting FORT trap depth. This AO amplitude control voltage versus time curve is referred to as the "lowering curve."

When the Agilent 33250A signal generators were first added to the system to control the process of forced evaporation, the process of programming them involved a number of steps. First, the user needed to modify the timing file to establish the desired lowering curve parameters (how long to lower the trap, etc.). Next, these same parameters needed to be updated in a Perl script. This Perl script would then be executed, and would produce several output files which contained the lowering curves in various formats. One of these formats was designed to be used with Agilent's software for controlling the 33250A signal generators. After the lowering curves were imported into the Agilent program, they were sent via GPIB cable to the signal generators. Once the lowering curves were stored in the signal generator's memory, the user would then need to walk to the signal generator and execute a series of keystrokes on the front panel to configure the generator to output the lowering curve once upon receiving an external trigger. This process would need to be completed for each signal generator.

Clearly, the sequence described in the previous paragraph could become tedious, especially when experimenting with lowering curve parameters to determine the optimum lowering sequence. *Arbitrary Waveforms.vi* was written to automate many of these steps. In Section C.3.1, I will describe the proper usage of *Arbitrary Waveforms.vi*, and in Section C.3.2 I will discuss the operation of the VI in

greater detail.

C.3.1 Using *Arbitrary Waveforms.vi*

The front panel of *Arbitrary Waveforms.vi* is displayed in Figure C.2. There are 8 subsections of the front panel and their outlines are denoted by the green boxes. Each of these subsections will be described in turn. The box in the upper left corner is labelled “GPIB Addresses” and contains two numeric controls with which the user can specify the GPIB address of the two Agilent 33250A arbitrary waveform generators which are used in the lowering curve sequence. Below this is a box containing two push buttons, one labelled “Disable Outputs” and the other labelled “Send.” Pressing the Disable Outputs button will disable the outputs of both arbitrary waveform generators while pressing the Send button will initiate the process of creating the lowering curves and sending them to the arbitrary waveform generators (this latter process will be described in greater detail throughout the remainder of this section). To the right of these boxes is a subsection named “File Path Specifications.” Here, the user can define the path to reach the root directory of the timing system, the path to the Perl script for the lowering curve, and the path to the timing file directory. These path constants should be changed only when alterations in the timing system control directory structure have been made. To the left of the File Path Specifications area is a section labelled “Input Parameters.” This panel contains the most critical input controls for the VI. The dropdown list “Lowering Curve Generator” allows the user to select the Perl script to be used for generating the lowering curves. The “Select Timing File” box allows the user to select the desired timing file using the numeric control at the left. The “Tau (sec),” “Recompress to,” “Offset (mV),”

“Tau 2 (sec),” and “Final Depth” controls are used to set critical parameters for the lowering curves. These controls will be discussed in greater detail in a moment. Below the “File Path Specifications” and “Input Parameters” boxes is a box with a horizontal row of round LED indicators with a status key shown at the left. Above each LED is a descriptive name of a process that occurs when the user presses the Send button. When the LED is dark green, the process has not begun. When the LED is blinking yellow, the associated process is in progress. When the LED switches to a bright green color, the process has been completed. When all of the LEDs have turned bright green, the process of sending the arbitrary waveforms has been completed. Below the box of LEDs is a section that contains two graphs showing the lowering curves being sent to the two arbitrary waveform generators. To the right of this box is a small box containing a push button labelled “VIEW NOTES.” Pressing this button will display a dialog box showing helpful hints on the usage of the VI. Above this box is a subsection of the front panel labelled “Lowering Parameters.” There are a number of indicators in this box which show the parameters associated with the lowering curve which were read from the user-selected timing file.

While this array of controls and indicators might seem daunting at first, the learning curve for using *Arbitrary Waveforms.vi* is not especially steep. The following discussion will help acclimate you to the VI while describing the basics of its operation. For a more detailed nuts-and-bolts discussion of the VI’s operation, please see Section C.3.2.

Before delving into too much detail, it is worth describing how the Agilent 33250A arbitrary waveform generators and *Arbitrary Waveforms.vi* fit into the grand scheme of an experimental sequence. The process of forced evaporation

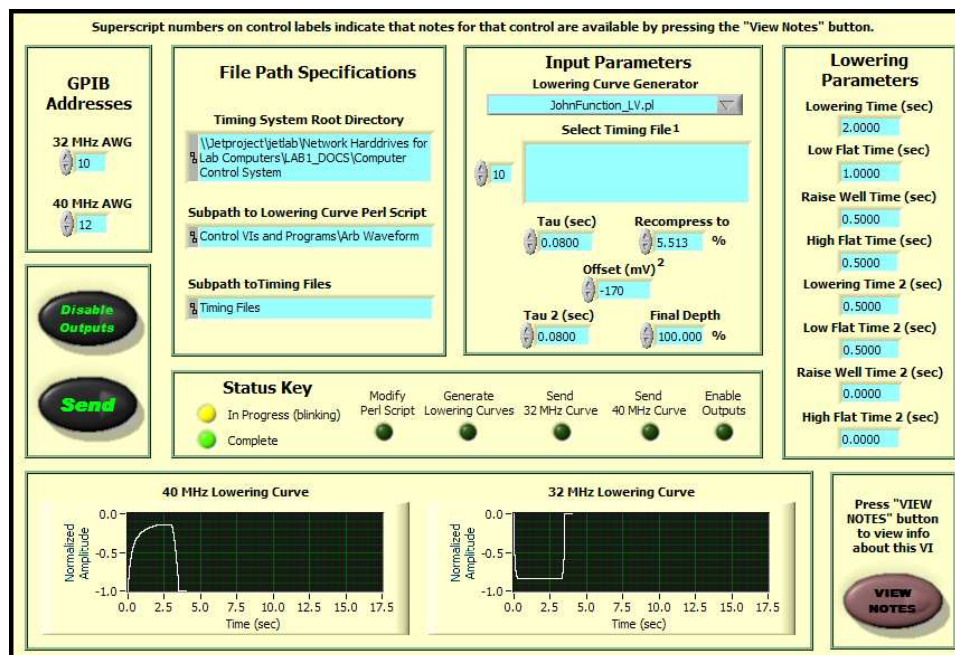


Figure C.2: Front panel of *Arbitrary Waveforms.vi*. The operation of the VI is described in the text.

involves a number of steps. First, the bias magnetic field must be set to the appropriate value; this is accomplished using the timing file and a power supply on the system. With the bias magnetic field at the desired value, the depth of the FORT can be lowered. This process is controlled by the 33250A arbitrary waveform generators. The generators receive an external trigger from the 32 channel breakout panel and begin to output waveforms which control the lowering (and reraising) of the FORT. These lowering curve waveforms must be stored in the arbitrary waveform generators ahead of time using *Arbitrary Waveforms.vi* (that is, the lowering curves must be sent to the arbitrary waveform generators before the timing file is run). Once the process of forced evaporation and FORT recompression is completed, other elements in the system handle the remaining experimental manipulations. Of course, these other devices on the experimental

system must have information regarding the lowering curve sequence and duration. As a result, many of the lowering curve parameters must be defined in the timing file, so all of the devices on the system have access to that information. However, there are some experimental parameters which are relevant to only the arbitrary waveform generators, and these parameters are defined on the front panel of *Arbitrary Waveforms.vi*. It is worthwhile outlining where each of these parameters are defined, as well as their meaning. We will begin with a discussion of each parameter's meaning.

A single lowering sequence consists of four time intervals, as seen in Figure C.3. The first time period is referred to as the “lowering time,” in which the trap depth $U(t)$ obeys

$$U(t) = U_0 \left(\frac{1}{1 + t/\tau_{lc}} \right)^{1.45}, \quad (\text{C.1})$$

where U_0 is the initial trap depth, t is the time elapsed from the start of lowering, and τ_{lc} is the lowering time constant [23, 24]. Suppose we lower the trap depth over a time period denoted by t_{lower} . Then $U(t_{lower})$ is the trap depth at the conclusion of the lowering sequence. The trap depth is maintained at this value for a length of time referred to as the “low flat time.” At the conclusion of “low flat time,” the trap depth is reraised to a user-specified value U_{final} over a time period defined by the “raise well time,” t_{raise} . An exponential curve governs the reraising process, such that if the time origin $t = 0$ corresponds to the start of reraising, we have

$$U(t) = U_{final} \exp \left\{ \frac{(t_{raise} - t) \ln [U(t_{lower})/U_{final}]}{t_{raise}} \right\}. \quad (\text{C.2})$$

Finally, the trap is held at depth U_{final} for a duration specified by “high flat

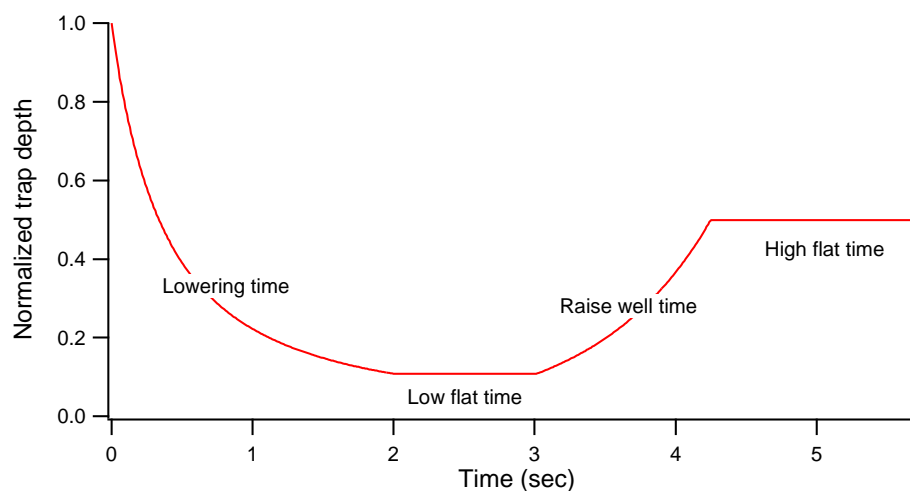


Figure C.3: The normalized trap depth is plotted as a function of time after the start of the forced evaporation sequence. For a single lowering curve sequence, there are four distinct time periods: the lowering time, the low flat time, the raise well time, and the high flat time.

time.”

The combination of “lowering time,” “low flat time,” “raise well time,” and “high flat time” constitutes a single lowering curve sequence. When this scheme for conducting forced evaporation was first implemented, a single lowering curve sequence was all that was required for each experimental cycle. Since that time, however, more complex experimental manipulations have required expanding the capabilities of the system to conduct two lowering curve sequences in succession. Such a lowering curve is displayed in Figure C.4. The second lowering and reraising sequence proceeds in exactly the same fashion as the first lowering and reraising sequence. The time dependence of the trap depth during “lowering time 2” is governed by (C.1) and the time dependence of the trap during “raise well time 2” is given by (C.2). Finally, it is important to note that the time constants and time durations used in the first and second lowering sequences are completely in-




Figure C.4: The normalized trap depth is plotted as a function of time after the start of forced evaporation. For the dual lowering curve sequence, there are eight distinct time periods, as labelled on the above graph.

dependent. Different lowering time constants can be used during “lowering time” and “lowering time 2,” just as different trap depths can be selected for “high flat time” and “high flat time 2.”

Now that the various portions of the lowering curve have been identified, I will outline where these various parameters are defined in the custom software. The durations of the different parts of the lowering curve sequence are defined in individual timing files. Since *Arbitrary Waveforms.vi* must have access to this information, there are precise naming conventions that must be used when writing the timing files. These names are nearly identical to the names shown in Figure C.4, with some slight modifications. As the timing files are processed by a Perl script, and as variables in Perl must be preceded by a \$ sign and cannot have spaces in them, the various lowering names must be slightly modified. The following is the name list that must be employed when writing timing files if those timing files are to be used successfully by *Arbitrary Waveforms.vi*:

\$lowering_time, \$low_flat_time, \$raise_well_time, \$high_flat_time, \$lowering_time2, \$low_flat_time2, \$raise_well_time2, and \$high_flat_time2.

As noted earlier, there are several lowering curve parameters that are needed by *Arbitrary Waveforms.vi* alone, and these parameters can be entered directly from the front panel of the VI. The dropdown list “Lowering Curve Generator” allows the user to select the Perl script to be used for generating the lowering curves. The user must also specify the desired timing file using the “Select Timing File” control. The control labelled “Tau (sec)” corresponds to the value of τ_{lc} used in (C.1) for the first lowering sequence, while “Tau 2 (sec)” is the corresponding lowering time constant for the second lowering curve sequence. The “Recompress to” control lets the user specify the desired trap depth during “high flat time” in percent (with 100% representing maximum trap depth). The “Final Depth” control lets the user set the desired trap depth during “high flat time 2.” Finally, to obtain the most effective cooling, it is convenient to have a control that permits the user to provide a small DC shift in the voltage which controls the trap depth. This degree of freedom is provided by the “Offset (mV)” control. By setting this control to a negative value, the trap depth during “low flat time” and “low flat time 2” can be made shallower.² Or, if the user wishes to have a deeper trap at the lowest well depth, “Offset (mV)” should be set to a positive value. Note that as there is only one “Offset” control, the user does not have the freedom to use different offsets for the two lowering curve sequences. At present, we find that setting “Offset (mV)” to a value between -120 and -180 mV provides the most effective evaporative cooling.

²In actuality, setting the Offset control to a negative value makes the trap shallower throughout the entire lowering curve sequence. However, the effect is most apparent at the lowest trap depths.

Now that most of the necessary background information has been developed, the following subsection will provide the sequence of steps one should follow to use *Arbitrary Waveforms.vi*.

Step-by-step instructions for using *Arbitrary Waveforms.vi*

The following is a step-by-step guide to using *Arbitrary Waveforms.vi*.

1. Begin by setting the \$lowering_time, \$low_flat_time, \$raise_well_time, \$high_flat_time, \$lowering_time2, \$low_flat_time2, \$raise_well_time2, and \$high_flat_time2 variables to the desired values from within the desired timing file. If you wish to have only one lowering and reraising sequence in the timing file, set all variables that end in 2 equal to zero. Be sure to save the timing file after the alterations are complete.
2. Open *Arbitrary Waveforms.vi* and make sure that it is running. Select the desired lowering curve from the drop down list on the front panel.
3. Select the timing file you modified in Step 1 in the “Select Timing File” control on the front panel.
4. On the front panel of the VI, enter the desired values for the following variables: “Tau (sec),” “Tau 2 (sec),” “Offset (mV),” “Recompress to,” and “Final Depth.” Note that if you are going to use only one lowering and reraising sequence, then “Tau 2 (sec)” and “Final Depth” will be irrelevant.
5. Double check that all of your timing parameters are set to the desired values, and then press the “Send” button on the front panel. The VI will read the timing parameters from the timing file and display them in the “Lowering

Parameters” pane on the front panel. Verify that these numbers match those you entered in the timing file.

6. The VI will output the lowering curves to the arbitrary waveform generators. You can monitor the progress of the VI by watching the blinking LEDs in the center of the front panel. When the process is complete, a dialog box will be displayed informing the user that the arbitrary waveform generators are ready. Pressing “OK” dismisses the dialog box and enables the outputs of the arbitrary waveform generators.
7. The lowering curves are now stored in the arbitrary waveform generators and the generators will output those waveforms upon receiving a trigger from the 32 channel breakout panel. You can now run the timing file from within *Alter Channels.vi* or *Automate Acquisition.vi*.

If you wish to disable the outputs of the arbitrary waveform generators, simply press the “Disable Outputs” button on the front panel of the VI. Finally, before moving on to a technical discussion of the operation of *Arbitrary Waveforms.vi* in Section C.3.2, I will outline the steps to take when you wish to add a lowering curve or a timing file to the available lists of files on the front panel.

Adding lowering curves to the drop-down list

Adding new lowering curves to the drop-down list in *Arbitrary Waveforms.vi* is a multi-step process which is described in the Notes section of the VI (which can be accessed by pressing the “VIEW NOTES” button on the VI’s front panel when the VI is running). The instructions for adding a new lowering curve are repeated here.

1. Stop the VI. On the front panel, select the last item in the dropdown list labelled “Lowering Curve Generator”. Next, right click on the drop-down list and select “Add Item After.” A text cursor will begin flashing in the drop down list. Type the name of the lowering curve you wish to add to the list and hit “Enter” on the numeric keypad.
2. Switch to the diagram, and look at frame #1 in the 10-frame sequence structure in the center of the diagram window. (If you cannot see this sequence structure, you might need to change the value of the very large case structure.) Inside the 10-frame sequence structure is a small case structure which lists the names of the lowering curves. Make sure you are looking at the highest numbered case in the case structure. Right click on the header in the case structure and select “Add Case After.” LabVIEW will automatically generate a name (in this case, a number) for the case you wish to add. Hit “Enter” to select this number. Now, place a string constant containing the name of the new lowering curve in the new case window and wire it to the string exit tunnel on the left side of the case structure.

The new lowering curve should now be available to you.

Adding new timing files to the list in *Arbitrary Waveforms.vi*

Adding a new timing file to the list of available timing files in *Arbitrary Waveforms.vi* is relatively straightforward. When the VI is first opened and the user presses the Run button, the VI will generate a list of all the timing files in the timing files directory. Consequently, if you created a new timing file after you started running *Arbitrary Waveforms.vi*, you will need to stop and restart the VI

in order for it to recognize the new timing file.

C.3.2 Operation of *Arbitrary Waveforms.vi*

The preceding subsections were designed to present the basic knowledge one needs to use *Arbitrary Waveforms.vi*. This subsection goes into greater detail regarding the operation of *Arbitrary Waveforms.vi* and is intended to be of assistance to those who wish to expand or modify the existing code. An explanation of the operation of the code is most easily conducted by considering the sequence of events that occur when the “Send” button is pressed on the front panel of the VI, so that will be the general method employed in this section. Please note that the details below are accurate for version 8.0 of *Arbitrary Waveforms.vi*. If someone fails to update the version number when saving changes to the VI following an upgrade, these notes might not be applicable.

If you look at the diagram for *Arbitrary Waveforms.vi*, you will note a very large sequence structure with two frames. In the zeroth frame is a simple bit of code which reads the list of timing files from the timing file directory. This is the first event that occurs when you first run the VI. Note also that this is the only time that the VI checks to see which timing files are available. As noted earlier, if you write a new timing file after you have started *Arbitrary Waveforms.vi*, you will need to stop and restart the VI before it will recognize the new timing file.

The next and final frame in the very large sequence structure contains the remainder of the code for the VI. Contained inside the very large frame of the sequence structure is a large While loop. The large While loop has three essential elements: a small case structure toward the bottom which allows the user to display a notes dialog after pressing “VIEW NOTES” on the front panel of the VI, a

small While loop that disables the outputs of the arbitrary waveform generators when the “Disable Outputs” button is depressed on the front panel, and a large While loop which contains the code that determines the sequence of events that occur when the user presses the “Send” button on the front panel. As the operation of the first two elements in the preceding list are relatively straightforward, I will concentrate on the contents of the large While loop which contains the code for sending the waveforms to the arbitrary waveform generators.

Inside the large While loop is a large case structure whose input is the “Send” button on the front panel of the VI. If the “Send” button is not pressed, no action is taken (the “False” instance of the case structure). However, when the “Send” button is pressed, and the “True” case of the case structure is accessed, a minute-long series of events occurs. The sequencing of these events is controlled by an 11-frame sequence structure. The following enumerated list details the actions of each frame in the sequence structure. The lead number in the following list corresponds to the numeric frame identifier shown at the top of the sequence structure.

0. Both arbitrary waveform generators are initialized using the sub-VIs designed for the Agilent 33250A waveform generators. Further, all of the indicator LEDs are reset to the dark green state. Finally, at a later stage, a new Perl script named Modified Lowering Curve.pl will be generated. This frame deletes this file that remained from the previous time the program was run.
1. Here the outputs are disabled for both arbitrary waveform generators. Also in this frame, the lowering curve Perl script specified by the user on the

front panel is read and converted into a string.

2. There are two sequence structures embedded in this frame of the larger sequence structure. In the small sequence structure in the lower left corner of this frame, the timing file specified by the user is read and converted to a string. This string is then searched for each of the timing parameters (`$lowering_time`, `$low_flat_time`, etc.). Note that the parameters from the timing file will not be identified properly unless they have precisely the names that are listed in this sequence structure. The string values of the timing parameters are then converted to numeric values and displayed on the front panel of the VI. The string values of the lowering curve parameters read from the timing file are also passed to the sequence structure in the upper portion of this frame. This sequence structure searches the string contents of the lowering curve Perl script for specific strings and replaces those strings with the values read from the timing file. The original Perl script which was read in Step 1 does not contain the numeric values for the lowering parameters. The locations in the Perl script where we wish to insert numeric values for the lowering parameters are marked by dummy strings with names such as “inputlowertime” and “inputlowflattime.” The LabVIEW code will replace the string “inputlowertime” in the Perl script with the numeric value of `$lowering_time` from the timing file. The next frame of the sequence structure replaces the string “inputlowflattime” with the numeric value of `$low_flat_time` read from the timing file, and so on. For the correspondence between the dummy strings in the lowering curve Perl script and the associated source of the numeric values, see Table C.1. Once all of the dummy strings have been replaced with numeric values drawn from

<i>Dummy string in Perl script</i>	<i>Variable name in timing file</i>	<i>Front panel control</i>
inputlowertime	\$lowering_time	–
inputlowflattime	\$low_flat_time	–
inputraisetime	\$raise_well_time	–
inputflattime	\$high_flat_time	–
inputloweringtime2	\$lowering_time2	–
inputlowflattime2	\$low_flat_time2	–
inputraisewelltime2	\$raise_well_time2	–
inputhighflattime2	\$high_flat_time2	–
inputtau	–	Tau (sec)
inputtau2	–	Tau 2 (sec)
inputfracdepth	–	Recompress to
inputfracdepth2	–	Final Depth

Table C.1: The above table gives the correspondence between the dummy variables in the lowering curve Perl script and the associated numeric values in the timing file or the front panel of Arbitrary Waveforms.vi. A dash indicates that information regarding that particular timing parameter is not located in the information source (timing file or front panel of VI) which labels that column.

the timing file or the front panel of the VI, the resulting string is passed to the next frame in the sequence structure.

3. The string containing the modified Perl script is written to a new Perl script named *Modified Lowering Curve.pl* which is placed in the same directory as the original Perl script.
4. Here, the modified Perl script (*Modified Lowering Curve.pl*) is executed. The execution of the Perl script results in the creation of several new files. Of these, the files *amp35wave.txt* and *amp40wave.txt* will be used in later frames in this sequence structure.

5. This frame forces a 0.5 second pause in the execution of the program to be certain that the amp35wave.txt and amp40wave.txt files have been generated.
6. This frame sends GPIB commands to the arbitrary waveform generator which controls the amplitude of the ≈ 32 MHz signal sent to the AO. Note that for historical reasons some of the files associated with the 32 MHz signal are labelled with “35” (amp35wave.txt, for example). At the top of the frame, the durations of the various lowering curve segments are added together to calculate the period of the lowering curve. This is converted to a frequency and added to an array of GPIB commands to be sent to the arbitrary waveform generator. Next, the upper and lower voltage limits are specified. The maximum voltage for the arbitrary waveform generator is specified by the variable \$stopamp35 in the lowering curve Perl script. Then, the text file amp35wave.txt (which was generated in Step 4) is read and the data from this file is converted to integer values which maximize the output resolution of the arbitrary waveform generator. Finally, the arbitrary waveform generator is configured to operate in burst mode, outputting the arbitrary waveform once upon receiving an external trigger. The output impedance of the arbitrary waveform generator is set to 2 k Ω .
7. This frame forces a 0.1 second pause between the sending of data to the first and second arbitrary waveform generators.
8. This frame conducts essentially the same steps for the 40 MHz generator as those done in Step 6 for the ≈ 32 MHz generator, with the following exceptions. The maximum magnitude of the 40 MHz control voltage is

specified by the variable \$stopamp40 in the lowering curve Perl script, and the text file amp40wave.txt is used as the source of the data points for the lowering curve. Finally, note that this is the frame where the “Offset (mV)” control comes into play. The amplitude of the lowering curve can be effectively stretched or compressed by a small amount to either raise or lower the trap depth.

9. This is the error-checking frame. Both arbitrary waveform generators are sent an error query. If either generator experienced an error, that error is reported to the user and the user is given the option of enabling the outputs of both generators (even though one or both experienced an error) or leaving the generators disabled. If no errors are encountered, the user is informed that both generators are ready to continue. The user can press “OK” to dismiss this dialog box, at which point both arbitrary waveform generators will be enabled.
10. Finally, some of the indicator LEDs are reset to their default colors.

In closing, the reader is cautioned once more that the preceding discussion is applicable for version 8.0 of *Arbitrary Waveforms.vi*. Newer versions might feature code which is not described in the preceding discussion.

Bibliography

- [1] A. Cho. Ultracold atoms spark a hot race. *Science*, 301:750, 2003.
- [2] G. P. Collins. The next big chill. *Scientific American*, page 26, October 2003.
- [3] J. E. Thomas and M. E. Gehm. Optically trapped Fermi gases. *American Scientist*, 92:238, 2004.
- [4] M. Tinkham. *Introduction to superconductivity*. McGraw-Hill, New York, 1966.
- [5] Q. Chen, J. Stajic, S. Tan, and K. Levin. BCS-BEC crossover: From high temperature superconductors to ultracold superfluids. *Physics Reports*, 412:1, 2005.
- [6] M. Holland, S. J. J. M. F. Kokkelmans, M. L. Chiofalo, and R. Walser. Resonance superfluidity in a quantum degenerate Fermi gas. *Phys. Rev. Lett.*, 87:120406, 2001.
- [7] S. J. J. M. F. Kokkelmans, J. N. Milstein, M. L. Chiofalo, R. Walser, and M. J. Holland. Resonance superfluidity: renormalization of resonance scattering theory. *Phys. Rev. A*, 65:053617, 2002.
- [8] J. Stajic, J. N. Milstein, Q. Chen, M. L. Chiofalo, M. J. Holland, and K. Levin. Nature of superfluidity in ultracold Fermi gases near Feshbach resonances. *Phys. Rev. A*, 69:063610, 2004.
- [9] A. Perali, P. Pieri, L. Pisani, and G. C. Strinati. BCS-BEC crossover at finite temperature for superfluid trapped Fermi atoms. *Phys. Rev. Lett.*, 92:220404, 2004.
- [10] Joseph Kinast, Andrey Turlapov, John E. Thomas, Qijin Chen, Jelena Stajic, and Kathryn Levin. Heat capacity of a strongly interacting Fermi gas. *Science*, 307:1296–1299, 2005.
- [11] B. Müller. Physics of the quark-gluon plasma. nucl-th/9211010, November 1992.

- [12] U. Heinz. The quark-gluon plasma at RHIC. *Nucl. Phys. A*, 721:30, 2003.
- [13] T. Ludlam. Experimental results from the early measurements at RHIC; hunting for the quark-gluon plasma. *Nuclear Physics A*, 750:9, 2005.
- [14] First three years of operation of RHIC. *Nuclear Physics A*, 757:1, 2005.
- [15] E. Shuryak. Why does the quark-gluon plasma at RHIC behave as a nearly ideal fluid? *Prog. Part. Nucl. Phys.*, 53:273, 2004.
- [16] K. M. O'Hara, S. L. Hemmer, M. E. Gehm, S. R. Granade, and J. E. Thomas. Observation of a strongly interacting degenerate Fermi gas of atoms. *Science*, 298:2179, 2002.
- [17] Private communication between John Thomas and Ed Shuryak.
- [18] M. E. Gehm, S. L. Hemmer, S. R. Granade, K. M. O'Hara, and J. E. Thomas. Mechanical stability of a strongly interacting Fermi gas of atoms. *Phys. Rev. A*, 68:011401(R), 2003.
- [19] J. Kinast, S. L. Hemmer, M.E. Gehm, A. Turlapov, and J. E. Thomas. Evidence for superfluidity in a resonantly interacting Fermi gas. *Phys. Rev. Lett.*, 92:150402, 2004.
- [20] J. Kinast, A. Turlapov, and J. E. Thomas. Damping of a unitary Fermi gas. *Phys. Rev. Lett.*, 94:170404, 2005.
- [21] M. Bartenstein, A. Altmeyer, S. Riedl, S. Jochim, C. Chin, J. Hecker Denschlag, and R. Grimm. Collective excitations of a degenerate gas at the BEC-BCS crossover. *Phys. Rev. Lett.*, 92:203201, 2004.
- [22] K. M. O'Hara. *Optical Trapping and Evaporative Cooling of Fermionic Atoms*. PhD thesis, Duke University, 2000.
- [23] S. R. Granade. *All-optical Production of a Degenerate Gas of ^6Li : Characterization of Degeneracy*. PhD thesis, Duke University, 2002.
- [24] M. E. Gehm. *Preparation of an Optically-trapped Degenerate Fermi gas of ^6Li : Finding the Route to Degeneracy*. PhD thesis, Duke University, 2003.
- [25] A. Einstein. Quantentheorie des einatomigen idealen gases. Zweite abhandlung. *Sitzungsberichte, Preussische Akademie der Wissenschaften*, page 3, 1925.

- [26] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell. Observation of Bose-Einstein condensation in a dilute atomic vapor. *Science*, 269:198, 1995.
- [27] C. C. Bradley, C. A. Sackett, J. J. Tollett, and R. G. Hulet. Evidence of Bose-Einstein condensation in an atomic gas with attractive interactions. *Phys. Rev. Lett.*, 75:1687, 1995.
- [28] K. B. Davis, M. O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle. Bose-Einstein condensation in a gas of sodium atoms. *Phys. Rev. Lett.*, 75:3969, 1995.
- [29] D. J. Griffiths. *Introduction to Quantum Mechanics*. Prentice Hall, Upper Saddle River, NJ, 1995.
- [30] R. Shankar. *Principles of Quantum Mechanics*. Plenum, New York, 2nd edition, 1994.
- [31] J. J. Sakurai. *Modern Quantum Mechanics*. Addison-Wesley, New York, revised edition, 1994.
- [32] C. Cohen-Tannoudji, B. Diu, and F. Laloë. *Quantum Mechanics*, volume One. John Wiley & Sons, New York, 1977.
- [33] D. A. Butts and D. S. Rokhsar. Trapped Fermi gases. *Phys. Rev. A*, 55:4346, 1997.
- [34] K. Huang and C. N. Yang. Quantum mechanical many-body problem with hard-sphere interaction. *Phys. Rev.*, 105:767, 1957.
- [35] R. K. Pathria. *Statistical Mechanics*. Butterworth-Heinemann, Boston, 2nd edition, 1996.
- [36] H. Heiselberg. Fermi systems with long scattering lengths. *Phys. Rev. A*, 63:043606, 2001.
- [37] G. A. Baker Jr. Neutron matter model. *Phys. Rev. C*, 60:054311, 1999.
- [38] J. Carlson, S.-Y. Chang, V. R. Pandharipande, and K. E. Schmidt. Superfluid Fermi gases with large scattering length. *Phys. Rev. Lett.*, 91:050401, 2003.
- [39] S. Y. Chang, V. R. Pandharipande, J. Carlson, and K. E. Schmidt. Quantum Monte-Carlo studies of superfluid Fermi gases. *Phys. Rev. A*, 70:043602, 2004.

- [40] A. Perali, P. Pieri, and G.C. Strinati. Quantitative comparison between theoretical predictions and experimental results for the BCS-BEC crossover. *Phys. Rev. Lett.*, 93:100404, 2004.
- [41] M. Bartenstein, A. Altmeyer, S. Riedl, S. Jochim, C. Chin, J. Hecker Denschlag, and R. Grimm. Crossover from a molecular Bose-Einstein condensate to a degenerate Fermi gas. *Phys. Rev. Lett.*, 92:120401, 2004.
- [42] T. Bourdel, L. Khaykovich, J. Cubizolles, J. Zhang, F. Chevy, M. Teichmann, L. Tarruell, S. J. J. M. F. Kokkelmans, and C. Salomon. Experimental study of the BEC-BCS crossover region in Lithium 6. *Phys. Rev. Lett.*, 93:050401, 2004.
- [43] G. B. Partridge, W. Li, R. I Kamar, Y. Liao, and R. G. Hulet. Pairing and phase separation in a polarized Fermi gas. *Science*, 10.1126/science.1122876, 2005.
- [44] T.-L. Ho. Universal thermodynamics of degenerate quantum gases in the unitarity limit. *Phys. Rev. Lett.*, 92:090402, 2004.
- [45] T.-L. Ho and E. J. Mueller. High temperature expansion applied to fermions near Feshbach resonance. *Phys. Rev. Lett.*, 92:160404, 2004.
- [46] H. Feshbach. A unified theory of nuclear reactions II. *Ann. Phys.*, 19:287, 1962.
- [47] A. J. Moerdijk, B. J. Verhaar, and A. Axelsson. Resonances in ultracold collisions of ${}^6\text{Li}$, ${}^7\text{Li}$ and ${}^{23}\text{Na}$. *Phys. Rev. A*, 51:4852, 1995.
- [48] M. Houbiers, H. T. C. Stoof, W. I. McAlexander, and R. G. Hulet. Elastic and inelastic collisions of ${}^6\text{Li}$ atoms in magnetic and optical traps. *Phys. Rev. A*, 57:R1497, 1998.
- [49] M. Bartenstein, A. Altmeyer, S. Riedl, R. Geursen, S. Jochim, C. Chin, J. H. Denschlag, R. Grimm, A. Simoni, E. Tiesinga, C.J. Williams, and P.S. Julienne. Precise determination of ${}^6\text{Li}$ cold collision parameters by radio-frequency spectroscopy on weakly bound molecules. *Phys. Rev. Lett.*, 94:103201, 2005.
- [50] K. E. Strecker, G. B. Partridge, and R. G. Hulet. Conversion of an atomic Fermi gas to a long-lived molecular Bose gas. *Phys. Rev. Lett.*, 91:080406, 2003.
- [51] S. R. Granade, M. E. Gehm, K. M. O'Hara, and J. E. Thomas. All-optical production of a degenerate Fermi gas. *Phys. Rev. Lett.*, 88(12):120405, 2002.

- [52] S. Jochim, M. Bartenstein, G. Hendl, J. Hecker Denschlag, R. Grimm, A. Mosk, and M. Weidemüller. Magnetic field control of elastic scattering in a cold gas of fermionic lithium atoms. *Phys. Rev. Lett.*, 89:273202, 2002.
- [53] B. DeMarco and D. S. Jin. Onset of Fermi degeneracy in a trapped atomic gas. *Science*, 285:1703, 1999.
- [54] A. G. Truscott, K. E. Strecker, W. I. McAlexander, G. B. Partridge, and R. G Hulet. Observation of Fermi pressure in a gas of trapped atoms. *Science*, 291:2570, 2001.
- [55] F. Schreck, L. Kaykovich, K. L. Corwin, G. Ferrari, T. Bourdel, J. Cubizolles, and C. Salomon. Quasipure Bose-Einstein condensate immersed in a Fermi sea. *Phys. Rev. Lett.*, 87:080403, 2001.
- [56] H. J. Lee, C. S. Adams, N. Davidson, B. Young, M. Weitz, M. Kasevich, and S. Chu. Dipole trapping, cooling in traps, and long coherence times. *Atomic Physics 14*, 323:258, 1995.
- [57] C. S. Adams and E. Riis. Laser cooling and trapping of neutral atoms. *Prog. Quant. Electr.*, 21:1, 1997.
- [58] T. A. Savard, K. M. O'Hara, and J. E. Thomas. Laser-noise-induced heating in far-off resonance optical traps. *Phys. Rev. A*, 56:R1095, 1997.
- [59] M. E. Gehm, K. M. O'Hara, T. A. Savard, and J. E. Thomas. Dynamics of noise-induced heating in atom traps. *Phys. Rev. A*, 58:3914, 1998.
- [60] K. M. O'Hara, S. R. Granade, M. E. Gehm, T. A. Savard, S. Bali, C. Freed, and J. E. Thomas. Ultrastable CO₂ laser trapping of lithium fermions. *Phys. Rev. Lett.*, 82:4204, 1999.
- [61] K. Dieckmann, C. A. Stan, S. Gupta, Z. Hadzibabic, C. H. Schunck, and W. Ketterle. Decay of an ultracold fermionic lithium gas near a Feshbach resonance. *Phys. Rev. Lett.*, 89:203201, 2002.
- [62] K. M. O'Hara, S. L. Hemmer, S. R. Granade, M. E. Gehm, J. E. Thomas, V. Venturi, E. Tiesinga, and C. J. Williams. Measurement of the zero crossing in a Feshbach resonance of fermionic ⁶Li. *Phys. Rev. A*, 66:041401(R), 2002.
- [63] S. Gupta, Z. Hadzibabic, J. R. Anglin, and W. Ketterle. Collisions in zero temperature Fermi gases. *Phys. Rev. Lett.*, 92:100401, 2004.

- [64] B. Jackson, P. Pedri, and S. Stringari. Collisions and expansion of an ultracold dilute Fermi gas. *Europhys. Lett.*, 67(4):524, 2004.
- [65] C. A. Regal and D. S. Jin. Measurement of positive and negative scattering lengths in a Fermi gas of atoms. *Phys. Rev. Lett.*, 90:230404, 2003.
- [66] T. Bourdel, J. Cubizolles, L. Khaykovich, K. M. F. Magalhaes, S. J. J. M. F. Kokkelmans, G. V. Shlyapnikov, and C. Salomon. Measurement of the interaction energy near a Feshbach resonance in a ${}^6\text{Li}$ Fermi gas. *Phys. Rev. Lett.*, 91:020402, 2003.
- [67] S. Gupta, Z. Hadzibabic, M. W. Zwierlein, C. A. Stan, K. Dieckmann, C. H. Schunck, E. G. M van Kempen, B. J. Verhaar, and W. Ketterle. Radio-frequency spectroscopy of ultracold fermions. *Science*, 300:1723, 2003.
- [68] M. W. Zwierlein, Z. Hadzibabic, S. Gupta, and W. Ketterle. Spectroscopic insensitivity to cold collisions in a two-state mixture of fermions. *Phys. Rev. Lett.*, 91:250404, 2003.
- [69] C. A. Regal, C. Ticknor, J. L. Bohn, and D. S. Jin. Creation of ultracold molecules from a Fermi gas of atoms. *Nature*, 424:47, 2003.
- [70] J. Cubizolles, T. Bourdel, S. J. J. M. F. Kokkelmans, G. V. Shlyapnikov, and C. Salomon. Production of long-lived ultracold Li_2 molecules from a Fermi gas. *Phys. Rev. Lett.*, 91:240401, 2003.
- [71] S. Jochim, M. Bartenstein, A. Altmeyer, G. Hendl, C. Chin, J. Hecker Denschlag, and R. Grimm. Pure gas of optically trapped molecules created from fermionic atoms. *Phys. Rev. Lett.*, 91:240402, 2003.
- [72] C. A. Regal, M. Greiner, and D. S. Jin. Lifetime of molecule-atom mixtures near a Feshbach resonance in ${}^{40}\text{K}$. *Phys. Rev. Lett.*, 92:083201, 2004.
- [73] S. Jochim, M. Bartenstein, A. Altmeyer, G. Hendl, S. Riedl, C. Chin, J. Hecker Denschlag, and R. Grimm. Bose-Einstein condensation of molecules. *Science*, 302:2101, 2003.
- [74] M. Greiner, C. A. Regal, and D. S. Jin. Emergence of a molecular Bose-Einstein condensate from a Fermi gas. *Nature*, 426:537, 2003.
- [75] M. W. Zweirlein, C. A. Stan, C. H. Schunck, S. M. F. Raupach, S. Gupta, Z. Hadzibabic, and W. Ketterle. Observation of Bose-Einstein condensation of molecules. *Phys. Rev. Lett.*, 91:250401, 2003.

- [76] C. A. Regal, M. Greiner, and D. S. Jin. Observation of resonance condensation of fermionic atom pairs. *Phys. Rev. Lett.*, 92:040403, 2004.
- [77] M. W. Zwierlein, C. A. Stan, C. H. Schunck, S. M. F. Raupach, A. J. Kerman, and W. Ketterle. Condensation of pairs of fermionic atoms near a Feshbach resonance. *Phys. Rev. Lett.*, 92:120403, 2004.
- [78] M. W. Zwierlein, C. H. Schunck, C. A. Stan, S. M. F. Raupach, and W. Ketterle. Formation dynamics of a fermion pair condensate. *Phys. Rev. Lett.*, 94:180401, 2005.
- [79] J. Kinast, A. Turlapov, and J. E. Thomas. Breakdown of hydrodynamics in the radial breathing mode of a strongly interacting Fermi gas. *Phys. Rev. A*, 70:051401(R), 2004.
- [80] C. Chin, M. Bartenstein, A. Altmeyer, S. Riedl, S. Jochim, J. Hecker Denschlag, and R. Grimm. Observation of the pairing gap in a strongly interacting Fermi gas. *Science*, 305:1128, 2004.
- [81] M. Greiner, C. A. Regal, and D. S. Jin. Probing the excitation spectrum of a Fermi gas in the BCS-BEC crossover regime. *Phys. Rev. Lett.*, 94:070403, 2005.
- [82] M. Greiner, C. A. Regal, C. Ticknor, J. L. Bohn, and D. S. Jin. Detection of spatial correlations in an ultracold gas of fermions. *Phys. Rev. Lett.*, 92:150405, 2004.
- [83] M. Greiner, C. A. Regal, J. T. Stewart, and D. S. Jin. Probing pair-correlated fermionic atoms through correlations in atom shot noise. *Phys. Rev. Lett.*, 94:110401, 2005.
- [84] C. H. Schunck, M. W. Zwierlein, C. A. Stan, S. M. F. Raupach, W. Ketterle, A. Simoni, E. Tiesinga, C. J. Williams, and P. S. Julienne. Feshbach resonances in fermionic ${}^6\text{Li}$. *Phys. Rev. A*, 71:045601, 2005.
- [85] J. E. Thomas, J. Kinast, and A. Turlapov. Virial theorem and universality in a unitary Fermi gas. *Phys. Rev. Lett.*, 95:120402, 2005.
- [86] G. B. Partridge, K. E. Strecker, R. I. Kamar, M. W. Jack, and R. G. Hulet. Molecular probe of pairing in the BEC-BCS crossover. *Phys. Rev. Lett.*, 95:020404, 2005.
- [87] C. A. Regal, M. Greiner, S. Giorgini, M. Holland, and D. S. Jin. Momentum distribution of a Fermi gas of atoms in the BCS-BEC crossover. *Phys. Rev. Lett.*, 95, 2005.

- [88] M. W. Zwierlein, A. Schirotzek, C. H. Schunck, and W. Ketterle. Fermionic superfluidity with imbalanced spin populations. *Science*, 10.1126/science.1122318, 2005.
- [89] C. A. Regal, C. Ticknor, J. L. Bohn, and D. S. Jin. Tuning p-wave interactions in an ultracold Fermi gas of atoms. *Phys. Rev. Lett.*, 90:053201, 2003.
- [90] J. Zhang, E. G. M. van Kempen, T. Bourdel, L. Khaykovich, J. Cubizolles, F. Chevy, M. Teichmann, L. Tarruell, S. J. J. M. F. Kokkelmans, and C. Salomon. P-wave Feshbach resonances of ultracold ${}^6\text{Li}$. *Phys. Rev. A*, 70:030702(R), 2004.
- [91] C. A. Stan, M. W. Zwierlein, C. H. Schunck, S. M. F. Raupach, and W. Ketterle. Observation of Feshbach resonances between two different atomic species. *Phys. Rev. Lett.*, 93:143001, 2004.
- [92] S. Inouye, J. Goldwin, M. L. Olsen, C. Ticknor, J. L. Bohn, and D. S. Jin. Observation of heteronuclear Feshbach resonances in a mixture of bosons and fermions. *Phys. Rev. Lett.*, 93:183201, 2004.
- [93] T. A. Savard. *Raman induced resonance imaging of trapped atoms*. PhD thesis, Duke University, 1998.
- [94] A. Yariv. *Quantum Electronics*. John Wiley and Sons, New York, 2nd edition, 1975.
- [95] K. M. O'Hara, M. E. Gehm, S. R. Granade, and J. E. Thomas. Scaling laws for evaporative cooling in time-dependent optical traps. *Phys. Rev. A*, 64:051403(R), 2001.
- [96] R. C. Weast, M. J. Astle, and W. H. Beyer, editors. *CRC Handbook of Chemistry and Physics*. CRC Press, Boca Raton, 1983.
- [97] C. A. Baird. Design and characterization of a multi-coil Zeeman slower. Master's thesis, Duke University, 1996.
- [98] J. R. Gardner. *Ultra-High Resolution Atom Imaging in a Light-Shift Gradient*. PhD thesis, Duke University, 1995.
- [99] M. L. Marable. *Adaptive resonance imaging and all-optical interferometry*. PhD thesis, Duke University, 1995.
- [100] N. W. Ashcroft and N. D. Mermin. *Solid State Physics*. W. B. Saunders, Philadelphia, 1976.

- [101] P. Pedri, D. Guéry-Odelin, and S. Stringari. Dynamics of a classical gas including dissipative and mean-field effects. *Phys. Rev. A*, 68:043608, 2003.
- [102] L. D. Landau and E. M. Lifshitz. *Mechanics*, volume One. Butterworth-Heinemann, Boston, 3rd edition, 1981.
- [103] W. Ketterle, D. S. Durfee, and D. M. Stamper-Kurn. *Proceedings of the International School of Physics - Enrico Fermi*, chapter Making, probing, and understanding Bose-Einstein condensates, page 67. IOS Press, 1999.
- [104] F. Roth. Measurement of the number of atoms and of the density of magneto-optically trapped cesium, using absorption imaging and fluorescence. Master's thesis, UmeåUniversitet, 2003.
- [105] E. Arimondo, M. Inguscio, and P. Violino. Experimental determinations of the hyperfine structure in the alkali atoms. *Rev. Mod. Phys.*, 49:31, 1977.
- [106] F. London. On the Bose-Einstein condensation. *Phys. Rev.*, 54:947, 1938.
- [107] D. Guéry-Odelin. Mean field effects in a trapped gas. *Phys. Rev. A*, 66:033613, 2002.
- [108] C. Menotti, P. Pedri, and S. Stringari. Expansion of an interacting Fermi gas. *Phys. Rev. Lett.*, 89(25):250402, 2002.
- [109] B. DeMarco. *Quantum Behavior of an Atomic Fermi Gas*. PhD thesis, University of Colorado, 2001.
- [110] Z. Hadzibabic, C. A. Stan, K. Dieckmann, S. Gupta, M.W. Zwierlein, A. Görlitz, and W. Ketterle. Two-species mixture of quantum degenerate Bose and Fermi gases. *Phys. Rev. Lett.*, 8:160401, 2002.
- [111] J. Stajic, Q. Chen, and K. Levin. Density profiles of strongly interacting trapped Fermi gases. *Phys. Rev. Lett.*, 94:060401, 2005.
- [112] M. Bartenstein. *From molecules to Cooper pairs: experiments in the BEC-BCS crossover*. PhD thesis, University of Innsbruck, 2005.
- [113] Y. He, Q. Chen, and K. Levin. Radio-frequency spectroscopy and the pairing gap in trapped Fermi gases. *Phys. Rev. A*, 72:011602(R), 2005.
- [114] Q. Chen, J. Stajic, and K. Levin. Thermodynamics of interacting fermions in atomic traps. *Phys. Rev. Lett.*, 95:260405, 2005.

- [115] H. Heiselberg. Collective modes of trapped gases at the BEC-BCS crossover. *Phys. Rev. Lett.*, 93:040402, 2004.
- [116] L. Vichi. Collisional damping of the collective oscillations of a trapped Fermi gas. *Journal of Low Temperature Physics*, 121:177, 2000.
- [117] The discrepancy between Bartenstein *et al*'s measurements of the radial breathing mode frequency and predictions for a hydrodynamic gas arose from an assumption that the optical trap used in [21] was cylindrically symmetric. Accounting for trap anisotropy in the two radial dimensions yielded improved agreement between measured frequencies and predictions. Private communication between John Thomas and Rudolf Grimm.
- [118] H. Hu, A. Minguzzi, X-J Liu, and M. P. Tosi. Collective modes and ballistic expansion of a Fermi gas in the BCS-BEC crossover. *Phys. Rev. Lett.*, 93:190403, 2004.
- [119] H. Heiselberg. Pairing gaps in atomic gases at the BCS-BEC crossover. *New Journal of Physics*, 6:137, 2004.
- [120] G. M. Falco and H. T. C. Stoof. Crossover temperature of Bose-Einstein condensation in an atomic Fermi gas. *Phys. Rev. Lett.*, 92:130401, 2004.
- [121] D. Guéry-Odelin, F. Zambelli, J. Dalibard, and S. Stringari. Collective oscillations of a classical gas confined in harmonic traps. *Phys. Rev. A*, 60:4851, 1999.
- [122] P. Massignan, G. M. Bruun, and H. Smith. Viscous relaxation and collective oscillations in a trapped Fermi gas near the unitarity limit. *Phys. Rev. A*, 71:033607, 2005.
- [123] M. E. Gehm, S. L. Hemmer, K. M. O'Hara, and J. E. Thomas. Unitarity-limited elastic collision rate in a harmonically trapped Fermi gas. *Phys. Rev. A*, 68:011603(R), 2003.
- [124] B. Jackson and C. S. Adams. Damping and revivals of collective oscillations in a finite temperature model of trapped Bose-Einstein condensation. *Phys. Rev. A*, 63:053606, 2001.
- [125] Private communication between John Thomas and Kathy Levin and Qijin Chen.

- [126] M. W. Zwierlein, J. R. Abo-Shaeer, A. Schirotzek, C. H. Schunck, and W. Ketterle. Vortices and superfluidity in a strongly interacting Fermi gas. *Nature*, 435:1047, 2005.
- [127] A. H. Carter. *Classical and Statistical Thermodynamics*. Prentice Hall, Upper Saddle River, NJ, 2001.

Biography

Joseph Michael Kinast was born August 19, 1978 in Pittsburgh, Pennsylvania. In 1997, he graduated as valedictorian from Bethel Park Senior High School in Bethel Park, Pennsylvania. Four years in New Jersey yielded a B. A. from Drew University (*Summa Cum Laude*) in May 2001, with a double major in physics and psychology. Later that same year, he enrolled in graduate school at Duke University. After joining John Thomas's research group in June 2002, he contributed to studies of degenerate, strongly interacting Fermi gases. During his time at Duke, he was awarded the Charles Townes and Fritz London Fellowships. In September 2004, he was awarded an A.M. in Physics, with the Ph.D. following in May 2006.

Publications

J. E. Thomas, J. Kinast, and A. Turlapov, "Virial theorem and universality in a unitary Fermi gas," *Physical Review Letters*, **95**, 120402 (2005).

J. Kinast, A. Turlapov, and J. E. Thomas, "Damping of a unitary Fermi gas," *Physical Review Letters*, **94**, 170404 (2005).

Joseph Kinast, Andrey Turlapov, John E. Thomas, Qijin Chen, Jelena Stajic, and Kathryn Levin, "Heat capacity of a strongly interacting Fermi gas," *Science*, **307**, 1296 (2005).

J. Kinast, A. Turlapov, and J. E. Thomas, "Breakdown of hydrodynamics in the

radial breathing mode of a strongly interacting Fermi gas,” *Physical Review A*, **70**, 051401(R) (2004).

J. Kinast, S. L. Hemmer, M.E. Gehm, A. Turlapov, and J. E. Thomas, “Evidence for superfluidity in a resonantly interacting Fermi gas,” *Physical Review Letters*, **92**, 150402 (2004).

J. E. Thomas, S. L. Hemmer, J. Kinast, A. Turlapov, M. E. Gehm, and K. M. O’Hara, “Dynamics of a highly degenerate, strongly-interacting Fermi gas of atoms,” *Journal of Low Temperature Physics*, **134**, 655 (2004).

Selected Presentations

J. E. Thomas, J. Kinast, and A. Turlapov, “Universal thermodynamics of a strongly-interacting Fermi gas,” *24th International Conference on Low Temperature Physics*, Orlando, FL, presented by J. E. Thomas, August 2005.

J. Kinast, A. Turlapov, and J. E. Thomas, “Thermodynamics of a strongly interacting Fermi gas,” *APS/DAMOP 2005*, Lincoln, NE, presented by J. Kinast, May 2005.

J. E. Thomas, J. Kinast, and A. Turlapov, “Thermodynamics of a strongly interacting Fermi gas,” *Strongly Interacting Quantum Gases Workshop*, Columbus, OH, presented by J. E. Thomas, April 2005.

S. Hemmer, M. Gehm, J. Kinast, K. O’Hara, and J. E. Thomas, “Measuring universal mean field interactions in a strongly interacting Fermi gas,” *CLEO/QELS 2003*, Baltimore, MD, poster presented by J. Kinast, May 2003.

J. E. Thomas, S. L. Hemmer, J. M. Kinast, A. V. Turlapov, M. E. Gehm,

and K. M. O'Hara, "Dynamics of a highly-degenerate, strongly-interacting Fermi gas," *XVI International Conference on Laser Spectroscopy*, Palm Cove, Far North Queensland, Australia, presented by J. E. Thomas, July 2003.

J. Supplee, D. Benjamin, J. Kinast, D. McGee, R. Martini, E. Whittaker, "Harmonic generation in a mechanical Duffing oscillator to illustrate nonlinear optical effects," *Townes Festival & Workshops on Quantum Optics*, Jackson Hole, Wyoming, presented by J. Supplee, August 2000.

J. Supplee, D. Benjamin, J. Kinast, D. McGee, R. Martini, E. Whittaker, "Harmonic generation - a mechanical educational demonstration," *APS/DAMOP 2000*, Storrs, Connecticut, presented by J. Supplee, June 2000.